Instruction type	Pipe	Stages						
Integer instruction	IF	ID	EX	MEM	WB			
FP instruction	IF	ID	EX	MEM	WB			
Integer instruction		IF	ID	EX	MEM	WB		
FP instruction		IF	ID	EX	MEM	WB		
Integer instruction			IF	ID	EX	MEM	WB	
FP instruction			IF	ID	EX	MEM	WB	
Integer instruction				IF	ID	EX	MEM	WB
FP instruction				IF	ID	EX	MEM	WB

FIGURE 6.46 Superscalar pipeline in operation. The integer and floating-point instructions are issued at the same time, and each executes at its own pace through the pipeline. This scheme will only improve the performance of programs with a fair amount of floating point.

By issuing an integer and a floating-point operation in parallel, the need for additional hardware is minimized—integer and floating-point operations use different register sets and different functional units. The only conflict arises when the integer instruction is a floating-point load, store, or move. This creates contention for the floating-point register ports and may also create a hazard if the floating-point operation uses the result of a floating-point load issued at the same time. Both problems could be solved by detecting this contention as a structural hazard and delaying the issue of the floating-point instruction. The contention could also be eliminated by providing two additional ports, a read and a write, on the floating-point register file. We would also need to add several additional bypass paths to avoid performance loss.

There is another difficulty that may limit the effectiveness of a superscalar pipeline. In our simple DLX pipeline, loads had a latency of one clock cycle; this prevented one instruction from using the result without stalling. In the superscalar pipeline, the result of a load instruction cannot be used on the same clock cycle or on the next clock cycle. This means that the next three instructions cannot use the load result without stalling; without extra ports, moves between the register sets are similarly affected. The branch delay also becomes three instructions. To effectively exploit the parallelism available in a superscalar machine, more ambitious compiler-scheduling techniques, as well as more complex instruction decoding, will need to be implemented. Loop unrolling helps generate larger straightline fragments for scheduling; more powerful compiler techniques are discussed near the end of this section.

Let's see how well loop unrolling and scheduling work on a superscalar version of DLX with the same delays in clock cycles.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 351

Example

How would the unrolled loop on page 317 be scheduled on a superscalar pipeline for DLX? To schedule it without any delays, we will need to unroll it to make five copies of the body.

Answer

The resulting code is shown in Figure 6.47.

	Integer instruction		FP instruction	Clock cycle
Loop:	LD	F0,0(R1)		1
	LD	F6,-8(R1)		2
	LD	F10,-16(R1)	ADDD F4,F0,F2	3
	LD	F14,-24(R1)	ADDD F8,F6,F2	4
	LD	F18,-32(R1)	ADDD F12,F10,F2	5
	SD	0(R1),F4	ADDD F16,F14,F2	6
	SD	-8(R1),F8	ADDD F20,F18,F2	7
	SD	-16(R1),F12		8
	SD	-24(R1),F16		9
	SUB	R1,R1,#40		10
	BNEZ	R1,LOOP		11
	SD	8(R1),F20		12

FIGURE 6.47 The unrolled and scheduled code as it would look on a superscalar DLX.

This unrolled superscalar loop now runs in 12 clock cycles per iteration, or 2.4 clock cycles per element, versus 3.5 for the scheduled and unrolled loop on the ordinary DLX pipeline. In this example, the performance of the superscalar DLX is limited by the balance between integer and floating-point computation. Every floating-point instruction is issued together with an integer instruction, but there are not enough floating-point instructions to keep the floating-point pipeline full. When scheduled, the original loop ran in 6 clock cycles per iteration. We have improved on that by a factor of 2.5, more than half of which came from loop unrolling, which took us from 6 to 3.5, with the rest coming from issuing more than one instruction per clock cycle.

Ideally, our superscalar machine will pick up two instructions and issue them both if the first is an integer and the second is a floating-point instruction. If they do not fit this pattern, which can be quickly detected, then they are issued sequentially. This points to one of the major advantages of a general superscalar machine: There is little impact on code density, and even unscheduled programs can be run. The number of issues and classes of instructions that can be issued together are the major factors that differentiate superscalar processors.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 352

Multiple Instruction Issue with Dynamic Scheduling

Multiple instruction issue can also be applied to dynamically scheduled machines. We could start with either the scoreboard scheme or Tomasulo's algorithm. Let's assume we want to extend Tomasulo's algorithm to support issuing two instructions per clock cycle, one integer and one floating point. We do not want to issue instructions in the queue out of order, since this makes the bookkeeping in the register file impossible. Rather, by employing data structures for the integer and floating-point registers, both types of instructions can be issued to their respective reservation stations, as long as the two instructions at the head of the instruction queue do not access the same register set. Unfortunately, this approach bars issuing two instructions with a dependence in the same clock cycle. This is, of course, true in the superscalar case, where it is clearly the compiler's problem. There are three approaches that can be used to achieve dual issue. First, we could use software scheduling to ensure that dependent instructions do not appear adjacent. However, this would require pipeline-scheduling software, thereby defeating one of the advantages of dynamically scheduled pipelines.

A second approach is to pipeline the instruction-issue stage so that it runs twice as fast as the basic clock rate. This permits updating the tables before processing the next instruction; then the two instructions can begin execution at once.

The third approach is based on the observation that if multiple instructions are not being issued to the same functional unit, then it will only be loads and stores that will create dependences among instructions that we wish to issue together. The need for reservation tables for loads and stores can be eliminated by using queues for the result of a load and for the source operand of a store. Since dynamic scheduling is most effective for loads and stores, while static scheduling is highly effective in register–register code sequences, we could use static scheduling to eliminate reservation stations completely and rely only on the queues for loads and stores. This style of machine organization has been called a *decoupled architecture*.

For simplicity, let us assume that we have pipelined the instruction issue logic so that we can issue two operations that are dependent but use different functional units. Let's see how this would work with our example.

Example

Consider the execution of our simple loop on a DLX pipeline extended with Tomasulo's algorithm and with multiple issue. Assume that both a floating-point and an integer operation can be issued on every clock cycle, even if they are related. The number of cycles of latency per instruction is the same. Assume that issue and write results take one cycle each, and that there is dynamic branch-prediction hardware. Create a table showing when each instruction issues, begins execution, and writes its result, for the first two iterations of the loop. Here is the original loop:

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 353

```
Loop:   LD    F0,0(R1)
        ADDD  F4,F0,F2
        SD    0(R1),F4
        SUB   R1,R1,#8
        BNEZ  R1,LOOP
```

Answer

The loop will be dynamically unwound and, whenever possible, instructions will be issued in pairs. The result is shown in Figure 6.48. The loop runs in $4 + \frac{7}{n}$ clock cycles per result for n iterations. For large n this approaches 4 clock cycles per result.

Iteration number	Instructions	Issues at clock-cycle number	Executes at clock-cycle number	Writes result at clock-cycle number
1	LD F0,0(R1)	1	2	4
1	ADDD F4,F0,F2	1	5	8
1	SD 0(R1),F4	2	9	
1	SUB R1,R1,#8	3	4	5
1	BNEZ R1,LOOP	4	5	
2	LD F0,0(R1)	5	6	8
2	ADDD F4,F0,F2	5	9	12
2	SD 0(R1),F4	6	13	
2	SUB R1,R1,#8	7	8	9
2	BNEZ R1,LOOP	8	9	

FIGURE 6.48 The time of issue, execution, and writing result for a dual-issue version of our Tomasulo pipeline. The write-result stage does not apply to either stores or branches, since they do not write any registers.

The number of dual issues is small because there is only one floating-point operation per iteration. The relative number of dual-issued instructions would be helped by the compiler partially unwinding the loop to reduce the instruction count by eliminating loop overhead. With that transformation, the loop would run as fast as on a superscalar machine. We will return to this transformation in Exercises 6.16 and 6.17.

The VLIW Approach

Our superscalar DLX machine can issue two instructions per clock cycle. That could perhaps be extended to three or at most four, but it becomes difficult to

determine whether three or four instructions can all issue simultaneously without knowing what order the instructions could be in when fetched and what dependencies might exist among them. An alternative is an LIW (*Long Instruction Word*) or VLIW (*Very Long Instruction Word*) architecture. VLIWs use multiple, independent functional units. Rather than attempting to issue multiple, independent instructions to the units, a VLIW packages the multiple operations into one very long instruction, hence the name. A VLIW instruction might include two integer operations, two floating-point operations, two memory references, and a branch. An instruction would have a set of fields for each functional unit—perhaps 16 to 24 bits per unit, yielding an instruction length of between 112 and 168 bits. To keep the functional units busy there must be enough work in a straightline code sequence to keep the instructions scheduled. This is accomplished by unrolling loops and scheduling code across basic blocks using a technique called *trace scheduling*. In addition to eliminating branches by unrolling loops, trace scheduling provides a method to move instructions across branch points. We will discuss trace scheduling more in the next section. For now, let's assume we have a technique to generate long, straightline code sequences for building up VLIW instructions.

Example

Suppose we have a VLIW that could issue two memory references, two FP operations, and one integer operation or branch in every clock cycle. Show an unrolled version of the vector sum loop for such a machine. Unroll as many times as necessary to eliminate any stalls. Ignore the branch-delay slot.

Answer

The code is shown in Figure 6.49. The loop has been unrolled 6 times, which eliminates stalls, and runs in 9 cycles. This yields a running rate of 7 results in 9 cycles, or 1.28 cycles per result.

Memory reference 1	Memory reference 2	FP operation 1	FP operation 2	Integer operation / branch
LD F0,0(R1)	LD F6,-8(R1)			
LD F10,-16(R1)	LD F14,-24(R1)			
LD F18,-32(R1)	LD F22,-40(R1)	ADDD F4,F0,F2	ADDD F8,F6,F2	
LD F26,-48(R1)		ADDD F12,F10,F2	ADDD F16,F14,F2	
		ADDD F20,F18,F2	ADDD F24,F22,F2	
SD 0(R1),F4	SD -8(R1),F8	ADDD F28,F26,F2		
SD -16(R1),F12	SD -24(R1),F16			
SD -32(R1),F20	SD -40(R1),F24			SUB R1,R1,#48
SD -0(R1),F28				BNEZ R1,LOOP

FIGURE 6.49 VLIW instructions that occupy the inner loop and replace the unrolled sequence. This code takes nine cycles assuming no branch delay; normally the branch would also be scheduled. The issue rate is 23 operations in 9 clock cycles, or 2.5 operations per cycle. The efficiency, the percentage of available slots that contained an operation, is about 60%. To achieve this issue rate requires a much larger number of registers than DLX would normally use in this loop.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 355

What are the limitations and costs of a VLIW approach? If we can issue 5 operations per clock cycle, why not 50? Three different limitations are encountered: limited parallelism, limited hardware resources, and code size explosion. The first is the simplest: There is a limited amount of parallelism available in instruction sequences. Unless loops are unrolled very large numbers of times, there may not be enough operations to fill the instructions. At first glance, it might appear that 5 instructions that could be executed in parallel would be sufficient to keep our VLIW completely busy. This, however, is not the case. Several of these functional units—the memory, the branch, and the floating-point units—will be pipelined, requiring a much larger number of operations that can be executed in parallel. For example, if the floating-point pipeline has 8 steps, the 2 operations being issued on a clock cycle cannot depend on any of the 14 operations already in the floating-point pipeline. Thus, we need to find a number of independent operations roughly equal to the average pipeline depth times the number of functional units. This means about 15 to 20 operations would be needed to keep a VLIW with 5 functional units busy.

The second cost, the hardware resources for a VLIW, seem quite straightforward; duplicating the floating-point and integer functional units is easy and cost scales linearly. However, there is a large increase in the memory- and register-file bandwidth. Even with a split floating-point and integer register file, our VLIW will require 5 read ports and 2 write ports on the integer register file and 4 read ports and 2 write ports on the floating-point register file. This bandwidth cannot be supported without some substantial cost in the size of the register file and possible degradation of clock speed. Our 5-unit VLIW also has 2 data memory ports. Furthermore, if we wanted to expand it, we would need to continue adding memory ports. Adding only arithmetic units would not help, since the machine would be starved for memory bandwidth. As the number of data memory ports grows, so does the complexity of the memory system. To allow multiple memory accesses in parallel, the memory must be broken into banks containing different addresses with the hope that the operations in a single instruction do not have conflicting accesses. A conflict will cause the entire machine to stall, since all the functional units must be kept synchronized. This same factor makes it extremely difficult to use data caches in a VLIW.

Finally, there is the problem of code size. There are two different elements that combine to increase code size substantially. First, generating enough operations in a straightline code fragment requires ambitiously unrolling loops, which increases code size. Second, whenever instructions are not full, the unused functional units translate to wasted bits in the instruction encoding. In Figure 6.49, we saw that only about 60% of the functional units were used; almost half of each instruction was empty. To combat this problem, clever encodings are sometimes used. For example, there may be only one large immediate field for use by any functional unit. Another technique is to compress the instructions in main memory and expand them when they are read into the cache or are decoded.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 356

The major challenge for these machines is to try to exploit large amounts of instruction-level parallelism. When the parallelism comes from unrolling simple loops, the original loop probably could have been run efficiently on a vector machine (see the next chapter). It is not clear that a VLIW is preferred over a vector machine for such applications; the costs are similar, and the vector machine is typically the same speed or faster. The open question in 1990 is whether there are large classes of applications that are not suitable for vector machines, but still offer enough parallelism to justify the VLIW approach rather than a simpler one, such as a superscalar machine.

Increasing Instruction-Level Parallelism with Software Pipelining and Trace Scheduling

We have already seen that one compiler technique, loop unrolling, is used to help exploit parallelism among instructions. Loop unrolling creates longer sequences of straightline code, which can be used to exploit more instruction-level parallelism. There are two other more general techniques that have been developed for this purpose: software pipelining and trace scheduling.

Software pipelining is a technique for reorganizing loops such that each iteration in the software-pipelined code is made from instruction sequences chosen from different iterations in the original code segment. This is most easily understood by looking at the scheduled code for the superscalar version of DLX. The scheduler essentially interleaves instructions from different loop iterations, putting together all the loads, then all the adds, then all the stores. A software-pipelined loop interleaves instructions from different iterations without unrolling the loop. This technique is the software counterpart to what Tomasulo's algorithm does in hardware. The software-pipelined loop would contain one load, one add, and one store, each from a different iteration. There is also some startup code that is needed before the loop begins as well as code to finish up after the loop is completed. We will ignore these in this discussion.

Example

Show a software-pipelined version of this loop:

```
Loop:  LD    F0,0(R1)
       ADDD  F4,F0,F2
       SD    0(R1),F4
       SUB   R1,R1,#8
       BNEZ  R1,LOOP
```

You may omit the start-up and clean-up code.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 357

Answer

Given the vector M in memory, and ignoring the start-up and finishing code, we have:

```
Loop:   SD   0(R1),F4    ;stores into M[i]
        ADDD F4,F0,F2    ;adds to M[i-1]
        LD   F0,-16(R1) ;loads M[i-2]
        BNEZ R1,LOOP
        SUB  R1,R1,#8    ;subtract in delay slot
```

This loop can be run at a rate of 5 cycles per result, ignoring the start-up and clean-up portions. Because the load fetches two array elements beyond the element count, the loop should run for two fewer iterations. This would be accomplished by decrementing R1 by 16 prior to the loop.

Software pipelining can be thought of as symbolic loop unrolling. Indeed, some of the algorithms for software pipelining use loop unrolling to figure out how to software pipeline the loop. The major advantage of software pipelining over straight loop unrolling is that software pipelining consumes less code space. Software pipelining and loop unrolling, in addition to yielding a better scheduled inner loop, each reduce a different type of overhead. Loop unrolling reduces the overhead of the loop—the branch and counter-update code. Software pipelining reduces the time when the loop is not running at peak speed to once per loop at the beginning and end. If we unroll a loop that does 100 iterations a constant number of times, say 4, we pay the overhead 100/4 = 25 times—every time the inner unrolled loop is reinitiated. Figure 6.50 shows this behavior graphically. Because these techniques attack two different types of overhead, the best performance comes from doing both.

The other technique used to generate additional parallelism is *trace scheduling*. This is particularly useful for VLIWs, for which the technique was originally developed. Trace scheduling is a combination of two separate processes. The first process, called *trace selection* tries to find the most likely sequence of operations to put together into a small number of instructions; this sequence is called a *trace*. Loop unrolling is used to generate long traces, since loop branches are taken with high probability. Once a trace is selected, the second process, called *trace compaction*, tries to squeeze the trace into a small number of wide instructions. Trace compaction attempts to move operations as early as it can in a sequence (trace), packing the operations into as few wide instructions as possible.

There are two different considerations in compacting a trace: data dependences, which force a partial order on operations, and branch points, which create places across which code cannot be easily moved. In essence, the code wants to be compacted into the shortest possible sequence that preserves the data dependences; branches are the main impediment to this process. The major advantage of trace scheduling over simpler pipeline-scheduling techniques is that it includes a method to move code across branches. Figure 6.51 shows a code fragment, which may be thought of as an iteration of an unrolled loop, and the trace selected.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 358

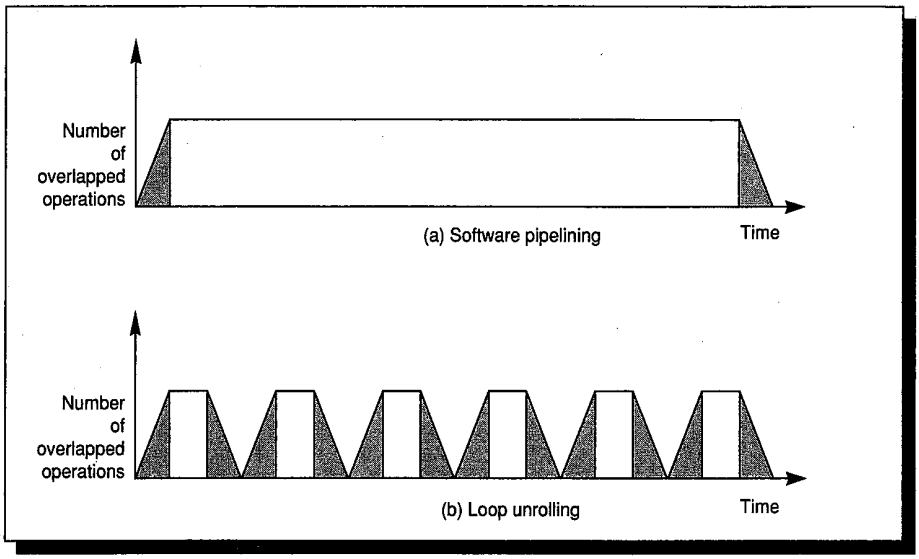

FIGURE 6.50 This shows the execution pattern for (a) a software-pipelined loop and (b) an unrolled loop. The shaded areas are the times when the loop is not running with maximum overlap or parallelism among instructions. This occurs once at loop beginning and once at the end for the software-pipelined loop. For the unrolled loop it occurs $\frac{m}{n}$ times if the loop has a total of m executions and is unrolled n times. Each block represents an unroll of n iterations. Increasing the number of unrolls will reduce the start-up and clean-up overhead.

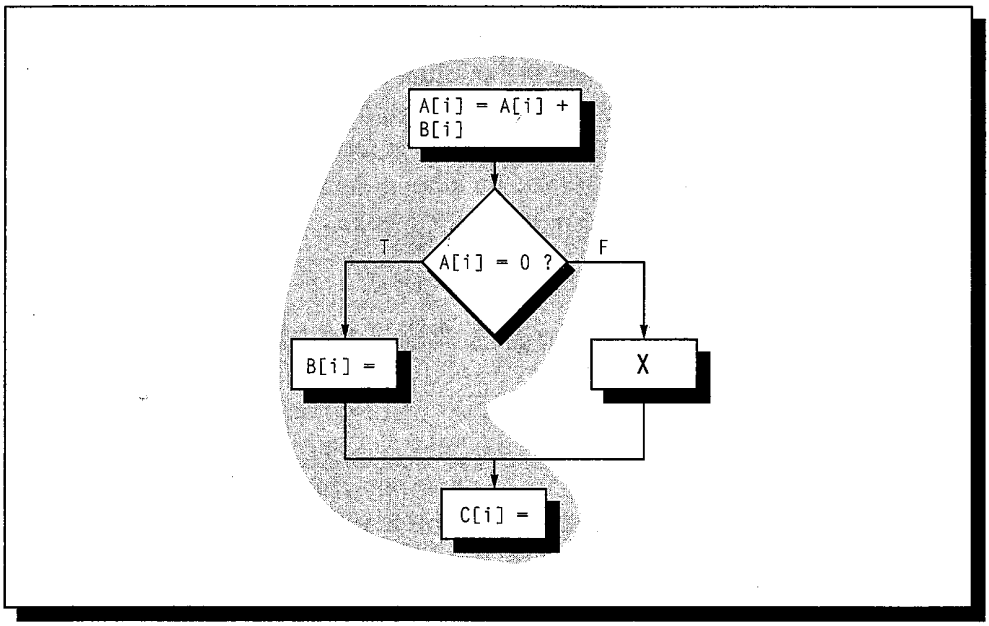

FIGURE 6.51 A code fragment and the trace selected shaded with gray. This trace would be selected first, if the probability of the true branch being taken were much higher than the probability of the false branch being taken. The branch from the decision (A[i]=0) to X is a branch **out** of the trace, and the branch from X to the assignment to C is a branch **into** the trace. These branches are what make compacting the trace difficult.

Once the trace is selected as shown in Figure 6.51, it must be compacted so as to fill the wide instruction word. Compacting the trace involves moving the assignments to variables B and C up to the block before the branch decision. Let's first consider the problem of moving the assignment to B. If the assignment to B is moved above the branch (and thus out of the trace), the code in X would be affected if it used B, since moving the assignment would change the value of B. Thus, to move the assignment to B, B must not be read in X. One could imagine more clever schemes if B were read in X—for example, making a shadow copy and updating B later. Such schemes are generally not used, both because they are complex to implement and because they will slow down the program if the trace selected is not optimal and the operations end up requiring additional instructions. Also, because the assignment to B is moved before the if test, for this schedule to be valid either X also assigns to B or B is not read after the if statement.

Moving the assignment to C up to before the first branch requires first moving it over the branch from X into the trace. To do this, a copy is made of the assignment to C on the branch into the trace. A check must still be done, as was done for B, to make sure that the assignment can be moved over the branch out of the trace. If C is successfully moved to before the first branch and the "false" direction of the branch—the branch off the trace—is taken, the assignment to C will have been done twice. This may be slower than the original code, depending on whether this operation or other moved operations create additional work in the main trace. Ironically, the more successful the trace-scheduling algorithm is in moving code across the branch, the higher the penalty for misprediction.

Loop unrolling, trace scheduling, and software pipelining all aim at trying to increase the amount of local instruction parallelism that can be exploited by a machine issuing more than one instruction on every clock cycle. The effectiveness of each of these techniques and their suitability for various architectural approaches are among the most significant open research areas in pipelined-processor design.

6.9 Putting It All Together: A Pipelined VAX

In this section we will examine the pipeline of the VAX 8600, a macropipelined VAX. This machine is described in detail by DeRosa et al. [1985] and Troiani et al. [1985]. The 8600 pipeline is a more dynamic structure than the DLX integer pipeline. This is because the processing steps may take multiple cycles in one stage of the pipeline. Additionally, the hazard detection is more complicated because of the possibility that stages progress independently and because instructions may modify registers before they complete. Techniques similar to those used in the DLX FP pipeline to handle variable-length instructions are used in the 8600 pipeline.

The 8600 is macropipelined—the pipeline understands the structure of VAX instructions and overlaps their execution, checking the hazards on the instruction

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 360

operands. By comparison, the VAX 8800 is micropipelined—microinstructions are overlapped and hazard detection occurs in the microprogram unit. A different issue of the Digital Technical Journal [Digital 1987] describes this machine, and Clark [1987] describes the pipeline and its performance. The designs are interesting to compare.

Figure 6.52 shows the 8600 partitioned into four major structural components. The MBox is responsible for address translation and memory access (see Chapter 8). The IBox is the heart of the 8600 pipeline; it is responsible for instruction fetch and decode, operand address calculation, and operand fetch. The EBox and FBox are responsible for execution of integer and floating-point operations, and their primary function is to implement the opcode portion of an instruction. (Because the FBox is optional, the EBox also contains microcode to do the floating point, albeit at much lower performance. The optional presence of the FBox further complicates the operand processing in the EBox.) Since the EBox and FBox are not pipelined, we will focus our attention primarily on the IBox. In explaining the IBox function we will refer to the EBox occasionally; usually the same comments apply to the FBox.

Figure 6.53 breaks the execution of a VAX instruction into four overlapped steps. The number of clock cycles per step may vary widely, though each step in the pipeline takes at least one clock.

A VAX instruction may take many clock cycles in a given step. For example, with multiple memory operands, the instruction will take multiple clock cycles in the Opfetch step. Because of this, an instruction that takes many cycles at a

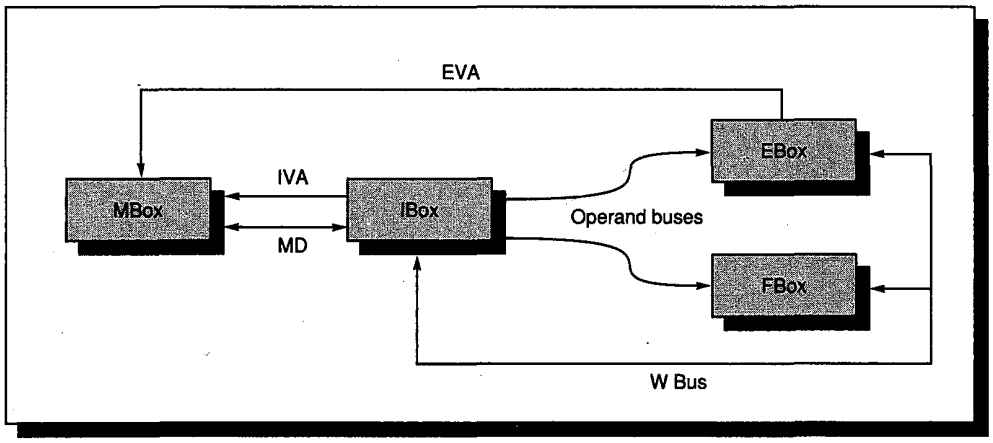

FIGURE 6.52 The basic structure of the 8600 consists of an MBox (responsible for memory access), IBox (handles instruction and operand processing), EBox (all opcode interpretation except floating point), and FBox (performs floating-point operations). These four units are connected by six major buses. The IVA and EVA carry the address for a memory access to the MBox from the IBox and EBox. The MD bus carries memory data to or from the MBox; all such data flows through the IBox. The EBox initiates memory access directly with the MBox only under unusual conditions (e.g., misaligned references). The operand buses carry operands from the IBox (where they are fetched from memory or registers) to the EBox and FBox. Finally, the W Bus carries results to be written from the EBox and FBox to the GPRs and to memory, via the IBox.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 361

Step	Function	Located in
1. Ifetch	Prefetch instruction bytes and decode them	IBox
2. Opfetch	Operand address calculation and fetch	IBox
3. Execution	Execute opcode and write result	EBox, FBox
4. Result store	Write result to memory or registers	EBox, IBox

FIGURE 6.53 The basic structure of the 8600 pipeline has four stages, each taking from 1 to a large number of clock cycles. Up to four VAX instructions are being processed at once.

stage may cause a back up in the pipeline; this back up may eventually reach the Ifetch step, where it will cause the pipeline to simply stop fetching instructions. Additionally, several resources (e.g., the W Bus and GPR ports) are contended for by multiple stages in the pipeline. In general, these problems are resolved on the fly using a fixed-priority scheme.

Operand Decode and Fetch

Much of the work in interpreting a VAX instruction is in the operand specifier and decode process, and this is the heart of the IBox. Substantial effort is devoted to decoding and fetching operands as fast as possible to keep instructions flowing through the pipeline. Figure 6.54 shows the number of cycles spent in Opfetch under ideal conditions (no cache misses or other stalls from the memory hierarchy) for each operand specifier. If the result is a register, the EBox stores

Specifier	Cycles
Literal or immediate	1
Register	1
Deferred	1
Displacement	1
PC-relative and absolute	1
Autodecrement	1
Autoincrement	2
Autoincrement deferred	5
Displacement deferred	4
PC-relative deferred	4

FIGURE 6.54 The minimum number of cycles spent in Opfetch by operand specifier. This shows the data for an operand of type byte, word, or longword that is read. Modified and written operands take an additional cycle, except for register mode and immediate or literal, where writes are not allowed. Quadword and octaword operands may take much longer. If any stalls are encountered, the cycle count will increase.

the result. If the result is a memory operand, Opfetch calculates the address and waits for the EBox to signal ready, then the IBox stores the result during the Result store step. If an instruction result is to be stored in memory, the EBox signals to the IBox when it enters the last cycle of execution for the instruction. This allows Opfetch to overlap the first cycle of a two-cycle memory write with the last cycle of execution (even if the operation only takes one cycle).

To maximize the performance of the machine, there are three copies of the GPRs—in the IBox, EBox, and FBox. A write is broadcast from the FBox, EBox, or IBox (in the case of autoincrement or autodecrement addressing) to the other two units, so that their copies of the registers can be updated.

Handling Data Dependences

Register hazards are tracked in Opfetch by maintaining a small table of registers that will be written. Whenever an instruction passes through Opfetch, its result register is marked as busy. If an instruction that uses that register arrives in Opfetch and sees the busy flag set, it stalls until the flag is cleared. This prevents RAW hazards. The busy flag is cleared when the register is written. Because there are only two stages after Opfetch (execute and write memory result), the busy flag can be implemented as a two-entry associative memory. Writes are maintained in order and always at the end of the pipeline, and all reads are done in Opfetch. This eliminates all explicit WAW and WAR hazards. The only possible remaining hazards are those that can occur on implicit operands, such as the registers written by a `MOVC3`. Hazards on implicit operands are prevented by explicit control in the microcode.

Opfetch optimizes the case when the last operand specifier is a register by processing the register operand specifier at the same time as the next-to-last specifier. In addition, when the result register of an instruction is the source operand of the next instruction, rather than stall the dependent instruction, Opfetch merely signals this relationship to the EBox, allowing execution to proceed without a stall. This is like the bypassing in our DLX pipeline.

Memory hazards between reads and writes are easily resolved because there is a single memory port, and the IBox decodes all operand addresses.

Handling Control Dependences

There are two aspects to handling branches in a VAX: synchronizing on the condition code and dealing with the branch hazard. Most of the branch processing is handled by the IBox. A predict-taken strategy is used; the following steps are taken when the IBox sees a branch:

1. Compute the branch target address, send it to the MBox, and initiate a fetch from the target address. Wait for the EBox to issue CCSYNC, which indicates that the condition codes will be available in the next clock cycle.

2. Evaluate the condition codes from the EBox to check the prediction. If the prediction was incorrect, the access initiated in the MBox is aborted. The current PC points at the next instruction or its first operand specifier.
3. Assuming the branch was taken, the IBox flushes the prefetch and decode stages and begins loading the instruction register and processing the new target stream. If the branch was not taken, the access to the potential target has already been killed and the pipeline can continue just using what is in the prefetch and decode stages.

Simple conditional branches (`BEQL`, `BNEQ`), the unconditional branches (`BRB`, `BRW`), and the computed branches (e.g., `AOBLEQ`) are handled by the IBox. The EBox handles more complex branches and also the instructions used for calls and returns.

An Example

To really understand how this pipeline works, let's look at how a code sequence executes. This example is somewhat simplified, but is sufficient to demonstrate the major pipeline interactions. The code sequence we will consider is as follows (remember that for consistency the result of the `ADDL3` is given first):

```
            ADDL3    R1,R2,56(R3)
            CMPL     45(R1),@54(R2)
            BEQL     target
            MOVL     ...
target:     SUBL3    ...
```

Figure 6.55 shows an annotated pipeline diagram of how these instructions would progress through the 8600 pipeline.

Dealing with Interrupts

The 8600 maintains three program counters so that instruction interruption and restart are possible. These program counters and what they designate are:

- Current Program Counter—points to the next byte to be processed and consumed in Opfetch.
- IBox Starting Address—points to the instruction currently in Opfetch.
- EBox Starting Address—points to the instruction executing in the EBox or FBox.

In addition, the prefetch unit keeps an address to prefetch from (the VIBA, Virtual Instruction Buffer Address), but this does not affect interrupt handling. When an exception is caused by a prefetch operation, the byte in the instruction buffer is marked. When Opfetch eventually asks for the byte, it will see the exception, and the Current Program Counter will have the address of the byte that caused the exception.

Instr.	Clock Cycle 1	2	3	4	5	6	7	8	9
ADDL3	**IF**: Fetch ADDL.	**IF**: Continue prefetch if space and MBox available.	**IF**: Decode R1.	**IF**: Decode R2. **OP**: Fetch R1.	**IF**: Decode 56(R3). **OP**: Fetch R2.	**OP**: Compute 56+(R3). **EX**: get first operand.	**OP**: Start write. **EX**: Add.	**WR**: Store.	
CMPL						**IF**: Decode 45(R1).		**IF**: Decode @54(R2). **OP**: Fetch 45(R1).	**OP**: Fetch 54(R2).
BEQL									**IF**: Decode BEQL displace.
SUBL									

Instr.	Clock Cycle 10	11	12	13	14	15	16	17	18
ADDL3									
CMPL	**OP**: stall. **EX**: get first operand.	**OP**: get indirect address.	**OP**: Fetch @54(R2).		**EX**: compare and set CC.				
BEQL				**OP**: Load VA.	**OP**: Fetch branch target.	**OP**: Fetch target +4; load VIBA; flush IBuffer.			
SUBL							**IF**: Decode SUBL3.	**OP**: Fetch first operand.	**OP**: Fetch second operand.

FIGURE 6.55 The VAX 8600 executing a code sequence. The top portion shows the events on clock ticks 1–9, while the bottom portion shows the events on clock ticks 10–18. The pipeline stages are abbreviated as IF (Instruction Fetch), OP (Opfetch), EX (Execution), and WR (Write Result) and are shown in bold. Each instruction passes through the 8600 pipeline as soon as the pipe stage is empty and the required data is available. Note that an instruction can be in both the IF and OP stages at the same time. This figure assumes that at the beginning of cycle 1, the prefetch buffer is empty. The prefetch in the IF stage continues to fetch instructions as long as there is room in the prefetch buffer and an available MBox cycle. It is omitted from the diagram for simplicity. The action "stall" indicates a stall for a memory operand during Opfetch. In total, the three VAX instructions executed take 15 cycles, assuming no stalls from the memory system. This sequence was chosen to demonstrate the functioning of the pipeline—it is not necessarily typical.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 365

These PCs are updated when an instruction enters the corresponding pipeline stage. Hence, if an interrupt occurs in a given stage, the PC can be set back to the beginning of that instruction. These PCs are needed because the length of VAX instructions is variable and can only be determined by finding the opcode byte.

In addition to restoring the starting address of the instruction that caused the interrupt, we must unwind any register updates done by addressing modes processed in Opfetch for instructions that are after the instruction that interrupts the processor. The IBox maintains a log of updates to the register file done on behalf of multiple instructions, as we did in Section 5.6. The effects of any changes are undone and the PC is restored. This allows the operating system to have a clean machine state to work from.

Final Remarks

The 8600 uses a four-step pipeline. The theoretical peak performance with the 80-ns clock is 12.5 million VAX instructions per second. Some simple sequences of instructions can actually attain this peak performance with a CPI of 1. Typically, the performance on integer code is about 1.75 million VAX instructions per second for a CPI of about 7. This yields about 3.5 times the performance of a VAX-11/780.

6.10 Fallacies and Pitfalls

Fallacy: Instruction set design has little impact on pipelining.

This is perhaps the most prominent misconception about pipelining and one that was widely held until recently. Many of the difficulties of pipelining arise because of instruction set complications. Here are some examples, many of which are mentioned in the chapter:

- Variable instruction lengths and running times can lead to imbalance among pipeline stages causing other stages to back up. They also severely complicate hazard detection and the maintenance of precise interrupts. Of course, there are exceptions to every rule. For example, caches cause instruction running times to vary when they miss; however, the performance advantages of caches make the added complexity acceptable. To minimize the complexity, most machines freeze the pipeline on a cache miss. Other machines try to continue running parts of the pipeline; though this is very complex, it may overcome some of the performance losses from cache misses.
- Sophisticated addressing modes can lead to different sorts of problems. Addressing modes that update registers, such as post autoincrement, complicate

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 366

hazard detection. They also slightly increase the complexity of instruction restart. Other addressing modes that require multiple memory accesses substantially complicate pipeline control and make it difficult to keep the pipeline flowing smoothly.

- Architectures that allow writes into the instruction space (self-modifying code) can cause trouble for pipelining (as well as for cache designs). For example, if an instruction in the pipeline can modify another instruction, we must constantly check if the address being written to by an instruction corresponds to the address of an instruction further on in the pipeline. If so, the pipeline must be flushed or the instruction in the pipeline somehow updated.
- Implicitly set condition codes increase the difficulty of finding when a branch has been decided and the difficulty of scheduling branch delays. The former problem occurs when the condition-code setting is not uniform, making it difficult to decide which instruction sets the condition code last. The latter problem occurs when the setting of the condition code is not under program control. This makes it hard to find instructions that can be scheduled between the condition evaluation and the branch. Many newer architectures avoid condition codes or set them explicitly under program control to eliminate the pipelining difficulties.

As a simple example, suppose the DLX instruction format were more complex, so that a separate, decode pipe stage were required before register fetch. This would increase the branch delay to two clock cycles. At best, the second branch-delay slot would be wasted at least as often as the first. Gross [1983] found that a second delay slot was only used half as often as the first. This would lead to a performance penalty for the second delay slot of more than 0.1 clock cycles per instruction.

Pitfall: Unexpected execution sequences may cause unexpected hazards.

At first glance, WAW hazards look like they should never occur because no compiler would ever generate two writes to the same register without an intervening read. But they can occur when the sequence was unexpected. For example, the first write might be in the delay slot of a taken branch when the scheduler thought the branch would not be taken. Here is the code sequence that could cause this:

```
        BNEZ    R1,foo
        DIVD    F0,F2,F4    ; moved into delay slot
                            ; from fall through
        .....
        .....
foo:    LD      F0,qrs
```

If the branch is taken, then before the DIVD can complete the LD will reach WB, causing a WAW hazard. The hardware must detect this and may stall the issue of the LD. Another way this can happen is if the second write is in a trap routine. This occurs when an instruction that traps and is writing results continues and completes after an instruction that writes the same register in the trap handler. The hardware must detect and prevent this as well.

Fallacy: Increasing the depth of pipelining always increases performance.

Two factors combine to limit the performance improvement gained by pipelining. Data dependences in the code mean that increasing the pipeline depth will increase the CPI, since a larger percentage of the cycles will become stalls. Second, clock skew and latch overhead combine to limit the decrease in clock period obtained by further pipelining. Figure 6.56 shows the tradeoff between pipeline depth and performance for the first 14 of the Livermore Loops (see Chapter 2, page 43). The performance flattens out when the pipeline depth reaches 4 and actually drops when the execution portion is pipelined 16 deep.

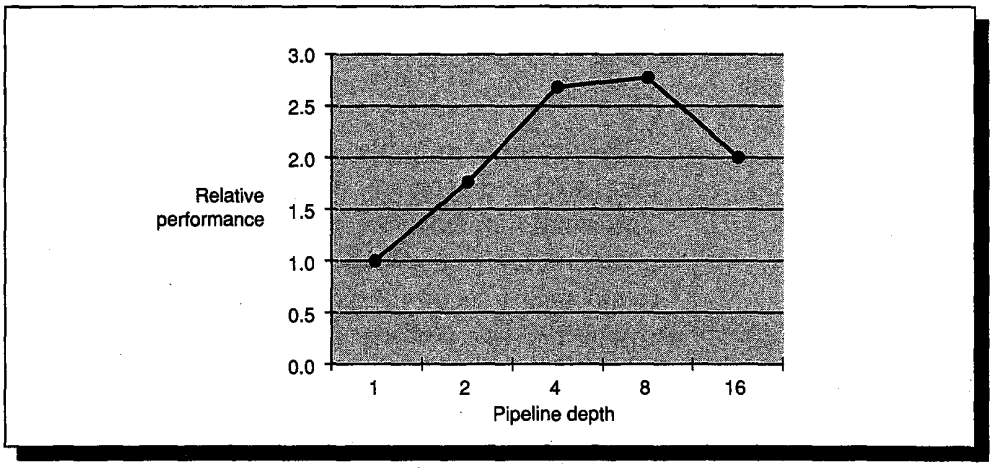

FIGURE 6.56 The depth of pipelining versus the speedup obtained. This data is based on Table 2 in Kunkel and Smith [1986]. The x axis shows the number of stages in the EX portion of the floating-point pipeline. A single-stage pipeline corresponds to 32 levels of logic, which might be appropriate for a single FP operation.

Pitfall: Evaluating a scheduler on the basis of unoptimized code.

Unoptimized code—containing redundant loads, stores, and other operations that might be eliminated by an optimizer—is much easier to schedule than "tight" optimized code. In GCC running on a DECstation 3100, the frequency of idle clock cycles increases by 18% from the unoptimized and scheduled code to the optimized and scheduled code. TeX shows a 20% increase for the same measurement. To fairly evaluate a scheduler you must use optimized code, since in the real system you will derive a good performance from other optimizations in addition to scheduling.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 368

Pitfall: Extensive pipelining can impact other aspects of a design, leading to overall lower cost/performance.

The best example of this phenomenon comes from two implementations of the VAX, the 8600 and the 8700. We discussed the instruction pipeline of the 8600 in Section 6.9. When the 8600 was initially delivered, it had a cycle time of 80 ns. Subsequently, a redesigned version, called the 8650, with a 55-ns clock was introduced. The 8700 has a much simpler pipeline that operates at the microinstruction level. The 8700 CPU is much smaller and has a faster clock rate, 45 ns. The overall outcome is that the 8650 has a CPI advantage of about 20%, but the 8700 has a clock rate that is about 20% faster. Thus, the 8700 achieves the same performance with much less hardware.

6.11 Concluding Remarks

Figure 6.57 shows how the various pipelining approaches affect both clock speed and CPI. This figure does not account for instruction-count differences. Since performance is clock speed divided by CPI (ignoring instruction-count differences), machines in the top left corner will be slowest, and machines in the bottom right corner will be fastest. However, the machines that move towards the lower right corner will probably achieve their maximum performance on the narrowest range of applications.

Machines that are *underpipelined* lump multiple DLX pipestages into one. The clock cannot be run as fast, and the CPI will be only marginally lower. The DLX pipeline achieves a CPI very close to 1 (ignoring memory-system stalls) at a reasonable clock speed. Architectural simplicity and efficient pipelining are two of the most important attributes of the RISC (Reduced Instruction Set Computer) machines. DLX constitutes an example of such a machine. We have chosen to use the term load/store architecture because the ideas apply to a broad range of machines, and not just to the machines that identify themselves as RISCs. Much of the discussion in the first part of this chapter centered around the key ideas developed by the RISC projects.

Machines with higher clock rates and deeper pipelines have been called *superpipelined*. Superpipelined machines are characterized by pipelining all functional units. A superpipelined version of DLX might have a 10-stage pipeline, rather than the 5-stage pipeline described earlier. Other than increasing the complexity of pipeline scheduling and pipeline control, superpipelined machines are not fundamentally different from the machines we have already examined in this chapter. Due to limited instruction-level parallelism, a superpipelined machine will have a slightly higher CPI than a DLX-style pipeline, but its advantage in clock cycle time should be larger than the disadvantage in CPI.

Superscalar processors can have clock cycle times very close to that of a DLX pipeline and maintain a smaller CPI. The VLIW machines can have a

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 369

substantially lower CPI, but tend to have a significantly higher clock cycle time for the reasons discussed in this chapter. The vector machines effectively use both techniques. They are usually superpipelined and have powerful vector operations that can be considered equivalent to issuing multiple independent operations on a machine like DLX. We will explore vector machines in detail in the next chapter.

Going out from the top left corner on either axis in Figure 6.57, the requirement to exploit more instruction-level parallelism increases; at the same time, of course, fewer programs will run at maximum speed.

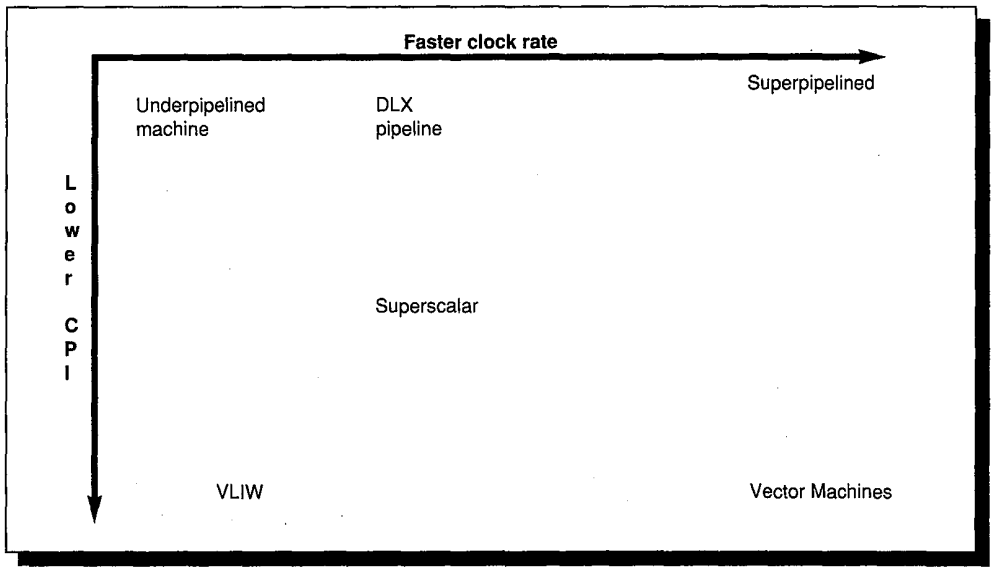

FIGURE 6.57 Increasing the instruction-issue rate lowers the CPI, while a deeper pipeline increases the clock rate. Various machines combine these techniques.

6.12 Historical Perspective and References

This section describes some of the major advances in pipelining and ends with some of the recent literature on high-performance pipelining.

The first general-purpose pipelined machine is considered to be Stretch, the IBM 7030. Stretch followed on the IBM 704 and had a goal of being 100 times faster than the 704. The goals were a stretch from the state of the art at that time—hence the nickname. The plan was to obtain a factor of 1.6 from overlapping fetch, decode, and execute, using a 4-stage pipeline. Bloch [1959] and Bucholtz [1962] describe the design and engineering tradeoffs, including the use of ALU bypasses.

In 1964 CDC delivered the first CDC 6600. The CDC 6600 was unique in many ways. In addition to introducing scoreboarding, the CDC 6600 was the first machine to make extensive use of multiple functional units. It also had

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 370

peripheral processors that used a timeshared pipeline. The interaction between pipelining and instruction set design was understood, and the instruction set was kept simple to promote pipelining. The CDC 6600 also used an advanced packaging technology. Thornton [1964] describes the pipeline and I/O processor architecture, including the concept of out-of-order instruction execution. Thornton's book [1970] provides an excellent description of the entire machine, from technology to architecture, and includes a foreword by Cray. (Unfortunately, this book is currently out of print.) The CDC 6600 also has an instruction scheduler for the FORTRAN compilers, described by Thorlin [1967].

The IBM 360/91 introduced many new concepts, including tagging of data, register renaming, dynamic detection of memory hazards, and generalized forwarding. Tomasulo's algorithm is described in his 1967 paper. Anderson, Sparacio, and Tomasulo [1967] describe other aspects of the machine, including the use of branch prediction. Patt and his colleagues have described an approach, called HPSm, that is an extension of Tomasulo's algorithm [Hwu and Patt 1986].

A series of general pipelining descriptions that appeared in the late 1970s and early 1980s provided most of the terminology and described most of the basic techniques used in simple pipelines. These surveys include Keller [1975], Ramamoorthy and Li [1977], Chen [1980], and Kogge's book [1981], devoted entirely to pipelining. Davidson and his colleagues [1971, 1975] developed the concept of pipeline reservation tables as a design methodology for multicycle pipelines with feedback (also described in Kogge [1981]). Many designers use a variation of these concepts, as we did in Figures 6.3 and 6.4.

The RISC machines refined the notion of compiler-scheduled pipelines in the early 1980s. The concepts of delayed branches and delayed loads—common in microprogramming—were extended into the high-level architecture. The Stanford MIPS architecture made the pipeline structure purposely visible to the compiler and allowed multiple operations per instruction. Schemes for scheduling the pipeline in the compiler were described by Sites [1979] for the Cray, by Hennessy and Gross [1983], (and in Gross's thesis [1983]) and by Gibbons and Muchnik [1986]. Rymarczyk [1982] describes the interlock conditions that programmers should be aware of for a 360-like machine; this paper also shows the complex interaction between pipelining and an instruction set not designed to be pipelined.

J. E. Smith and his colleagues have written a number of papers examining instruction issue, interrupt handling, and pipeline depth for high-speed scalar machines. Kunkel and Smith [1986] evaluate the impact of pipeline overhead and dependences on the choice of optimal pipeline depth; they also have an excellent discussion of latch design and its impact on pipelining. Smith and Plezkun [1988] evaluate a variety of techniques for preserving precise interrupts, including the future file concept mentioned in Section 6.6. Weiss and Smith [1984] evaluate a variety of hardware pipeline scheduling and instruction-issue techniques.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 371

Dynamic hardware branch-prediction schemes are described by J. E. Smith [1981] and by A. Smith and Lee [1984]. Ditzel [1987] describes a novel branch-target buffer for CRISP. McFarling and Hennessy [1986] is a quantitative comparison of a variety of compile-time and run-time branch-prediction schemes.

A series of early papers, including Tjaden and Flynn [1970] and Foster and Riseman [1972], concluded that only small amounts of parallelism could be available at the instruction level without investing an enormous amount of hardware. These papers dampened the appeal of multiple instruction issue for more than ten years. Nicolau and Fisher [1984] published a paper asserting the presence of large amounts of potential instruction-level parallelism.

Charlesworth [1981] reports on the Floating Point Systems AP-120B, one of the first wide-instruction machines containing multiple operations per instruction. Floating Point Systems applied the concept of software pipelining—albeit by hand, rather than with a compiler—by writing assembly language libraries to use the machine efficiently. Weiss and J. E. Smith [1987] compare software pipelining versus loop unrolling as techniques for scheduling code on a pipelined machine. Lam [1988] presents algorithms for software pipelining and evaluates their use on Warp, a wide-instruction-word machine. Along with his colleagues at Yale, Fisher [1983] proposed creating a machine with a very wide instruction (512 bits), and named this type of machine a VLIW. Code was generated for the machine using trace scheduling, which Fisher [1981] had developed originally for generating horizontal microcode. The implementation of trace scheduling for the Yale machine is described by Fisher, et. al. [1984] and by Ellis [1986]. The Multiflow machine (see Colwell et. al. [1987]) commercialized the concepts developed at Yale.

Several researchers proposed techniques for multiple instruction issue. Agerwala and Cocke [1987] proposed this approach as an extension of the RISC ideas, and coined the name "superscalar." IBM described a machine based on these ideas in late 1989 (see Bakoglu et al. [1989]). In 1990, the IBM was announced as the RS/6000. The implementation can issue up to four instructions per clock. A good description of the machine, its background, and software appears in IBM [1990]. The Apollo DN 10000 and the Intel i860 both offer multiple instruction issue, though the requirements for multiple issue are rather rigid. The Intel i860 should probably be considered a LIW machine because the program must explicitly indicate whether instruction pairs should be dual issued. Although the pairs are ordinary instructions, there are substantial limitations on what can appear as a member of a dual-issued pair. The Intel 960CA and Tandem Cyclone are examples of superscalar machines with complex instruction sets.

J. E. Smith and his colleagues at Wisconsin [1984] proposed the decoupled approach that included multiple issue with dynamic pipeline scheduling. The Astronautics ZS-1 described by Smith et al. [1987] embodies this approach and uses queues to connect the load/store unit and the operation units. J. E. Smith [1989] also describes the advantages of dynamic scheduling and compares that approach to static scheduling. Dehnert, Hsu, and Bratt [1989] explain the

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 372

architecture and performance of the Cydrome Cydra 5, a machine with a wide instruction word that provides dynamic register renaming. The Cydra 5 is a unique blend of hardware and software aimed at extracting instruction-level parallelism.

Recently there have been a number of papers exploring the tradeoffs among alternative pipelining approaches. Jouppi and Wall [1989] examine the performance differences between superpipelined and superscalar systems, concluding that their performance is similar, but that superpipelined machines may require less hardware to achieve the same performance. Sohi and Vajapeyam [1989] give measurements of available parallelism for wide-instruction-word machines. Smith, Johnson, and Horowitz [1989] recount studies of available instruction-level parallelism in nonscientific code using an ambitious hardware scheme that allows multiple-instruction execution.

References

AGERWALA, T. AND J. COCKE [1987]. "High performance reduced instruction set processors," IBM Tech. Rep. (March).

ANDERSON, D. W., F. J. SPARACIO, AND R. M. TOMASULO [1967]. "The IBM 360 Model 91: Machine philosophy and instruction handling," *IBM J. of Research and Development* 11:1 (January) 8–24.

BAKOGLU, H. B., G. F. GROHOSKI, L. E. THATCHER, J. A. KAHLE, C. R. MOORE, D. P. TUTTLE, W. E. MAULE, W. R. HARDELL, D. A. HICKS, M. NGUYEN PHU, R. K. MONTOYE, W. T. GLOVER , AND S. DHAWAN [1989]. "IBM second-generation RISC machine organization," Proc. Int'l Conf. on Computer Design, IEEE (October) Rye, N.Y., 138–142.

BLOCH, E. [1959]. "The engineering design of the Stretch computer," *Proc. Fall Joint Computer Conf.*, 48–59.

BUCHOLTZ, W. [1962]. *Planning a Computer System: Project Stretch,* McGraw-Hill, New York.

CHARLESWORTH, A. E. [1981]. "An approach to scientific array processing: The architecture design of the AP-120B/FPS-164 family," *Computer* 14:12 (December) 12–30.

CHEN, T. C. [1980]. "Overlap and parallel processing" in *Introduction to Computer Architecture,* H. Stone, ed., Science Research Associates, Chicago, 427–486.

CLARK, D. W. [1987]. "Pipelining and performance in the VAX 8800 processor," *Proc. Second Conf. on Architectural Support for Programming Languages and Operating Systems,* IEEE/ACM (March), Palo Alto, Calif., 173–177.

COLWELL, R. P., R. P. NIX, J. J. O'DONNELL, D. B. PAPWORTH, AND B. K. RODMAN [1987]. "A VLIW architecture for a trace scheduling compiler," *Proc. Second Conf. on Architectural Support for Programming Languages and Operating Systems,* IEEE/ACM (March), Palo Alto, Calif., 180–192.

DAVIDSON, E. S. [1971]. "The design and control of pipelined function generators," *Proc. Conf. on Systems, Networks, and Computers,* IEEE (January), Oaxtepec, Mexico, 19–21.

DAVIDSON, E. S., A. T. THOMAS, L. E. SHAR, AND J. H. PATEL [1975]. "Effective control for pipelined processors," *COMPCON, IEEE* (March), San Francisco, 181–184.

DEHNERT, J. C., P. Y.-T. HSU, AND J. P. BRATT [1989]. "Overlapped loop support on the Cydra 5," *Proc. Third Conf. on Architectural Support for Programming Languages and Operating Systems* (April), IEEE/ACM, Boston, 26–39.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 373

DEROSA, J., R. GLACKEMEYER, AND T. KNIGHT [1985]. "Design and implementation of the VAX 8600 pipeline," *Computer* 18:5 (May) 38–48.

DIGITAL EQUIPMENT CORPORATION [1987]. *Digital Technical J.* 4 (March), Hudson, Mass. (This entire issue is devoted to the VAX 8800 processor.)

DITZEL, D. R. AND H. R. MCLELLAN [1987]. "Branch folding in the CRISP microprocessor: Reducing the branch delay to zero," *Proc. 14th Symposium on Computer Architecture* (June), Pittsburgh, 2–7.

EARLE, J. G. [1965]. "Latched carry-save adder," *IBM Technical Disclosure Bull.* 7 (March) 909–910.

ELLIS, J. R., [1986]. *Bulldog: A Compiler for VLIW Architectures*, The MIT Press,1986.

EMER, J. S. AND D. W CLARK [1984]. "A characterization of processor performance in the VAX-11/780," *Proc. 11th Symposium on Computer Architecture* (June), Ann Arbor, Mich., 301–310.

FISHER, J. A. [1981]. "Trace Scheduling: A Technique for Global Microcode Compaction," *IEEE Trans. on Computers* 30:7 (July), 478-490.

FISHER, J. A. [1983]. "Very long instruction word architectures and ELI-512," *Proc. Tenth Symposium on Computer Architecture* (June), Stockholm, Sweden., 140-150.

FISHER J. A., J. R. ELLIS, J. C. RUTTENBERG, AND A. NICOLAU [1984]. "Parallel processing: A smart compiler and a dumb machine," *Proc. SIGPLAN Conf. on Compiler Construction* (June), Palo Alto, CA, 11-16.

FOSTER, C. C. AND E. M. RISEMAN [1972]. "Percolation of code to enhance parallel dispatching and execution," *IEEE Trans. on Computers* C-21:12 (December) 1411–1415.

GIBBONS, P. B. AND S. S. MUCHNIK [1986]. "Efficient Instruction Scheduling for a Pipelined Processor," *SIGPLAN '86 Symposium on Compiler Construction, ACM* (June), Palo Alto, CA, 11-16.

GROSS, T. R. [1983]. *Code Optimization of Pipeline Constraints,* Ph.D. Thesis (December), Computer Systems Lab., Stanford Univ.

HENNESSY, J. L. AND T. R. GROSS [1983]. "Postpass code optimization of pipeline constraints," *ACM Trans. on Programming Languages and Systems* 5:3 (July) 422-448

HWU, W.-M. AND Y. PATT [1986]. "HPSm, a high performance restricted data flow architecture having minimum functionality," *Proc. 13th Symposium on Computer Architecture* (June), Tokyo, 297–307.

IBM [1990]. "The IBM RISC System/6000 processor," collection of papers, *IBM Jour. of Research and Development* 34:1, (January), 119 pages.

JOUPPI N. P. AND D. W. WALL [1989]. "Available instruction-level parallelism for superscalar and superpipelined machines," *Proc. Third Conf. on Architectural Support for Programming Languages and Operating Systems,* IEEE/ACM (April), Boston, 272–282.

KELLER R. M. [1975]. "Look-ahead processors," *ACM Computing Surveys* 7:4 (December) 177–195.

KOGGE, P. M. [1981]. *The Architecture of Pipelined Computers,* McGraw-Hill, New York.

KUNKEL, S. R. AND J. E. SMITH [1986]. "Optimal pipelining in supercomputers," *Proc. 13th Symposium on Computer Architecture* (June), Tokyo, 404–414.

LAM, M. [1988]. "Software pipelining: An effective scheduling technique for VLIW machines," *SIGPLAN Conf. on Programming Language Design and Implementation,* ACM (June), Atlanta, Ga., 318–328.

MCFARLING, S. AND J. HENNESSY [1986]. "Reducing the cost of branches," *Proc. 13th Symposium on Computer Architecture* (June), Tokyo, 396–403.

NICOLAU, A. AND J. A. FISHER [1984]. "Measuring the parallelism available for very long instruction work architectures," *IEEE Trans. on Computers* C-33:11 (November) 968–976.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 374

RAMAMOORTHY, C. V. AND H. F. LI [1977]. "Pipeline architecture," *ACM Computing Surveys* 9:1 (March) 61–102.

RYMARCZYK, J. [1982]. "Coding guidelines for pipelined processors," *Proc. Symposium on Architectural Support for Programming Languages and Operating Systems,* IEEE/ACM (March), Palo Alto, Calif., 12–19.

SITES, R. [1979]. *Instruction Ordering for the CRAY-1 Computer,* Tech. Rep. 78-CS-023 (July), Dept. of Computer Science, Univ. of Calif., San Diego.

SMITH, A. AND J. LEE [1984]. "Branch prediction strategies and branch target buffer design," *Computer* 17:1 (January) 6–22.

SMITH, J. E. [1981]. "A study of branch prediction strategies," *Proc. Eighth Symposium on Computer Architecture* (May), Minneapolis, 135–148.

SMITH, J. E. [1984]. "Decoupled access/execute computer architectures," *ACM Trans. on Computer Systems* 2:4 (November), 289–308.

SMITH, J. E. [1989]. "Dynamic instruction scheduling and the Astronautics ZS-1," *Computer* 22:7 (July) 21–35.

SMITH, J. E. AND A. R. PLEZKUN [1988]. "Implementing precise interrupts in pipelined processors," *IEEE Trans. on Computers* 37:5 (May) 562–573.

SMITH, J. E., G. E. DERMER, B. D. VANDERWARN, S. D. KLINGER, C. M. ROZEWSKI, D. L. FOWLER, K. R. SCIDMORE, J. P. LAUDON [1987]. "The ZS-1 central processor," *Proc. Second Conf. on Architectural Support for Programming Languages and Operating Systems,* IEEE/ACM (March), Palo Alto, Calif., 199–204.

SMITH, M. D., M. JOHNSON, AND M. A. HOROWITZ [1989]. "Limits on multiple instruction issue," *Proc. Third Conf. on Architectural Support for Programming Languages and Operating Systems,* IEEE/ACM (April), Boston, Mass., 290–302.

SOHI , G. S., AND S. VAJAPEYAM [1989]. "Tradeoffs in instruction format design for horizontal architectures," *Proc. Third Conf. on Architectural Support for Programming Languages and Operating Systems,* IEEE/ACM (April), Boston, Mass. 15–25.

THORLIN, J. F. [1967]. "Code generation for PIE (parallel instruction execution) computers," *Spring Joint Computer Conf.* (April), Atlantic City, N.J.

THORNTON, J. E. [1964]. "Parallel operation in the Control Data 6600," *Proc. Fall Joint Computer Conf.* 26, 33–40.

THORNTON, J. E. [1970]. *Design of a Computer, the Control Data 6600,* Scott, Foresman, Glenview, Ill.

TJADEN, G. S. AND M. J. FLYNN [1970]. "Detection and parallel execution of independent instructions," *IEEE Trans. on Computers* C-19:10 (October) 889–895.

TOMASULO, R. M. [1967]. "An efficient algorithm for exploiting multiple arithmetic units," *IBM J. of Research and Development* 11:1 (January) 25–33.

TROIANI, M., S. S. CHING, N. N. QUAYNOR, J. E. BLOEM, AND F. C. COLON OSORIO [1985]. "The VAX 8600 I Box, a pipelined implementation of the VAX architecture," *Digital Technical J.* 1 (August) 4–19.

WEISS, S. AND J. E. SMITH [1984]. "Instruction issue logic for pipelined supercomputers," *Proc. 11th Symposium on Computer Architecture* (June), Ann Arbor, Mich., 110–118.

WEISS, S. AND J. E. SMITH [1987]. "A study of scalar compilation techniques for pipelined supercomputers," *Proc. Second Conf. on Architectural Support for Programming Languages and Operating Systems* (March), IEEE/ACM, Palo Alto, Calif., 105–109.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 375

EXERCISES

6.1 [12/12/15/20/15/15] <6.2–6.4> Consider an architecture with two instruction formats: a register–register format and a register–memory format. There is a single memory addressing mode (offset + base register).

There is a set of ALU operations with format:

ALUop Rdest, $Rsrc_1$, $Rsrc_2$

or

ALUop Rdest, $Rsrc_1$, MEM

Where the ALUop is one of the following: Add, Subtract, And, Or, Load ($Rsrc_1$ ignored), Store (Rdest ignored). Rsrc or Rdest are registers. MEM is a base register and offset pair and is a source for any ALUop, except a store instruction where it is the destination.

Branches use a full compare of two registers and are PC-relative. Assume that this machine is pipelined so that a new instruction is started every clock cycle. The following pipeline structure—similar to that used in the VAX 8800 micropipeline—is used:

IF	RF	ALU1	MEM	ALU2	WB						
	IF	RF	ALU1	MEM	ALU2	WB					
		IF	RF	ALU1	MEM	ALU2	WB				
			IF	RF	ALU1	MEM	ALU2	WB			
				IF	RF	ALU1	MEM	ALU2	WB		
					IF	RF	ALU1	MEM	ALU2	WB	

The first ALU stage is used for effective address calculation for memory references and branches. The second ALU cycle is used for operations and branch comparison. RF is both a decode and register-fetch cycle. Assume reading in RF and writing in WB occur as in Figure 6.8 (page 262).

a. [12] Find the number of adders needed, counting any adder or incrementer; show a combination of instructions and pipe stages that justify this answer. You need only give one combination that maximizes the adder count.

b. [12] Find the number of register read and write ports and memory read and write ports required. Show that your answer is correct by showing a combination of instructions and pipeline stage indicating the instruction and the number of read ports and write ports required for that instruction.

c. [15] Determine any *data forwarding* between the two separate ALUs used for the ALU1 and ALU2 pipe stages. Put in all forwarding of ALU to ALU needed to avoid or reduce stalls. Show the relationship between the two instructions involved in forwarding.

d. [20] Show any other data-forwarding requirements for the units listed below by giving an example of the source instruction and destination instruction of the forwarding. Each example should show the maximum separation of the two instructions. How many instructions can each example forward across? You need only consider the following units: MDR_{in} (memory data in register), MDR_{out} (memory-data register for outgoing data), ALU_1, and ALU_2. Include any forwarding that is required to prevent or reduce stalls.

e. [15] Give an example of all remaining hazards after all forwarding of parts C and D above has been implemented. What is the maximum number of stalls for each hazard?

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 376

f. [15] Show all control hazard types by example and state the length of the stall. The control hazards should be resolved as early as possible (but not using a delayed branch).

6.2 [12] <6.1–6.4> A machine is called "underpipelined" if additional levels of pipelining can be added without changing the pipeline-stall behavior appreciably. Suppose that the DLX pipeline was changed to four stages by merging ID and EX and lengthening the clock cycle by 50%. How much faster would the conventional DLX pipeline be versus the underpipelined DLX on integer code only? Make sure you include the effect of any change in pipeline stalls using the data in Figure 6.24 (page 278).

6.3 [15] <6.2–6.4> We know that a four-deep pipelined implementation has the following hazard frequencies and stall requirements between an instruction i and its successors:

$i + 1$ (and not on $i + 2$)	20%	2 cycle stall
$i + 2$	5%	1 cycle stall

Assume that the clock rate of the pipelined machine is four times the clock rate of the nonpipelined implementation. What is the effective performance increase from pipelining if we ignore the effect of hazards? What is the effective performance increase from pipelining if we account for the effect of pipelining hazards?

6.4 [15] <6.3> Suppose the branch frequencies (as percentages of all instructions) are as follows:

Conditional branches	20%
Jumps and calls	5%
Conditional branches	60% are taken

We are examining a four-deep pipeline where the branch is resolved at the end of the second cycle for unconditional branches, and at the end of the third cycle for conditional branches. Assuming that only the first pipe stage can always be done independent of whether the branch goes and ignoring other pipeline stalls, how much faster would the machine be without any branch hazards?

6.5 [20] <6.4> Several designers have proposed the concept of canceling branches (also called squashing or nullifying), as a way to improve the performance of delayed branches. (Several of the machines discussed in Appendix E have this capability.) The idea is to allow the branch to indicate that the instruction in the delay slot should be aborted if the branch is mispredicted. The advantage of canceling branches is that the delay slot can **always** be filled, since the branch can abort the contents of the delay slot if mispredicted. The compiler need not worry about whether the instruction is OK to execute when the branch is mispredicted.

A simple version of canceling branches cancels if the branch is not taken; assume this type of canceling branch. Use the data in Figure 6.18 (page 272) for branch frequency. Assume that 27% of the branch-delay slots are filled using strategy (a) of Figure 6.20 (page 274) with standard delayed branches, and that the rest of the slots are filled using canceling branches and strategy (b). Using the taken/not taken data for Spice from Figure 3.22 on page 107, show the effectiveness of this scheme with canceling branches for Spice using the same format as the graph in Figure 6.22 (page 276). How much faster on Spice would a machine with canceling branches run, assuming there is no clock-speed penalty compared to a machine with only delayed branches? Assume CPI without branch stalls is 1.

6.6 [20/15/20] <6.2–6.4> Suppose that we have the following pipeline layout:

Stage	Function
1	Instruction fetch
2	Operand decode
3	Execution or memory access (branch resolution)

All data dependences are between the register written in Stage 3 of instruction i and a register read in Stage 2 of instruction $i + 1$, before instruction i has completed. The probability of such an interlock occurring is $1/p$.

We are considering a change in the machine organization that would write back the result of an instruction during an effective 4th pipe stage. This would decrease the length of the clock cycle by d (i.e., if the length of the clock cycle was T, it is now T–d). The probability of a dependence between instruction i and instruction $i+2$ is p^{-2}. (Assume that the value of p^{-1} excludes instructions that would interlock on $i+2$.) The branch would also be resolved during the fourth stage.

a. [20] Considering only the data hazard, find the lower bound on d that makes this a profitable change. Assume that each result has exactly one use and that the basic clock cycle has length T.

b. [15] Suppose that the probability of an interlock between i and $i+n$ were $0.3 - 0.1n$ for $1 \le n \le 3$. What increase in the clock rate is needed so that this change improves performance?

c. [20] Now assume that we have used forwarding to eliminate the extra hazard introduced by the change. That is, for all *data* hazards the pipeline length is *effectively* 3. This design may still not be worthwhile because of the impact of control hazards coming from a four-stage versus a three-stage pipeline. Assume that only Stage 1 of the pipeline can be safely executed before we decide whether a branch goes or not and that all branches are conditional. We want to know what the impact of branch hazards can be before this longer pipeline does not yield high performance. Find an upper bound on the percent of conditional branches in programs in terms of the ratio of d to the original clock-cycle time, so that the longer pipeline has better performance. What if d is a 10% reduction, what is the maximum percentage of conditional branches, before we lose with this longer pipeline? Assume the taken-branch frequency for conditional branches is 60%.

6.7 [12] <6.7> A shortcoming of the scoreboard approach occurs when multiple functional units that share input buses are waiting for a single result. The units cannot start simultaneously, but must serialize. This is not true in Tomasulo's algorithm. Give a code sequence that uses no more than 10 instructions and shows this problem. Use the FP latencies from Figure 6.29 (page 289) and the same functional units in both examples. Indicate where the Tomasulo approach can continue, but the scoreboard approach must stall.

6.8 [15] <6.7> Tomasulo's algorithm also has a disadvantage versus the scoreboard: only one result can complete per clock, due to the CDB. Using the FP latencies from Figure 6.29 (page 289) and the same functional units in both cases, find a code sequence of no more than 10 instructions where scoreboard does not stall, but Tomasulo's algorithm must. Indicate where this occurs in your sequence.

6.9 [15] <6.7> Suppose we have a deeply pipelined machine, for which we implement a branch-target buffer for the conditional branches only. Assume that the misprediction

penalty is always 4 cycles and the buffer miss penalty is always 3 cycles. Assume 90% hit rate and 90% accuracy, and the branch statistics in Figure 6.18 (page 272). How much faster is the machine with the branch-target buffer versus a machine that has a fixed 2-cycle branch penalty? Assume a base CPI without branch stalls of 1.

6.10 [15] <6.7> Some designers have proposed using branch-target buffers to obtain a zero-delay unconditional branch (see Ditzel and McLellan [1987]). The buffer simply caches the target instruction rather than the target PC. On an unconditional branch that hits in the branch-target buffer, the target instruction is fetched and sent to the pipeline in place of the unconditional branch. Assuming a 90% hit rate, a base CPI of 1, and the data in Figure 6.18 (page 272), how much improvement is gained by this enhancement versus a machine whose effective CPI is 1.1.

6.11–6.19 For these problems we will look at how a common vector loop runs on a variety of pipelined versions of DLX. The loop is the so-called SAXPY loop (discussed extensively in Chapter 7). The loop implements the vector operation Y = a*X+Y for a vector of length 100. Here is the DLX code for the loop:

```
foo:  LD     F2,0(R1)    ;load X(i)
      MULTD  F4,F2,F0    ;multiply a*X(i)
      LD     F6,0(R2)    ;load Y(i)
      ADDD   F6,F4,F6    ;add aX(i) + Y(i)
      SD     0(R2),F6    ;store Y(i)
      ADDI   R1,R1,8     ;increment X index
      ADDI   R2,R2,8     ;increment Y index
      SGTI   R3,R1,done  ;test if done
      BEQZ   R3,foo      ; loop if not done
```

For these problems, assume that the integer operations issue and complete in one clock cycle and that their results are fully bypassed. Ignore the branch delay. You will use the FP latencies shown in Figure 6.29 (page 289) unless stated otherwise. Assume the FP units are not pipelined unless the problem states otherwise.

6.11 [20] <6.2–6.6> For this problem use the pipeline constraints shown in Figure 6.29 (page 289). Show the number of stall cycles for each instruction and what clock cycle the instruction begins execution (i.e., enters its first EX cycle) on the first iteration of the loop. How many clock cycles does each loop iteration take?

6.12 [22] <6.7> Using the DLX code for SAXPY above, show the state of the scoreboard tables (as in Figure 6.32) when the `SGTI` instruction reaches Write result. Assume that issue and read operands each take a cycle. Assume that there are three integer functional units and they take only a single execution cycle (including loads and stores). Assume the functional unit count described in Section 6.7 with the FP latencies of Figure 6.29. The branch should not be included in the scoreboard.

6.13 [22] <6.7> Use the DLX code for SAXPY above and the latencies of Figure 6.29. Assuming Tomasulo's algorithm for the hardware with the functional units described in Section 6.7, show the state of the reservation stations and register-status tables (as in

Figure 6.37) when the SGTI writes its result on the CDB. Make the same assumptions about latencies and functional units as Exercise 6.12.

6.14 [22] <6.7> Using the DLX code for SAXPY above, assume a scoreboard with the functional units described in the algorithm for the hardware, plus three integer functional units (also used for load/store). Assume the following latencies in clock cycles:

FP multiply	10
FP add	6
FP load/store	2
All integer operations	1

Show the state of the scoreboard (as in Figure 6.32) when the branch issues for the second time. Assume the branch was correctly predicted taken and took one cycle. How many clock cycles does each loop iteration take? You may ignore any register port/bus conflicts.

6.15 [25] <6.7> Use the DLX code for SAXPY above. Assume Tomasulo's algorithm for the hardware using the functional-unit count shown in Section 6.7. Assume the following latencies in clock cycles:

FP multiply	10
FP add	6
FP load/store	2
All integer operations	1

Show the state of the reservation stations and register status tables (as in Figure 6.37) when the branch is executed for the second time. Assume the branch was correctly predicted as taken. How many clock cycles does each loop iteration take?

6.16 [22] <6.8> Unwind the DLX code for SAXPY three times and schedule it for the standard DLX pipeline. Assume the FP latencies of Figure 6.29. When unwinding, you should optimize the code as in Section 6.8. Significant reordering of the code will be needed to maximize performance. What is the speedup over the original loop?

6.17 [25] <6.8> Assume a superscalar architecture that can issue any two independent operations in a clock cycle (including two integer operations). Unwind the DLX code for SAXPY three times and schedule it assuming the FP latencies of Figure 6.29. Assume one fully-pipelined copy of each functional unit (e.g., FP adder, FP multiplier). How many clock cycles will each iteration on the original code take? When unwinding, you should optimize the code as in Section 6.8. What is the speedup versus the original code?

6.18 [25] <6.8> In a superpipelined machine, rather than have multiple functional units, we would fully pipeline all the units. Suppose we designed a superpipelined DLX that had twice the clock rate of our standard DLX pipeline and could issue any two unrelated operations in the same time that the normal DLX pipeline issued one operation. Unroll the DLX SAXPY code three times and schedule it for this superpipelined machine assuming the FP latencies of Figure 6.29. How many clock cycles does each loop iteration take? Remember that these clock cycles are half as long as those on a standard DLX pipeline or a superscalar DLX.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 380

6.19 [20] <6.8> Start with the SAXPY code and the machine used in Figure 6.49. Unroll the SAXPY loop three times, performing simple optimizations (as on page 315). Fill in a table like Figure 6.49 for the unrolled loop. How many clock cycles does each loop iteration take?

6.20 [35] <6.1–6.4> Change the DLX instruction simulator to be pipelined. Measure the frequency of empty branch-delay slots, the frequency of load delays, and the frequency of FP stalls for a variety of integer and FP programs. Also, measure the frequency of forwarding operations. Determine what the performance impact of eliminating forwarding and stalling would be.

6.21 [35] <6.6> Using a DLX simulator, create a DLX pipeline simulator. Explore the impact of lengthening the FP pipelines, assuming both fully pipelined and nonpipelined FP units. How does clustering of FP operations affect the results? Which FP units are most susceptible to changes in the FP pipeline length?

6.22 [40] <6.4–6.6> Write an instruction scheduler for DLX that works on DLX assembly language. Evaluate your scheduler using either profiles of programs or with a pipeline simulator. If the DLX C compiler does optimization, evaluate your scheduler's performance both with and without optimization.

6.23 [35] <6.4–6.6> Write a DLX pipeline simulator that uses Tomasulo's algorithm with the functional units described. Evaluate the performance of this machine compared to the straightforward DLX pipeline.

6.24 [Discussion] <6.7> Dynamic instruction scheduling requires a considerable investment in hardware. In return, this capability allows the hardware to run programs that could not be run at full speed with only compile-time, static scheduling. What tradeoffs should be taken into account in trying to decide between a dynamically and a statically scheduled scheme? What sort of situations in both hardware technology and program characteristics are likely to favor one approach or the other?

6.25 [Discussion] <6.7> There is a subtle problem that must be considered when implementing Tomasulo's algorithm. It might be called the "two ships passing in the night problem." What happens if an instruction is being passed to a reservation station during the same clock period as one of its operands is going onto the common data bus? Before an instruction is in a reservation station, the operands are fetched from the register file; but once it is in the station, the operands are always obtained from the CDB. Since the instruction and its operand tag are in transit to the reservation station, the tag cannot be matched against the tag on the CDB. So there is a possibility that the instruction will then sit in the reservation station forever waiting for its operand, which it just missed. How might this problem be solved? You might consider subdividing one of the steps in the algorithm into multiple parts. (This intriguing problem is courtesy of J. E. Smith.)

6.26 [Discussion] <6.8> Discuss the advantages and disadvantages of a superscalar implementation, a superpipelined implementation, and a VLIW approach in the context of DLX. What levels of instruction-level parallelism favor each approach? What other concerns would you consider in choosing which type of machine to build?

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 381

I'm certainly not inventing vector machines. There are three kinds that I know of existing today. They are represented by the Illiac-IV, the (CDC) Star machine, and the TI (ASC) machine. Those three were all pioneering machines. . . . One of the problems of being a pioneer is you always make mistakes and I never, never want to be a pioneer. It's always best to come second when you can look at the mistakes the pioneers made.

Seymour Cray, Public Lecture at Lawrence Livermore Laboratories on the Introduction of the CRAY-1 (1976)

7.1	**Why Vector Machines?**	**351**
7.2	**Basic Vector Architecture**	**353**
7.3	**Two Real-World Issues: Vector Length and Stride**	**364**
7.4	**A Simple Model for Vector Performance**	**369**
7.5	**Compiler Technology for Vector Machines**	**371**
7.6	**Enhancing Vector Performance**	**377**
7.7	**Putting It All Together: Evaluating the Performance of Vector Processors**	**383**
7.8	**Fallacies and Pitfalls**	**390**
7.9	**Concluding Remarks**	**392**
7.10	**Historical Perspective and References**	**393**
	Exercises	**397**

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.

7 Vector Processors

7.1 Why Vector Machines?

In the last chapter we looked at pipelining in detail and saw that pipeline scheduling, issuing multiple instructions per clock cycle, and more deeply pipelining a processor could as much as double the performance of a machine. Yet there are limits on the performance improvement that pipelining can achieve. These limits are set by two primary factors:

- Clock cycle time—The clock cycle time can be decreased by making the pipelines deeper, but a deeper pipeline will increase the pipeline dependences and result in a higher CPI. At some point, each increase in pipeline depth has a corresponding increase in CPI. As we saw in Section 6.10, very deep pipelining can slow down a processor.
- Instruction fetch and decode rate—This limitation, sometimes called the *Flynn bottleneck* (based on Flynn [1966]), prevents fetching and issuing of more than a few instructions per clock cycle. We saw that for most pipelined machines the average number of instruction issues per clock was less than one.

The dual limitations imposed by deeper pipelines and issuing multiple instructions can be viewed from the standpoint of either clock rate or CPI: It is just as

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 383

difficult to schedule a pipeline that is n times deeper as it is to schedule a machine that issues n instructions per clock cycle.

High-speed, pipelined machines are particularly useful for large scientific and engineering applications. A high-speed pipelined machine will usually use a cache to avoid forcing memory reference instructions to have very long latency. However, big, long-running, scientific programs often have very large active data sets that are often accessed with low locality, yielding poor performance from the memory hierarchy. The resulting impact is a decrease in cache performance. This problem could be overcome by not caching these structures if it were possible to determine the memory-access patterns and pipeline the accesses efficiently. Compiler assistance may help address this problem in the future (see Section 10.7).

Vector machines provide high-level operations that work on *vectors*—linear arrays of numbers. A typical vector operation might add two 64-entry, floating-point vectors to obtain a single 64-entry vector result. The vector instruction is equivalent to an entire loop, with each iteration computing one of the 64 elements of the result, updating the indices, and branching back to the beginning.

Vector operations have several important properties that solve most of the problems mentioned above:

- The computation of each result is independent of the computation of previous results, allowing a very deep pipeline *without* generating any data hazards. Essentially, the absence of data hazards was determined by the compiler or programmer when they decided that a vector instruction could be used.
- A single vector instruction specifies a great deal of work—it is equivalent to executing an entire loop. Thus, the instruction bandwidth requirement is reduced, and the Flynn bottleneck is considerably mitigated.
- Vector instructions that access memory have a known access pattern. If the vector's elements are all adjacent, then fetching the vector from a set of heavily interleaved memory banks works very well. The high latency of initiating a main memory access versus accessing a cache is amortized because a single access is initiated for the entire vector rather than to a single word. Thus, the cost of the latency to main memory is seen only once for the entire vector, rather than once for each word of the vector.
- Because an entire loop is replaced by a vector instruction whose behavior is predetermined, control hazards that would normally arise from the loop branch are nonexistent.

For these reasons, vector operations can be made faster than a sequence of scalar operations on the same number of data items, and designers are motivated to include vector units if the applications domain can use them frequently.

As mentioned above, vector machines pipeline the operations on the individual elements of a vector. The pipeline includes not only the arithmetic operations (multiplication, addition, and so on), but also memory accesses and effective

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 384

address calculations. In addition, most high-end vector machines allow multiple vector operations to be done at the same time, creating parallelism among the operations on different elements. In this chapter, we focus on vector machines that gain performance by pipelining and instruction overlap. In Chapter 10, we discuss parallel machines that operate on many elements in parallel rather than in pipelined fashion.

7.2 Basic Vector Architecture

A vector machine typically consists of an ordinary pipelined scalar unit plus a vector unit. All functional units within the vector unit have a latency of several clock cycles. This allows a shorter clock cycle time and is compatible with long-running, vector operations that can be deeply pipelined without generating hazards. Most vector machines allow the vectors to be dealt with as floating-point numbers (FP), as integers, or as logical data, though we will focus on floating point. The scalar unit is basically no different from the type of pipelined CPU discussed in Chapter 6.

There are two primary types of vector architectures: vector-register machines and memory–memory vector machines. In a *vector-register machine*, all vector operations—except load and store—are among the vector registers. These machines are the vector counterpart of a load/store architecture. All major vector machines being shipped in 1990 use a vector-register architecture; these include the Cray Research machines (CRAY-1, CRAY-2, X-MP, and Y-MP), the Japanese supercomputers (NEC SX/2, Fujitsu VP200, and the Hitachi S820), and the mini-supercomputers (Convex C-1 and C-2). In a *memory–memory vector machine* all vector operations are memory to memory. The first vector machines were of this type, as were CDC's machines. From this point on we will focus on vector-register architectures only; we will briefly return to memory–memory vector architectures at the end of the chapter (Section 7.8) to discuss why they have not been as successful as vector-register architectures.

We begin with a vector-register machine consisting of the primary components shown in Figure 7.1 (page 354). This machine, which is loosely based on the CRAY-1, is the foundation for discussion throughout most of this chapter. We will call it DLXV; its integer portion is DLX, and its vector portion is the logical vector extension of DLX. The rest of this section examines how the basic architecture of DLXV relates to other machines.

The primary components of the instruction set architecture of DLXV are:

- Vector registers—Each vector register is a fixed-length bank holding a single vector. DLXV has eight vector registers, and each vector register holds 64 doublewords. Each vector register must have at least two read ports and one write port in DLXV. This will allow a high degree of overlap among vector operations to different vector registers. (The CRAY-1 manages to implement the register file with only a single port per register using some clever implementation techniques.)

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 385

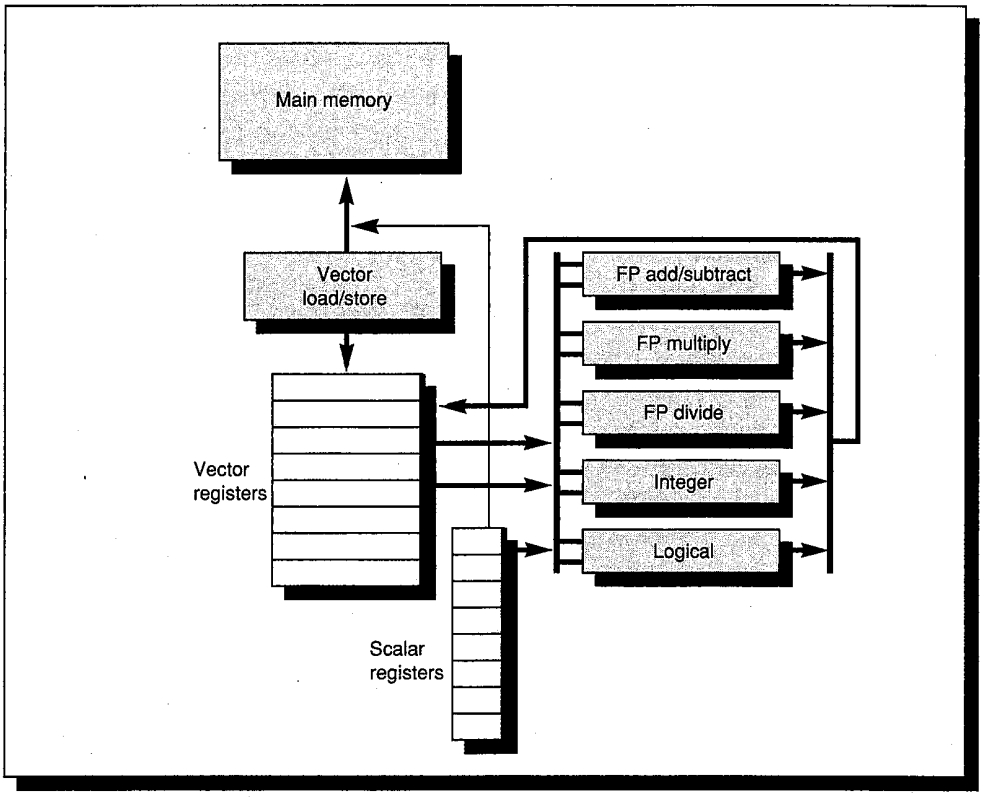

FIGURE 7.1 The basic structure of a vector-register architecture, DLXV. This machine has a scalar architecture just like DLX. There are also eight 64-element vector registers, and all the functional units are vector functional units. Special vector operations and vector loads and stores are defined. We show vector units for logical and integer operations. These are included so that DLXV looks like a standard vector machine, which usually includes these units. However, we will not be discussing these units except in the Exercises. In Section 7.6 we add chaining, which will require additional interconnect capability.

- Vector functional units—Each unit is fully pipelined and can start a new operation on every clock cycle. A control unit is needed to detect hazards, both on conflicts for the functional units (structural hazards) and on conflicts for register accesses (data hazards). DLXV has five functional units, as shown in Figure 7.1. For simplicity, we will focus exclusively on the floating-point functional units.
- Vector load/store unit—A vector memory unit that loads or stores a vector to or from memory. The DLXV vector loads and stores are fully pipelined, so that words can be moved between the vector registers and memory with a bandwidth of one word per clock cycle, after an initial latency.
- A set of scalar registers—These can also provide data as input to the vector functional units, as well as compute addresses to pass to the vector load/store unit. These are the normal 32 general-purpose registers and 32 floating-point registers of DLX.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 386

Figure 7.2 shows the characteristics of some typical vector machines, including the size and count of the registers, the number and types of functional units, and the number of load/store units.

In DLXV, the vector operation has the same name as the DLX name with the letter "V" appended. These are double-precision, floating-point, vector operations. (We have omitted single-precision FP operations and integer and logical operations for simplicity.) Thus, `ADDV` is an add of two double-precision vectors. The vector operations take as their input either a pair of vector registers (`ADDV`) or a vector register and a scalar register designated by appending "SV" (`ADDSV`). In the latter case, the value in the scalar register is used as the input for all operations—the operation `ADDSV` will add the contents of a scalar register to each element in a vector register. Vector operations always have a vector destination register. The names `LV` and `SV` denote vector load and vector store, and load or store an entire vector of double-precision data. One operand is

Machine	Year announced	Vector registers	Elements per vector register (64-bit elements)	Vector functional units	Vector load / store units
CRAY-1	1976	8	64	6: add, multiply, reciprocal, integer add, logical, shift	1
CRAY X-MP CRAY Y-MP	1983 1988	8	64	8: FP add, FP multiply, FP reciprocal, integer add, 2 logical, shift, population count/parity	2 loads 1 store
CRAY-2	1985	8	64	5: FP add, FP multiply, FP reciprocal/sqrt, integer (add shift, population count), logical	1
Fujitsu VP100/200	1982	8–256	32–1024	3: FP or integer add/logical, multiply, divide	2
Hitachi S810/820	1983	32	256	4: 2 integer add/logical, 1 multiply-add and 1 multiply/divide–add unit	4
Convex C-1	1985	8	128	4: multiply, add, divide, integer/logical	1
NEC SX/2	1984	8 + 8192	256 variable	16: 4 integer add/logical, 4 FP multiply/divide, 4 FP add, 4 shift	8
DLXV	1990	8	64	5: multiply, divide, add, integer add, logical	1

FIGURE 7.2 Characteristics of several vector-register architectures. The vector functional units include all operation units used by the vector instructions. The functional units are floating point unless stated otherwise. If the machine is a multiprocessor, the entries correspond to the characteristics of one processor. Each vector load/store unit represents the ability to do an independent, overlapped transfer to or from the vector registers. The Fujitsu VP200's vector registers are configurable: The size and count of the 8K 64-bit entries may be varied inversely to one another (e.g., 8 registers each 1K elements long, or 128 registers each 64 elements long). The NEC SX/2 has 8 fixed registers of length 256, plus 8K of configurable 64-bit registers. The reciprocal unit on the CRAY machines is used to do division (and square root on the CRAY-2). Add pipelines perform floating-point add and subtract. The multiply/divide–add unit on the Hitachi S810/200 performs an FP multiply or divide followed by an add or subtract (while the multiply-add unit performs a multiply followed by an add or subtract). Note that most machines use the vector FP multiply and divide units for vector integer multiply and divide, just like DLX, and several of the machines use the same units for FP scalar and FP vector operations.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 387

the vector register to be loaded or stored; the other operand, which is a DLX general-purpose register, is the starting address of the vector in memory. Figure 7.3 lists the DLXV vector instructions. In addition to the vector registers, we need two additional special-purpose registers: the vector-length and vector-mask registers. We will discuss these registers and their purpose in Sections 7.3 and 7.6, respectively.

Vector instruction	Operands	Function
ADDV	V1,V2,V3	Add elements of V2 and V3, then put each result in V1.
ADDSV	V1,F0,V2	Add F0 to each element of V2, then put each result in V1.
SUBV	V1,V2,V3	Subtract elements of V3 from V2, then put each result in V1.
SUBVS	V1,V2,F0	Subtract F0 from elements of V2, then put each result in V1.
SUBSV	V1,F0,V2	Subtract elements of V2 from F0, then put each result in V1.
MULTV	V1,V2,V3	Multiply elements of V2 and V3, then put each result in V1.
MULTSV	V1,F0,V2	Multiply F0 by each element of V2, then put each result in V1.
DIVV	V1,V2,V3	Divide elements of V2 by V3, then put each result in V1.
DIVVS	V1,V2,F0	Divide elements of V2 by F0, then put each result in V1.
DIVSV	V1,F0,V2	Divide F0 by elements of V2, then put each result in V1.
LV	V1,R1	Load vector register V1 from memory starting at address R1.
SV	R1,V1	Store vector register V1 into memory starting at address R1.
LVWS	V1,(R1,R2)	Load V1 from address at R1 with stride in R2, i.e., R1+i*R2.
SVWS	(R1,R2),V1	Store V1 from address at R1 with stride in R2, i.e., R1+i*R2.
LVI	V1,(R1+V2)	Load V1 with vector whose elements are at R1+V2(i), i.e., V2 is an index.
SVI	(R1+V2),V1	Store V1 with vector whose elements are at R1+V2(i), i.e., V2 is an index.
CVI	V1,R1	Create an index vector by storing the values 0,1*R1,2*R1,...,63*R1 into V1.
S_V S_SV	V1,V2 F0,V1	Compare (EQ, NE, GT, LT, GE, LE) the elements in V1 and V2. If condition is true put a 1 in the corresponding bit vector; otherwise put 0. Put resulting bit vector in vector-mask register (VM). The instruction S_SV performs the same compare but using a scalar value as one operand.
POP	R1,VM	Count the 1s in the vector-mask register and store count in R1.
CVM		Set the vector-mask register to all 1s.
MOVI2S	VLR,R1	Move contents of R1 to the vector-length register.
MOVS2I	R1,VLR	Move the contents of the vector-length register to R1.
MOVF2S	VM,F0	Move contents of F0 to the vector-mask register.
MOVS2F	F0,VM	Move contents of vector-mask register to F0.

FIGURE 7.3 The DLXV vector instructions. Only the double-precision FP operations are shown. In addition to the vector registers there are two special registers VLR (discussed in Section 7.3) and VM (discussed in Section 7.6). The operations with stride are explained in Section 7.3, and the use of the index creation and indexed load/store operations are explained in Section 7.6.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 388

A vector machine is best understood by looking at a vector loop on DLXV. Let's take a typical vector problem, which will be used throughout this chapter:

```
Y = a * X + Y
```

`X` and `Y` are vectors, initially resident in memory, and `a` is a scalar. This is the so-called SAXPY or DAXPY (Single-precision or Double-precision A*X Plus Y) loop that forms the inner loop of the Linpack benchmark. Linpack is a collection of linear algrebra routines; the Gaussian elimination portion of Linpack is the segment used as a benchmark. SAXPY represents a small piece of the program, though it takes most of the time in the benchmark.

For now, let us assume that the number of elements, or length, of a vector register (64) matches the length of the vector operation we are interested in. (This restriction will be lifted shortly.)

Example

Show the code for DLX and DLXV for the DAXPY loop. Assume that the starting addresses of X and Y are in `Rx` and `Ry`, respectively.

Answer

Here is the DLX code.

```
      LD      F0,a
      ADDI    R4,Rx,#512   ;last address to load
loop:
      LD      F2,0(Rx)     ;load X(i)
      MULTD   F2,F0,F2     ;a*X(i)
      LD      F4,0(Ry)     ;load Y(i)
      ADDD    F4,F2,F4     ;a*X(i) + Y(i)
      SD      F4,0(Ry)     ;store into Y(i)
      ADDI    Rx,Rx,#8     ;increment index to X
      ADDI    Ry,Ry,#8     ;increment index to Y
      SUB     R20,R4,Rx    ;compute bound
      BNZ     R20,loop     ;check if done
```

Here is the code for DLXV for DAXPY.

```
      LD      F0,a         ;load scalar a
      LV      V1,Rx        ;load vector X
      MULTSV  V2,F0,V1     ;vector-scalar multiply
      LV      V3,Ry        ;load vector Y
      ADDV    V4,V2,V3     ;add
      SV      Ry,V4        ;store the result
```

There are some interesting comparisons between the two code segments in the example above. The most dramatic is that the vector machine greatly reduces the dynamic instruction bandwidth, executing only 6 instructions versus almost 600 for DLX. This reduction occurs both because the vector operations work on

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 389

64 elements, and because the overhead instructions that constitute nearly half the loop on DLX are not present in the DLXV code.

Another important difference is the frequency of pipeline interlocks. In the straightforward DLX code every `ADDD` must wait for a `MULTD`, and every `SD` must wait for the `ADDD`. On the vector machine, each vector instruction operates on all the vector elements independently. Thus, pipeline stalls are required only once per vector operation, rather than once per vector element. In this example, the pipeline-stall frequency on DLX will be about 64 times higher than it is on DLXV. The pipeline stalls can be eliminated on DLX by using software pipelining or loop unrolling (as we saw in Chapter 6, Section 6.8). However, the large difference in instruction bandwidth cannot be reduced.

Vector Start-up Time and Initiation Rate

Let's investigate the running time of this vector code on DLXV. The running time of each vector operation in the loop has two components—the *start-up time* and the *initiation rate*. The start-up time comes from the pipelining latency of the vector operation and is principally determined by how deep the pipeline is for the functional unit used. For example, a latency of 10 clock cycles means both that the operation takes 10 clock cycles and that the pipeline is 10 deep. (In discussions of the performance of vector operations, clock cycles are customarily used as the metric.) The initiation rate is the time per result once a vector instruction is running; this rate is usually one per clock cycle for individual operations, though some supercomputers have vector operations that can produce 2 or more results per clock, and others have units that may not be fully pipelined. The *completion rate* must at least equal the initiation rate—otherwise there is no place to put results. Hence, the time to complete a single vector operation of length n is:

$$\text{Start-up time} + n * \text{Initiation rate}$$

Example

Suppose the start-up time for a vector multiply is 10 clock cycles. After start-up the initiation rate is one per clock cycle. What is the number of clock cycles per result (i.e., one element of the vector) for a 64-element vector?

Answer

$$\text{Clock cycles per result} = \frac{\text{Total time}}{\text{Vector length}}$$

$$= \frac{\text{Start-up time} + 64 * \text{Initiation rate}}{64}$$

$$= \frac{10 + 64}{64} = 1.16 \text{ clock cycles.}$$

Figure 7.4 shows the effect of start-up time and initiation rate on vector performance. The effect of increasing start-up time on a slow-running vector is

small, while the same increase in start-up time on a system with an initiation rate of one per clock decreases performance by a factor of nearly two.

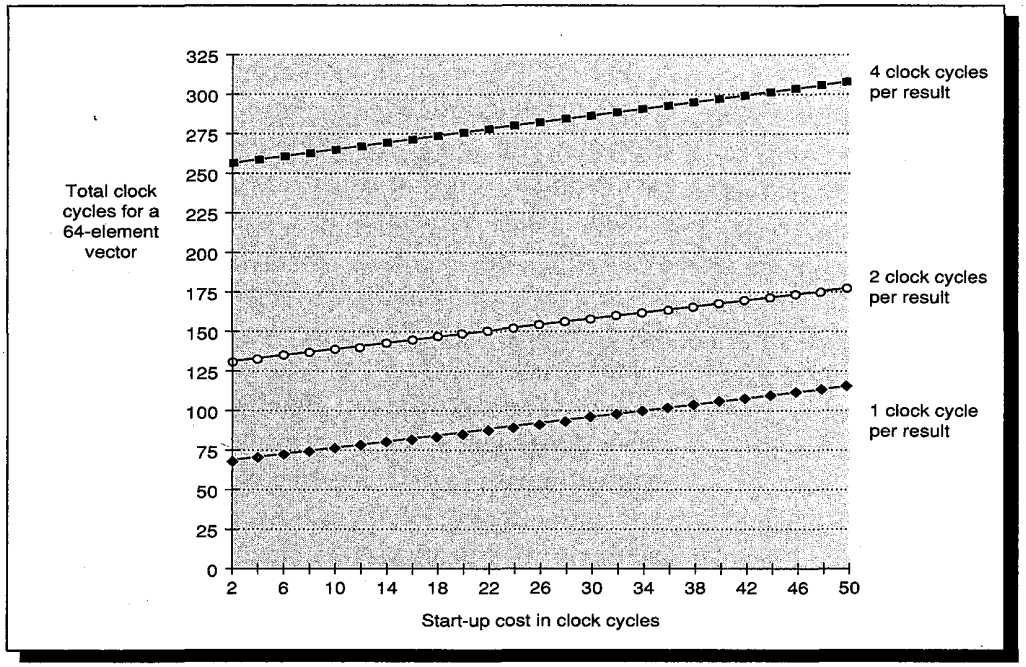

FIGURE 7.4 Total running time increases with start-up cost from 2 to 50 clock cycles per operation on the x axis. The impact of start-up time is much greater for fast-running than for slow-running vectors. The operation running at one clock cycle per result increases its run time by 75%, while the operation running at four clock cycles per result increases by less than 20%.

What determines the start-up and initiation rates? Let's first consider the operations that do not involve a memory access. For register–register operations the start-up time (in clock cycles) is equal to the depth of the functional unit pipeline, since this is the time to get the first result. In the earlier example, the depth of 10 gave a start-up time of 10 clock cycles. In the next few sections, we will see that there are other costs involved that increase the start-up time. The initiation rate is determined by how often the corresponding vector functional unit can accept an operand. If it is fully pipelined, then it can start an operation on new operands every clock cycle, yielding an initiation rate of one per clock (as in the earlier example).

Start-up time for an operation comprises the total latency for the functional unit implementing that operation. If the initiation rate is to be kept at 1 clock per result, then

$$\text{Pipeline depth} = \left\lceil \frac{\text{Total functional unit time}}{\text{Clock cycle time}} \right\rceil$$

For example, if an operation takes 10 clock cycles, it must be pipelined 10 deep to achieve an initiation rate of one per clock. Pipeline depth, then, is determined

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 391

by the complexity of the operation and the clock cycle time of the machine. The pipeline depths of functional units vary widely—from 2 to 20 stages is not uncommon—though the most heavily used units have start-up times of 4 to 8 clocks.

For DLXV, we will choose the same pipeline depths as the CRAY-1. All functional units are fully pipelined. Pipeline depths are six clock cycles for floating-point add and seven clock cycles for floating-point multiply. If a vector computation depends on an uncompleted computation and will need to be stalled, it adds an extra 4-clock-cycle start-up penalty. This penalty is typical on vector machines and arises due to the lack of bypassing: the penalty is the time to write and then read the operands and is only seen when there is a dependence. Thus, back-to-back dependent vector operations will see the full latency of a vector operation. On DLXV, as on most vector machines, independent vector operations using different functional units can issue without any penalty or delay. Independent vector operations may also be fully overlapped, and each instruction issue only takes one clock. Thus, when the operations are independent and different, DLXV can overlap vector operations, just as DLX can overlap integer and floating-point operations.

Because DLXV is fully pipelined, the initiation rate for a vector instruction is always 1. However, a sequence of vector operations will not be able to run at that rate, due to start-up costs. The term *sustained rate* is applied to this situation and refers to the time per element for a collection of related vector operations. Here an element is not the result of a single vector operation, but one result of a series of vector operations. The time per element, then, is the time required for each operation to produce an element. For example, in the SAXPY loop, the sustained rate will be the time to compute and store one element of the result vector Y.

Example

For a vector length of 64 on DLXV and the following two vector instructions, what is the sustained rate for the sequence, and the effective number of floating-point operations per clock for the sequence?

```
MULTV V1,V2,V3
ADDV  V4,V5,V6
```

Answer

Let's look at the start and completion times of these independent operations (remember that the start-up times are 7 cycles for multiply and 6 cycles for add):

Operation	Start	Complete
MULTV	0	7 + 64 = 71
ADDV	1	1 + 6 + 64 =71

The sustained rate is one element per clock—remember that sustained rate requires all vector operations to produce a result. The sequence executes 128

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 392

FLOPs (FLoating-point OPerations) in 71 clock cycles, for a rate of 1.8 FLOPs per clock. A vector machine can sustain a throughput of more than one operation per clock cycle by issuing independent vector operations to different vector functional units.

The behavior of the load/store vector unit is significantly more complicated. The start-up time for a load is the time to get the first word from memory into a register. If the rest of the vector can be supplied without stalling, then the vector initiation rate is equal to the rate at which new words are fetched or stored. Typically, penalties for start-ups on load/store units are higher than for functional units—up to 50 clock cycles on some machines. For DLXV we will assume a low start-up time of 12 clock cycles, since the CRAY-1 and CRAY X-MP have load/store start-up times of between 9 and 17 clock cycles. For stores, we will not usually care about the start-up time, since stores do not directly produce results. However, when an instruction must wait for a store to complete (as a load might have to with only one memory pipeline), the load may see part or all of the 12-cycle latency of a store. Figure 7.5 summarizes the start-up penalties for DLXV vector operations.

Operation	Start-up penalty
Vector add	6
Vector multiply	7
Vector divide	20
Vector load	12

FIGURE 7.5 Start-up penalties on DLXV. These are the start-up penalties in clock cycles for DLXV vector operations. When a vector instruction depends on another vector instruction that has not completed at the time the second vector instruction issues, the start-up penalty is increased by 4 clock cycles.

To maintain an initiation rate of one word fetched or stored per clock, the memory system must be capable of producing or accepting this much data. This is usually done by creating multiple *memory banks*. Each memory bank is like a small, separate memory that can access different addresses in parallel with other banks. The words are then transferred from the memory at the maximum rate (one per clock in DLXV).

There are two possible implementation techniques for memory banks. One approach is to synchronize all the banks and to access them in parallel, latching the result in each bank. Once the result is latched, the next access can begin while the words are transferred. An alternative implementation technique uses independent bank phasing. On the first access, all the banks are accessed in parallel, and then the words are transferred one at a time from the banks. Once a

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 393

bank has transmitted or stored its data, it begins the next access immediately. The first approach (synchronized accesses) requires more latches, but has simpler control than an approach that uses independent bank phasing. The concept of memory banks is similar to but not identical to interleaving, as we will see in Figure 7.6. We discuss interleaving extensively in Chapter 8, Section 8.4.

Assuming each bank is one double-precision-word wide, if an initiation rate of one per clock is to be maintained, the following must hold:

Number of memory banks ≥ Memory-bank access time in clock cycles

To see why this relationship exists, think about a vector load of 64 double-precision words. Let the addresses of the vector elements be given by k_i, where

k_i = Starting address of the vector + (i-1) * Distance between vector elements.

For double-precision vector elements that are adjacent, the distance between elements will be 8 bytes. The addresses of the vector elements to be accessed by a bank will be the values of k_i such that

k_i mod number of banks = Bank number

Let's look at the first access by each bank. After a time equal to the memory-access time, all the memory banks will have fetched a double-precision word, and the words can begin returning to the vector registers. (This requires, of course, that the accesses be aligned on doubleword boundaries.) Words are sent serially from the banks, starting with the bank fetching from the lowest address. If the banks are synchronized, the next accesses start immediately; if the banks are phased, then the next access begins after an element is transmitted from the bank. In either case, a bank begins its next access at a byte address that is (8 * number of banks) higher than the last byte address. Because the memory-access time in clock cycles is less than the number of memory banks and because the words are transferred from the banks in round-robin order at a rate of one transfer per clock cycle, a bank will complete the next access before its turn to transmit data comes again. To simplify addressing, the number of memory banks is usually made a power of two. As we will see shortly, designers will probably want to have more than the minimum number of required banks so as to minimize memory stalls.

Example

Suppose we want to fetch a vector of 64 elements starting at byte address of 136, and a memory access takes 6 clocks. How many memory banks must we have? With what addresses are the banks accessed? When will the various elements arrive at the CPU?

Answer

Six clocks per access require at least 6 banks, but because we want the number of banks to be a power of two, we choose to have 8 banks. Figure 7.6 shows what byte addresses each bank accesses within each time period. Remember that a bank begins a new access as soon as it has completed the old access.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 394

Beginning at clock no.	Bank 0	1	2	3	4	5	6	7
0	192	136	144	152	160	168	176	184
6	256	200	208	216	224	232	240	248
14	320	264	272	280	288	296	304	312
22	384	328	336	344	352	360	368	376

FIGURE 7.6 Memory addresses (in bytes) by bank number and time slot at which access begins. The exact time when a bank transmits its data is given by the address it accesses minus the starting address divided by 8 plus the memory latency (6 clocks). It is important to observe that Bank 0 accesses a word in the next block (i.e., it accesses 192 rather than 128 and then 256 rather than 192, and so on). If Bank 0 were to start at the lower address we would require an extra cycle to transmit the data, and we would transmit one value unnecessarily. While this problem is not severe for this example, if we had 64 banks, up to 63 unnecessary clock cycles and transfers could occur. The fact that Bank 0 does not access a word in the same block of 8 distinguishes this type of memory system from interleaved memory. Normally, interleaved memory systems combine the bank address and the base starting address by concatenation rather than addition. Also, interleaved memories are almost always implemented with synchronized access. Memory banks require address latches for each bank, which are not normally needed in a system with only interleaving.

Figure 7.7 shows the timing for the first few sets of accesses for an 8–bank system with a 6–clock-cycle access latency. Two important observations about these two figures are these: First, notice that the exact address fetched by a bank is largely determined by the lower-order bits in the bank number; however, the initial access to a bank is always within 8 doublewords of the initial address. Second, notice that once the initial latency is overcome (6 clocks in this case), the pattern is to access a bank every n clock cycles, where n is the total number of banks (n=8 in this case).

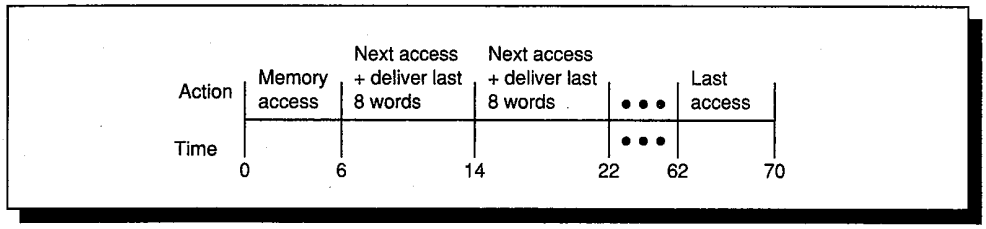

FIGURE 7.7 Access timing for the first 64 double-precision words of the load. After the 6–clock-cycle initial latency, 8 double-precision words are returned every 8 clock cycles.

The number of banks in the memory system and the pipeline depth in the functional units are essentially counterparts, since they determine the initiation rates for operations using these units. The processor cannot access memory faster than the memory cycle time. Thus, if memory is built from DRAM, where cycle time is about twice the access time, the processor will usually need twice as many banks as the computations above would give. This characteristic of DRAM is discussed further in Chapter 8, Section 8.4.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 395

7.3 Two Real-World Issues: Vector Length and Stride

This section deals with two issues that transpire in real programs. These are what to do when the vector length in a program is not exactly 64, and how to deal with nonadjacent elements in vectors when a matrix is laid out in memory. First, let's deal with the issue of vector length.

Vector-Length Control

A vector-register machine has a natural vector length determined by the number of elements in each vector register. This length, which is 64 for DLXV, is unlikely to match the real vector length in a program. Moreover, in a real program the length of a particular vector operation is often unknown at compile time. In fact, a single piece of code may require different vector lengths. For example, consider this code:

```
          do 10 i = 1,n
10          Y(i) = a * X(i) + Y(i)
```

The size of all the vector operations depends on n, which may not even be known until run-time! The value of n might also be a parameter to the procedure and therefore be subject to change during execution.

The solution to these problems is to create a *vector-length register* (VLR). The VLR controls the length of any vector operation, including a vector load or store. The value in the VLR, however, cannot be any greater than the length of the vector registers. This solves our problem as long as the real length is less than the *maximum vector length* (MVL) defined by the machine.

What if the value of n is not known at compile time, and thus may be greater than MVL? To tackle this problem, a technique called *strip mining* is used. Strip mining is the generation of code such that each vector operation is done for a size less than or equal to the MVL. The strip-mined version of the SAXPY loop written in FORTRAN, the major language used for scientific applications, is shown with C-style comments:

```
      low = 1
      VL = (n mod MVL)  /*find the odd size piece*/
      do 1 j = 0,(n / MVL)  /*outer loop*/
         do 10 i = low,low+VL-1  /*runs for length VL*/
            Y(i) = a*X(i) + Y(i)  /*main operation*/
10       continue
         low = low+VL  /*start of next vector*/
         VL = MVL  /*reset the length to max*/
1     continue
```

The term `n / MVL` represents truncating integer division (which is what FORTRAN does) and is used throughout this section. The effect of this loop is to block the vector into segments which are then processed by the inner loop. The length of the first segment is (n mod MVL) and all subsequent segments are of length MVL. This is depicted in Figure 7.8.

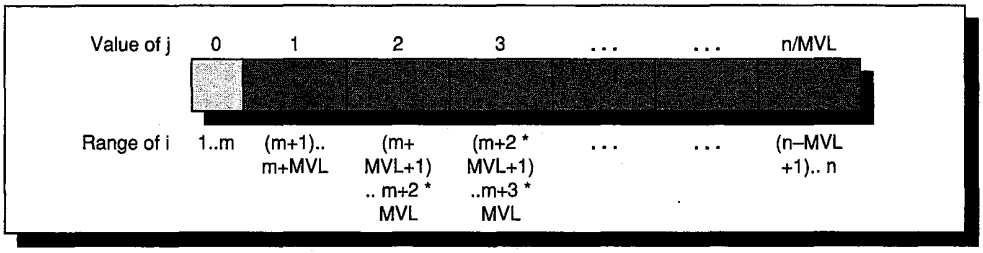

FIGURE 7.8 A vector of arbitrary length processed with strip mining. All blocks but the first are of length MVL, utilizing the full power of the vector machine. In this figure, the variable *m* is used for the expression (n mod MVL).

The inner loop of the code above is vectorizable with length VL, which is equal to either (n mod MVL) or MVL. The VLR register must be set twice—once at each place where the variable VL in the code is assigned. With multiple vector operations executing in parallel, the hardware must copy the value of VLR when a vector operation issues, in case VLR is changed for a subsequent vector operation.

In the previous section, start-up overhead could be computed independently for each vector operation. With strip mining, a significant percentage of the start-up cost will be the strip-mining overhead itself; and, therefore, computing the start-up overhead will be more complex.

Let's see how significant these added overheads are. Consider a simple loop:

```
      do 10 i  =  1,n
10      A(i)   =  B(i)
```

The compiler will generate two nested loops for this code, just as our earlier example does. The inner loop contains a sequence of two vector operations, `LV` (load vector) followed by `SV` (store vector). Each loop iteration of the original vector operation would require two clocks if there were no start-up penalties of any kind. The start-up penalties consist of two types: vector start-up overhead and strip-mining overhead. For DLXV the vector start-up overhead is 12 clock cycles for the vector load plus a 4-clock-cycle delay because the store depends on the load, for a total of 16 clock cycles. We can ignore the store latency, since nothing depends on it. Figure 7.9 (page 366) shows the impact of the vector start-up cost alone as the vector grows from length 1 to length 64. This start-up cost can decrease the throughput rate by a factor of as much as 9, depending on the vector length.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 397

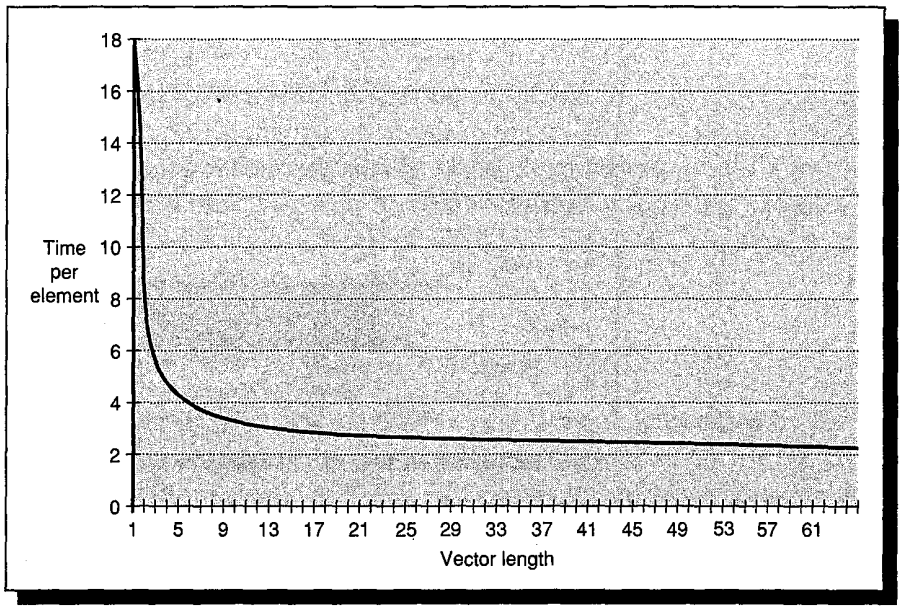

FIGURE 7.9 The impact of just the vector start-up cost on a loop consisting of a vector assignment. For short vectors, the impact of the 16-cycle start-up cost is enormous, decreasing performance by up to nine times. The strip-mining overhead has not been included.

In Section 7.4, we will see a unified performance model that incorporates all the start-up and overhead costs. First, let's examine how to implement vectors with nonsequential memory accesses.

Vector Stride

The second problem this section addresses is that the position in memory of adjacent elements in a vector may not be sequential. Consider the straightforward code for matrix multiply:

```
      do 10 i = 1,100
         do 10 j = 1,100
            A(i,j) = 0.0
            do 10 k = 1,100
10             A(i,j) = A(i,j)+B(i,k)*C(k,j)
```

At the statement labeled 10 we could vectorize the multiplication of each row of B with each column of C and strip-mine the inner loop with k as the index variable. To do so, we must consider how adjacent elements in B and adjacent elements in C are addressed. When an array is allocated memory it is linearized and must be laid out in either *row-major* or *column-major* order. Row-major order, used by most languages except FORTRAN, lays out the rows first, making elements B(i,j) and B(i,j+1) adjacent. Column-major order, used by FORTRAN,

makes B(i,j) and B(i+1,j) adjacent. Figure 7.10 illustrates these two alternatives. Let's look at the accesses to B and C in the inner loop of the matrix multiply. In FORTRAN, the accesses to the elements of B will be nonadjacent in memory, and each iteration will access an element that is separated by an entire row of the array. In this case, the elements of B that are accessed by iterations in the inner loop are separated by the row size times 8 (the number of bytes per entry) for a total of 800 bytes.

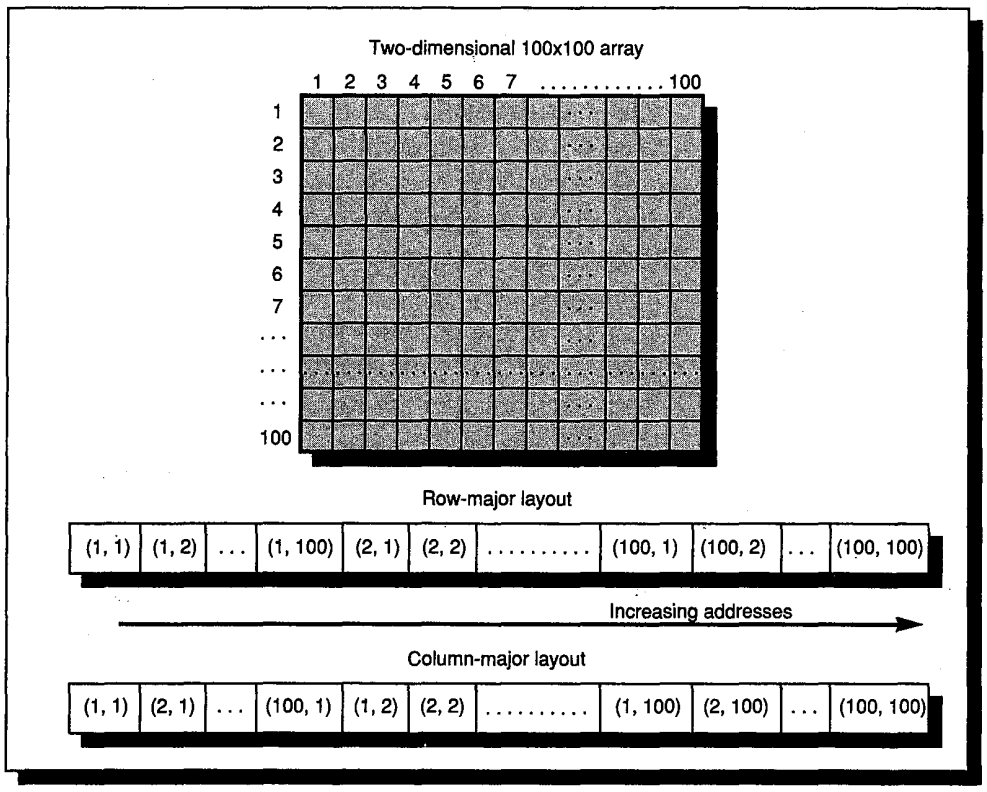

FIGURE 7.10 Matrix for a two-dimensional array and corresponding layouts in one-dimensional storage. In row-major order, successive row elements are adjacent in storage, while in column-major order, successive column elements are adjacent. It is easy to imagine extending this to arrays with more dimensions.

This distance separating elements that are to be merged into a single vector is called the *stride*. In the current example, using column-major layout for the matrices means that matrix C has a stride of 1, or 1 doubleword (8 bytes), separating successive elements, and matrix B has a stride of 100, or 100 doublewords (800 bytes).

Once a vector is loaded into a vector register it acts as if it had logically adjacent elements. This enables a vector-register machine to handle strides greater than one, called *nonunit strides*, by making more general vector-load and vector-store operations. For example, if we could load a row of B into a vector register, we could then treat the row as logically adjacent.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 399

Thus, it is desirable for the vector load and store operations to specify a stride in addition to a starting address. On a DLXV, where the addressable unit is a byte, the stride for our example would be 800. The value must be computed dynamically, since the size of the matrix may not be known at compile time, or—just like vector length—may change for different executions of the same statement. The vector stride, like the vector starting address, can be put in a general-purpose register, where it is used for the life of the vector operation. Then the DLXV instruction `LVWS` (Load Vector With Stride) can be used to fetch the vector into a vector register. Likewise, when a nonunit stride vector is being stored, `SVWS` (Store Vector With Stride) can be used. In some vector machines the loads and stores always have a stride value stored in a register, so there is only a single instruction.

Memory-unit complications can occur from supporting strides greater than one. Earlier, we saw that a vector-memory operation could proceed at full speed if the number of memory banks was at least as large as the memory-access time in clock cycles. However, once nonunit strides are introduced it becomes possible to request accesses from the same bank at a higher rate than the memory-access time. This situation is called *memory-bank conflict* and results in each load seeing a larger portion of the memory-access time. A memory-bank conflict occurs whenever the same bank is asked to do an access before it has completed another. Thus, a bank conflict, and hence a stall, will occur if:

$$\frac{\text{Least common multiple (Stride,Number of banks)}}{\text{Stride}} < \text{Memory-access latency}$$

Example

Suppose we have 16 memory banks with an access time of 12 clocks. How long will it take to complete a 64-element vector load with a stride of 1? With a stride of 32?

Answer

Since the number of banks is larger than the load latency, for a stride of 1, the load will take 12 + 64 = 76 clock cycles, or 1.2 clocks per element. The worst possible stride is a value that is a multiple of the number of memory banks, as in this case with a stride of 32 and 16 memory banks. Every access to memory will collide with the previous one. This leads to an access time of 12 clock cycles per element and a total time for the vector load of 768 clock cycles.

Memory bank conflicts will not occur if the stride and number of banks are relatively prime with respect to each other and there are enough banks to avoid conflicts in the unit-stride case. Increasing the number of memory banks to a number greater than the minimum to prevent stalls with a stride of length 1 will decrease the stall frequency for some other strides. For example, with 64 banks, a stride of 32 will stall on every other access, rather than every access. If we originally had a stride of 8 and 16 banks, every other access would stall; while with 64 banks, a stride of 8 will stall on every eighth access. If we have multiple memory pipelines, we will also need more banks to prevent conflicts. In the

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 400

1990s, most vector supercomputers have at least 64 banks, and some have as many as 512.

7.4 A Simple Model for Vector Performance

This section presents a model for understanding the performance of a vectorized loop. There are three key components of the running time of a strip-mined loop whose body is a sequence of vector instructions:

1. The time for each vector operation in the loop to process one element, ignoring the start-up costs, which we call $T_{element}$. The vector sequence often has a single result, in which case $T_{element}$ is the time to produce an element in that result. If the vector sequence produces multiple results, $T_{element}$ is the time to produce one element in each result. This time depends only on the execution of vector instructions. We will see an example shortly.

2. The overhead for each strip-mined block of vector instructions. This overhead consists of the cost of executing the scalar code for strip mining of each block, T_{loop}, plus the vector start-up cost for each block, T_{start}.

3. The overhead from computing the starting addresses and setting up the vector control. This occurs once for the entire vector operation. This time, T_{base}, consists solely of scalar overhead instructions.

These components can be used to state the total running time for a vector sequence operating on a vector of length n, which we will call T_n:

$$T_n = T_{base} + \left\lceil \frac{n}{MVL} \right\rceil * (T_{loop} + T_{start}) + n * T_{element}$$

The values of T_{start} and T_{loop} are both compiler and machine dependent, while the value of $T_{element}$ depends mainly on the hardware. The exact vector sequence affects all three values; the effect on $T_{element}$ is probably the most pronounced, with T_{start} and T_{loop} less affected.

For simplicity, we will use constant values for T_{base} and for T_{loop} on DLXV. Based on a variety of measurements of CRAY-1 vector execution, the values chosen are 10 for T_{base} and 15 for T_{loop}. At first glance, you might think that these values, especially T_{loop}, are too small. The overhead in each loop requires: setting up the vector starting addresses and the strides, incrementing counters, and executing a loop branch. However, these scalar instructions can be overlapped with the vector instructions, minimizing the time spent on these overhead functions. The values of T_{base} and T_{loop} of course depend on the loop structure, but the dependence is slight compared to the connection between the vector code and the values of $T_{element}$ and T_{start}.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 401

Example

What is the execution time for the vector operation A = B * s, where s is a scalar and the length of the vectors A and B is 200?

Answer

Here is the strip-mined DLXV code, assuming the addresses of A and B are initially in `Ra` and `Rb`, and s is in `Fs`:

```
      ADDI    R2,R0,#1600 ;no. bytes in vector
      ADD     R2,R2,Ra    ;end of A vector
      ADDI    R1,R0,#8    ;strip-mined length
      MOVI2S  VLR,R1      ;load vector length
      ADDI    R1,R0,#64   ;length in bytes
      ADDI    R3,R0,#64   ;vector length of other pieces
loop: LV      V1,Rb       ;load B
      MULTSV  V2,Fs,V1    ;vector * scalar
      SV      Ra,V2       ;store A
      ADD     Ra,Ra,R1    ;next segment of A
      ADD     Rb,Rb,R1    ;next segment of B
      ADDI    R1,R0,#512  ;full vector length (bytes)
      MOVI2S  VLR,R3      ;set length to 64
      SUB     R4,R2,Ra    ;at the end of A?
      BNZ     R4,LOOP     ;if not, go back
```

From this code, we can see that: $T_{element} = 3$, for the load, multiply and store of each value of the vector. Furthermore, our assumptions for DLXV are $T_{loop} = 15$ and $T_{base} = 10$. Let's use our basic formula:

$$T_n = T_{base} + \left\lceil \frac{n}{MVL} \right\rceil * (T_{loop} + T_{start}) + n * T_{element}$$

$$T_{200} = 10 + (4) * (15 + T_{start}) + 200 * 3$$

$$T_{200} = 10 + 4 * (15 + T_{start}) + 600 = 670 + 4 * T_{start}$$

The value of T_{start} is the sum of

- The vector load start-up of 12 clock cycles,
- The 4–clock-cycle stall due to the dependence between the load and multiply,
- A 7–clock-cycle start-up for the multiply, plus
- A 4–clock-cycle stall due to the dependence between the multiply and store.

Thus, the value of T_{start} is given by:

$$T_{start} = 12 + 4 + 7 + 4 = 27$$

So, the overall value becomes

$$T_{200} = 670 + 4 * 27 = 778$$

The execution time per element with all start-up costs is then $\frac{778}{200} = 3.9$, compared with an ideal case of 3.

Figure 7.11 shows the overhead and effective rates per element for the above example (A = B*s) with various vector lengths. Compared to the simpler model of start-up, illustrated in Figure 7.9 on page 366, we see that the overhead accounting for all sources is higher. In this example, the vector start-up cost, which is what is plotted in Figure 7.9, accounts for only about half the total overhead per element.

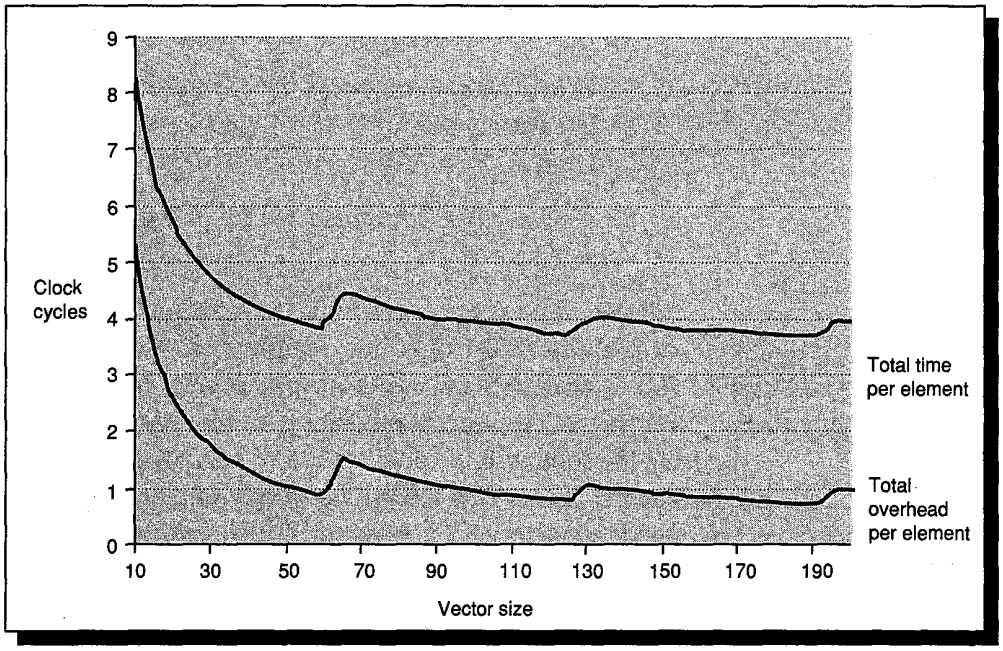

FIGURE 7.11 This shows the total execution time per element and the total overhead time per element, versus the vector length for the example on page 370. For short vectors the total start-up time is more than one-half of the total time, while for long vectors it reduces to about one-third of the total time. The sudden jumps occur when the vector length crosses a multiple of 64, forcing another iteration of the strip-mining code and execution of a set of vector instructions. These operations increase T_n by $T_{loop} + T_{start}$.

7.5 Compiler Technology for Vector Machines

To make effective use of a vector machine a compiler must be able to recognize that a loop (or part of a loop) is vectorizable and generate the appropriate vector code. This involves determining what dependences exist among the operands in the loop. For now, we will consider only dependences that occur when an operand is written at one point and read at a later point. These correspond to RAW (read after write—see page 264) hazards. Consider a loop like this one:

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 403

```
        do 10 i=1,100
1          A(i+1) = A(i) + B(i)
2          B(i+1) = B(i) + A(i+1)
10      continue
```

Call the numbered statements 1 and 2 in the loop body S1 and S2, respectively. The possible different types of dependences are

1. S1 uses a value computed by S1 in an earlier iteration. This is true for S1 since iteration i+1 uses the value A(i) that was computed in iteration i as A(i+1). The same is true of S2 for B(i) and B(i+1).
2. S1 uses a value computed by S2 in an earlier iteration. This is true since S1 uses the value of B(i+1) in iteration i+1 that is computed by S2 in iteration i.
3. S2 uses a value computed by S1 in the same iteration. This is true for the value A(i+1).

Because the vector operations are pipelined and the latency may be quite long, an early iteration may not complete before a later iteration begins: Thus, the values that will be written by the early iteration may not have been written before the later iteration begins. Consequently, if situation 1 or 2 exists, vectorizing the loop will introduce a RAW hazard—a hazard that a vector machine does not check for. This means that if any of the three dependences in situation 1 and 2 exist, the loop is not vectorizable, and the compiler will not generate vector instructions for this code. In situation 3, the normal hazard-detection hardware could handle the situation. A loop containing only dependences like those in situation 3 can therefore be vectorized, as we will see soon. The dependences in the first two situations, which involve the use of values computed on earlier loop iterations, are called *loop-carried dependences*.

The first task of the compiler is to determine whether there are any loop-carried dependences within the loop body. The compiler accomplishes this with a dependence-analysis algorithm. Because the statements in the loop body involve arrays, dependence analysis is complex. (If there weren't arrays, there would be nothing to vectorize.) The simplest case occurs when an array name appears only on one side of an assignment statement. Take, for example, this variation of our earlier loop:

```
        do   10 i=1,100
             A(i) = B(i) + C(i)
             D(i) = A(i) * E(i)
10      continue
```

If the arrays A, B, C, D, and E are different, then no loop-carried dependence can exist. There is a dependence between the two statements for the vector A. If the compiler realized that there were two accesses to A, it might try not to reload A

the second statement, instead doing the vector multiply using the result register from the vector add. In this case, the processor would see the potential RAW hazard and stall the issue of the vector multiply. If the compiler stored A and reloaded it, then the loads and stores would occur in order, yielding correct execution.

Often the same name appears as both a source and destination within a loop, as it did in the SAXPY loop. There, Y appears on both sides of the assignment:

```
      do 10 i=1,100
         Y(i) = a*X(i) + Y(i)
10    continue
```

In this case there is still no loop-carried dependence because the assignment to Y does not depend on a value of Y computed in an earlier iteration. However, the following loop, which is called a *recurrence*, does contain a loop-carried dependence:

```
      do 10 i=2,100
         Y(i) = Y(i-1) + Y(i)
10    continue
```

The dependence can be seen by unwinding the loop: In iteration j the value of Y(j–1) is used, but that element is stored in iteration j–1, creating a loop-carried dependence.

How does the compiler detect dependences in general? Suppose we have written to an array element with index value $a * i + b$ and accessed with index value $c * i + d$, where i is the for-loop index variable that runs from m to n. A dependence exists if two conditions hold:

1. There are two iteration indices, j and k, both within the limits of the for loop.
2. The loop stores into an array element indexed by $a*j+b$ and later fetches from that **same** array element when it is indexed by $c*k+d$. That is, $a*j+b = c*k+d$.

In general, we may not be able to determine whether a dependence exists at compile time. For example, the values of a, b, c, and d may not be known, making it impossible to tell if a dependence exists. In other cases, the dependence testing may be very expensive but decidable at compile time. For example, the accesses may depend on the iteration indices of multiply nested loops. Many programs do not contain these complex structures, but instead contain simple indices where a, b, c, and d are all constants. For these cases, it is possible to devise reasonable tests for dependence.

A simple and sufficient test used to detect dependences is the *greatest common divisor*, or GCD. It is based on the observation that if a loop-carried dependence exists, then GCD (c,a) must **divide** (d–b). (Remember that an integer, x, *divides* another integer, y, if there is no remainder when we do the division $\frac{y}{x}$ and

get an integer result.) The GCD test is sufficient to guarantee that no dependence exists (see Exercise 7.10); however, there are cases where the GCD test succeeds, but no dependence exists. For example, this can arise because the GCD test does not take the loop bounds into account. A more complex test is the Banerjee test, named after U. Banerjee [1979], that accounts for loop bounds, but is still not exact. An exact test can always be done by solving equations for integer values, but this can be expensive for complex loop structures.

Example

Use the GCD test to determine whether dependences exist in the following loop:

```
      do 10 i=1,100
10      X(2*i+3) = X(2*i) * 5.0
```

Answer

Given the values a=2, b=3, c=2, and d=0, then GCD(a,c) = 2, and d–b = –3. Since 2 does not divide –3, no dependence is possible.

A *true data dependence* arises from a RAW hazard and will prevent vectorization of the loop as a single vector sequence. There are cases where the loop can be vectorized as two separate vector sequences (see Exercise 7.11). There are also dependences corresponding to a WAR (write after read) hazard, called an *antidependence*, and to a WAW (write after write) hazard, called an *output dependence*. Antidependences and output dependences are not true data dependences. They are name conflicts and can be eliminated by renaming of registers in the compiler in a method similar to how Tomasulo's algorithm renames registers at run time (see Section 6.7 in Chapter 6). Vectorizing compilers often use compile-time renaming to eliminate antidependences and output dependences.

Example

The following loop has an antidependence (WAR) and an output dependence (WAW). Find all the true dependences, output dependences, and antidependences, and eliminate the output dependences and antidependences by renaming.

```
      do 10 i=1,100
1         Y(i) = X(i) / s
2         X(i) = X(i) + s
3         Z(i) = Y(i) + s
4         Y(i) = s - Y(i)
10    continue
```

Answer

There are true dependences from statement 1 to statement 3 and from statement 1 to statement 4 because of `Y(i)`. These are not loop carried, so they will not prevent vectorization. However, the dependences will force statements 3 and 4 to wait for statement 1 to complete, even though statements 3 and 4 use a different functional unit than statement 1. In the next section we will see a technique for eliminating this serialization.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 406

There is an antidependence from statement 1 to statement 2, and an output dependence from statement 1 to statement 4. The following version of the loop eliminates these false (or pseudo) dependences.

```
        do 10 i=1,100
C       Y renamed to T to remove output dependence
1           T(i) = X(i) / s
C       X renamed to X1 to remove antidependence
2           X1(i) = X(i) + s
3           Z(i) = T(i) + s
4           Y(i) = s - T(i)
10      continue
```

After the loop the variable X has been renamed X1. In code that follows the loop, the compiler can simply replace the name X by X1. Renaming does not require an actual copy operation; it can be done by substituting names or by register allocation.

Besides deciding which loops are vectorizable, the compiler must generate strip-mining code and allocate vector registers. Most vectorization transformations are done at the source level, although some optimizations involve coordinating high-level source transformations with lower-level, machine-dependent transformations. Efficient allocation of vector registers is such an optimization and is perhaps the most difficult optimization—one that many vectorizing compilers do not attempt.

Effectiveness of Vectorization Techniques

Two factors affect the success with which a program can be run in vector mode. The first factor is the structure of the program itself: do the loops have true data dependences, or can they be restructured so as not to have such dependences? This factor is influenced by the algorithms chosen and, to some extent, how they are coded. The second factor is the capability of the compiler. While no compiler can vectorize a loop where no parallelism among the loop iterations exists, there is tremendous variation in the ability of compilers to determine whether a loop can be vectorized.

As an indication of the level of vectorization that can be achieved in scientific programs, let's look at the vectorization levels observed for the Perfect Club benchmarks, discussed in Section 2.7 of Chapter 2. These benchmarks are large, real scientific applications. Figure 7.12 (page 376) shows the percentage of floating-point operations in each benchmark and the percentage executed in vector mode on the CRAY X-MP. The wide variation in level of vectorization has been observed by several studies of the performance of applications on

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 407

vector machines. While better compilers might improve the level of vectorization in some of these programs, most will require rewriting to achieve significant increases in vectorization. For example, let's look at our version of the Spice benchmark in detail. In Spice with the input chosen we found that only 3.7% of the floating-point operations are executed in vector mode on the CRAY X-MP, and the vector version runs only 0.5% faster than the scalar version. Clearly, a new program or a significant rewrite will be needed to obtain the benefits of a vector machine on Spice.

Benchmark name	FP operations	FP operations executed in vector mode
ADM	23%	68%
DYFESM	26%	95%
FLO52	41%	100%
MDG	28%	27%
MG3D	31%	86%
OCEAN	28%	58%
QCD	14%	1%
SPICE	16%	7%
TRACK	9%	23%
TRFD	22%	10%

FIGURE 7.12 Level of vectorization among the Perfect Club benchmarks when executed on the CRAY X-MP. The first column contains the percentage of operations that are floating point, while the second contains the percentage of FP operations executed in vector instructions. Note that this run of Spice with different inputs shows a higher vectorization ratio.

There is also tremendous variation in how well compilers do in vectorizing programs. As a summary of the state of vectorizing compilers, consider the data in Figure 7.13, which shows the extent of vectorization for different machines using a test suite of 100 hand-written FORTRAN kernels. The kernels were designed to test vectorization capability and can all be vectorized by hand; we will see several examples of these loops in the Exercises.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 408

Machine	Compiler	Completely vectorized	Partially vectorized	Not vectorized
Ardent Titan-1	FORTRAN V1.0	62	6	32
CDC CYBER-205	VAST-2 V2.21	62	5	33
Convex C-series	FC5.0	69	5	26
CRAY X-MP	CFT77 V3.0	69	3	28
CRAY X-MP	CFT V1.15	50	1	49
CRAY-2	CFT2 V3.1a	27	1	72
ETA-10	FTN 77 V1.0	62	7	31
Hitachi S810/820	FORT77/HAP V20-2B	67	4	29
IBM 3090/VF	VS FORTRAN V2.4	52	4	44
NEC SX/2	FORTRAN77 / SX V.040	66	5	29
Stellar GS 1000	F77 prerelease	48	11	41

FIGURE 7.13 Result of applying vectorizing compilers to the 100 FORTRAN test kernels. For each machine we indicate how many loops were completely vectorized, partially vectorized, and unvectorized. These loops were collected by Callahan, Dongarra, and Levine [1988]. The machines shown are those mentioned at some point in this chapter. Two different compilers for the CRAY X-MP show the large dependence on compiler technology.

7.6 Enhancing Vector Performance

Three techniques for improving the performance of vector machines are discussed in this section. The first deals with making a sequence of dependent vector operations run faster. The other two deal with expanding the class of loops that can be run in vector mode. The first technique, chaining, originated in the CRAY-1, but is now supported on many vector machines. The techniques discussed in the second and third parts of this section are taken from a variety of machines and are, in general, more extensive than the capabilities provided on the CRAY-1 or CRAY X-MP architectures.

Chaining—The Concept of Forwarding Extended to Vector Registers

Consider the simple vector sequence

```
MULTV     V1,V2,V3
ADDV      V4,V1,V5
```

In DLXV as it currently stands these two instructions run in time equal to

$$T_{element} * \text{Vector length} + \text{Start-up time}_{ADDV} + \text{stall time} + \text{Start-up time}_{MULTV}$$
$$= 2 * \text{Vector length} + 6 + 4 + 7$$
$$= 2 * \text{Vector length} + 17$$

Because of the dependence, the `MULTV` must complete before the `ADDV` can begin. However, if the vector register, `V1` in this case, is treated not as a single entity but as a group of individual registers, then the pipelining concept of forwarding can be extended to work on individual elements of a vector. This idea, which will allow the `ADDV` to start earlier in this example, is called *chaining*. Chaining allows a vector operation to start as soon as the individual elements of its vector source operand become available: The results from the first functional unit in the chain are forwarded to the second functional unit. (Of course, they must be different units to avoid using the same unit twice per clock!) In a chained sequence the initiation rate is equal to one per clock cycle if the functional units in the chained operations are all fully pipelined. Even though the operations depend on one another, chaining allows the operations to proceed in parallel on separate elements of the vector. A sustained rate (ignoring start-up) of two floating-point operations per clock cycle can be achieved, even though the operations are dependent!

The total running time for the above sequence becomes

$$\text{Vector length} + \text{Start-up time}_{ADDV} + \text{Start-up time}_{MULTV}$$

Figure 7.14 shows the timing of a chained and an unchained version of the above pair of vector instructions with a vector length of 64. In Figure 7.14, the total time for chained operation is 77 clock cycles. With 128 floating-point operations done in that time, 1.7 FLOPs per clock cycle are obtained, versus a total time of 145 clock cycles or 0.9 FLOPs per clock cycle for the unchained version.

We will see in Section 7.7 that chaining plays a major role in boosting vector performance.

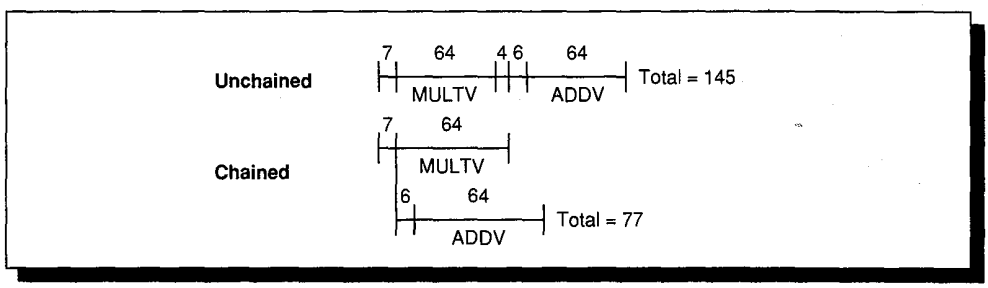

FIGURE 7.14 Timings for a sequence of dependent vector operations `ADDV` and `MULTV`, both unchained and chained. The 4–clock-cycle delay comes from a stall for dependence, described earlier; the 6– and 7–clock-cycle delays are the latency of the adder and multiplier.

Conditionally Executed Statements and Sparse Matrices

In the last section, we saw that many programs only achieved low to moderate levels of vectorization. Because of Amdahl's Law, the speedup on such programs will be very limited. Two reasons why higher levels of vectorization are not achieved are the presence of conditionals (if statements) inside loops and the use of sparse matrices. Programs that contain if statements in loops cannot be run in vector mode using the techniques we have discussed so far because the if statements introduce control flow into a loop. Likewise, sparse matrices cannot be efficiently implemented using any of the capabilities we have seen so far; this is a major factor in the lack of vectorization for Spice. This section discusses techniques that allow programs with these structures to execute in vector mode. Let's start with conditional execution.

Consider the following loop:

```
      do 100 i = 1, 64
           if (A(i) .ne. 0) then
                A(i) = A(i) - B(i)
           endif
100   continue
```

This loop cannot normally be vectorized because of the conditional execution of the body. However, if the inner loop could be run for the iterations for which A(i) ≠ 0, then the subtraction could be vectorized.

Vector-mask control helps us do this. The *vector-mask control* takes a Boolean vector of length MVL. When the *vector-mask register* is loaded with the result of a vector test, any vector instructions to be executed operate only on the vector elements whose corresponding entries in the vector-mask register are 1. The entries in the destination vector register that correspond to a 0 in the mask register are unaffected by the vector operation. Clearing the vector-mask register sets it to all 1s, making subsequent vector instructions operate on all vector elements. The following code can now be used for the above loop, assuming that the starting addresses of A and B are in Ra and Rb respectively:

```
LV     V1,Ra        ;load vector A into V1
LV     V2,Rb        ;load vector B
LD     F0,#0        ;load FP zero into F0
SNESV  F0,V1        ;sets the VM to 1 if V1(i)≠F0
SUBV   V1,V1,V2     ;subtract under vector mask
CVM                 ;set the vector mask to all 1s
SV     Ra,V1        ;store the result in A
```

Most modern vector machines provide vector-mask control. The vector-mask capability described here is available on some machines, but others allow the use of the vector mask with only a small number of instructions.

Using a vector-mask register does, however, have disadvantages. First, execution time is not decreased, even though some elements in the vector are not operated on. Second, in some vector machines the vector mask serves only to disable the storing of the result into the destination register, and the actual operation still occurs. Thus, if the operation in the above example were a divide rather than a subtract and the test was on B rather than A, false floating-point exceptions might result since the operation was actually done. Machines that mask the operation as well as the result store avoid this problem.

Now, let's turn to sparse matrices; later we will show another method for handling conditional execution. We have dealt with vectors in which the elements are separated by a constant stride. If an application called for a sparse matrix, we might see code that looks like:

```
        do    100 i = 1,n
100           A(K(i)) = A(K(i)) + C(M(i))
```

This code implements a sparse vector sum on the arrays A and C, using index vectors K and M to designate to the nonzero elements of A and C. (A and C must have the same number of nonzero elements—n of them.) Another common representation for sparse matrices uses a bit vector to say which elements exist, and often both representations exist in the same program. Sparse matrices are found in many codes, and there are many ways to implement them, depending on the data structure used in the program.

The primary mechanism for supporting sparse matrices is scatter-gather operations using index vectors. A *gather* operation takes an *index vector*, and fetches the vector whose elements are at the addresses given by adding a base address to the offsets given in the index vector. The result is a nonsparse vector in a vector register. After these elements are operated on in dense form, the sparse vector can be stored in expanded form by a *scatter* store, using the same index vector. Hardware support for such operations is called *scatter-gather* and appeared on the CDC STAR-100. The instructions `LVI` (Load Vector Indexed) and `SVI` (Store Vector Indexed) provide these operations in DLXV. For example, assuming that `Ra`, `Rc`, `Rk`, and `Rm` contain the starting addresses of the vectors in the above sequence, the inner loop of the sequence can be coded with vector instructions such as:

```
LV   Vk,Rk          ;load K
LVI  Va,(Ra+Vk)     ;load A(K(I))
LV   Vm,Rm          ;load M
LVI  Vc,(Rc+Vm)     ;load C(M(I))
ADDV Va,Va,Vc       ;add them
SVI  (Ra+Vk),Va     ;store A(K(I))
```

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 412

This technique allows code with sparse matrices to be run in vector mode. The source code above would **never** be automatically vectorized by a compiler because the compiler cannot know that the elements of K are distinct values, and thus that no dependences exist. Instead, a programmer directive would tell the compiler that it could run the loop in vector mode.

A scatter/gather capability is included on many of the newest supercomputers. Such operations rarely run at one element per clock, but they are still much faster than the alternative, which may be a scalar loop. If the sparsity properties of a matrix change, a new index vector must be computed. Many machines provide support for computing the index vector quickly. The `CVI` (Create Vector Index) instruction in DLXV creates an index vector given a stride (m), where the values in the index vector are $0,m,2*m,...,63*m$. Some machines provide an instruction to create a compressed index vector whose entries correspond to the positions with a 1 in the mask register. Other vector architectures provide a method to compress a vector. In DLXV, we define the `CVI` instruction to always create a compressed index vector using the vector mask. When the vector mask is all ones a standard index vector will be created.

The indexed loads/stores and the CVI instruction provide an alternative method to support conditional execution. Here is a vector sequence that implements the loop we saw on page 379:

```
LV      V1,Ra        ;load vector A into V1
LD      F0,#0        ;load FP zero into F0
SNESV   F0,V1        ;sets the VM to 1 if V1(i)≠F0
ADDI    Rc,#8
CVI     V2,Rc        ;generates indices in V2
POP     R1,VM        ;find the number of 1's in VM
MOVI2S  VLR,R1       ;load vector length register
CVM
LVI     V3,(Ra+V2)   ;load the nonzero A elements
LVI     V4,(Rb+V2)   ;load corresponding B elements
SUBV    V3,V3,V4     ;do the subtract
SVI     (Ra+V2),V3   ;store A back
```

Whether the implementation using scatter/gather is better than the conditionally executed version depends on the frequency with which the condition holds and the cost of the operations. Ignoring chaining, the running time of the first version (on page 379) is $5n + c_1$. The running time of the second version using indexed loads and stores with a running time of one element per clock is $4n + 4*f*n + c_2$, where f is the fraction of elements for which the condition is true (i.e., $A \neq 0$). If we assume that the values of c_1 and c_2 are comparable, or that they are much smaller than n, we can find when this second technique is better.

$$\text{Time}_1 = 5n$$

$$\text{Time}_2 = 4n + 4*f*n$$

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 413

We want $\text{Time}_1 \geq \text{Time}_2$, so

$$5n \geq 4n + 4*f*n$$

$$\frac{1}{4} \geq f$$

That is, the second method is faster if less than one-quarter of the elements are nonzero. In many cases the frequency of execution is much lower. If the index vector can be reused, or if the number of vector statements within the if statement grows, the advantage of the scatter/gather approach will increase sharply.

Vector Reduction

As we saw in Section 7.5, some loop structures are not easily vectorized. One common structure is a *reduction*—a loop that reduces an array to a single value by repeated application of an operation. This is a special case of a recurrence. A common example occurs in dot product:

```
      dot = 0.0
      do 10 i=1,64
10       dot = dot + A(i) * B(i)
```

This loop has an obvious loop-carried dependence (on `dot`) and cannot be vectorized in a straightforward fashion. The first thing a good vectorizing compiler would do is split the loop to separate out the vectorizable portion and the recurrence and perhaps rewrite the loop as:

```
      do 10 i=1,64
10       dot(i) =  A(i) * B(i)

      do 20 i=2,64
20       dot(1) = dot(1) + dot(i)
```

The variable `dot` has been expanded into a vector; this transformation is called *scalar expansion.*

One simple scheme for compiling the loop with the recurrence is to add sequences of progressively shorter vectors—two 32-element vectors, then two 16-element vectors, and so on. This technique has been called *recursive doubling*. It is faster than doing all the operations in scalar mode. Many vector machines provide hardware assist for doing reductions, as we will see next.

Example

Show how the FORTRAN code would look for execution of the second loop in the code fragment above using recursive doubling.

Answer

Here is the code:

```
            len = 32
            do 100 j=1,6
                  do 10 i=1,len
10                      dot(i) = dot(i) + dot(i+len)
                  len = len / 2
100         continue
```

When the loop is done, the sum is in dot(1).

In some vector machines, the vector registers are addressable, and another technique, sometimes called partial sums, can be used. This is discussed in Exercise 7.12. There is an important caveat in the use of vector techniques for reduction. To make reduction work, we are relying on the associativity of the operator being used for the reduction. Because of rounding and finite range, however, floating-point arithmetic is not strictly associative. For this reason, most compilers require the programmer to indicate whether associativity can be used to more efficiently compile reductions.

7.7 Putting It All Together: Evaluating the Performance of Vector Processors

In this section we look at different measures of performance for vector machines and what they tell us about the machine. To determine the performance of a machine on a vector problem we must look at the start-up cost and the sustained rate. The simplest and best way to report the performance of a vector machine on a loop is to give the execution time of the vector loop. For vector loops people often give the MFLOPS (Millions FLoating point Operations Per Second) rating rather than execution time. We use the notation R_n for the MFLOPS rating on a vector of length n. Using the measurements T_n (time) or R_n (rate) is equivalent if the number of FLOPs is agreed upon (see Chapter 2, Section 2.2, page 35 for an extensive discussion on MFLOPS). In any event, either measurement should include the overhead.

In this section we examine the performance of DLXV on our SAXPY loop by looking at performance from different viewpoints. We will continue to compute the execution time of a vector loop using the equation developed in Section 7.4. At the same time, we will look at different ways to measure performance using the computed time. The constant values for T_{loop} and T_{base} used in this section introduce some small amount of error, which will be ignored.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 415

Measures of Vector Performance

Because vector length is so important in establishing the performance of a machine, length-related measures are often applied in addition to time and MFLOPs. These length-related measures tend to vary dramatically across different machines and are interesting to compare. (Remember, though, that **time** is always the measure of interest when comparing the relative speed of two machines.) Three of the most important length-related measures are:

R_∞—The MFLOPS rate on an infinite-length vector. Although this measure may be of interest when estimating peak performance, real problems do not have unlimited vector lengths, and the overhead penalties encountered in real problems will be larger. (R_n is the MFLOPS rate for a vector of length n.)

$N_{1/2}$—The vector length needed to reach one-half of R_∞. This is a good measure of the impact of overhead.

N_v—The vector length needed to make vector mode faster than scalar mode. This measures both overhead and the speed of scalars relative to vectors.

Let's look at these measures for our SAXPY problem running on DLXV. When chained, the inner loop of the SAXPY code looks like this (assuming that `Rx` and `Ry` hold starting addresses):

```
LV      V1,Rx       ;load the vector X
MULTSV  V2,S1,V1    ;vector*scalar-chained to LV X
LV      V3,Ry       ;vector load Y
ADDV    V4,V2,V3    ;sum aX + Y, chained to LV Y
SV      Ry,V4       ;store the vector Y
```

Recall our performance equation for the execution time of a vector loop with n elements, T_n:

$$T_n = T_{base} + \left\lceil \frac{n}{MVL} \right\rceil * (T_{loop} + T_{start}) + n * T_{element}$$

Since there are three memory references and only one memory pipeline, the value of $T_{element}$ must be at least 3, and chaining allows it to be exactly 3. If $T_{element}$ were a complete indication of performance, the loop would run at a MFLOPS rate of $\frac{2}{3}$ * clock rate (since there are 2 FLOPS per iteration). Thus, based only on the $T_{element}$ time, an 80-MHz DLXV would run this loop at 53 MFLOPS. But the Linpack benchmark, whose core is this computation, runs at only 13 MFLOPS (without some sophisticated compiler optimization we discuss in the Exercises) on an 80-MHz CRAY-1, DLXV's cousin! Let's see what accounts for the difference.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 416

The Peak Performance of DLXV on SAXPY

First, we should determine what the peak performance, R_∞, really is, since we know it differs from the ideal 53-MFLOPS rate. Figure 7.15 shows the timing within each block of strip-mined code.

Operation		Starts at clock number	Completes at clock number	Comment
LV	V1,Rx	0	12 + 64 = 76	Simple latency
MULTV	a,V1	12 + 1 = 13	13 + 7 + 64 = 84	Chained to LV
LV	V2,Ry	76 + 1 = 77	77 + 12 + 64 = 153	Starts after first LV done (memory contention)
ADDV	V3,V1,V2	77 + 1 + 12 = 90	90 + 6 + 64 = 160	Chained to MULTV and LV
SV	Ry,V3	160 + 1 + 4 =165	165 + 12 + 64 = 241	Must wait on ADDV; not chained (memory contention)

FIGURE 7.15 The SAXPY loop when chained in DLXV. There are three distinct types of delays: 4–clock-cycle delays when a nonchained dependence occurs, latency delays that occur when waiting for a result for the pipeline (6 for add, 7 for multiply, and 12 for memory access), and delays due to contention for the memory pipeline. The last cause is what makes the time per element at least 3 clocks.

From the data in Figure 7.15 and the value of $T_{element}$, we know that

$$T_{start} = 241 - 64 * T_{element} = 241 - 192 = 49$$

This value is equal to the sum of the latencies of the functional units: 12 + 7 + 12 + 6 + 12 = 49.

Using MVL = 64, T_{loop} = 15, T_{base} = 10, and $T_{element}$ = 3 in the performance equation, the time for an n-element operation is

$$T_n = 10 + \left\lceil \frac{n}{64} \right\rceil * (15 + 49) + 3n$$

$$T_n = 10 + n + 64 + 3n = 4n + 74$$

The sustained rate is actually over 4 clock cycles per iteration, rather than the theoretical rate of 3 clocks per iteration, which ignores overhead. The major part of the difference is the cost of the overhead for each block of 64 elements. The basic start-up overhead, T_{base}, adds only $\frac{10}{n}$ to the time for each element. This overhead disappears with long vectors.

We can now compute R_∞ for an 80-MHz clock as

$$R_\infty = \lim_{n \to \infty} \left(\frac{\text{Operations per iteration} * \text{Clock rate}}{\text{Clock cycles per iteration}} \right)$$

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 417

The numerator is independent of n, hence

$$R_\infty = \frac{\text{Operations per iteration} * \text{Clock rate}}{\lim_{n\to\infty} (\text{Clock cycles per iteration})}$$

$$\lim_{n\to\infty} (\text{Clock cycles per iteration}) = \lim_{n\to\infty} \left(\frac{T_n}{n}\right) = \lim_{n\to\infty} \left(\frac{4n+74}{n}\right) = 4$$

$$R_\infty = \frac{2 * 80 \text{ MHz}}{4} = 40 \text{ MFLOPS}$$

Sustained Performance of Linpack on DLXV

The Linpack benchmark is a Gaussian elimination on a 100x100 matrix. Thus, the vector element lengths range from 99 down to 1. A vector of length k is used k times. Thus, the average vector length is given by:

$$\frac{\sum_{i=1}^{99} i^2}{\sum_{i=1}^{99} i} = 66.3$$

Now we can obtain an accurate estimate of the performance of SAXPY using a vector length of 66.

$$T_{66} = 10 + 2 * (15 + 49) + 66 * 3 = 10 + 128 + 198 = 336$$

$$R_{66} = \frac{2 * 66 * 80}{336} \text{MFLOPS} = 31.4 \text{ MFLOPS}$$

In reality, Linpack does not spend all its time in the inner loop. The benchmark's actual performance can be found by taking the weighted harmonic mean of the MFLOPS ratings inside the inner loop (31.4 MFLOPS) and outside that loop (about 0.5 MFLOPS). We can compute the weighting factors by knowing the percentage of the time inside the inner loop after vectorization.

The percentage in the inner loop after vectorization can be obtained using Amdahl's Law if we know the percentage in scalar and the speedup from vectorization. In scalar mode, about 75% of the execution time is spent in the inner loop, and the speedup from vectorization is about 5 times. With this information the percentage of time in the inner loop after vectorization can be computed:

$$\text{Total relative time after vectorization} = \frac{0.75}{5} + 0.25$$

$$= 0.15 + 0.25 = 0.40$$

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 418

Percentage of time in inner loop after vectorization $= \frac{0.15}{0.40} = 37.5\%$

The remaining 62.5% of the time is spent outside the main loop. Thus, the overall MFLOPS rating is

$$\text{Percentage}_{\text{inner}} * \text{MFLOPS}_{\text{inner}} + \text{Percentage}_{\text{other}} * \text{MFLOPS}_{\text{other}}$$

$$= 37.5\% * 31.4 + 62.5\% * 0.5 = 12.1 \text{ MFLOPS}$$

This is comparable to the rate at which the CRAY-1 runs this benchmark.

Example

What is $N_{1/2}$ for just the inner loop of SAXPY for DLXV with an 80-MHz clock?

Answer

Using R_∞ as the peak rate, we want to know the vector length that will achieve about 20 MFLOPS. So,

$$\frac{\text{Clock cycles}}{\text{Iteration}} = \frac{\frac{\text{FLOPS}}{\text{Iteration}} * \frac{\text{Clocks}}{\text{Second}}}{\frac{\text{FLOPS}}{\text{Second}}}$$

$$= \frac{2 * 80 \text{ MHz}}{20 \text{ MFLOPS}} = 8$$

Hence, a rate of 20 MFLOPS means that a loop iteration completes every 8 clock cycles on average, or that $\frac{T_n}{n} = 8$. Using our equation and assuming that $n \leq 64$,

$$T_n = 10 + 1 * 64 + 3 * n$$

Substituting for T_n in the first equation, we obtain

$$8n = 74 + 3 * n$$

$$5n = 74$$

$$n = 14.8$$

So $N_{1/2} = 15$; that is, a vector of length 15 gives approximately one-half the peak performance for the SAXPY loop on DLXV.

Example

What is the vector length, N_v, such that the vector operation runs faster than the scalar?

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 419

Answer

Again, we know that $N_v < 64$. The time to do one iteration in scalar mode can be estimated as 10 + 12 + 12 + 7 + 6 = 47 clocks, where 10 is the estimate of the loop overhead, known to be somewhat less than the strip-mining loop overhead. In the last problem, we showed that this vector loop runs in vector mode in time $T_n = 74 + 3*n$ clock cycles for a vector of length ≤ 64. Therefore,

$$74 + 3n = 47n$$

$$n = \frac{74}{44}$$

$$N_v = 2$$

For the SAXPY loop, vector mode is faster than scalar as long as the vector has at least two elements. This number is surprisingly small, as we will see in the next section (Fallacies and Pitfalls).

SAXPY Performance on an Enhanced DLXV

SAXPY, like many vector problems, is memory limited. Consequently, performance could be improved by adding more memory-access pipelines. This is the major architectural difference between the CRAY X-MP and the CRAY-1. The CRAY X-MP has three memory pipelines, compared to the CRAY-1's single memory pipeline, and the X-MP has more flexible chaining. How does this affect performance?

Example

What would be the value of T_{66} for SAXPY on DLXV if we added two more memory pipelines?

Answer

Figure 7.16 is a version of Figure 7.15 (page 385), adjusted for multiple memory pipelines.

Operation		Starts at clock number	Completes at clock number	Comment
LV	V1,Rx	0	12 + 64 = 76	Simple latency
MULTV	a,V1	12 + 1 = 13	13 + 7 + 64 = 84	Chained to LV
LV	V2,Ry	2	2 + 12 + 64 = 78	Starts immediately
ADDV	V3,V1,V2	13 + 1 + 7 = 21	21 + 6 + 64 = 91	Chained to MULTV and LV
SV	Ry,V3	21 + 1 + 6 = 28	28 + 12 + 64 = 104	Chained to ADDV

FIGURE 7.16 The SAXPY loop when chained in DLXV with three memory pipelines. The only delays are latency delays that occur when waiting for a result for the pipeline (6 for add, 7 for multiply, and 12 for each memory access).

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 420

With three memory pipelines, the performance is greatly improved. Here's our standard performance equation:

$$T_n = T_{base} + \left\lceil \frac{n}{MVL} \right\rceil * (T_{loop} + T_{start}) + n * T_{element}$$

With three memory pipelines the value of $T_{element}$ becomes 1, so that

$$T_{start} = 104 - 64 * T_{element} = 104 - 64 = 40$$

The reduction in stalls reduces the start-up penalty for each sequence. The values of T_{loop} and T_{base}, 15 and 10, remain the same. Therefore, for an average vector length of 66, we have:

$$T_{66} = T_{base} + \left\lceil \frac{66}{64} \right\rceil * (T_{loop} + T_{start}) + 66 * T_{element}$$

$$T_{66} = 10 + 2 * (15 + 40) + 66 * 1 = 186$$

With three memory pipelines, we have reduced the clock-cycle count for sustained performance from 336 to 186, a factor of 1.8. Note the effect of Amdahl's Law: We improved the theoretical peak rate, as measured by $T_{element}$, by a factor of 3, but only achieved an overall improvement of a factor of 1.8 in sustained performance. Because the speedup outside the inner loop is likely to be less than 1.8, the overall improvement in run time for the benchmark will also be less.

Another improvement could come from allowing the start-up of one loop iteration before another completes. This requires that one vector operation be allowed to begin using a functional unit, before another operation has completed. This complicates the instruction issue logic substantially, but has the advantage that the start-up overhead will only occur once, independent of the vector length. On a long vector the overhead per block ($T_{loop} + T_{start}$) can be completely amortized. In this way a machine with vector registers can have both low start-up overhead for short vectors and high peak performance for very long vectors.

Example

What would be the values of R_∞ and T_{66} for SAXPY on DLXV if we added two more memory pipelines and allowed the strip-mining and start-up overhead to be fully overlapped?

Answer

$$R_\infty = \lim_{n \to \infty} \left(\frac{\text{Operations per iteration} * \text{Clock rate}}{\text{Clock cycles per iteration}} \right)$$

$$\lim_{n \to \infty} (\text{Clock cycles per iteration}) = \lim_{n \to \infty} \left(\frac{T_n}{n} \right)$$

Since $T_n = n + 40 + 10 + 15 = n + 65$,

$$\lim_{n\to\infty}\left(\frac{T_n}{n}\right) = \lim_{n\to\infty}\left(\frac{n+65}{n}\right) = 1$$

$$R_\infty = \frac{2 * 80\ \text{MHz}}{1} = 160\ \text{MFLOPS}$$

Thus, adding the extra memory pipelines and more flexible issue logic yields an improvement in peak performance of a factor of 4. However, $T_{66} = 131$, so for shorter vectors, the sustained performance improvement is about 40%.

In summary, we have examined several measures of vector performance. Theoretical peak performance can be calculated based purely on the value of $T_{element}$ as

$$\frac{\text{Number of FLOPS per iteration} * \text{Clock rate}}{T_{element}}$$

By including the loop overhead, we can calculate values for peak performance for an infinite-length vector (R_∞), and also for sustained performance R_n for a vector of length n, which is computed as:

$$R_n = \frac{\text{Number of FLOPS per iteration} * n * \text{Clock rate}}{T_n}$$

Using these measures we also can find $N_{1/2}$ and N_v, which give us another way of looking at the start-up overhead for vectors and the ratio of vector to scalar speed. A wide variety of measures of performance of vector machines are useful in understanding the wide range of performance that applications may see on a vector machine.

7.8 Fallacies and Pitfalls

Pitfall: Concentrating on peak performance and ignoring start-up overhead.

Early vector machines such as the TI ASC and the CDC STAR-100 had long start-up times. For some vector problems, N_v could be greater than 100! Today, the Japanese supercomputers often have higher sustained rates than the Cray Research machines. But with start-up overheads that are 50–100% higher, the faster sustained rates often provide no real advantage. On the CYBER-205 the start-up overhead for SAXPY is 158 clock cycles, substantially increasing the break-even point. With a single vector unit, which contains 2 memory pipelines, the CYBER-205 can sustain a rate of 2 clocks per iteration. The time for SAXPY for a vector of length n is therefore roughly $158 + 2n$. If the clock rates

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 422

of the CRAY-1 and the CYBER-205 were identical, the CRAY-1 would be faster until $n > 64$. Because the CRAY-1 clock is also faster (even though the 205 is newer), the crossover point is over 100. Comparing a four-vector-pipeline CYBER-205 (the maximum-size machine) to the CRAY X-MP that was delivered shortly after the 205, the 205 completes two results per clock cycle—twice as fast as the X-MP. However, vectors must be longer than about 200 for the CYBER-205 to be faster. The problem of start-up overhead has been the major difficulty for the memory–memory vector architectures.

Pitfall: Increasing vector performance, without comparable increases in scalar performance.

This is another area where Seymour Cray rewrote the rules. Many of the early vector machines had comparatively slow scalar units (as well as large start-up overheads). Even today, machines with higher peak vector performance, can be outperformed by a machine with lower vector performance but better scalar performance. Good scalar performance keeps down overhead costs (strip mining, for example) and reduces the impact of Amdahl's Law. A good example of this comes from comparing a fast scalar machine and a vector machine with lower scalar performance. The Livermore FORTRAN kernels are a collection of 24 scientific kernels with varying degrees of vectorization (see Chapter 2; Section 2.2). Figure 7.17 shows the performance of two different machines on this benchmark. Despite the vector machine's higher peak performance, its low scalar performance makes it slower than a fast scalar machine. The next fallacy is closely related.

Machine	Minimum rate for any loop	Maximum rate for any loop	Harmonic mean of all 24 loops
MIPS M/120-5	0.80 MFLOPS	3.89 MFLOPS	1.85 MFLOPS
Stardent-1500	0.41 MFLOPS	10.08 MFLOPS	1.72 MFLOPS

FIGURE 7.17 Performance measurements for the Livermore FORTRAN kernels on two different machines. Both the MIPS M/120-5 and the Stardent-1500 (formerly the Ardent Titan-1) use a 16.7-MHz MIPS R2000 chip for the main CPU. The Stardent-1500 uses its vector unit for scalar FP and has about half the scalar performance (as measured by the minimum rate) of the MIPS M/120, which uses the MIPS R2010 FP chip. The vector machine is more than a factor of 2.5 times faster for a highly vectorizable loop (maximum rate). However, the lower scalar performance of the Stardent-1500 negates the higher vector performance when total performance is measured by the harmonic mean on all 24 loops.

*Fallacy: The scalar performance of the **best** supercomputers is low.*

The supercomputers from Cray Research have always had good scalar performance. Measurements of the CRAY Y-MP running (the nonvectorizable) Spice benchmark show this. When our Spice benchmark is run on the CRAY Y-MP in scalar mode it executes 665 million instructions, with a CPI of 4.1. By comparison, the DECstation 3100 executes 738 million instructions with a CPI of 2.1.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 423

Although the DECstation uses fewer cycles, the Y-MP uses fewer instructions and is much faster overall, since it has a clock cycle one-tenth as long.

Fallacy: You can get vector performance without providing memory bandwidth.

As we saw with the SAXPY loop, memory bandwidth is quite important. SAXPY requires 1.5 memory references per floating-point operation, and this ratio is typical of many scientific codes. Even if the floating-point operations took no time, a CRAY-1 could not increase the performance of the vector sequence used, since it is memory limited. Recently, the CRAY-1 performance on Linpack has jumped because the compiler used clever transformations to change the computation so that values could be kept in the vector registers. This lowered the number of memory references per FLOP and improved the performance by nearly a factor of 2! Thus, the memory bandwidth on the CRAY-1 became sufficient for a loop that formerly required more bandwidth.

7.9 Concluding Remarks

In the late 1980s rapid performance increases in efficiently pipelined scalar machines lead to a dramatic closing of the gap between vector supercomputers, costing millions of dollars, and fast, pipelined, VLSI microprocessors costing less than $100,000. The basic reason for this was the rapidly decreasing CPI of the scalar machines.

For scientific programs, an interesting counterpart to CPI is clock cycles per FLOP, or CPF. We saw in this chapter that for vector machines this number was typically in the range of 2 (for a CRAY X-MP style machine) to 4 (for a CRAY-1 style machine). In the last chapter, we saw that the pipelined machine varied from about 6 (for DLX) down to about 2.5 (for a superscalar DLX with no memory system losses running a SAXPY-type loop).

Recent trends in vector machine design have focused on high peak-vector performance and multiprocessing. Meanwhile, high-speed scalar machines concentrate on keeping the ratio of peak to sustained performance near one. Thus, if the peak rates advance comparably, the sustained rates of the scalar machines will advance more quickly, and the scalar machines will continue to close the CPF gap. These multiple-issue scalar machines can rival or exceed the performance of vector machines with comparable clock speeds, especially for levels of vectorization below 70%. Furthermore, the differences in clock rate are largely technology driven—the low-end, microprocessor-based vector machines have clock rates comparable to the pipelined machines using microprocessor technology. (In fact, they often use the same microprocessors!) In the future, we can expect high-speed pipelined scalar machines to be built with clock rates that will rival those of the current vector supercomputers. However, the vector machines

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 424

should retain a performance advantage for problems with very long vectors that can use multiple memory pipelines and achieve performance close to the peak.

The 1990s will be interesting as the pipelined scalar machines that exploit more instruction-level parallelism and are usually much cheaper (because their peak performance and hence total hardware is much less) begin to offer performance levels for many applications that are difficult to distinguish from those of vector machines.

7.10 Historical Perspective and References

The first vector machines were the CDC STAR-100 (see Hintz and Tate [1972]) and the TI ASC (see Watson [1972]), both announced in 1972. Both were memory–memory vector machines. They had relatively slow scalar units—the STAR used the same units for scalars and vectors—making the scalar pipeline extremely deep. Both machines had high start-up overhead and worked on vectors of several hundred to several thousand elements. The crossover between scalar and vector could be over 50 elements. It appears that not enough attention was paid to the role of Amdahl's Law on these two machines.

Cray, who worked on the 6600 and the 7600 at CDC, founded Cray Research and introduced the CRAY-1 in 1976 (see Russell [1978]). The CRAY-1 used a vector-register architecture to significantly lower start-up overhead. He also had efficient support for nonunit stride and invented chaining. Most importantly, the CRAY-1 was also the fastest scalar machine in the world at that time. This matching of good scalar and vector performance was probably the most significant factor in making the CRAY-1 a success. Some customers bought the machine primarily for its outstanding scalar performance. Many subsequent vector machines are based on the architecture of this first commercially successful vector machine. Baskett and Keller [1977] is a good evaluation of the CRAY-1.

In 1981, CDC started shipping the CYBER-205 (see Lincoln [1982]). The 205 had the same basic architecture as the STAR, but offered improved performance all around as well as expansibility of the vector unit with up to four vector pipelines, each with multiple functional units and a wide load/store pipe that provided multiple words per clock. The peak performance of the CYBER-205 greatly exceeded the performance of the CRAY-1. However, on real programs, the performance difference was much smaller.

The CDC STAR machine and its descendant, the CYBER-205, were memory–memory vector machines. To keep the hardware simple and support the high bandwidth requirements (up to 3 memory references per FLOP), these machines did not efficiently handle nonunit stride. While most loops have unit stride, a nonunit stride loop had poor performance on these machines because memory-to-memory data movements were required to gather together (and scatter back) the nonadjacent vector elements.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 425

Schneck [1987] described several of the early pipelined machines (e.g., Stretch) through the first vector machines including the 205 and CRAY-1. Dongarra [1986] did another good survey, focusing on more recent machines.

In 1983, Cray shipped the first CRAY X-MP (see Chen [1983]). With an improved clock rate (9.5 ns versus 12.5 on the CRAY-1), better chaining support, and multiple memory pipelines, this machine maintained the Cray Research lead in supercomputers. The CRAY-2, a completely new design configurable with up to four processors, was introduced later. It has a much faster clock than the X-MP, but also much deeper pipelines. The CRAY-2 lacks chaining, has an enormous memory latency, and has only one memory pipe per processor. In general, it is only faster than the CRAY X-MP on problems that require its very large main memory.

In 1983, the Japanese computer vendors entered the supercomputer marketplace, starting with the Fujitsu VP100 and VP200 (Miura and Uchida [1983]), and later expanding to include the Hitachi S810, and the NEC SX/2 (see Watanabe [1987]). These machines have proved to be close to the CRAY X-MP in performance. In general, these three machines have much higher peak performance than the CRAY X-MP, though because of large start-up overhead, their typical performance is often lower than the CRAY X-MP (see Figure 2.24 in Chapter 2). The CRAY X-MP favored a multiple-processor approach, first offering a two-processor version and later a four-processor machine. In contrast, the three Japanese machines had expandable vector capabilities. In 1988, Cray Research introduced the CRAY Y-MP—a bigger and faster version of the X-MP. The Y-MP allows up to 8 processors and lowers the cycle time to 6 ns. With a full complement of 8 processors, the Y-MP is generally the fastest supercomputer, though the single-processor Japanese supercomputers may be faster than a one-processor Y-MP. In late 1989 Cray Research was split into two companies, both aimed at building high-end machines available in the early 1990s. Seymour Cray continues to head the spin-off, which is now called Cray Computer Corporation.

In the early 1980s, CDC spun out a group, called ETA, to build a new supercomputer, the ETA-10, capable of 10 GigaFLOPs. The ETA machine delivered in the late 1980s (see Fazio [1987]) used low-temperature CMOS in a configuration with up to 10 processors. Each processor retained the memory–memory architecture based on the CYBER-205. Although the ETA-10 achieved enormous peak performance, its scalar speed was not comparable. In 1989 CDC, the first supercomputer vendor, closed ETA and left the supercomputer design business.

In 1986, IBM introduced the System/370 vector architecture (see Moore et al. [1987]) and its first implementation in the 3090 Vector Facility. The architecture extends the System/370 architecture with 171 vector instructions. The 3090/VF is integrated into the 3090 CPU. Unlike most other vector machines, the 3090/VF routes its vectors through the cache.

The 1980s also saw the arrival of smaller-scale vector machines, called minisupercomputers. Priced at roughly one-tenth the cost of a supercomputer ($0.5 to

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 426

$1 million versus $5 to $10 million), these machines caught on quickly. Although many companies joined the market, the two companies that have been most successful are Convex and Alliant. Convex started with a uniprocessor vector machine (C-1) and now offers a small multiprocessor (C-2); they emphasize Cray software capability. Alliant [1987] has concentrated more on the multiprocessor aspects; they build an eight-processor machine, with each processor offering vector capability.

The basis for modern vectorizing compiler technology and the notion of data dependence was developed by Kuck and his colleagues [1974] at the University of Illinois. Banerjee [1979] developed the test named after him. Padua and Wolf [1986] gave a good overview of vectorizing compiler technology.

Benchmark studies of various supercomputers including attempts to understand the performance differences have been undertaken by Lubeck, Moore and Mendez [1985], Bucher [1983], and Jordan [1987]. In Chapter 2, we discussed several benchmark suites aimed at scientific usage and often employed for supercomputer benchmarking, including Linpack, the Lawrence Livermore Laboratories FORTRAN kernels, and the Perfect Club suite.

In the late 1980s, graphics supercomputers arrived on the market from Stellar [Sporer, Moss, and Mathais 1988] and Ardent [Miranker, Rubenstein, and Sanguinetti 1988]. The Stellar machine used a timeshared pipeline to allow high-speed vector processing and efficient multitasking. This approach was used earlier in a machine designed by B. J. Smith [1981] called the HEP and built by Denelcor in the mid-1980s. This approach does not yield high-speed scalar performance, as evident in the scalar benchmarks of the Stellar machine. The Ardent machine combines a RISC processor (the MIPS R2000) with a custom vector unit. These vector machines, which cost about $100K, brought vector capabilities to a new potential market. In late 1989, Stellar and Ardent were merged to form Stardent, and the Ardent architecture is being shipped from the combined company.

From this overview we can see the progress vector machines have made. In less than 20 years they have gone from unproven, new architectures to playing a significant role in the goal to provide engineers and scientists with ever larger amounts of computing power.

References

ALLIANT COMPUTER SYSTEMS CORP. [1987]. *Alliant FX/Series: Product Summary* (June), Acton, Mass.

BANERJEE, U. [1979]. *Speedup of Ordinary Programs*, Ph.D. Thesis, Dept. of Computer Science, Univ. of Illinois at Urbana-Champaign (October).

BASKETT, F. AND T. W. KELLER [1977]. "An Evaluation of the CRAY-1 Computer," in *High Speed Computer and Algorithm Organization*, Kuck, D. J., Lawrie, D. H. and A. H. Sameh, eds., Academic Press, 71-84.

BUCHER, I. Y. [1983]. "The computational speed of supercomputers," *Proc. SIGMETRICS Conf. on Measuring and Modeling of Computer Systems*, ACM (August) 151–165.

CALLAHAN, D., J. DONGARRA, AND D. LEVINE [1988]. "Vectorizing compilers: A test suite and results," *Supercomputing '88*, ACM/IEEE (November), Orlando, Fla., 98–105.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 427

CHEN, S. [1983]. "Large-scale and high-speed multiprocessor system for scientific applications," *Proc. NATO Advanced Research Work on High Speed Computing* (June); also in K. Hwang, ed., "Supercomputers: Design and applications," *IEEE* (August) 1984.

DONGARRA, J. J. [1986]. "A survey of high performance computers," *COMPCON, IEEE* (March) 8–11.

FAZIO, D. [1987]. "It's really much more fun building a supercomputer than it is simply inventing one," *COMPCON, IEEE* (February) 102-105.

FLYNN, M. J. [1966]. "Very high-speed computing systems," *Proc. IEEE* 54:12 (December) 1901–1909.

HINTZ, R. G. AND D. P. TATE [1972]. "Control data STAR-100 processor design," *COMPCON, IEEE* (September) 1–4.

JORDAN, K. E. [1987]. "Performance comparison of large-scale scientific computers: Scalar mainframes, mainframes with vector facilities, and supercomputers," *Computer* 20:3 (March) 10–23.

KUCK, D., P. P. BUDNIK, S.-C. CHEN, D. H. LAWRIE, R. A. TOWLE, R. E. STREBENDT, E. W. DAVIS, JR., J. HAN, P. W. KRASKA, Y. MURAOKA [1974]. "Measurements of parallelism in ordinary FORTRAN programs," *Computer* 7:1 (January) 37–46.

LINCOLN, N. R. [1982]. "Technology and design trade offs in the creation of a modern supercomputer," *IEEE Trans. on Computers* C-31:5 (May) 363–376.

LUBECK, O., J. MOORE, AND R. MENDEZ [1985]. "A benchmark comparison of three supercomputers: Fujitsu VP-200, Hitachi S810/20, and CRAY X-MP/2," *Computer* 18:1 (January) 10–29.

MIRANKER, G. S., J. RUBENSTEIN, AND J. SANGUINETTI [1988]. "Squeezing a Cray-class supercomputer into a single-user package," *COMPCON, IEEE* (March) 452–456.

MIURA, K. AND K. UCHIDA [1983]. "FACOM vector processing system: VP100/200," *Proc. NATO Advanced Research Work on High Speed Computing* (June); also in K. Hwang, ed., "Supercomputers: Design and applications," *IEEE* (August 1984) 59–73.

MOORE, B., A. PADEGS, R. SMITH, AND W. BUCHOLZ [1987]. "Concepts of the System/370 vector architecture," *Proc. 14th Symposium on Computer Architecture* (June), ACM/IEEE, Pittsburgh, Pa., 282–292.

PADUA, D. AND M. WOLFE [1986]. "Advanced compiler optimizations for supercomputers," *Comm. ACM* 29:12 (December) 1184–1201.

RUSSELL, R. M. [1978]. "The CRAY-1 computer system," *Comm. of the ACM* 21:1 (January) 63–72.

SCHNECK, P. B. [1987]. *Supercomputer Architecture,* Kluwer Academic Publishers, Norwell, Mass.

SMITH, B. J. [1981]. "Architecture and applications of the HEP multiprocessor system," *Real-Time Signal Processing IV* 298 (August) 241–248.

SPORER, M., F. H. MOSS AND C. J. MATHAIS [1988]. "An introduction to the architecture of the Stellar Graphics supercomputer," *COMPCON, IEEE* (March) 464–467.

WATANABE, T. [1987]. "Architecture and performance of the NEC supercomputer SX system," *Parallel Computing* 5, 247–255.

WATSON, W. J. [1972]. "The TI ASC–A highly modular and flexible super computer architecture," *Proc. AFIPS Fall Joint Computer Conf.,* 221–228.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 428

EXERCISES

In these Exercises assume DLXV has a clock rate of 80 MHz and that T_{base} = 10 and T_{loop} = 15. Also assume that the store latency is always included in the running time.

7.1 [10] <7.1–7.2> Write a DLXV vector sequence that achieves the peak MFLOPS performance of the machine (use the functional unit and instruction description in Section 7.2). Assuming an 80-MHz clock rate, what is the peak MFLOPS?

7.2 [20/15/15] <7.1–7.6> Consider the following vector code run on an 80-MHz version of DLXV for a fixed vector length of 64:

```
LV      V1,Ra
MULTV   V2,V1,V3
ADDV    V4,V1,V3
SV      Rb,V2
SV      Rc,V4
```

Ignore all strip-mining overhead, but assume that the store latency must be included in the time to perform the loop. The entire sequence produces 64 results.

a. [20] Assuming no chaining and a single memory pipeline, how many clock cycles per result (including both stores as one result) does this vector sequence require?

b. [15] If the vector sequence is chained, how many clock cycles per result does this sequence require?

c. [15] Suppose DLXV had three memory pipelines and chaining. If there were no bank conflicts in the accesses for the above loop, how many clock cycles are required per result for this sequence?

7.3 [20/20/15/15/20/20/20] <7.2–7.7> Consider the following FORTRAN code:

```
      do 10 i=1,n
         A(i) = A(i) + B(i)
         B(i) = x * B(i)
10    continue
```

Use the techniques of Section 7.7 to estimate performance throughout this exercise assuming an 80-MHz version of DLXV.

a. [20] Write the best DLXV vector code for the inner portion of the loop. Assume x is in F0 and the addresses of `A` and `B` are in `Ra` and `Rb`, respectively.

b. [20] Find the total time for this loop on DLXV (T_{100}). What is the MFLOP rating for the loop (R_{100})?

c. [15] Find R_∞ for this loop.

d. [15] Find $N_{1/2}$ for this loop.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 429

e. [20] Find N_v for this loop. Assume the scalar code has been pipeline scheduled so that each memory reference takes six cycles and each FP operation takes 3 cycles. Assume the scalar overhead is also T_{loop}.

f. [20] Assume DLXV has two memory pipelines. Write vector code that takes advantage of the second memory pipeline.

g. [20] Compute T_{100} and R_{100} for DLX with two memory pipelines.

7.4 [20/10] <7.3> Suppose we have a version of DLXV with eight memory banks (each a doubleword wide) and a memory-access time of eight cycles.

a. [20] If a load vector of length 64 is executed with a stride of 20 doublewords, how many cycles will the load take to complete?

b. [10] What percentage of the memory bandwidth do you achieve on a 64-element load at stride 20 versus stride 1?

7.5 [12/12/20] <7.4–7.7> Consider the following loop:

```
      C = 0.0
      do 10 i=1,64
          A(i) = A(i) + B(i)
          C = C + A(i)
10    continue
```

a. [12] Split the loop into two loops: one with no dependence and one with a dependence. Write these loops in FORTRAN—as a source-to-source transformation. This optimization is called *loop fission.*

b. [12] Write the DLXV vector code for the loop without a dependence.

c. [20] Write the DLXV code to evaluate the dependent loop using recursive doubling.

7.6 [20/15/20/20] <7.5–7.7> The compiled Linpack performance of the CRAY-1 (designed in 1976) was almost doubled by a better compiler in 1989. Let's look at a simple example of how this might occur. Consider the "SAXPY–like" loop (where *k* is a parameter to the procedure containing the loop):

```
      do 10 i=1,64
          do 10 j=1,64
          Y(k,j) = a*X(i,j) + Y(k,j)
10    continue
```

a. [20] Write the **straightforward** code sequence for just the inner loop in DLXV vector instructions.

b. [15] Using the techniques of Section 7.7, estimate the performance of this code on DLXV by finding T_{64} in clock cycles. You may assume that T_{base} applies once and T_{loop} of overhead is incurred for each iteration of the outer loop. What limits the performance?

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 430

c. [20] Rewrite the DLXV code to reduce the performance limitation; show the resulting inner loop in DLXV vector instructions. (Hint: think about what establishes $T_{element}$; can you affect it?) Find the total time for the resulting sequence.

d. [20] Estimate the performance of your new version using the techniques of Section 7.7 and finding T_{64}.

7.7 [15/15/25] <7.6> Consider the following code.

```
      do 10 i=1,64
          if (B(i) .ne. 0) then
              A(i) = A(i) / B(i)
          endif
10    continue
```

Assume that the addresses of `A` and `B` are in `Ra` and `Rb`, respectively, and that F0 contains 0.

a. [15] Write the DLXV code for this loop using the vector-mask capability.

b. [15] Write the DLXV code for this loop using scatter/gather.

c. [25] Estimate the performance (T_{100} in clock cycles) of these two vector loops assuming a divide latency of 20 cycles. Assume that all vector instructions run at one result per clock, independent of the setting of the vector-mask register. Assume that 50% of the entries of B are 0. Considering hardware costs, which would you build if the above loop **was** typical?

7.8 [15/20/15/15] <7.1–7.7> In Figure 2.24 of Chapter 2 (page 75), we saw that the difference between peak and sustained performance could be large: For one problem, a Hitachi S810 had a peak speed twice as high as the CRAY X-MP, while for another more realistic problem the CRAY X-MP was twice as fast as the Hitachi machine. Let's examine why this might occur using two versions of DLXV and the following code sequences:

```
C       Code sequence 1
        do 10 i=1,10000
           A(i) = x * A(i) + y * A(i)
10      continue

C       Code sequence 2
        do 10 i=1,100
           A(i) = x * A(i)
10      continue
```

Assume there is a version of DLXV (call it DLXVII) that has two copies of every floating-point functional unit with full chaining among them. Assume that both DLXV and DLXVII have two load/store units. Because of the extra functional units and the increased complexity of assigning operations to units, all the overheads (T_{base}, T_{loop}, and the start-up overheads per vector operation) are doubled.

a. [15] Find the number of clock cycles for code sequence 1 on DLXV.

b. [20] Find the number of clock cycles on code sequence 1 for DLXVII. How does this compare to DLXV?

c. [15] Find the number of clock cycles on code sequence 2 for DLXV.

d. [15] Find the number of clock cycles on code sequence 2 for DLXVII. How does this compare to DLXV?

7.9 [15/15/20] <7.5> In this problem we will examine some of the vector loop tests discussed in Section 7.5 and summarized in Figure 7.13 (page 377).

a. [15] Here is a simple code fragment:

```
          do 400 i = 2,100,2
             a(i-1) = a(50*i+1)
400       continue
```

To use the GCD test this loop must first be "normalized"—written so that the index starts at 1 and increments by 1 on every iteration. Write a normalized version of the loop (change the indices as needed), then use the GCD test to see if it vectorizes.

b. [15] Here is another loop:

```
          do 400 i = 2,100,2
             a(i) = a(i-1)
400       continue
```

Normalize the loop and use the GCD test to detect a dependence. Is there a real dependence in this loop?

c. [20] Here is a tricky piece of code with two-dimensional arrays. Can it be vectorized? If so, how? Rewrite the **source** code so that it is clear that the loop can be vectorized, if possible.

```
          do 290 j = 2,n
             do 290 i = 2,j
                aa(i,j)=aa(i-1,j)*aa(i-1,j)+bb(i,j)
290       continue
```

7.10 [25] <7.5> Show that if for two array elements A(a*i +b) and A(c*i+d) there is a true dependence, then GCD(c,a) divides (d–b).

7.11 [12/15] <7.5> Consider the following loop:

```
          do 10 i = 2,n
             A(i) = B
10           C(i) = A(i-1)
```

a. [12] Show there is a loop-carried dependence in this code fragment.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 432

b. [15] Rewrite the code in Fortran so that it can be vectorized as two separate vector sequences.

7.12 [25] <7.6> Because the difference between vector and scalar modes is so large on a supercomputer and the machines often cost tens of millions of dollars, programmers are frequently willing to go to extraordinary effort to achieve good performance. This often includes tricky assembly language programming. An interesting problem is to write a vectorizable sort for floating-point numbers—a task sometimes required in scientific code. Choose a sorting algorithm and write a version for DLXV that uses vector operations as much as possible. (Hint: One good choice is quicksort where the vector compares and compress/expand capability can be used.)

7.13 [25] <7.6> In some vector machines, the vector registers are addressable, and the operands to a vector operation may be two different parts of the same vector register. This allows another solution for the reduction shown on page 382. The key idea in partial sums is to reduce the vector to m sums where m is the total latency through the vector functional unit including the operand read and write times. Assume that the DLXV vector registers are addressable (e.g., you can initiate a vector operation with the operand V1(16), indicating that the input operand began with element 16). Also, assume that the total latency for adds including operand read and write is eight cycles. Write a DLXV code sequence that reduces the contents of V1 to eight partial sums. It can be done with one vector operation.

7.14 [40] <7.2–7.6> Extend the DLX simulator to be a DLXV simulator including the ability to count clock cycles. Write some short benchmark programs in DLX and DLXV assembly language. Measure the speedup on DLXV, the percentage of vectorization, and usage of the functional units.

7.15 [50] <7.5> Modify the DLX compiler to include a dependence checker. Run some scientific code and loops through it and measure what percentage of the statements could be vectorized.

7.16 [Discussion] Some proponents of vector machines might argue that the vector processors have provided the best path to ever-increasing amounts of computer power by focusing their attention on boosting peak vector performance. Others would argue that the emphasis on peak performance is misplaced because an increasing percentage of the programs are dominated by nonvector performance. (Remember Amdahl's Law?) The proponents would respond that programmers should work to make their programs vectorizable. What do you think about this argument?

7.17 [Discussion] Consider the points raised in the Concluding Remarks (Section 7.9). This topic—the relative advantages of pipelined scalar machines versus FP vector machines—is the source of much debate in the early 1990s. What advantages do you see for each side? What would you do in this situation?

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 433

Ideally one would desire an indefinitely large memory capacity such that any particular . . . word would be immediately available. . . . We are . . . forced to recognize the possibility of constructing a hierarchy of memories, each of which has greater capacity than the preceding but which is less quickly accessible.

A. W. Burks, H. H. Goldstine, and J. von Neumann,
Preliminary Discussion of the Logical Design of an Electronic Computing Instrument (1946)

8.1	**Introduction: Principle of Locality**	**403**
8.2	**General Principles of Memory Hierarchy**	**404**
8.3	**Caches**	**408**
8.4	**Main Memory**	**425**
8.5	**Virtual Memory**	**432**
8.6	**Protection and Examples of Virtual Memory**	**438**
8.7	**More Optimizations Based on Program Behavior**	**449**
8.8	**Advanced Topics—Improving Cache-Memory Performance**	**454**
8.9	**Putting It All Together: The VAX-11/780 Memory Hierarchy**	**475**
8.10	**Fallacies and Pitfalls**	**480**
8.11	**Concluding Remarks**	**484**
8.12	**Historical Perspective and References**	**485**
	Exercises	**490**

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.

8 Memory-Hierarchy Design

8.1 Introduction: Principle of Locality

Computer pioneers correctly predicted that programmers would want unlimited amounts of fast memory. As the 90/10 rule in the first chapter predicts, most programs fortunately do not access all code or data uniformly (see Section 1.3, pages 8–12). The 90/10 rule can be restated as the *principle of locality*. This hypothesis, which holds that all programs favor a portion of their address space at any instant of time, has two dimensions:

- *Temporal locality* (locality in time)—If an item is referenced, it will tend to be referenced again soon.
- *Spatial locality* (locality in space)—If an item is referenced, nearby items will tend to be referenced soon.

A memory hierarchy is a natural reaction to locality and technology. The principle of locality and the guideline that smaller hardware is faster yield the concept of a hierarchy based on different speeds and sizes. Since slower memory is cheaper, a memory hierarchy is organized into several levels—each smaller, faster, and more expensive per byte than the level below. The levels of the hierarchy subset one another; all data in one level is also found in the level below, and all data in that lower level is found in the one below it, and so on until we reach the bottom of the hierarchy.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 435

This chapter includes a half-dozen examples that demonstrate how taking advantage of the principle of locality can improve performance. All these strategies map addresses from a larger memory to a smaller but faster memory. As part of address mapping, the memory hierarchy is usually given the responsibility of address checking; protection schemes used for doing this are covered in this chapter. Later we will explore advanced memory hierarchy topics and trace a memory access through three levels of memory on the VAX-11/780.

8.2 General Principles of Memory Hierarchy

Before proceeding with examples of the memory hierarchy, let's define some general terms applicable to all memory hierarchies. A memory hierarchy normally consists of many levels, but it is managed between two adjacent levels at a time. The *upper* level—the one closer to the processor—is smaller and faster than the *lower* level (see Figure 8.1). The minimum unit of information that can be either present or not present in the two-level hierarchy is called a *block*. The size of a block may be either fixed or variable. If it is fixed, the memory size is a multiple of that block size. Most of this chapter will be concerned with fixed block sizes, although a variable block design is discussed in Section 8.6.

Success or failure of an access to the upper level is designated as a hit or a miss: A *hit* is a memory access found in the upper level, while a *miss* means it is not found in that level. *Hit rate*, or hit ratio—like a batting average—is the fraction of memory accesses found in the upper level. This is sometimes represented as a percentage. *Miss rate* (1.0 – hit rate) is the fraction of memory accesses not found in the upper level.

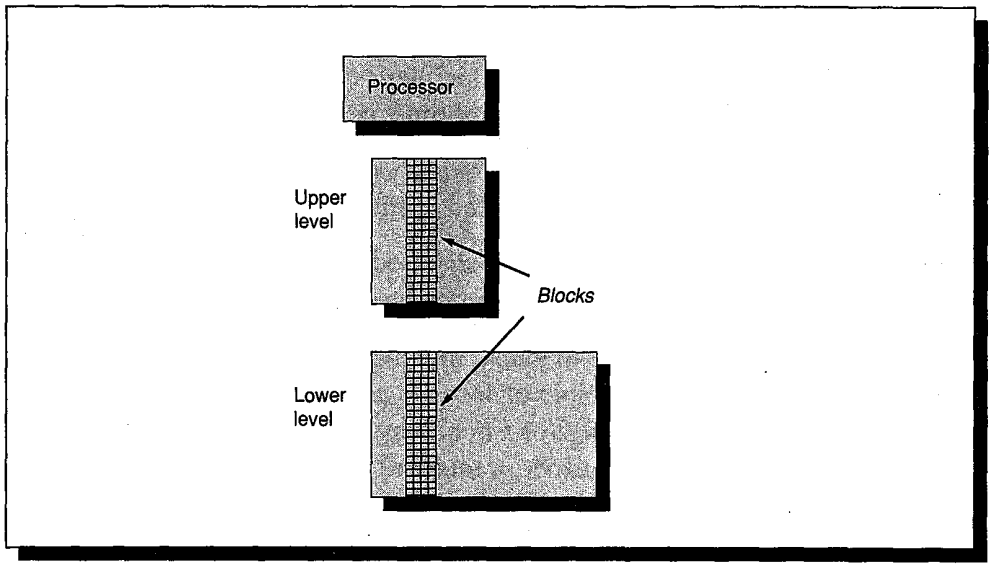

FIGURE 8.1 Every pair of levels in the memory hierarchy can be thought of as having an upper and lower level. Within each level the unit of information that is present or not is called a *block*.

Since performance is the major reason for having a memory hierarchy, the speed of hits and misses is important. *Hit time* is the time to access the upper level of the memory hierarchy, which includes the time to determine whether the access is a hit or a miss. *Miss penalty* is the time to replace a block in the upper level with the corresponding block from the lower level, plus the time to deliver this block to the requesting device (normally the CPU). The miss penalty is further divided into two components: *access time*—the time to access the first word of a block on a miss; and *transfer time*—the additional time to transfer the remaining words in the block. Access time is related to the latency of the lower-level memory, while transfer time is related to the bandwidth between the lower-level and upper-level memories. (Sometimes *access latency* is used to mean access time.)

The memory address is divided into pieces that access each part of the hierarchy. The *block-frame address* is the higher-order piece of the address that identifies a block at that level of the hierarchy (see Figure 8.2). The *block-offset address* is the lower-order piece of the address and identifies an item within a block. The size of the block-offset address is $\log_2$ (size of block); the size of the block-frame address is then the size of the full address at this level less the size of the block-offset address.

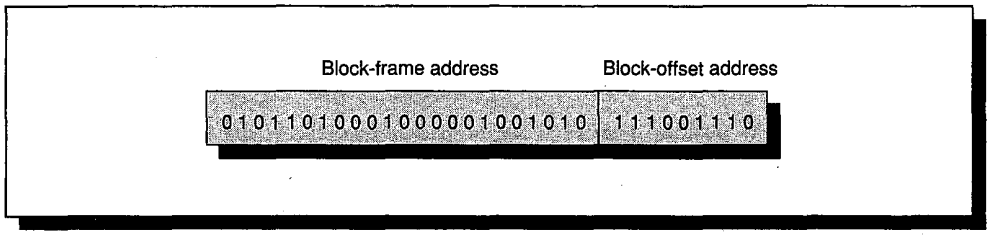

FIGURE 8.2 Example of the frame address and offset address portions of a 32-bit lower-level memory address. In this case the block size is 512, making the size of the offset address 9 bits and the size of the block-frame address 23 bits.

Evaluating Performance of a Memory Hierarchy

Because instruction count is independent of the hardware, it is tempting to evaluate CPU performance using that number. As we saw in Chapters 2 and 4, however, such indirect performance measures have waylaid many a computer designer. The corresponding temptation for evaluating memory-hierarchy performance is to concentrate on miss rate, for it, too, is independent of the speed of the hardware. As we shall see, miss rate can be just as misleading as instruction count. A better measure of memory-hierarchy performance is the average time to access memory:

$$\text{Average memory-access time} = \text{Hit time} + \text{Miss rate} * \text{Miss penalty}$$

The components of average access time can be measured either in absolute time—say, 10 nanoseconds on a hit—or in the number of clock cycles that the

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 437

CPU waits for the memory—such as a miss penalty of 12 clock cycles. Remember that average memory-access time is still an indirect measure of performance; so while it is a better measure than miss rate, it is not a substitute for execution time.

The relationship of block size to miss penalty and miss rate is shown abstractly in Figure 8.3. These representations assume that the size of the upper-level memory does not change. The access-time portion of the miss penalty is not affected by block size, but the transfer time does increase with block size. If access time is large, initially there will be little additional miss penalty relative to access time as block size increases. However, increasing block size means fewer blocks in the upper-level memory. Increasing block size lowers the miss rate until the reduced misses of larger blocks (spatial locality) are outweighed by the increased misses as the number of blocks shrinks (temporal locality).

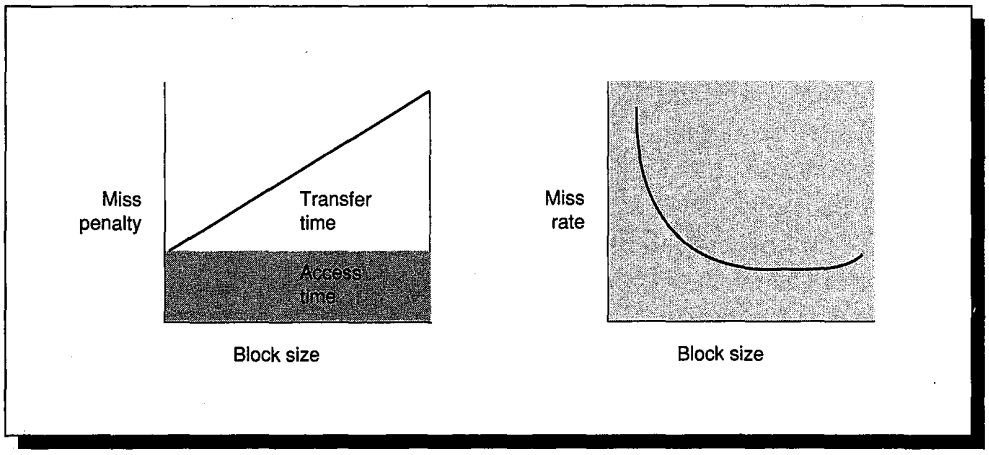

FIGURE 8.3 Block size versus miss penalty and miss rate. The transfer-time portion of the miss penalty obviously grows with increasing block size. For a fixed-size upper-level memory, miss rates fall with increasing block size until so much of the block is not used that it displaces useful information in the upper level, and miss rates begin to rise. The point on the curve on the right where miss rates begin to rise with increasing block size is sometimes called the *pollution point.*

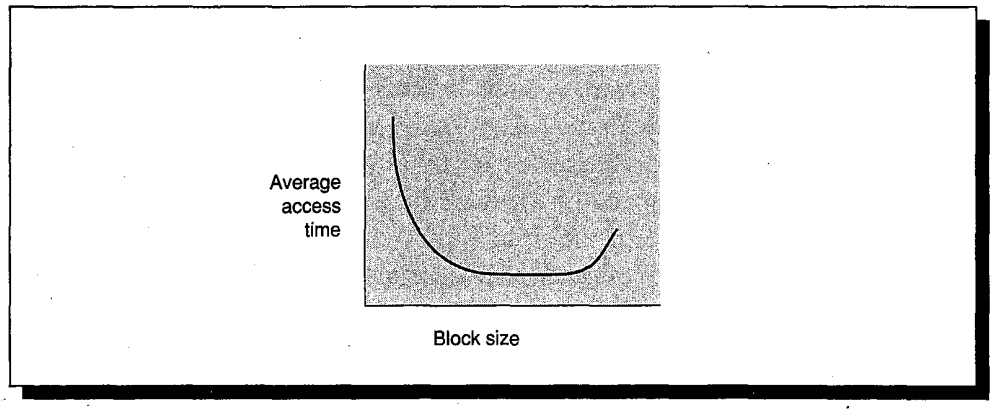

FIGURE 8.4 The relationship between average memory-access time and block size.

The goal of a memory hierarchy is to reduce execution time, not misses. Hence, computer designers favor a block size with the lowest average access time rather than the lowest miss rate. This is related to the product of miss rate and miss penalty, as Figure 8.4 shows abstractly. Of course, overall CPU performance is the ultimate performance test, so care must be taken when reducing average memory-access time to be sure that changes to clock cycle time and CPI improve overall performance as well as average memory-access time.

Implications of a Memory Hierarchy to the CPU

Processors designed without a memory hierarchy are simpler because memory accesses always take the same amount of time. Misses in a memory hierarchy mean that the CPU must be able to handle variable memory-access times. If the miss penalty is on the order of tens of clock cycles, the processor normally waits for the memory transfer to complete. On the other hand, if the miss penalty is thousands of processor clock cycles, it is too wasteful to let the CPU sit idle; in this case, the CPU is interrupted and used for another process during the miss handling. Thus, avoiding the overhead of a long miss penalty means any memory access can result in a CPU interrupt. This also means the CPU must be able to recover any memory address that can cause such an interrupt, so that the system can know what to transfer to satisfy the miss (see Section 5.6). When the memory transfer is complete, the original process is restored, and the instruction that missed is retried.

The processor must also have some mechanism to determine whether or not information is in the top level of the memory hierarchy. This check happens on every memory access and affects hit time; maintaining acceptable performance usually requires the check to be implemented in hardware. The final implication of a memory hierarchy is that the computer must have a mechanism to transfer blocks between upper- and lower-level memory. If the block transfer is tens of clock cycles, it is controlled by hardware; if it is thousands of clock cycles, it can be controlled by software.

Four Questions for Classifying Memory Hierarchies

The fundamental principles that drive all memory hierarchies allow us to use terms that transcend the levels we are talking about. These same principles allow us to pose four questions about any level of the hierarchy:

Q1: Where can a block be placed in the upper level? (*Block placement*)

Q2: How is a block found if it is in the upper level? (*Block identification*)

Q3: Which block should be replaced on a miss? (*Block replacement*)

Q4: What happens on a write? (*Write strategy*)

These questions will help us gain an understanding of the different tradeoffs demanded by the relationships of memories at different levels of a hierarchy.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 439

8.3 Caches

Cache: a safe place for hiding or storing things.

Webster's New World Dictionary of the American Language, Second College Edition (1976)

Cache is the name first chosen to represent the level of the memory hierarchy between the CPU and main memory, and that is the dominant use of the term. While the concept of caches is younger than the IBM 360 architecture, caches appear today in every class of computer and in some computers more than once. In fact, the word has become so popular that it has replaced "buffer" in many computer-science circles.

The general terms defined in the prior section can be used for caches, although the word *line* is often used instead of block. Figure 8.5 shows the typical range of memory-hierarchy parameters for caches.

Block (line) size	4 – 128 bytes
Hit time	1 – 4 clock cycles (normally 1)
Miss penalty	8 – 32 clock cycles
(Access time)	(6 – 10 clock cycles)
(Transfer time)	(2 – 22 clock cycles)
Miss rate	1% – 20%
Cache size	1 KB – 256 KB

FIGURE 8.5 Typical values of key memory-hierarchy parameters for caches in 1990 workstations and minicomputers.

Now let's examine caches in more detail by answering the four memory-hierarchy questions.

Q1: Where Can a Block Be Placed in a Cache?

Restrictions on where a block is placed create three categories of cache organization:

- If each block has only one place it can appear in the cache, the cache is said to be *direct mapped.* The mapping is usually (block-frame address) modulo (number of blocks in cache).
- If a block can be placed anywhere in the cache, the cache is said to be *fully associative.*

- If a block can be placed in a restricted set of places in the cache, the cache is said to be *set associative*. A *set* is a group of two or more blocks in the cache. A block is first mapped onto a set, and then the block can be placed anywhere within the set. The set is usually chosen by bit selection; that is, (block-frame address) modulo (number of **sets** in cache). If there are n blocks in a set, the cache placement is called *n-way set associative*.

The range of caches from direct mapped to fully associative is really a continuum of levels of set associativity: Direct mapped is simply one-way set associative and a fully associative cache with m blocks could be called m-way set associative. Figure 8.6 shows where block 12 can be placed in a cache according to the block-placement policy.

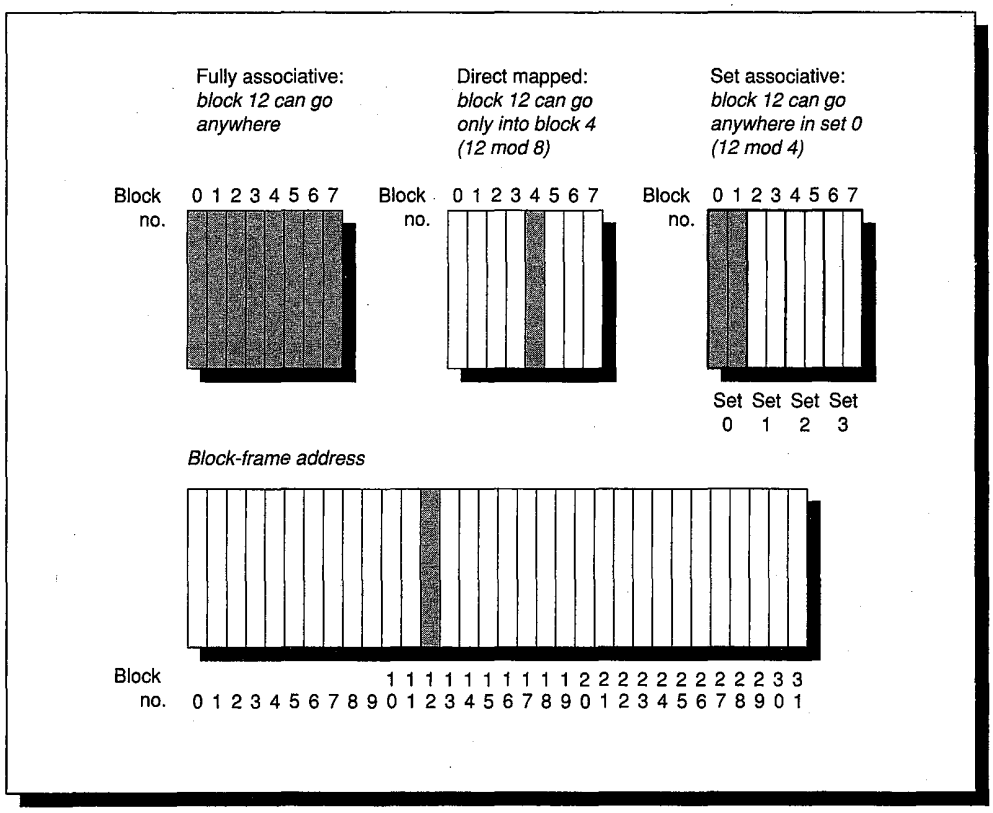

FIGURE 8.6 The cache has 8 blocks, while memory has 32 blocks. The set-associative organization has 4 sets with 2 blocks per set, called two-way set associative. (Real caches contain hundreds of blocks and real memories contain hundreds of thousands of blocks.) Assume that there is nothing in the cache and that the block-frame address in question identifies lower-level block 12. The three options for caches are shown left to right. In fully associative, block 12 from the lower level can go into any of the 8 blocks of the cache. With direct mapped, block 12 can only be placed into block 4 (12 modulo 8). Set associative, which has some of both features, allows the block to be placed anywhere in set 0 (12 modulo 4). With two blocks per set, this means block 12 can be placed either in block 0 or block 1 of the cache.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 441

Q2: How Is a Block Found If It Is in the Cache?

Caches include an address tag on each block that gives the block-frame address. The tag of every cache block that might contain the desired information is checked to see if it matches the block-frame address from the CPU. Figure 8.7 gives an example. Because speed is of the essence, all possible tags are searched in parallel; serial search would make set associativity counterproductive.

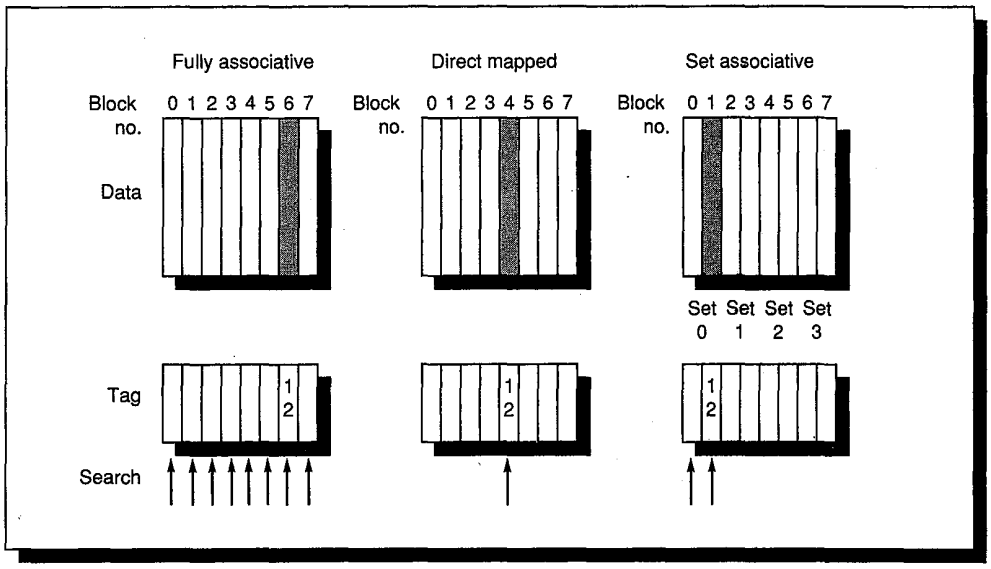

FIGURE 8.7 In fully associative placement, the block for block-frame address 12 can appear in any of the 8 blocks; thus, all 8 tags must be searched. The desired data is found in cache block 6 in this example. In direct-mapped placement there is only one cache block where memory block 12 can be found. In set-associative placement, with 4 sets, memory block 12 must be in set 0 (12 mod 4); thus, the tags of cache blocks 0 and 1 are checked. In this case the data is found in cache block 1. Speed of cache access dictates that searching must be performed in parallel for fully associative and set-associative mappings.

There must be a way to know that a cache block does not have valid information. The most common procedure is to add a *valid bit* to the tag to say whether or not this entry contains a valid address. If the bit is not set, there cannot be a match on this address.

A common omission in finding the cost of caches is to forget the cost of the tag memory. One tag is required for each block. An advantage of increasing block sizes is that the tag overhead per cache entry becomes a smaller fraction of the total cost of the cache.

Before proceeding to the next question, let's explore the relationship of a CPU address to the cache. Figure 8.8 shows how an address is divided into three fields to find data in a set-associative cache: the *block-offset* field used to select the desired data from the block, the *index* field used to select the set, and the *tag* field used for the comparison. While the comparison could be made on more of the address than the tag, there is no need:

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 442

- Checking the index would be redundant, since it was used to select the set to be checked (an address stored in set 0, for example, must have 0 in the index field or it couldn't be stored in set 0).
- The offset is unnecessary in the comparison because all block offsets match and the entire block is present or not.

If the total size is kept the same, increasing associativity increases the number of blocks per set, thereby decreasing the size of the index and increasing the size of the tag. That is, the tag/index boundary in Figure 8.8 moves to the right with increasing associativity.

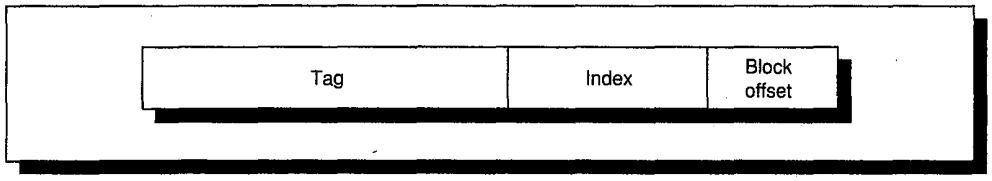

FIGURE 8.8 The 3 portions of an address in a set-associative or direct-mapped cache. The tag is used to check all the blocks in the set and the index is used to select the set. The block offset is the address of the desired data within the block.

Q3: Which Block Should Be Replaced on a Cache Miss?

If the choice were between a block that has valid data and a block that doesn't, then it would be easy to select which block to replace. Alas, the high hit rate of caches means that the overwhelming decision is between blocks that have valid data.

A benefit of direct-mapped placement is that hardware decisions are simplified. In fact, so simple that there is no choice: Only one block is checked for a hit, and only that block can be replaced. With fully associative or set-associative placement, there are several blocks to choose from on a miss. There are two primary strategies employed for selecting which block to replace:

- *Random*—To spread allocation uniformly, candidate blocks are randomly selected. Some systems use a scheme for spreading data across a set of blocks in a pseudorandomized manner to get reproducible behavior, which is particularly useful during hardware debugging.
- *Least-recently used* (LRU)—To reduce the chance of throwing out information that will be needed soon, accesses to blocks are recorded. The block replaced is the one that has been unused for the longest time. This makes use of a corollary of temporal locality: If recently used blocks are likely to be used again, then the best candidate for disposal is the least recently used. Figure 8.9 (page 412) shows which block is the least-recently used for a sequence of block-frame addresses in a fully associative memory hierarchy.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 443

A virtue of random is that it is simple to build in hardware. As the number of blocks to keep track of increases, LRU becomes increasingly expensive and is frequently only approximated. Figure 8.10 shows the difference in miss rates between LRU and random replacement. Replacement policy plays a greater role in smaller caches than in larger caches where there are more choices of what to replace.

Block-frame addresses		3	2	1	0	0	2	3	1	3	0
LRU block number	0	0	0	0	3	3	3	1	0	0	2

FIGURE 8.9 Least-recently used blocks for a sequence of block-frame addresses in a fully associative memory hierarchy. This assumes that there are 4 blocks and that in the beginning the LRU block is number 0. The LRU block number is shown below each new block reference. Another policy, *First-in–first-out* (FIFO), simply discards the block that was used N unique accesses before, independent of its reference pattern in the last N – 1 references. Random replacement generally outperforms FIFO and it is easier to implement.

Associativity:	2-way		4-way		8-way	
Size	LRU	Random	LRU	Random	LRU	Random
16 KB	5.18%	5.69%	4.67%	5.29%	4.39%	4.96%
64 KB	1.88%	2.01%	1.54%	1.66%	1.39%	1.53%
256 KB	1.15%	1.17%	1.13%	1.13%	1.12%	1.12%

FIGURE 8.10 Miss rates comparing least-recently used versus random replacement for several sizes and associativities. This data was collected for a block size of 16 bytes using one of the VAX traces containing user and operating system code (SAVE0). This trace is included in the software supplement for course use. There is little difference between LRU and random for larger size caches in this trace.

Q4: What Happens on a Write?

Reads dominate cache accesses. All instruction accesses are reads, and most instructions don't write to memory. Figure 4.34 (page 181) suggests a mix of 9% stores and 17% loads for four DLX programs, making writes less than 10% of the memory traffic. Making the common case fast means optimizing caches for reads, but Amdahl's Law reminds us that high-performance designs cannot neglect the speed of writes.

Fortunately, the common case is also the easy case to make fast. The block can be read at the same time that the tag is read and compared, so the block read begins as soon as the block-frame address is available. If the read is a hit, the block is passed on to the CPU immediately. If it is a miss, there is no benefit—but also no harm.

Such is not the case for writes. The processor specifies the size of the write, usually between 1 and 8 bytes; only that portion of a block can be changed. In general this means a read-modify-write sequence of operations on the block: read the original block, modify one portion, and write the new block value. Moreover, modifying a block cannot begin until the tag is checked to see if it is a hit. Because tag checking cannot occur in parallel, then, writes normally take longer than reads.

Thus, it is the write policies that distinguish many cache designs. There are two basic options when writing to the cache:

- *Write through* (or *store through*)—The information is written to both the block in the cache **and** to the block in the lower-level memory.
- *Write back* (also called *copy back* or *store in*)—The information is written only to the block in the cache. The modified cache block is written to main memory only when it is replaced.

Write-back cache blocks are called *clean* or *dirty*, depending on whether the information in the cache differs from that in lower-level memory. To reduce the frequency of writing back blocks on replacement, a feature called the *dirty bit* is commonly used. This status bit indicates whether or not the block was modified while in the cache. If it wasn't, the block is not written, since the lower level has the same information as the cache.

Both write back and write through have their advantages. With write back, writes occur at the speed of the cache memory, and multiple writes within a block require only one write to the lower-level memory. Since every write doesn't go to memory, write back uses less memory bandwidth, making write back attractive in multiprocessors. With write through, read misses don't result in writes to the lower level, and write through is easier to implement than write back. Write through also has the advantage that main memory has the most current copy of the data. This is important in multiprocessors and for I/O, which we shall examine in Section 8.8. Hence, multiprocessors want write back to reduce the memory traffic per processor and write through to keep the cache and memory consistent.

When the CPU must wait for writes to complete during write throughs, the CPU is said to *write stall*. A common optimization to reduce write stalls is a *write buffer*, which allows the processor to continue while the memory is updated. As we shall see in Section 8.8, write stalls can occur even with write buffers.

There are two options on a write miss:

- *Write allocate* (also called *fetch on write*)—The block is loaded, followed by the write-hit actions above. This is similar to a read miss.
- *No write allocate* (also called *write around*)—The block is modified in the lower level and not loaded into the cache.

While either write-miss policy could be used with write through or write back, generally write-back caches use write allocate (hoping that subsequent writes to that block will be captured by the cache) and write-through caches often use no write allocate (since subsequent writes to that block will still have to go to memory).

An Example Cache: The VAX-11/780 Cache

To give substance to these ideas, Figure 8.11 shows the organization of the cache on the VAX-11/780. The cache contains 8192 bytes of data in 8-byte blocks with two-way–set-associative placement, random replacement, write through with a one-word write buffer, and no write allocate on a write miss.

Let's trace a cache hit through the steps of a hit as labeled in Figure 8.11. (The five steps are shown as circled numbers.) The address coming into the cache is divided into two fields: the 29-bit block-frame address and 3-bit block offset. The block-frame address is further divided into an address tag and cache index. Step 1 shows this division.

The cache index selects the set to be tested to see if the block is in the cache. (A set is one block from each bank in Figure 8.11.) The size of the index depends on cache size, block size, and set associativity. In this case, a 9-bit index results:

$$\frac{\text{Blocks}}{\text{Bank}} = \frac{\text{Cache size}}{\text{Block size} * \text{Set associativity}} = \frac{8192}{8 * 2} = 512 = 2^9$$

In a two-way–set-associative cache, the index is sent to both banks. This is step 2.

After reading an address tag from each bank, the tag portion of the block-frame address is compared to the tags. This is step 3 in the figure. To be sure the tag contains valid information, the valid bit must be set, or the results of the comparison are ignored.

Assuming one of the tags does match, a 2:1 multiplexer (step 4) is set to select the block from the matching set. Why can't both tags match? It is the job of the replacement algorithm to make sure that an address appears in only one block. To reduce the hit time, the data is read at the same time as the address tags; thus, by the time the block multiplexer is ready, the data is also ready.

This step is needed in set-associative caches, but it can be omitted from direct-mapped caches since there is no selection to be made. The multiplexer used in this step can be on the critical timing path, endangering the clock cycle time of the CPU. (The example on pages 418–419 and the fallacy on page 481 explore the trade-off of lower miss rates and higher clock cycle time.)

In the final step the word is sent to the CPU. All five steps occur within a single CPU clock cycle.

What happens on a miss? The cache sends a stall signal to the CPU telling it to wait, and two words (eight bytes) are read from memory. That takes 6 clock cycles on the VAX-11/780 (ignoring bus interference). When the data arrives,

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 446

the cache must pick a block to replace; the VAX-11/780 selects one of the two blocks at random. Replacing a block means updating the data, the address tag, and the valid bit. Once this is done, the cache goes through a regular hit cycle and returns the data to the CPU.

Writes are more complicated in the VAX-11/780, as they are in any cache. If the word to be written is in the cache, the first four steps are the same. The next step is to write the data in the block, then write the changed-data portion into the

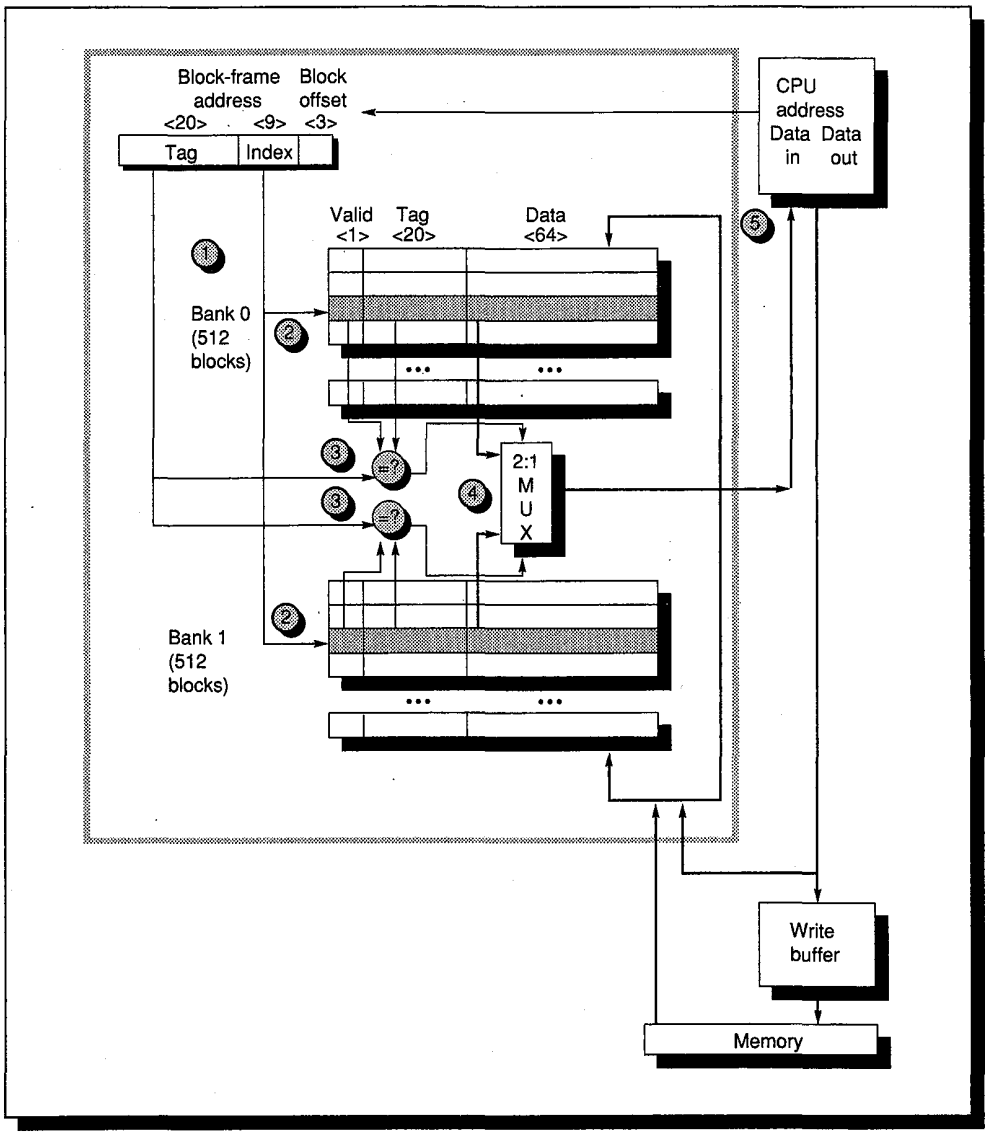

FIGURE 8.11 The organization of the VAX-11/780 cache. The 8-KB cache is two-way set associative with 8-byte blocks. It has 512 sets with two blocks per set; the set is selected by the 9-bit index. The five steps of a read hit, shown as circled numbers in order of occurrence, label this organization. The line from memory to the cache is used on a miss to load the cache. Multiplexing as found in step 4 is not needed in a direct-mapped cache. Note that the offset is connected to chip select of the data SRAMs to allow the proper words to be sent to the 2:1 multiplexer.

cache. The VAX-11/780 uses no write allocate. Consequently, on a write miss the CPU writes "around" the cache to lower-level memory and does not affect the cache.

Since this is a write-through cache, the process isn't yet over. The word is also sent to a one-word write buffer. If the write buffer is empty, the word and full address are written in the buffer, and we are finished. The CPU continues working while the write buffer writes the word to memory. If the buffer is full, the cache (and CPU) must wait until the buffer is empty.

Cache Performance

CPU time can be divided into the clock cycles the CPU spends executing the program and the clock cycles the CPU spends waiting for the memory system. Thus,

$$\text{CPU time} = (\text{CPU-execution clock cycles} + \text{Memory-stall clock cycles}) * \text{Clock cycle time}$$

To simplify evaluation of cache alternatives, sometimes designers assume that all memory stalls are due to the cache. This is true for many machines; on machines where this is not true, the cache still dominates stalls that are not exclusively due to the cache. We use this simplifying assumption here, but it is important to account for **all** memory stalls when calculating final performance!

The formula above raises the question whether the clock cycles for a cache access should be considered part of CPU-execution clock cycles or part of memory-stall clock cycles. While either convention is defensible, the most widely accepted is to include hit clock cycles in CPU-execution clock cycles.

Memory-stall clock cycles can then be defined in terms of the number of memory accesses per program, miss penalty (in clock cycles), and miss rate for reads and writes:

$$\text{Memory-stall clock cycles} = \frac{\text{Reads}}{\text{Program}} * \text{Read miss rate} * \text{Read miss penalty}$$

$$+ \frac{\text{Writes}}{\text{Program}} * \text{Write miss rate} * \text{Write miss penalty}$$

We simplify the complete formula by combining the reads and writes together:

$$\text{Memory-stall clock cycles} = \frac{\text{Memory accessess}}{\text{Program}} * \text{Miss rate} * \text{Miss penalty}$$

Factoring instruction count (IC) from execution time and memory stall cycles, we now get a CPU-time formula that includes memory accesses per instruction, miss rate, and miss penalty:

$$\text{CPU time} = \text{IC} * \left(\text{CPI}_{\text{Execution}} + \frac{\text{Memory accesses}}{\text{Instruction}} * \text{Miss rate} * \text{Miss penalty}\right) * \text{Clock cycle time}$$

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 448

Some designers prefer measuring miss rate as *misses per instruction* rather than misses per memory reference:

$$\frac{\text{Misses}}{\text{Instruction}} = \frac{\text{Memory accesses}}{\text{Instruction}} * \text{Miss rate}$$

The advantage of this measure is that it is independent of the hardware implementation. For example, the VAX-11/780 instruction unit can make repeated references to a single byte (see Section 8.7), which can artificially reduce the miss rate if measured as misses per memory reference rather than per instruction executed. The drawback is that this measure is architecture dependent, thus it is most popular with architects working with a single computer family. They then use this version of the CPU-time formula:

$$\text{CPU time} = \text{IC} * \left(\text{CPI}_{\text{Execution}} + \frac{\text{Misses}}{\text{Instruction}} * \text{Miss penalty}\right) * \text{Clock cycle time}$$

We can now explore the consequences of caches on performance.

Example

Let's use the VAX-11/780 as a first example. The cache miss penalty is 6 clock cycles, and all instructions normally take 8.5 clock cycles (ignoring memory stalls). Assume the miss rate is 11%, and there is an average of 3.0 memory references per instruction. What is the impact on performance when behavior of the cache is included?

Answer

$$\text{CPU time} = \text{IC} * \left(\text{CPI}_{\text{Execution}} + \frac{\text{Memory-stall clock cycles}}{\text{Instruction}}\right) * \text{Clock cycle time}$$

The performance, including cache misses, is

$$\text{CPU time}_{\text{with cache}} = \text{IC} * (8.5 + 3.0 * 11\% * 6) * \text{Clock cycle time}$$

$$= \text{Instruction count} * 10.5 * \text{Clock cycle time}$$

The clock cycle time and instruction count are the same, with or without a cache, so CPU time increases with CPI from 8.5 to 10.5. Hence, the impact of the memory hierarchy is to stretch the CPU time by 24%.

Example

Let's now calculate the impact on performance when behavior of the cache is included on a machine with a lower CPI. Assume that the cache miss penalty is 10 clock cycles and, on average, instructions take 1.5 clock cycles; the miss rate is 11%, and there is an average of 1.4 memory references per instruction.

Answer

$$\text{CPU time} = \text{IC} * \left(\text{CPI}_{\text{Execution}} + \frac{\text{Memory-stall clock cycles}}{\text{Instruction}}\right) * \text{Clock cycle time}$$

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 449

Making the same assumptions as in the previous example on cache hits, the performance, including cache misses, is

$$\text{CPU time}_{\text{with cache}} = \text{IC} * (1.5 + 1.4*11\%*10) * \text{Clock cycle time}$$

$$= \text{Instruction count}*3.0*\text{Clock cycle time}$$

The clock cycle time and instruction count are the same, with or without a cache, so CPU time increases with CPI from 1.5 to 3.0. Including cache behavior doubles execution time.

As these examples illustrate, cache-behavior penalties range from significant to enormous. Furthermore, cache misses have a double-barreled impact on a CPU with a low CPI and a fast clock:

1. The lower the CPI, the more pronounced the impact is.
2. Independent of the CPU, main memories have similar memory-access times, since they are built from the same memory chips. When calculating CPI, the cache miss penalty is measured in **CPU** clock cycles needed for a miss. Therefore, a higher CPU clock rate leads to a larger miss penalty, even if main memories are the same speed.

The importance of the cache for CPUs with low CPI and high clock rates is thus greater; and, consequently, greater is the danger of neglecting cache behavior in assessing performance of such machines.

While minimizing average memory-access time is a reasonable goal and we will use it in much of this chapter, keep in mind that the final goal is to reduce CPU execution time.

Example

What is the impact of two different cache organizations on the performance of a CPU? Assume that the CPI is normally 1.5 with a clock cycle time of 20 ns, that there are 1.3 memory references per instruction, and that the size of both caches is 64 KB. One cache is direct mapped and the other is two-way set associative. Since the speed of the CPU is tied directly to the speed of the caches, assume the CPU clock cycle time must be stretched 8.5% to accommodate the selection multiplexer of the set-associative cache (step 4 in Figure 8.11 on page 415.) To the first approximation, the cache miss penalty is 200 ns for either cache organization. (In practice it must be rounded up or down to an integer number of clock cycles.) First, calculate the average memory-access time, and then CPU performance.

Answer

Figure 8.12 on page 421 shows that the miss rate of a direct-mapped 64-KB cache is 3.9% and the miss rate for a two-way–set-associative cache of the same size is 3.0%. Average memory-access time is

$$\text{Average memory-access time} = \text{Hit time} + \text{Miss rate} * \text{Miss penalty}$$

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 450

Thus, the time for each organization is

$$\text{Average memory-access time}_{1\text{-way}} = 20 + .039*200 = 27.8 \text{ ns}$$

$$\text{Average memory-access time}_{2\text{-way}} = 20*1.085 + .030*200 = 27.7 \text{ ns}$$

The average memory-access time is better for the two-way–set-associative cache.

CPU performance is

$$\text{CPU time} = \text{IC} * \left(\text{CPI}_{\text{Execution}} + \frac{\text{Misses}}{\text{Instruction}} * \text{Miss penalty}\right) * \text{Clock cycle time}$$

$$= \text{IC} * \left(\text{CPI}_{\text{Execution}} * \text{Clock cycle time} + \frac{\text{Memory accesses}}{\text{Instruction}} * \text{Miss rate} * \text{Miss penalty} * \text{Clock cycle time}\right)$$

Substituting 200ns for (Miss penalty * Clock cycle time), the performance of each cache organization is

$$\text{CPU time}_{1\text{-way}} = \text{IC}*(1.5*20 + 1.3*0.039*200) = 40.1*\text{IC}$$

$$\text{CPU time}_{2\text{-way}} = \text{IC}*(1.5*20*1.085 + 1.3*0.030*200) = 40.4*\text{IC}$$

and relative performance is

$$\frac{\text{CPU time}_{2\text{-way}}}{\text{CPU time}_{1\text{-way}}} = \frac{40.4 * \text{Instruction count}}{40.1 * \text{Instruction count}}$$

In contrast to the results of average access-time comparison, the direct-mapped cache leads to slightly better performance. Since CPU time is our bottom-line evaluation (and direct mapped is simpler to build), the preferred cache is direct mapped in this example. (See the fallacy on page 481 for more on this kind of trade-off.)

The Three Sources of Cache Misses: Compulsory, Capacity, and Conflicts

An intuitive model of cache behavior attributes all misses to one of three sources:

- *Compulsory*—The first access to a block is not in the cache, so the block must be brought into the cache. These are also called *cold start misses* or *first reference misses.*
- *Capacity*—If the cache cannot contain all the blocks needed during execution of a program, capacity misses will occur due to blocks being discarded and later retrieved.

- *Conflict*—If the block-placement strategy is set associative or direct mapped, conflict misses (in addition to compulsory and capacity misses) will occur because a block can be discarded and later retrieved if too many blocks map to its set. These are also called *collision misses*.

Figure 8.12 shows the relative frequency of cache misses, broken down by the "three Cs." To show the benefit of associativity, conflict misses are divided into misses caused by each decrease in associativity. The categories are labeled *n*-way, meaning the misses caused by going to the lower level of associativity from the next one above. Here are the four categories:

8-way: from fully associative (no conflicts) to 8-way associative

4-way: from 8-way associative to 4-way associative

2-way: from 4-way associative to 2-way associative

1-way: from 2-way associative to 1-way associative (direct mapped)

Figure 8.13 (page 422) presents the same data graphically. The top graph shows absolute miss rates; the bottom graph plots percentage of all the misses by cache size.

Having identified the three Cs, what can a computer designer do about them? Conceptually, conflicts are the easiest: Fully associative placement avoids all conflict misses. Associativity is expensive in hardware, however, and may slow access time (see the example above or the second fallacy in Section 8.10), leading to lower overall performance. There is little to be done about capacity except to buy larger memory chips. If the upper-level memory is much smaller than what is needed for a program, and a significant percentage of the time is spent moving data between two levels in the hierarchy, the memory hierarchy is said to *thrash*. Because so many replacements are required, thrashing means the machine runs close to the speed of the lower-level memory, or maybe even slower due to the miss overhead. Making blocks larger reduces the number of compulsory misses, but it can increase conflict misses.

The three C's give insight into the cause of misses, but this simple model has its limits. For example, increasing cache size reduces conflict misses as well as capacity misses, since a larger cache spreads out references. Thus, a miss might move from one category to the other as parameters change. Three C's ignore replacement policy, since it is difficult to model and since, in general, it is of less significance. In specific circumstances the replacement policy can actually lead to anomalous behavior, such as poorer miss rates for larger associativity, which is directly contradictory to the three C's model.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 452

Cache size	Degree associative	Total miss rate	Miss-rate components (relative percent) (*Sum = 100% of total miss rate*)					
			Compulsory		Capacity		Conflict	
1 KB	1-way	0.191	0.009	5%	0.141	73%	0.042	22%
1 KB	2-way	0.161	0.009	6%	0.141	87%	0.012	7%
1 KB	4-way	0.152	0.009	6%	0.141	92%	0.003	2%
1 KB	8-way	0.149	0.009	6%	0.141	94%	0.000	0%
2 KB	1-way	0.148	0.009	6%	0.103	70%	0.036	24%
2 KB	2-way	0.122	0.009	7%	0.103	84%	0.010	8%
2 KB	4-way	0.115	0.009	8%	0.103	90%	0.003	2%
2 KB	8-way	0.113	0.009	8%	0.103	91%	0.001	1%
4 KB	1-way	0.109	0.009	8%	0.073	67%	0.027	25%
4 KB	2-way	0.095	0.009	9%	0.073	77%	0.013	14%
4 KB	4-way	0.087	0.009	10%	0.073	84%	0.005	6%
4 KB	8-way	0.084	0.009	11%	0.073	87%	0.002	3%
8 KB	1-way	0.087	0.009	10%	0.052	60%	0.026	30%
8 KB	2-way	0.069	0.009	13%	0.052	75%	0.008	12%
8 KB	4-way	0.065	0.009	14%	0.052	80%	0.004	6%
8 KB	8-way	0.063	0.009	14%	0.052	83%	0.002	3%
16 KB	1-way	0.066	0.009	14%	0.038	57%	0.019	29%
16 KB	2-way	0.054	0.009	17%	0.038	70%	0.007	13%
16 KB	4-way	0.049	0.009	18%	0.038	76%	0.003	6%
16 KB	8-way	0.048	0.009	19%	0.038	78%	0.001	3%
32 KB	1-way	0.050	0.009	18%	0.028	55%	0.013	27%
32 KB	2-way	0.041	0.009	22%	0.028	68%	0.004	11%
32 KB	4-way	0.038	0.009	23%	0.028	73%	0.001	4%
32 KB	8-way	0.038	0.009	24%	0.028	74%	0.001	2%
64 KB	1-way	0.039	0.009	23%	0.019	50%	0.011	27%
64 KB	2-way	0.030	0.009	30%	0.019	65%	0.002	5%
64 KB	4-way	0.028	0.009	32%	0.019	68%	0.000	0%
64 KB	8-way	0.028	0.009	32%	0.019	68%	0.000	0%
128 KB	1-way	0.026	0.009	34%	0.004	16%	0.013	50%
128 KB	2-way	0.020	0.009	46%	0.004	21%	0.006	33%
128 KB	4-way	0.016	0.009	55%	0.004	25%	0.003	20%
128 KB	8-way	0.015	0.009	59%	0.004	27%	0.002	14%

FIGURE 8.12 Total miss rate for each size cache and percentage of each according to the "three Cs." Compulsory misses are independent of cache size, while capacity misses decrease as capacity increases. Hill [1987] measured this trace using 32-byte blocks and LRU replacement. It was generated on a VAX-11 running Ultrix by mixing three systems' traces, using a multiprogramming workload and three user traces. The total length was just over a million addresses; the largest piece of data referenced during the trace was 221 KB. Figure 8.13 (page 422) shows the same information graphically. Note that the 2:1 cache rule of thumb (inside front cover) is supported by the statistics in this table: a direct-mapped cache of size N has about the same miss rate as a 2-way–set-associative cache of size N/2.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 453

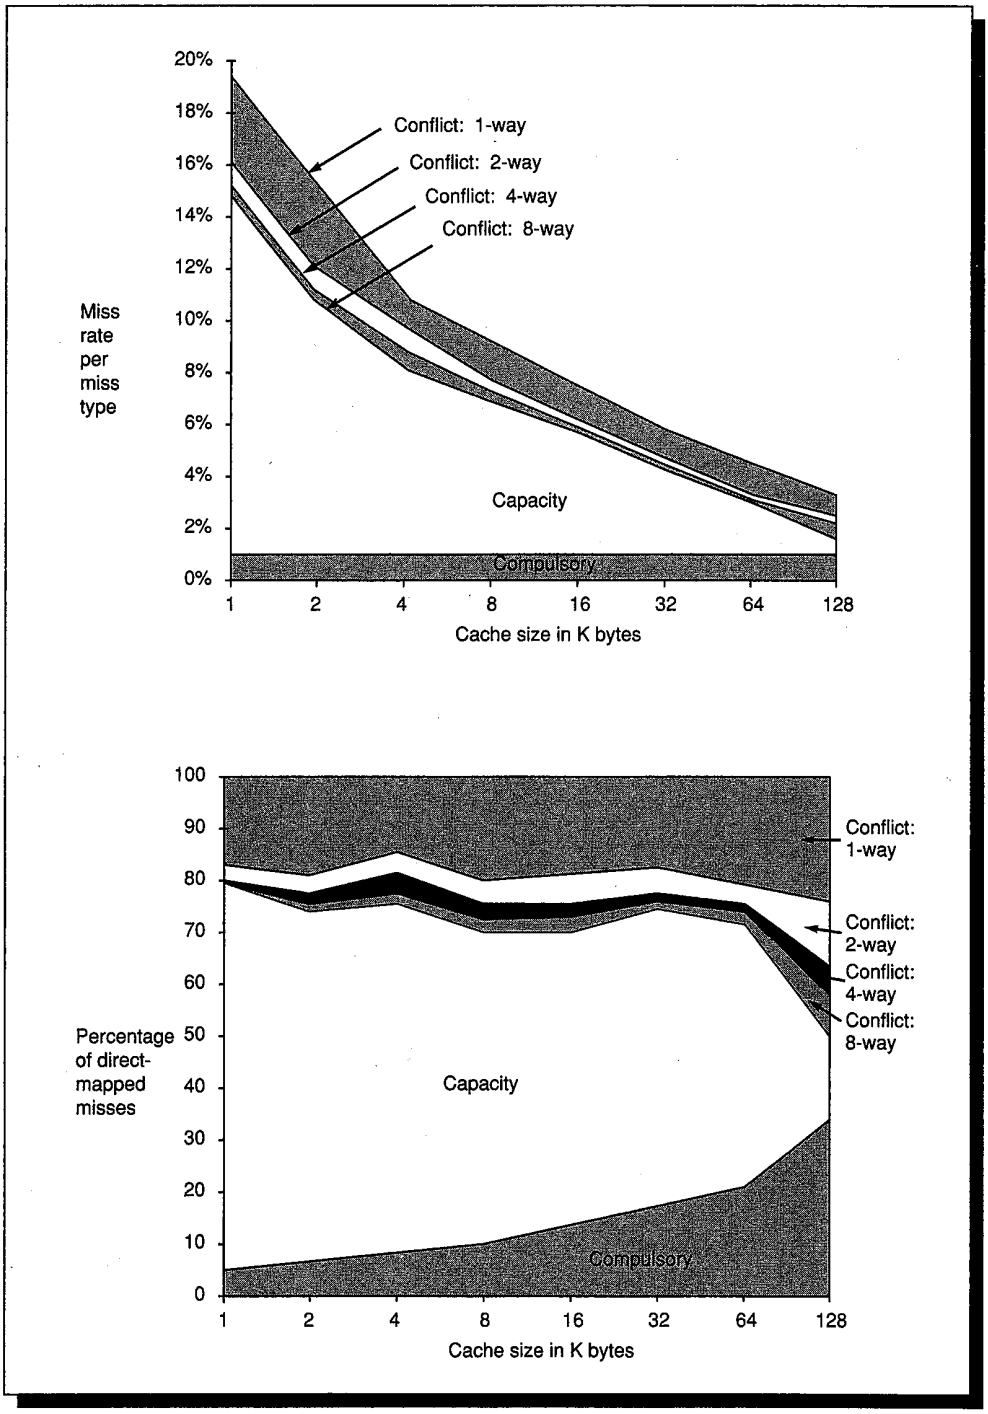

FIGURE 8.13 Total miss rate (top) and distribution of miss rate (bottom) for each size cache according to three Cs for the data in Figure 8.12 (page 421). The top diagram is the actual miss rates, while the bottom diagram is scaled to the direct-mapped miss ratio.

Choices for Block Sizes in Caches

Figures 8.3 and 8.4 (page 406) showed the abstract tradeoff of block size versus miss rate and memory-access time. Figures 8.14 and 8.15 (page 424) show the specific numbers for a set of programs and cache sizes. Larger block sizes reduce compulsory misses, as the principle of spatial locality suggests. At the same time, larger blocks also reduce the number of blocks in the cache, increasing conflict misses.

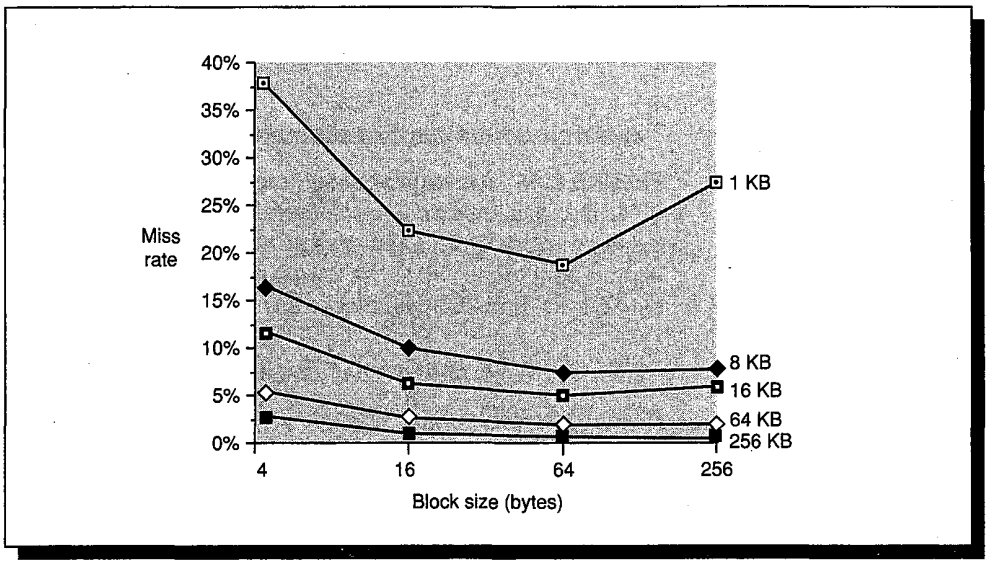

FIGURE 8.14 Miss rate versus block size. Note that for a 1-KB cache, 256-byte blocks have a higher miss rate than either 16- or 64-byte blocks. (The smallest block is 4 bytes.) In this particular example, the cache would have to be 256 KB in order for increasing block size to always result in decreased misses. This data was collected for a direct-mapped cache using one of the VAX traces containing user and operating system code, which is distributed with this book (SAVE0).

Instruction-Only or Data-Only Caches Versus Unified Caches

Unlike other levels of the memory hierarchy, caches are sometimes divided into instruction-only and data-only caches. Caches that can contain either instructions or data are *unified* caches, or *mixed* caches. The CPU knows whether it is issuing an instruction address or a data address, so there can be separate ports for both, thereby doubling the bandwidth between the cache and the CPU. (Section 6.4 in Chapter 6 shows the advantages of dual memory ports for pipelined execution.) Separate caches also offers the opportunity of optimizing each cache separately: different capacities, block sizes, and associativities may lead to better performance. Splitting thus affects the cost and performance far beyond what is indicated by the change in miss rates. We limit our discussion to that point now simply to show how miss rates for instructions differ from miss rates for data.

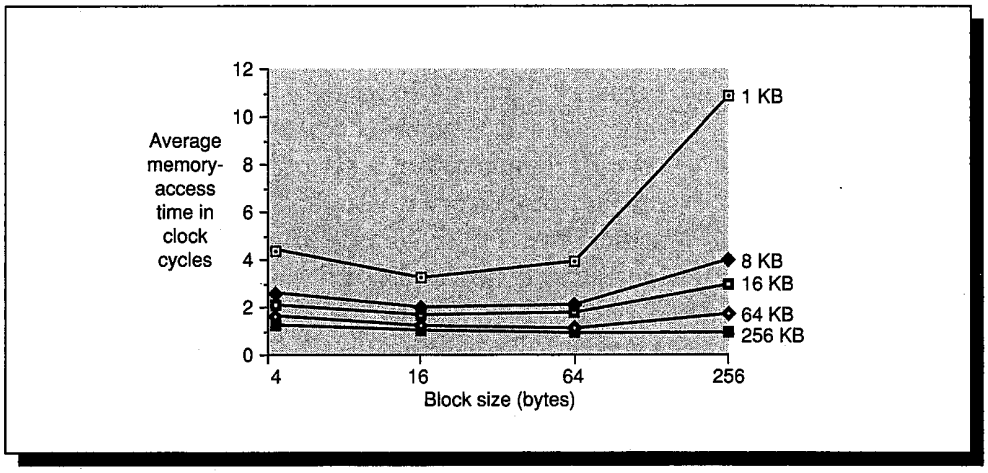

FIGURE 8.15 Average access time versus block size using the miss rates in Figure 8.14. This assumes an 8–clock-cycle latency and that the memory and bus can transfer 4 bytes per clock cycle. On a miss all the blocks are loaded into the cache before the requested word is sent to the CPU. The lowest average memory-access time is either for 16-byte or 64-byte blocks, and 256-byte blocks are better than 4-byte blocks only for the largest cache.

Figure 8.16 shows that instruction-only caches have lower miss rates than data-only caches. Separating instructions and data removes misses due to conflicts between instruction blocks and data blocks, but the split also fixes the cache space devoted to each type. A fair comparison of separate instruction and data caches to unified caches requires the total cache size to be the same. Therefore, a separate 1-KB instruction cache and 1-KB data cache should be compared to a unified 2-KB cache. Calculating the average miss rate with separate instruction-only and data-only caches necessitates knowing the percentage of memory references to each cache.

Size	Instruction only	Data only	Unified
0.25 KB	22.2%	26.8%	28.6%
0.50 KB	17.9%	20.9%	23.9%
1 KB	14.3%	16.0%	19.0%
2 KB	11.6%	11.8%	14.9%
4 KB	8.6%	8.7%	11.2%
8 KB	5.8%	6.8%	8.3%
16 KB	3.6%	5.3%	5.9%
32 KB	2.2%	4.0%	4.3%
64 KB	1.4%	2.8%	2.9%
128 KB	1.0%	2.1%	1.9%
256 KB	0.9%	1.9%	1.6%

FIGURE 8.16 Miss rates for instruction-only, data-only, and unified caches of different sizes. The data are for a 2-way–associative cache using LRU replacement with 16-byte blocks for an average of user/system traces on the VAX-11 and system traces on the IBM 370 [Hill 1987]. The percentage of instruction references in these traces is about 53%.

Example

Which has the lower miss rate: a 16-KB instruction cache with a 16-KB data cache or a 32-KB unified cache? Assume 53% of the references are instructions.

Answer

As stated in the legend of Figure 8.16, 53% of the memory accesses are instruction references. Thus, the overall miss rate for the split caches is

$$53\% * 3.6\% + 47\% * 5.3\% = 4.4\%$$

A 32-KB unified cache has a slightly lower miss rate of 4.3%.

8.4 Main Memory

... the one single development that put computers on their feet was the invention of a reliable form of memory, namely, the core memory, ... Its cost was reasonable, it was reliable and, because it was reliable, it could in due course be made large.

Maurice Wilkes, *Memoirs of a Computer Pioneer* (1985, p. 209)

Provided there is only one level of cache, main memory is the next level down in the hierarchy. Main memory satisfies the demands of caches and vector units, and serves as the I/O interface as it is the destination of input as well as the source for output. Unlike caches, performance measures of main memory emphasize both latency and bandwidth. Generally, main memory latency (which affects the cache miss penalty) is the primary concern of the cache, while main-memory bandwidth is the primary concern of I/O and vector units. As cache blocks grow from 4-8 bytes to 64–256 bytes, main memory bandwidth becomes important to caches as well. The relationship of main memory and I/O is discussed in Chapter 9.

Memory latency is traditionally quoted using two measures—access time and cycle time. *Access time* is the time between when a read is requested and when the desired word arrives, while *cycle time* is the minimum time between requests to memory. In the 1970s, as DRAMs grew in capacity the cost of a package with all the necessary address lines became an issue. The solution was to multiplex the address lines, thereby cutting the number of address pins in half. The top half of the address comes first, during the *row-access strobe*, or RAS. This is followed by the second half of the address during the *column-access strobe*, or CAS. These names come from the internal chip organization, for the memory is organized as a rectangular matrix addressed by rows and columns.

An additional requirement of DRAMs derives from the property signified by its first letter, D, for dynamic. Every DRAM must have every row accessed within a certain time window, such as 2 milliseconds, or the information in the DRAM can be lost. This requirement means that the memory system is

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 457

occasionally unavailable because it is sending a signal telling every chip to refresh. The cost of a refresh is typically a full memory access (RAS and CAS) for each row of the DRAM. Since the memory matrix in a DRAM is likely to be square, the number of steps in a refresh is usually the square root of the DRAM capacity.

In contrast to DRAMs are SRAMs—the first letter standing for "static." The dynamic nature of the circuits for DRAM require data to be written back after being read, hence the difference between the access time and the cycle time and also the need to refresh. SRAMs use more circuits per bit to prevent the information from being disturbed when read. Thus, unlike DRAMs, there is no difference between access time and cycle time and there is no need to refresh SRAM. In DRAM designs the emphasis is on capacity, while SRAM designs are concerned with both capacity **and** speed. (Because of this concern, SRAM address lines are not multiplexed.) For memories designed in comparable technologies, the capacity of DRAMs is roughly 16 times that of SRAMs, and the cycle time of SRAMs is 8 to 16 times faster than DRAMs.

The main memory of virtually every computer sold in the last decade is composed of semiconductor DRAMs (and virtually all caches use SRAM). Amdahl suggested a rule of thumb that memory capacity should grow linearly with CPU speed to keep a balanced system (see Section 1.4), and CPU designers rely on DRAMs to supply that demand: they expect a four-fold improvement in capacity every three years. Unfortunately, the performance of DRAMs is growing at a much slower rate. Figure 8.17 shows a performance improvement in row-access time of about 22% per generation, or 7% per year. As noted in Chapter 1, CPU performance improved 18% to 35% per year prior to 1985, and since that time has jumped to 50% to 100% per year. Figure 8.18 plots these optimistic and pessimistic CPU performance projections against the steady 7% performance improvement in DRAM speeds.

Year of introduction	Chip size	Row access (RAS) Slowest DRAM	Row access (RAS) Fastest DRAM	Column access (CAS)	Cycle time
1980	64 Kbit	180 ns	150 ns	75 ns	250 ns
1983	256 Kbit	150 ns	120 ns	50 ns	220 ns
1986	1 Mbit	120 ns	100 ns	25 ns	190 ns
1989	4 Mbit	100 ns	80 ns	20 ns	165 ns
1992?	16 Mbit	≈85 ns	≈65 ns	≈15 ns	≈140 ns

FIGURE 8.17 Times of fast and slow DRAMs with each generation. The improvement by a factor of two in column access accompanied the switch from NMOS DRAMs to CMOS DRAMs. With three years per generation, the performance improvement of row access time is about 7% per year. Data in the last row represent predicted performance for 16-Mbit DRAMs, which are not yet available.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 458

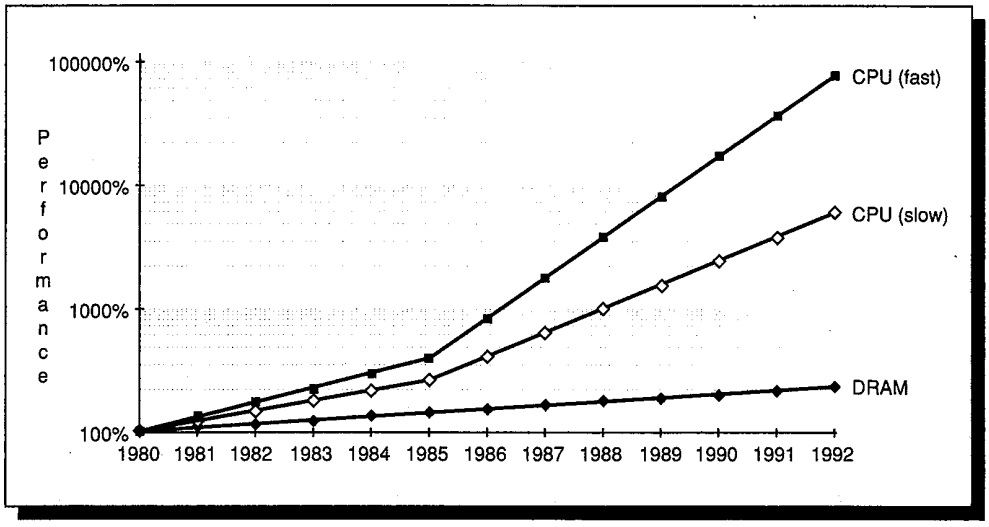

FIGURE 8.18 Starting with 1980 performance as a baseline, the performance of DRAMs and CPUs are plotted over time. The DRAM baseline is 64 KB in 1980, with three years to the next generation. The slow CPU line assumes a 19% improvement per year until 1985 and a 50% improvement thereafter. The fast CPU line assumes a 26% performance improvement between 1980 and 1985 and 100% per year thereafter. Note that the vertical axis must be on a logarithmic scale to record the size of the CPU–DRAM performance gap.

The CPU–DRAM performance gap is clearly a problem on the horizon—Amdahl's Law warns us what will happen if we ignore one portion of the computation while trying to speed up the rest. Section 8.8 will describe what can be done with cache organization to reduce this performance gap, but simply making caches larger cannot eliminate it. Innovative organizations of main memory are needed as well. In the rest of this section we will examine techniques for organizing memory to improve performance, including techniques especially for DRAMs.

Organizations for Improving Main Memory Performance

While it is generally easier to improve memory bandwidth with new organizations than it is to reduce latency, a bandwidth improvement does allow cache-block size to increase without a corresponding increase in the miss penalty.

Let's illustrate these organizations with the case of satisfying a cache miss. Assume the performance of the basic memory organization is

1 clock cycle to send the address

6 clock cycles for the access time per word

1 clock cycle to send a word of data

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 459

Given a cache block of four words, the miss penalty is 32 clock cycles, with a memory bandwidth of one-half byte per clock cycle.

Figure 8.19 shows some of the options to faster memory systems. The simplest approach to increasing memory bandwidth, then, is to make the memory wider.

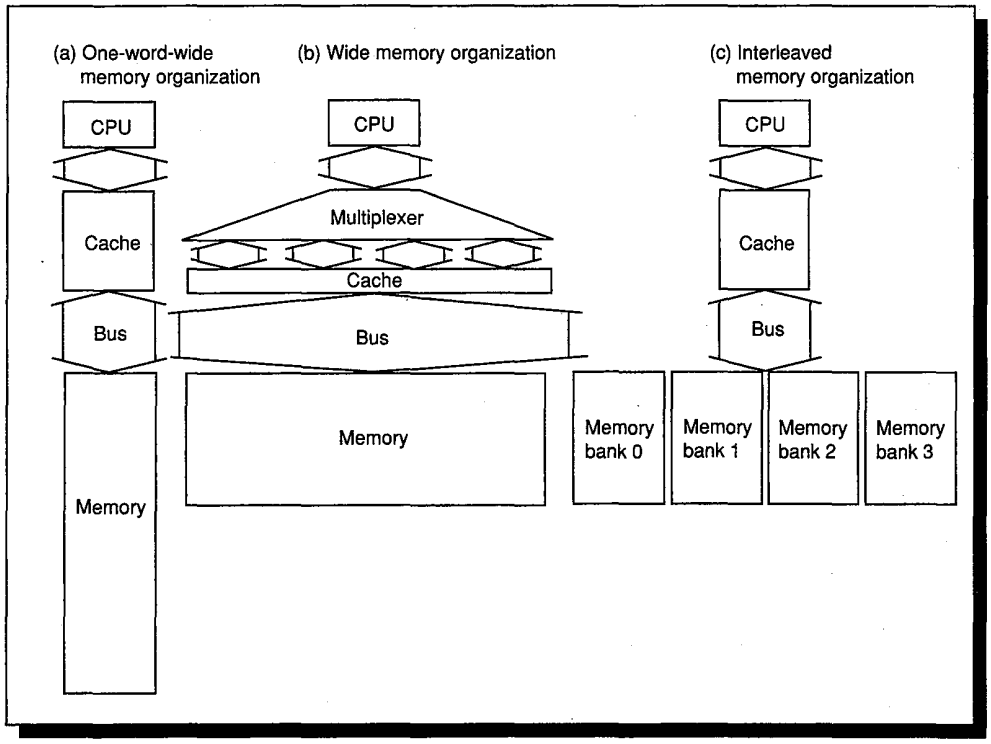

FIGURE 8.19 Three examples of bus width, memory width, and memory interleaving to achieve higher memory bandwidth. (a) is the simplest design, with everything the width of one word; (b) shows a wider memory, bus, and cache; while (c) shows a narrow bus and cache with an interleaved memory.

Wider Main Memory

Caches are often organized with a width of one word because most CPU accesses are that size. Main memory, in turn, is one word wide to match the width of the cache. Doubling or quadrupling the width of the memory will therefore double or quadruple the memory bandwidth. With a main memory width of two words the miss penalty in our example would drop from 4*8 or 32 clock cycles to 2*8 or 16 clock cycles. At four words wide the miss penalty is just 1*8 clock cycles. The bandwidth is then one byte per clock cycle at two words wide and two bytes per clock cycle when the memory is four words wide.

There is cost in the wider bus. The CPU will still access the cache a word at a time, so there now needs to be a multiplexer between the cache and the CPU—and that multiplexer may be on the critical timing path. (If the cache is faster

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 460

than the bus, however, the multiplexer can be placed between the cache and the bus.) Another drawback is that since main memory is traditionally expansible by the customer, the minimum increment is doubled or quadrupled. Finally, memories with error correction have difficulties with writes to a portion of the protected block (e.g., a write of a byte); the rest of the data must be read so that the new error correction code can be calculated and stored when the data is written. If the error correction is done over the full width, the wider memory will increase the frequency of such "read-modify-write" sequences because more writes become partial block writes. Many designs of wider memory have separate error correction every 32 bits since most writes are that size. One example of wider main memory was a computer whose cache, bus, and memory were all 512 bits wide.

Interleaved Memory

Memory chips can be organized in banks to read or write multiple words at a time rather than a single word. The banks are one word wide so that the width of the bus and the cache need not change, but sending addresses to several banks permits them all to read simultaneously. For example, sending an address to four banks (with access times shown on page 427) yields a miss penalty of 1+6+4*1 or 11 clock cycles, giving a bandwidth of about 1.5 bytes per clock cycle. Banks are also valuable on writes. While back-to-back writes would normally have to wait for earlier writes to finish, banks allow one clock cycle for each write, provided the writes are not destined to the same bank.

The mapping of addresses to banks affects the behavior of the memory system. The example above assumes the addresses of the four banks are interleaved at the word level—bank 0 has all words whose address modulo 4 is 0, bank 1 has all words whose address modulo 4 is 1, and so on. This mapping is referred to as the *interleaving factor*; *interleaved memory* normally means banks of memory that are word interleaved. This optimizes sequential memory accesses. A cache-read miss is an ideal match to word-interleaved memory, as the words in a block are read sequentially. Write-back caches make writes as well as reads sequential, getting even more efficiency from interleaved memory.

Example

What can interleaving and a wide memory buy? Consider the following description of a machine and its cache performance:

Block size = 1 word

Memory bus width = 1 word

Miss rate = 15%

Memory accesses per instruction = 1.2

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 461

Cache miss penalty = 8 cycles (as above)

Average cycles per instruction (ignoring cache misses) = 2

If we change the block size to two words, the miss rate falls to 10%, and a four-word block has a miss rate of 5%. What is the improvement in performance of interleaving two ways and four ways versus doubling the width of memory and the bus, assuming the access times on page 427.

Answer

The CPI for the base machine using one-word blocks is

$$2 + (1.2*15\%*8) = 3.44$$

Since the clock cycle time and instruction count won't change in this example, we can calculate performance improvement by just comparing CPI.

Increasing the block size to two words gives the following options:

32-bit bus and memory, no interleaving = 2 + (1.2*10%*2*8) = 3.92

32-bit bus and memory, interleaving = 2 + (1.2*10%*(1+6+2)) = 3.08

64-bit bus and memory, no interleaving = 2 + (1.2*10%*1*8) = 2.96

Thus, doubling the block size slows down the straightforward implementation (3.92 versus 3.44), while interleaving or wider memory is 12% or 16% faster, respectively. If we increase the block size to four, the following is obtained:

32-bit bus and memory, no interleaving = 2 + (1.2*5%*4*8) = 3.92

32-bit bus and memory, interleaving = 2 + (1.2*5%*(1+6+4)) = 2.66

64-bit bus and memory, no interleaving = 2 + (1.2*5%*2*8) = 2.96

Again, the larger block hurts performance for the simple case, although the interleaved 32-bit memory is now fastest—29% versus 16% for the wider memory and bus.

The original motivation for memory banks was interleaving sequential accesses. A further reason is to allow multiple independent accesses. Multiple memory controllers allow banks (or sets of word-interleaved banks) to operate independently. For example, an input device may use one controller and its memory, the cache may use another, and a vector unit may use a third. To reduce the chances of conflicts many banks are needed; the NEC SX/3, for instance, has up to 128 banks.

As capacity per memory chip increases, there are fewer chips in the same-sized memory system, making multiple banks much more expensive. For example, a 16-MB main memory takes 512 memory chips of 256 K (262,144) x 1 bits, easily organized into 16 banks of 32 memory chips. But it takes only 32 4-M (4,194,304) x 1-bit memory chips for 16 MB, making one bank the limit. This is the main disadvantage of interleaved memory banks. Even though the

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 462

Amdahl/Case rule of thumb for balanced computer systems recommends increasing memory capacity with increasing CPU performance, the 60% growth in DRAM capacity exceeded the rate of increase in CPU performance in the past (page 17 of Chapter 1). If the rate of increase of CPU speeds seen in the late 1980s can be maintained (Figure 8.18, page 427) and these systems follow the Amdahl/Case rule of thumb, then the number of chips may not be reduced.

A second disadvantage of interleaving is again the difficulty of main memory expansion. Since memory-control hardware will likely need equal-sized banks, doubling the main memory will probably be the minimum increment.

DRAM-Specific Interleaving for Improving Main Memory Performance

DRAM access times are divided into row access and column access. DRAMs buffer a row of bits inside the DRAM for the column access. This row is usually the square root of the DRAM size—1024 bits for 1 Mbit, 2048 for 4 Mbits, and so on. All DRAMs come with optional timing signals that allow repeated accesses to the buffer without a row-access time. There are three versions for this optimization:

- *Nibble mode*—The DRAM can supply three extra bits from sequential locations for every row access.
- *Page mode*—The buffer acts like a SRAM; by changing column address, random bits can be accessed in the buffer until the next row access or refresh time.
- *Static column*—Very similar to page mode, except that it's not necessary to hit the column-access strobe line every time the column address changes; this option has been nicknamed SCRAM, for static column DRAM.

Starting with the 1-Mbit DRAMs, most dies can perform any of the three options, with the optimization selected at the time the die is packaged by choosing which pads to wire up. These operations change the definition of cycle time for DRAMs. Figure 8.20 (page 432) shows the traditional cycle time plus the fastest speed between accesses in the optimized mode.

The advantage of these optimizations is that they use the circuitry already on the DRAMs, adding little cost to the system while achieving almost a fourfold improvement in bandwidth. For example, nibble mode was designed to take advantage of the same program behavior as interleaved memory. The chip reads four bits at a time internally, supplying four bits externally in the time of four optimized cycles. Unless the bus transfer time is faster than the optimized cycle time, the cost of four-way interleaved memory is only more complicated timing control. Page mode and static column could also be used to get even higher interleaving with slightly more complex control. DRAMs also tend to have weak tristate buffers, implying traditional interleaving with more memory chips must include buffer chips for each memory bank.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 463

Chip size	Row access Slowest DRAM	Row access Fastest DRAM	Column access	Cycle time	Optimized time nibble, page, static column
64 Kbits	180 ns	150 ns	75 ns	250 ns	150 ns
256 Kbits	150 ns	120 ns	50 ns	220 ns	100 ns
1 Mbits	120 ns	100 ns	25 ns	190 ns	50 ns
4 Mbits	100 ns	80 ns	20 ns	165 ns	40 ns
16 Mbits	≈85 ns	≈65 ns	≈15 ns	≈140 ns	≈30 ns

FIGURE 8.20 DRAM cycle time for the optimized accesses. This is Figure 8.17 (page 426) with a column added to show the optimized cycle time for the three modes. Starting with the 1-Mbit DRAM, optimized cycle time is about four times faster than unoptimized cycle time. It is so much faster that page mode was renamed *fast page mode*. The optimized cycle time is the same no matter which of the 3 optimized modes is selected.

Thus, the authors expect that most main memory systems in the future will use such techniques to reduce the CPU–DRAM performance gap. Unlike traditional interleaved memories, there are no disadvantages using these DRAM modes as DRAMs scale upward in capacity, nor is there the problem of the minimum expansion increment in main memory.

One possibility that recently arrived is DRAMs that do not multiplex the address lines. At the cost of a larger package, a full random access falls between a row-access time and a column-access time in Figure 8.20. If unencoded DRAMs can stay close to the price per bit of the high volume encoded DRAMs, the computer architect will have another option in his bag of tricks for memory design.

8.5 Virtual Memory

... a system has been devised to make the core drum combination appear to the programmer as a single level store, the requisite transfers taking place automatically.

Kilburn et al. [1962]

At any instant in time computers are running multiple processes, each with its own address space. (Processes are described in the next section.) It would be too expensive to dedicate a full-address-space worth of memory for each process, especially since many processes use only a small part of their address space. Hence, there must be a means of sharing a smaller amount of physical memory between many processes. One way to do this, *virtual memory*, divides physical memory into blocks and allocates them to different processes. Inherent in such an approach must be a *protection* scheme that restricts a process to the blocks

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 464

belonging just to that process. Most forms of virtual memory also reduce the time to start a program, since not all code and data need be in physical memory before a program can begin.

While virtual memory is essential for current computers, sharing is not the reason virtual memory was invented. In former days if a program became too large for physical memory, it was up to the programmer to make it fit. Programmers divided programs into pieces and then identified the pieces that were mutually exclusive. These *overlays* were loaded or unloaded under user program control during execution, with the programmer ensuring that the program never tried to access more physical main memory in the machine. As one can well imagine, this responsibility eroded programmer productivity. Virtual memory, invented to relieve programmers of this burden, automatically managed the two levels of the memory hierarchy represented by main memory and secondary storage.

In addition to sharing protected memory space and automatically managing the memory hierarchy, virtual memory also simplifies loading the program for execution. Called *relocation*, this procedure allows the same program to run in any location in physical memory. (Prior to the popularity of virtual memory, machines would include a relocation register just for that purpose.) An alternative to a hardware solution would be software that changed all addresses in a program each time it was run.

Several general memory-hierarchy terms from Section 8.3 apply to virtual memory, while some other terms are different. *Page* or *segment* is used for block, and *page fault*, or *address fault*, is used for miss. With virtual memory, the CPU produces *virtual addresses* that are translated by a combination of hardware and software to *physical addresses*, which can be used to access main memory. This process is called *memory mapping* or *address translation*. Today, the two memory hierarchy levels controlled by virtual memory are DRAMs and magnetic disks. Figure 8.21 shows a typical range of memory hierarchy parameters for virtual memory.

Block (page) size	512 – 8192 bytes
Hit time	1–10 clock cycles
Miss penalty	100,000 – 600,000 clock cycles
(Access time)	(100,000–500,000 clock cycles)
(Transfer time)	(10,000–100,000 clock cycles)
Miss rate	0.00001%–0.001%
Main memory size	4 MB – 2048 MB

FIGURE 8.21 Typical ranges of parameters for virtual memory. These figures, contrasted with the values for caches in Figure 8.5 (page 408), represent increases of 10 to 100,000 times.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 465

There are further differences between caches and virtual memory beyond those quantitative ones seen by comparing Figure 8.21 (page 433) to Figure 8.5 (page 408):

- Replacement on cache misses is primarily controlled by hardware, while virtual memory replacement is primarily controlled by the operating system; the longer miss penalty means the operating system can afford to get involved and spend more time deciding what to replace.
- The size of the processor address determines the size of virtual memory, but the cache size is normally independent of the processor address.
- In addition to acting as the lower-level memory for main memory in the hierarchy, secondary storage is also used for the file system that is not normally part of the address space; most of secondary storage is in fact taken up by the file system.

Virtual memory encompasses several related techniques. Virtual memory systems can be categorized into two classes: those with fixed-size blocks, called *pages*, and those with variable size blocks, called *segments*. Pages are typically fixed at 512 to 8192 bytes, while segment size varies. The largest segment supported on any machine ranges from 2^{16} bytes up to 2^{32} bytes; the smallest segment is one byte.

The decision to use paged virtual memory versus segmented virtual memory affects the CPU. Paged addressing has a single, fixed-size address divided into page number and offset within a page, analogous to cache addressing. A single address does not work for segmented addresses; the variable size of segments requires one word for a segment number and one word for an offset within a segment, for a total of two words. An unsegmented address space is simpler for the compiler.

The pros and cons of these two approaches have been well documented in operating systems textbooks; these are summarized in Figure 8.22. Because of the replacement problem (the third line of the figure), few machines today use pure segmentation. Some machines use a hybrid approach, called *paged segments*, in which a segment is an integral number of pages. This simplifies replacement because memory need not be contiguous, and the full segments need not be in main memory.

We are now ready to answer the four memory-hierarchy questions for virtual memory.

Q1: Where Can a Block Be Placed in Main Memory?

The miss penalty for virtual memory involves access to a rotating magnetic storage device and is therefore quite high. Given the choice of lower miss rates or a simpler placement algorithm, operating systems designers always pick lower miss rates because of the horrendous cost of a miss. Thus, operating systems allow blocks to be placed anywhere in main memory. According to the

terminology in Figure 8.6 (page 409), this strategy would be labeled fully associative.

Q2: How Is a Block Found If It Is in Main Memory?

Both paging and segmentation rely on a data structure that is indexed by the page or segment number. This data structure contains the physical address of the block. For paging, the offset is simply concatenated to this physical page address (see Figure 8.23, page 436). For segmentation, the offset is added to the segment's physical address to obtain the final virtual address.

	Page	Segment
Words per address	One	Two (segment and offset)
Programmer visible?	Invisible to application programmer	May be visible to application programmer
Replacing a block	Trivial (all blocks are the same size)	Hard (must find contiguous, variable-size, unused portion of main memory)
Memory use inefficiency	Internal fragmentation (unused portion of page)	*External fragmentation* (unused pieces of main memory)
Efficient disk traffic	Yes (adjust page size to balance access time and transfer time)	Not always (small segments may transfer just a few bytes)

FIGURE 8.22 Paging versus segmentation. Both can waste memory, depending on the block size and how well the segments fit together in main memory. Programming languages with unrestricted pointers require both the segment and the address to be passed. A hybrid approach, called *paged segments*, shoots for the best of both worlds: segments are composed of pages, so replacing a block is easy, yet a segment may be treated as a logical unit.

This data structure containing the physical page addresses usually takes the form of a *page table*. Indexed by the virtual page number, the size of the table is the number of pages in the virtual-address space. Given a 28-bit virtual address, 4 KB pages, and 4 bytes per page-table entry, the size of the page table would be 256 KB. To reduce the size of this data structure, some machines apply a hashing function to the virtual address so that the data structure need only be the size of the number of **physical** pages in main memory; this number would be much smaller than the number of virtual pages. Such a structure is called an *inverted page table*. Using the example above, a 64-MB physical memory would only need 64 KB (4*64 MB/4 KB) for an inverted page table.

To reduce address translation time, computers use a cache dedicated to these address translations, called a translation-lookaside buffer, or simply translation buffer. They are described in more detail shortly.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 467

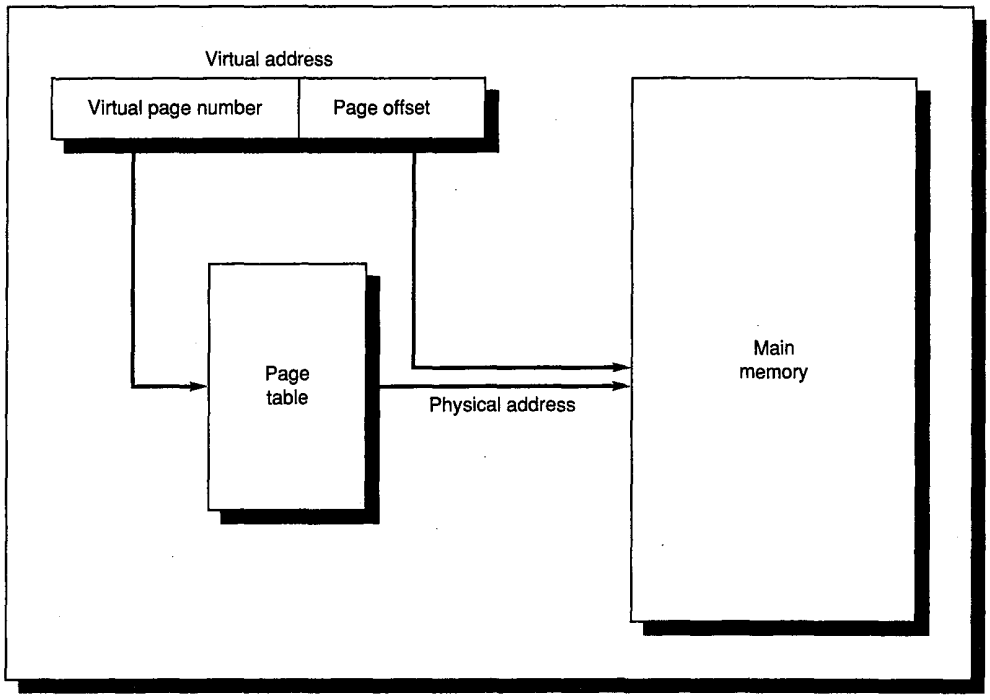

FIGURE 8.23 The mapping of a virtual address to a physical address via a page table.

Q3: Which Block Should Be Replaced on a Virtual Memory Miss?

As mentioned above, the overriding operating system guideline is minimizing page faults. Consistent with this, almost all operating systems try to replace the least-recently used (LRU) block, because that is the one least likely to be needed. To help the operating system estimate LRU, many machines provide a *use bit* or *reference bit*, which is set whenever a page is accessed. The operating system periodically clears the use bits and later records them so it can determine which pages were touched during a particular time period. By keeping track in this way, the operating system can select a page that is among the least-recently referenced.

Q4: What Happens on a Write?

The level below main memory contains rotating magnetic disks that take hundreds of thousands of clock cycles to access. Because of the great discrepancy in access time, no one has yet built a virtual memory operating system that can write through main memory straight to disk on every store by the CPU. (This remark should not be interpreted as an opportunity to become famous by being the first to build one!) Thus, the write strategy is always write back. Since the cost of an unnecessary access to the next-lower level is so high, virtual memory systems include a dirty bit so that the only blocks written to disk are those that have been altered since they were loaded from the disk.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 468

Selecting a Page Size

The most obvious architectural parameter is the page size. Choosing the page is a question of balancing forces that favor a larger page size versus those favoring a smaller size. The following favor a larger size:

- The size of the page table is inversely proportional to the page size; memory (or other resources used for the memory map) can therefore be saved by making the pages bigger.
- Transferring larger pages to or from secondary storage, possibly over a network, is more efficient than transferring smaller pages.

(The larger page size may also help in address translation of cache addresses; see Section 8.8.)

The main motivation for a smaller page size is conserving storage. A small page size will result in less wasted storage when a contiguous region of virtual memory is not equal in size to a multiple of the page size. The term for this unused memory in a page is *internal fragmentation.* Assuming that each process has three primary segments (text, heap, and stack), the average wasted storage per process will be 1.5 times the page size. This is negligible for machines with megabytes of memory and page sizes in the range of 2 KB to 8 KB. Of course, when the page sizes become very large (more than 32 KB), lots of storage (both main and secondary) may be wasted, as well as I/O bandwidth. A final concern is process start-up time; many processes are small, so larger page sizes would lengthen the time to invoke a process.

Techniques for Fast Address Translation

Page tables are usually so large that they are stored in main memory and often paged themselves. This means that every memory access takes at least twice as long, with one memory access to obtain the physical address and a second access to get the data. This cost is far too dear.

One remedy is to remember the last translation, so that the mapping process is skipped if the current address refers to the same page as the last one. A more general solution is to again rely on the principle of locality; if the references have locality, then the address translations for the references must also have locality. By keeping these address translations in a special cache, a memory access rarely requires a second access to translate the data. This special address translation cache is referred to as a *translation-lookaside buffer* or TLB, also called a "translation buffer," or TB. A TLB entry is like a cache entry where the tag holds portions of the virtual address and the data portion holds a physical page-frame number, protection field, use bit, and dirty bit. To change the physical page-frame number or protection of an entry in the page table the operating system must make sure the old entry is not in the TLB; otherwise, the

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 469

system won't behave properly. Note that this dirty bit means the corresponding **page** is dirty, not that the address translation in the TLB is dirty nor that a particular block in the data cache is dirty. Figure 8.24 shows typical parameters for TLBs.

Block size	4 – 8 bytes (1 page-table entry)
Hit time	1 clock cycle
Miss penalty	10 – 30 clock cycles
Miss rate	0.1% – 2%
TLB size	32 – 8192 bytes

FIGURE 8.24 Typical values of key memory-hierarchy parameters for TLBs. TLBs are simply caches for the virtual-to-physical address translations found in the page tables.

One architectural challenge stems from the difficulty of combining caches with virtual memory. The virtual address must first go through the TLB **before** the physical address can access the cache, meaning that the cache hit time must be stretched to allow for address translation (or the pipeline could be stretched as in Chapter 6). One way to reduce hit time is to access the cache with the page offset, the portion of the virtual address that does not need to be translated. While the cache address tags are being read, the virtual portion of the address (the page-frame address) is sent to the TLB to be translated. The address comparison is then between the physical address from the TLB and the cache tag. Since the TLB is usually smaller and faster than the cache-address-tag memory, simultaneous TLB reading need not slow down cache hit times. The drawback with this scheme is that a direct-mapped cache can be no bigger than a page. Another option, virtually addressed caches, is discussed in Section 8.8.

8.6 Protection and Examples of Virtual Memory

The invention of multiprogramming led to new demands for protection and sharing between programs. These are closely tied to virtual memory in computers today, and so we cover the topic here along with two examples of virtual memory.

Multiprogramming lead to the concept of a *process*. Metaphorically, a process is a program's breathing air and living space; that is, a running program plus any state needed to continue running the program. Timesharing means sharing the CPU and memory with several users at the same time to give the appearance that every user has his own machine. Thus, at any instant it must be possible to switch from one process to another. This is called a *process switch* or *context switch*. Figure 8.25 shows the frequency of these switches on the VAX 8700.

Instructions between process switches	19,353
Clock cycles between process switches	170,113
Time between process switches	7.7 ms

FIGURE 8.25 Frequency of process switches on VAX 8700 for timesharing workload. Most switching occurs on interrupts caused by I/O events or by the interval timer (see Figure 5.10, page 216). Since neither the latency of the I/O device nor the timer is affected by the speed of the CPU clock, faster machines generally execute more clock cycles and instructions between process switches.

A process must operate correctly whether it executes continuously from start to finish or is interrupted repeatedly and switched with other processes. The responsibility for maintaining correct process behavior is shared by the computer designer, who must ensure that the CPU portion of the process state can be saved and restored, and the operating system designer, who must guarantee that processes do not interfere with each others' computations. The safest way to protect the state of one process from another would be to copy the current information to disk. But a process switch would then take seconds—far too long for a timesharing environment. The problem is solved by operating systems partitioning main memory so that several different processes have their state in memory at the same time. This means that the operating system designer needs help from the computer designer to provide protection so that one process cannot modify another. Besides protection, the computers also provide for sharing of code and data between processes, to allow communication between processes or to save memory by reducing the number of copies of identical information.

Protecting Processes

The simplest protection mechanism is a pair of registers that checks every address to be sure that it falls between the two limits traditionally called *base* and *bound*. An address is valid if

$$\text{Base} \leq \text{Address} \leq \text{Bound}$$

In some systems the address is considered an unsigned number that is always added to the base, so the valid test is just

$$(\text{Base} + \text{Address}) \leq \text{Bound}$$

For user processes to be protected from each other, they can't change the base and bounds registers, yet the operating system must be able to change the registers so that it can switch processes. Hence, the computer designer has three more responsibilities in helping the operating system designer protect processes from each other:

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 471

1. Provide at least two modes indicating whether the running process is a user process or an operating system process, sometimes called a *kernel* process, a *supervisor* process or an *executive* process.

2. Provide a portion of the CPU state that a user process can use but not write. This includes the base/bound registers, a user/supervisor mode bit(s), and the interrupt enable/disable bit. Users are prevented from writing this state because the operating system cannot control user processes if users can change the address-range checks, disable interrupts, or give themselves supervisor privileges.

3. Provide mechanisms whereby the CPU can go from user mode to supervisor mode and vice versa. The first direction is typically accomplished by a *system call*, implemented as a special instruction that transfers control to a dedicated location in supervisor code space. The PC from the point of the system call is saved, and the CPU is placed in supervisor mode. The return to user mode is like a subroutine return that restores the previous user/supervisor mode.

Base and bound constitute the minimum protection system. Virtual memory provides an alternative to this simple model. As we have seen, the CPU address must go through a mapping from virtual to physical address. This provides the opportunity for the hardware to check further for errors in the program or to protect processes from each other. The simplest way of doing this is to add access permission flags to each page or segment. For example, since few programs today intentionally modify their own code, an operating system can detect accidental writes to code by offering read-only protection to pages. This can be extended by adding a user/kernel bit to prevent a user program from trying to access pages that belong to the kernel. As long as the CPU provides a read/write signal and a user/kernel signal, it is easy for the address translation hardware to detect stray memory accesses before they can do damage. As seen in Section 5.6 of Chapter 5, such reckless behavior interrupts the CPU. Obviously, user programs cannot be allowed to modify the page table.

Protection can be escalated, depending on the apprehension of the computer designer or the purchaser. Rings added to the CPU-protection structure expand memory-access protection from two levels (user and kernel) to many more. Like a military classification system of top secret, secret, classified, and unclassified, concentric *rings* of security levels allow the most trusted to access anything, the second most trusted to access everything except the innermost level, and so on down to "civilian" programs which are the least trusted and, hence, have the most limited range of accesses. There may also be restrictions on the entrance point between the levels. The 80286 protection structure, which uses rings, is described later in this section. It is not clear today whether rings are an improvement on the simple system of user and kernel modes.

As the designer's apprehension escalates to trepidation, these simple rings may not suffice. The fact that a program in the inner sanctum can access anything calls for a new classification system. Instead of a military model, the

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 472

analogy of this next model is to keys and locks: A program can't unlock access to the data unless it has the key. For these keys, or *capabilities*, to be useful, the hardware and operating system must be able to explicitly pass them from one program to another without allowing a program itself to forge them. Such checking requires a great deal of hardware support.

A Paged Virtual Memory Example: VAX-11 Memory Management and the VAX-11/780 TLB

The VAX architecture uses a combination of segmentation and paging. This combination provides protection while minimizing page-table size. The address space is first divided into two segments: process (bit 31 = 0) and system (bit 31=1). Every process has its own private space and shares system space with every other process. The process address space is further subdivided into two regions called P0 and P1, using bit 30 to distinguish them. Area P0 (bit 30 = 0) grows from address 0 upward while P1 (bit 30 = 1) grows downward to 0. Figure 8.26 shows the layout of P0 and P1. The two segments can grow until one exceeds its 2^{30} address-space size and its virtual memory is exhausted. Many systems today use some such combination of predivided segments and paging. The approach provides many advantages: Segmentation divides system and process address space and conserves page-table space, while paging provides virtual memory, relocation, and protection.

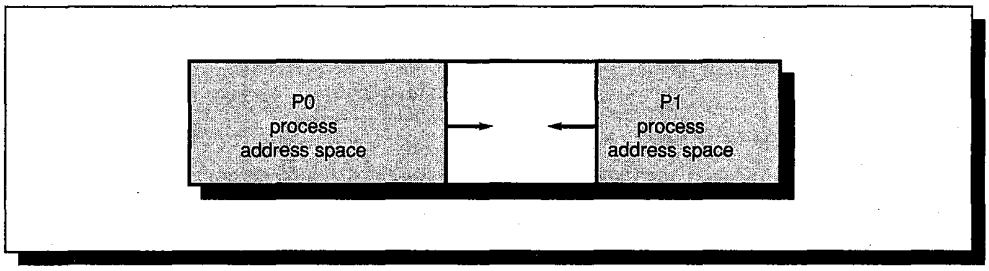

FIGURE 8.26 The organization of P0 and P1 in the VAX. This is the process half of the address space, selected with a 0 in bit 31 of a virtual address. Bit 30 of the address divides P0 and P1. Operating systems put the text and heap areas into P0 and a downward growing stack into P1.

To conserve page-table space, each of the three regions—P0 process, P1 process, and system—is provided with a pair of base-bound registers that indicate the start and limit of the page table for each region. The alternative would be to have a single page table that covers the full address space, independent of the program's actual size. The small size of the VAX pages—512 bytes, yielding large page tables—makes such conservation especially important.

Figure 8.27 (page 442) shows the mapping of a VAX address. The two most-significant bits of an address select which segment or base-bound–register pair

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 473

to use in selecting a page table and checking the reference. A one in the first bit selects the system page table, whose base and length are found respectively in the system base register and in the system length register. A zero in the first bit of an address (as in the figure) selects page table P0 or P1, found by the P0 or P1 base registers and checked by the P0 or P1 limit (bound) registers. The P0 and P1 page tables are in the system-space virtual memory, while the system page table is in physical memory.

This offers an interesting way to conserve physical memory. Since the P0 and P1 page tables are also in virtual memory, this means the page tables can be paged. Just as some code and data can remain on disk during program execution, the page-table translation entries for that code and data can remain on disk until they are used. This is especially important for programs whose memory size varies dynamically during execution, as page tables can be increased as P0 or P1 space grows. In the worst case, then, a process page fault can result in a second page fault bringing in the missing piece of the process page table needed to complete the address translation. What prevents all pages tables from being

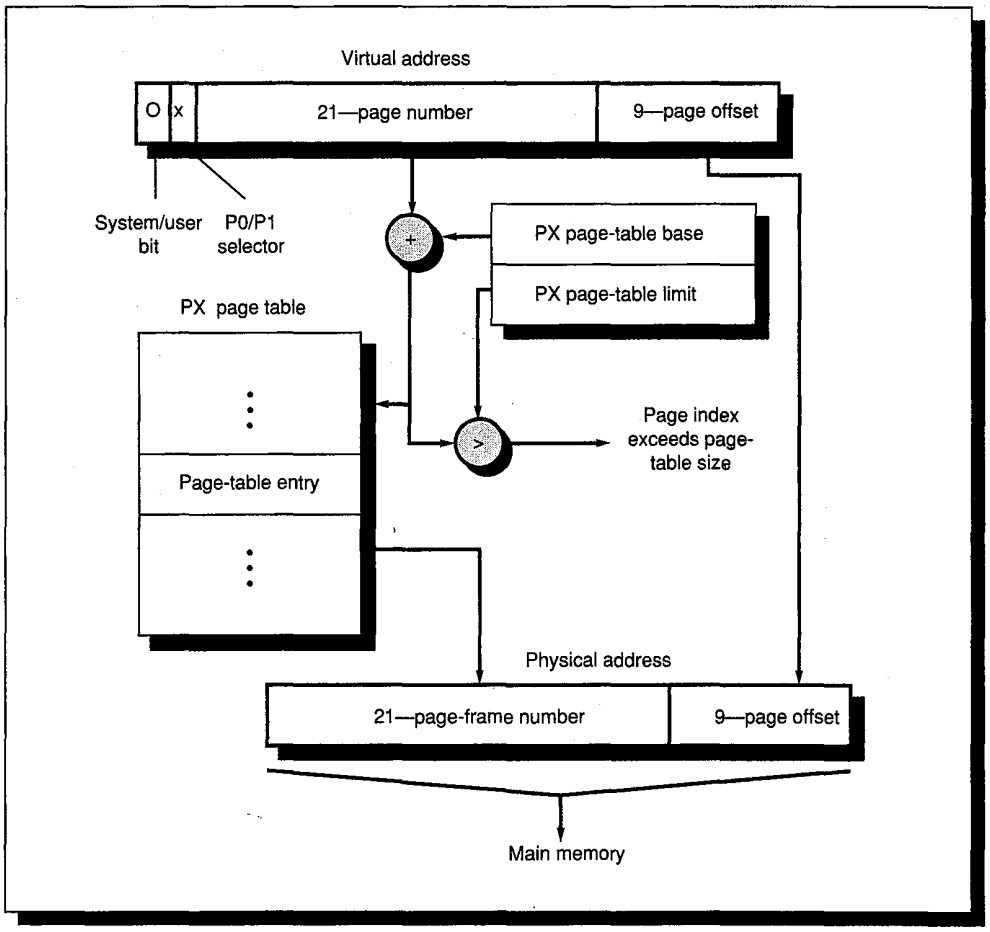

FIGURE 8.27 The mapping of a VAX virtual address. PX refers to either P0 or P1.

migrated to secondary storage? Some system page tables are loaded into physical memory when the operating system is booted and are prevented from migrating to disk. Thus, eventually a series of faults must cross an address stored in the system page table that is "frozen" into main memory.

While this explains translation of legal addresses, what prevents the user from creating illegal address translations and getting into mischief? The page tables themselves are protected from being written to by user programs. Thus, the user can try any virtual address, but by controlling the page-table entries the operating system controls what physical memory is accessed. Sharing of memory between processes is accomplished by having a page-table entry in each address space point to the same physical-memory page.

A *page-table entry* (PTE) on the VAX is straightforward. Other than the physical page-frame number these are the only architecture-defined fields:

M—the *modify bit* indicating the page is dirty

V—the *valid bit* indicating this PTE has a valid address

PROT—four protection bits

Note that there is no reference or use bit. Hence, a page-replacement algorithm such as LRU must rely on the modify bit or some software technique to measure usage. Rather than simply a kernel/user protection structure, the VAX uses a four-level structure consisting of kernel, executive, supervisor, and user. The four protection bits in the PTE contain 16 encodings of selected combinations of no access, read-only access, and read-write access, with the four security levels. For example, 1001 means read-write access for kernel and executive-level processes, read access for supervisor-level processes, and no access for user-level processes. To further isolate these four levels, each has its own stack and its own copy of the stack pointer (R15).

The first implementation of this architecture was the VAX-11/780, which employs a TLB to reduce address-translation time. Figure 8.28 shows the key parameters of this TLB.

Block size	1 PTE (4 bytes)
Hit time	1 clock cycle
Miss penalty (average)	22 clock cycles
Miss rate	1% – 2%
Cache size	128 PTEs (512 bytes)
Block selection	Random
Write strategy	(Not applicable)
Block placement	2-way set associative

FIGURE 8.28 Memory hierarchy parameters of the VAX-11/780 TLB.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 475

Figure 8.29 shows the VAX-11/780 TLB organization, with each step of a translation labeled. The TLB uses two-way–set-associative placement; thus, the translation begins (steps 1 and 2) by sending a portion of the virtual address ("index") to both sets to select the two tags that are to be compared. Of course, the tag must be marked valid to allow a match. At the same time, the type of memory access is checked for a violation (also in step 2) against protection information in the TLB.

For reasons similar to those in the cache case, there is no need to include the 9 bits of the VAX page offset in the TLB; nor is there reason to include the 6 address bits to index the TLB. The remaining bits are used in the comparison (step 3). The matching address tag sends the corresponding physical address through the multiplexer (step 4). The page offset is then combined with the physical page frame to form a full physical address (step 5).

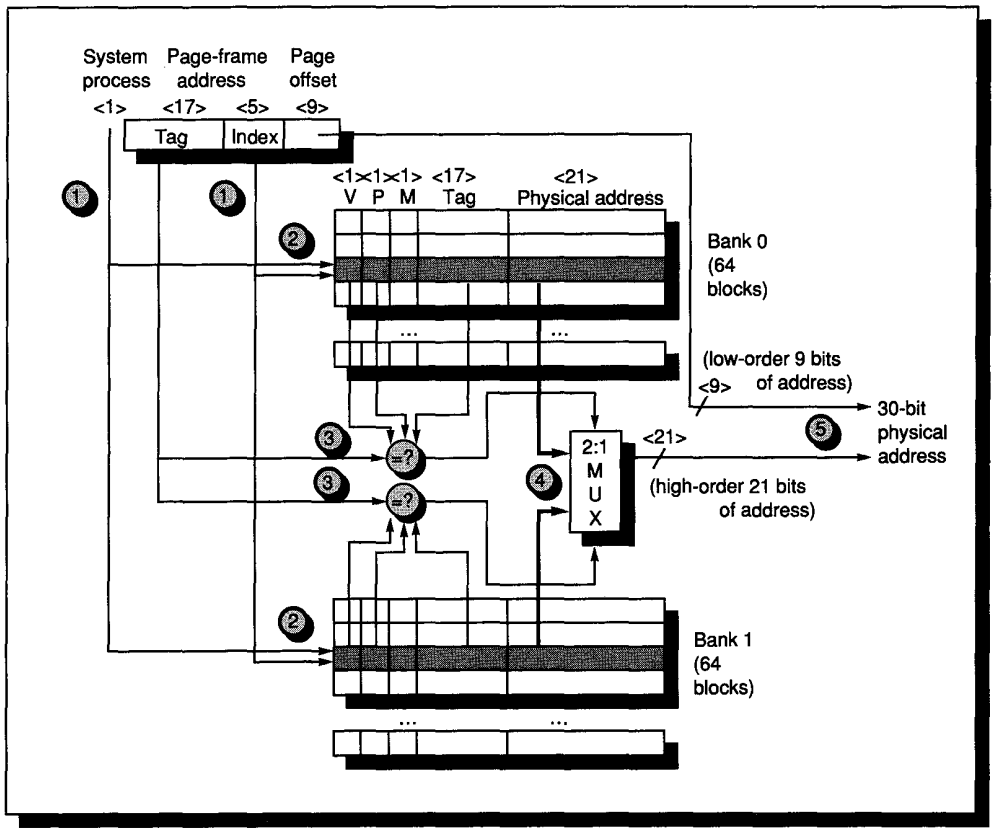

FIGURE 8.29 Operation of the VAX-11/780 TLB during address translation. The five steps of a TLB hit are shown as circled numbers.

There is one unusual feature of the VAX-11/780 TLB: The TLB is further subdivided to make sure the process portion of the address occupies no more than 50% of the TLB entries. The top 32 entries of each bank are reserved for system space, and the bottom 32 are reserved for process space. The most

significant bit of the address is used to select the appropriate half of the TLB (step 1). Since the system portion of the address space is the same for all processes, a process switch invalidates only the lower 32 entries of each bank for the VAX-11/780 TLB. This restriction had two goals. The first was to reduce the process-switch time by reducing the number of TLB entries that had to be invalidated; the second was to improve performance by preventing the system or user process from throwing out the other's translations when process switches were frequent. Splitting the TLB will usually lead to higher overall TLB miss rate, but may reduce the peak TLB miss rate in heavily process-switching environments.

A Segmented Virtual Memory Example: Protection in the Intel 80286/80386

The second system is the most dangerous system a man ever designs. . . . The general tendency is to over-design the second system, using all the ideas and frills that were cautiously sidetracked on the first one.

F. P. Brooks, Jr., *The Mythical Man-Month* (1975)

The original 8086 used segments for addressing, yet it provided nothing for virtual memory or for protection. Segments had base registers but no bound registers and no access checks; and before a segment register could be loaded the corresponding segment had to be in physical memory. Intel's dedication to virtual memory and protection is evident in subsequent models, with a few fields extended to support larger addresses.

Like the VAX, the 80286 has four levels of protection. The innermost level (0) corresponds to VAX kernel mode, and the outermost level (3) corresponds to VAX user mode. The 80286 also follows the VAX by having separate stacks for each level to avoid security breaches between the levels. There are also data structures analogous to VAX page tables that contain the physical addresses for segments, as well as a list of checks to be made on translated addresses.

The Intel designers did not stop there. The 80286 divides the address space, allowing both the operating system and the user access to the full space. The 80286 user can call an operating system routine in this space and even pass parameters to it retaining full protection. This is not a trivial action, since the stack for the operating system is different from the user's stack. Moreover, the 80286 allows the operating system to maintain the protection level of the called routine for the parameters that are passed to it. This potential loophole in protection is prevented by not allowing the user to ask the operating system to access something indirectly that he would not have been able to access himself. Such security loopholes are called *Trojan horses.*

The 80286 designers were guided by the principle of trusting the operating system as little as possible, while supporting sharing and protection. As an example of the use of such protected sharing, suppose a payroll program writes checks and also updates the year-to-date information on total salary and benefits payments. Thus, we want to give the program the ability to read the salary and

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 477

year-to-date information and modify the year-to-date information but not the salary. We shall see the mechanism to support such features shortly. In the rest of this section we will look at the big picture of the 80286 protection and examine its motivation. Readers interested in the detailed picture can find it in a comprehensive book by Crawford and Gelsinger [1987].

Adding Bounds Checking and Memory Mapping

The first step in enhancing the 80286 was getting the segmented addressing to check bounds as well as supply a base. Rather than a base address, as in the 8086, segment registers in the 80286 contain an index to a virtual memory data structure called a *descriptor table*. Descriptor tables play the role of page tables in the VAX. On the 80286 the equivalent of a page-table entry is a *segment descriptor*. It contains fields found in PTEs:

A *present bit*—equivalent to the PTE valid bit, used to indicate this is a valid translation

A *base field*—equivalent to a page-frame address, containing the physical address of the first byte of the segment

An *access bit*—like the reference bit or use bit in some architectures that is helpful for replacement algorithms

An *attributes field*—like the protection field in the VAX PTE, which specifies the valid operations and protection levels for operations that use this segment

There is also a *limit field*, not found in paged systems, which establishes the upper bound of valid offsets for this segment. Figure 8.30 shows examples of 80286 segment descriptors.

Adding Sharing and Protection

The Intel designers' next step was to provide for protected sharing. Like the VAX, half of the address space is shared by all processes and half is unique to each process, called *global address space* and *local address space*, respectively. Each half is given a descriptor table with the appropriate name. A descriptor pointing to a shared segment is placed in the global-descriptor table, while a descriptor for a private segment is placed in the local-descriptor table.

A program loads an 80286 segment register with an index to the table **and** a bit saying which table it desires. The operation is checked according to the attributes in the descriptor, the physical address being formed by adding the offset in the CPU to the base in the descriptor, provided the offset is less than the limit field. Unlike the encoding of operations and levels in the VAX PTE, every segment descriptor has a separate two-bit field to give the legal access level of this segment. A violation occurs only if the program tries to use a segment with a lower protection level in the segment descriptor.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 478

We can now show how to invoke the payroll program to update the year-to-date information without allowing it to update salaries. The program could be given a descriptor to the information that has the writable field clear, meaning it can read but not write the data. A trusted program can then be supplied that will only write the year-to-date information and is given a descriptor with the writable field set (Figure 8.30). The payroll program invokes the trusted code using a code-segment descriptor with the conforming field set (Figure 8.30). This means the called program takes on the privilege level of the code being called rather than the privilege level of the caller. Hence, the payroll program can read the salaries and call a trusted program to update the year-to-date totals, yet the payroll program cannot modify the salaries. If a Trojan horse exists in this system, to be effective it must be located in the trusted code whose only job is to update the year-to-date information. The argument for this style of protection is that limiting the scope of the vulnerability enhances security.

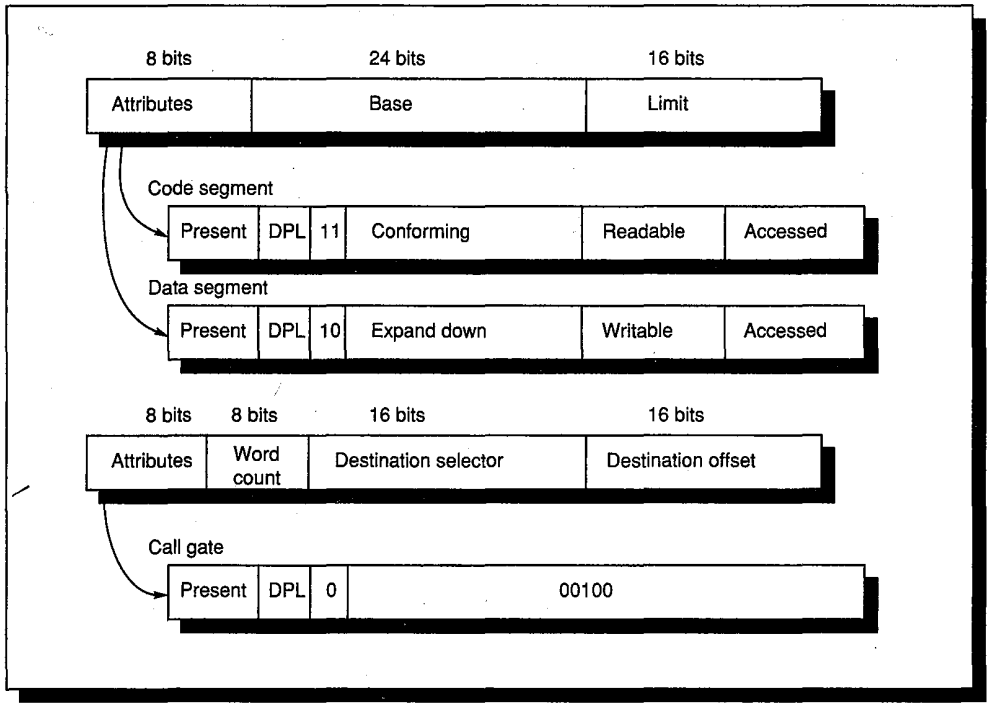

FIGURE 8.30 The 80286 segment descriptors are all 48 bits long and are distinguished by bits in the attributes field. *Base*, *limit*, *present*, *readable*, and *writable* are all self-explanatory. DPL means *descriptor privilege level*—this is checked against the code privilege level to see if the access will be allowed. *Conforming* says the code takes on the privilege level of the code being called rather than the privilege level of the caller; it is used for library routines. The *expand-down field* flips the check to let the base field be the high-water mark and the limit field be the low-water mark. As one might expect, this is used for stack segments that grow down. *Word count* controls the number of words copied from the current stack to the new stack on a call gate. The other two fields of the call-gate descriptor, *destination selector* and *destination offset*, select the descriptor of the destination of the call and the offset into it. There are many more than these three segment descriptors in the 80286. The principal change in the 80386 was to lengthen the base by eight bits and the limit by four bits.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 479

Adding Safe Calls from User to OS Gates and Inheriting Protection Level for Parameters

Allowing the user to jump into the operating system is a bold step. How, then, can a hardware designer increase the chances of a safe system without trusting the operating system or any other piece of code? The 80286 approach is to restrict where the user can enter a piece of code, to safely place parameters on the proper stack, and to make sure the user parameters don't get the protection level of the called code.

To restrict entry into others' code, the 80286 provides a special segment descriptor, or *call gate*, identified by a bit in the attributes field. Unlike other descriptors, call gates are full physical addresses of an object in memory; the offset supplied by the CPU is ignored. As stated above, their purpose is to prevent the user from randomly jumping anywhere into a protected or more- privileged code segment. In our programming example, this means the only place the payroll program can invoke the trusted code is at the proper boundary. This is needed to make conforming segments work as intended.

What happens if caller and callee are "mutually suspicious," so that neither trusts each other? The solution is found in the word-count field in the bottom descriptor in Figure 8.30 (page 447). When a call instruction invokes a call-gate descriptor, the descriptor will copy the number of words specified in the descriptor from the local stack onto the stack corresponding to the level of this segment. This allows the user to pass parameters by first pushing them onto the local stack. The hardware then safely transfers them onto the correct stack. A return from a call gate will pop the parameters off both stacks and copy any return values to the proper stack.

This still leaves open the potential loophole of having the operating system use the user's address, passed as parameters, with the operating system's security level, instead of with the user's level. The 80286 solves this problem by dedicating two bits in every CPU segment register to the *requested protection level*. When an operating system routine is invoked, it can execute an instruction that sets this two-bit field in all address parameters with the protection level of the user that called the routine. Thus, when these address parameters are loaded into the segment registers, they will set the requested protection level to the proper value. The 80286 hardware then uses the requested protection level to prevent any foolishness: No segment can be accessed from the system routine using those parameters if it has a more-privileged protection level than requested.

Summary: Protection on the VAX Versus the 80286

If the 80286 protection model looks harder to build than the VAX model, that's because it is. This effort must be especially frustrating for the 80286 engineers, since most customers just use the 80286 as a fast 8086 and don't exploit the elaborate protection mechanism. Also, the fact that the protection model is a

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 480

mismatch for the simple paging protection of UNIX means it will be used only by someone writing an operating system specially for this computer. OS/2 from Microsoft is the best candidate, but only time will tell whether the performance cost of such protection is justified for a personal-computer operating system. Two questions remain: Will the considerable protection-engineering effort, which must be borne by each generation of the 80x86 family, be put to good use, and will it prove any safer in practice than a paging system?

8.7 More Optimizations Based on Program Behavior

Making the frequent case fast is the inspiration for almost all inventions aimed at improving performance. In this section are two more examples of hardware optimized to program behavior. The first fetches instructions before they are needed, and the second avoids saving registers to memory on procedure calls.

Instruction-Prefetch Buffers

Many machines use an *instruction-prefetch buffer* to take advantage of the normal sequential execution of instructions. Typically, an instruction buffer contains two to eight sequential instructions; as each instruction is consumed by the CPU, a subsequent instruction word is prefetched. Prefetching only makes sense if the memory system can deliver instructions much faster than the CPU can consume them; otherwise the buffer cannot get ahead of the CPU. This can be accomplished by having a wider path that fetches more than one instruction at a time, or by simply having a faster memory system than the CPU. The drawback to instruction buffers is that they increase memory traffic by requesting words of instructions that may never be needed by the CPU, as is the case when a branch is taken. Instruction-prefetch buffers are also useful for aligning variable-sized instructions.

The 8-byte instruction-prefetch buffer (IB) of the VAX-11/780, shown in Figure 8.31 (page 450), will serve as an example. The opcode of the current instruction is in the high-order byte of the IB; as pieces of the instruction are consumed, the whole buffer is shifted to the left by the appropriate amount. The left-most byte can correspond to any byte address, while the rest of the bytes in the IB must be sequential. The Vs in the figure represent a valid bit per byte of the instruction buffer and indicate the sequential bytes that contain valid instructions.

The IB tries to stay ahead of the PC. Whenever at least one byte is free in the IB, a read is requested for an aligned 32-bit word that contains that byte; only 32-bit words are prefetched from the memory. When the 32-bit prefetched word arrives, the IB loads as much of it as it has space for. A 32-bit instruction word therefore takes between one and four fetches from memory, depending on luck.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 481

When the PC changes due to a branch or interrupt, the IB may have prefetched one or two unneeded instructions. The PC change causes all the valid bits to be turned off, and the IB is reloaded. Section 8.9 examines the performance impact of the IB.

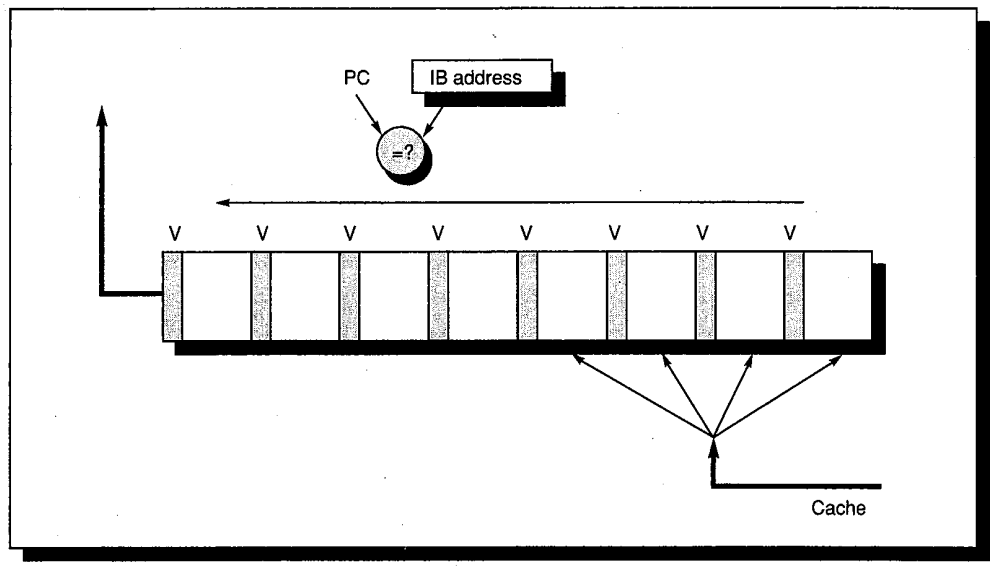

FIGURE 8.31 The VAX-11/780 instruction-prefetch buffer. Every byte has a valid bit to determine the number of consecutive bytes that have valid instructions. The instruction decoder can read the top four bytes of the buffer in a single clock cycle.

Registers and Register Windows

Figures 3.28 and 3.29 (pages 117–118) in Chapter 3 show that saving registers on procedure calls and restoring them on returns can account for 5% to 40% of the data memory references. As an alternative, several banks of registers can be used, with a new one allocated on each call. Although this could limit the depth of procedure calls, the limitation is avoided by operating the banks as a circular buffer, providing unlimited depth. This technique has been termed *register windows*.

Figure 8.32 shows the essence of the idea. On the x axis is time, measured in procedure calls or returns; on the y axis is the depth or nesting of procedure calls. Each call moves down the y axis, and each return moves up. The boxes show memory being accessed to save some of the buffer, either when it is full and is followed by a call (*window overflow*) or when it is empty and is followed by a return (*window underflow*). The figure shows eight window overflows and two window underflows during this section of program execution. Over the life of the program the number of overflows and underflows will equalize.

One might well ask what the trade-off is between buffer size and overflows or underflows. Figure 8.33 shows the shape of the curve for several programs written in several programming languages. The knee of the curve seems to be six to eight banks. While this holds for most programs, the optimization is based on

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 482

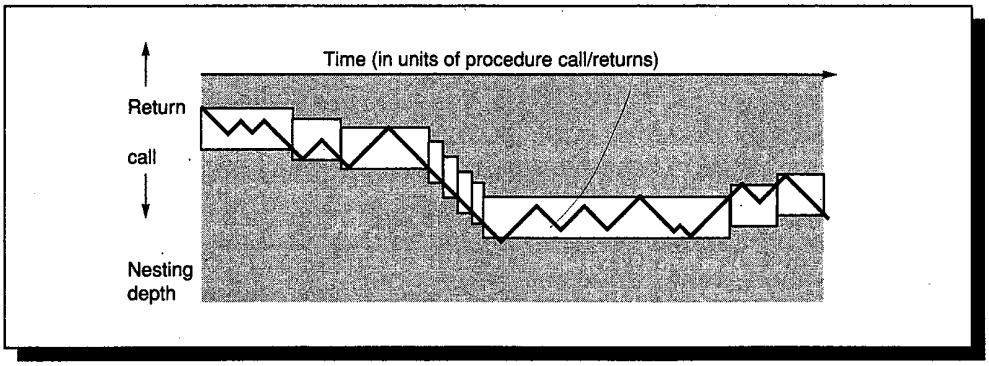

FIGURE 8.32 Change in procedure nesting depth over time. The boxes show procedure calls and returns inside the buffer before a window overflow or underflow. The program starts with three calls, a return, a call, a return, three calls, and then a window overflow.

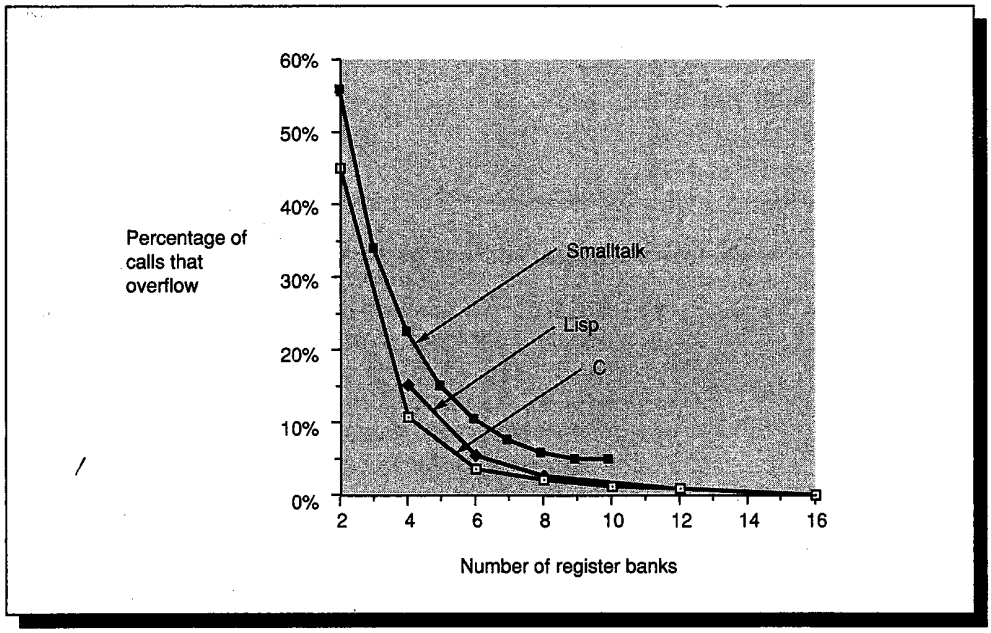

FIGURE 8.33 Number of banks or windows of registers versus overflow rate for several programs in C, LISP, and Smalltalk. The programs measured for C include a C compiler, a Pascal interpreter, troff, a sort program, and a few UNIX utilities [Halbert and Kessler 1980].The LISP measurements include a circuit simulator, a theorem prover, and several small LISP benchmarks [Taylor et al. 1986]. The Smalltalk programs come from the Smalltalk macro benchmarks [McCall 1983] which include a compiler, browser, and decompiler [Blakken 1983 and Ungar 1987].

program-specific patterns of calls and returns that might be quite different in some other programs. The worst case for register windows would be hundreds of calls followed by hundreds of returns. This would make Figure 8.32 look like seismograph output during an earthquake, and the performance impact would be just as devastating!

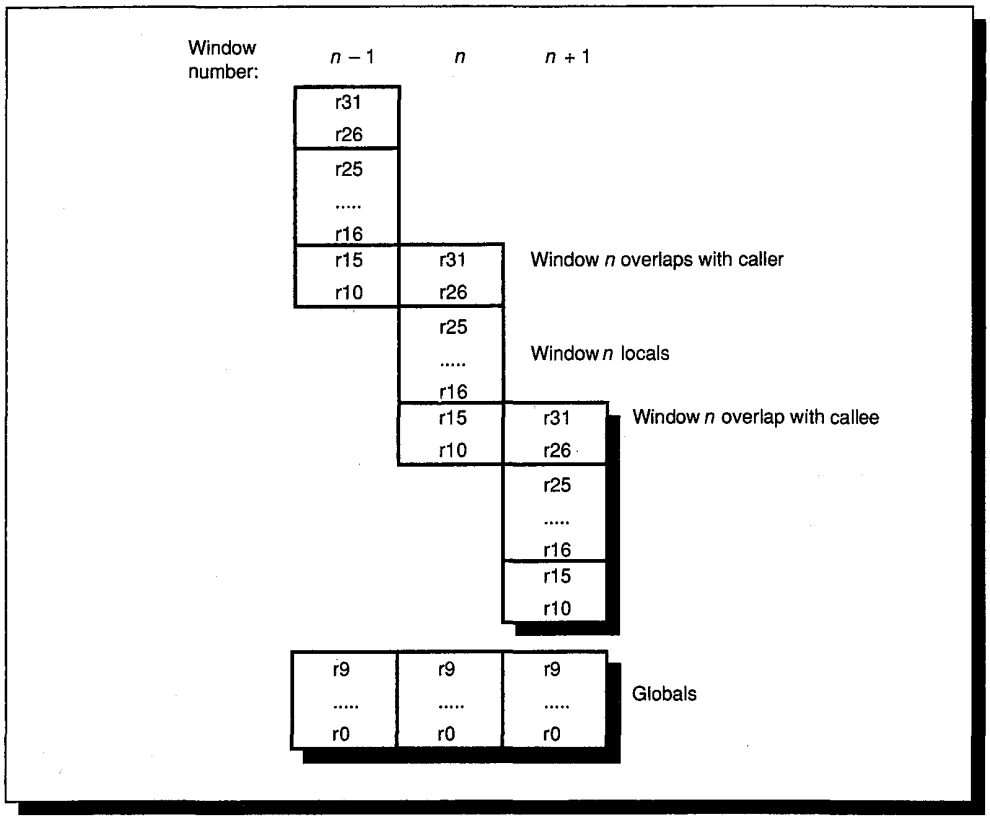

FIGURE 8.34 Parameters can be passed in registers if there are common registers between two banks or windows. This scheme divides registers into globals, which don't change on a procedure call, and locals, which do change. By having an overlap between locals for adjacent procedure calls and renumbering the registers on a call, the outgoing parameters of the caller become the incoming parameters of the callee. For example, a value placed in register 15 before a call is in register 31 after the call.

The difficulty of passing parameters in registers presents a drawback: If each procedure has its own unique set of registers, then nothing is common. This can be overcome by overlapping the register banks or windows such that there is a common area in which to pass parameters. Figure 8.34 shows one such design. Six registers overlap each window, with R15 to R10 of the caller's registers becoming R31 to R26 after the call. Ten registers are not included in the windows, so there are 16 (32 – 10 – 6) registers per window even though each procedure sees 32 registers at a time.

From Figure 8.33 we can estimate the percentage of calls that overflow the windows or returns that underflow them, but to understand the impact on performance we must know the cost an overflow or underflow. With an overlapping register design, like the one on SPARC, the cost is saving 16 registers on an overflow (or restoring 16 registers on an underflow) plus the cost of interrupt. On the Sun 4 today it takes about 60 clock cycles for an overflow or underflow.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 484

The Pros and Cons of Register Windows

Depending on the application, programming language, and user practices, the compiler can close the gap between machines with and without register windows. Most machines, for example, have separate floating-point registers, which means that floating-point-intensive programs will be unaffected by register windows. Also, many data references are to objects that cannot be allocated in registers, like arrays or structures (see Figures 3.28 and 3.29 on pages 117–118 of Chapter 3).

An optimization called *interprocedural register allocation* allows more intelligent allocation of registers across procedure boundaries. Unfortunately, interprocedural register allocation works best when procedures are compiled or linked at the same time. Long compilation and link time do not match the emphasis on a rapid debug-edit-compile cycle in current dynamic languages like LISP and Smalltalk. Interprocedural register allocation is not generally applicable to object-oriented languages like Objective C and Smalltalk because in the dynamic equivalent of a procedure call the compiler doesn't know which procedure will be invoked on such calls. Register windows also simplify some compiler decisions, since there is no extra cost in using a register that will not be saved or restored separately.

	GCC	TeX
Percentage of DLX instructions call or return	1.8%	3.6%
Registers stored per call	2.3	3.2
Loads DLX	3,928,710	2,811,545
Loads SPARC	3,313,317	2,736,979
Ratio loads DLX / SPARC	1.20	1.03
Stores DLX	2,037,226	1,974,078
Stores SPARC	1,246,538	1,401,186
Ratio stores DLX / SPARC	1.60	1.41

FIGURE 8.35 Benefits of register windows on loads and stores for non–floating-point programs. The first row shows the percentage of DLX instructions executed that are calls or returns. The second row shows the average number of register saves and restores per call on the DLX architecture with optimization level O2. The following rows show the total number of loads and stores for each optimization and for the SPARC architecture, which has register windows. The data below includes the loads and stores due to window overflow and window underflow. GCC executes about 20% more loads and 60% more stores on DLX than on a machine with register windows, while TeX executes about 3% more loads and 41% more stores. These savings correspond to about 7% of the instruction count for GCC and 5% for TeX. How this translates into memory-system performance depends on the details of the rest of the memory hierarchy. Interprocedural register allocation closes this gap. For example, using O3 optimization on TeX reduces the number of DLX loads by 5% to 2,671,631 and the number of stores by 10% to 1,791,831. Note that the inputs for these programs were not the same as those used in Chapters 2 or 4. (Spice was not included because register windows offer no benefit for floating-point programs.)

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 485

The danger of register windows is that the larger number of registers could slow down the clock rate. So far, this has not been the case for commercial machines. The SPARC architecture (with register windows) and the MIPS R2000 architecture (without) are contemporary machines built in several technologies. The SPARC clock rate has not been slower than MIPS for implementations in similar technologies, probably because cache-access times dominate register-access times in implementations to date of either architecture. A second concern is the impact of register windows on process-switch time. Sun Microsystems has found that UNIX operating system vagaries dominate process-switch time, and less than 20% of the process-switch time is spent on saving or restoring registers. Figure 8.35 (page 453) compares some measures of the benefits of register windows on our benchmark programs.

8.8 Advanced Topics—Improving Cache-Memory Performance

This section covers advanced topics in cache memories, going through new ideas at a much quicker pace than previous sections. The central points of this chapter are not lost if this section is skipped; in fact, the Putting It All Together section that follows is independent of this material.

The increasing gap between CPU and main memory speeds has attracted the attention of many architects. After making some easy decisions in the beginning, the architect faces a threefold dilemma when attempting to further reduce average access time:

- Increasing block size doesn't improve average access time; the lower miss rate doesn't offset the higher miss penalty.
- Making the cache bigger would make it slower, jeopardizing the CPU clock rate.
- Making the cache more associative would also make it slower, again jeopardizing the CPU clock rate.

Moreover, the miss rate calculated from user programs paints too rosy a picture. Figure 8.36 shows the real cache miss rate for a running program, including the operating system code invoked by the programs. This reveals the average access time to be worse than expected.

This section covers a plethora of techniques for improving cache performance: subblock placement, write buffers, out-of-order fetching, virtually addressed caches, two-level caches, and issues relating to cache coherency. The cache-coherency sections include an example of the stale-data problem, a survey of coherency alternatives, an example cache protocol, a synchronization algorithm used in cache coherent multiprocessors, a timeline showing multiprocessor synchronization, and comments about the impact of memory consistency on parallel processors.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 486

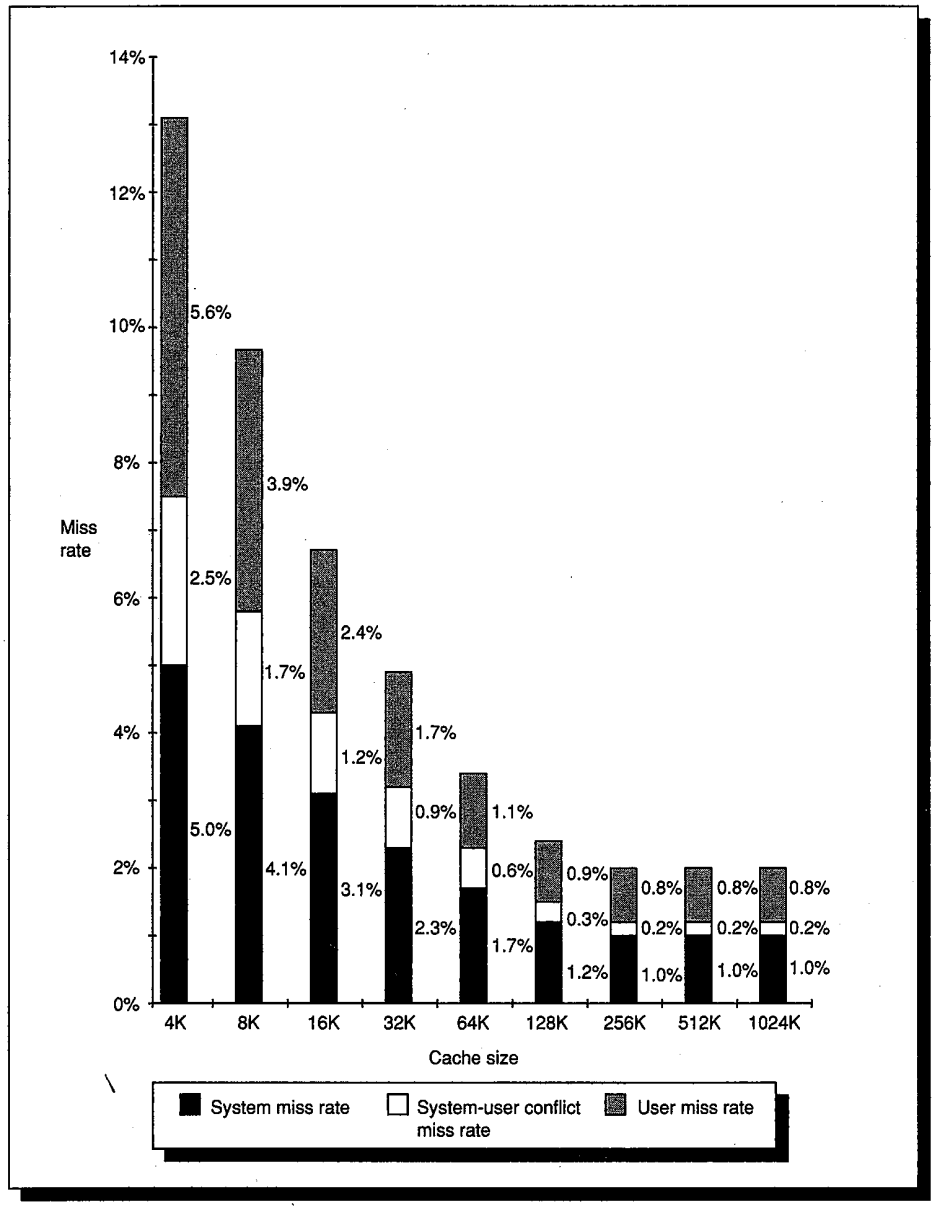

FIGURE 8.36 The miss rate of a program, including the operating system code it invokes, versus cache size. The top category is what would be measured from a user trace; the bottom category is the miss rate for the operating system code; and the middle category is the miss rate due to conflicts between the user code and system code. Agarwal [1987] collected these statistics for the Ultrix operating system running on a VAX, assuming direct-mapped caches with a block size of 16 bytes.

Reducing Hit Times—Making Writes Faster

As mentioned before, writes usually take more than one clock cycle because the tag must be checked before writing the data. There are two ways to do faster writes.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 487

The first, used on the VAX 8800, pipelines the writes for a write-through cache. Tags and data are split so that they can be addressed independently. As usual, the cache compares the tag with the current write address. The difference is that the memory access during this comparison uses the address and data from the **previous** write. Therefore, writes can be performed back to back at one per clock cycle because the CPU does not have to wait for the write to the cache if the first stage is a hit. The 8800 pipeline does not affect read hits—the second stage of the write occurs during the first stage of the next write or during a cache miss.

Another way of reducing writes to one clock cycle involves caches that must be direct mapped, using a technique known as *subblock placement*. Like the VAX-11/780 instruction buffer, there is a valid bit on units smaller than the full block, called *subblocks*. The valid bits specify some parts of the block as valid and some parts as invalid. A match of the tag doesn't mean the word is necessarily in the cache, as the valid bits for that word must also be on. Figure 8.37 gives an example. Note that for caches with subblock placement a block can no longer be defined as the minimum unit transferred between cache and memory. For such caches a block is defined as the unit of information associated with an address tag.

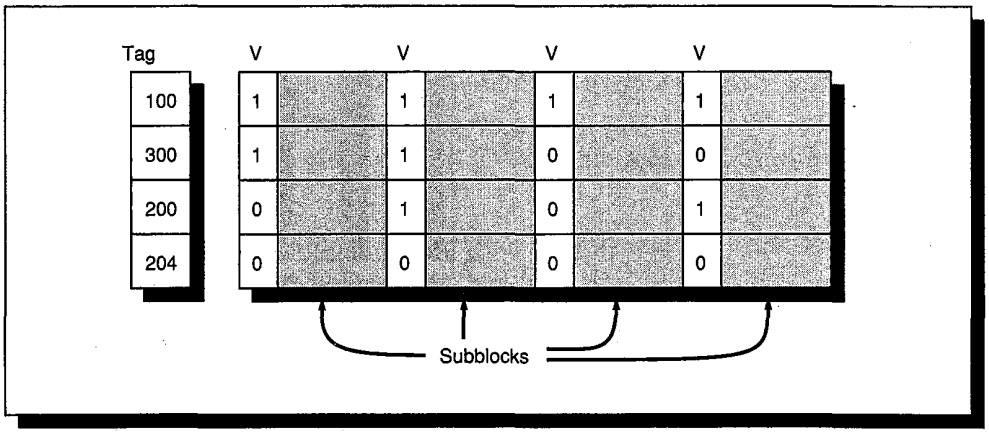

FIGURE 8.37 In this example there are four subblocks per block. In the first block (top) all the valid bits are on, equivalent to the valid bit being on for a block in a normal cache. In the last block (bottom), the opposite is true; no valid bits are on. In the second block, locations 300 and 301 are valid and will be hits, while locations 302 and 303 will be misses. For the third block, locations 201 and 203 are hits. If, instead of this organization, there were 16 blocks the size of the subblock, 16 tags would be needed instead of 4.

Subblock placement was invented to reduce the long miss penalty of large blocks (since only a part of a large block need be read) and to reduce the tag storage for small caches. It can also help write hits by **always** writing the word (no matter what happens with the tag match), turning the valid bit on, and then sending the word to memory. Let's look at the cases to see why this trick works:

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 488

- *Tag match and valid bit already set.* Writing the block was the proper action, and nothing was lost by setting the valid bit on again.
- *Tag match and valid bit not set.* The tag match means that this is the proper block; writing the data into the block makes it appropriate to turn the valid bit on.
- *Tag mismatch.* This is a miss and will modify the data portion of the block. However, as this is a write-through cache, no harm was done; memory still has an up-to-date copy of the old value. Only the tag to the address of the write need be changed because the valid bit has already been set. If the block size is one word and the store instruction is writing one word, then the write is complete. When the block is larger than a word or if the instruction is a byte or halfword store, then either the rest of the valid bits are turned off (allocating the subblock without fetching the rest of the block) or memory is requested to send the missing part of the block (write allocate).

This trick isn't possible with a write-back cache because the only valid copy of the data may be in the block, and it could be overwritten before checking the tag.

Reducing Miss Penalty—Making Write Misses Faster

Now that we have seen how to make write hits faster, let's look at write misses. With a write-through cache the most important improvement is a write buffer (page 416) of the proper size (see the fallacy on page 482 in Section 8.10). Write buffers, however, do complicate things in that they might have the updated value of a location needed on a read miss.

Example

Look at this code sequence:

```
SW   512(R0),R3   ; M[512] ← R3 (cache index 0)
LW   R1,1024(R0)  ; R1 ← M[1024] (cache index 0)
LW   R2,512(R0)   ; R2 ← M[512] (cache index 0)
```

Assume a direct-mapped cache that maps 512 and 1024 to the same block, and a four-word write buffer. Will R3 always equal R2?

Answer

Let's follow the cache to see the danger. The data in R3 is placed into the write buffer after the store. The following load uses the same cache index and is therefore a miss. We then try to load the data from location 512 into register R2; this also results in a miss. If the write buffer hasn't completed writing to location 512 in memory, the read of location 512 will put the old, wrong value into the cache block, and then into R2. Without proper precautions, R3 would not be equal to R2!

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 489

The simplest way out of this dilemma is for the read miss to wait until the write buffer is empty. However, a write buffer of a few words in a write-through cache will almost always have data in the buffer on a miss, thereby increasing the read miss penalty. The designers of the MIPS M/1000 estimated that waiting for a four-word buffer to empty would have increased the average read miss penalty by 50%. The alternative is to check the contents of the write buffer on a read miss, and if there are no conflicts and the memory system is available, let the read miss continue.

The cost of writes in a write-back cache can also be reduced. By just adding a full block buffer to store a dirty block, the read can happen first. After the new data is loaded into the block, the CPU continues execution. The buffer then writes in parallel with the CPU. Similar to the situation above, if a read miss occurs the CPU can stall until the buffer is empty.

Reducing Miss Penalty—Making Read Misses Faster

Making writes faster is helpful, but it is reads that dominate cache accesses. The strategy to making read misses faster is to be impatient: Don't wait for the full block to be loaded before sending the requested word to the CPU. Here are two specific strategies:

- *Early restart*—As soon as the requested word of the block arrives, send it to the CPU and let the CPU continue execution.
- *Out-of-order fetch*—Request the missed word first from memory and send it to the CPU as soon as it arrives; let the CPU continue execution while filling the rest of the words in the block. Out-of-order fetch is also called *wrapped* fetch.

Alas, these read tricks are not as important as they sound. Spatial locality—the reason for big blocks in the first place—dictates that the next cache request is likely to be to the same block. Also, handling another request while trying to fill the rest of a block quickly gets complicated.

A more subtle reason why out-of-order fetch will not be as rewarding as one might think is that not all the words of a block have an equal likelihood of being accessed first. With a 16-word block in an instruction cache, for example, the average block entry point is 2.8 words from the left-most byte. If entries were evenly distributed, the average would be 8 words. The high-order word is the most likely one, due to sequential accesses from prior blocks on instruction fetches and sequentially stepping through arrays for data caches.

For pipelined machines that allow out-of-order completion using a scoreboard or Tomasulo-style control (Section 6.7 of Chapter 6), the CPU need not stall on a cache miss, offering another way to reduce memory stalls. Spatial locality suggests this optimization (called a *lock-up free cache*) may be limited in practice, since again the next reference is likely to be to the same block.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 490

Making Cache Hits Faster—Virtually Addressed Caches

Miss penalty is an important part of average access time, but hit time affects both the average access time and the clock rate of the CPU. Helping the hit time may therefore help everything. A solution mentioned earlier is to use the physical part of the address to index the cache while sending the virtual address through the TLB. The limitation is that a direct-mapped cache can be no bigger than the page size. To allow large cache sizes with the 4-KB pages in the System/370, IBM uses high associativity so that they can still access the cache with a physical index. The IBM 3033, for example, is 16-way set associative, even though studies show there is little benefit to miss rates above 4-way set associativity.

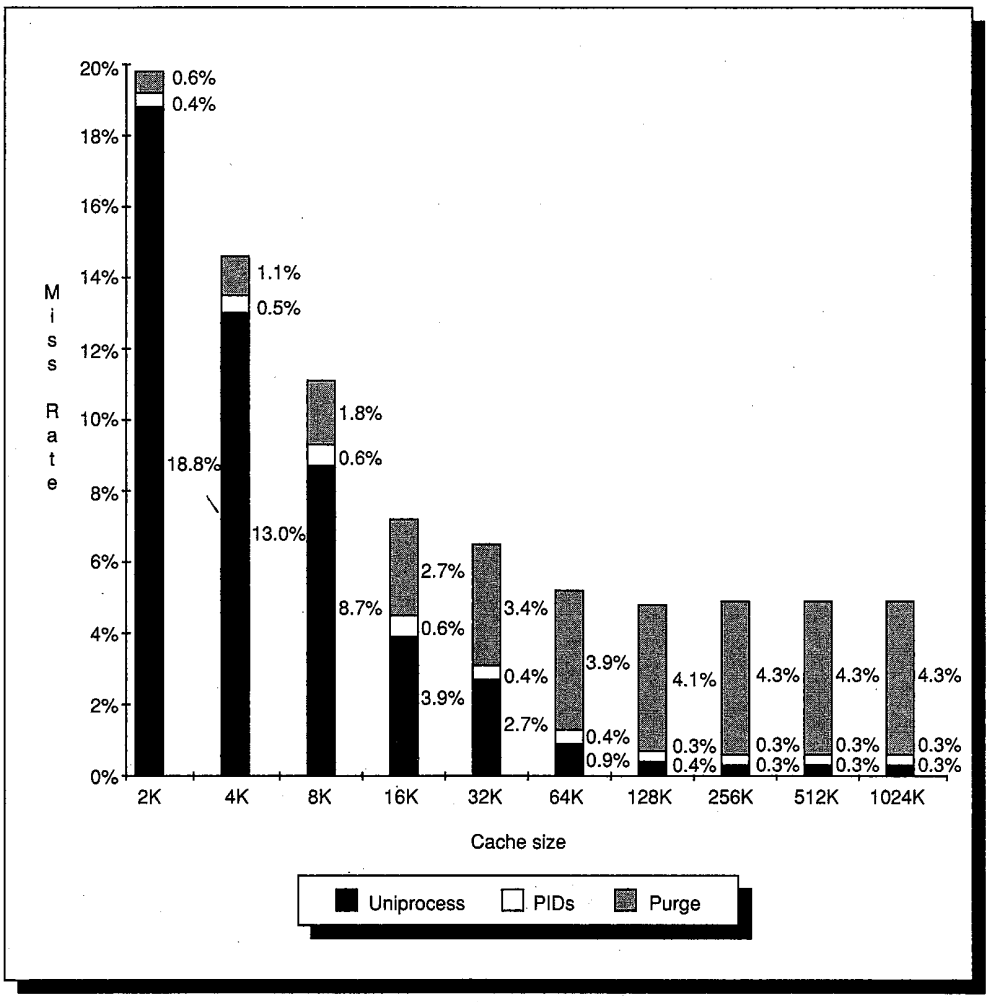

FIGURE 8.38 Miss rate versus cache size of a program measured three ways: without process switches (uniprocess), with process switches using a process-identifier tag (PIDs), and with process switches but without PIDs (purge). PIDs increase the uniprocess absolute miss rate by 0.3 to 0.6 and save 0.6 to 4.3 over purging. Agarwal [1987] collected these statistics for the Ultrix operating system running on a VAX, assuming direct-mapped caches with a block size of 16 bytes.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 491

One scheme for fast cache hits without this size restriction is go to a more heavily pipelined memory access where the TLB is just one step of the pipeline. The TLB is a distinct unit that is smaller than the cache, and thus easily pipelined. This scheme doesn't change memory latency, but relies on the efficiency of the CPU pipeline to achieve higher memory bandwidth.

Another alternative is to match on virtual addresses directly. Such caches are termed *virtual caches*. This eliminates the TLB translation time from a cache hit. Why doesn't everyone build virtually addressed caches? One reason is that every time a process is switched, the virtual addresses refer to different physical addresses, requiring the cache to be flushed. Figure 8.38 (page 459) shows the impact on miss rates of this flushing. One solution is to increase the width of the cache-address tag with a *process-identifier tag* (PID). If the operating system assigns these tags to processes, it only need flush the cache when a PID is recycled (the PID provides protection). Figure 8.38 shows that improvement.

Another reason why virtual caches are not more universally adopted has to do with operating systems and user programs that use two different virtual addresses for the same physical address. These duplicate addresses, called *synonyms* or *aliases*, could result in two copies of the same data in a virtual cache; if one is modified, the other will have the wrong value. With a physical cache this wouldn't happen, since the accesses would first be translated to the same physical cache block. There are hardware schemes, called *anti-aliasing*, that can guarantee every cache block a unique physical address, but software can make this much easier by forcing aliases to share some address bits. The version of UNIX from Sun Microsystems, for example, requires all aliases to be identical in the last 18 bits of their addresses. Thus, a direct-mapped cache that is 2^{18} (256K) bytes or smaller can never have duplicate physical addresses for blocks. This requirement also simplifies anti-aliasing hardware for larger caches or for set-associative caches. (Of course, the best software solution from the hardware designers perspective is to do away with aliases!)

The final area of concern with virtual addresses is I/O. I/O typically uses physical addresses and thus would require mapping to virtual addresses to interact with a virtual cache. (The impact of I/O on caches is further discussed below.)

Reducing Miss Penalty—Two-Level Caches

Let's return our attention to miss penalty. CPUs are getting faster and main memories are getting larger, but slower relative to the faster CPUs. The question facing the architect is: Should I make the cache faster to keep pace with the speed of CPUs, or make the cache larger to overcome the widening gap between the CPU and main memory? One answer is: Both. By adding another level of cache between the original cache and memory, the first-level cache can be small enough to match the clock cycle time of the CPU while the second-level cache can be large enough to capture many accesses that would go to main memory.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 492

Definitions for a second level of cache are not always straightforward. Let's start with the definition of *average memory-access time* for a two-level cache. Using the subscripts L1 and L2 to refer respectively to a first-level and a second-level cache, the original formula is

$$\text{Average memory-access time} = \text{Hit time}_{L1} + \text{Miss rate}_{L1} * \text{Miss penalty}_{L1}$$

and

$$\text{Miss penalty}_{L1} = \text{Hit time}_{L2} + \text{Miss rate}_{L2} * \text{Miss penalty}_{L2}$$

so

$$\text{Average memory-access time} = \text{Hit time}_{L1} + \text{Miss rate}_{L1} * (\text{Hit time}_{L2} + \text{Miss rate}_{L2} * \text{Miss penalty}_{L2})$$

In this formula, the success of the second-level miss rate is measured on the leftovers from the first-level cache. To avoid ambiguity, these terms are adopted here for a two-level cache system:

- *Local miss rate*—The number of misses in the cache divided by the total number of memory accesses to this cache; this is miss rate_{L2} above.
- *Global miss rate*—The number of misses in the cache divided by the total number of memory accesses generated by the CPU; using the terms above, this is miss rate_{L1} * miss rate_{L2}.

Example

Suppose that in 1000 memory references there are 40 misses in the first-level cache and 20 misses in the second-level cache. What are the various miss rates?

Answer

The miss rate for the first-level cache is 40/1000 or 4%. The local miss rate for the second-level cache is 20/40 or 50%. The global miss rate of the second-level cache is 20/1000 or 2%.

Figure 8.39 (page 462) and Figure 8.40 (page 463) show how miss rates and relative execution time change with the size of a second-level cache. Figure 8.41 (page 463) shows typical parameters of second-level caches.

With these definitions in place, we can consider the parameters of second-level caches. The foremost difference between the two levels is that the speed of the first-level cache affects the clock rate of the CPU, while the speed of the second-level cache only affects the miss penalty of the first-level cache. Thus, we can consider many alternatives in the second-level cache that would be ill chosen for the first-level cache. There is but one consideration for the design of the second-level cache: Will it lower the average memory-access–time portion of the CPI?

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 493

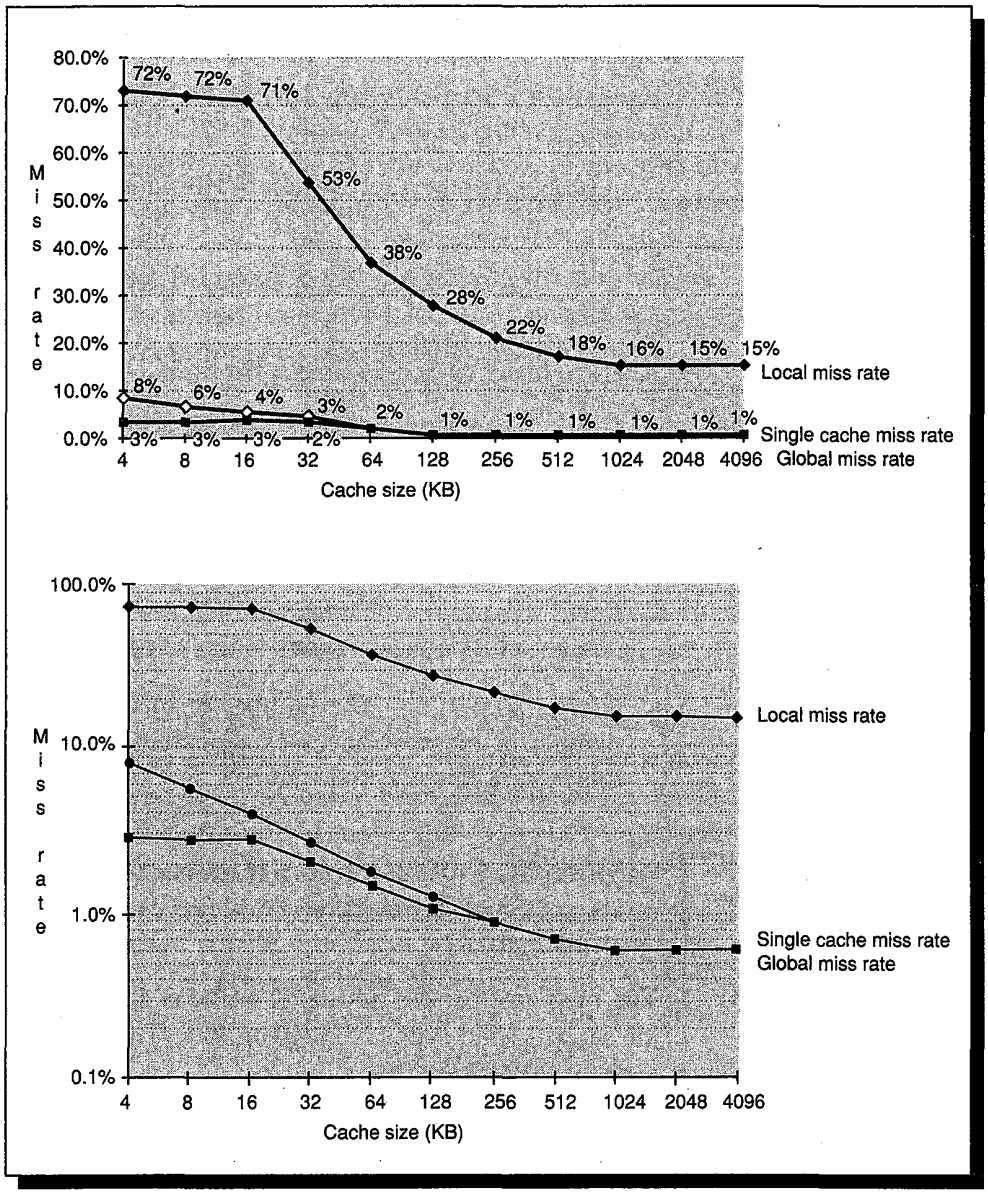

FIGURE 8.39 Miss rates versus cache size. The top graph shows the results plotted on a linear scale as we have done with earlier figures, while the bottom graph shows the results plotted on a log scale. As miss rates shrink the log scale makes the differences easier to follow. The miss rate of a single-level cache versus size is plotted against the local miss rate and global miss rate of a second-level cache using a 32-KB first-level cache. Second-level caches **smaller** than the 32-KB first level have high miss rates (at least for similar block sizes), as this figure illustrates. After 256 KB the single cache and global miss rates are virtually identical. Przybylski [1990] collected these data using traces available with this book: four traces from the VAX system and user programs and four user programs from the MIPS R2000 that were randomly interleaved to duplicate the effect of process switches.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 494

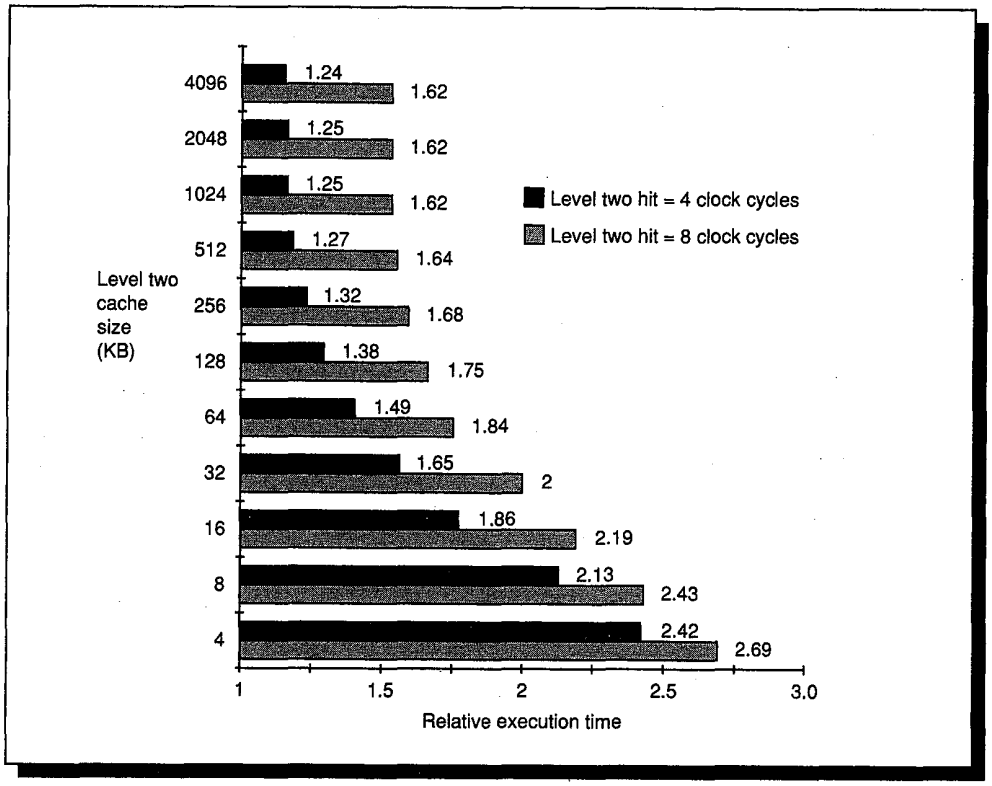

FIGURE 8.40 Relative execution time by second-level–cache size. Przybylski [1990] collected these data using a 32-KB, first-level, write-back cache, varying the size of the second-level cache. The two bars are for different clock cycles for a level two cache hit. The reference execution time of 1.00 is for a 4096-KB, second-level cache with a one–clock-cycle latency on a second-level hit. He used four traces from the VAX system and user programs (available with this book) and four user programs from the MIPS R2000 that were randomly interleaved to duplicate the effect of process switches.

Block (line) size	32 – 256 bytes
Hit time	4 – 10 clock cycles
Miss penalty	30 – 80 clock cycles
(Access time)	(14 – 18 clock cycles)
(Transfer time)	(16 – 64 clock cycles)
Local miss rate	15% – 30%
Cache size	256 KB – 4 MB

FIGURE 8.41 Typical values of key memory-hierarchy parameters for second-level caches.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 495

The initial choice for second-level caches is size. Since everything in the first-level cache is likely to be in the second-level cache, the second-level cache should be bigger. If second-level caches are just a little bigger, the local miss rate will be high. This observation inspires design of huge second-level caches—the size of main memory in recent computers! If the second-level cache is much larger than the first-level cache, then the global miss rate is about the same as a single-level cache of the same size (see Figure 8.39, page 462). Large size means that the second-level cache may have practically no capacity misses, leaving compulsory and a few conflict misses for our attention. One question is whether set associativity makes more sense for second-level caches.

Example

Given the data below, what is the impact of second-level–cache associativity on the miss penalty?

- Two-way set associativity increases hit time by 10% of a CPU clock cycle
- Hit time$_{L2}$ for direct mapped = 4 clock cycles
- Local miss rate$_{L2}$ for direct mapped = 25%
- Local miss rate$_{L2}$ for two-way set associative = 20%
- Miss penalty$_{L2}$ = 30 clock cycles

Answer

For a direct-mapped, second-level cache, the first-level–cache miss penalty is

$$\text{Miss penalty}_{L1} = 4 + 25\% * 30 = 11.5 \text{ clock cycles}$$

Adding the cost of associativity increases the hit cost only 0.1 clock cycles, making the new first-level–cache miss penalty

$$\text{Miss penalty}_{L1} = 4.1 + 20\% * 30 = 10.1 \text{ clock cycles}$$

In reality, second-level caches are almost always synchronized with the first-level cache and CPU. Accordingly, the second-level hit time must be an integral number of clock cycles. If we are lucky, we can shave the second-level hit time to four cycles; if not, we can round up to five cycles. Either choice is an improvement over the direct-mapped, second-level cache:

$$\text{Miss penalty}_{L1} = 4 + 20\% * 30 = 10.0 \text{ clock cycles}$$

$$\text{Miss penalty}_{L1} = 5 + 20\% * 30 = 11.0 \text{ clock cycles}$$

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 496

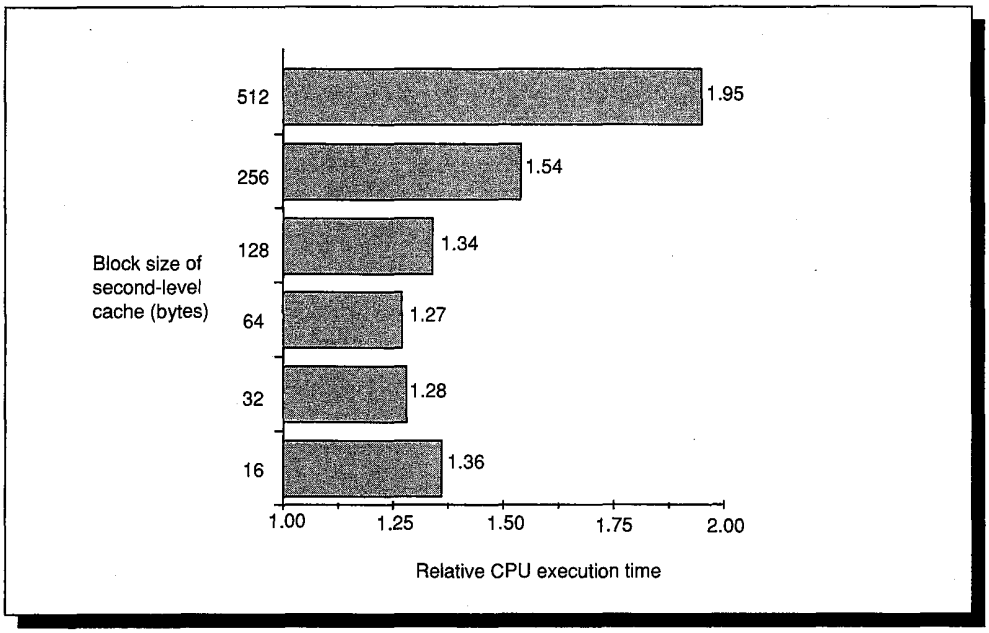

FIGURE 8.42 Relative execution time by block size for a two-level cache. Przybylski [1990] collected these data using a 512-KB second-level cache. He used four traces from the VAX system and user programs (available with this book) and four user programs from the MIPS R2000 that were randomly interleaved to duplicate the effect of process switches.

Higher associativity is worth considering because it has small impact on the second-level hit time and because so much of the average access time is due to misses. However, for these very large caches the benefits of associativity diminish because larger size has eliminated many conflict misses.

As long as spatial locality holds there may be a benefit in increasing block size. Increasing block size can increase conflict misses with small caches since there may not be enough places to put data, therefore increasing miss rate. Because this is not an issue in large, second-level caches, and because memory-access time is relatively longer, larger block sizes are popular. Figure 8.42 shows the variation in execution time as the second-level block size changes.

One final consideration concerns whether all data in the first-level cache is always in the second-level cache. If so, the second-level cache is said to have the *multilevel inclusion property*. Inclusion is desirable because consistency between I/O and caches (or between caches in a multiprocessor) can be determined just by checking the second-level cache.

The drawback to this natural inclusion is that the lower average memory-access times can suggest smaller blocks for the smaller first-level cache and larger blocks for the larger second-level cache. Inclusion can still be maintained in this case with a little extra work on a second-level miss: The second-level cache must invalidate all first-level blocks that map onto the second-level block to be replaced, causing a slightly higher first-level miss rate.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 497

Reducing Miss Rate by Reducing Cache Flushes—I/O

Although there is little more that can improve CPU execution time, there are issues in cache design to improve system performance, particularly for input/output. Because of caches, data can be found in memory or in the cache. As long as the CPU is the sole device changing or reading the data and the cache stands between the CPU and memory, there is little danger in the CPU seeing the old or *stale* copy. I/O means the opportunity exists for other devices to cause copies to be inconsistent or for other devices to read the stale copies. Figure 8.43 illustrates the problem. This is generally referred to as the *cache-coherency* problem.

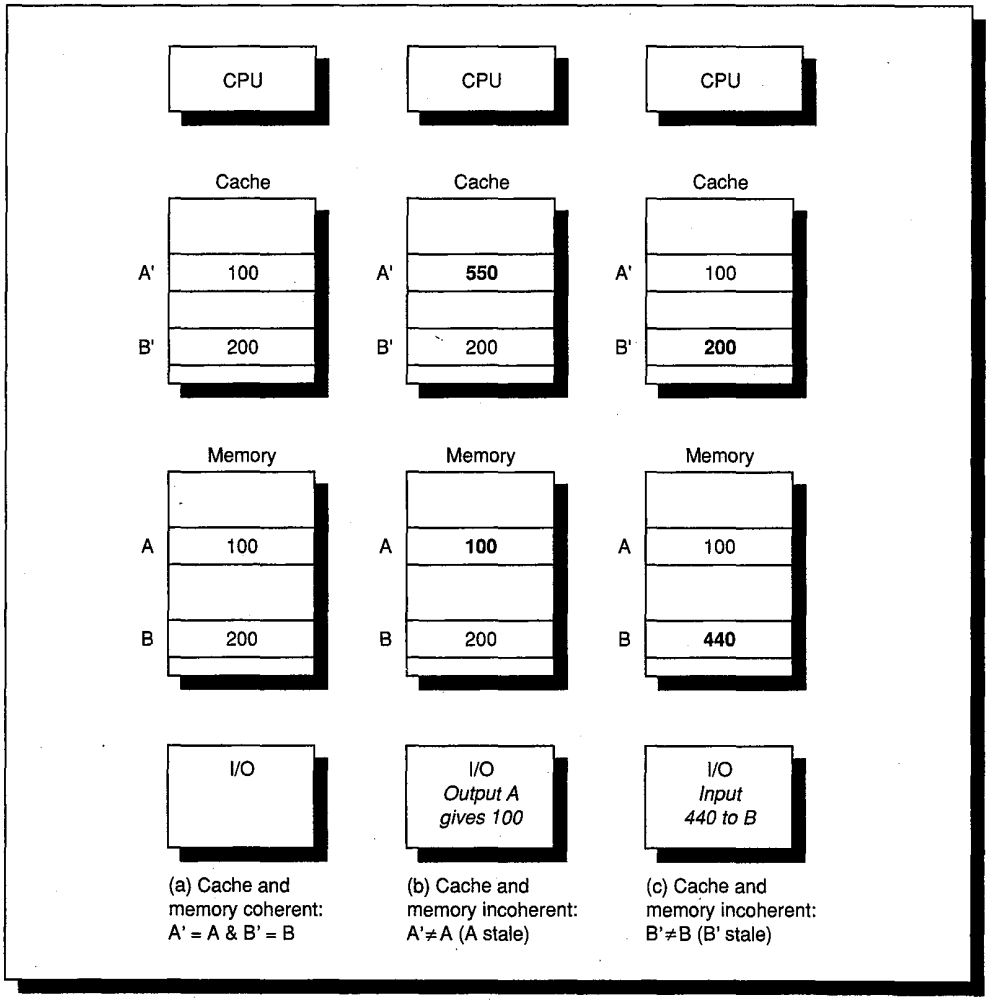

FIGURE 8.43 The cache-coherency problem. A' and B' refer to the cached copies of A and B in memory. (a) shows cache and main memory in a coherent state. In (b) we assume a write-back cache when the CPU writes 550 into A. Now A' has the value but the value in memory has the old, stale value of 100. If an output used the value of A from memory, it would get the stale data. In (c) the I/O system inputs 440 into the memory copy of B, so now B' in the cache has the old, stale data.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 498

The question is this: Where does the I/O occur in the computer—between the I/O device and the cache or between the I/O device and main memory? If input puts data into the cache and output reads data from the cache, both I/O and the CPU see the same data, and the problem is solved. The difficulty in this approach is that it interferes with the CPU. I/O competing with the CPU for cache access will cause the CPU to stall for I/O. Input will also interfere with the cache by displacing some information with the new data that is unlikely to be accessed by the CPU soon. For example, on a page fault the CPU may need to access a few words in a page, but a program is not likely to access every word of the page if it were loaded into the cache.

The goal for the I/O system in a computer with a cache is to prevent the stale-data problem while interfering with the CPU as little as possible. Many systems, therefore, prefer that I/O occur directly to main memory, acting as an I/O buffer. If a write-through cache is used, then memory has an up-to-date copy of the information, and there is no stale-data issue for output. (This is the reason many machines use write through.) Input requires some extra work. The software solution is to guarantee that no blocks of the I/O buffer designated for input are in the cache. In one approach, a buffer page is marked as noncacheable; the operating system always inputs to such a page. In another approach, the operating system flushes the buffer addresses from the cache after the input occurs. A hardware solution is to check the I/O addresses on input to see if they are in the cache. If so, the cache entries are invalidated to avoid stale data. All these approaches can also be used for output with write-back caches. More about this is found in the next chapter.

Reducing Bus Traffic—Multiprocessor Cache Coherency

The cache-coherency problem applies to multiprocessors as well as I/O. Unlike I/O, where multiple data copies is a rare event—one to be avoided whenever possible—a program running on multiple processors will want to have copies of the same data in several caches. Performance of a multiprocessor program depends on the performance of the system when sharing data. The protocols to maintain coherency for multiple processors are called *cache-coherency protocols*. There are two classes of protocols followed to maintain cache coherency:

- *Directory based*—The information about one block of physical memory is kept in just one location.
- *Snooping*—Every cache that has a copy of the data from a block of physical memory also has a copy of the information about it. These caches are usually on a shared-memory bus, and all cache controllers monitor or *snoop* on the bus to determine whether or not they have a copy of the shared block.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 499

In directory-based protocols there is logically a single directory that keeps the state of every block in main memory. Information in the directory can include which caches have copies of the block, whether it is dirty, and so on. Of course directory entries can be distributed so that different requests can go to different memories, thereby reducing contention. However, they retain the characteristic that the sharing status of a block is always in a single known location.

Snooping protocols became popular with multiprocessors using microprocessors and caches on a shared memory because they can use a preexisting physical connection: the bus to memory. Snooping has an edge over directory protocols in that the coherency information is proportional to the number of blocks in a cache rather than the number of blocks in main memory. Directories, on the other hand, do not require a single bus going to all caches and, hence, may scale to more processors.

The coherency problem is for a processor to have exclusive access to write an object and to have the most recent copy when reading an object. Thus, both directory-based and snooping protocols must locate all the caches that share the object to be written. The consequence of a write to shared data is either to invalidate all other copies or to broadcast the write to the shared copies. Because of write-back caches, coherency protocols must also help read misses determine who has the most up-to-date value.

For the remainder of this section we concentrate on snooping caches; the same ideas apply to directory-based caches except the state of the caches is tracked differently, and caches are involved only if the directory says they have a copy of a block whose status must change.

Sharing information is added to the status bits already in a cache block for snooping protocols, and that information is used in monitoring bus activities. On a read miss all caches check to see if they have a copy of the requested block and take the appropriate action, such as supplying the data to the cache that missed. Similarly, on a write all caches check to see if they have a copy and then act, perhaps invalidating their copy or changing their copy to the new value.

Since every bus transaction checks cache-address tags, one might assume that it interferes with the CPU. It would, were it not for duplicating the address-tag portion of the cache (not the whole cache) to get an extra read port for snooping. This way, snooping interferes with the CPU's access to the cache only when there is a coherency problem (although on a miss with snooping the CPU must arbitrate with the bus to change the snoop tags as well as the normal tags). When a coherency operation occurs in the cache the CPU will likely stall, since the cache is unavailable. In multilevel caches, if the coherency check can be limited to the lower cache because of multilevel inclusion, duplicating the address tags will probably not be necessary.

Snooping protocols are of two types, depending on what happens on a write:

- *Write invalidate*—The writing processor causes all copies in other caches to be invalidated before changing its local copy; it is then free to update the data until another processor asks for it. The writing processor issues an invalida-

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 500

tion signal over the bus, and all caches check to see if they have a copy; if so, they must invalidate the block containing the word. Thus, this scheme allows multiple readers but only a single writer.

- *Write broadcast*—Rather than invalidate every block that is shared, the writing processor broadcasts the new data over the bus; all copies are then updated with the new value. This scheme continuously broadcasts writes to shared data while write invalidate deletes all other copies so that there is only one local copy for subsequent writes. Write-broadcast protocols usually allow blocks to be tagged as shared (broadcast) or private (local). One way to think of this protocol is it acts like a write-through cache for shared data (broadcasting to other caches) and a write-back cache for private data (the modified data leaves the cache only on a miss).

Most cache-based multiprocessors use write back caches because it reduces bus traffic and thereby allows more processors on a single bus. Write-back caches use either invalidation or broadcast, and numerous variations exist for both alternatives (see the next section). So far, there is no consensus on which is the superior scheme. Some programs have less coherency overhead with write invalidate, and some with write broadcast. A later section shows how synchronization can be implemented in coherency-based multiprocessors; the accesses for synchronization seem to favor write broadcast.

One early insight has been that block size plays an important role in cache coherency. Take, for example, the case of snooping on a second-level cache with a block size of eight words, and a single word is alternatively written and read by two processors. Whether write invalidation or write broadcast is used, the protocol that only broadcasts or sends a word has an advantage over a scheme that transfers the full block. Another concern of large blocks is called *false sharing*: two different shared variables are located in the same cache block, causing the block to be exchanged between processors even though the processors are accessing different variables. Compiler research is working to reduce cache miss rates by allocating data with high processor locality to the same blocks. Success in this field could increase the desirability of large blocks for multiprocessors.

Measurements to date indicate that shared data has lower spatial and temporal locality than observed for other types of data, independent of the coherency policy.

An Example Protocol

To illustrate the complexities of a cache-coherency protocol, Figure 8.44 (page 470) shows a finite-state transition diagram for a write-invalidation protocol based on write- back policy. The three states of the protocol are duplicated to represent transitions based on CPU actions, as opposed to transitions based on bus operations. This is done only for purposes of this figure; there is only one finite-state machine per cache, with stimuli coming either from the attached CPU or from the bus.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 501

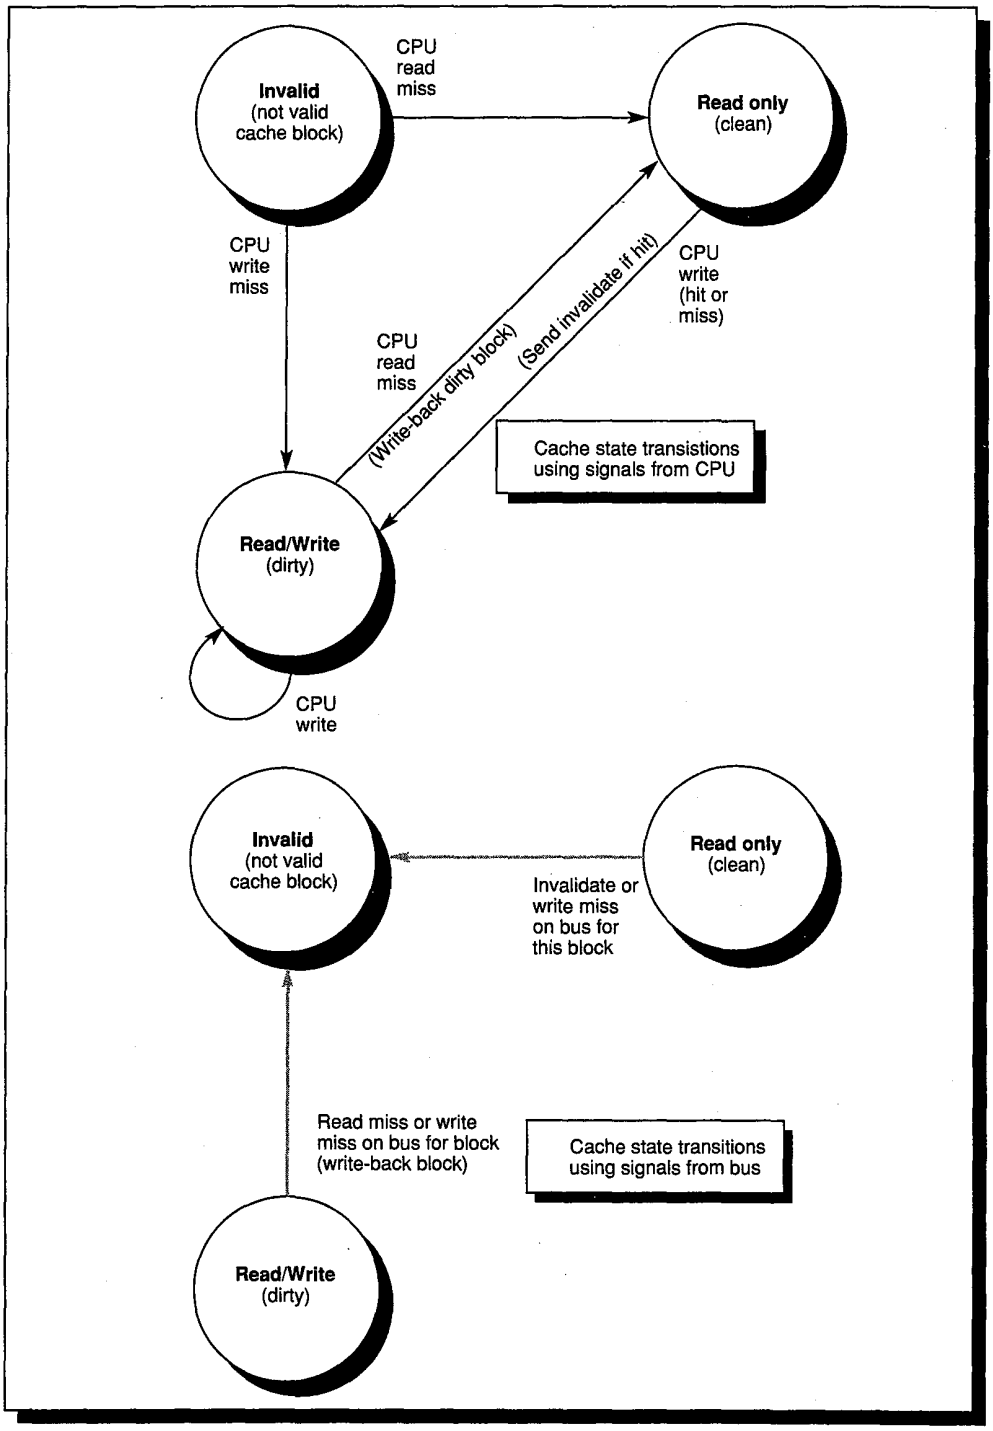

FIGURE 8.44 A write-invalidate, cache-coherency protocol. The upper part of the diagram shows state transitions based on actions of the CPU associated with this cache; the lower part shows transitions based on operations on the bus. There is only one state machine in a cache, although there are two represented here to clarify when a transition occurs. The black arrows and states would be in a normal cache, with the gray arrows added to get cache coherency. In contrast to what is shown here, some protocols call writes to clean data a "write miss," so that there is no separate signal for invalidation.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 502

Transitions happen on read misses, write misses, or write hits; read hits do not change cache state. When the CPU has a read miss, it will change the state of that block to Read only and write back the old block if it was in the Read/Write state (dirty). All the caches snoop on the read miss to see if this block is in their cache. If one has a copy and it is in the Read/Write state, then the block is written to memory and that block is changed to the invalid state. (An optimization not shown in the figure would be to change the state of that block to Read only.) When a CPU writes into a block, that block goes to the Read/Write state. If the write was a hit, an invalidate signal goes out over the bus. Because caches monitor the bus, all check to see if they have a copy of that block; if they do, they invalidate it. If the write was a miss, all caches with copies go to the invalid state.

As you might imagine, there are many variations on cache coherency that are much more complicated than this simple model. The variations include whether or not the other caches try to supply the block if they have a copy, whether or not the block must be invalidated on a read miss, as well as write invalidate versus write broadcast as discussed above. Figure 8.45 summarizes several snooping cache-coherency protocols.

Name	Category	Memory-write policy	Unique feature
Write Once	Write invalidate	Write back after first write	
Synapse N+1	Write invalidate	Write back	Explicit memory ownership
Berkeley	Write invalidate	Write back	Owned shared state
Illinois	Write invalidate	Write back	Clean private state; can supply data from any cache with a clean copy
Firefly	Write broadcast	Write back for private, Write through for shared	Memory updated on broadcast
Dragon	Write broadcast	Write back for private, Write through for shared	Memory not updated on broadcast

FIGURE 8.45 Six snooping protocols summarized. Archibald and Baer [1986] use these names to describe the six protocols, and Eggers [1989] summarizes the similarities and differences as shown above. Figure 8.44 (page 470) is simpler than any of these protocols.

Synchronization Using Coherency

One of the major requirements of a shared-memory multiprocessor is being able to coordinate processes that are working on a common task. Typically, a programmer will use *lock variables* to synchronize the processes.

The difficulty for the architect of a multiprocessor is to provide a mechanism to decide which processor gets the lock and to provide the operation that locks a variable. Arbitration is easy for shared-bus multiprocessors, since the bus is the only path to memory: The processor that gets the bus locks out all other processors from memory. If the CPU and bus provide an atomic swap operation, programmers can create locks with the proper semantics. The adjective *atomic* is

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 503

key, for it means that a processor can both read a location **and** set it to the locked value in the same bus operation, preventing any other processor from reading or writing memory.

Figure 8.46 shows a typical procedure for locking a variable using an atomic swap instruction. Assume that 0 means unlocked and 1 means locked. A processor first reads the lock variable to test its state. A processor keeps reading and testing until the value indicates that the lock is unlocked. The processor then races against all other processes that were similarly "spin waiting" to see who

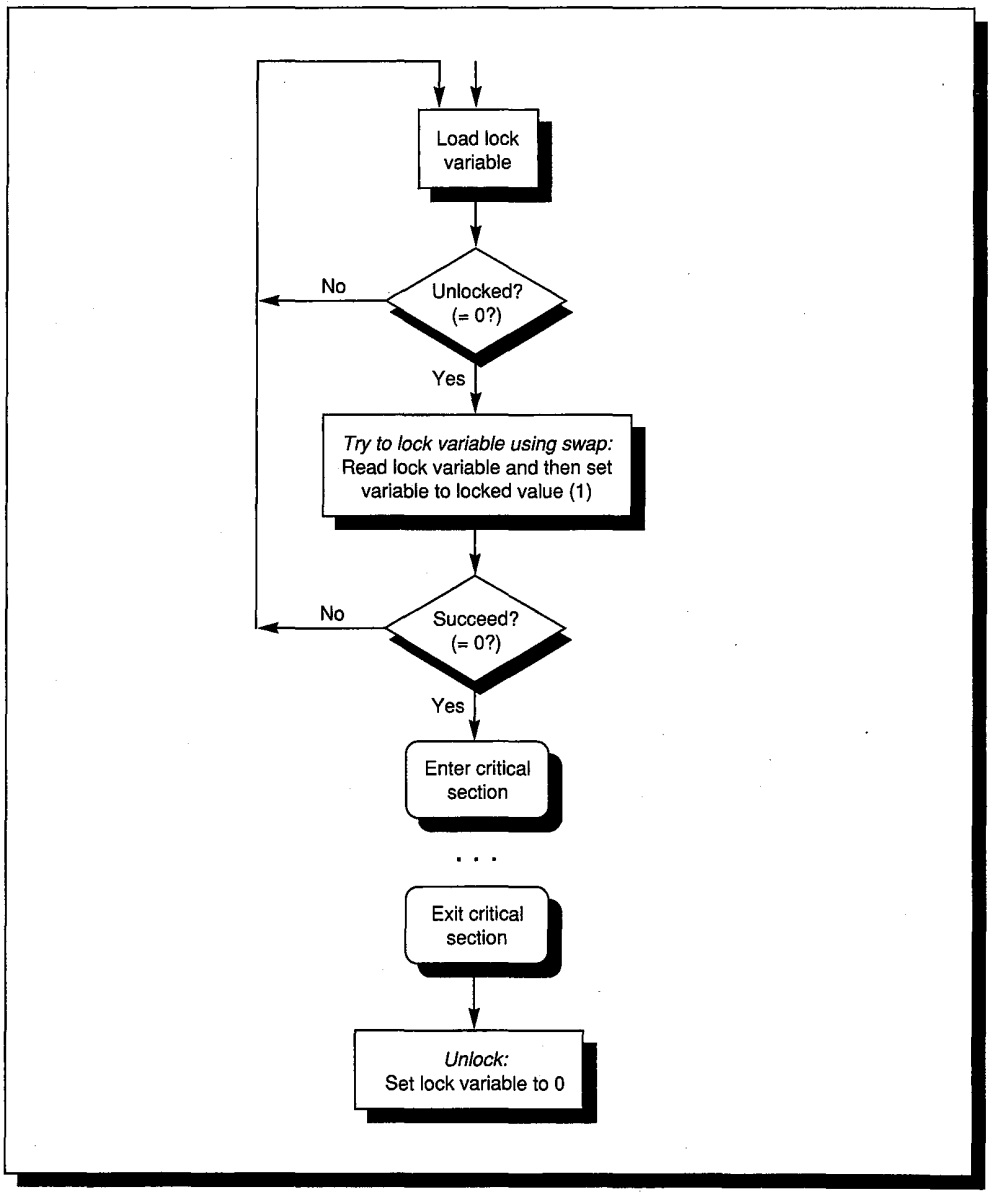

FIGURE 8.46 Steps to acquire a lock to synchronize processes and then to release the lock on exit from the key section of code.

Step	Processor P0	Processor P1	Processor P2	Bus activity
1	Has lock	Spins, testing if lock = 0	Spins, testing if lock = 0	None
2	Set lock to 0 and 0 sent over bus			Write invalidate of lock variable from P0
3		Cache miss	Cache miss	Bus decides to service P2 cache miss
4		(Waits while bus busy)	Lock = 0	Cache miss for P2 satisfied
5		Lock = 0	Swap: read lock and set to 1	Cache miss for P1 satisfied
6		Swap: read lock and set to 1	Value from swap = 0 and 1 sent over bus	Write invalidate of lock variable from P2
7		Value from swap = 1 and 1 sent over bus	Enter critical section	Write invalidate of lock variable from P1
8		Spins, testing if lock = 0		None

FIGURE 8.47 Cache-coherency steps and bus traffic for three processors, P0, P1, and P2. This figure assumes write-invalidate coherency. P0 starts with the lock (step 1). P0 exits and unlocks the lock (step 2). P1 and P2 race to see which reads the unlocked value during the swap (steps 3-5). P2 wins and enters the critical section (steps 6 and 7), while P1 spins and waits (steps 7 and 8).

can lock the variable first. All processes use a swap instruction that reads the old value and stores a 1 into the lock variable. The single winner will see the 0, and the losers will see a 1 that was placed there by the winner. (The losers will continue to set the variable to the locked value, but that doesn't matter.) The winning processor executes the code after the lock and then stores a 0 into the lock variable when it exits, starting the race all over again. Testing the old value and then setting to a new value is why the atomic swap instruction is called *test and set* in some instruction sets.

Let's examine how the "spin lock" scheme of Figure 8.46 works with bus-based cache coherency. One advantage of this algorithm is that it allows processors to spin wait on a local copy of the lock in their caches. This reduces the amount of bus traffic versus lock algorithms that loop trying to perform a test and set. (Figure 8.47 shows the bus and cache operations for multiple processes trying to lock a variable.) Once the processor with the lock stores a 0 into the lock, all other caches see that store and invalidate their copy of the lock variable. They then get the new value for the lock of 0. (With write-broadcast cache coherency as on page 469, the caches would update their copy rather than first invalidate and then load from memory.) This new value starts the race to see who can set the lock first. The winner gets the bus and stores a 1 into the lock; the other caches replace their copy of the lock variable containing 0 with a 1. They read that the variable is already locked and must return to testing and spinning. This scheme has difficulty scaling up to many processors because of the communication traffic generated when the lock is released.

Models of Memory Consistency

When we introduce cache coherency to maintain the consistency of multiple copies of an object, we raise a new question: How consistent must the values seen by two processors be kept? The problem is best understood with an example: Here are two code segments from processes P1 and P2 shown side by side:

```
P1:    A = 0;                    P2:    B = 0;
         .....                           .....
       A = 1;                          B = 1;
L1:    if (B == 0)   ...   L2:    if (A == 0)  ...
```

Assume the processes are running on different processors, and that locations A and B are originally cached by both processors with the initial value of 0. If memory is always consistent, it will be impossible for **both** if statements (labeled L1 and L2) to evaluate their conditions as true (either A=1 or B=1). But suppose write invalidates have a delay, and the processor is allowed to continue during this delay, then it is possible that both P1 and P2 have not seen the invalidations for B and A (respectively) **before** they attempt to read the values. The question that is raised by this example is: How consistent a picture of memory must different processors see?

One approach, called sequential consistency, requires that the result of any execution is the same as if the accesses of each processor were kept in order and the accesses among different processors were arbitrarily interleaved. In this case, the apparent anomaly in the above example cannot occur. Implementing sequential consistency usually requires a processor to delay any memory access until all the invalidations caused by all previous writes are completed. Although this model presents a simple programming paradigm, it reduces potential performance, especially in a machine with a large number of processors, or long interconnect delays.

Alternative models provide a weaker model of memory consistency. For example, the programmer may be required to use synchronization instructions to order memory accesses to the same variable. Now, instead of delaying all accesses until invalidations complete, only synchronization accesses need to be delayed.

Whether programmers expect sequential consistency or some weaker form of consistency is still an open issue in 1990. The example above would work "correctly" with sequential consistency, but not with a weaker model. For weak consistency to produce the same results as sequential consistency, the program would have to be modified to include synchronization operations that order the accesses to variables A and B. It is natural to expect synchronization if you want processes to see the latest data independent of execution rates. Some machines choose to implement sequential consistency as the programming model, while others opt for a weaker consistency. In the future, as attempts are made to build larger multiprocessors, the issue of memory consistency will become increasingly performance critical.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 506

8.9 Putting It All Together: The VAX-11/780 Memory Hierarchy

The challenge for the memory-hierarchy designer is in choosing parameters that work well together, not in inventing new techniques or simulating a cache in a well-understood configuration. A full example using the VAX-11/780 memory hierarchy is presented here in detail to illuminate the interactions. Although VAX-11/780 is not a very recent machine, measurements and design documentation are available on all aspects of its memory hierarchy. Figure 8.48 gives the overall picture.

Let's start with an instruction fetch just after a branch, when the instruction prefetch buffer is empty. The virtual address in the PC is first sent to the TLB. The most significant bit and the lower five bits of the page-frame address index an entry in each bank of the TLB. Including the most-significant bit, used to distinguish system space from process space, guarantees that half of each bank contains system translations and half contains process translations. The addresses in the tags are compared to see if the entry is a match to the page address requested by the TLB. If the valid bit of the entry is not set then there is no match no matter what the tag comparison says, and a miss is indicated.

If there is a match, the physical address is formed by concatenating the physical page-frame address of the TLB page-table entry with the page-offset portion of the address. To save time, the portion of the TLB containing the PTE is read at the same time as the tags, and a 2:1 multiplexer controlled by the tag-matching logic picks the proper PTE. While the address is being formed, the protection bits of the PTE are checked. Since this is an instruction fetch, there is no problem as long as the page can be read by a process at this level. If there are no protection violations, this physical address is sent to the cache.

At the same time the physical address is sent to the cache, two registers in the CPU instruction-prefetch buffer get the new values. The *virtual-instruction-buffer address* register (VIBA) is given the virtual page frame of the PC, and the *physical-instruction-buffer address* register (PIBA) is given the corresponding physical address. This trick, which was originally used in the first machine with virtual memory, avoids the instruction-prefetch buffer's accessing the TLB as long as the instructions are from the same page. The PIBA is actually given the PC address plus 4, so that it can begin prefetching the next instruction. It continues trying to prefetch ahead of the PC until a jump (a frequent occurrence in the VAX) or until the PIBA tries to cross a page boundary; in either case the VIBA and PIBA are no longer used for translating instruction addresses.

Meanwhile, the cache has just received the physical address of the instruction. With 8-byte blocks, a two-way–set-associative cache, and 512 blocks per set, nine bits of the address are needed to index both banks simultaneously. The partial addresses in the tags are compared with the corresponding bits of the physical PC address to see if there is a match. Of course, there are valid bits in each tag that must be turned on, or there can be no match.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 507

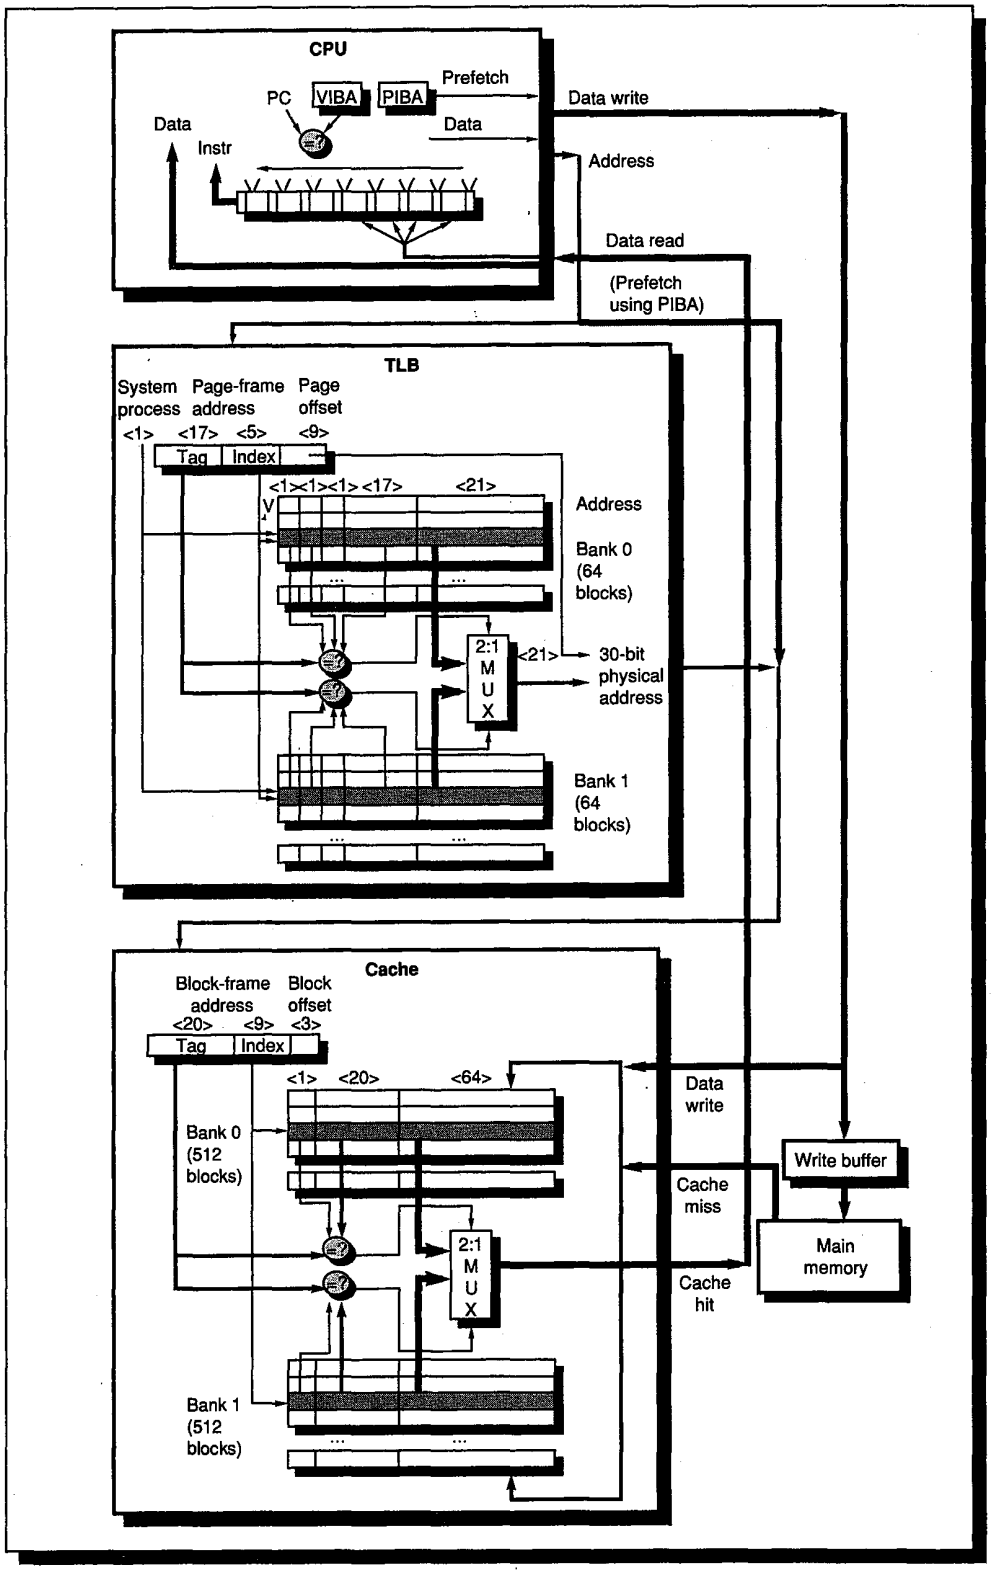

FIGURE 8.48 The overall picture of the VAX-11/780 memory hierarchy. Individual components can be seen in greater detail in Figures 8.11 (page 415), 8.29 (page 444), and 8.31 (page 450).

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 508

If there is a match, the lower bits of the physical PC address select the word from the cache block to be sent to the instruction-prefetch unit. Once again, reading data and tags together obviates any additional time delay.

When the word arrives at the prefetch unit, it is placed in the high-order four bytes of the buffer, and those bytes are marked valid. The PIBA immediately begins accessing the cache with the PC address plus 4 to prefetch the next word. As mentioned above, as long as the page-frame address in the PC matches the VIBA, the PIBA bypasses the TLB and goes directly to the cache.

Let's assume this instruction writes a register into memory. The first step will be to send the effective memory address to the TLB for translation. Since this is a write, the modify bit of the matching PTE must also be turned on; this results in a microcode-level trap of the instruction storing the register if the modify bit isn't set already, taking another clock cycle to write the new value in the TLB. The physical address is then sent to the cache. We then go through the same process as before (excluding the read), except that this time it takes an extra clock cycle to modify the portion of the block selected by the write and to write it back into the cache.

In a write-through cache the data must be written to main memory. To avoid the seven-cycle delay of main memory on every write, the VAX-11/780 uses a one-word write buffer. If the buffer is empty, the word is written and the CPU is given the signal to continue. If it is full, the CPU stalls until the buffer is empty.

How well does the 780 work? The bottom line in this evaluation is the percentage of time lost while the CPU is waiting for the memory hierarchy. In one timesharing workload the average number of clock cycles per 780 instruction is 10.6 clock cycles. The breakdown by category is

Compute: 7.3 clock cycles

Read: 0.8 clock cycles

Read stall: 1.0 clock cycles

Write: 0.4 clock cycles

Write stall: 0.4 clock cycles

Instruction-prefetch–buffer stall: 0.7 clock cycles

About 20% of the time the VAX-11/780 stalls while waiting for memory. When the base CPI is 8.5 (compute + read + write), 2.1 clock cycles for the memory hierarchy (read stall + write stall + prefetch stall) may be satisfactory, but it would devastate the performance of a machine with a CPI of 1 to 2.

Let's analyze each unit of the 780 memory hierarchy. An instruction-prefetch–buffer stall means that the buffer is empty, waiting for the cache to supply instructions because of a cache miss, a branch, too many data accesses (they have priority), not enough bytes to decode the instruction, or some combination of the above. The PIBA loadings due to branches versus page crossings vary with the benchmark, but branching is the cause 64% to 91% of the time

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 509

(median = 76%). The prefetch unit references the cache 2.2 times on average per VAX instruction. The average instruction size is 3.8 bytes, making the effective size of the average prefetch just 1.7 bytes.

Example

Figure 3.33 in Chapter 3 (page 123) shows that the VAX executes many fewer bytes of instructions than DLX. This ignores the instruction-prefetch buffer. How much should we increase the instruction bytes fetched from the cache to include the effect of prefetching?

Answer

We can answer this in a couple of ways. Every prefetch access to the cache actually returns 4 bytes, and the average VAX instruction size is 3.8 bytes; the increase could therefore be

$$\frac{2.2 * 4}{3.8} = 2.32$$

since the prefetch unit references the cache 2.2 times per instruction. This suggests that the bytes fetched from the cache should be increased by 132%. Because the same code may be fetched multiple times by the prefetcher, however, the bandwidth between the cache and memory may not change since the prefetcher cannot cause cache misses.

The question can also be answered in terms of the number of bytes discarded because of a taken branch. About 25% of instructions change the PC on the VAX, and there could be from zero to eight bytes in the prefetch unit when a branch is taken. Assuming an optimistic two bytes, we get a 13% increase:

$$\frac{3.8 + (25\% * 2)}{3.8} = 1.13$$

Assuming six bytes, we get a 39% increase:

$$\frac{3.8 + (25\% * 6)}{3.8} = 1.39$$

While the variable size of VAX instructions does improve the bytes fetched in comparison to DLX, a fairer evaluation of the VAX would increase the bytes fetched from the cache by at least 13% to 39%.

With the instruction-prefetch buffer performing many translations via the PIBA and VIBA, how should TLB misses be measured? The TLB instruction and data-stream miss rates provide one definition:

$$\text{TLB instruction-stream miss rate} = \frac{\text{Misses caused by IB}}{\text{Reloadings of PIBA}}$$

$$\text{TLB data-stream miss rate} = \frac{\text{Misses}}{\text{Requests for 32-bit words of data}}$$

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 510

The data-stream definition means references to data objects larger than four bytes count as multiple accesses, as do accesses to unaligned data. Figure 8.49 shows the TLB miss rates.

TLB miss rates	Instruction stream	Data stream	Total
Process	0.7 %	0.6 %	0.7 %
System	15.4 %	5.4 %	7.2 %
Total	3.5 %	1.6 %	1.9 %

FIGURE 8.49 Miss rates for the VAX-11/780 TLB, ignoring the impact of instructions not translated by the TLB. This data was measured on a different timesharing workload than earlier VAX measurements [Clark and Emer 1985].

Overall references to the TLB after filtering by the PIBA are divided into 20% user instruction stream, 62% user data stream, 3% system instruction stream, and 15% system data stream. To account for the filtering of addresses by the PIBA optimization, TLB misses can also be counted as a rate per instruction executed, as in Figure 8.50.

TLB misses per 100 instructions	Instruction stream	Data stream	Total
Process	0.18	0.50	0.68
System	0.62	1.03	1.65
Total	0.80	1.53	2.33

FIGURE 8.50 Misses per hundred instructions for the VAX-11/780 TLB. Unlike Figure 8.49, this overall TLB evaluation accounts for the effect of the PIBA.

The VAX TLB spends on average 21.6 clock cycles on a miss (including 3.5 clock cycles for cache misses for some page-table entries), adding a total of 0.7 clock cycles per instruction for TLB misses to the average instruction. Thus, about a third of the memory-system stalls are due to TLB misses.

The same study by Emer and Clark [1984] showed a significant variation on cache miss rates:

- Data-stream, cache miss rates varied over the day from 12% to 25%, with a mean of 17%.
- Instruction-buffer–stream, cache miss rates varied from 4% to 13%, with a mean of 8%.
- The distribution of accesses to the cache from the CPU was instruction-prefetch–buffer–stream reads, 68%, data-stream reads, 20%, and data-stream writes, 12%. Calculated per instruction, there are about 2.2 references from the instruction-prefetch buffer, 0.8 data reads per instruction, and 0.4 data writes per instruction.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 511

Example

According to the VAX-11/780 Architecture Handbook, for the workload measured in 1978 the TLB miss rate was about 3%. What do the measurements say for the timesharing workload measured in 1984?

Answer

Assuming just one memory reference to get the average VAX instruction of 3.8 bytes, the miss rate is 1%:

$$\frac{\frac{2.3 \text{ TLB misses}}{100 \text{ instructions}}}{\frac{1+0.8+0.4 \text{ references}}{\text{Instruction}}} = \frac{2.3}{100*2.2} = 0.01$$

Including the VIBA-PIBA, Figure 8.49 on page 479 shows a 1.9% miss rate.

Example

According to the VAX-11/780 Architecture Handbook, for the workload measured in 1978 the cache miss rate was about 5%. What do the measurements say for the timesharing workload measured in 1984?

Answer

The cache miss rate varies. The mean miss rate is

$$68\%*8\% + 20\%*17\% + 12\%*17\% = 11\%$$

In the best case, the answer is

$$68\%*4\% + 20\%*12\% + 12\%*12\% = 7\%$$

In the worst case,

$$68\%*13\% + 20\%*25\% + 12\%*25\% = 17\%$$

8.10 Fallacies and Pitfalls

As the most naturally quantitative of the computer architecture disciplines, memory hierarchy would seem to be less vulnerable to fallacies and pitfalls. Yet the authors were limited here not by lack of warnings, but by space.

Pitfall: Too small an address space.

Just five years after DEC and Carnegie-Mellon University collaborated to design the new PDP-11 computer family, it was apparent that their creation had a fatal flaw. An architecture announced by IBM six years **before** the PDP-11 is still thriving, with minor modifications, 25 years later. And the DEC VAX, criticized for including unnecessary functions, has sold 100,000 units since the PDP-11 went out of production. Why?

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 512

The fatal flaw of the PDP-11 was the size of its addresses as compared to the IBM 360 and the VAX. Address size limits the program length, since the size of a program and the amount of data needed by the program must be less than $2^{\text{address size}}$. The reason the address size is so hard to change is that it determines the minimum width of anything that can contain an address: PC, register, memory word, and effective-address arithmetic. If there is no plan to expand the address from the start, then the chances of successfully changing address size are so slim that it normally means the end of that computer family. Bell and Strecker [1976] put it like this:

There is only one mistake that can be made in computer design that is difficult to recover from—not having enough address bits for memory addressing and memory management. The PDP-11 followed the unbroken tradition of nearly every known computer. [p. 2]

A partial list of successful machines that eventually starved to death for lack of address bits includes the PDP-8, PDP-10, PDP-11, Intel 8080, Intel 8086, Intel 80186, Intel 80286, AMI 6502, Zilog Z80, CRAY-1, and CRAY X-MP.

Fallacy: Given the hardware resources, the computer designer who selects a set-associative cache over a direct-mapped cache of the same size will get a faster computer.

The question here is whether the extra logic of the set-associative cache affects the hit time, and therefore possibly the CPU clock rate. (See Figure 8.11.) If it does affect hit time, then the question is whether the advantage in lower miss rate offsets the slower hit time. In the mid-1980s many recognized this danger and selected direct-mapped placement; for example, the MIPS M/500, Sun 3/260, and VAX 8800. Hill [1988] makes an eloquent case for direct-mapped caches, including lower costs, faster hit times, and therefore smaller average access times for large, direct-mapped caches. Direct-mapped caches also allow the data read to be sent to the CPU and used even before hit/miss is determined, particularly useful with a pipelined CPU. Hill found about a 10% difference in hit times for TTL or ECL board-level caches and 2% difference for custom CMOS caches, with an absolute change in the miss rates of less than 1% for large caches. Since a direct-mapped cache hit can be accessed faster and hit time typically sets the clock cycle time of the processor, a CPU with a direct-mapped cache can be as fast as or faster than a CPU with a two-way–set-associative cache of the same size. Przybylski, Horowitz, and Hennessy [1988] show several examples of such tradeoffs.

Fallacy: A memory system can be designed using traces from a different architecture.

Figure 8.51 (page 482) shows instruction and data cache miss rates for the same programs on two different architectures. This data is from the first portion of execution of Spice on DLX and the VAX. The shift from data accesses in the

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 513

VAX to instruction accesses on DLX seen in Figure 3.33 (page 123) of Chapter 3 is reflected here: 61% of the VAX references and 52% of the misses are to data. Note that while DLX has only three-quarters of the absolute number of data misses, its data miss **rate** is three times higher.

	VAX	DLX
Instruction references	576,169	918,537
Instruction misses	2,033	3,188
Instruction miss rate	0.4%	0.3%
Data references	923,831	264,453
Data misses	2,200	1,595
Data miss rate	0.2%	0.6%
Total references	1,500,000	1,182,990
Percentage of instructions of total references	38%	78%
Total misses	4,233	4,782
Percentage of instruction misses of total misses	48%	67%
Average miss rate	0.3%	0.4%

FIGURE 8.51 Miss rates for VAX and DLX for an initial phase of Spice. The simulation assumes separate instruction and data caches. Each cache is direct mapped, uses 16-byte blocks, and contains 64 KB. Both use write through with write allocate. (Note that unlike Chapter 2, this data was collected using the F77 compiler and was for a portion of the Spice program).

Pitfall: Basing the size of the write buffer on the speed of memory and the average mix of writes.

This seems like a reasonable approach:

$$\text{Write-buffer size} = \frac{\text{Memory references}}{\text{Clock cycle}} * \text{Write percentage} * \text{Clock cycles to write memory}$$

If there is one memory reference per clock cycle, 10% of the memory references are writes, and writing a word of memory takes 10 cycles, then a one-word buffer is added (1*10%*10=1). Calculating for the VAX-11/780 using data from the last section,

$$\frac{3.4 \text{ memory references}}{10.6 \text{ clock cycles}} * \frac{0.4 \text{ writes}}{3.4 \text{ memory references}} * \frac{6 \text{ clock cycles}}{\text{Write}} = 0.22$$

Thus, a one-word buffer seems sufficient.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 514

The pitfall is that when writes come close together, the CPU must stall until the prior write is completed. The single-word write buffer of the VAX-11/780 is the major reason for its write stalling (about 20% of all stalls). The proper question to ask is how large a buffer is needed to keep CPU write stalls to a small amount. The impact of write-buffer size can be established by simulation or estimated with a queuing model.

Pitfall: Extending an address space by adding segments on top of a flat address space.

During the 1970s, many programs grew to the point they couldn't address all of the code and data with just a 16-bit address. Machines were then revised to offer 32-bit addresses, either through a flat 32-bit address space or by adding 16 bits of segment to the existing 16-bit address. From the point of view of marketing, adding segments solves the addressing problem. Unfortunately, there is trouble any time a programming language wants an address that is larger than one segment, such as indices for large arrays, unrestricted pointers, or reference parameters. Moreover, adding segments can turn every address into two words—one for the segment number and one for the segment offset—causing problems in the use of addresses in registers. In the 1990s, 32-bit addresses will be exhausted, and it will be interesting to see if history will repeat itself on the consequences of going to larger flat addresses versus adding segments.

Fallacy: Caches are as fast as registers.

This fallacy is important, because if caches were as fast as registers, there would be no need for registers. Without registers there would be no need for a register allocator, and so compilers could be simpler. The fallacy is difficult to prove quantitatively, yet example after example can be cited. Lampson [1982] summarized this experience:

A register bank is faster than a cache, both because it is smaller, and because the address mechanism is much simpler. Designers of high performance machines have typically found it is possible to read one register and write another in a single cycle, while two cycles [latency] are needed for a cache access. ... Also, since there are not too many registers it is feasible to duplicate or triplicate them, so that several registers can be read out simultaneously. [p. 74]

As mentioned in Chapter 3, the short addresses of registers allow more compact instruction encoding. It seems to the authors that the deterministic access of multiported register banks will always offer lower latency or higher bandwidth, or both, when compared to the nondeterministic access of caches.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 515

8.11 Concluding Remarks

The difficulty of building a memory system to keep pace with faster CPUs is underscored by the fact that the raw material for main memory is the same as that found in the cheapest computer. It is the principle of locality that saves us here—its soundness is demonstrated at all levels of the memory hierarchy in current computers, from disks to instruction buffers.

	Register windows	Instruction-prefetch buffer	TLB	First-level cache	Second-level cache	Virtual memory
Block size	64 bytes	1 byte	4 – 8 (1 PTE)	4 – 128 bytes	32 – 256 bytes	512 – 8192 bytes
Hit time	1 clock cycle	1 clock cycle	1 clock cycle	1 – 4 clock cycles	4 – 10 clock cycles	1 – 10 clock cycles
Miss penalty	32 – 64 clock cycles	2 – 6 clock cycles	10 – 30 clock cycles	8 – 32 clock cycles	30 – 80 clock cycles	100,000 – 600,000 clock cycles
Miss rate (local)	1% – 3%	10% – 25%	0.1% – 2%	1% – 20%	15% – 30%	0.00001% – 0.001%
Size	512 bytes	6 – 12 bytes	32 – 8192 (8 – 1024 PTEs)	1 KB – 256 KB	256 KB – 4 MB	4 MB – 2048 MB
Backing store	First-level cache	First-level cache	First-level cache	Second-level cache	Static-column DRAM	Disks
Q1: block placement	Circular buffer	N.A. (Queue)	Set associative	Direct mapped	Set associative	Fully associative
Q2: block identification	2 registers: high and low	Valid bits + 1 register	Tag/ block	Tag/ block	Tag/ block	Table
Q3: block replacement	First in– first out	N.A. (Queue)	Random	N.A. (Direct mapped)	Random	LRU
Q4: write strategy	Write back	Flush on write to instruction buffer (if possible)	Flush on write to page table	Write through or write back	Write through or write back	Write back

FIGURE 8.52 Summary of the memory-hierarchy examples in this chapter.

Misses in every level can be categorized by three causes—compulsory, capacity, and conflict—and different techniques work for each case. Figure 8.52 summarizes the attributes of the memory-hierarchy examples described in this chapter.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 516

There tends to be a knee in the curve of memory-hierarchy cost/performance: Above that knee is wasted performance and below that knee is wasted hardware. Architects find that knee by simulation and quantitative analysis.

8.12 Historical Perspective and References

While the pioneers of computing knew of the need for a memory hierarchy and coined the term, the automatic management of two levels was first proposed by Kilburn, et al. [1962] and demonstrated with the Atlas computer at the University of Manchester. This was the year **before** the IBM 360 was announced. While IBM planned for its introduction with the next generation (System/370), the operating system wasn't up to the challenge in 1970. Virtual memory was announced for the 370 family in 1972, and it was for this machine that the term "translation-lookaside buffer" was coined (see Case and Padegs [1978]). The only computers today without virtual memory are a few supercomputers and personal computers.

Both the Atlas and the IBM 360 provided protection on pages, and over time machines evolved more elaborate mechanisms. The most elaborate mechanism was capabilities, which reached its highest interest in the late 1970s and early 1980s [Fabry 1974 and Wulf, Levin, and Harbison 1981]. Wilkes [1982], one of the early workers on capabilities, had this to say about capabilities:

Anyone who has been concerned with an implementation of the type just described [capability system], or has tried to explain one to others, is likely to feel that complexity has got out of hand. It is particularly disappointing that the attractive idea of capabilities being tickets that can be freely handed around has become lost ....

Compared with a conventional computer system, there will inevitably be a cost to be met in providing a system in which the domains of protection are small and frequently changed. This cost will manifest itself in terms of additional hardware, decreased runtime speed, and increased memory occupancy. It is at present an open question whether, by adoption of the capability approach, the cost can be reduced to reasonable proportions.

Today there is little interest in capabilities either from the operating systems or the computer architecture communities, although there is growing interest in protection and security.

Bell and Strecker [1976] reflected on the PDP-11 and identified a small address space as the only architectural mistake that is difficult to recover from. At the time of the creation of PDP-11, core memories were increasing at a very slow rate, and the competition from 100 other minicomputer companies meant that DEC might not have a cost-competitive product if every address had to go through the 16-bit datapath twice. Hence, the decision to add just 4 more address

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 517

bits than the predecessor of the PDP-11. The architects of the IBM 360 were aware of the importance of address size and planned for the architecture to extend to 32 bits of address. Only 24 bits were used in the IBM 360, however, because the low-end 360 models would have been even slower with the larger addresses. Unfortunately, the architects didn't reveal their plans to the software people, and the expansion effort was foiled by programmers who stored extra information in the upper eight "unused" address bits.

A few years after the Atlas paper, Wilkes published the first paper describing the concept of a cache [1965]:

The use is discussed of a fast core memory of, say, 32,000 words as slave to a slower core memory of, say, one million words in such a way that in practical cases the effective access time is nearer that of the fast memory than that of the slow memory. [p. 270]

This two-page paper describes a direct-mapped cache. While this is the first publication on caches, the first implementation was probably a direct-mapped instruction cache built at the University of Cambridge. It was based on tunnel diode memory, the fastest form of memory available at the time. Wilkes states that G. Scarott suggested the idea of a cache memory.

Subsequent to that publication, IBM started a project that led to the first commercial machine with a cache, the IBM 360/85 [Liptay 1968]. Gibson [1967] describes how to measure program behavior as memory traffic as well as miss rate and shows how the miss rate varies between programs. Using a sample of 20 programs (each with 3,000,000 references!), Gibson also relied on average memory-access time to compare systems with and without caches. This was over 20 years ago, and yet many used miss rates until recently.

Conti, Gibson, and Pitkowsky [1968] describe the resulting performance of the 360/85. The 360/91 outperforms the 360/85 on only 3 of the 11 programs in the paper, even though the 360/85 has a slower clock cycle time (80 ns versus 60 ns), smaller memory interleaving (4 versus 16), and a slower main memory (1.04 μsec versus 0.75 μsec). This is the first paper to use the term "cache." Strecker [1976] published the first comparative cache-design paper examining caches for the PDP-11. Smith [1982] later published a thorough survey paper, using the terms "spatial locality" and "temporal locality"; this paper has served as a reference for many computer designers. While most studies have relied on simulations, Clark [1983] used a hardware monitor to record cache misses of the VAX-11/780 over several days. Section 8.9 reports these findings, along with the work Clark did with Emer on TLBs [1984, 1985]. A similar study was performed on the VAX 8800 [Clark et al. 1988]. Agarwal, Sites, and Horowitz [1986] changed the microcode of a VAX to make traces of system and user code. These traces are used in this book (and are available through the publisher). Hill [1987] proposed the three Cs used in Section 8.4 to explain cache misses. Caches remain an active area of research, as Smith [1986] has recorded in his extensive bibliography.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 518

Many of the ideas in the advanced cache section have only been tried recently. The inclusion of caches on microprocessors such as the Motorola 68020 gave rise to two-level cache machines; the Sun 3/260 in 1986 was perhaps the first. In 1988, the Silicon Graphics 4D/240 had two levels of caches for data and instructions, with the second level added primarily for cache coherency to allow four-way multiprocessing. The MIPS RC 6280 is probably the first machine to go to two-level caches for the reasons given on page 465 [Roberts, Taylor, and Layman 1990]. Goodman and Chiang [1984] were the first to publish an investigation of static-column DRAM in a memory hierarchy, while Kelly [1988] refined the idea by using virtual addresses. Goodman [1987] showed that aliases can be handled at cache-miss time, and Wang, Baer, and Levy [1989] show that the extra control for this does not look too bad for two levels of cache.

In comparison to the other ideas in the advanced section, cache-coherency research is much older. Tang [1976] published the first cache-coherency protocol using directories, and this approach was implemented in the IBM 3081. Censier and Feautrier [1978] describe a technique with status tags in memory. The first machine to use snooping caches was the Synapse N+1 [Frank 1984]; the first publication on snooping caches was by Goodman [1983]. Archibald and Baer [1986] survey the wide variety of schemes for cache coherency. References on the protocols mentioned in their paper and in Figure 8.45 are Frank [1984] for Synapse; Goodman [1983] for Write Once; Katz et al. [1985] for Berkeley; McCreight [1984] for Dragon; Papamarcos and Patel [1984] for Illinois; and Thacker and Stewart [1987] for Firefly. Baer and Wang [1988] discuss multilevel inclusion. Eggers's [1989] nomenclature for categorizing snooping caches is adopted in this text. Chapter 10, Section 10.7 mentions the use of prefetching to improve cache performance, and Kroft [1981] describes the design of a cache that allows the cache to service subsequent requests while the requested data is prefetched. Przybylski [1990] and the dissertations by Agarwal [1987], Eggers [1989], and Hill [1987] investigate many aspects of the advanced cache topics in more depth.

Papers on another use of locality, register windows or stack caches, are by Patterson and Sequin [1981], Ditzel and McClellan [1982], and Lampson [1982]. Sites wrote an earlier paper [1979] suggesting one way to use the expanding resources of VLSI was to get higher performance by using a lot of registers, and these schemes are one interpretation of that recommendation.

References

AGARWAL, A. [1987]. *Analysis of Cache Performance for Operating Systems and Multiprogramming*, Ph.D. Thesis, Stanford Univ., Tech. Rep. No. CSL-TR-87-332 (May).

AGARWAL, A., R. L. SITES, AND M. HOROWITZ [1986]. "ATUM: A new technique for capturing address traces using microcode," *Proc. 13th Annual Symposium on Computer Architecture* (June 2–5), Tokyo, Japan, 119–127.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 519

ARCHIBALD, J. AND J.-L. BAER [1986]. "Cache coherence protocols: Evaluation using a multiprocessor simulation model," *ACM Trans. on Computer Systems* 4:4 (November) 273–298.

BAER, J.-L. AND W.-H. WANG [1988]. "On the inclusion property for multi-level cache hierarchies," *Proc. 15th Annual Symposium on Computer Architecture* (May–June), Honolulu, 73–80.

BELL , C. G. AND W. D. STRECKER [1976]. "Computer structures: What have we learned from the PDP-11?," *Proc. Third Annual Symposium on Computer Architecture* (January), Pittsburgh, Penn., 1–14.

BLAKKEN, J. [1983]. "Register windows for SOAR," in *Smalltalk On A RISC: Architectural Investigations,* Proc. of CS 292R (April) 126–140, University of California.

CASE, R.P. AND A. PADEGS [1978]. "The architecture of the IBM System/370," *Communications of the ACM* 21:1, 73–96. Also appears in D. P. Siewiorek, C. G. Bell, and A. Newell, *Computer Structures: Principles and Examples* (1982), McGraw-Hill, New York, 830–855.

CENSIER, L. M. AND P. FEAUTRIER [1978]. "A new solution to the coherence problem in multicache systems," *IEEE Trans. on Computers* C-27:12 (December) 1112–1118.

CLARK, D. W. [1983]. "Cache performance of the VAX-11/780," *ACM Trans. on Computer Systems* 1:1, 24–37.

CLARK, D. W. AND J. S. EMER [1985]. "Performance of the VAX-11/780 translation buffer: Simulation and measurement," *ACM Trans. on Computer Systems* 3:1, 31–62.

CLARK, D. W, P. J. BANNON, AND J. B. KELLER [1988]. "Measuring VAX 8800 Performance with a Histogram hardware monitor," *Proc. 15th Annual Symposium on Computer Architecture* (May–June), Honolulu, Hawaii, 176–185.

CONTI, C., D. H. GIBSON, AND S. H. PITOWSKY [1968]. "Structural aspects of the System/360 Model 85, part I: General organization," *IBM Systems J.* 7:1, 2–14.

CRAWFORD, J. H AND P. P. GELSINGER [1987]. *Programming the 80386*, Sybex, Alameda, Calif.

DITZEL, D. R., AND H.R. MCCLELLAN [1982]. "Register allocation for free: The C machine stack cache" *Symposium on Architectural Support for Programming Languages and Operating Systems* (March 1–3), Palo Alto, Calif., 48–56.

EGGERS, S. [1989]. *Simulation Analysis of Data Sharing in Shared Memory Multiprocessors* , Ph. D. Thesis, Univ. of California, Berkeley, Computer Science Division Tech. Rep. UCB/CSD 89/501 (April).

EMER, J. S. AND D. W. CLARK [1984]. "A characterization of processor performance of the VAX-11/780," *Proc. 11th Annual Symposium on Computer Architecture* (June), Ann Arbor, Mich., 301–310.

FABRY, R. S. [1974]. "Capability based addressing," *Comm. ACM* 17:7 (July) 403–412.

FRANK, S. J. [1984]. "Tightly coupled multiprocessor systems speed memory access times," *Electronics* 57:1 (January) 164–169.

GIBSON, D. H. [1967]. "Considerations in block–oriented systems design," *AFIPS Conf. Proc.* 30, SJCC, 75–80.

GOODMAN, J. R. [1983]. "Using cache memory to reduce processor memory traffic," *Proc. Tenth Annual Symposium on Computer Architecture* (June 5–7), Stockholm, Sweden, 124–131.

GOODMAN, J. R. and M.-C. Chiang [1984]. "The use of static column RAM as a memory hierarchy," *Proc. 11th Annual Symposium on Computer Architecture* (June 5–7), Ann Arbor, Mich., 167–174.

GOODMAN, J. R. [1987]. "Coherency for multiprocessor virtual address caches," *Proc. Second Int'l Conf. on Architectural Support for Programming Languages and Operating Systems*, Palo Alto, Calif., 71–81.

HALBERT, D. C. AND P. B. KESSLER [1980]. "Windows of overlapping register frames," *CS 292R Final Reports* (June) 82–100.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 520

HILL, M. D. [1987]. *Aspects of Cache Memory and Instruction Buffer Performance*, Ph. D. Thesis, Univ. of California at Berkeley Computer Science Division, Tech. Rep. UCB/CSD 87/381 (November).

HILL, M. D. [1988]. "A case for direct mapped caches," *Computer* 21:12 (December) 25–40.

HUGUET, M. AND T. LANG [1985]. "A reduced register file for RISC architectures," *Computer Architecture News* 13:4 (September) 22–31.

KATZ, R., S. EGGERS, D. A. WOOD, C. PERKINS, AND R. G. SHELDON [1985]. "Implementing a cache consistency protocol," *Proc. 12th Annual Symposium on Computer Architecture*, 276–283.

KELLY, E. [1988]. "'SCRAM Cache' in Sun-4/110 beats traditional caches," *Sun Technology* 1:3 (Summer) 19–21.

KILBURN, T., D. B. G. EDWARDS, M. J. LANIGAN, F. H. SUMNER [1962]. "One-level storage system," *IRE Transactions on Electronic Computers* EC-11 (April) 223–235. Also appears in D. P. Siewiorek, C. G. Bell, and A. Newell, *Computer Structures: Principles and Examples* (1982), McGraw-Hill, New York, 135–148.

KROFT, D. [1981]. "Lockup-free instruction fetch/prefetch cache organization," *Proc. Eighth Annual Symposium on Computer Architecture* (May 12–14), Minneapolis, Minn., 81–87.

LAMPSON, B. W. [1982]. "Fast procedure calls," *Symposium on Architectural Support for Programming Languages and Operating Systems* (March 1–3), Palo Alto, Calif., 66–75.

LIPTAY, J. S. [1968]. "Structural aspects of the System/360 Model 85, part II: The cache," *IBM Systems J.* 7:1, 15–21.

MCCALL, K. [1983]. "The Smalltalk-80 benchmarks," *Smalltalk 80: Bits of History, Words of Advice*, G. Krasner, ed., Addison-Wesley, Reading, Mass., 153–174.

MCCREIGHT, E. [1984]. "The Dragon computer system: An early overview," Tech. Rep. Xerox Corp. (September).

MCFARLING, S. [1989]. "Program optimization for instruction caches," *Proc. Third International Conf. on Architectural Support for Programming Languages and Operating Systems* (April 3–6), Boston, Mass., 183–191.

PAPAMARCOS, M. AND J. PATEL [1984]. "A low coherence solution for multiprocessors with private cache memories," *Proc. of the 11th Annual Symposium on Computer Architecture* (June), Ann Arbor, Mich., 348–354.

PRZYBYLSKI, S. A. [1990]. *Cache Design: A Performance-Directed Approach*, Morgan Kaufmann Publishers, San Mateo, Calif.

PRZYBYLSKI, S. A., M. HOROWITZ, AND J. L. HENNESSY [1988]. "Performance tradeoffs in cache design," *Proc. 15th Annual Symposium on Computer Architecture* (May–June), Honolulu, Hawaii, 290–298.

ROBERTS, D., G. TAYLOR, AND T. LAYMAN [1990]. "An ECL RISC microprocessor designed for two-level cache," *IEEE Compcon* (February).

SAMPLES, A. D. AND P. N. HILFINGER [1988]. "Code reorganization for instruction caches," Tech. Rep. UCB/CSD 88/447 (October), Univ. of Calif., Berkeley.

SITES, R. L., [1979]. "How to use 1000 registers," *Caltech Conf. on VLSI* (January).

SMITH, A. J. [1982]. "Cache memories," *Computing Surveys* 14:3 (September) 473–530.

SMITH, A. J. [1986]. "Bibliography and readings on CPU cache memories and related topics," *Computer Architecture News* (January) 22–42.

SMITH, J. E. AND J. R. GOODMAN [1983]. "A study of instruction cache organizations and replacement policies," *Proc. Tenth Annual Symposium on Computer Architecture* (June 5–7), Stockholm, Sweden,, 132–137.

STRECKER, W. D. [1976]. "Cache memories for the PDP-11?," *Proc. Third Annual Symposium on Computer Architecture* (January), Pittsburgh, Penn., 155–158.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 521

TANG, C. K. [1976]. "Cache system design in the tightly coupled multiprocessor system," *Proc. 1976 AFIPS National Computer Conf.*, 749–753.

TAYLOR, G. S., P. N. HILFINGER, J. R. LARUS, D. A. PATTERSON, AND B. G. ZORN [1986]. "Evaluation of the SPUR Lisp architecture," *Proc. 13th Annual Symposium on Computer Architecture* (June 2–5), Tokyo, Japan, 444–452.

THACKER, C. P. AND L. C. STEWART [1987]. "Firefly: a multiprocessor workstation," *Proc. Second Int'l Conf. on Architectural Support for Programming Languages and Operating Systems*, Palo Alto, Calif., 164–172.

UNGAR, D. M. [1987]. *The Design of a High Performance Smalltalk System*, The MIT Press Distinguished Dissertation Series, Cambridge, Mass.

WANG, W.-H., J.-L. BAER, AND H. M. LEVY [1989]. "Organization and performance of a two-level virtual-real cache hierarchy," *Proc. 16th Annual Symposium on Computer Architecture* (May 28–June 1), Jerusalem, Israel , 140–148.

WILKES, M. [1965]. "Slave memories and dynamic storage allocation," *IEEE Trans. Electronic Computers* EC-14:2 (April) 270–271.

WILKES, M. V. [1982]. "Hardware support for memory protection: Capability implementations," *Proc. Symposium on Architectural Support for Programming Languages and Operating Systems* (March 1–3), Palo Alto, Calif., 107–116.

WULF, W. A., R. LEVIN AND S. P. HARBISON [1981]. *Hydra/C.mmp: An Experimental Computer System,* McGraw-Hill, New York.

EXERCISES

8.1 [15/15/12/12] <2.2,8.4> Let's try to show how you can make *unfair* benchmarks. Here are two machines with the same processor and main memory but different cache organizations. Assume the miss time is 10 times a cache-hit time for both machines. Assume writing a 32-bit word takes 5 times as long as a cache hit (for the write-through cache), and that writing a whole 16-byte block takes 10 times as long as a cache-read hit. (for the write-back cache). The caches are unified; that is, they contain both instructions and data.

Cache A: 64 sets, 2 elements per set, each block is 16 bytes, and it uses write through.

Cache B: 128 sets, 1 element per set, each block is 16 bytes, and it uses write back.

a. [15] Describe a program that makes machine A run as much faster as possible than machine B. (Be sure to state any further assumptions you need, if any.)

b. [15] Describe a program that makes machine B run as much faster as possible than machine A. (Be sure to state any further assumptions you need, if any.)

c. [12] Approximately how much faster is the program in Part a on machine A than machine B?

d. [12] Approximately how much faster is the program in Part b on machine B than machine A?

8.2 [20] <2.2,6.4,8.4> To simplify pipelined execution, some machines insert NOP instructions rather than interlock the pipeline (see pages 273–275 in Chapter 6). Ignoring cache misses, assume that the Spice code takes 2,000,000 clocks in either case (since the version without NOPS still interlocks, which takes an extra clock each time.) Figure 8.53

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 522

shows data collected for a portion of Spice execution with a 64-KB, direct-mapped, instruction-only cache with one-word blocks.

	With NOPS	Without NOPS	Ratio with/without
Total references	1,500,000	1,180,000	1.27
Cache misses	34,153	24,908	1.37
Miss rate	2.28	2.10	1.09

FIGURE 8.53 Spice miss rates with and without NOPs.

The conclusion of a study based on Figure 8.53 was that a 9% increase in the miss rate of the program with NOPS will have a small but measurable impact on performance. What is the actual impact on performance assuming a 10-clock miss penalty?

8.3 [15/15] <8.4> You purchased an Acme computer with the following features:

1. 90% of all memory accesses are found in the cache;
2. Each cache block is two words, and the whole block is read on any miss;
3. The processor sends references to its cache at the rate of 10^7 words per second;
4. 25% of the references of (3) are writes;
5. Assume that the bus can support 10^7 words per second, reads or writes;
6. The bus reads or writes a single word at a time (the bus cannot read or write two words at once);
7. Assume at any one time, 30% of the blocks in the cache have been modified;
8. The cache uses write allocate on a write miss.

You are considering adding a peripheral to the bus, and you want to know how much of the bus bandwidth is already used. Calculate the percentage of bus bandwidth used on the average in the two cases below. The percentage is called the *traffic ratio* in the literature. Be sure to state your assumptions.

a. [15] The cache is write through.

b. [15] The cache is write back.

8.4 [20] <8.4> One drawback to the write-back scheme is that writes will probably take two cycles. During the first cycle, we detect whether a hit will occur, and during the second (assuming a hit) we actually write the data. Let's assume that 50% of the blocks are dirty for a write-back cache. Using statistics for loads and stores from DLX in Figure C.4 in Appendix C, estimate the performance of a write-through cache with a one-cycle write versus a write-back cache with a two-cycle write for each of the programs. For this question, assume that the write buffer for write through will never stall the CPU (no penalty). Assume a cache hit takes 1 clock cycle, the cache miss penalty is 10 clock

cycles, and a block write from the cache to main memory takes 10 clock cycles. Finally, assume the instruction-cache miss rate is 2% and the data-cache miss rate is 4%.

8.5 [15/20/10] <8.4> To save development time, the Sun 3/280 and the Sun 4/280 used identical memory systems, even though the CPUs were quite different. Assume the same case exists for a new machine, one board using a VAX CPU and the other a DLX CPU. For now assume the miss-rate information in Figure 8.12 and 8.16 (pages 421 and 424) apply to both architectures. Use the average column in Figure C.4 in Appendix C as needed for DLX instruction mix, and the caption of Figure 8.16 (page 424) for VAX instruction/data mix. Assume the following:

Miss penalty is 12 clock cycles.

A perfect write buffer that never stalls the CPU.

The base CPI assuming a perfect memory system is 6.0 for the VAX and 1.5 for DLX.

A unified cache adds 1 extra clock cycle to each load and store of DLX (since there is a single memory port) but not for the VAX.

You are considering three options:

1. A 4-way–set-associative unified cache of 64 KB.

2. Two 2-way–set-associative caches of 32 KB each, one for instructions and one for data.

3. A direct-mapped unified cache of 128 KB. Assume that clock rate is 10% faster in this case since the mapping is direct and the CPU address does not need to drive two caches, nor does the data bus need to be multiplexed. This faster clock rate increases the miss penalty to 13 clock cycles.

a. [15] What is the average memory-access time in clock cycles for each organization?

b. [20] What is the CPI for each machine and cache organization?

c. [10] What cache organization gives the best average performance for the two CPUs?

8.6 [25/15] <2.3,8.4,8.8> Some microprocessors have custom single-chip caches as companions to the CPU. For example, the Motorola 88100 CPU can have up to 8 of the 88200 cache chips. These chips tend to be more expensive than off-the-shelf static RAM chips. The MIPS R3000 includes a comparator on the CPU chip so that cache tags and data can be built from off-the-shelf static RAMs.

a. [25] Using the program that analyzes cache miss rates how many 16K-by-4 cache RAMs must the R3000 use to get the same performance as two 88200 chips? Both designs use separate instruction and data caches. The MIPS design assumes a block size of 16 bytes with subblock placement for each word. The cache is write through with a 4-word write buffer. The Motorola 88200 is 4-way set associative with 16 KB per chip and a 16-byte block using LRU replacement.

b. [15] Here is the data on the price of each chip (quantity 1 as of 8/1/89):

Motorola 88100: $697

Motorola 88200: $875

MIPS R3000 (25 MHz): $300

MIPS R3010 FPU (25 MHz): $350

16K by 4 SRAM (for 25 MHz R3000): $21

Which system will be cheaper and by how much?

8.7 [15/25/15/15] <2.3,8.4> The Intel i860 has its caches on chip and its die size is 1.2 cm*1.2 cm. It has a 2-way–set-associative, 4-KB instruction cache and a 2-way–set-associative, 8-KB data cache using write through or write back. Both caches use 32-byte blocks. There are no write buffers or process identifiers to reduce cache flushing. The i860 also includes a 64-entry, 4-way–set-associative TLB to map its 4-KB pages. Address translation occurs before the caches are accessed. The Cypress 7C601 CPU chip size is 0.8 cm by 0.7 cm and has no on-board cache—a cache controller chip (7C604) and two 16K * 16 cache chips (7C157) are offered to form a 64-KB unified cache. The controller includes a TLB with 64 entries managed fully associatively with 4096 process identifiers to reduce flushing. It supports 32-byte blocks with direct-mapped placement, and either write through or write back. There is a one-block write buffer for write back and a four-word write buffer for write through. The chip sizes are 1.0 cm by 0.9 cm for the 7C604 and 0.8 cm by 0.7 cm. for the 7C157.

a. [15] Using the cost model of Chapter 2, what is the cost of the Cypress chip set versus the Intel chip? (Use Figure 2.11 on page 62 to determine chip costs by finding the closest die size in that table to the Intel and Cypress die area.)

b. [25] Use the DLX cache traces and cache simulator to determine the average memory-access time for each cache organization. Assume a miss takes 6 clocks latency plus 1 clock for each 32-bit word. Assume both systems run at the same clock rate and use write allocate.

c. [15] What is the comparative cost/performance of these chips using average memory-access time as the measure?

d. [15] What is the percent increase in cost of a color workstation that uses the more expensive chips?

8.8 [25/10/15] <8.4> The CRAY X-MP instruction buffers can be thought of as an instruction-only cache. The total size is 1 KB, broken into 4 blocks of 256 bytes per block. The cache is fully associative and uses a first-in/first-out replacement policy. The access time on a miss is 10 clock cycles, with the transfer time of 64 bytes every clock cycle. The X-MP takes 1 clock cycle on a hit. Use the cache simulator and the DLX traces to determine:

a. [25] Instruction miss rate

b. [10] Average instruction memory-access time measured in clock cycles

c. [15] What does the CPI of the CRAY X-MP have to be for the portion due to instruction cache misses to be 10% or less?

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 525

8.9 [25] <8.4> Traces from a single process give too-high estimates for caches used in a multiprocess environment. Write a program that merges the uniprocess DLX traces into a single reference stream. Use the process-switch statistics in Figure 8.25 (page 439) as the average process-switch rate with an exponential distribution about that mean. (Use number of clock cycles rather than instructions, and assume the CPI of DLX is 1.5.) Use the cache simulator on the original traces and the merged trace. What is the miss rate for each assuming a 64-KB direct-mapped cache with 16-byte blocks? (There is a process-identified tag in the cache tag so that the cache doesn't have to be flushed on each switch.)

8.10 [25] <8.4> One approach to reducing misses is to prefetch the next block. A simple but effective strategy is when block i is referenced to make sure block i+1 is in the cache, and if not, to prefetch it. Do you think prefetching is more or less effective with increasing block size? Why? Is it more or less effective with increasing cache size? Why? Use statistics from the cache simulator and the traces to support your conclusion.

8.11 [20/25] <8.4> Smith and Goodman [1983] found that for a **small-instruction–only** cache, a cache using direct mapping could consistently outperform one using fully associative with LRU replacement.

a. [20] Explain why this would be possible. (Hint: you can't explain this with the 3C model because it ignores replacement policy.)

b. [25] Use the cache simulator to see if their results hold for the traces.

8.12 [Discussion] <8.4> If you look at conflict misses for a given associativity in Figure 8.12, as capacity increases the conflict misses go up and down. For example, for 2-way–set-associative mapping the miss rate for 2-KB cache is .010, a 4-KB cache is .013, and an 8-KB cache is .008. Why in the world would this happen?

8.13 [30] <8.5> Use the cache simulator and traces to calculate the effectiveness of a 4-bank versus 8-bank interleaved memory. Assume each word transfer takes one clock on the bus and a random access is 8 clocks. Measure the bank conflicts and memory bandwidth for these cases:

a. No cache and no write buffer.

b. A 64-KB, direct-mapped, write-though cache with four-word blocks.

c. A 64-KB, direct-mapped, write-back cache with four-word blocks.

d. A 64-KB, direct-mapped, write-though cache with four-word blocks but the "interleaving" comes from a page-mode DRAM.

e. A 64-KB, direct-mapped, write-back cache with four-word blocks but the "interleaving" comes from a page mode DRAM.

8.14 [20] <8.6> If the base CPI with a perfect memory system is 1.5, what is the CPI for these cache organizations? Use Figure 8.12 (page 421):

a. Direct-mapped, 16-KB unified cache using write back.

b. Two-way–set-associative, 16-KB unified cache using write back.

c. Direct-mapped, 32-KB unified cache using write back.

Assume the memory latency is 6 clocks, the transfer rate is 4 bytes per clock cycle and that 50% of the transfers are dirty. There are 16 bytes per block and 20% of the instructions are data-transfer instructions. The caches fetch words of the block in address order and the CPUs stall until all words of the block arrive. There is no write buffer. Add to the assumptions above a TLB that takes 20 clock cycles on a TLB miss. A TLB does not slow down a cache hit. For the TLB, make the simplifying assumption that 1% of all references aren't found in TLB, either when addresses come directly from the CPU or when addresses come from cache misses. What is the impact on performance of the TLB if the cache above is physical or virtual?

8.15 [30] <3.8,8.9> The example in Section 8.9 (page 478) refines the instructions fetched into the CPU from the cache due to the instruction-prefetch buffer. How does this increase of 13% to 39% in instruction words fetched affect the difference in the instruction words fetched from DLX versus VAX? The extra instruction fetches of the VAX hurt only when they bring something into the cache that is not used before it is displaced, while DLX would seem to need a larger cache for its larger program. Write a simulator emulating the instruction-prefetch buffer to measure the increase in cache misses using the VAX address traces and see if prefetching is a significant increase in cache misses.

8.16 [25–40] <8.7> Study the impact of adding register windows to DLX. This study can range from simply estimating the register-traffic savings to modifying the DLX compiler and simulator to measure costs and benefits directly.

8.17 [10] <8.8> Data General described the design of a three-level cache for an ECL implementation of the 88000 architecture. What is the formula for average access time for a three-level cache?

8.18 [20] <8.8> What is the performance loss for a four-way multiprocessor with I/O devices? Suppose 1% of all data references to the cache cause invalidation to the other data caches and that all CPUs stall four clocks on an invalidation. Assume a 64-KB, direct-mapped cache for data and a 64-KB, direct-mapped cache for instructions with a block size of 32 bytes yields a 1% miss rate for instructions and a 2% miss rate for data, with 20% of all CPU memory references being for data. The CPI of the CPU is 1.5 with a perfect memory system and it takes 10 clocks on a cache miss whether the data is dirty or clean.

8.19 [25] <8.8> Use the traces to calculate the effectiveness of early restart and out-of-order fetch. What is the distribution of first accesses to a block as block size increases from 2 words to 64 words by factors of two for:

a. A 64-KB, instruction-only cache?

b. A 64-KB, data-only cache?

c. A 128-KB unified cache?

Assume direct-mapped placement.

8.20 [30] <8.8> Use the cache simulator and traces with a program you write yourself to compare the effectiveness schemes for fast writes:

a. 1-word buffer and the CPU stalls on a data-read cache miss with a write-through cache.

b. 4-word buffer and the CPU stalls on a data-read cache miss with a write-through cache.

c. 4-word buffer and the CPU stalls on a data-read cache miss only if there is a potential conflict in the addresses with a write-through cache.

d. A write-back cache that writes dirty data first and then loads the missed block.

e. A write-back cache with a one-block write buffer that loads the miss data first and then stalls the CPU on a clean miss if the write buffer is not empty.

f. A write-back cache with a one-block write buffer that loads the miss data first and then stalls the CPU on a clean miss only if the write buffer is not empty and there is a potential conflict in the addresses.

Assume a 64-KB, direct-mapped cache for data and a 64-KB, direct-mapped cache for instructions with a block size of 32 bytes. The CPI of the CPU is 1.5 with a perfect memory system and it takes 14 clocks on a cache miss and 7 clocks to write a single word to memory.

8.21 [30] <8.8> Use the cache simulator and traces with a program you write yourself to create a two-level cache simulator. Use this program to see at what cache size is the global miss rate of a second-level cache approximately the same as a single-level cache of the same capacity.

8.22 [Discussion] <8.6> Some people have argued that with increasing capacity of memory storage per chip, virtual memory is an idea whose time has passed, and they expect to see it dropped from future computers. Find reasons for and against this argument.

8.23 [Discussion] <8.6> So far, few computer systems take advantage of the extra security available with gates and rings found in a machine like the Intel 80286. Construct some scenario whereby the computer industry would switch over to this model of protection.

8.24 [Discussion] <8.4> Recent research has tried to use compilers to improve cache performance (see McFarling [1989] and Samples and Hilfinger [1988]):

a. Which of the 3C's are compilers trying to improve and which are they not? Why?

b. Which mapping is best for compiler improvement? Why?

8.25 [Discussion] <8.3> Many times a new technology has been invented that is expected to make a major change to the memory hierarchy. For the sake of this question, let's suppose that biological computer technology becomes a reality. Suppose biological

memory technology has an unusual characteristic: It is as fast as the fastest semiconductor DRAMs, and it can be randomly accessed; but it only costs as much as magnetic-disk memory. It has the further advantage of not being any slower no matter how big it is. The only drawback is that you can only Write it Once, but you can Read it Many times. Thus it is called a "WORM" memory. Because of the way it is manufactured, the WORM- memory module can be easily replaced. See if you can come up with several new ideas to take advantage of WORMs to build better computers using "bio-technology."

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 529

I/O certainly has been lagging in the last decade.

Seymour Cray, Public Lecture (1976)

Also, I/O needs a lot of work.

David Kuck, Keynote Address,
15th Annual Symposium on Computer Architecture (1988)

9.1	**Introduction**	**499**
9.2	**Predicting System Performance**	**501**
9.3	**I/O Performance Measures**	**506**
9.4	**Types of I/O Devices**	**512**
9.5	**Buses—Connecting I/O Devices to CPU/Memory**	**528**
9.6	**Interfacing to the CPU**	**533**
9.7	**Interfacing to an Operating System**	**535**
9.8	**Designing an I/O System**	**539**
9.9	**Putting It All Together: The IBM 3990 Storage Subsystem**	**546**
9.10	**Fallacies and Pitfalls**	**554**
9.11	**Concluding Remarks**	**559**
9.12	**Historical Perspective and References**	**560**
	Exercises	**563**

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.

9 Input/Output

9.1 Introduction

Input/output has been the orphan of computer architecture. Historically neglected by CPU enthusiasts, the prejudice against I/O is institutionalized in the most widely used performance measure, CPU time (page 35). Whether a computer has the best or the worst I/O system in the world cannot be measured by CPU time, which by definition ignores I/O. The second class citizenship of I/O is even apparent in the label "peripheral" applied to I/O devices.

This attitude is contradicted by common sense. A computer without I/O devices is like a car without wheels—you can't get very far without them. And while CPU time is interesting, response time—the time between when the user types a command and when she gets results—is surely a better measure of performance. The customer who pays for a computer cares about response time, even if the CPU designer doesn't. Finally, as rapid improvements in CPU performance compress traditional classes of computers together, it is I/O that serves to distinguish them:

- The difference between a mainframe computer and a minicomputer is that a mainframe can support many more terminals and disks.
- The difference between a minicomputer and a workstation is that a workstation has a screen, a keyboard, and a mouse.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 531

- The difference between a file server and a workstation is that a file server has disks and tape units but no screen, keyboard, or mouse.
- The difference between a workstation and a personal computer is that workstations are always connected together on a network.

It may come to pass that computers from high-end workstations to low-end supercomputers will use the same "super-microprocessors." Differences in cost and performance would be determined only by the memory and I/O systems (and the number of processors).

I/O's revenge is at hand. Suppose we have a difference between CPU time and response time of 10%, and we speed up the CPU by a factor of 10, while neglecting I/O. Amdahl's Law tells us that we will get a speedup of only 5 times, with half the potential of the CPU wasted. Similarly, making the CPU 100 times faster without improving the I/O would obtain a speedup of only 10 times, squandering 90% of the potential. If, as predicted in Chapter 1, performance of CPUs improves at 50% to 100% per year, and I/O does not improve, every task will become I/O bound. There would be no reason to buy faster CPUs—and no jobs for CPU designers.

While this single chapter cannot fully vindicate I/O, it may at least atone for some of the sins of the past and restore some balance.

Are CPUs Ever Idle?

Some suggest that the prejudice is well founded. I/O speed doesn't matter, they argue, since there is always another process to run while one process waits for a peripheral.

There are several points to make in reply. First, this is an argument that performance is measured as throughput—more tasks per hour—rather than as response time. Plainly, if users didn't care about response time, interactive software never would have been invented, and there would be no workstations today. (The next section gives experimental evidence on the importance of response time.) It may also be expensive to rely on processes while waiting for I/O, since main memory must be larger or else the paging traffic from process switching would actually increase I/O. Furthermore, with desktop computing there is only one person per CPU, and thus fewer processes than in timesharing; many times the only waiting process is the human being! And some applications, such as transaction processing (Section 9.3), place strict limits on response time as part of the performance analysis.

But let's accept the argument at face value and explore it further. Suppose the difference between response time and CPU time today is 10%, and a CPU that is ten times faster can be achieved without changing I/O performance. A process will then spend 50% of its time waiting for I/O, and two processes will have to be perfectly aligned to avoid CPU stalls while waiting for I/O. Any further CPU improvement will only increase CPU idle time.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 532

Thus, I/O throughput can limit system throughput, just as I/O response time limits system response time. Let's see how to predict performance for the whole system.

9.2 Predicting System Performance

System performance is limited by the slowest part of the path between CPU and I/O devices. The performance of a system can be limited by the speed of any of these pieces of the path, shown in Figure 9.1:

- The CPU
- The cache memory
- The main memory
- The memory–I/O bus
- The I/O controller or I/O channel
- The I/O device
- The speed of the I/O software
- The efficiency of the software's use of the I/O device

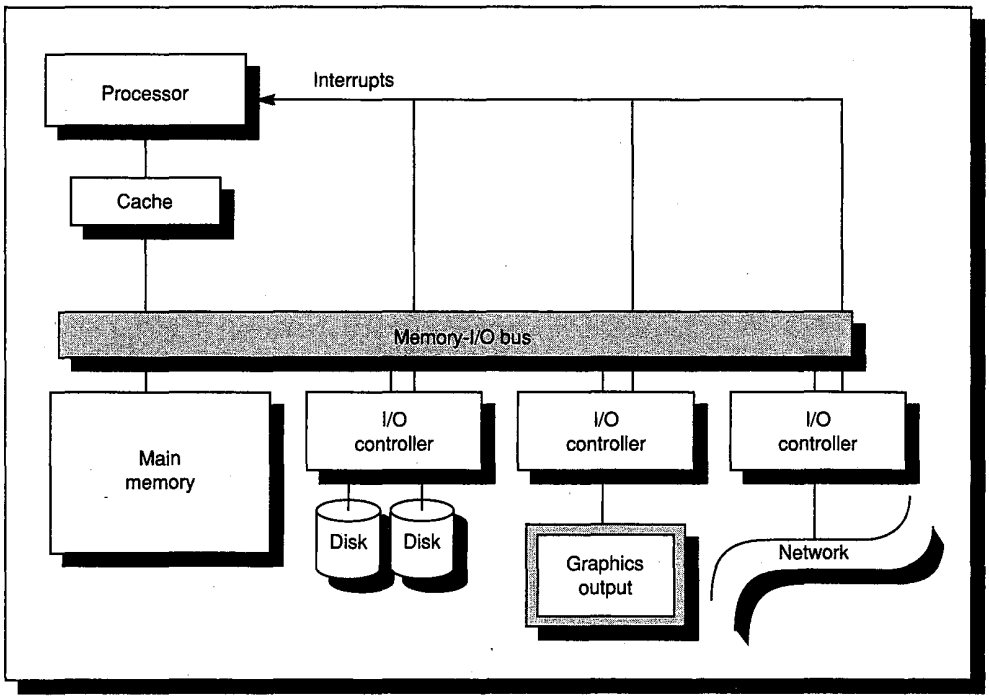

FIGURE 9.1 Typical collection of I/O devices on a computer.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 533

If the system is not balanced, the high performance of some components may be lost due to the low performance of one link in the chain. The art of I/O design is to configure a system such that the speeds of all components are matched.

In earlier chapters we have assumed that the fastest CPU is the single object of our desire, but CPU performance is not the same as system performance. For example, suppose we have two workloads, A and B. Both workloads take 10 seconds to run. Workload A does so little I/O that it is not worth mentioning. Workload B keeps I/O devices busy four seconds, and this time is completely overlapped with CPU activities. Suppose the CPU is replaced by a newer model with five times the performance. Intuitively, we realize that workload A takes two seconds—fully five times faster—but workload B is I/O bound and cannot take less than four seconds. Figure 9.2 illustrates our intuition.

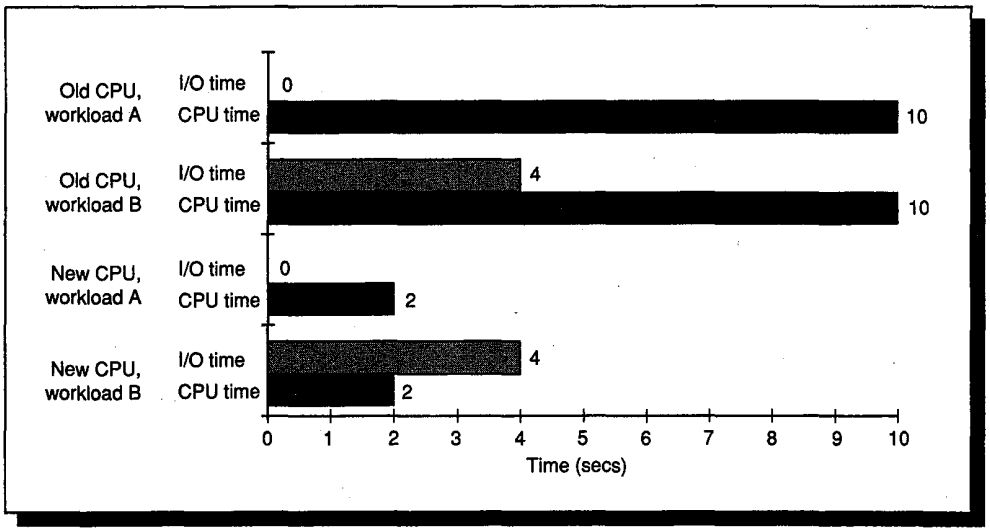

FIGURE 9.2 The overlapped execution of the two workloads with the original CPU and then a CPU with five times the performance. We can see that the elapsed time for workload A is indeed 1/5 of the time with the new CPU, but it is limited to four seconds in workload B because I/O speed is not improved.

Determining the performance of such cases requires a new formula. The elapsed execution time of a workload can be broken into three pieces

$$\text{Time}_{\text{workload}} = \text{Time}_{\text{CPU}} + \text{Time}_{\text{I/O}} - \text{Time}_{\text{overlap}}$$

where Time_{CPU} means the time the CPU is busy, $\text{Time}_{\text{I/O}}$ means the time the I/O system is busy, and $\text{Time}_{\text{overlap}}$ means the time both the CPU and the I/O system are busy. Using workload B with the old CPU in Figure 9.2 as an example, the times in seconds are:

10 for $\text{Time}_{\text{workload}}$,

10 for Time_{CPU},

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 534

4 for $\text{Time}_{I/O}$, and

4 for $\text{Time}_{overlap}$.

Assuming we speed up only the CPU, one way to calculate the time to execute the workload is:

$$\text{Time}_{workload} = \frac{\text{Time}_{CPU}}{\text{Speedup}_{CPU}} + \text{Time}_{I/O} - \frac{\text{Time}_{overlap}}{\text{Speedup}_{CPU}}$$

Since the CPU time is shrunk, it stands to reason that the overlap time is also shrunk. The system speedup when we want to improve I/O is equivalent:

$$\text{Time}_{workload} = \text{Time}_{CPU} + \frac{\text{Time}_{I/O}}{\text{Speedup}_{I/O}} - \frac{\text{Time}_{overlap}}{\text{Speedup}_{I/O}}$$

Let's try an example before explaining a limitation of these formulas.

Example

One workload takes 50 seconds to run, with the CPU being busy 30 seconds and the I/O being busy 30 seconds. How much time will the workload take if we replace the CPU with one that has four times the performance?

Answer

The total elapsed time is 50 seconds, yet the sum of CPU time and I/O time is 60 seconds. Thus the overlap time must be 10 seconds. Plugging into the formula:

$$\text{Time}_{workload} = \frac{\text{Time}_{CPU}}{\text{Speedup}_{CPU}} + \text{Time}_{I/O} - \frac{\text{Time}_{overlap}}{\text{Speedup}_{CPU}} = \frac{30}{4} + 30 - \frac{10}{4} = 35$$

This example uncovers a complication with this formula: How much of the time that the workload is busy on the faster CPU is overlapped with I/O? Figure 9.3 (page 504) shows three options. Depending on the resulting overlap after speedup, the time for the workload varies from 30 to 37.5 seconds.

In reality we can't know which is correct without measuring the workload on the faster CPU to see what overlap occurs. The formulas above assume option (c) in Figure 9.3; the overlap scales by the same speedup as the CPU, so we will call it Time_{scaled} (rather than $\text{Time}_{workload}$). Maximum overlap assumes that as much of the overlap as possible is maintained, but that the new overlap cannot be larger than the original overlap or the CPU time after speedup. Minimum overlap assumes that as much of the overlap as possible is eliminated, but that the overlap time will not shrink by more than the time removed from the CPU or I/O time. If we introduce the abbreviations $\text{New}_{CPU} = \text{Time}_{CPU} / \text{Speedup}_{CPU}$ and $\text{New}_{I/O} = \text{Time}_{I/O} / \text{Speedup}_{I/O}$, the time of the workload for maximum overlap (Time_{best}) and minimum overlap (Time_{worst}) can be written as:

$$\text{Time}_{best} = \text{New}_{CPU} + \text{Time}_{I/O} - \text{Minimum}\ (\text{Time}_{overlap}, \text{New}_{CPU})$$

$$\text{Time}_{worst} = \text{New}_{CPU} + \text{Time}_{I/O} - \text{Maximum}\ (0, \text{Time}_{overlap} - (\text{Time}_{CPU} - \text{New}_{CPU}))$$

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 535

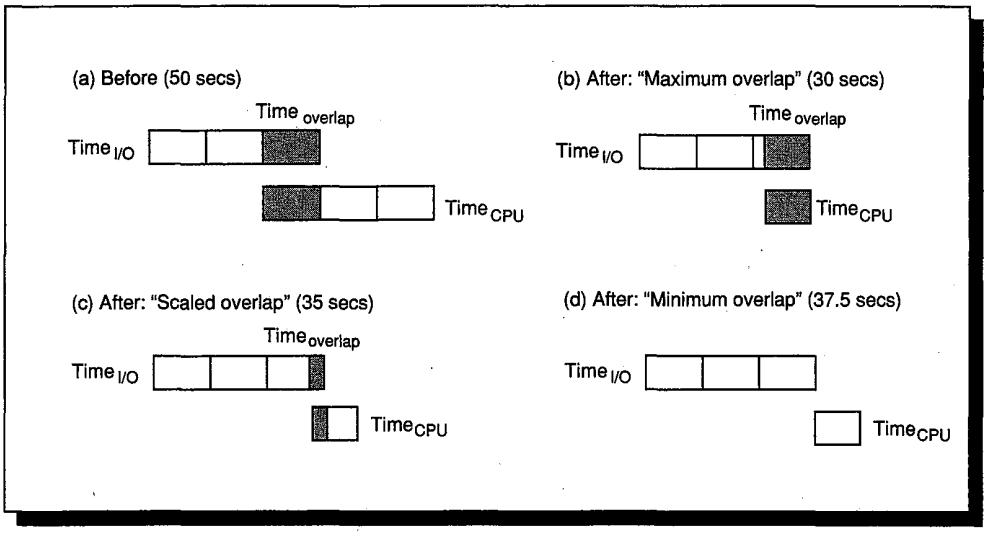

FIGURE 9.3 The original overlap in the example above (a) and three interpretations of overlap after speedup. Each block represents 10 seconds, except that the block for the new CPU time is 7.5 seconds. The overlapped portions of Time_{CPU} and $\text{Time}_{I/O}$ are shaded. (b) shows the new Time_{CPU} overlapping completely with I/O, giving a time of the workload of 30 seconds. (c) shows the overlap of the Time_{CPU} is scaled with Speedup_{CPU}, giving a total of 35 seconds, with 2.5 seconds of overlapped execution. (d) shows no overlap with I/O, so the total is 37.5 seconds.

Example

Calculate the three time predictions for workload B in Figure 9.2

Answer

$$\text{Time}_{best} = \frac{10}{5} + 4 - \text{Minimum}\left(\frac{10}{5}, 4\right) = 2 + 4 - 2 = 4$$

$$\text{Time}_{scaled} = \frac{10}{5} + 4 - \frac{4}{5} = 2 + 4 - 0.8 = 5.2$$

$$\text{Time}_{worst} = \frac{10}{5} + 4 - \text{Maximum}\ (0, 4 - (10 - \frac{10}{5})) = 2 + 4 - 0 = 6$$

Sometimes changes will be made to both the CPU and the I/O system. The formulas become:

$$\text{Time}_{scaled} = \text{New}_{CPU} + \text{New}_{I/O} - \frac{\text{Time}_{overlap}}{\text{Maximum}(\text{Speedup}_{CPU}, \text{Speedup}_{I/O})}$$

$$\text{Time}_{best} = \text{New}_{CPU} + \text{New}_{I/O} - \text{Minimum}(\text{Time}_{overlap}, \text{New}_{CPU}, \text{New}_{I/O})$$

$$\text{Time}_{worst} = \text{New}_{CPU} + \text{New}_{I/O} - \text{Max}\ (0, \text{Time}_{overlap} - \text{Max}\ (\text{Time}_{CPU} - \text{New}_{CPU}, \text{Time}_{I/O} - \text{New}_{I/O}))$$

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 536

The formula for scaled overlap says that the overlap period is reduced by the larger of the two speedups. The formula for maximum overlap ($Time_{best}$) says that as much overlap as possible is retained, but the new overlap cannot be larger than the original overlap or the CPU or I/O time after speedup. Finally, the formula for minimum overlap ($Time_{worst}$) says that the overlap is reduced by the larger of the time removed from the CPU time and the time removed from the I/O time (but that the overlap time cannot be less than 0). Figure 9.4 shows the three examples of speedup where both the I/O and CPU are improved.

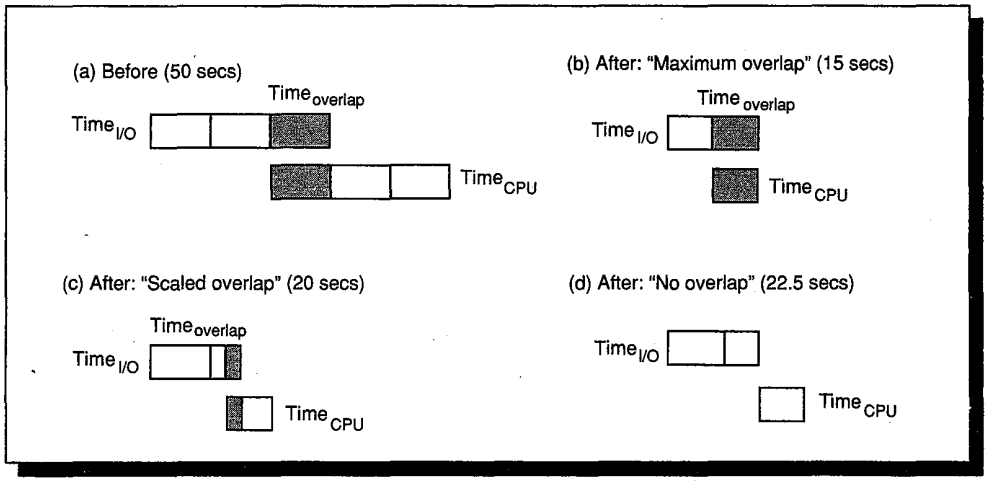

FIGURE 9.4 Time for workload in Figure 9.3(a) with $Speedup_{CPU}$ = 4 and $Speedup_{I/O}$ = 2.

Let's look at a detailed example showing speedup of both the CPU and I/O.

Example

Suppose a workload on the current systems takes 64 seconds. The CPU is busy the whole time, and the channels connecting the I/O devices to the CPU are busy 36 seconds. The computer manager is considering two upgrade options: either a single CPU that has twice the performance, or two CPUs that have twice the throughput and twice as many channels. The time of the actual I/O devices is so small it can be ignored. For the dual CPU option assume that the workload can be evenly spread between the CPUs and channels. What is the performance improvement for each option?

Answer

Since there is no change to the I/O system with the single faster CPU, time for the workload assuming scaled overlap is then simply

$$Time_{scaled} = \frac{Time_{CPU}}{Speedup_{CPU}} + Time_{I/O} - \frac{Time_{overlap}}{Speedup_{CPU}}$$

$$= \frac{64}{2} + 36 - \frac{36}{2} = 32 + 36 - 18 = 50$$

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 537

For the dual CPU with more channels,

$$\text{Time}_{\text{scaled}} = \frac{\text{Time}_{\text{CPU}}}{\text{Speedup}_{\text{CPU}}} + \frac{\text{Time}_{\text{I/O}}}{\text{Speedup}_{\text{I/O}}} - \frac{\text{Time}_{\text{overlap}}}{\text{Maximum}(\text{Speedup}_{\text{CPU}}, \text{Speedup}_{\text{I/O}})}$$

$$= \frac{64}{2} + \frac{36}{2} - \frac{36}{\text{Maximum}(2\ ,\ 2)} = 32 + 18 - 18 = 32$$

Assuming scaled overlap, the dual CPU is more than 50% faster. Using best-case scaling, the dual CPU is 13% faster, while worst-case scaling suggests it is 39% faster.

As these examples demonstrate, we need improvement in I/O performance to match the improvement in CPU performance if we are to achieve faster computer systems. We can now examine metrics of I/O devices to understand how to improve their performance and thus the whole system.

9.3 I/O Performance Measures

I/O performance has measures that have no counterparts in CPU design. One of these is diversity: Which I/O devices can connect to the computer system? Another is capacity: How many I/O devices can connect to a computer system?

In addition to these unique measures, the traditional measures of performance, response time and throughput also apply to I/O. (I/O throughput is sometimes called "I/O bandwidth" and response time is sometimes called "latency.") The next two figures offer insight into how response time and throughput trade off against each other. Figure 9.5 shows the simple producer-server model. The producer creates tasks to be performed and places them in the queue; the server takes tasks from the queue and performs them.

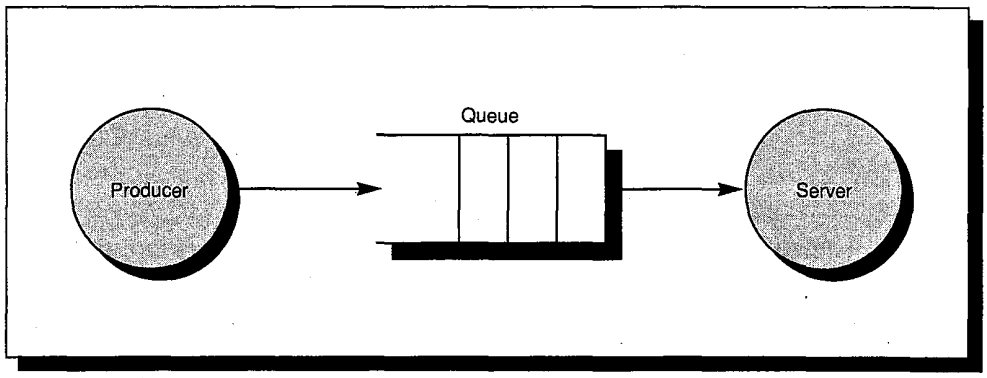

FIGURE 9.5 The traditional producer-server model of response time and throughput. Response time begins when a task is placed in the queue and ends when it is completed by the server. Throughput is the number of tasks completed by the server in unit time.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 538

Response time is defined as the time a task takes from the moment it is placed in the queue until the server finishes the task. Throughput is simply the average number of tasks completed by the server over a time period. To get the highest possible throughput, the server should never be idle, and thus the queue should never be empty. Response time, on the other hand, counts time spent in the queue and is therefore minimized by the queue being empty.

Another measure of I/O performance is the interference of I/O with CPU execution. Transferring data may interfere with the execution of another process. There is also overhead due to handling I/O interrupts. Our concern here is how many more clock cycles a process will take because of I/O for another process.

Throughput Versus Response Time

Figure 9.6 shows throughput versus response time (or latency), for a typical I/O system. The knee of the curve is the area where a little more throughput results in much longer response time or, conversely, a little shorter response time results in much lower throughput.

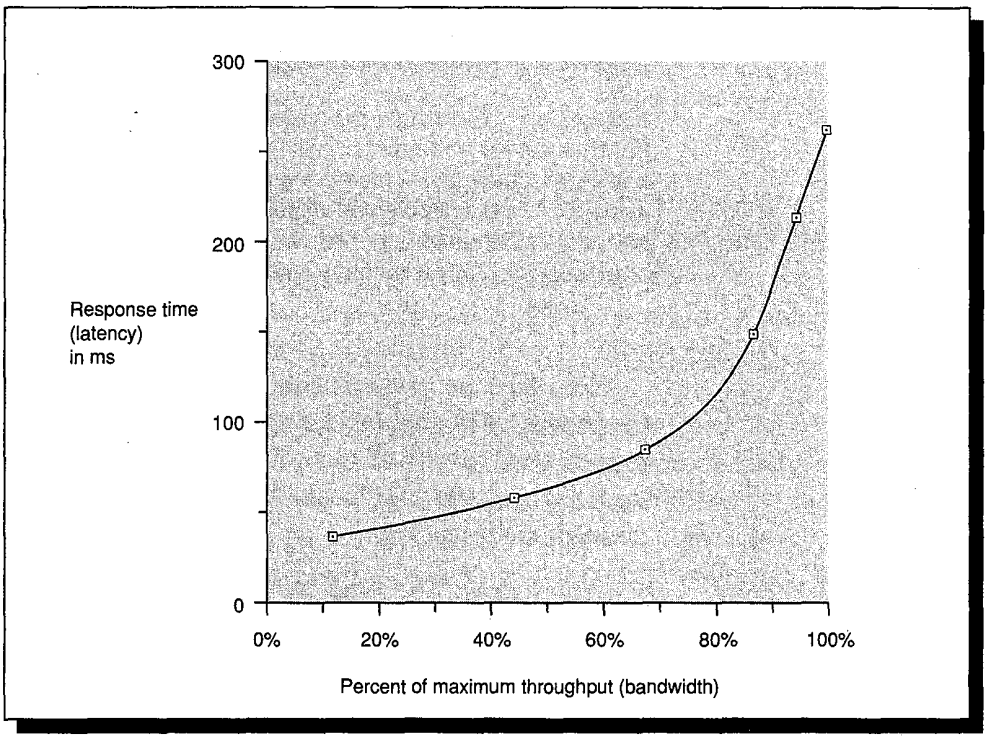

FIGURE 9.6 Throughput versus response time. Latency is normally reported as response time. Note that absolute minimum response time achieves only 11% of the throughput while the response time for 100% throughput takes seven times the minimum response time. Chen [1989] collected these data for an array of magnetic disks.

Life would be simpler if improving performance always meant improvements in both response time and throughput. Adding more servers, as in Figure 9.7, increases throughput: By spreading data across two disks instead of one, tasks may be serviced in parallel. Alas, this does not help response time, unless the workload is held constant and the time in the queues is reduced because of more resources.

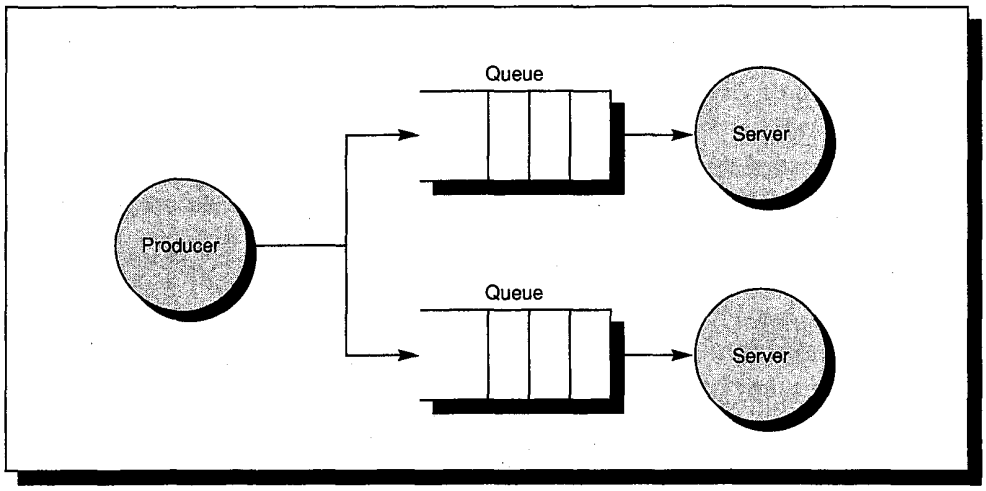

FIGURE 9.7 The single-producer, single-server model of Figure 9.5 is extended with another server and queue. This increases I/O system throughput and takes less time to service producer tasks. Increasing the number of servers is a common technique in I/O systems. There is a potential imbalance problem with two queues; unless data is placed perfectly in the queues, sometimes one server will be idle with an empty queue while the other server is busy with many tasks in its queue.

How does the architect balance these conflicting demands? If the computer is interacting with human beings, Figure 9.8 suggests an answer. This figure presents the results of two studies of interactive environments, one keyboard oriented and one graphical. An interaction or *transaction* with a computer is divided into three parts:

1. *Entry time*: The time for the user to enter the command. In the graphics system in Figure 9.8 it took 0.25 seconds on average to enter the command versus 4.0 seconds for the conventional system.
2. *System response time*: The time between when the user enters the command and the complete response is displayed.
3. *Think time*: The time from the reception of the response until the user begins to enter the next command.

The sum of these three parts is called the *transaction time*. Several studies report that user productivity is inversely proportional to transaction time; transactions per hour measures the work completed per hour by the user.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 540

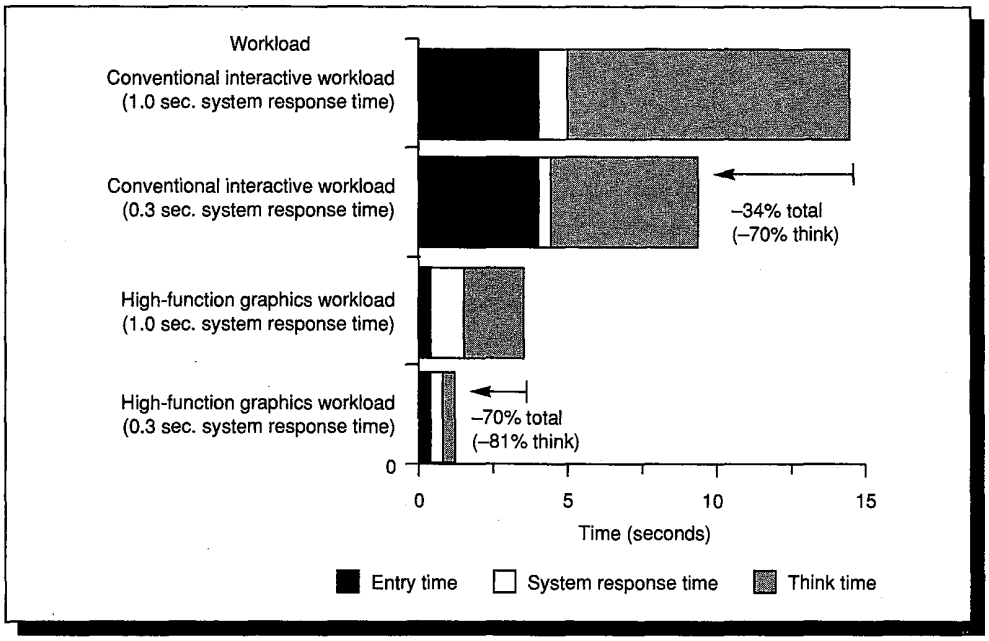

FIGURE 9.8 A user transaction with an interactive computer divided into entry time, system response time, and user think time for a conventional system and graphics system. The entry times are the same independent of system response time. The entry time was 4 seconds for the conventional system and 0.25 seconds for the graphics system. (From Brady [1986].)

The results in Figure 9.8 show that reduction in response time actually decreases transaction time by more than just the response time reduction: Cutting system response time by 0.7 seconds saves 4.9 seconds (34%) from the conventional transaction and 2.0 seconds (70%) from the graphics transaction. This implausible result is explained by human nature; people need less time to think when given a faster response.

Whether these results are explained as a better match to the human attention span or getting people "on a roll," several studies report this behavior. In fact, as computer responses drop below a second, productivity seems to make a more than linear jump. Figure 9.9 (page 510) compares transactions per hour (the inverse of transaction time) of a novice, an average engineer, and an expert performing physical design tasks at graphics displays. System response time magnified talent: a novice with subsecond response time was as productive as an experienced professional with slower response, and the experienced engineer in turn could outperform the expert with a similar advantage in response time. In all cases the number of transactions per hour jumps more than linearly with subsecond response time.

Since humans may be able to get much more work done per day with better response time, it is possible to attach an economic benefit to the customer of lowering response time into the subsecond range [IBM 1982], thereby helping the architect decide how to tip the balance between response time and throughput.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 541

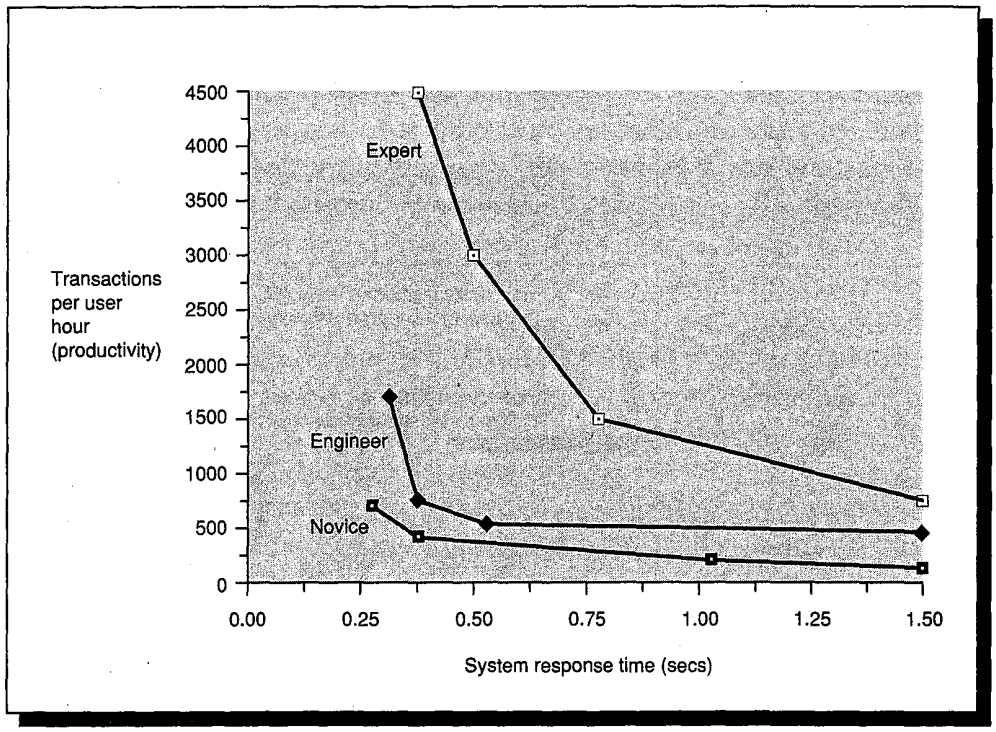

FIGURE 9.9 Transactions per hour versus computer response time for a novice, experienced engineer, and expert doing physical design on a graphics system. Transactions per hour is a measure of productivity. (From IBM [1982].)

Examples of Measurements of I/O Performance—Magnetic Disks

Benchmarks are needed to evaluate I/O performance, just as they are needed to evaluate CPU performance. We begin with benchmarks for magnetic disks. Three traditional applications of disks are with large-scale scientific problems, transaction processing, and file systems.

Supercomputer I/O Benchmarks

Supercomputer I/O is dominated by accesses to large files on magnetic disks. For example, Bucher and Hayes [1980] benchmarked supercomputer I/O using 8-MB sequential file transfers. Many supercomputer installations run batch jobs, each of which may last for hours. In these situations, I/O consists of one large read followed by writes to snapshot the state of the computation should the computer crash. As a result, supercomputer I/O in many cases consists of more output than input. Some models of Cray Research computers have such limited main memory that programmers must break their programs into overlays and swap them to disk (see Section 8.5 of Chapter 8), which also causes large sequential transfers. Thus, the overriding supercomputer I/O measure is data

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 542

throughput: number of bytes per second that can be transferred between supercomputer main memory and disks during large transfers.

Transaction Processing I/O Benchmarks

In contrast, *transaction processing* (TP) is chiefly concerned with *I/O rate*: the number of disk accesses per second, as opposed to *data rate*, measured as bytes of data per second. TP generally involves changes to a large body of shared information from many terminals, with the TP system guaranteeing proper behavior on a failure. If, for example, a bank's computer fails when a customer withdraws money, the TP system would guarantee that the account is debited if the customer received the money and that the account is unchanged if the money was not received. Airline reservations systems as well as banks are traditional customers for TP.

Two dozen members of the TP community conspired to form a benchmark for the industry and, to avoid the wrath of their legal departments, published the report anonymously [1985]. This benchmark, called DebitCredit, simulates bank tellers and has as its bottom line the number of debit/credit transactions per second (TPS); in 1990, the TPS for high-end machines is about 300. The DebitCredit performs the operation of a customer depositing or withdrawing money. The performance measurement is the peak TPS, with 95% of the transactions having less than a one-second response time. The DebitCredit computes the cost per TPS, based on the five-year cost of the computer-system hardware and software. Disk I/O for DebitCredit is random reads and writes of 100-byte records along with occasional sequential writes.

Depending on how cleverly the transaction-processing system is designed, each transaction results in between two and ten disk I/Os and takes between 5,000 and 20,000 CPU instructions per disk I/O. The variation largely depends on the efficiency of the transaction processing software, although in part it depends on the extent to which disk accesses can be avoided by keeping information in main memory. The benchmark requires that for TPS to increase, the number of tellers and the size of the account file must also increase. Figure 9.10 shows this unusual relationship in which more TPS requires more users.

TPS	Number of ATMs	Account-file size
10	1,000	0.1 GB
100	10,000	1.0 GB
1,000	100,000	10.0 GB
10,000	1,000,000	100.0 GB

FIGURE 9.10 Relationship among TPS, tellers, and account-file size. The DebitCredit benchmark requires that the computer system handle more tellers and larger account files before it can claim a higher transaction-per-second milestone. The benchmark is supposed to include "terminal handling" overhead, but this metric is sometimes ignored.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 543

This is to ensure that the benchmark really measures disk I/O; otherwise a large main memory dedicated to a database cache with a small number of accounts would unfairly yield a very high TPS. (Another perspective is the number of accounts must grow since a person is not likely to use the bank more frequently just because the bank has a faster computer!)

File System I/O Benchmarks

File systems, for which disks are mainly used in timesharing systems, have a different access pattern. Ousterhout et al. [1985] measured a UNIX file system and found that 80% of accesses to files of less than 10 KB and 90% of **all** file accesses were sequential. The distribution by type of file access was 67% reads, 27% writes, and 6% read-write accesses. In 1988, Howard et al. [1988] proposed a file-system benchmark that is becoming popular. Their paper describes five phases of the benchmark, using 70 files with a total size of 200 KB:

MakeDir—Constructs a target subtree that is identical in structure to the source subtree.

Copy—Copies every file from the source subtree to the target subtree.

ScanDir—Recursively traverses the target subtree and examines the status of every file in it. It does not actually read the contents of any file.

ReadAll—Scans every byte of every file in the target subtree once.

Make—Compiles and links all the files in the target subtree. [p. 55]

The file-system measurements of Howard et al. [1988], like those of Ousterhout et al. [1985], found the ratio of disk reads to writes to be about 2:1. This benchmark reflects that measure.

9.4 Types of I/O Devices

Now that we have covered measurements of I/O performance, let's describe the devices themselves. While the computing model has changed little since 1950, I/O devices have become rich and diverse. Three characteristics are useful in organizing this disparate conglomeration:

- *Behavior*—input (read once), output (write only, cannot be read), or storage (can be reread and usually rewritten)
- *Partner*—either a human or a machine is at the other end of the I/O device, either feeding data on input or reading data on output
- *Data rate*—the peak rate at which data can be transferred between the I/O device and the main memory or CPU

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 544

Using these characteristics, a keyboard is an input device used by a human with a peak data rate of about 10 bytes per second. Figure 9.11 shows some of the I/O devices connected to computers.

The advantage of designing I/O devices for humans is that the performance target is fixed. Figure 9.12 shows the I/O performance of people.

Device	Behavior	Partner	Data rate (KB/sec)
Keyboard	Input	Human	0.01
Mouse	Input	Human	0.02
Voice input	Input	Human	0.02
Scanner	Input	Human	200.00
Voice output	Output	Human	0.60
Line printer	Output	Human	1.00
Laser printer	Output	Human	100.00
Graphics display	Output	Human	30,000.00
(CPU to frame buffer)	Output	Human	200.00
Network-terminal	Input or output	Machine	0.05
Network-LAN	Input or output	Machine	200.00
Optical disk	Storage	Machine	500.00
Magnetic tape	Storage	Machine	2,000.00
Magnetic disk	Storage	Machine	2,000.00

FIGURE 9.11 Examples of I/O devices categorized by behavior, partner, and data rate. This is the raw data rate of the device rather than the rate an application would see. Storage devices can be further distinguished by whether they support sequential access (e.g., tapes) or random access (e.g., disks). Note that networks can act either as input or output devices but, unlike storage, cannot reread the same information.

Human organ	I/O rate (KB/sec)	I/O latency (ms)
Ear	8.000–60.000	10
Eye—reading text	0.030–0.375	10
Eye—pattern recognition	125.000	10
Hand—typing	0.010–0.020	100
Voice	0.003–0.015	100

FIGURE 9.12 Peak I/O rates for people. Input via seeing patterns is our highest I/O rate; hence the popularity of graphic output devices. Maberly [1966] says the average reading speed is 28 bytes per second and the maximum is 375 bytes per second. The telephone company sets a 170-ms limit to the time between when an operator pushes a button to accept a call until a voice path must be established. The phone company transmits voice at 8 KB per second. (None of these parameters are expected to change, unless anabolic steroids become a breakfast supplement!)

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 545

To put the data rates of each device into perspective, Figure 9.13 shows the relative peak memory bandwidth needed to support each device, assuming a computer had exactly one of each device transferring at its peak rate.

Rather than discuss the characteristics of all I/O devices, we will concentrate on the three devices with the highest data rates: magnetic disks, graphics displays, and local area networks. These are also the devices that have the highest leverage on user productivity. In this chapter we are not talking about floppy disks, but the original "hard" disks. These magnetic disks are what IBM calls *DASD*s, for *Direct-Access Storage Devices.*

Magnetic Disks

I think Silicon Valley was misnamed. If you look back at the dollars shipped in products in the last decade there has been more revenue from magnetic disks than from silicon. They ought to rename the place Iron Oxide Valley.

Al Hoagland, one of the pioneers of magnetic disks (1982)

In spite of repeated attacks by new technologies, magnetic disks have dominated secondary storage since 1965. Magnetic disks play two roles in computer systems:

- Long-term, nonvolatile storage for files, even when no programs are running
- A level of the memory hierarchy below main memory used for virtual memory during program execution (see Section 8.5 in Chapter 8)

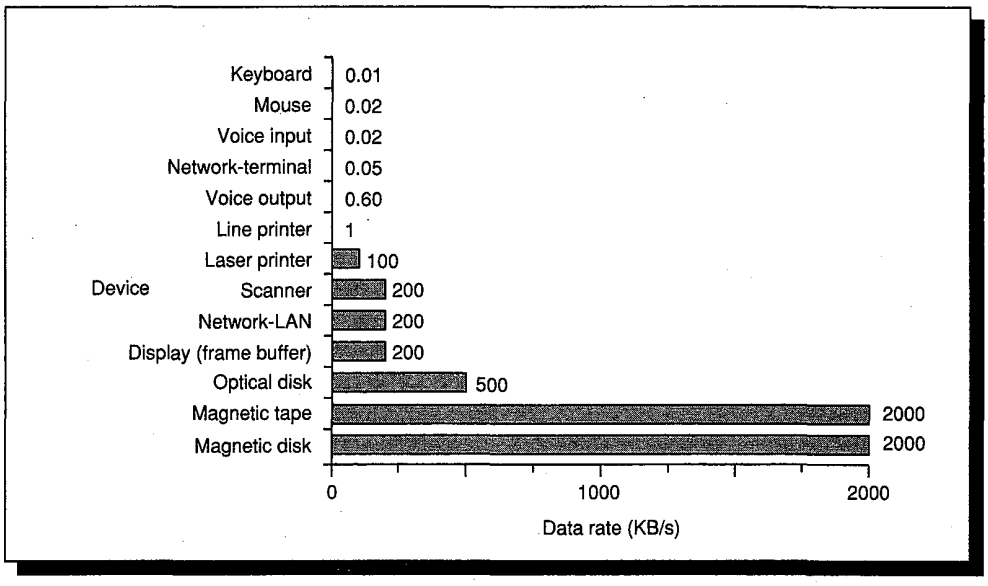

FIGURE 9.13 I/O devices sorted from lowest data rate to highest. The data rate for the graphics display is from the CPU to the frame buffer because the CPU isn't involved in the transfer from the frame buffer to the display (see Graphics Displays subsection below).

As descriptions of magnetic disks can be found in countless books, we will only list the key characteristics with the terms illustrated in Figure 9.14. A magnetic disk consists of a collection of platters (1 to 20), rotating on a spindle at about 3600 revolutions per minute (RPM). These platters are metal disks covered with magnetic recording material on both sides. Disk diameters vary by a factor of five, from 14 to 2.5 inches. Traditionally, the widest disks have the highest performance, and the smallest disks have the lowest cost per disk drive.

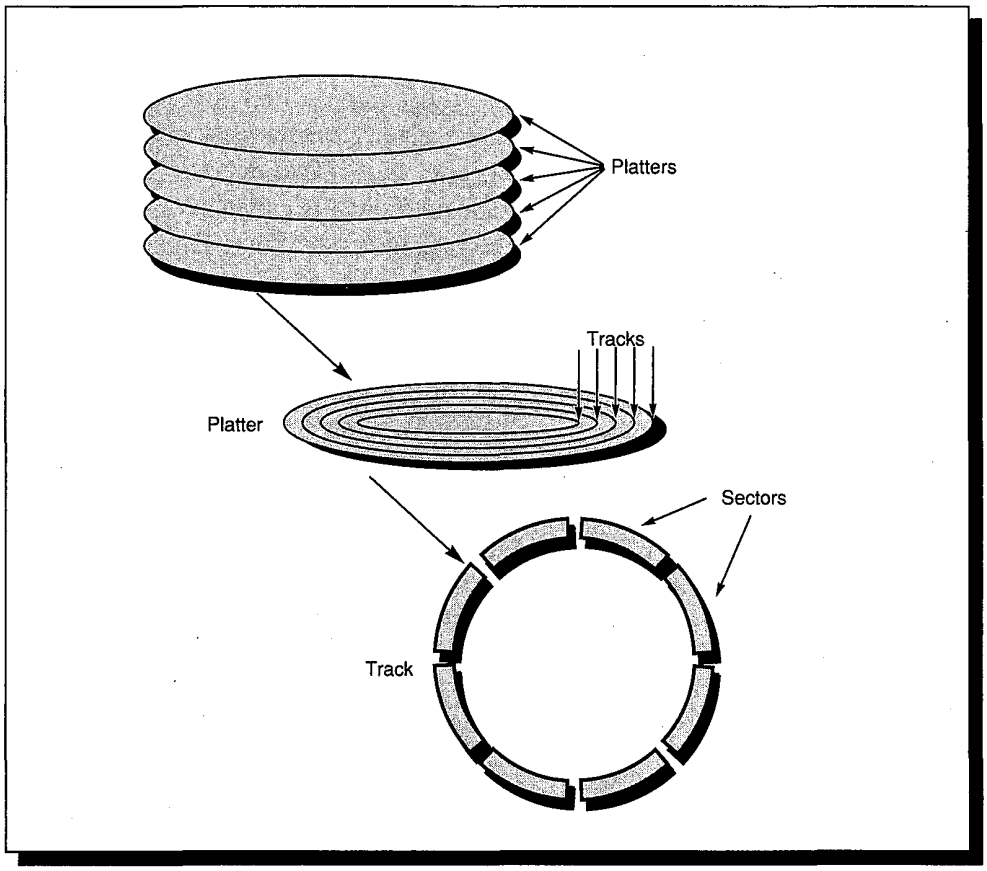

FIGURE 9.14 Disks are organized into platters, tracks, and sectors. Both sides of a platter are coated so that information can be stored on both surfaces.

Each disk surface is divided into concentric circles, designated *tracks*. There are typically 500 to 2000 tracks per surface. Each track in turn is divided into *sectors* that contain the information; each track might have 32 sectors. The sector is the smallest unit that can be read or written. The sequence recorded on the magnetic media is a sector number, a gap, the information for that sector including error correction code, a gap, the sector number of the next sector, and so on. Traditionally all tracks have the same number of sectors; the outer tracks, which are longer, record information at a lower density than the inner tracks. Recording more sectors on the outer tracks than on the inner tracks, called

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 547

constant bit density, is becoming more widespread with the advent of intelligent interface standards such as SCSI (see Section 9.5). IBM mainframe disks allow users to select the size of the sectors, while almost all other systems fix the size of the sector.

To read and write information into a sector, a movable *arm* containing a *read/write head* is located over each surface. Bits are recorded using a run-length limited code, which improves the recording density of the magnetic media. The arms for each surface are connected together and move in conjunction, so that every arm is over the same track of every surface. The term *cylinder* is used to refer to all the tracks under the arms at a given point on all surfaces.

To read or write a sector, the disk controller sends a command to move the arm over the proper track. This operation is called a *seek*, and the time to move the arm to the desired track is called *seek time*. Average seek time is the subject of considerable misunderstanding. Disk manufacturers report minimum seek time, maximum seek time, and average seek time in the manuals. The first two are easy to measure, but average was open to wide interpretation. The industry decided to calculate average seek time as the sum of the time for all possible seeks divided by the number of possible seeks. Average seek times are advertised to be 12 ms to 20 ms, but depending on the application and operating system the actual average seek time may be only 25% to 33% of the advertised number, due to locality of disk references. Section 9.10 has a detailed example.

The time for the requested sector to rotate under the head is the *rotation latency* or *rotational delay*. Most disks rotate at 3600 RPM, and an average latency to the desired information is halfway around the disk; the average rotation time for most disks is therefore

$$\text{Average rotation time} = \frac{0.5}{3600 \text{ RPM}} = 0.0083 \text{ sec} = 8.3 \text{ ms}$$

The next component of a disk access, *transfer time,* is the time to transfer a block of bits, typically a sector, under the read-write head. This is a function of the block size, rotation speed, recording density of a track, and speed of the electronics connecting disk to computer. Transfer rates in 1990 are typically 1 to 4 MB per second.

In addition to the disk drive, there is usually also a device called a *disk controller*. Between the disk controller and main memory is a hierarchy of controllers and data paths, whose complexity varies with the cost of the computer (see Section 9.9). Since the transfer time is often a small portion of a full disk access, the controller in higher performance systems disconnects the data paths from the disks while they are seeking so that other disks can transfer their data to memory.

Thus, the final component of disk-access time is *controller time*, which is the overhead the controller imposes in performing an I/O access. When referring to performance of a disk in a computer system, the time spent waiting for a disk to become free (*queueing delay*) is added to this time.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 548

Example

What is the average time to read or write a 512-byte sector for a typical disk today? The advertised average seek time is 20 ms, the transfer rate is 1MB/sec, and the controller overhead is 2 ms. Assume the disk is idle so that there is no queuing delay.

Answer

Average disk access is equal to average seek time + average rotational delay + transfer time + controller overhead. Using the calculated, average seek time, the answer is

$$20 \text{ ms} + 8.3 \text{ ms} + \frac{0.5 \text{ KB}}{1.0 \text{ MB/sec}} + 2 \text{ ms} = 20 + 8.3 + 0.5 + 2 = 30.8 \text{ ms}$$

Assuming the measured, average seek time is 25% of the calculated number, the answer is

$$5 \text{ ms} + 8.3 \text{ ms} + 0.5 \text{ ms} + 2 \text{ ms} = 15.8 \text{ ms}$$

Figure 9.15 shows characteristics of magnetic disks for four manufacturers. Large-diameter drives have many more megabytes to amortize the cost of electronics, so the traditional wisdom was that they had the lowest cost per megabyte. But this advantage is offset for the small drives by the much higher sales volume, which lowers manufacturing costs: 1990 OEM prices are $2 to $3

Characteristics	IBM 3380	Fujitsu M2361A	Imprimis Wren IV	Conner CP3100
Disk diameter (inches)	14	10.5	5.25	3.5
Formatted data capacity (MB)	7500	600	344	100
MTTF (hours)	52,000	20,000	40,000	30,000
Number of arms/box	4	1	1	1
Maximum I/Os/second/arm	50	40	35	30
Typical I/Os/second/arm	30	24	28	20
Maximum I/Os/second/box	200	40	35	30
Typical I/Os/second/box	120	24	28	20
Transfer rate (MB/sec)	3	2.5	1.5	1
Power/box (W)	1,650	640	35	10
MB/W	1.1	0.9	9.8	10.0
Volume (cu. ft.)	24	3.4	0.1	.03
MB/cu. ft.	310	180	3440	3330

FIGURE 9.15 Characteristics of magnetic disks from four manufacturers. Comparison of IBM 3380 disk model AK4 for mainframe computers, Fujitsu M2361A "Super Eagle" disk for minicomputers, Imprimis Wren IV disk for workstations, and Conner Peripherals CP3100 disk for personal computers. Maximum I/Os/second signifies maximum number of average seeks and average rotates for a single sector access. (Table from Katz, Patterson, and Gibson [1990].)

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 549

per megabyte, almost independent of width. The small drives also have advantages in power and volume. The price of a megabyte of disk storage in 1990 is 10 to 30 times cheaper than the price of a megabyte of DRAM in a system.

The Future of Magnetic Disks

The disk industry has concentrated on improving the capacity of disks. Improvement in capacity is customarily expressed as *areal density*, measured in bits per square inch:

$$\text{Areal density} = \frac{\text{Tracks}}{\text{Inch}} \text{ on a disk surface} * \frac{\text{Bits}}{\text{Inch}} \text{ on a track}$$

Areal density can be predicted according to the *maximum areal density* (MAD) formula:

$$\text{MAD} = 10^{(\text{year}-1971)/10} \text{ million bits per square inch}$$

Thus, storage density improves by a factor of 10 every decade, doubling density every three years.

Cost per megabyte has dropped consistently at 20% to 25% per year, with smaller drives playing the larger role in this improvement. Because it is easier to

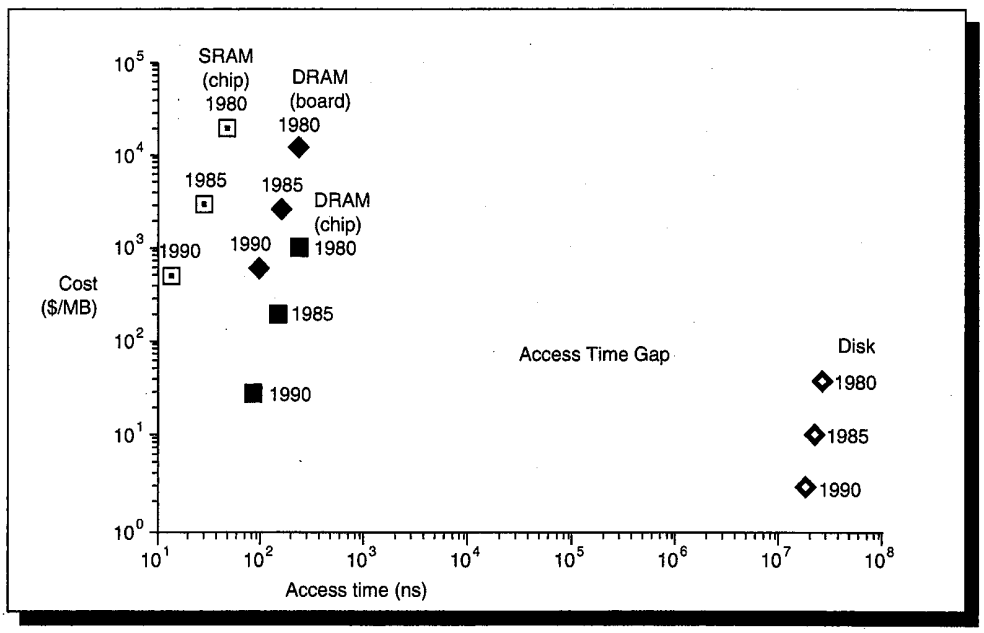

FIGURE 9.16 Cost versus access time for SRAM, DRAM, and magnetic disk in 1980, 1985, and 1990. (Note the difference in cost between a DRAM chip and DRAM chips packaged on a board and ready to plug into a computer.) The two-order-of-magnitude gap in cost and five-order-of-magnitude gap in access times between semiconductor memory and rotating magnetic disk has inspired a host of competing technologies to try to fill it. So far, such attempts have been made obsolete before production by improvements in magnetic disks, DRAMs, or both.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 550

spin the smaller mass, smaller diameter disks save power as well as volume. Smaller drives also have fewer cylinders so the seek distances are shorter. In 1990, 5.25-inch or 3.5-inch drives are probably the leading technology, while the future may see even smaller drives. We can expect significant savings in volume and power, but little in speed. Increasing density (bits per inch on a track) has improved transfer times, and there has been some small improvement in seek speed. Rotation speeds have been steady at 3600 RPM for a decade, but some manufacturers plan to go to 5400 RPM in the early 1990s.

As mentioned earlier, magnetic disks have been challenged many times for supremacy of secondary storage. One reason has been the fabled *Access Time Gap* as shown in Figure 9.16. Many a scientist has tried to invent a technology to fill that gap. Let's look at some of the recent attempts.

Using DRAMs as Disks

A current challenger to disks for dominance of secondary storage is *solid state disks* (SSDs), built from DRAMs with a battery to make the system nonvolatile; and *expanded storage* (ES), a large memory that allows only block transfers to or from main memory. ES acts like a software-controlled cache (the CPU stalls during the block transfer) while SSD involves the operating system just like a transfer from magnetic disks. The advantages of SSD and ES are trivial seek times, higher potential transfer rate, and possibly higher reliability. Unlike just a larger main memory, SSDs and ESs are autonomous: They require special commands to access their storage, and thus are "safe" from some software errors that write over main memory. The block-access nature of SSD and ES allows error correction to be spread over more words, which means lower cost or greater error recovery. For example, IBM's ES uses the greater error recovery to allow it to be constructed from less reliable (and less expensive) DRAMs without sacrificing product availability. SSDs, unlike main memory and ES, may be shared by multiple CPUs because they function as separate units. Placing DRAMs in an I/O device rather than memory is also one way to get around the address-space limits of the current 32-bit computers. The disadvantage of SSD and ES is cost, which is at least ten times per megabyte the cost of magnetic disks.

Optical Disks

Another challenger to magnetic disks is *optical compact disks* or CDs. The *CD/ROM* is removable and inexpensive to manufacture, but it is a read-only media. The newer CD/writable is also removable, but has a high cost per megabyte and low performance. A common misperception about *write-once optical disks* is that once they are written, the information cannot be destroyed; in fact, write once means one reliable write and then a "fuzzy" bitwise ORing of the previous and new data.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 551

So far, magnetic disk challengers have never had a product to market at the right time. By the time a new product ships, disks have made advances as predicted by MAD formula, and costs have dropped accordingly. Optical disks, however, may have the potential to compete with new tape technologies for archival storage.

Disk Arrays

One other future candidate for optimizing storage is not a new technology, but a new organization of disk storage—arrays of small and inexpensive disks. The argument for arrays is that since price per megabyte is independent of disk size, potential throughput can be increased by having many disk drives and, hence, many disk arms. Simply spreading data over multiple disks automatically forces accesses to several disks. (While arrays improve throughput, latency is not necessarily improved.) The drawback to arrays is that with more devices, reliability drops: N devices generally have $1/N$ the reliability of a single device.

Reliability and Availability

This brings us to two terms that are often confused—reliability and availability. The term reliability is commonly used incorrectly to mean availability; if something breaks, but the user can still use the system, it seems as if the system still "works," and hence it seems more reliable. Here is the proper distinction:

Reliability—is anything broken?

Availability—is the system still available to the user?

Adding hardware can therefore improve availability (for example, ECC on memory), but it cannot improve reliability (the DRAM is still broken). Reliability can only be improved by bettering environmental conditions, by building from more reliable components, or by building with fewer components. Another term, *data integrity*, refers to always reporting when information is lost when a failure occurs; this is very important to some applications.

So, while a disk array can never be more reliable than a smaller number of larger disks when each disk has the same failure rate, availability can be improved by adding redundant disks. That is, if a single disk fails, the lost information can be reconstructed from redundant information. The only danger is in getting another disk failure between the time a disk fails and the time it is replaced (termed *mean time to repair* or MTTR). Since the *mean time to failure* (MTTF) of disks is three to five years, and the MTTR is measured in hours, redundancy can make the availability of 100 disks much higher than that of a single disk.

Since disk failures are self-identifying, information can be reconstructed from just parity: The good disks plus the parity disk can be used to calculate the information that is on the failed disk. Hence, the cost of higher availability is

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 552

$1/N$, where N is the number of disks protected by parity. Just as direct-mapped associative placement in caches can be considered a special case of set-associative placement (see Section 8.4), the *mirroring* or *shadowing* of disks can be considered the special case of one data disk and one parity disk (N=1). Parity can be accomplished by duplicating the data, so mirrored disks have the advantage of simplifying parity calculation. Duplicating data also means that the controller can improve read performance by reading from the disk of the pair that has the shortest seek distance, although this optimization is at the cost of write performance because the arms of the pair of disks are no longer always over the same track. Of course, the redundancy of $N = 1$ has the highest overhead for increasing disk availability.

The higher throughput, measured either as megabytes per second or as I/Os per second, and the ability to recover from failures make disk arrays attractive. When combined with the advantages of smaller volume and lower power of small-diameter drives, redundant arrays of small or inexpensive drives may play a larger role in future disk systems. The current drawback is the added complexity of a controller for disk arrays.

Graphics Displays

Through computer displays I have landed an airplane on the deck of a moving carrier, observed a nuclear particle hit a potential well, flown in a rocket at nearly the speed of light and watched a computer reveal its innermost workings.

Ivan Sutherland (the "father" of computer graphics), quoted in "Computer Software for Graphics," *Scientific American* (1984)

While magnetic disks may dominate throughput and cost of I/O devices, the most fascinating I/O device is the graphics display. Based on television technology, a *raster cathode ray tube* (CRT) *display* scans an image out one line at a time, 30 to 60 times per second. At this *refresh rate* the human eye doesn't notice a "flicker" on the screen. The image is composed of a matrix of picture elements, or *pixels*, which can be represented as a matrix of bits, called a *bit map*. Depending on size of screen and resolution, the display matrix consists of 340*512 to 1560*1280 pixels. For black and white displays, often 0 is black and 1 is white. For displays that support over 100 different shades of black and white, sometimes called *gray-scale* displays, 8 bits per pixel are required. A color display might use 8 bits for each of the three primary colors (red, blue, and green), for 24 bits per pixel.

The hardware support for graphics consists mainly of a *raster refresh buffer*, or *frame buffer*, to store the bit map. The image to be represented on screen is stored into the frame buffer, and the bit pattern per pixel is read out to the graphics display at the refresh rate. Figure 9.17 (page 522) shows a frame buffer with four bits per pixel and Figure 9.18 (page 522) shows how the buffer is connected to the bus.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 553

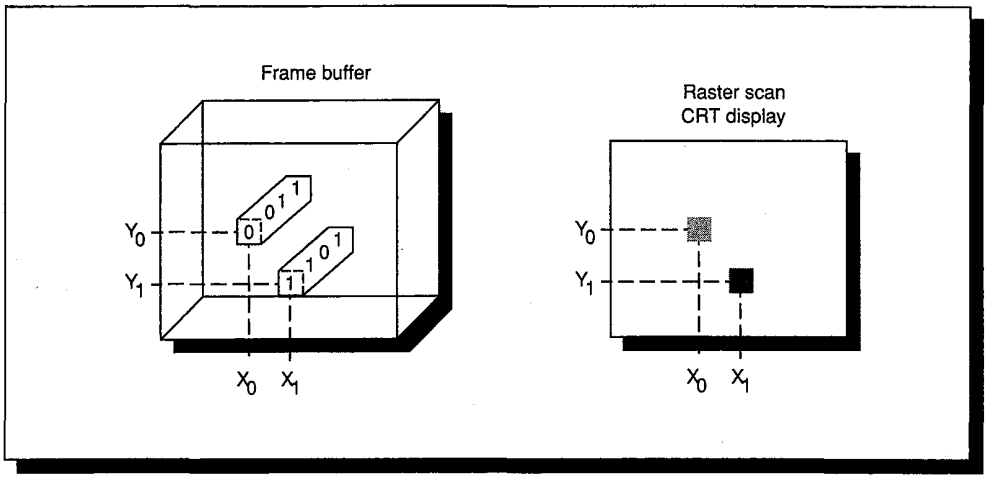

FIGURE 9.17 Each coordinate in the frame buffer on the left determines the shade of the corresponding coordinate for the raster scan CRT display on the right. Pixel (x_0,y_0) contains the bit pattern 0011, which is a lighter shade of gray on the screen than the bit pattern 1101 in pixel (x_1,y_1).

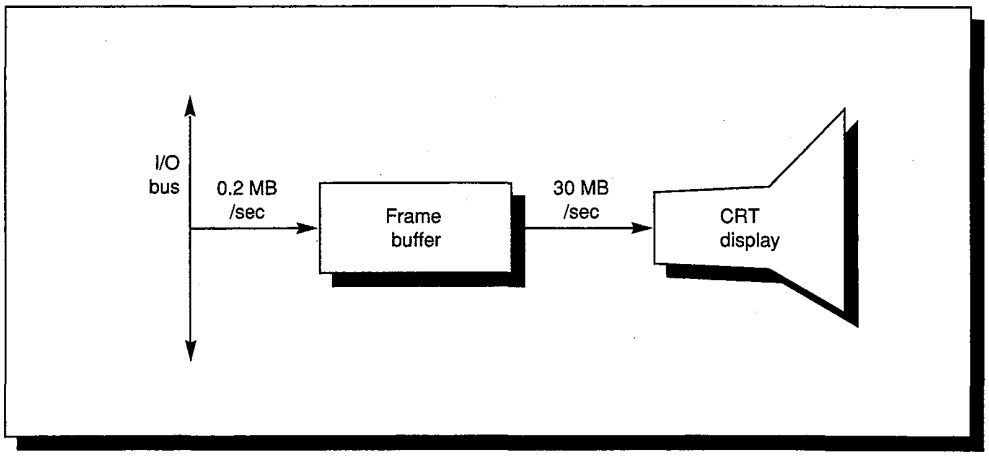

FIGURE 9.18 The frame buffer is connected to both the I/O bus and the display. Because of the high data rate from the buffer to the display, the frame buffer is frequently dual ported.

The goal of the bit map is to faithfully represent what is on the screen. As the computer switches from one image to another, the screen may look "splotchy" during the change. Here are two ways of dealing with this:

- Change the frame buffer only during the "vertical blanking interval." This is the time the gun in the raster CRT display takes to go back to the upper-left-hand corner before starting to paint the pixels of the next image. This takes 1 to 2 ms of every 16 ms at the 60-Hz refresh rate each time the screen is painted.

- If the vertical blanking interval is not long enough, the frame buffer can be double buffered, so that one is read while the other is being written. This way, images in sequence (as in animation) are drawn in alternate frame buffers. Double buffering, of course, doubles the cost of the memory in the frame buffer.

From the point of view of the CPU, graphics is logically output only. But the frame buffer is capable of being read as well as written, permitting operations to be performed directly on the screen images. These operations are called *bit blts*, for bit block transfer. Bit blts are commonly used for operations such as moving a window or changing the shape of the cursor. A current debate in graphics architecture is whether reading the frame buffer is limited to the operating system or should user programs be able to read it as well.

Cost of Computer Graphics

The CRT monitor itself is based on television technology and is sensitive to consumer demand. Today prices vary from $100 for a black-and-white monitor to $15,000 for a large studio color monitor, not including memory. The amount of memory in a frame buffer depends directly on the size of the screen and the bits per pixel:

$$340*512*1 \text{ bits} = 21.5 \text{ KB}$$

$$1280*1024*24 \text{ bits} = 3840 \text{ KB}$$

(By the way, this bottom dimension is the proposed size for high-definition television.) Note that the memory cost is doubled if double buffering is used.

To reduce costs of a color frame buffer, many systems use a two-level representation that takes advantage of the fact that few pictures need the full pallet of possible colors (see Figure 9.19 on page 524).

The intermediate level contains the full color width of, say, 24 bits and a large collection of the possible colors that can appear on the screen—256 different colors, for example. While this collection is large, it is still much smaller than 2^{24}. This intermediary table has been variously named a *color map*, *color table*, or *video look-up table*. Each pixel need have only enough bits to indicate a color in the color map. As a simple example, Figure 9.19 uses a 4-word color map, which means the frame buffer needs only 2 bits per pixel. The savings for a full-sized color display with a 256-color map is

$$1280*1024*24 - (1280*1024*8 + 256*24)$$

$$= 3{,}840 \text{ KB} - (1280 \text{ KB} + .75 \text{ KB}) \approx 2560 \text{ KB}$$

This amounts to a threefold reduction in memory size. In 1990 a 256- by 24-bit color map and an analog interface to a color CRT fit in a single chip.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 555

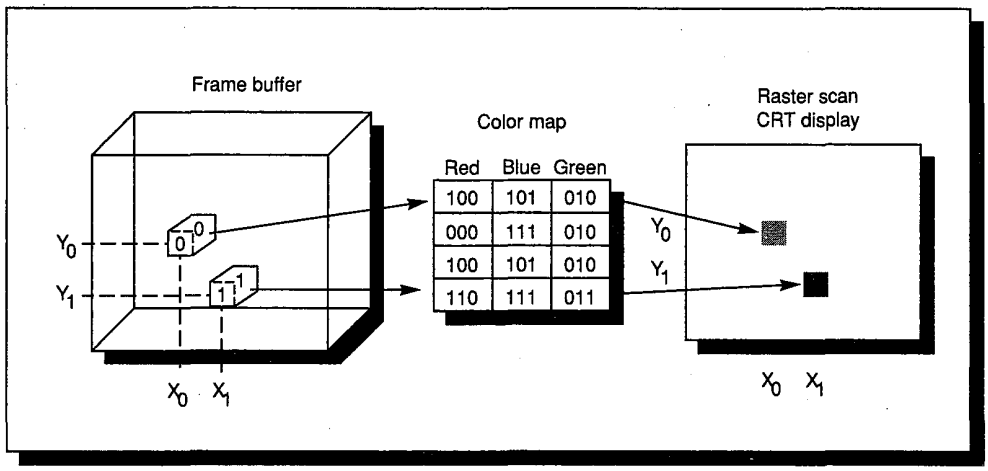

FIGURE 9.19 An example of a color map to reduce the cost of the frame buffer. Suppose only nine bits per color are needed. Rather than store the full nine bits per pixel in the frame buffer, just enough bits per pixel are stored to index the table containing the unique colors in a picture. Only the color map has the nine bits for the colors in the display. Near photographic color pictures can be produced with about 125 colors using the right shades of the color spectrum; but at least 24 bits are needed to get the right shades! The color map is loaded by the application program, offering each picture its own palette of colors to chose from.

Performance Demands of Graphics Displays

The performance of graphics is determined by the frequency an application needs new images and by the quality of those images. The amount of information transferred from memory to the frame buffer depends on complexity of image, with a full color display requiring almost four megabytes. The transfer rate depends on the speed with which the image should be changed as well as the amount of information. Animation requires at least 15 changes per second for movement to appear smooth on a screen. For interactive graphics, the time to update the frame buffer measures the effectiveness of the application; for people to feel comfortable the total reaction time must be less than a second (see Figure 9.9, page 510). With a drawing system, the portion of the screen one is working on must change almost immediately, as human visual perception is on the order of 0.02 seconds. Figure 9.20 shows some sample graphics tasks and their performance requirements. Note that the frame buffer must have enough bandwidth to refresh the display **and** to allow the CPU to change the image being refreshed.

The high data rate—and the large market of graphics displays—has made a dual-ported DRAM chip popular. This chip has a serial I/O port and internal shift register that is connected to the display in a graphics application in addition to the traditional randomly addressed data port. This chip is so widely used in frame buffers that it is called a *video DRAM*.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 556

Graphics tasks	Bandwidth requirements
Text editor—Scrolling text in window means moving all bits in half the frame buffer about 10 times per second.	0.8 MB/sec
VLSI design—Moving a portion of the design means moving all bits in half of a color frame buffer in less than 0.1 second.	6.3 MB/sec
Television commercial—Showing movie-quality images means changing 24 times per second.	90.0 MB/sec
Visualization of scientific data—About the same as a television commercial.	90.0 MB/sec

FIGURE 9.20 Graphics tasks and their performance requirements. VLSI design uses 8 bits of color while the television commercial and visualization use 24 bits. Bandwidth is measured at the frame buffer.

Future Directions in Graphics Displays

It is safe to predict that people will want better pictures in the future. They will want, for example, more lines on a screen and more bits per inch on a line to make sharper images, more bits per color to make more colorful images, and more bandwidth to allow animation.

To simplify the display of three-dimensional images, a z dimension per pixel can be added to the x and y coordinates. It says where the pixel is located from the viewer along a z axis (e.g., into the CRT). A 3D image starts with z set to the furthest possible location from the viewer and the color set to the background color. To get a proper 3D perspective, the z coordinate stored with the pixel in the frame buffer is checked before placing a color in a pixel. If the new color is closer, the old color is replaced and the z coordinate is updated; if it is further away, the new color is discarded. This scheme is called a *z buffer* approach to *hidden surface elimination*. It adds at least 8 bits per pixel, plus the performance cost of reading and comparing before writing a pixel. The Silicon Graphics 4D series of graphics workstations uses 16 bits for the z dimension in its pixels, meaning objects are assigned a 16-bit number to show how close they are to the viewer.

The increasing number of bits per DRAM chip reduces the number of chips needed in the frame buffer, as well as the number of chips that can simultaneously transfer bits to the screen. This is why video DRAMS are so popular. As capacity increases, the serial ports of video DRAMs will have to become faster and wider to match the demands of future graphics systems.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 557

Networks

There is an old network saying: Bandwidth problems can be cured with money. Latency problems are harder because the speed of light is fixed—you can't bribe God.

David Clark, M.I.T.

Networks are the backbone of current computer systems; a new machine without an optional network interface would be ridiculed. By connecting computers electronically, networked computers have these advantages:

- *Communication*—Information is exchanged between computers at high speeds.
- *Resource sharing*—Rather than each machine having its own I/O devices, devices can be shared by computers on the network.
- *Nonlocal access*—By connecting I/O devices over long distances, users need not be near the computer they are using.

Figure 9.21 shows the characteristics of networks. These characteristics are illustrated below with three examples.

Distance	0.01 to 10,000 kilometers
Speed	0.001 MB/sec to 100 MB/sec
Topology	Bus, ring, star, tree
Shared lines	None (point-to-point) or shared (multidrop)

FIGURE 9.21 Range of network characteristics.

The RS232 standard provides a 0.3- to 19.2-Kbits-per-second *terminal network*. A central computer connects to many terminals over slow but cheap dedicated wires. These point-to-point connections form a star from the central computer, with each terminal ranging from 10 to 100 meters in distance from the computer.

The *local area network*, or LAN, is what is commonly meant today when people mention a network, and *Ethernet* is what most people mean when they mention a LAN. (Ethernet has in fact become such a common term that it is often used as a generic term for LAN.) The Ethernet is essentially a 10,000 Kbits-per-second bus that has no central control. Messages or *packets* are sent over the Ethernet in blocks that vary from 128 bytes to 1530 bytes and take 0.1 ms and 1.5 ms to send, respectively. Since there is no central control, all nodes "listen" to see if there is a message for that node. Without a central arbiter to decide who gets the bus, a computer first listens to make sure it doesn't send a message while another message is on the network. If the network is idle the node tries to send. Of course, some other node may decide to send at the same instant. Luckily, the computer can detect any resulting collisions by listening to what is

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 558

sent. (Mixed messages will sound like garbage.) To avoid repeated head-on collisions, each node whose packet was trashed backs off a random time before resending. If Ethernets do not have high utilization, this simple approach to arbitration works well. Many LANs become overloaded through poor capacity planning, and response time and throughput can degrade rapidly at higher utilization.

The success of LANs has led to multiples of them at a single site. Connecting computers to separate Ethernets becomes necessary at a certain point because there is a limit to the number of nodes that can be active on a bus if effective communication speeds are to be achieved; one limit is 1024 nodes per Ethernet. There is also a physical limit to the distance of an Ethernet, usually about 1 kilometer. To allow Ethernets to work together, two kinds of devices have been created:

- A *bridge* connects two Ethernets. There are still two independent buses that can simultaneously send messages, but the bridge acts as a filter, allowing only those messages from nodes on one bus to nodes on the other bus to cross over the bridge.
- A *gateway* typically connects several Ethernets. It receives a message, looks up the destination address in a table, and then routes the message over the appropriate network to the proper node. This *routing table* can be changed during execution to reflect the state of the networks. Some use the term *router* instead of gateway since it is closer to the function performed.

When Ethernets are connected together with gateways they form an *Internet.*

Long-haul networks cover distances of 10 to 10,000 kilometers. The first and most famous long-haul network was the ARPANET (named after its funding agency, the Advanced Research Projects Agency of the U.S. government). It transferred at 50 Kbits per second and used point-to-point dedicated lines leased from telephone companies. The host computer talked to an *interface message processor* (IMP), which communicated over the telephone lines. The IMP took information and broke it into 1-Kbit packets. At each hop the packet was stored and then forwarded to the proper IMP according to the address in the packet. The destination IMP reassembled the packets into a message and then gave it to the host. *Fragmentation and reassembly*, as it was called, was done to reduce the latency due to the *store and forward delay*. Most networks today use this *packet switched* approach, where packets are individually routed from source to destination. Figure 9.22 (page 528) summarizes the performance, distance, and costs of these various networks.

While these networks have been presented here as alternatives, a computer system is really a hierarchy of networks, as Figure 9.23 (page 528) shows. To deal with this hierarchy of networks connecting machines that communicate differently, there must be a standard software interface to handle messages. These are called *protocols*, and are typically layered to interface with different levels of software in computer systems. The overhead of these protocols can eat up a significant portion of the network bandwidth.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 559

Just as with disks in Figure 9.6 (page 507), there is a tradeoff of latency and throughput in networks. Small messages give the lowest latency in most networks, but they also result in lower network bandwidth; similarly, a network can achieve higher bandwidth at the cost of longer latency.

Network	Performance (Kbits / sec)	Distance (km)	Cable cost	Connect to network cost	Connector to computer cost
RS232	19	0.1	$0.25 /foot	$1–$5 /connector	$5 /serial port chip
Ethernet	10,000	1	$1–$5 /foot	$100 /transceiver	$50 /Ethernet interface chip
ARPANET	50	10,000	$10,000 /month	$50,000–$100,000/ IMP	$5,000–$10,000 /IMP connection

FIGURE 9.22 The performance, maximum distance, and costs of three example networks. An Internet is simply multiple Ethernets and a bridge, which costs about $2,000 to $5,000, or a gateway, which costs about $20,000 to $50,000.

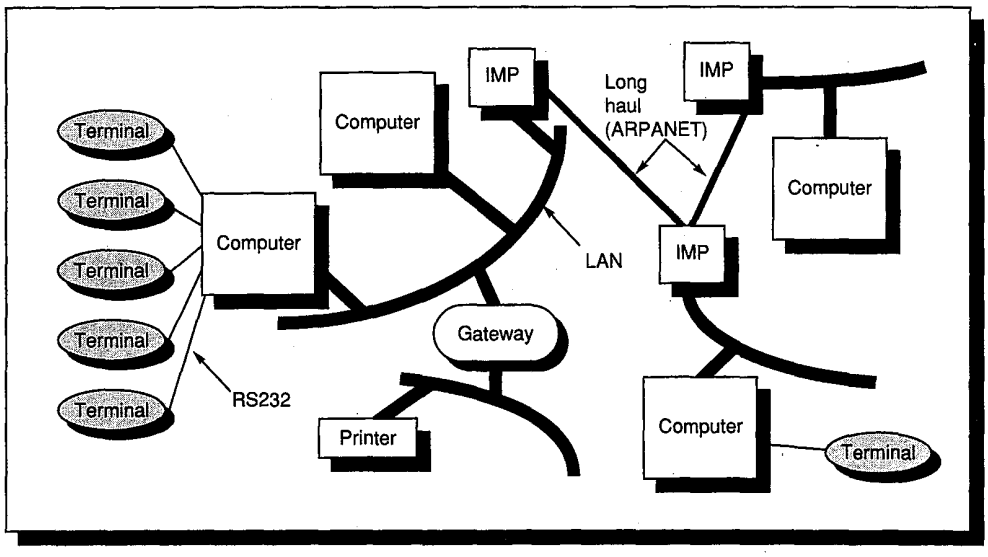

FIGURE 9.23 A computer system today participates in a hierarchy of networks. Ideally, the user is not aware of what network is being used in performing tasks. The gateway routes packets to a particular network, a network routes packets to a particular host computer, and the host computer routes packets to a particular process.

9.5 Buses—Connecting I/O Devices to CPU/Memory

In a computer system, the various subsystems must have interfaces to one another; for instance, the memory and CPU need to communicate, as well as the CPU and I/O devices. This is commonly done with a *bus*. The bus serves as a

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 560

shared communication link between the subsystems. The two major advantages of the bus organization are low cost and versatility. By defining a single interconnection scheme, new devices can easily be added, and peripherals may even be ported between computer systems that use a common bus. The cost is low, since a single set of wires is shared multiple ways.

The major disadvantage of a bus is that it creates a communication bottleneck, possibly limiting the maximum I/O throughput. When I/O must pass through a central bus this bandwidth limitation is as real as—and sometimes more severe than—memory bandwidth. In commercial systems, where I/O is very frequent, and in supercomputers, where the necessary I/O rates are very high because the CPU performance is high, designing a bus system capable of meeting the demands of the processor is a major challenge.

One reason bus design is so difficult is that the maximum bus speed is largely limited by physical factors: the length of the bus and the number of devices (and, hence, bus loading). These physical limits prevent arbitrary bus speedup. The desire for high I/O rates (low latency) and high I/O throughput can also lead to conflicting design requirements.

Buses are traditionally classified as *CPU–memory buses* or *I/O buses*. I/O buses may be lengthy, may have many types of devices connected to them, have a wide range in the data bandwidth of the devices connected to them (see Figure 9.1 on page 501), and normally follow a bus standard. CPU–memory buses, on the other hand, are short, generally high speed, and matched to the memory system to maximize memory–CPU bandwidth. During the design phase, the designer of a CPU–memory bus knows all the types of devices that must connect together, while the I/O bus designer must accept devices varying in latency and bandwidth capabilities. To lower costs, some computers have a single bus for both memory and I/O devices.

Let's consider a typical bus transaction. A *bus transaction* includes two parts: sending the address and receiving or sending the data. Bus transactions are usually defined by what they do to memory: A *read* transaction transfers data *from* memory (to either the CPU or an I/O device), and a *write* transaction writes data to the memory. In a read transaction, the address is first sent down the bus to the memory, together with the appropriate control signals indicating a read. The memory responds by returning the data on the bus with the appropriate control signals. A write transaction requires that the CPU or I/O device send both address and data and requires no return of data. Usually the CPU must wait between sending the address and receiving the data on a read, but the CPU often does not wait on writes.

The design of a bus presents several options, as Figure 9.24 (page 530) shows. Like the rest of the computer system, decisions will depend on cost and performance goals. The first three options in the figure are clear choices—separate address and data lines, wider data lines, and multiple-word transfers all give higher performance at more cost.

The next item in the table concerns the number of *bus masters*. These are devices that can initiate a read or write transaction; the CPU, for instance, is al-

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 561

ways a bus master. A bus has multiple masters when there are multiple CPUs or when I/O devices can initiate a bus transaction. If there are multiple masters, an arbitration scheme is required among the masters to decide who gets the bus next. Arbitration is often a fixed priority, as is the case with daisy-chained devices or an approximately fair scheme that randomly chooses which master gets the bus.

With multiple masters a bus can offer higher bandwidth by going to packets, as opposed to holding the bus for the full transaction. This technique is designated *split transactions*. (Some systems call this ability *connect/disconnect* or a *pipelined bus*.) The read transaction is broken into a read-request transaction that contains the address, and a memory-reply transaction that contains the data. Each transaction must now be tagged so that the CPU and memory can tell what is what. Split transactions make the bus available for other masters while the memory reads the words from the requested address. It also normally means that the CPU must arbitrate for the bus to send the data and the memory must arbitrate for the bus to return the data. Thus, a split-transaction bus has higher bandwidth, but it usually has higher latency than a bus that is held during the complete transaction.

The final item, clocking, concerns whether a bus is synchronous or asynchronous. If a bus is *synchronous* it includes a clock in the control lines and a fixed protocol for address and data relative to the clock. Since little or no logic is needed to decide what to do next, these buses can be both fast and inexpensive. However, they have two major disadvantages. Everything on the bus must run at the same clock rate, and because of clock-skew problems, synchronous buses cannot be long. CPU–memory buses are typically synchronous.

An *asynchronous* bus, on the other hand, is not clocked. Instead, self-timed, handshaking protocols are used between bus sender and receiver. This scheme makes it much easier to accommodate a wide variety of devices and to lengthen the bus without worrying about clock skew or synchronization problems. If a synchronous bus can be used, it is usually faster than an asynchronous bus because of the overhead of synchronizing the bus for each transaction. The choice of synchronous versus asynchronous bus has implications not only for data bandwidth but also for an I/O system's capacity in terms of physical

Option	High performance	Low cost
Bus width	Separate address and data lines	Multiplex address and data lines
Data width	Wider is faster (e.g., 32 bits)	Narrower is cheaper (e.g., 8 bits)
Transfer size	Multiple words has less bus overhead	Single-word transfer is simpler
Bus masters	Multiple (requires arbitration)	Single master (no arbitration)
Split transaction?	Yes—separate Request and Reply packets gets higher bandwidth (needs multiple masters)	No—continuous connection is cheaper and has lower latency
Clocking	Synchronous	Asynchronous

FIGURE 9.24 The main options for a bus. The advantage of separate address and data buses is primarily on writes.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 562

distance and number of devices that can be connected to the bus; asynchronous buses scale better with technological changes. I/O buses are typically asynchronous. Figure 9.25 suggests the relationship of when to use one over the other.

Bus Standards

The number and variety of I/O devices are not fixed on most computer systems, permitting customers to tailor computers to their needs. As the interface to which devices are connected, the I/O bus can also be considered an expansion bus for adding I/O devices over time. Standards that let the computer designer and I/O-device designer work independently, therefore, play a large role in determining the choice of buses. As long as both the computer-system designer and the I/O-device designer meet the requirements, any I/O device can connect to any computer. In fact, an I/O bus standard is the document that defines how to connect them.

Machines sometimes grow to be so popular that their I/O buses become de facto standards; examples are the PDP-11 Unibus and the IBM PC-AT Bus. Once many I/O devices have been built for the popular machine, other computer designers will build their I/O interface so that those devices can plug into their machines as well. Sometimes standards also come from an explicit standards effort on the part of I/O device makers. The *intelligent peripheral interface* (IPI)

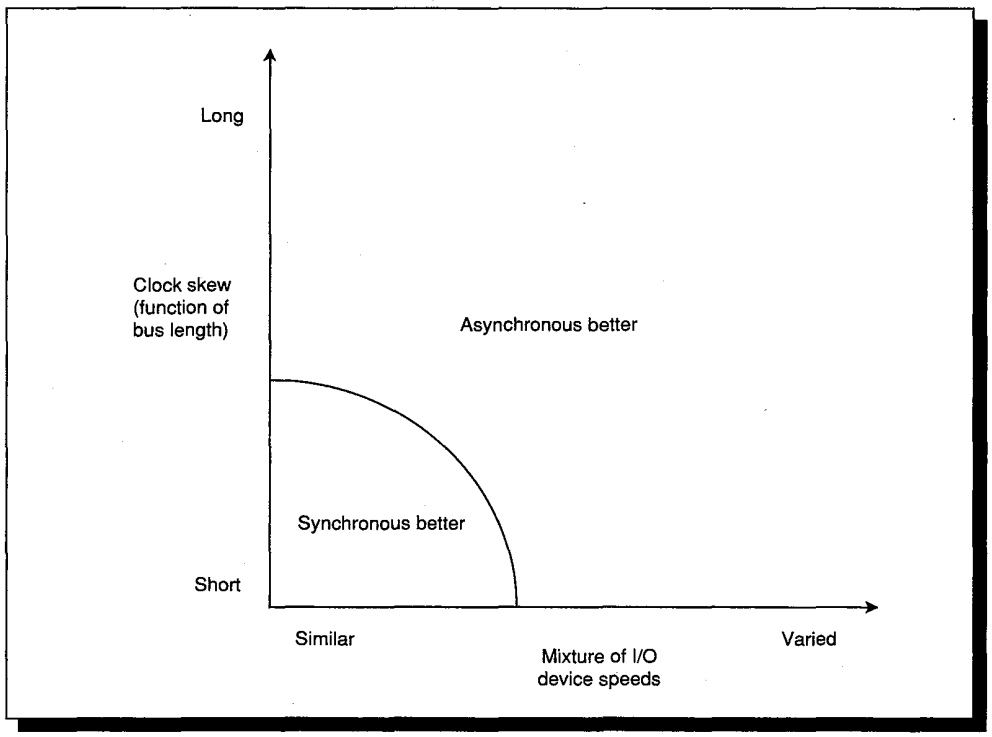

FIGURE 9.25 Preferred bus type as a function of length/clock skew and variation in I/O device speed. Synchronous is best when the distance is short and the I/O devices on the bus all transfer at similar speeds.

and Ethernet are examples of standards from cooperation of manufacturers. If standards are successful, they are eventually blessed by a sanctioning body like ANSI or IEEE. Occasionally, a bus standard comes top-down directly from a standards committee—the FutureBus is one example.

Figure 9.26 summarizes characteristics of several bus standards. Note that the bandwidth entries in the figure are not listed as single numbers for the CPU–memory buses (VME, FutureBus, and Multibus II). Because of the bus overhead, the size of the transfer affects bandwidth significantly. Since the bus usually transfers to or from memory, the speed of the memory also affects the bandwidth. For example, with infinite transfer size and infinitely fast (0 ns) memory, FutureBus is 240% faster than VME, but FutureBus is only about 20% faster than VME for single-word transfers from a 150-ns memory.

	VME bus	FutureBus	Multibus II	IPI	SCSI
Bus width (signals)	128	96	96	16	8
Address/data multiplexed?	Not multiplexed	Multiplexed	Multiplexed	N/A	N/A
Data width (primary)	16 to 32 bits	32 bits	32 bits	16 bits	8 bits
Transfer size	Single or multiple	Single or multiple	Single or multiple	Single or multiple	Single or multiple
Number of bus masters	Multiple	Multiple	Multiple	Single	Multiple
Split transaction?	No	Optional	Optional	Optional	Optional
Clocking	Asynchronous	Asynchronous	Synchronous	Asynchronous	Either
Bandwidth, 0-ns access memory, single word	25.0 MB/sec	37.0 MB/sec	20.0 MB/sec	25.0 MB/sec	5.0 MB/sec or 1.5 MB/sec
Bandwidth, 150-ns access memory, single word	12.9 MB/sec	15.5 MB/sec	10.0 MB/sec	25.0 MB/sec	5.0 MB/sec or 1.5 MB/sec
Bandwidth, 0-ns access memory, multiple words (infinite block length)	27.9 MB/sec	95.2 MB/sec	40.0 MB/sec	25.0 MB/sec	5.0 MB/sec or 1.5 MB/sec
Bandwidth, 150-ns access memory, multiple words (infinite block length)	13.6 MB/sec	20.8 MB/sec	13.3 MB/sec	25.0 MB/sec	5.0 MB/sec or 1.5 MB/sec
Maximum number of devices	21	20	21	8	7
Maximum bus length	0.5 meter	0.5 meter	0.5 meter	50 meters	25 meters
Standard	IEEE 1014	IEEE 896.1	ANSI/IEEE 1296	ANSI X3.129	ANSI X3.131

FIGURE 9.26 Information on five bus standards. The first three were defined originally as CPU–memory buses and the last two as I/O buses. For the CPU–memory buses the bandwidth calculations assume a fully loaded bus and are given for both single-word transfers and block transfers of unlimited length; measurements are shown both ignoring memory latency and assuming 150-ns access time. Bandwidth assumes the average distance of a transfer is one-third of the backplane length. (Data in the first three columns is from Borrill [1986].) The bandwidth for the I/O buses is given as their maximum data-transfer rate. The SCSI standard offers either asynchronous or synchronous I/O; the asynchronous version transfers at 1.5 MB/sec and the synchronous at 5 MB/sec.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 564

9.6 Interfacing to the CPU

Having described I/O devices and looked at some of the issues of the connecting bus, we are ready to discuss the CPU end of the interface. The first question is how the physical connection of the I/O bus should be made. The two choices are connecting it to memory or to the cache. In the following section we will discuss the pros and cons of connecting an I/O bus directly to the cache; in this section we examine the more usual case in which the I/O bus is connected to the main memory bus. Figure 9.27 shows a typical organization. In low-cost systems, the I/O bus **is** the memory bus; this means an I/O command on the bus could interfere with a CPU instruction fetch, for example.

Once the physical interface is chosen, the question becomes how does the CPU address an I/O device that it needs to send or receive data. The most common practice is called *memory-mapped* I/O. In this scheme, portions of the address space are assigned to I/O devices. Reads and writes to those addresses may cause data to be transferred; some portion of the I/O space may also be set aside for device control, so commands to the device are just accesses to those memory-mapped addresses. The alternative practice is to use dedicated I/O opcodes in the CPU. In this case, the CPU sends a signal that this address is for I/O devices. Examples of computers with I/O instructions are the Intel 80x86 and the IBM 370 computers. No matter which addressing scheme is selected, each I/O device has registers to provide status and control information. Either

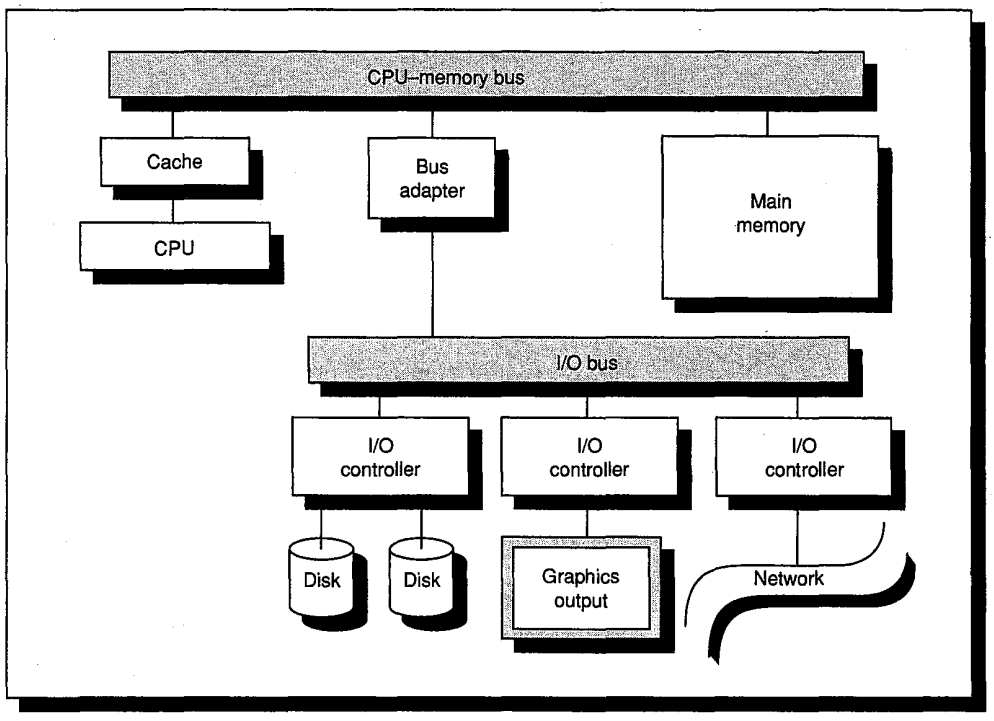

FIGURE 9.27 A typical interface of I/O devices and an I/O bus to the CPU–memory bus.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 565

through loads and stores in memory-mapped I/O or through special instructions, the CPU sets flags to determine the operation the I/O device will perform.

I/O is rarely a single operation. For example, the DEC LP11 line printer has two I/O device registers: one for status information and one for data to be printed. The status register contains a *done bit*, set by the printer when it has printed a character, and an *error bit*, indicating that the printer is jammed or out of paper. Each byte of data to be printed is put into the data register; the CPU must then wait until the printer sets the done bit before it can place another character in the buffer.

This simple interface, in which the CPU periodically checks status bits to see if it is time for the next I/O operation, is called *polling*. As one might expect, the fact that CPUs are so much faster than I/O devices means polling may waste a lot of CPU time. This was recognized long ago, leading to the invention of interrupts to notify the CPU when it is time to do something for the I/O device. *Interrupt-driven* I/O, used by most systems for at least some devices, allows the CPU to work on some other process while waiting on the I/O device. For example, the LP11 has a mode that allows it to interrupt the CPU whenever the done bit or error bit is set. In general-purpose applications, interrupt driven I/O is the key to multitasking operating systems and good response times.

The drawback to interrupts is the operating system overhead on each event. In real-time applications with hundreds of I/O events per second, this overhead can be intolerable. One hybrid solution for real-time systems is to use a clock to periodically interrupt the CPU, at which time the CPU polls all I/O devices.

Delegating I/O Responsibility from the CPU

Interrupt-driven I/O relieves the CPU from waiting for every I/O event, but there are still many CPU cycles spent in transferring data. Transferring a disk block of 2048 words, for instance, would require at least 2048 loads and 2048 stores, as well as the overhead for the interrupt. Since I/O events so often involve block transfers, *direct memory access* (DMA) hardware is added to many computer systems to allow transfers of numbers of words without intervention by the CPU.

DMA is a specialized processor that transfers data between memory and an I/O device, while the CPU goes on with other tasks. Thus, it is external to the CPU and must act as a master on the bus. The CPU first sets up the DMA registers, which contain a memory address and number of bytes to be transferred. Once the DMA transfer is complete, the controller interrupts the CPU. There may be multiple DMA devices in a computer system; for example, DMA is frequently part of the controller for an I/O device.

Increasing the intelligence of the DMA device can further unburden the CPU. Devices called *I/O processors*, (or *I/O controllers*, or *channel controllers*) operate from either fixed programs or from programs downloaded by the operating system. The operating system typically sets up a queue of *I/O control*

blocks that contain information such as data location (source and destination) and data size. The I/O processor then takes items from the queue, doing everything requested and sending a single interrupt when the task specified in the I/O control blocks is complete. Whereas the LP11 line printer would cause 4800 interrupts to print a 60-line by 80-character page, an I/O processor could save 4799 of those interrupts.

I/O processors can be compared to multiprocessors in that they facilitate several processes executing simultaneously in the computer system. I/O processors are less general than CPUs, however, since they have dedicated tasks, and thus parallelism is also much more limited. Also, an I/O processor doesn't normally change information, as a CPU does, but just moves information from one place to another.

9.7 Interfacing to an Operating System

In a manner analogous to the way compilers use an instruction set (see Section 3.7 of Chapter 3), operating systems control what I/O techniques implemented by the hardware will actually be used. For example, many I/O controllers used in early UNIX systems were 16-bit microprocessors. To avoid problems with 16-bit addresses in controllers, UNIX was changed to limit the maximum I/O transfer to 63 KB or less; at the time of this book's publication, that limit is still in effect. Thus, a new I/O controller designed to efficiently transfer 1-MB files would never see more than 63 KB at a time under UNIX, no matter how large the files.

Caches Cause Problems for Operating Systems—Stale Data

The prevalence of caches in computer systems has added to the responsibilities of the operating system. Caches imply the possibility of two copies of the data—one each for cache and main memory—while virtual memory can result in three copies—for cache, memory and disk. This brings up the possibility of *stale data*: the CPU or I/O system could modify one copy without updating the other copies (see Section 8.8 in Chapter 8). Either the operating system or the hardware must make sure that the CPU reads the most recently input data and that I/O outputs the correct data, in the presence of caches and virtual memory. Whether the stale-data problem arises depends in part on where the I/O is connected to the computer. If it is connected to the CPU cache, as shown in Figure 9.28 (page 536), there is no stale-data problem; all I/O devices and the CPU see the most accurate version in the cache, and existing mechanisms in the memory hierarchy ensure that other copies of the data will be updated. The side effect is lost CPU performance, since I/O will replace blocks in the cache with data that are unlikely to be needed by the process running in the CPU at the time of the

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 567

transfer. In other words, all I/O data goes through the cache but little of it is referenced. This arrangement also requires arbitration between CPU and I/O to decide who accesses the cache. If I/O is connected to memory, as in Figure 9.27 (page 533), then it doesn't interfere with CPU, provided the CPU has a cache. In this situation, however, the stale-data problem occurs. Alternatively, I/O can just invalidate data—either all data that might match (no tag check) or only data that matches.

There are two parts to the stale-data problem:

1. The I/O system sees stale data on output because memory is not up to date.
2 The CPU sees stale data in the cache on input after the I/O system has updated memory.

The first dilemma is how to output correct data if there is a cache and I/O is connected to memory. A write-through cache solves this by ensuring that memory will have the same data as the cache. A write-back cache requires the operating system to flush output addresses to make sure they are not in the cache. This takes time, even if the data is not in the cache, since address checks are sequential. Alternatively, the hardware can check cache tags during output to see if they are in a write-back cache, and only interact with the cache if the output tries to read data that is in the cache.

The second problem is ensuring that the cache won't have stale data after input. The operating system can guarantee that the input data area can't possibly

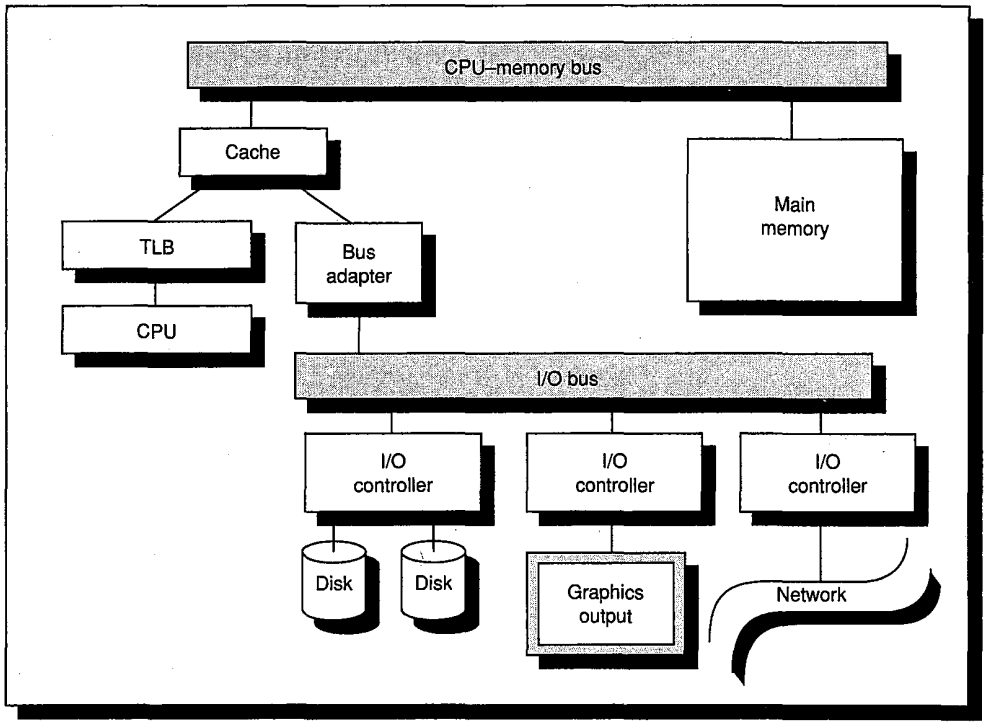

FIGURE 9.28 Example of I/O connected directly to the cache.

be in the cache. If it can't guarantee this, the operating system flushes input addresses to make sure they are not in the cache. Again, this takes time, whether or not the input addresses are in the cache. As before, extra hardware can be added to check tags during an input and invalidate the data if there is a conflict. These problems are basically the same as cache coherency in a multiprocessor, discussed in Section 8.8 of Chapter 8; I/O can be thought of as a second dedicated processor in a multiprocessor.

DMA and Virtual Memory

Given the use of virtual memory, there is the matter of whether DMA should transfer using virtual addresses or physical addresses. Here are some problems with DMA using physically mapped I/O:

- Transferring a buffer that is larger than one page will cause problems, since the pages in the buffer will not usually be mapped to sequential pages in physical memory.
- Suppose DMA is ongoing between memory and a frame buffer, and the operating system removes some of the pages from memory (or relocates them). The DMA would then be transferring data to or from the wrong page of memory.

One answer to these questions is *virtual DMA*. It allows the DMA to use virtual addresses that are mapped to physical addresses during the DMA. Thus, a buffer must be sequential in virtual memory but the pages can be scattered in physical memory. The operating system could update the address tables of a DMA if a process is moved using virtual DMA, or the operating system could "lock" the pages in memory until the DMA is complete. Figure 9.29 (page 538) shows address-translation registers added to the DMA device.

Caches Helping Operating Systems—File or Disk Caches

While the invention of caches made the life of the operating systems designer more difficult, operating systems designers' concern for performance led them to cache-like optimizations, using main memory as a "cache" for disk traffic to improve I/O performance. The impact of using main memory as a buffer or cache for file or disk accesses is demonstrated in Figure 9.30 (page 538). It shows the change in disk I/Os for a cacheless system measured as miss rate (see Section 8.2 in Chapter 8). File caches or disk caches change the number of disk I/Os and the mix of reads and writes; depending on cache size and write policy, between 50% to 70% of all disk accesses could become writes with such caches. Without file or disk caches, between 15% and 33% of all accesses are writes, depending on the environment.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 569

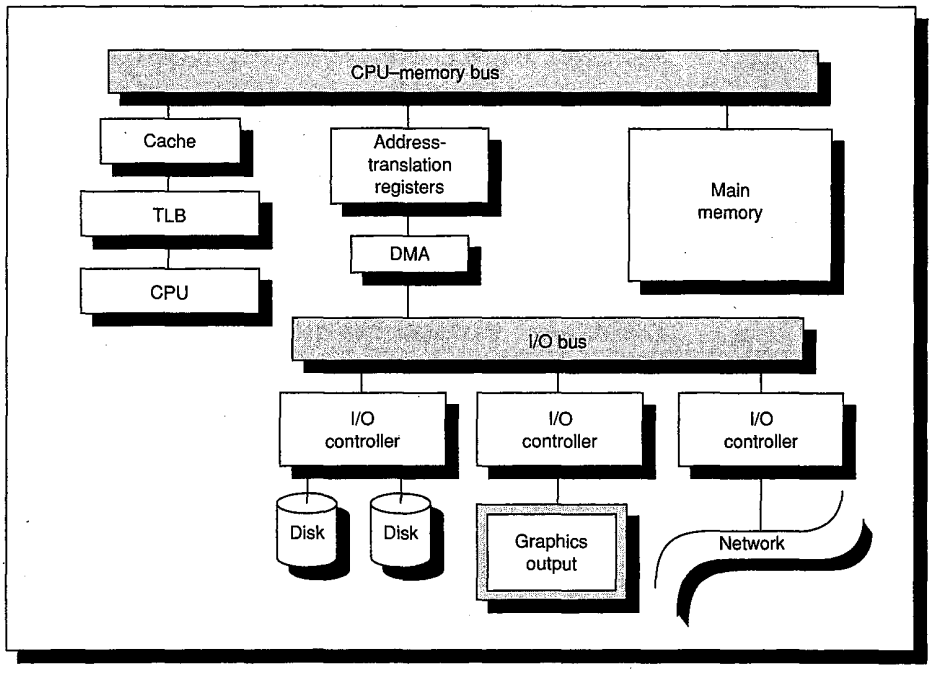

FIGURE 9.29 Virtual DMA requires a register for each page to be transferred in the DMA controller, showing the protection bits and the physical page corresponding to each virtual page.

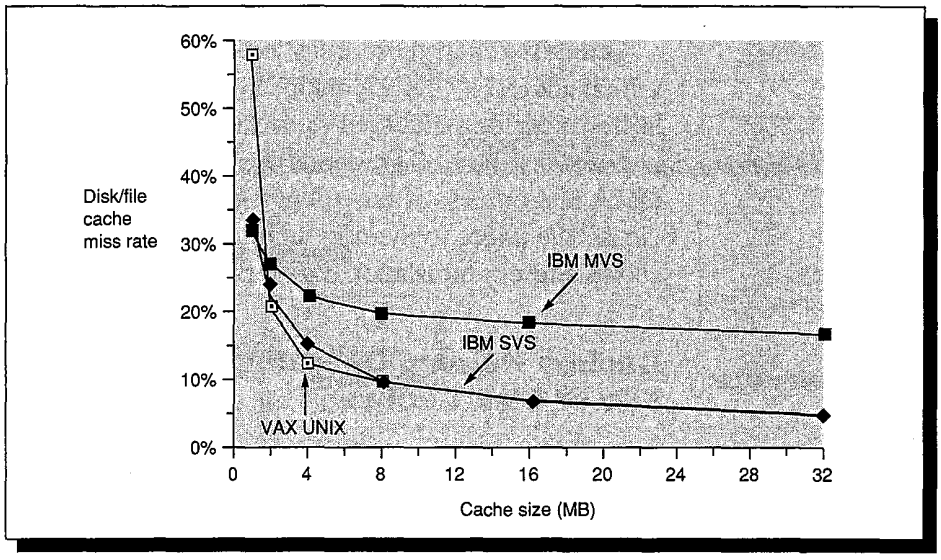

FIGURE 9.30 The effectiveness of a file cache or disk cache on reducing disk I/Os versus cache size. Ousterhout et al. [1985] collected the VAX UNIX data on VAX-11/785s with 8 MB to 16 MB of main memory, running 4.2 BSD UNIX using a 16-KB block size. Smith [1985] collected the IBM SVS and IBM MVS traces on IBM 370/168 using a one-track block size (which varied from 7294 bytes to 19254 bytes, depending on the disk). The difference between a file cache and a disk cache is that the file cache uses logical block numbers while a disk cache uses addresses that have been mapped to the physical sector and track on a disk. This difference is similar to the difference between a virtually addressed and a physically addressed cache (see Section 8.8 in Chapter 8).

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 570

9.8 Designing an I/O System

The art of I/O is finding a design that meets goals for cost and variety of devices while avoiding bottlenecks to I/O performance. This means that components must be balanced between main memory and the I/O device because performance—and hence effective cost/performance—can only be as good as the weakest link in the I/O chain. The architect must also plan for expansion so that customers can tailor the I/O to their applications. This expansibility, both in numbers and types of I/O devices, has its costs in longer backplanes, larger power supplies to support I/O devices, and larger cabinets.

In designing an I/O system, analyze performance, cost, and capacity using varying I/O connection schemes and different numbers of I/O devices of each type. Here is a series of six steps to follow in designing an I/O system. The answers in each step may be dictated by market requirements or simply by cost/performance goals.

1. List the different types of I/O devices to be connected to the machine, or a list of standard buses that the machine will support.
2. List the physical requirements for each I/O device. This includes volume, power, connectors, bus slots, expansion cabinets, and so on.
3. List the cost of each I/O device, including the portion of cost of any controller needed for this device.
4. Record the CPU resource demands of each I/O device. This should include:

   Clock cycles for instructions used to initiate an I/O, to support operation of an I/O device (such as handling interrupts), and complete I/O

   CPU clock stalls due to waiting for I/O to finish using the memory, bus, or cache

   CPU clock cycles to recover from an I/O activity, such as a cache flush
5. List the memory and I/O bus resource demands of each I/O device. Even when the CPU is not using memory, the bandwidth of main memory and the I/O bus are limited.
6. The final step is establishing performance of the different ways to organize these I/O devices. Performance can only be properly evaluated with simulation, though it may be estimated using queuing theory.

You then select the best organization, given your performance and cost goals.

Cost and performance goals affect the selection of the I/O scheme and physical design. Performance can be measured either as megabytes per second or I/Os per second, depending on the needs of the application. For high performance, the only limits should be speed of I/O devices, number of I/O devices, and speed of memory and CPU. For low cost, the only expenses should be those

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 571

for the I/O devices themselves and for cabling to the CPU. Cost/performance design, of course, tries for the best of both worlds.

To make these ideas clearer, let's go through several examples.

Example

First, let's look at the impact on the CPU of reading a disk page directly into the cache. Make the following assumptions:

Each page is 8 KB and the cache-block size is 16 bytes.

The addresses corresponding to the new page are **not** in the cache.

The CPU will not access any of the data in the new page.

90% of the blocks that were displaced from the cache will be read in again, and each will cause a miss.

The cache uses write back, and 50% of the blocks are dirty on average.

The I/O system buffers a full cache block before writing to the cache (this is called a *speed-matching buffer*, matching transfer bandwidth of the I/O system and memory).

The accesses and misses are spread uniformly to all cache blocks.

There is no other interference between the CPU and I/O for the cache slots.

There are 15,000 misses every one million clock cycles when there is **no** I/O.

The miss penalty is 15 clock cycles, plus 15 more cycles to write the block if it was dirty.

Assuming one page is brought in every one million clock cycles, what is the impact on performance?

Answer

Each page fills 8192/16 or 512 blocks. I/O transfers do not cause cache misses on their own because entire cache blocks are transferred. However, they do displace blocks already in the cache. If half of the displaced blocks are dirty it takes 256*15 clock cycles to write them back to memory. There are also misses from 90% of the blocks displaced in the cache because they are referenced later, adding another 90%*512, or 461 misses. Since this data was placed into the cache from the I/O system, all these blocks are dirty and will need to be written back when replaced. Thus, the total is 256*15 + 461*30 more clock cycles than the original 1,000,000 + 15,000*15. This turns into a 1% decrease in performance:

$$\frac{256*15 + 461*30}{1000000+15000*15} = \frac{17670}{1225000} = 0.014$$

Now let's look at the cost/performance of different I/O organizations. A simple way to perform this analysis is to look at maximum throughput assuming

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 572

that resources can be used at 100% of their maximum rate without side effects from interference. A later example takes a more realistic view.

Example

Given the following performance and cost information:

a 50-MIPS CPU costing $50,000

an 8-byte-wide memory with a 200-ns cycle time

80 MB/sec I/O bus with room for 20 SCSI buses and controllers

SCSI buses that can transfer 4 MB/sec and support up to 7 disks per bus (these are also called SCSI *strings*)

a $2500 SCSI controller that adds 2 milliseconds (ms) of overhead to perform a disk I/O

an operating system that uses 10,000 CPU instructions for a disk I/O

a choice of a large disk containing 4 GB or a small disk containing 1 GB, each costing $3 per MB

both disks rotate at 3600 RPM, have a 12-ms average seek time, and can transfer 2MB/sec

the storage capacity must be 100 GB, and

the average I/O size is 8 KB

Evaluate the cost per I/O per second (IOPS) of using small or large drives. Assume that every disk I/O requires an average seek and average rotational delay. Use the optimistic assumption that all devices can be used at 100% of capacity and that the workload is evenly divided between all disks.

Answer

I/O performance is limited by the weakest link in the chain, so we evaluate the maximum performance of each link in the I/O chain for each organization to determine the maximum performance of that organization.

Let's start by calculating the maximum number of IOPS for the CPU, main memory, and I/O bus. The CPU I/O performance is determined by the speed of the CPU and the number of instructions to perform a disk I/O:

$$\text{Maximum IOPS for CPU} = \frac{50 \text{ MIPS}}{10000 \text{ instructions per I/O}} = 5000$$

The maximum performance of the memory system is determined by the memory cycle time, the width of the memory, and the size of the I/O transfers:

$$\text{Maximum IOPS for main memory} = \frac{(1/200 \text{ ns})*8}{8 \text{ KB per I/O}} \approx 5000$$

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 573

The I/O bus maximum performance is limited by the bus bandwidth and the size of the I/O:

$$\text{Maximum IOPS for the I/O bus} = \frac{80 \text{ MB/sec}}{8 \text{ KB per I/O}} \approx 10000$$

Thus, no matter which disk is selected, the CPU and main memory limits the maximum performance to no more than 5000 IOPS.

Now its time to look at the performance of the next link in the I/O chain, the SCSI controllers. The time to transfer 8 KB over the SCSI bus is

$$\text{SCSI bus transfer time} = \frac{8 \text{ KB}}{4 \text{ MB/sec}} = 2 \text{ ms}$$

Adding the 2-ms SCSI controller overhead means 4 ms per I/O, making the maximum rate per controller

$$\text{Maximum IOPS per SCSI controller} = \frac{1}{4 \text{ ms}} = 250 \text{ IOPS}$$

All the organizations will use several controllers, so 250 IOPS is not the limit for the whole system.

The final link in the chain is the disks themselves. The time for an average disk I/O is

$$\text{I/O time} = 12 \text{ ms} + \frac{0.5}{3600 \text{ RPM}} + \frac{8 \text{ KB}}{2 \text{ MB/sec}} = 12+8.3+4 = 24.3 \text{ ms}$$

so the disk performance is

$$\text{Maximum IOPS (using average seeks) per disk} = \frac{1}{24.3 \text{ ms}} \approx 41 \text{ IOPS}$$

The number of disks in each organization depends on the size of each disk: 100 GB can be either 25 4-GB disks or 100 1-GB disks. The maximum number of I/Os for all the disks is:

$$\text{Maximum IOPS for 25 4-GB disks} = 25 * 41 = 1025$$

$$\text{Maximum IOPS for 100 1-GB disks} = 100 * 41 = 4100$$

Thus, provided there are enough SCSI strings, the disks become the new limit to maximum performance: 1025 IOPS for the 4-GB disks and 4100 for the 1-GB disks.

While we have determined the performance of each link of the I/O chain, we still have to determine how many SCSI buses and controllers to use and how many disks to connect to each controller, as this may further limit maximum performance. The I/O bus is limited to 20 SCSI controllers and the SCSI

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 574

standard limits disks to 7 per SCSI string. The minimum number of controllers is for the 4-GB disks

$$\text{Minimum number of SCSI strings for 25 4-GB disks} = \frac{25}{7} \text{ or } 4$$

and for 1-GB disks

$$\text{Minimum number of SCSI strings for 100 1-GB disks} = \frac{100}{7} \text{ or } 15$$

We can calculate the maximum IOPS for each configuration:

$$\text{Maximum IOPS for 4 SCSI strings} = 4 * 250 = 1000 \text{ IOPS}$$

$$\text{Maximum IOPS for 15 SCSI strings} = 15 * 250 = 3750 \text{ IOPS}$$

The maximum performance of this number of controllers is slightly lower than the disk I/O throughput, so let's also calculate the number of controllers so they don't become a bottleneck. One way is to find the number of disks they can support per string:

$$\text{Number of disks per SCSI string at full bandwidth} = \frac{250}{41} = 6.1 \text{ or } 6$$

and then calculate the number of strings:

$$\text{Number of SCSI strings for full bandwidth 4-GB disks} = \frac{25}{6} = 4.1 \text{ or } 5$$

$$\text{Number of SCSI strings for full bandwidth 1-GB disks} = \frac{100}{6} = 16.7 \text{ or } 17$$

This establishes the performance of four organizations: 25 4-GB disks with 4 or 5 SCSI strings and 100 1-GB disks with 15 to 17 SCSI strings. The maximum performance of each option is limited by the bottleneck (in boldface):

4-GB disks, 4 strings = Min(5000,5000,10000,1025,**1000**) = 1000 IOPS

4-GB disks, 5 strings = Min(5000,5000,10000,**1025**,1250) = 1025 IOPS

1-GB disks, 15 strings = Min(5000,5000,10000,4100,**3750**) = 3750 IOPS

1-GB disks, 17 strings = Min(5000,5000,10000,**4100**,4250) = 4100 IOPS

We can now calculate the cost for each organization:

4-GB disks, 4 strings = \$50,000 + 4*\$2,500 + 25 * (4096*\$3) = \$367,200

4-GB disks, 5 strings = \$50,000 + 5*\$2,500 + 25 * (4096*\$3) = \$369,700

1-GB disks, 15 strings = \$50,000 + 15*\$2,500 + 100 * (1024*\$3) = \$394,700

1-GB disks, 17 strings = \$50,000 + 17*\$2,500 + 100 * (1024*\$3) = \$399,700

Finally, the cost per IOPS for each of the four configurations is \$367, \$361, \$105, and \$97, respectively. Calculating maximum number of average I/Os per second assuming 100% utilization of the critical resources, the best cost/performance is the organization with the small disks and the largest number of controllers. The small disks have 3.4 to 3.8 times better cost/performance than the large disks in this example. The only drawback is that the larger number of disks will affect system availability unless some form of redundancy is added (see pages 520–521).

This above example assumed that resources can be used 100%. It is instructive to see what is the bottleneck in each organization.

Example

For the organizations in the last example, calculate the percentage of utilization of each resource in the computer system.

Answer

Figure 9.31 gives the answer.

Resource	4-GB disks, 4 strings	4-GB disks, 5 strings	1-GB disks, 15 strings	1-GB disks, 17 strings
CPU	20%	21%	75%	82%
Memory	20%	21%	75%	82%
I/O bus	10%	10%	38%	41%
SCSI buses	100%	82%	100%	96%
Disks	98%	100%	91%	100%

FIGURE 9.31 The percentage of utilization of each resource given the four organizations in the previous example. Either the SCSI buses or the disks are the bottleneck.

In reality buses cannot deliver close to 100% of bandwidth without severe increase in latency and reduction in throughput due to contention. A variety of rules of thumb have been evolved to guide I/O designs:

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 576

No I/O bus should be utilized more than 75% to 80%;

No disk string should be utilized more than 40%;

No disk arm should be seeking more than 60% of the time.

Example

Recalculate performance in the example above using these rules of thumb, and show the utilization of each component. Are there other organizations that follow these guidelines and improve performance?

Answer

Figure 9.31 shows that the I/O bus is far below the suggested guidelines, so we concentrate on the utilization of seek and SCSI bus. The utilization of seek time per disk is

$$\frac{\text{Time of average seek}}{\text{Time between I/Os}} = \frac{12 \text{ ms}}{\frac{1}{41 \text{ IOPS}}} = \frac{12}{24} = 50\%$$

which is below the rule of thumb. The biggest impact is on the SCSI bus:

$$\text{Suggested IOPS per SCSI string} = \frac{1}{4 \text{ ms}} * 40\% = 100 \text{ IOPS.}$$

With this data we can recalculate IOPS for each organization:

4-GB disks, 4 strings = Min(5000,5000,7500,1025,**400**) = 400 IOPS

4-GB disks, 5 strings = Min(5000,5000,7500,1025,**500**) = 500 IOPS

1-GB disks, 15 strings = Min(5000,5000,7500,4100,**1500**) = 1500 IOPS

1-GB disks, 17 strings = Min(5000,5000,7500,4100,**1700**) = 1700 IOPS

Under these assumptions, the small disks have about 3.0 to 4.2 times the performance of the large disks.

Clearly, the string bandwidth is the bottleneck now. The number of disks per string that would not exceed the guideline is

$$\text{Number of disks per SCSI string at full bandwidth} = \frac{100}{41} = 2.4 \text{ or } 2$$

and the ideal number of strings is

$$\text{Number of SCSI strings for full bandwidth 4-GB disks} = \frac{25}{2} = 12.5 \text{ or } 13$$

$$\text{Number of SCSI strings for full bandwidth 1-GB disks} = \frac{100}{2} = 50$$

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 577

This suggestion is fine for 4-GB disks, but the I/O bus is limited to 20 SCSI controllers and strings so that becomes the limit for 1-GB disks:

4-GB disks, 13 strings = Min(5000,5000,7500,**1025**,1300) = 1025 IOPS

1-GB disks, 20 strings = Min(5000,5000,7500,4100,**2000**) = 2000 IOPS

We can now calculate the cost for each organization:

4-GB disks, 13 strings = $50,000 + 13*$2,500 + 25 * (4096*$3) = $389,700

1-GB disks, 20 strings = $50,000 + 20*$2,500 + 100 * (1024*$3) = $407,200

In this case the small disks cost 5% more yet have about twice the performance of the large disks. The utilization of each resource is shown in Figure 9.32. It shows that following the rule of thumb of 40% string utilization sets the performance limit in all but one case.

Resource	4-GB disks, 4 strings	4-GB disks, 5 strings	1-GB disks, 15 strings	1-GB disks, 17 strings	4-GB disks, 13 strings	1-GB disks, 20 strings
CPU	8%	10%	30%	34%	21%	40%
Memory	8%	10%	30%	34%	21%	40%
I/O bus	5%	7%	20%	23%	14%	27%
SCSI buses	40%	40%	40%	40%	32%	40%
Disks	39%	49%	37%	41%	100%	49%
Seek utilization	19%	24%	18%	20%	49%	24%
IOPS	400	500	1500	1700	1025	2000

FIGURE 9.32 The percentage of utilization of each resource given the six organizations in this example, which tries to limit utilization of key resources to the rules of thumb given above.

9.9 Putting It All Together: The IBM 3990 Storage Subsystem

If computer architects were polled to select the leading company in I/O design, IBM would win hands down. A good deal of IBM's mainframe business is commercial applications, known to be I/O intensive. While there are graphic devices and networks that can be connected to an IBM mainframe, IBM's reputation comes from disk performance. It is on this aspect that we concentrate in this section.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 578

The IBM 360/370 I/O architecture has evolved over a period of 25 years. Initially, the I/O system was general purpose, and no special attention was paid to any particular device. As it became clear that magnetic disks were the chief consumers of I/O, the IBM 360 was tailored to support fast disk I/O. IBM's dominant philosophy is to choose latency over throughput whenever it makes a difference. IBM almost never uses a large buffer outside the CPU; their goal is to set up a clear path from main memory to the I/O device so that when a device is ready, nothing can get in the way. Perhaps IBM followed a corollary to the quote on page 526: you can buy bandwidth, but you need to design for latency. As a secondary philosophy, the CPU is unburdened as much as possible to allow the CPU to continue with computation while others perform the desired I/O activities.

The example for this section is the high-end IBM 3090 CPU and the 3990 Storage Subsystem. The IBM 3090, models 3090/100 to 3090/600, can contain one to six CPUs. This 18.5-ns-clock-cycle machine has a 16-way interleaved memory that can transfer eight bytes every clock cycle on each of two (3090/100) or four (3090/600) buses. Each 3090 processor has a 64-KB, 4-way–set-associative, write-back cache, and the cache supports pipelined access taking two cycles. Each CPU is rated about 30 IBM MIPS (see page 78), giving at most 180 MIPS to the IBM 3090/600. Surveys of IBM mainframe installations suggest a rule of thumb of about 4 GB of disk storage per MIPS of CPU power (see Section 9.12).

It is only fair warning to say that IBM terminology may not be self-evident, although the ideas are not difficult. Remember that this I/O architecture has evolved since 1964. While there may well be ideas that IBM wouldn't include if they were to start anew, they are able to make this scheme work, and make it work well.

The 3990 I/O Subsystem Data-Transfer Hierarchy and Control Hierarchy

The I/O subsystem is divided into two hierarchies:

1. Control—This hierarchy of controllers negotiates a path through a maze of possible connections between the memory and the I/O device and controls the timing of the transfer.
2. Data—This hierarchy of connections is the path over which data flows between memory and the I/O device.

After going over each of the hierarchies, we trace a disk read to help understand the function of each component.

For simplicity, we begin by discussing the data-transfer hierarchy, shown in Figure 9.33 (page 548). This figure shows one section of the hierarchy that contains up to 64 large IBM disks; using 64 of the recently announced IBM 3390 disks, this piece could connect to over one trillion bytes of storage! Yet this

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 579

piece represents only one-sixth of the capacity of the IBM 3090/600 CPU. This ability to expand from a small I/O system to hundreds of disks and terabytes of storage is what gives IBM mainframes their reputation in the I/O world.

The best-known member of the data hierarchy is the *channel*. The channel is nothing more than 50 wires that connect two levels on the I/O hierarchy together. Only 18 of the 50 wires are used for transferring data (8 data plus 1 parity in each direction), while the rest are for control information. For years the maximum data rate was 3 MB per second, but it recently was raised to 4.5 MB per second. Up to 48 channels can be connected to a 3090/100 CPU, and up to

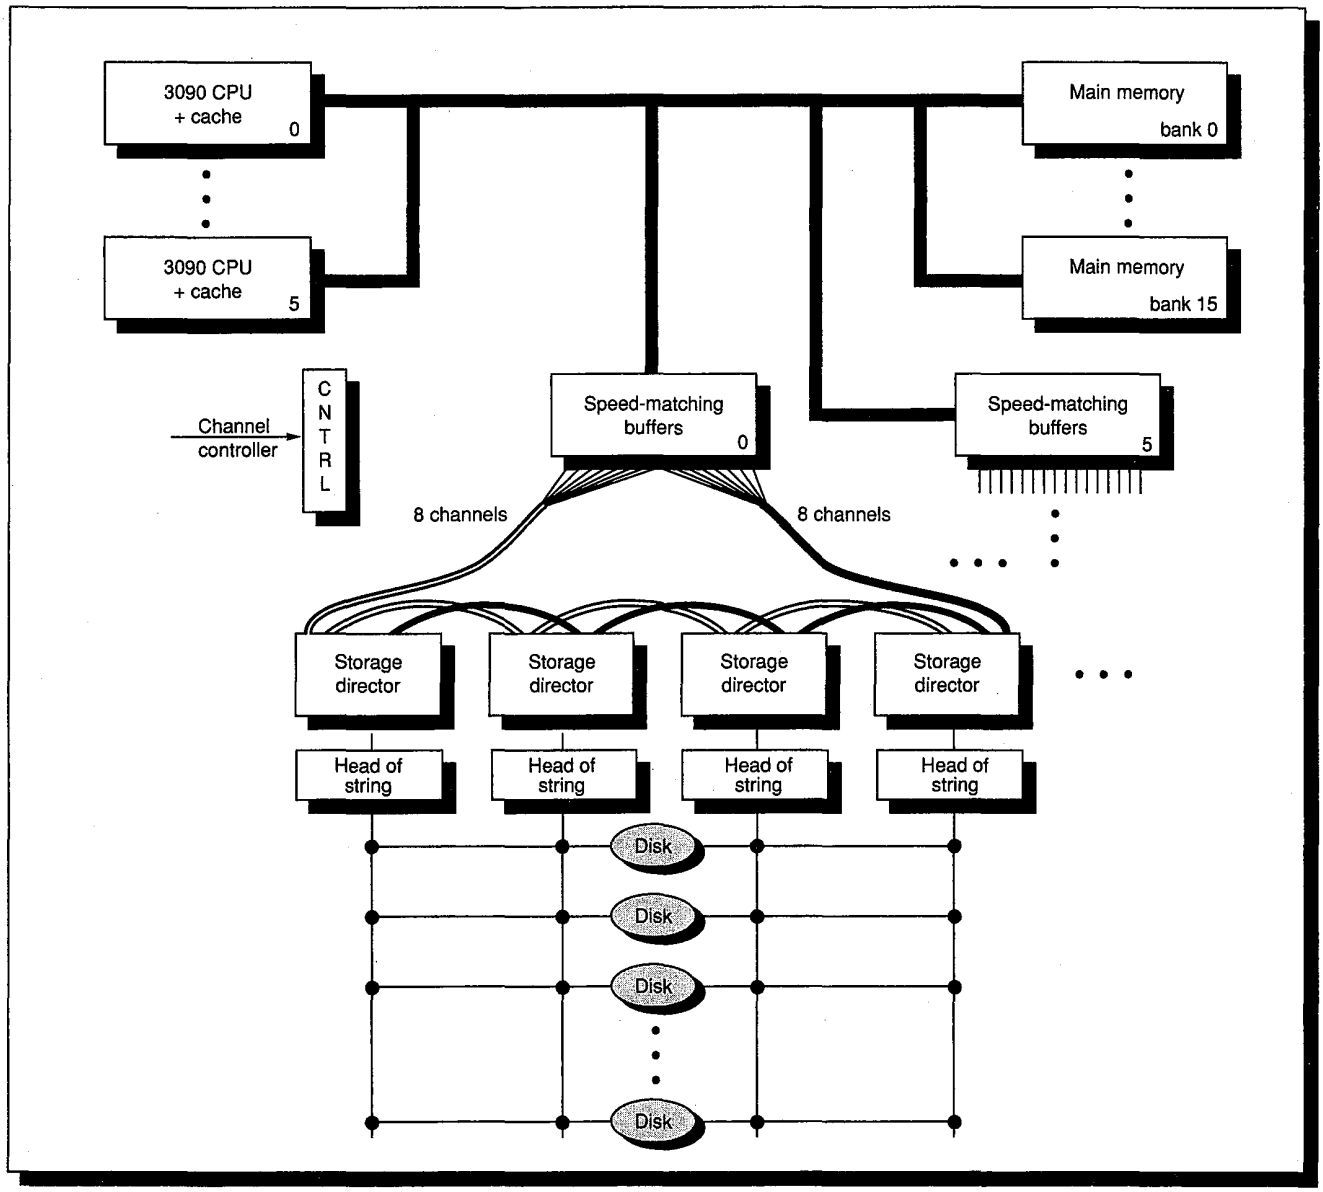

FIGURE 9.33 The data-transfer hierarchy in the IBM 3990 I/O Subsystem. Note that all the channels are connected to all the storage directors. The disks at the bottom represent the quad-ported IBM 3380 disk drives, with the maximum of 64 disks. The collection of disks on the same path to the head-of-string controller is called a *string*.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 580

96 channels to a 3090/600. Because they are "multiprogrammed," channels can actually service several disks. For historical reasons, IBM calls this *block multiplexing*.

Channels are connected to the 3090 main memory via two *speed-matching buffers*, which funnel all the channels into a single port to main memory. Such buffers simply match the bandwidth of the I/O device to the bandwidth of the memory system. There are two 8-byte buffers per channel.

The next level down the data hierarchy is the *storage director*. This is an intermediary device that allows the many channels to talk to many different I/O devices. Four to sixteen channels go to the storage director depending on the model, and two or four paths come out the bottom to the disks. These are called *two-path strings* or *four-path strings* in IBM parlance. Thus, each storage director can talk to any of the disks using one of the strings. At the top of each string is the *head of string*, and all communication between disks and control units must pass through it.

At the bottom of the datapath hierarchy are the disk devices themselves. To increase availability, disk devices like the IBM 3380 provide four paths to connect to the storage director; if one path fails, the device can still be connected.

The redundant paths from main memory to the I/O device not only improve availability, but also can improve performance. Since the IBM philosophy is to avoid large buffers, the path from the I/O device to main memory must remain connected until the transfer is complete. If there were a single hierarchical path from devices to the speed-matching buffer, only one I/O device in a subtree could transfer at a time. Instead, the multiple paths allow multiple devices to transfer simultaneously through the storage director and into memory.

The task of setting up the datapath connection is that of the control hierarchy. Figure 9.34 shows both the control and data hierarchies of the 3990 I/O subsystem. The new device is the I/O processor. The 3090 channel controller and I/O processor are load/store machines similar to DLX, except that there is no memory hierarchy. In the next subsection we see how the two hierarchies work together to read a disk sector.

Tracing a Disk Read in the IBM 3990 I/O Subsystem

The 12 steps below trace a sector read from an IBM 3380 disk. Each of the 12 steps is labeled on a drawing of the full hierarchy in Figure 9.34 (page 550).

1. The user sets up a data structure in memory containing the operations that should occur during this I/O event. This data structure is termed an *I/O control block*, or IOCB, which also points to a list of channel control words (CCWs). This list is called a *channel program*. Normally, the operating system provides the channel program, but some users write their own. The operating system checks the IOCB for protection violations before the I/O can continue.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 581

2. The CPU executes a `START SUBCHANNEL` instruction. The actual request is defined in the channel program. A channel program to read a record might look like Figure 9.35.

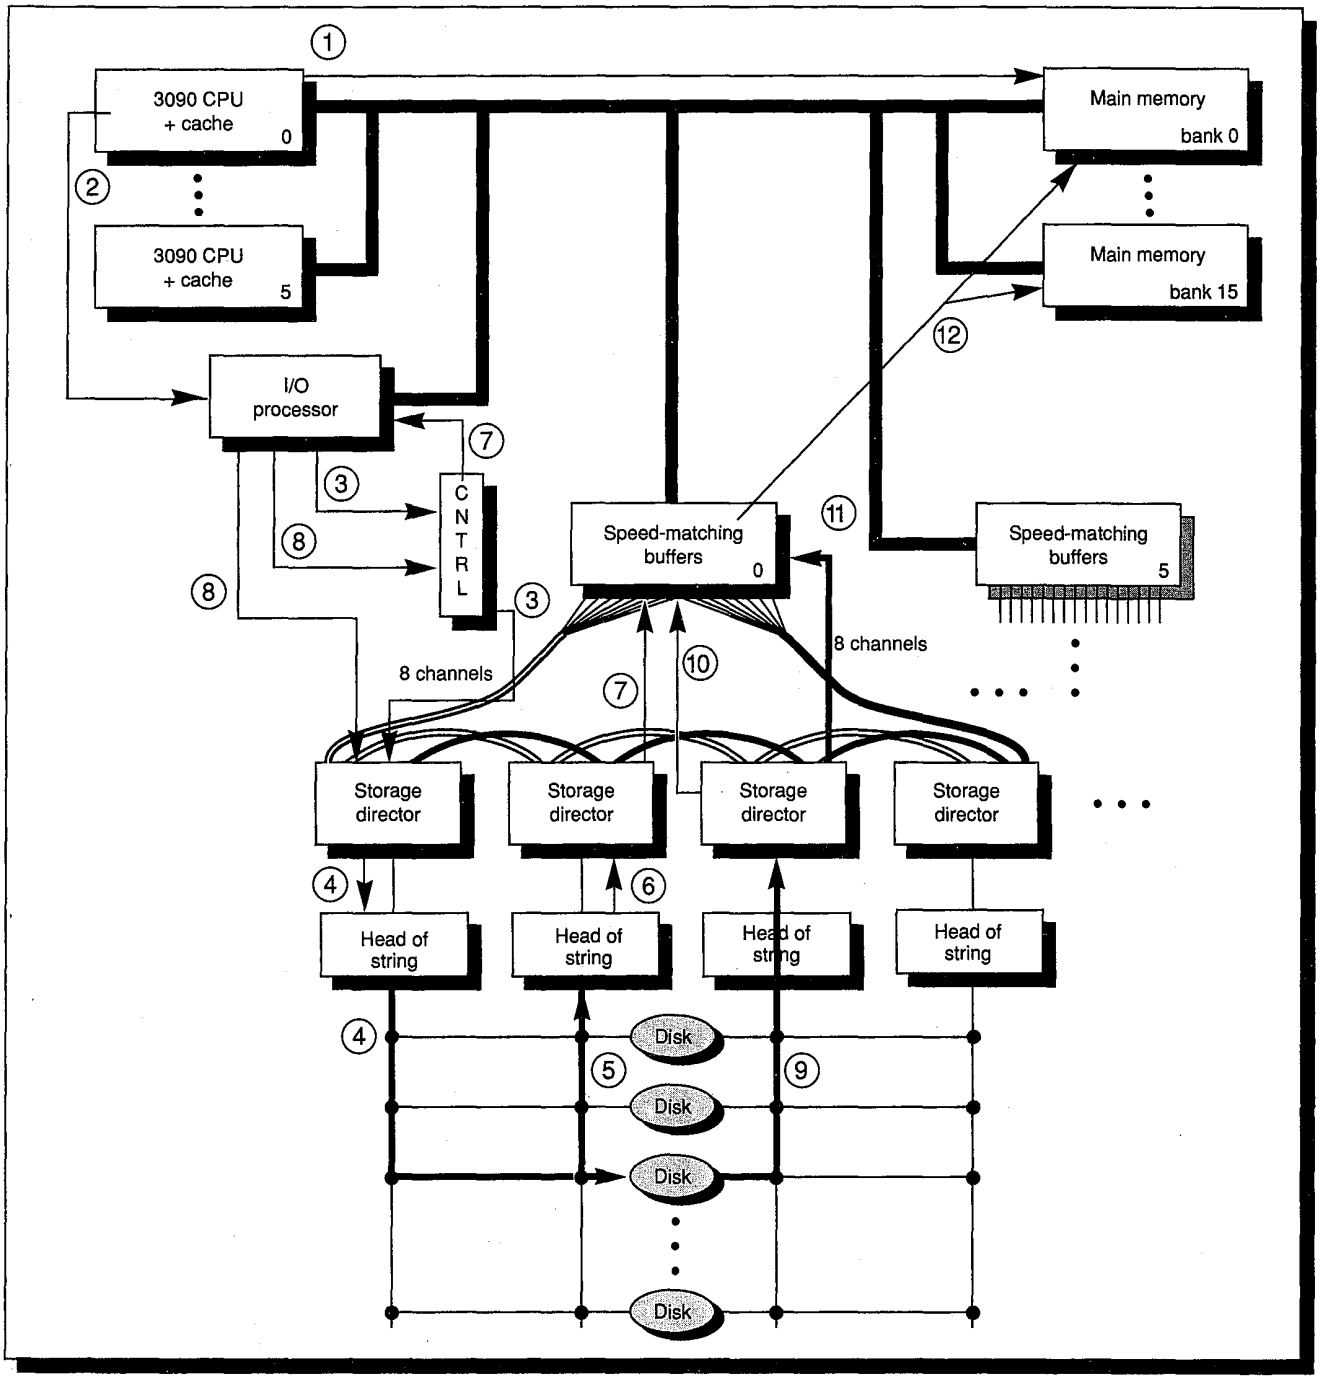

FIGURE 9.34 The control and data hierarchies in the IBM 3990 I/O Subsystem labeled with the 12 steps to read a sector from disk. The only new box over Figure 9.33 (page 548) is the I/O processor.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 582

Location	CCW	Comment
CCW1:	Define Extent	Transfers a 16-byte parameter to the storage director. The channel sees this as a write data transfer.
CCW2:	Locate Record	Transfers a 16-byte parameter to the storage director as above. The parameter identifies the operation (read in this case) plus seek, sector number, and record ID. The channel again sees this as a write data transfer.
CCW3:	Read Data	Transfers the desired disk data to the channel and then to the main memory.

FIGURE 9.35 A channel program to perform a disk read, consisting of three channel command words (CCWs). The operating system checks for virtual memory access violations of CCWs by simulating them to check for violations. These instructions are linked so that only one START SUBCHANNEL instruction is needed.

3. The I/O processor uses the control wires of one of the channels to tell the storage director which disk is to be accessed and the disk address to be read. The channel is then released.

4. The storage director sends a SEEK command to the head-of-string controller and the head-of-string controller connects to the desired disk, telling it to seek to the appropriate track, and then disconnects. The disconnect occurs between CCW2 and CCW3 in Figure 9.35.

Upon completion of these first four steps of the read, the arm on the disk seeks the correct track on the correct IBM 3380 disk drive. Other I/O operations can use the control and data hierarchy while this disk is seeking and the data is rotating under the read head. The I/O processor thus acts like a multiprogrammed system, working on other requests while waiting for an I/O event to complete.

An interesting question arises: When there are multiple uses for a single disk, what prevents another seek from screwing up the works before the original request can continue with the I/O event in progress? The answer is the disk appears busy to the programs in the 3090 between the time a START SUBCHANNEL instruction starts a channel program (step 2) and the end of that channel program. An attempt to execute another START SUBCHANNEL instruction would receive busy status from the channel or from the disk device.

After both the seek completes and the disk rotates to the desired point relative to the read head, the disk reconnects to a channel. To determine the rotational position of the 3380 disk, IBM provides rotational positional sensing (RPS), a feature that gives early warning when the data will rotate under the read head. IBM essentially extends the seek time to include some of the rotation time, thereby tying up the datapath as little as possible. Then the I/O can continue:

5. When the disk completes the seek and rotates to the correct position, it contacts the head-of-string controller.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 583

6. The head-of-string controller looks for a free storage director to send the signal that the disk is on the right track.

7. The storage director looks for a free channel so that it can use the control wires to tell the I/O processor that the disk is on the right track.

8. The I/O processor simultaneously contacts the storage director and I/O device (the IBM 3380 disk) to give the OK to transfer data, and tells the channel controller where to put the information in main memory when it arrives at the channel.

There is now a direct path between the I/O device and memory and the transfer can begin:

9. When the disk is ready to transfer, it sends the data at 3 megabytes per second over a bit-serial line to the storage director.

10. The storage director collects 16 bytes in one of two buffers and sends the information on to the channel controller.

11. The channel controller has a pair of 16-byte buffers per storage director and sends 16 bytes over a 3-MB or 4.5-MB per second, 8-bit-wide datapath to the speed-matching buffers.

12. The speed-matching buffers take the information coming in from all channels. There are two 8-byte buffers per channel that send 8 bytes at a time to the appropriate locations in main memory.

Since nothing is free in computer design, one might expect there to be a cost in anticipating the rotational delay using RPS. Sometimes a free path cannot be established in the time available due to other I/O activity, resulting in an *RPS miss*. An RPS miss means the 3990 I/O Subsystem must either:

- Wait another full rotation—16.7 ms—before the data is back under the head, or
- Break down the hierarchical datapath and start all over again!

Lots of RPS misses can ruin response times.

As mentioned above, the IBM I/O system evolved over many years, and Figure 9.36 shows the change in response time for a few of those changes. The first improvement concerns the path for data after reconnection. Before the System/370-XA, the data path through the channels and storage director (steps 5 through 12) had to be the same as the path taken to request the seek (steps 1 through 4). The 370-XA allows the path after reconnection to be different, and this option is called *dynamic path reconnection* (DPR). This change reduced the time waiting for the channel path and the time waiting for disks (queueing delay), yielding a reduction in the total average response time of 17%. The second change in Figure 9.36 involved a new disk design. Improvements to the

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 584

microcode control of the 3380D made slight improvements in seek time plus removed a restriction that disk arms that were on the same internal path were prevented from operating at the same time. IBM calls this option *Device Level Select* (DLS). This change reduced internal path delays to 0. This had little impact since there was not much time waiting on internal delays because customers intentionally placed data on disks trying to avoid internal path delays. This second change reduced response time another 9%. The final change was addition of a 32-MB write-through disk cache to a 3380D, called the IBM 3880-23. The disk cache reduced average rotational latency, seek time, and queueing delays, giving another 41% reduction in response time.

One indication of the effectiveness of DPR is the number of disk devices connected to a string. Studies of IBM systems using DPR, which average 16 disk devices per string versus 12 without DPR, suggest dynamic reconnect allows a higher I/O rate with comparable response time [Henly and McNutt 1989].

Summary of the IBM 3990 I/O Subsystem

Goals for I/O systems consist of supporting the following:

- Low cost
- A variety of types of I/O devices

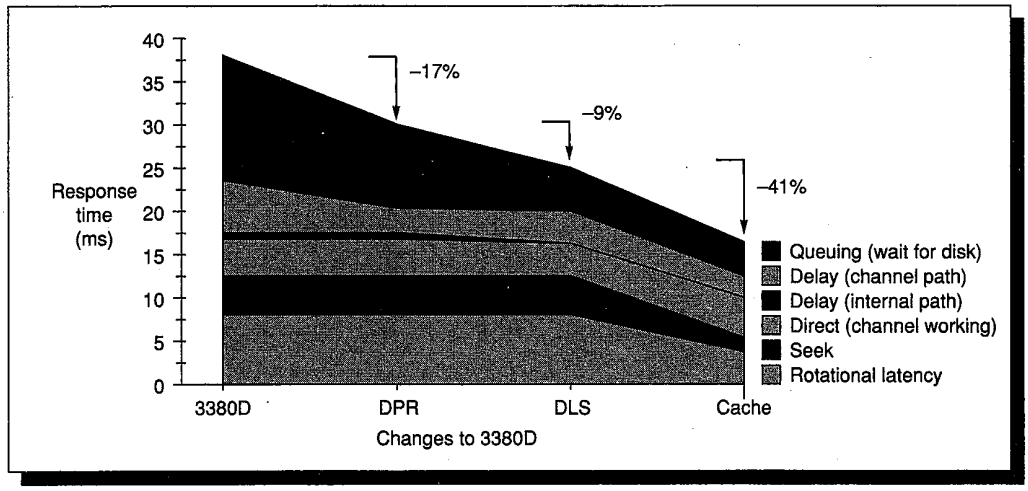

FIGURE 9.36 Changes in response time with improvements in 3380D broken into six categories [Friesenborg and Wicks 1985]. Queueing delay refers to the time when the program waits for another program to finish with the disk. Channel-path delay is the time the operation waits due to the channel path and storage director being busy with another task. Internal-path delay is similar to channel-path delay except it refers to internal paths in the 3380D. Direct means the time the channel path is busy with the operation. Seek time and rotational latency are the standard definitions. Robinson and Blount [1986] report in the study of the 3880-23 that the read hit rate for the 32-MB write-through cache in some large systems averages about 90%, with reads accounting for 92% of the disk accesses.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 585

- A large number of I/O devices at a time
- High performance
- Low latency

Substantial expendability and lower latency are hard to get at the same time. IBM channel-based systems achieve the third and fourth goals by utilizing hierarchical data paths to connect a large number of devices. The many devices and parallel paths allow simultaneous transfers and, thus, high throughput. By avoiding large buffers and providing enough extra paths to minimize delay from congestion, channels offer low-latency I/O as well. To maximize use of the hierarchy, IBM uses rotational positional sensing to extend the time that other tasks can use the hierarchy during an I/O operation.

Therefore, a key to performance of the IBM I/O subsystem is the number of rotational positional misses and congestion on the channel paths. A rule of thumb is that the single-path channels should be no more than 30% utilized and the quad-path channels should be no more than 60% utilized, or too many rotational positional misses will result. This I/O architecture dominates the industry, yet it would be interesting to see what, if anything, IBM would do differently if given a clean slate.

9.10 Fallacies and Pitfalls

Fallacy: I/O plays a small role in supercomputer design

The goal of the Illiac IV was to be the world's fastest computer. It may not have achieved that goal, but it showed I/O as the Achilles' Heel of high-performance machines. In some tasks, more time was spent in loading data than in computing. Amdahl's Law demonstrated the importance of high performance in all the parts of a high-speed computer. (In fact, Amdahl made his comment in reaction to claims for performance through parallelism made on behalf of the Illiac IV.) The Illiac IV had a very fast transfer rate (60 MB/sec), but very small, fixed-head disks (12-MB capacity). Since they were not large enough, more storage was provided on a separate computer. This led to two ways of measuring I/O overhead:

Warm start—Assuming the data is on the fast, small disks, I/O overhead is the time to load the Illiac IV memory from those disks.

Cold start—Assuming the data is in on the other computer, I/O overhead must include the time to first transfer the data to the Illiac IV fast disks.

Figure 9.37 shows ten applications written for the Illiac IV in 1979. Assuming warm starts, the supercomputer was busy 78% of the time and waiting for I/O 22% of the time; assuming cold starts, it was busy 59% of the time and waiting for I/O 41% of the time.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 586

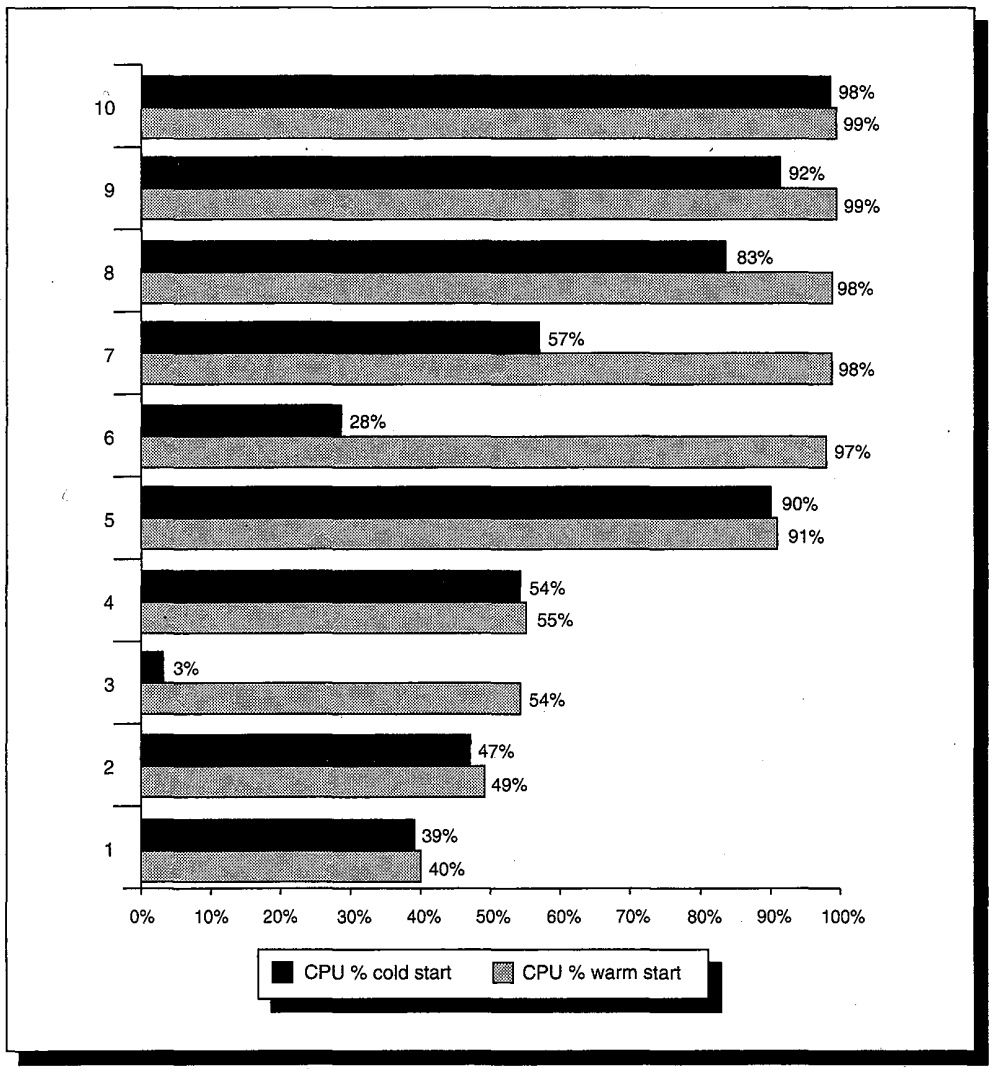

FIGURE 9.37 Feierback and Stevenson [1979] summarized the important Illiac IV applications and the percentage of time spent computing versus waiting for I/O. The arithmetic means of the 10 programs are 78% computing for warm start and 59% computing for cold start.

Pitfall: Moving functions from the CPU to the I/O processor to improve performance.

There are many examples of this pitfall, although I/O processors can enhance performance. A problem inherent with a family of computers is that the migration of an I/O feature usually changes the instruction set architecture or system architecture in a programmer-visible way, causing all future machines to have to live with a decision that made sense in the past. If CPUs are improved in cost/performance more rapidly than the I/O processor (and this will likely be the case) then moving the function may result in a slower machine in the next CPU.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 587

The most telling example comes from the IBM 360. It was decided that the performance of the ISAM system, an early database system, would improve if some of the record searching occurred in the disk controller itself. A key field was associated with each record, and the device searched each key as the disk rotated until it found a match. It would then transfer the desired record. For the disk to find the key, there had to be an extra gap in the track. This scheme is applicable to searches through indices as well as data.

The speed a track can be searched is limited by the speed of the disk and of the number of keys that can be packed on a track. On an IBM 3330 disk the key is typically 10 characters, but the total gap between records is equivalent to 191 characters if there were a key. (The gap is only 135 characters if there is no key, since there is no need for an extra gap for the key.) If we assume the data is also 10 characters and the track has nothing else on it, then a 13165-byte track can contain

$$\frac{13165}{191+10+10} = 62 \text{ key-data records}$$

This performance is

$$\frac{16.7 \text{ ms (1 revolution)}}{62} \approx .25 \text{ ms/key search}$$

In place of this scheme, we could put several key-data pairs in a single block and have smaller inter-record gaps. Assuming there are 15 key-data pairs per block and the track has nothing else on it, then

$$\frac{13165}{135+15*(10+10)} = \frac{13165}{135+300} = 30 \textbf{ blocks} \text{ of key-data pairs}$$

The revised performance is then

$$\frac{16.7 \text{ ms (1 revolution)}}{30*15} \approx .04 \text{ ms/key search}$$

Yet as CPUs got faster, the CPU time for a search was trivial. While the strategy made early machines faster, programs that use the search-key operation in the I/O processor run six times slower on today's machines!

Fallacy: Comparing the price of media versus the price of the packaged system.

This happens most frequently when new memory technologies are compared to magnetic disks. For example, comparing the DRAM-chip price to magnetic-disk packaged price in Figure 9.16 (page 518) suggests the difference is less than a factor of 10, but its much greater when the price of packaging DRAM is included. A common mistake with removable media is to compare the media cost not including the drive to read the media. For example, optical media costs

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 588

only $1 per MB in 1990, but including the cost of the optical drive may bring the price closer to $6 per MB.

Fallacy: The time of an average seek of a disk in a computer system is the time for a seek of one-third the number of cylinders.

This fallacy comes from confusing the way manufacturers market disks with the expected performance and with the false assumption that seek times are linear in distance. The 1/3 distance rule of thumb comes from calculating the distance of a seek from one random location to another random location, not including the current cylinder and assuming there are a large number of cylinders. In the past, manufacturers listed the seek of this distance to offer a consistent basis for comparison. (As mentioned on page 516, today they calculate the "average" by timing all seeks and dividing by the number.) Assuming (incorrectly) that seek time is linear in distance, and using the manufacturers reported minimum and "average" seek times, a common technique to predict seek time is:

$$\text{Time}_{\text{seek}} = \text{Time}_{\text{minimum}} + \frac{\text{Distance}}{\text{Distance}_{\text{average}}} * (\text{Time}_{\text{average}} - \text{Time}_{\text{minimum}})$$

The fallacy concerning seek time is twofold. First, seek time is **not** linear with distance; the arm must accelerate to overcome inertia, reach its maximum traveling speed, decelerate as it reaches the requested position, and then wait to allow the arm to stop vibrating (settle time). Moreover, in recent disks sometimes the arm must pause to control vibrations. Figure 9.38 (page 558) plots time versus seek distance for an example disk. It also shows the error in the simple seek-time formula above. For short seeks, the acceleration phase plays a larger role than the maximum traveling speed, and this phase is typically modeled as the square root of the distance. Figure 9.39 (page 558) shows accurate formulas used to model the seek time versus distance for two disks.

The second problem is the average in the product specification would only be true if there was no locality to disk activity. Fortunately, there is both temporal and spatial locality (page 403 in Chapter 8): disk blocks get used more than once and disk blocks near the current cylinder are more likely to be used than those farther away. For example, Figure 9.40 (page 559) shows sample measurements of seek distances for two workloads: a UNIX timesharing workload and a business-processing workload. Notice the high percentage of disk accesses to the same cylinder, labeled distance 0 in the graphs, in both workloads.

Thus, this fallacy couldn't be more misleading. The Exercises debunk this fallacy in more detail.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 589

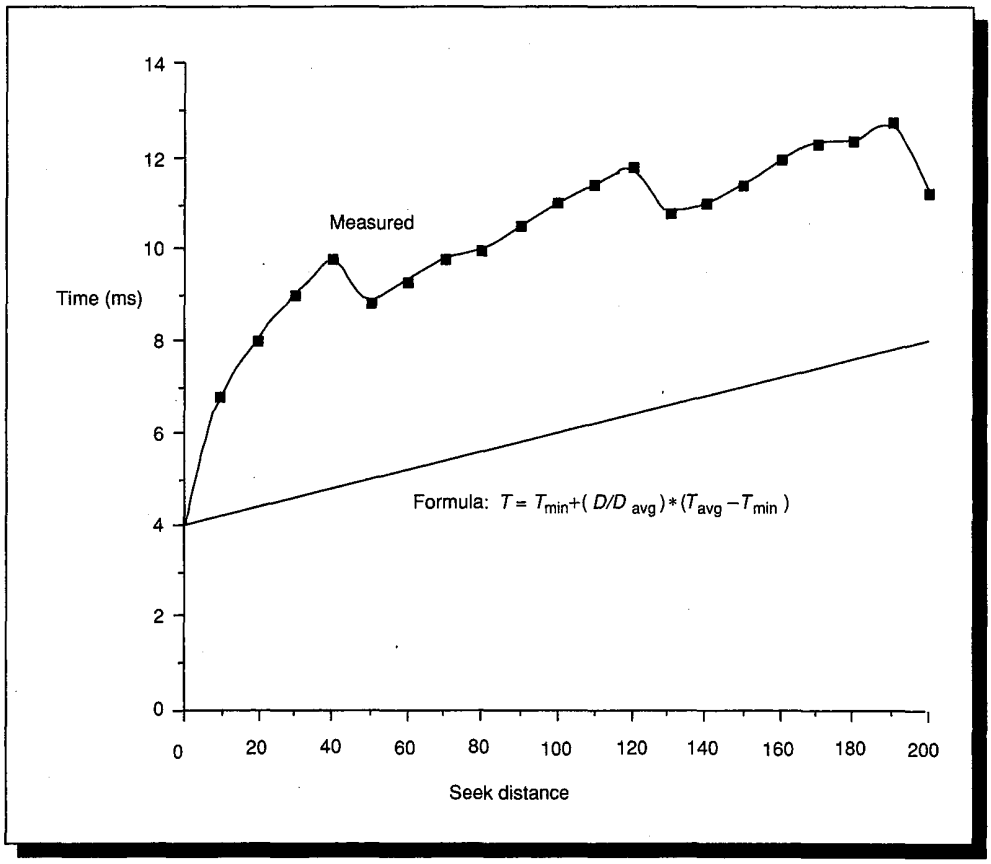

FIGURE 9.38 Seek time versus seek distance for the first 200 cylinders. The Imprimis Sabre 97209 contains 1.2 GB using 1635 cylinders and has the IPI-2 interface [Imprimis 1989]. This is an 8-inch disk. Note that longer seeks can take **less** time than shorter seeks. For example, a 40-cylinder seek takes almost 10 ms, while a 50-cylinder seek takes less than 9 ms.

IBM 3380D Range for formula		Formulas	IBM 3380J Range for formula		Formulas
≥	≤		≥	≤	
1	50	$1.9 + \sqrt{\text{Distance}} - \frac{\text{Distance}}{50}$	1	50	$2.48 + \sqrt{\text{Distance}} - \frac{\text{Distance}}{20}$
51	100	8.1 + 0.044 * (Distance–50)	51	130	7.28 + 0.0320 * (Distance–50)
101	500	10.3 + 0.025 * (Distance–100)	131	500	10.08 + 0.0166 * (Distance–130)
501	884	20.4 + 0.017 * (Distance–500)	501	884	16.00 + 0.0114 * (Distance–500)

FIGURE 9.39 Formulas for seek time in ms for two IBM disks. Thisquen [1988] measured these disks and proposed these formulas to model them. The two columns on the left show the range of seek distances in cylinders to which each formula applies. Each disk has 885 cylinders, so the maximum seek is 884.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 590

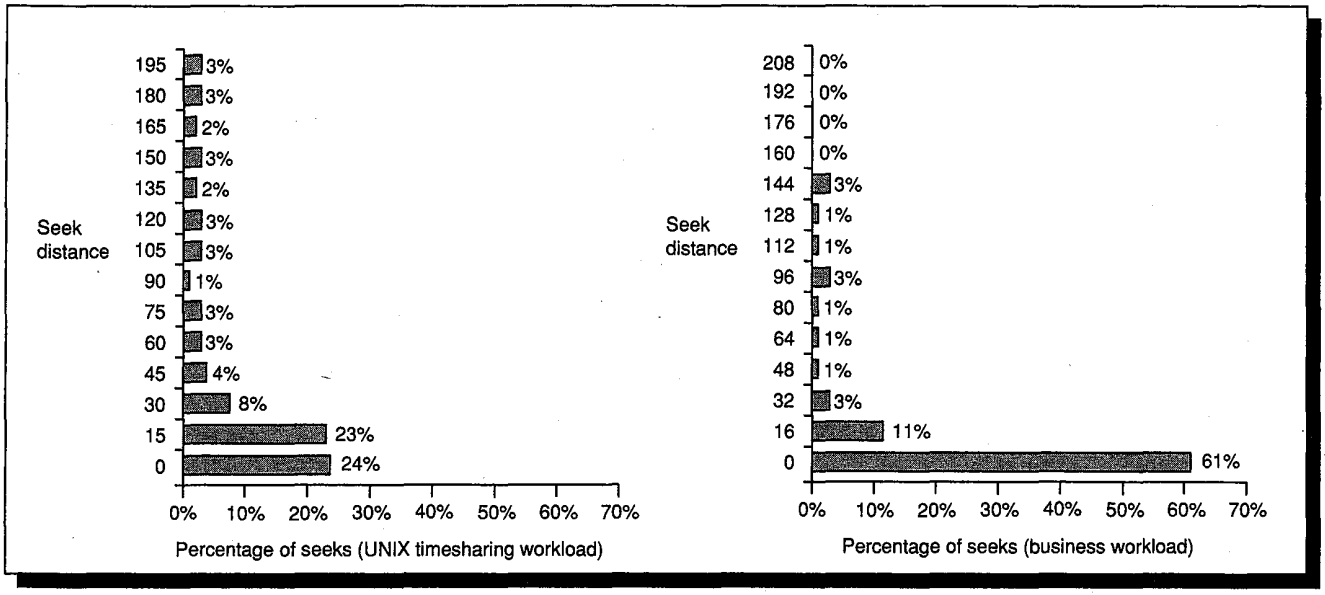

FIGURE 9.40 Sample measurements of seek distances for two systems. The left measurements were taken on a UNIX timesharing system. The right measurements were taken from a business processing application in which the disk seek activity was scheduled. Seek distance of 0 means the access was made to the same cylinder. The rest of the numbers show the collective percentage for distances up between numbers on the y axis. For example, 11% for the bar labeled 16 in the business graph means that the percentage of seeks between 1 and 16 cylinders was 11%. The UNIX measurements stopped at 200 cylinders, but this captured 85% of the accesses. The total was 1000 cylinders. The business measurements tracked all 816 cylinders of the disks. The only seek distances with 1% or greater of the seeks that are not in the graph are 224 with 4% and 304, 336, 512, and 624 each having 1%. This total is 94%, with the difference being small but nonzero distances in other categories. The measurements are courtesy of Dave Anderson of Imprimis.

9.11 Concluding Remarks

I/O systems are judged by the variety of I/O devices, the maximum number of I/O devices, cost, and performance, measured both in latency and in throughput. These common goals lead to widely varying schemes, with some relying extensively on buffering and some avoiding buffering at all costs. If one is clearly better than the other, it is not obvious today. Perhaps this situation is like the instruction set debates of the 1980s, and the strengths and weaknesses of the alternatives will become apparent in the 1990s.

According to Amdahl's Law, ignorance of I/O will lead to wasted performance as CPUs get faster. Disk performance is growing at 4% to 6% per year, while CPUs are growing at a much faster rate. The future demands for I/O include better algorithms, better organizations, and more caching in a struggle to keep pace.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 591

9.12 Historical Perspective and References

The forerunner of today's workstations was the Alto developed at Xerox Palo Alto Research Center in 1974 [Thacker et al. 1982]. This machine reversed traditional wisdom, making instruction set interpretation take back seat to the display: the display used half the memory bandwidth of the Alto. In addition to the bit-mapped display, this historic machine had the first Ethernet [Metcalfe and Boggs 1976] and the first laser printer. It also had a mouse, invented earlier by Doug Engelbart of SRI, and a removable cartridge disk. The 16-bit CPU implemented an instruction set similar to the Data General Nova and offered writable control store (see Chapter 5, Section 5.8). In fact, a single microprogrammable engine drove the graphics display, mouse, disks, network, and, when there was nothing else to do, interpreted the instruction set.

The attraction of a personal computer is that you don't have to share it with anyone. This means response time is predictable, unlike timesharing systems. Early experiments in the importance of fast response time were performed by Doherty and Kelisky [1979]. They showed that if computer-system response time increased a second that user think time did also. Thadhani [1981] showed a jump in productivity as computer response times dropped to a second and another jump as they dropped to a half-second. His results inspired a flock of studies, and they supported his observations [IBM 1982]. In fact, some studies were started to disprove his results! Brady [1986] proposed differentiating entry time from think time (since entry time was becoming significant when the two were lumped together) and provided a cognitive model to explain the more than linear relationship between computer response time and user think time.

The ubiquitous microprocessor has inspired not only personal computers in the 1970s, but the current trend to moving controller functions into I/O devices in the late 1980s and 1990s. For example, microcoded routines in a central CPU made sense for the Alto in 1975, but technological changes soon made separate microprogrammable controller I/O devices economical. These were then replaced by the application-specific integrated circuits. I/O devices continued this trend by moving controllers into the devices themselves. These are called *intelligent devices*, and some bus standards (e.g., IPI and SCSI) have been created just for these devices. Intelligent devices can relax the timing constraints by handling many of the low-level tasks and queuing the results. For example, many SCSI-compatible disk drives include a track buffer on the disk itself, supporting read ahead and connect/disconnect. Thus, on a SCSI string some disks can be seeking and others loading their track buffer while one is transferring data from its buffer over the SCSI bus.

Speaking of buses, the first multivendor bus may have been the PDP-11 Unibus in 1970. DEC encouraged other companies to build devices that would plug into their bus, and many companies did. A more recent example is SCSI,

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 592

which stands for *small computer systems interface*. This bus, originally called SASI, was invented by Shugart and was later standardized by the IEEE. Sometimes buses are developed in academia; the *NuBus* was developed by Steve Ward and his colleagues at MIT and used by several companies. Alas, this open-door policy on buses is in contrast to companies with proprietary buses using patented interfaces, thereby preventing competition from plug-compatible vendors. This practice also raises costs and lowers availability of I/O devices that plug into proprietary buses, since such devices must have an interface designed just for that bus. Levy [1978] has a nice survey on issues in buses.

We must also give a few references to specific I/O devices. Readers interested in the ARPANET should see Kahn [1972]. As mentioned in one of the section quotes, the father of computer graphics is Ivan Sutherland, who received the ACM Turing Award in 1988. Sutherland's Sketchpad system [1963] set the standard for today's interfaces and displays. See Foley and Van Dam [1982] and Newman and Sproull [1979] for more on computer graphics. Scranton, Thompson, and Hunter [1983] were among the first to report the myths concerning seek times and distances for magnetic disks.

Comments on the future of disks can be found in several sources. Goldstein [1987] projects the capacity and I/O rates for IBM mainframe installations in 1995, suggesting that the ratio is no less than 3.7 GB per IBM mainframe MIPS today, and that will grow to 4.5 GB per MIPS in 1995. Frank [1987] speculated on the physical recording density, proposing the MAD formula on disk growth that we used in Section 9.4. Katz, Patterson, and Gibson [1990] survey current high-performance disks and I/O systems and speculate about future systems. The possibility of achieving higher-performance I/O systems using collections of disks is found in papers by Kim [1986], Salem and Garcia-Molina [1986], and Patterson, Gibson, and Katz [1987].

Looking backward rather than forward, the first machine to extend interrupts from detecting arithmetic abnormalities to detecting asynchronous I/O events is credited as the NBS DYSEAC in 1954 [Leiner and Alexander 1954]. The following year the first machine with DMA was operational, the IBM SAGE. Just as today's DMA, the SAGE had address counters that performed block transfers in parallel with CPU operations. The first I/O channel may have been on the IBM 709 in 1957 [Bashe et al. 1981 and 1986]. Smotherman [1989] explores the history of I/O in more depth.

References

ANON ET AL. [1985]. "A measure of transaction processing power," Tandem Tech. Rep. TR 85.2. Also appeared in *Datamation*, April 1, 1985.

BASHE, C. J., W. BUCHHOLZ, G .V. HAWKINS, J .L. INGRAM, AND N. ROCHESTER [1981]. "The architecture of IBM's early computers," *IBM J. of Research and Development* 25:5 (September) 363–375.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 593

BASHE, C. J., L. R. JOHNSON, J. H. PALMER, AND E. W. PUGH [1986]. *IBM's Early Computers*, MIT Press, Cambridge, Mass.

BORRILL, P. L. [1986]. "32-bit buses–An objective comparison," *Proc. Buscon 1986 West*, San Jose, Calif., 138–145.

BRADY, J. T. [1986]. "A theory of productivity in the creative process," *IEEE CG&A* (May) 25–34.

BUCHER, I. V. AND A. H. HAYES [1980]. "I/O Performance measurement on Cray-1 and CDC 7000 computers," *Proc. Computer Performance Evaluation Users Group, 16th Meeting*, NBS 500-65, 245–254.

CHEN, P. [1989]. *An Evaluation of Redundant Arrays of Inexpensive Disks Using an Amdahl 5890*, M. S. Thesis, Computer Science Division, Tech. Rep. UCB/CSD 89/506.

DOHERTY, W. J. AND R. P. KELISKY [1979]. "Managing VM/CMS systems for user effectiveness," *IBM Systems J.* 18:1, 143–166.

FEIERBACK, G AND D. STEVENSON [1979]. "The Illiac-IV," in *Infotech State of the Art Report on Supercomptuers*, Maidenhead, England. This data also appears in D. P. Siewiorek, C. G. Bell, and A. Newell, *Computer Structures: Principles and Examples* (1982), McGraw-Hill, New York, 268–269.

FOLEY, J. D. AND A. VAN DAM [1982]. *Fundamentals of Interactive Computer Graphics*, Addison-Wesley, Reading, Mass.

FRANK, P. D. [1987]. "Advances in Head Technology," presentation at *Challenges in Winchester Technology* (December 15), Santa Clara Univ.

FRIESENBORG, S. E. AND R. J. WICKS [1985]. "DASD expectations: The 3380, 3380-23, and MVS/XA," Tech. Bulletin GG22-9363-02 (July 10), Washington Systems Center.

GOLDSTEIN, S. [1987]. "Storage performance—an eight year outlook," Tech. Rep. TR 03.308-1 (October), Santa Teresa Laboratory, IBM, San Jose, Calif.

HENLY, M. AND B. MCNUTT [1989]. "DASD I/O characteristics: A comparison of MVS to VM," Tech. Rep. TR 02.1550 (May), IBM, General Products Division, San Jose, Calif.

HOWARD, J. H. ET AL. [1988]. "Scale and performance in a distributed file system," *ACM Trans. on Computer Systems* 6:1, 51–81.

IBM [1982]. *The Economic Value of Rapid Response Time*, GE20-0752-0 White Plains, N.Y., 11–82.

IMPRIMIS [1989]. "Imprimis Product Specification, 97209 Sabre Disk Drive IPI-2 Interface 1.2 GB," Document No. 64402302 (May).

KAHN, R. E. [1972]. "Resource-sharing computer communication networks," *Proc. IEEE* 60:11 (November) 1397-1407.

KATZ, R. H., D. A. PATTERSON, AND G. A. GIBSON [1990]. "Disk system architectures for high performance computing," *Proc. IEEE* 78:2 (February).

KIM, M. Y. [1986]. "Synchronized disk interleaving," *IEEE Trans. on Computers* C-35:11 (November).

LEINER, A. L. [1954]. "System specifications for the DYSEAC," *J. ACM* 1:2 (April) 57–81.

LEINER, A. L. AND S. N. ALEXANDER [1954]. "System organization of the DYSEAC," *IRE Trans. of Electronic Computers* EC-3:1 (March) 1–10.

LEVY, J. V. [1978]. "Buses: The skeleton of computer structures," in *Computer Engineering: A DEC View of Hardware Systems Design*, C. G. Bell, J. C. Mudge, and J. E. McNamara, eds., Digital Press, Bedford, Mass.

MABERLY, N. C. [1966]. *Mastering Speed Reading*, New American Library, Inc., New York.

METCALFE, R. M. AND D. R. BOGGS [1976]. "Ethernet: Distributed packet switching for local computer networks," *Comm. ACM* 19:7 (July) 395–404.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 594

NEWMAN, W. N. AND R. F. SPROULL [1979]. *Principles of Interactive Computer Graphics*, 2nd ed., McGraw-Hill, New York.

OUSTERHOUT, J. K. ET AL. [1985]. "A trace-driven analysis of the UNIX 4.2 BSD file system," *Proc. Tenth ACM Symposium on Operating Systems Principles*, Orcas Island, Wash., 15–24.

PATTERSON, D. A., G. A. GIBSON, AND R. H. KATZ [1987]. "A case for redundant arrays of inexpensive disks (RAID)," Tech. Rep. UCB/CSD 87/391, Univ. of Calif. Also appeared in *ACM SIGMOD Conf. Proc.*, Chicago, Illinois, June 1–3, 1988, 109–116.

ROBINSON, B. AND L. BLOUNT [1986]. "The VM/HPO 3880-23 performance results," IBM Tech. Bulletin, GG66-0247-00 (April), Washington Systems Center, Gathersburg, Md.

SALEM, K. AND H. GARCIA-MOLINA [1986]. "Disk striping," *IEEE 1986 Int'l Conf. on Data Engineering*.

SCRANTON, R. A., D. A. THOMPSON, AND D. W. HUNTER [1983]. "The access time myth," Tech. Rep. RC 10197 (45223) (September 21), IBM, Yorktown Heights, N.Y.

SMITH, A. J. [1985]. "Disk cache—miss ratio analysis and design considerations," *ACM Trans. on Computer Systems* 3:3 (August) 161–203.

SMOTHERMAN , M. [1989]. "A sequencing-based taxonomy of I/O systems and review of historical machines," *Computer Architecture News* 17:5 (September) 5–15.

SUTHERLAND, I. E. [1963]. "Sketchpad: A man-machine graphical communication system," *Spring Joint Computer Conf.* 329.

THACKER, C. P., E. M. MCCREIGHT, B. W. LAMPSON, R. F. SPROULL, AND D. R. BOGGS [1982]. "Alto: A personal computer," in *Computer Structures: Principles and Examples*, D. P. Siewiorek, C. G. Bell, and A. Newell, eds., McGraw-Hill, New York, 549–572.

THADHANI, A. J. [1981]. "Interactive user productivity," *IBM Systems J.* 20:4, 407–423.

THISQUEN, J. [1988]. "Seek time measurements," *Amdahl Peripheral Products Division Tech. Rep.* (May).

EXERCISES

9.1 <9.10> [10/25/10] Using the formulas in Figure 9.39 (page 558):

a. [10] Calculate the seek time for moving the arm one-third of the cylinders for both disks.

b. [25] Write a program to calculate the "average" seek time by estimating the time for all possible seeks using these formulas and then dividing by the number of seeks.

c. [10] How close does (a) approximate (b)?

9.2 <9.10> [15/20] Using the formulas in Figure 9.39 (page 558) and the statistics in Figure 9.40 (page 559), calculate the average seek distance and the average seek time on the IBM 3380J. Use the midpoint of a range as the seek distance. For example, use 98 as the seek distance for the entry representing 91–105 in Figure 9.40. For the business workload, just ignore the missing 5% of the seeks. For the UNIX workload, assume the missing 15% of the seeks have an average distance of 300 cylinders.

a. [15] If you were misled by the fallacy, you might calculate the average distance as 884/3. What is the measured distance for each workload?

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 595

b. [20] The time to seek 884/3 cylinders on the IBM 3380J is about 12.8 ms. What is the average seek time for each workload on the IBM 3380J using the measurements?

9.3 <1.4,8.4,9.4> [20/10/Discussion] Assume the improvements in density of DRAMs and magnetic disks continue as predicted in Figure 1.5 (page 17). Assuming that the improvement in cost per megabyte tracks the density improvements and that 1990 is the start of the 4-megabit DRAM generation, when will the cost per megabyte of DRAM equal the cost per megabyte of magnetic disk given:

- The cost difference in 1990 is that DRAM is 10 times more expensive.
- The cost difference in 1990 is that DRAM is 30 times more expensive.

a. [20] Which generation of DRAM chip—measured in bits per chip—will reach equity for each cost difference assumption? What year will that occur?

b. [10] What will be the difference in cost in the previous generation?

c. [Discussion] Do you think the cost difference in the previous generation is sufficient to prevent disks being replaced by DRAMs?

9.4 <9.2> [12/12/12] Assume a workload takes 100 seconds total, with the CPU taking 70 seconds and I/O taking 50 seconds.

a. [12] Assume that the floating-point unit is responsible for 25 seconds of the CPU time. You are considering a floating-point accelerator that goes five times faster. What is the time of the workload for maximum overlap, scaled overlap, and no overlap?

b. [12] Assume that seek and rotational delay of magnetic disks are responsible for 10 seconds of the I/O time. You are considering replacing the magnetic disks with solid state disks that will remove all the seek and rotational delay. What is the time of the workload for maximum overlap, scaled overlap, and no overlap?

c. [12] What is the time of the workload for scaled overlap if you make both changes?

9.5–9.9 Transaction-processing performance. The I/O bus and memory system of a computer are capable of sustaining 100 MB/sec without interfering with the performance of an 80-MIPS CPU (costing $50,000). Here are the assumptions about the software:

- Each transaction requires 2 disk reads plus 2 disk writes.
- The operating system uses 15,000 instructions for each disk read or write.
- The database software executes 40,000 instructions to process a transaction.
- The transfer size is 100 bytes.

You have a choice of two different types of disks:

- A 2.5-inch disk that stores 100 MB and costs $500.
- A 3.5-inch disk that stores 250 MB and costs $1250.
- Either disk in the system can support on average 30 disk reads or writes per second.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 596

Answer the questions below using the TP-1 benchmark in Section 9.3. Assume that the requests are spread evenly to all the disks, that there is no waiting time due to busy disks, and that the account file must be large enough to handle 1000 TPS according to the benchmark ground rules.

9.5 <9.3,9.4> [20] How many TP-1 transactions per second are possible with each disk organization, assuming that each uses the minimum number of disks to hold the account file?

9.6 <9.3,9.4> [15] What is the system cost per transaction per second of each alternative for TP-1?

9.7 <9.3,9.4> [15] How fast a CPU makes the 100 MB/sec I/O bus a bottleneck for TP-1? (Assume that you can continue to add disks.)

9.8 <9.3,9.4> [15] As manager of MTP (Mega TP), you are deciding whether to spend your development money building a faster CPU or improve the performance of the software. The database group says they can reduce a transaction to 1 disk read and 1 disk write and cut the database instructions per transaction to 30,000. The hardware group can build a faster CPU that sells for the same amount of the slower CPU with the same development budget. (Assume you can add as many disks as needed to get higher performance.) How much faster does the CPU have to be to match the performance gain of the software improvement?

9.9 <9.3,9.4> [15/15] The MTP I/O group was listening at the door during the software presentation. They argue that advancing technology will allow CPUs to get faster without significant investment, but that the cost of the system will be dominated by disks if they don't develop new faster 2.5-inch disks. Assume the next CPU is 100% faster at the same cost and that the new disks have the same capacity as the old ones.

a. [15] Given the new CPU and the old software, what will be the cost of a system with enough old 2.5-inch disks so that they do not limit the TPS of the system ?

b. [15] Now assume you have as many new disks as you had old 2.5 inch disks in the original design. How fast must the new disks be (I/Os per second) to achieve the same TPS rate with the new CPU as the system in part a? What will the system cost?

9.10 <9.4> [20/20/20] Assume that we have the following two magnetic-disk configurations: a single disk and an array of four disks. Each disk has 20 surfaces, 885 tracks per surface with 16 sectors/track, each sector holds 1K bytes, and it revolves at 3600 RPM. Using the seek-time formula, for the IBM 3380D in Figure 9.39 (page 558). The time to switch between surfaces is the same as to move the arm one track. In the disk array all the spindles are synchronized—sector 0 in every disk rotates under the head at the exact same time—and the arms on all four disks are always over the same track. The data is "striped" across all 4 disks, so four consecutive sectors on a single disk system will be spread one sector per disk in the array. The delay of the disk controller is 2 ms per transaction, either for a single disk or for the array. Assume the performance of the I/O system is limited only by the disks and that there is a path to each disk in the array.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 597

Compare the performance in both I/Os per second and megabytes per second of these two disk organizations assuming the following request patterns:

a. [20] Random reads of 4 KB of sequential sectors. Assume the 4 KB are aligned under the same arm on each disk in the array.

b. [20] Reads of 4 KB of sequential sectors where the average seek distance is 10 tracks. Assume the 4 KB are aligned under the same arm on each disk in the array.

c. [20] Random reads of 1 MB of sequential sectors. (If it matters, assume the disk controller allows the sectors to arrive in any order.)

9.11 [20] <9.4> Assume that we have one disk defined as in Exercise 9.9. Assume that we read the next sector after any read and that *all* read requests are one sector in length. We store the extra sectors that were read ahead in a *disk cache*. Assume that the probability of receiving a request for the sector we read ahead at some time in the future (before it must be discarded because the disk-cache buffer fills) is 0.1. Assume that we must still pay the controller overhead on a disk-cache read hit, and the transfer time for the disk cache is 250 ns per word. Is the read-ahead strategy faster? (Hint: Solve the problem in the steady state by assuming that the disk cache contains the appropriate information and a request has just missed.)

9.12–9.14 Assume the following information about our DLX machine:

Loads 2 cycles

Stores 2 cycles

All other instructions are 1 cycle. Use the summary instruction mix information in Figure C.4 in Appendix C on DLX for GCC.

Here are the cache statistics for a write-through cache:

- Each cache block is four words, and the whole block is read on any miss.
- Cache miss takes 13 cycles.
- Write through takes 6 cycles to complete, and there is no write buffer.

Here are the cache statistics for a write-back cache:

- Each cache block is four words, and the whole block is read on any miss.
- Cache miss takes 13 cycles for a clean block and 21 cycles for a dirty block.
- Assume that on a miss, 30% of the time the block is dirty.

Assume that the bus

- is only busy during transfers,
- transfers on average 1 word / clock cycle, and
- must read or write a single word at a time (it is not faster to read or write two at once).

9.12 [20/10/20/20] <9.4,9.5,9.6> Assume that DMA I/O can take place simultaneously with CPU cache hits. Also assume that the operating system can guarantee that there will be no stale-data problem in the cache due to I/O. The sector size is 1 KB.

a. [20] Assume the cache miss rate is 5%. On the average, what percentage of the bus is used for each cache write policy? This measured is called the *traffic ratio* in cache studies.

b. [10] If the bus can be loaded up to 80% of capacity without suffering severe performance penalties, how much memory bandwidth is available for I/O for each cache write policy? The cache miss rate is still 5%.

c. [20] Assume that a disk sector read takes 1000 clock cycles to initiate a read, 100,000 clock cycles to find the data on the disk, and 1000 clock cycles for the DMA to transfer the data to memory. How many disk reads can occur per million instructions executed for each write policy? How does this change if the cache miss rate is cut in half?

d. [20] Now you can have any number of disks. Assuming ideal scheduling of disk accesses, what is the maximum number of sector reads that can occur per million instructions executed?

9.13 [20/20] <9.4,9.5> Most machines today have a separate frame buffer to update the screen to avoid slowing down the memory system. An interesting issue is the percentage of the memory bandwidth that would be used if there were no frame buffer. Assume that all accesses to the memory are the size of a full cache block and they all take the time of a cache miss. The refresh rate is 60 Hz. Using the information in Section 9.4, calculate the memory traffic for the following graphics devices:

1. A 340 by 540 black-and-white display.
2. A 1280 by 1024 color display with 24 bits of color.
3. A 1280 by 1024 color display using a 256-word color map.

Assume the clock rate of the CPU is 60 MHz.

a. [20] What percentage of the memory/bus bandwidth do each of the three displays consume?

b. [20] Suppose instead of the bus and main memory being 32 bits wide that both are 512 bits wide. How long should a memory access take now using the wider bus? What percentage of memory bandwidth is now used by each display?

9.14 [20] <9.4,9.9> The IBM 3990 I/O Subsystem storage director can have a large cache for reads and writes. Assume the cache costs the same as four 3380D disks. What hit rate must the cache achieve to get the same performance as four more 3380D disks? (See Figure 9.15 (page 517) for 3380 performance.) Assume the cache could support 5000 I/Os per second if everything hit the cache.

9.15 [50] <9.3, 9.4> Take your favorite computer and write three programs that achieve the following:

1. Maximum bandwidth to and from disks

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 599

2. Maximum bandwidth to a frame buffer

3. Maximum bandwidth to and from the local area network

What is the percentage of the bandwidth that you achieve compared to what the I/O device manufacturer claims? Also record CPU utilization in each case for the programs running separately. Next run all three together and see what percentage of maximum bandwidth you achieve for three I/O devices as well as the CPU utilization. Try to determine why one gets a larger percentage than the others.

9.16 [40] <9.2> The system speedup formulas are limited to one or two types of devices. Derive simple to use formulas for unlimited numbers of devices, using as many different assumptions on overlap that you can handle.

9.17 [Discussion] <9.2> What are arguments for predicting system performance using maximum overlap, scaled overlap, and nonoverlap? Construct scenarios where each one seems most likely and other scenarios where each interpretation is nonsensical.

9.18 [Discussion] <9.11> What are the advantages and disadvantages of a minimal buffer I/O system like that used by IBM versus a maximal buffer I/O system on I/O system cost/performance?

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 600

The turning away from the conventional organization came in the middle 1960's, when the law of diminishing returns began to take effect in the effort to increase the operational speed of a computer. . . . Electronic circuits are ultimately limited in their speed of operation by the speed of light . . . and many of the circuits were already operating in the nanosecond range.

Bouknight et al. [1972]

. . . sequential computers are approaching a fundamental physical limit on their potential computational power. Such a limit is the speed of light . . .

A. L. DeCegama, *The Technology of Parallel Processing, Volume I* (1989)

... today's machines ... are nearing an impasse as technologies approach the speed of light. Even if the components of a sequential processor could be made to work this fast, the best that could be expected is no more than a few million instructions per second.

Mitchell [1989]

10.1	**Introduction**	**571**
10.2	**Flynn Classification of Computers**	**572**
10.3	**SIMD Computers—Single Instruction Stream, Multiple Data Streams**	**572**
10.4	**MIMD Computers—Multiple Instruction Streams, Multiple Data Streams**	**574**
10.5	**The Roads to El Dorado**	**576**
10.6	**Special-Purpose Processors**	**580**
10.7	**Future Directions for Compilers**	**581**
10.8	**Putting It All Together: The Sequent Symmetry Multiprocessor**	**582**
10.9	**Fallacies and Pitfalls**	**585**
10.10	**Concluding Remarks—Evolution Versus Revolution in Computer Architecture**	**587**
10.11	**Historical Perspective and References**	**588**
	Exercises	**592**

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.

10 Future Directions

10.1 Introduction

In the first nine chapters we limited ourselves to ideas that have proven themselves in the marketplace. Yet the principles of these chapters can be found in the first paper on stored-program computers. The quotes on the facing page suggest that the days of the traditional computer are numbered. For a dated model of computation it has surely demonstrated its viability! Today it is improving in performance faster than at any time in its history, and the improvement in cost and performance since 1950 has been five orders of magnitude. Had the transportation industry kept pace with these advances, we could travel from San Francisco to New York in one minute for one dollar!

In this last chapter we abandon our conservative perspective and speculate about the future of computer architecture and compilers. The goal of innovative designs is dramatic improvements in cost/performance, or highly scalable performance with good cost/performance. Many of the ideas covered here have led to machines that are beginning to compete in the computer marketplace today. Some of them may not be around for the next edition of this book, while others may need their own chapters.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 603

10.2 Flynn Classification of Computers

Flynn [1966] proposed a simple model of categorizing all computers. He looked at the parallelism in the instruction and data streams called for by the instructions at the most constrained component of the machine, and placed all computers in one of four categories:

1. *Single instruction stream, single data stream* (SISD, the uniprocessor)
2. *Single instruction stream, multiple data streams* (SIMD)
3. *Multiple instruction streams, single data stream* (MISD)
4. *Multiple instruction streams, multiple data streams* (MIMD)

This is a coarse model, as some machines are hybrids of these categories. Yet in this chapter we stick with this classic model because it is simple, easy to understand, gives a good first approximation, and—perhaps because of ease of understanding—is also the most widely used scheme.

Your first question about the model should be, "Single or multiple compared to what?" A machine that can add a 32-bit number in one clock cycle would seem to have multiple data streams when compared to a bit-serial computer that takes 32 clock cycles for the same operation. Flynn chose popular computers of that day, the IBM 704 and IBM 7090, as the model of SISD, although today any of the machines in Chapter 4 would serve as the example.

Having thus established the reference point for SISD, the next class is SIMD.

10.3 SIMD Computers—Single Instruction Stream, Multiple Data Streams

The cost of a general multiprocessor is, however, very high and further design options were considered which would decrease the cost without seriously degrading the power or efficiency of the system. The options consist of recentralizing one of the three major components.... Centralizing the [control unit] gives rise to the basic organization of [an]... array processor such as the Illiac IV.

Bouknight et al. [1972]

We have already seen typical instructions for a SIMD machine, yet the machine is not SIMD. The vector instructions of Chapter 7 operate on several data elements within a single instruction, executing in pipelined fashion in a single functional unit. Unlike SIMD, many functional units are not being invoked by a single instruction. A true SIMD would have, say, 64 data streams simultaneously going to 64 ALUs to form 64 sums within the same clock cycle.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 604

The virtues of SIMD are that all the parallel execution units are synchronized and that they all respond to a single instruction from a single PC. From a programmer's perspective, this is close to the already familiar SISD. The original motivation for SIMD was to amortize the cost of the control unit over dozens of execution units. A more recently observed advantage is the reduced size of program memory—SIMD needs only one copy of the code being simultaneously executed, while MIMD needs a copy in every processor. Hence, the cost of program memory for a large number of execution units is less for SIMD.

Like vector machines, real SIMD computers have a mixture of SISD and SIMD instructions. There is a SISD host computer to perform operations such as branches or address calculation that do not need massive parallelism. The SIMD instructions are broadcast to all the execution units, each of which has its own set of registers. Also, as in vector machines, individual execution units can be disabled during a SIMD instruction. Unlike vector machines, massively parallel SIMD machines rely on interconnection or communication networks to exchange data between processing elements.

SIMD works best when vector instructions work best—in dealing with arrays in for-loops. Hence, to have the opportunity for massive parallelism in SIMD there must be massive amounts of data, or *data parallelism*. SIMD is at its weakest in case statements, where each execution unit must perform a different operation on its data, depending on what data it has. The execution units with the wrong data are disabled so that the proper units can continue. Such situations essentially run at 1/*n*th performance, where n is the number of cases.

The basic tradeoff in SIMD machines is performance of a processor versus number of processors. The machines in the marketplace today emphasize a large degree of parallelism over performance of the individual processors. The Connection Machine 2, for example, offers 65,536 single bit-wide processors while the ILLIAC IV had 64 64-bit processors.

While MISD fills out Flynn's classification, it is difficult to envision. A single instruction stream is simpler than multiple instruction streams, but multiple instruction streams with multiple data streams are easier to imagine than multiple instructions with a single data stream. A few of the architectures we have covered might be considered MISD: superscalar and VLIW architectures of Chapter 6 (Section 6.8) often have a single data stream and multiple instructions, although these machines have a single program counter. Perhaps closer to the mark are the decoupled architectures (pages 321–322), which have two instruction streams with independent program counters and a single data stream. Systolic architectures, covered in Section 10.6, might also be considered MISD.

While we can find examples of SIMD and MISD, their number is dwarfed by the multitude of MIMD machines.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 605

10.4 MIMD Computers—Multiple Instruction Streams, Multiple Data Streams

Multis are a new class of computers based on multiple microprocessors. The small size, low cost, and high performance of microprocessors allow design and construction of computer structures that offer significant advantages in manufacture, price-performance ratio, and reliability over traditional computer families.... Multis are likely to be the basis for the next, the fifth, generation of computers.

Bell [1985, 463]

Practically since the first working computer, architects have been striving for the El Dorado of computer design: To compose a powerful computer by simply connecting many existing smaller ones. The user orders as many CPUs as he can afford and gets a commensurate amount of performance. Other advantages of MIMD may be highest absolute performance, faster than the largest uniprocessor, and highest reliability/availability (page 520) via redundancy.

For decades, computer designers have been looking for the missing piece of the puzzle that allows this speedup to happen, as if by magic. People are heard making statements that begin "Now that computers have dropped to such a low price..." or "This new interconnection scheme will overcome the scaling problem, so..." or "As this new programming language becomes widespread...," and end with "MIMDs will (finally) dominate computing."

With so many attempts to use parallelism, there are a few terms that are useful to know when discussing MIMDs. The principal division is that which delineates how information is shared. *Shared-memory* processors offer the programmer a single memory address that all processors can access; cache-coherent multiprocessors are shared-memory machines (see Sections 8.8 and 10.8). Processes communicate through shared variables in memory, with loads and stores capable of accessing any memory location. Synchronization must be available to coordinate processes. An alternative model to sharing data is where processes communicate by sending messages. As an extreme example, processes on different workstations communicate by sending messages over a local area network. This communication distinction is so fundamental that Bell suggests the term *multiprocessor* be limited to MIMDs that can communicate via shared memory, while MIMDs that can only communicate via explicit message passing should be called *multicomputers*. Since a portion of a shared memory could be used for messages, most multiprocessors can efficiently execute message-passing software. A multicomputer might be able to simulate shared memory by sending a message for every load or store, but presumably this would run excruciatingly slowly. Thus, Bell's distinction is based on the underlying hardware and program execution model, reflected in the performance of shared-memory communication, as opposed to the software that might run on a machine. Message-passing docents question the *scalability* of multiprocessors, while

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 606

shared-memory advocates question the programmability of multicomputers. The next section examines this debate further.

The good news is that after many assaults, MIMD has established a beachhead. Today it is generally agreed that a multiprocessor may be more effective for a timesharing workload than a SISD. No single program takes less CPU time, but more independent tasks can be completed per hour—a throughput versus latency argument. Not only are start-up companies like Encore and Sequent selling small-scale multiprocessors, but the high-end machines from IBM, DEC, and Cray Research are multiprocessors. This means multiprocessors now embody a significant market, responsible for a majority of the mainframes and virtually all supercomputers. The only disappointment to computer architects is that shared memory is practically irrelevant for user programs run on the machine, with the operating system being the only benefactor. The development of a multiprocessor's operating system, particularly its resource manager, is simplified by shared memory.

The bad news is that it remains to be seen how many important applications run faster on MIMDs. The difficulty has not lain in the prices of SISDs, in flaws in topologies of interconnection networks, or in programming languages; but in the lack of applications software that have been reprogrammed to take advantage of many processors to complete important tasks sooner. Since it has been even harder to find applications that can take advantage of many processors, the challenge is greater for large scale MIMDs. When the positive gains from timesharing are combined with the scarcity of highly parallel applications, we can appreciate the predicament facing computer architects designing large-scale MIMDs that do not support timesharing.

But why is this so? Why should it be so much harder to develop MIMD programs than sequential programs? One reason is that it is hard to write MIMD programs that achieve close to linear speedup as the number of processors dedicated to the task increases. As an analogy, think of the communication overhead for a task done by one person versus the overhead for a task done by a committee, especially as the size of the group increases. While n people may have the potential to finish any task n times faster, the communication overhead for the group can prevent it from achieving this; this becomes especially hard as n increases. (Imagine the change in communication overhead going from 10 people to 1,000 people to 1,000,000.) Another reason for the difficulty in writing parallel programs is how much the programmer must know about the hardware. On a uniprocessor, the high-level language programmer writes his program ignoring the underlying machine organization—that's the job of the compiler. For a multiprocessor today, the programmer had better know the underlying hardware and organization if he is to write fast and scalable programs. This intimacy also makes portable parallel programs rare. Though this second obstacle may lessen over time, it is now the biggest challenge facing computer science. Finally, from Chapter 1 comes Amdahl's Law (page 8) to remind us that even small parts of a program must be parallelized to reach the full

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 607

potential. Thus, coming close to linear speedup involves inventing new algorithms that are inherently parallel.

Example

Suppose you want to achieve linear speedup with 100 processors. What fraction of the original computation can be sequential?

Answer

Amdahl's Law is

$$\text{Speedup} = \frac{1}{(1-\text{Fraction}_{\text{enhanced}}) + \frac{\text{Fraction}_{\text{enhanced}}}{\text{Speedup}_{\text{enhanced}}}}$$

Substituting for the goal of linear speedup with 100 processors gives:

$$100 = \frac{1}{(1-\text{Fraction}_{\text{enhanced}}) + \frac{\text{Fraction}_{\text{enhanced}}}{100}}$$

Solving for percentage converted to enhanced mode:

$$100 - 100 * \text{Fraction}_{\text{enhanced}} + 1 * \text{Fraction}_{\text{enhanced}} = 1$$

$$-99 * \text{Fraction}_{\text{enhanced}} = -99$$

$$\text{Fraction}_{\text{enhanced}} = 1$$

Thus, to achieve linear speedup with 100 processors, **none** of the original computation can be sequential. Put another way, to get a speedup of 99 from 100 processors means the sequential fraction of the original program had to be about 0.0001.

The example above demonstrates the need for new algorithms. This underlines the authors' belief that major successes in using large-scale parallel machines of the 1990s are possible for those who understand applications, algorithms, and architecture.

10.5 The Roads to El Dorado

Figure 10.1 shows the state of the industry, plotting number of processors versus performance of an individual processor. The massive parallelism question is whether taking the high road or the low road in Figure 10.1 will get us to El Dorado. Currently we don't know enough about parallel programming and applications to be able to quantitatively trade-off number of processors versus performance per processor to achieve the best cost/performance.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 608

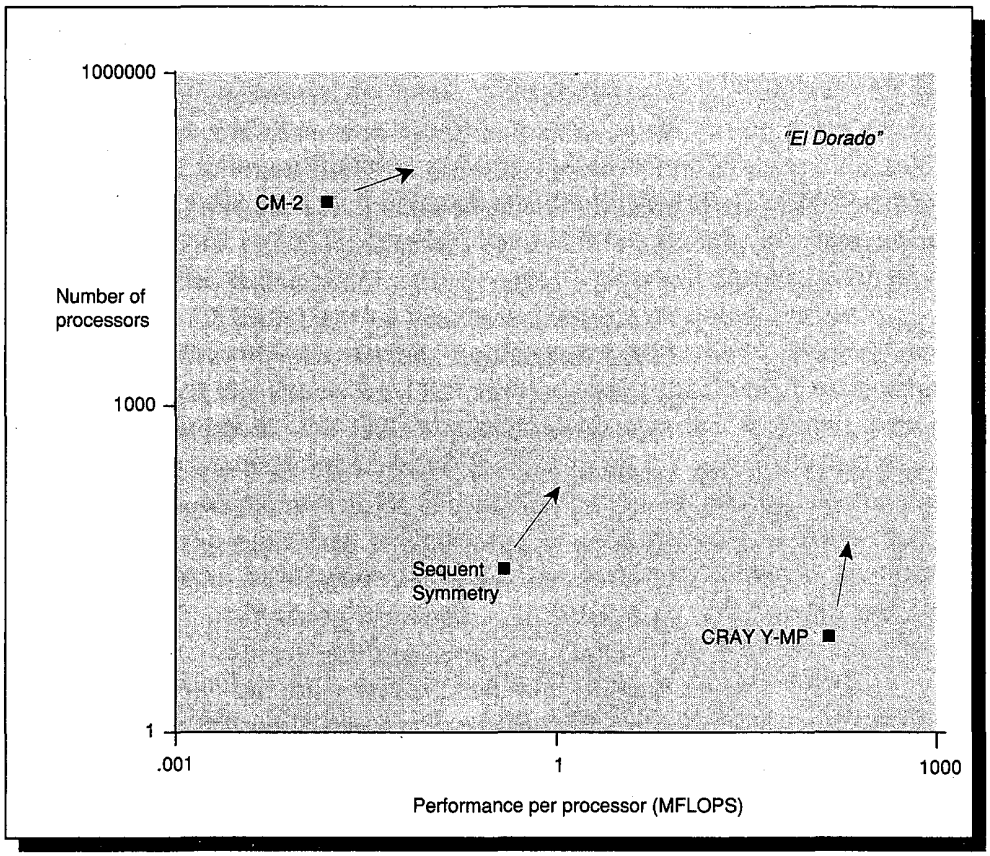

FIGURE 10.1 Danny Hillis, architect of the Connection Machines, has used a figure similar to this to illustrate the multiprocessor industry. (Hillis's x axis was processor width rather than processor performance.) Processor performance on this graph is approximated by the MFLOPS rating of a **single** processor for the DAXPY procedure of the Linpack benchmark for a 1000 x 1000 matrix. Generally, it is easier for programmers when moving to the right , while moving up is easier for the hardware designer because there is more hardware replication. The massive parallelism question is, "Which is the quickest path to the upper right corner?" The computer design question is, "Which has the best cost/performance or is more scalable for equivalent cost/performance?"

It is interesting to note that very different changes are required to improve performance depending on whether you take the low road or the high road in this figure. Since most programs are written in high-level languages, moving along the horizontal direction (increasing performance per processor) is almost entirely a matter of improving the hardware. The applications are unchanged, with compilers adapting them to the more powerful processor. Hence, increasing processor performance versus number of processors is easier for the applications software. Improving performance by moving in the vertical direction (increasing parallelism), on the other hand, may involve significant changes to applications, since programming ten processors may be very different from programming a thousand, and different yet again from programming a million. (But going from

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 609

100 to 101 is probably not different.) An advantage of the vertical path to performance is that the hardware may be simply replicated—the processors in particular, but also the hardware of the interconnection switch. Hence, increasing number of processors versus processor performance results in more hardware replication. An advantage of the low road is that it is much more likely that there will be a market at the various points along the way to El Dorado. In addition, those who take the high road must grapple with Amdahl's Law.

This brings us to a fundamental debate about the organization of memory in large-scale machines of the future. The debate unfortunately often centers on a false dichotomy: *shared memory* versus *distributed memory*. Shared memory means a single address space, implying implicit communication. The real opposite to a shared address is *multiple private address spaces*, implying explicit communication. Distributed memory refers to the location of the memory. If physical memory is divided into modules with some placed near each processor (which allows faster access time to that memory), then physical memory is distributed. The real opposite of distributed memory is *centralized memory*, where access time to a physical memory location is the same for all processors.

Clearly shared address versus multiple address and distributed memory versus centralized memory are orthogonal issues: SIMDs or MIMDs can have a shared address and a distributed physical memory or multiple private address spaces and a centralized physical memory (although this last combination would be unusual). Figure 10.2 categorizes several machines by these axes. The proper debates concerning the future are the pros and cons of a single address and the pros and cons of distributed memory.

The single address debate is closely tied to the model of communication, since shared-address machines must offer implicit communication (possibly

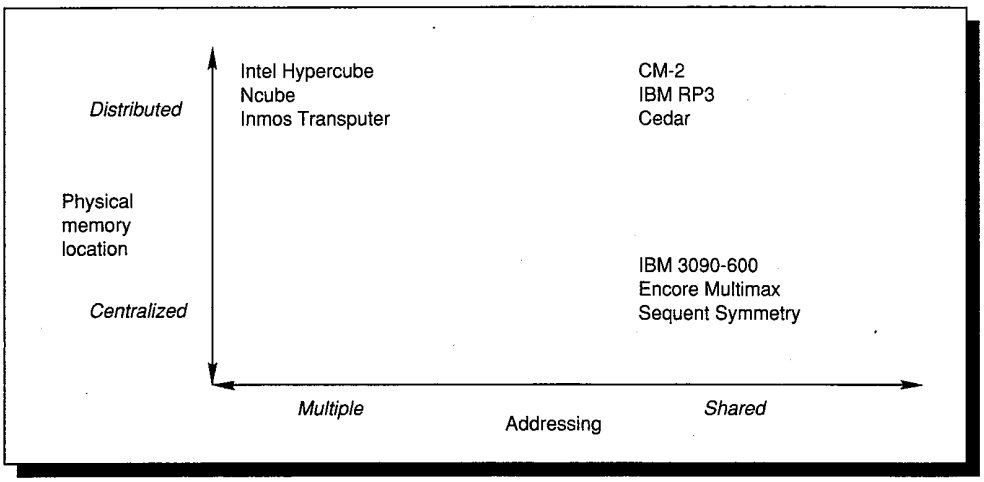

FIGURE 10.2 Parallel processors placed according to centralized versus distributed memory and shared versus multiple addressing. In general it is easier for software for machines on the shared side of the addressing axis and it is easier to build larger-scale machines on the distributed end of the vertical access. These machines in the graph are described in Section 10.11.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 610

part of any memory access) and multiple-address machines must have explicit communication. (It is not quite that simple since some shared-address machines also offer explicit communication in various forms.) "Implicitists" knock "explicitists" for advocating machines that are harder to program when it is already hard to find applications: Why make the programmer's life more difficult when software is the linchpin of large-scale parallelism? One reply is that if memory is distributed, as processors get faster the time to remote memory will be so long—say 50 to 100 clock cycles—the compiler or programmer must be aware he is writing for a large-scale parallel machine no matter which communication scheme is used. Explicit communication also offers the possibility of hiding the cost of communication by overlapping it with computation. The implicitist reply is that using hardware rather than explicit instructions reduces the overhead of communication. Moreover, a single address means processes can use pointers and communicate data only if the pointer is dereferenced, while explicit communication means the data must be sent in the presence of pointers since the data **might** be accessed. The explicitist rebuttal is the owner of the data can send the data, traversing a properly designed network only once, while in shared-memory machines a processor requests the data and then the owner returns it, requiring two trips over the communications network.

Distributed-memory advocates argue that no matter how much caching is placed in front of a single central memory, it has limited bandwidth, and thus, limits the number of processors. Central-memory advocates raise the question of efficiency: If there is not enough parallelism to use many processors, then why distribute memory? Centralists also point out that distributed memory increases the difficulty of programming, since now the programmer or the compiler must decide how to lay out the data in the physical memory modules so as to reduce communication. Hence, distributed memory introduces the concept of data elements being near a processor (the module taking less time to access) or far (in other memory modules).

We can now explain a difficulty of the distributed versus centralized dichotomy. Every processor will likely have a cache, which is in some sense a distributed memory no matter how main memory is organized. Even with caches, the latency of a miss and the effective bandwidth for satisfying cache requests can be improved if data is allocated to the memory module near the appropriate cache. Hence, there is still a distinction between centralized and distributed main memory in the presence of caches.

As you can imagine, these debates continue back and forth, practically interminably. Fortunately, in computer architecture such disagreements are settled by measurements rather than polemics. Thus, time will be the judge of these issues, but your authors will be the judge of a bet inspired by these debates (see page 590 in 10.11).

The real issues for future machines are these: Do problems and algorithms with sufficient parallelism exist? And can people be trained or compilers be written to exploit such parallelism?

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 611

10.6 Special-Purpose Processors

In addition to exploring parallelism, many designers today are exploring special-purpose computers. With the increasing sophistication of *computer-aided design* software and increasing capacity per chip comes the opportunity of quickly building a chip that does one thing well at low cost. Real-time speech recognition and image processing are examples. Such special-purpose devices, or *coprocessors*, frequently act in conjunction with the CPU. There are two types in the coprocessor trend: digital signal processors and systolic arrays.

Digital signal processors (or DSPs) are not derived from the traditional model of computing, and tend to look like horizontal microprogrammed machines (see page 212) or VLIW machines (see pages 322–325). They tend to solve real-time problems, essentially having an infinite-input data stream. There has been little emphasis on compiling from programming languages such as C, but that is starting to change. As DSPs bend to the demands of programming languages, it will be interesting to see how they differ from traditional microprocessors.

Systolic arrays evolved from attempts to get more efficient computing bandwidth from silicon. Systolic arrays can be thought of as a method for designing special-purpose computers to balance resources, I/O bandwidth, and computation. Relying on pipelining, data flows in stages from memory through an array of computation units and back to memory, as suggested in Figure 10.3. Recently, systolic-array research has moved away from many, dedicated special-purpose chips to fewer, more powerful chips that are programmable.

The authors expect an increasing role for special-purpose computers in the 1990s because they offer both higher performance and lower cost for dedicated functions such as real-time speech recognition and image processing. The consumer marketplace seems the most likely candidate, given its high volume and sensitivity to cost.

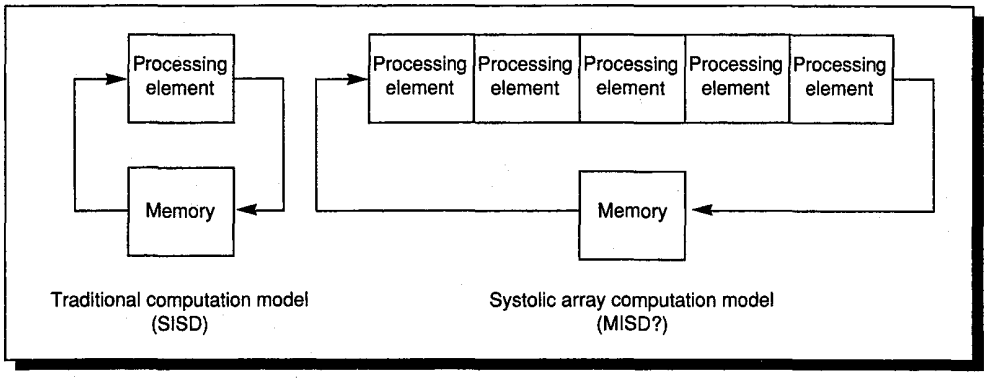

FIGURE 10.3 The systolic architecture gets its name from the heart rhythmically pumping blood. Data arrives at a processing element at regular intervals, where it is modified and passed to the next element, and so on, until it circulates back to memory. Some consider systolic arrays an example of MISD.

10.7 Future Directions for Compilers

Compilers of the future have two challenges on machines of the future:

- Lay out of data to reduce memory hierarchy and communication overhead, and
- Exploitation of parallelism.

Programs of the future will spend a larger percentage of the execution time waiting for the memory hierarchy as the gap grows between the clock cycle time of processors and the access time of main memory (see Figure 8.18, page 427). Compilers that arrange code and data so as to reduce cache misses may lead to larger performance improvements than traditional optimizations of today. Further improvements are possible with the possibility of prefetching data into a cache before it is needed by the program. One interesting proposition is by extending existing programming languages with array operations a programmer can express parallelism with calculations on entire arrays at a time, leaving it up to the compiler to lay out the data into processors to reduce the amount of communication. For example, the proposed extension to FORTRAN 77 called FORTRAN 8X includes array extensions. The hope is that the programmer's task might even be simpler than with SISD machines where array operations must be specified with loops. The range of programs that such a compiler can handle efficiently and the number of hints a programmer must supply on where to place data will determine the practical value of this proposal.

In addition to reducing the costs of memory access and communication, compilers may change performance by factors of two or three by utilizing parallelism available in the processor. Figure 2.25 (page 75) shows the Perfect Club benchmarks operate at only 1% of peak performance, clearly suggesting many opportunities for software. More specifically, the superscalar machines of Chapter 6 (pages 318–320) typically achieve a speedup of less than 2 using today's compilers, even through the potential performance improvement of executing 4 instructions at once is 4. From Chapter 7 we see that vector machines typically achieve a vectorization rate of 40% to 70%, delivering a speedup of 1.5 to 2.5, where a vectorization rate of 90% could achieve a speedup over 5. And current compilers for multiprocessors are considered successful if they achieve a speedup 3 for a single program when the potential from 8 processors is 8. Figure 10.4 (page 582) shows the potential improvement in performance of a larger percentage of the work executing in the higher-performance mode for each of these categories. Since we can expect multiple processors in machines where each processor has vector or superscalar features, the potential speedup of these factors may be multiplied together.

While this opportunity exists for compilers, we do not want to belittle its difficulty. Parallelizing compilers have been under development since 1975 but progress has been slow. These problems are hard, especially for the "dusty deck"

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 613

challenge of running existing programs. Success has been limited to programs where the parallelism is available in the algorithm and expressed in the program and to machines with a small number of processors. Significant progress may eventually require new programming languages as well as smarter compilers!

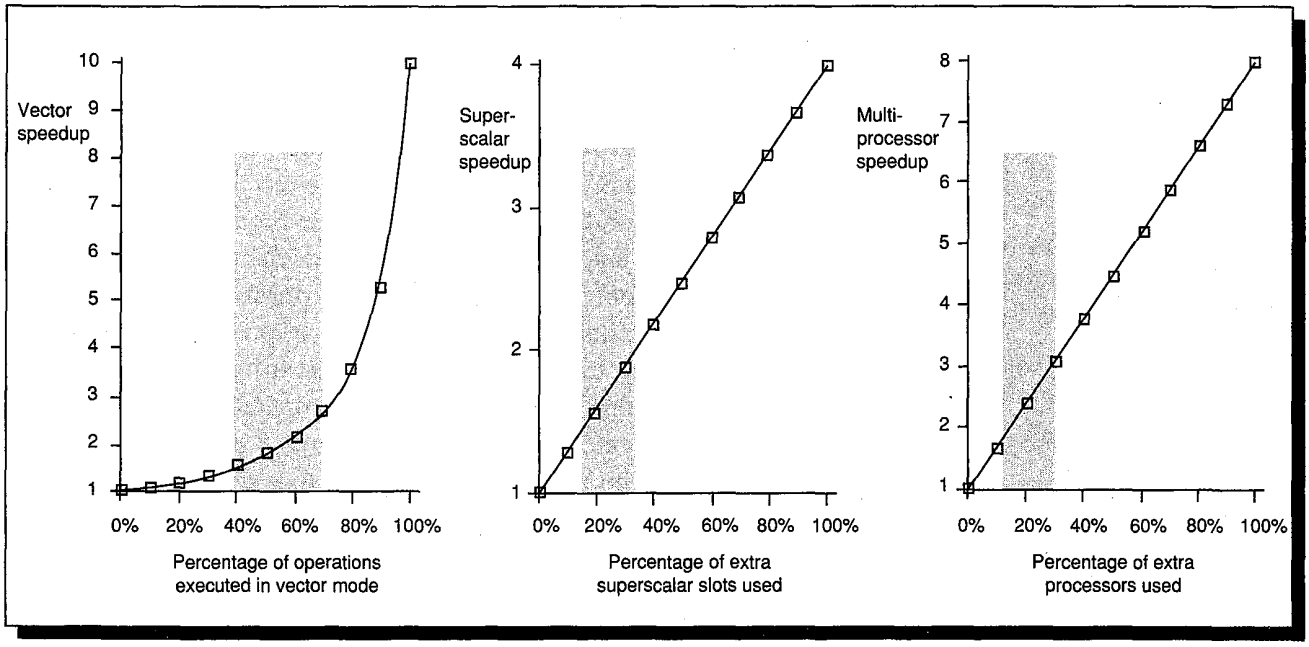

FIGURE 10.4 Potential for performance improvement by compilers transforming more of the computation into the faster mode. The leftmost graph shows the percentage of operations executed in vector mode, while the other graphs show the percentage of the potential speedup in use on average: percentage of four instructions used per cycle in superscalar and percentage of time all eight processors were utilized in the multiprocessor. The gray area shows the range of utilization typically found in programs using current compilers.

10.8 Putting It All Together: The Sequent Symmetry Multiprocessor

The high performance and low cost of the microprocessor inspired renewed interest in multiprocessors in the 1980s. Several microprocessors can be placed on a common bus because:

they are much smaller than multichip processors,

caches can lower bus traffic, and

coherency protocols can keep caches and memory consistent.

Traffic per processor and the bus bandwidth determine the number of processors in such a multiprocessor.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 614

Several research projects and companies investigated these shared-bus multiprocessors. One example is Sequent Corporation, founded to build multiprocessors based on standard microprocessors, and the UNIX operating system. The first-generation system was the Balance 8000, offered in 1984 with 2 to 12 National 32032 microprocessors, a 32-bit split transaction bus that multiplexed address and data, and one 8-KB, 2-way–set-associative, write-through cache per processor. Each cache watched the bus to maintain coherency using write through with invalidate. (See Sections 8.4, 8.8, and 9.4 for a review of these terms.) The sustained bandwidth of the main memory and bus is 26.7 MB/sec. Two years later Sequent upgraded to the Balance 21000, offering up to 30 National 32032 microprocessors with the same memory system and bus.

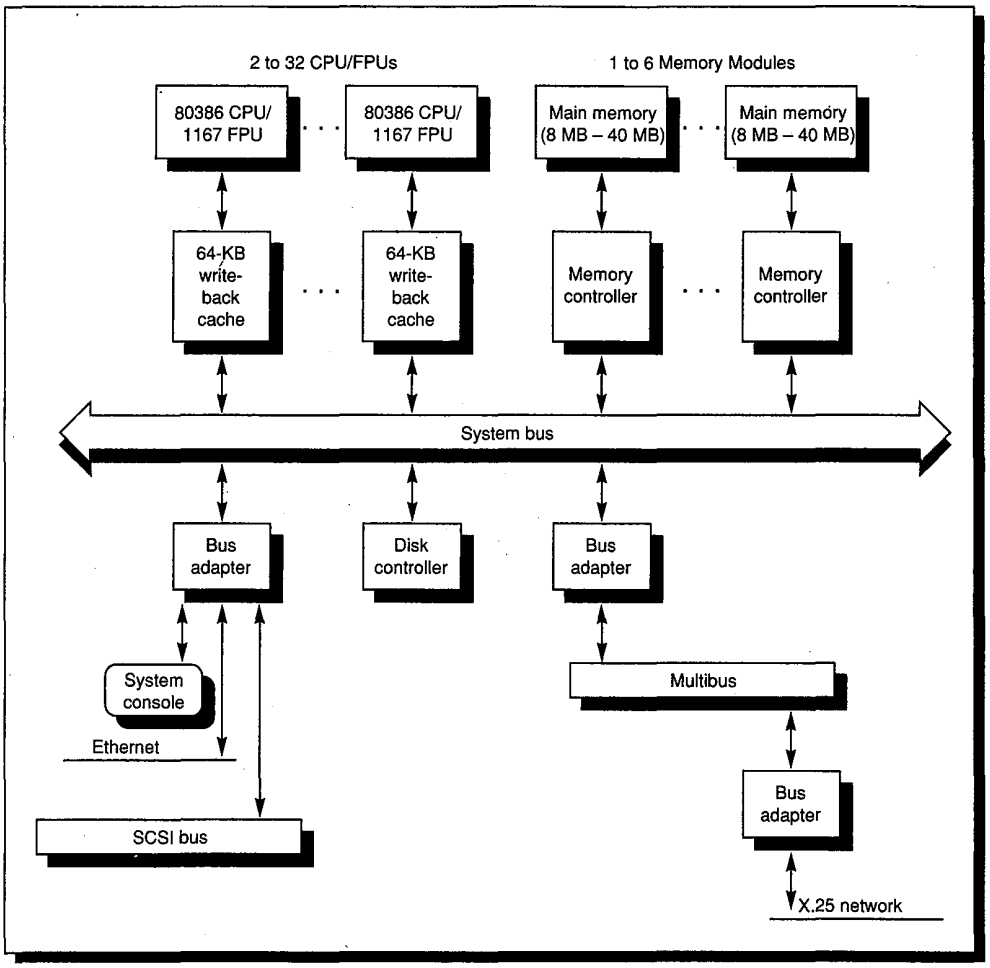

FIGURE 10.5 The Sequent Symmetry multiprocessor has up to 30 microprocessors, each with 64 KB of 2-way set associative, write-back caches connected over the shared system bus. Up to six memory controllers also talk to this 64-bit-wide bus, plus some interfaces for I/O. In addition to a special-purpose disk controller, there is an interface for the system console, Ethernet network, and SCSI I/O bus (see Chapter 9), as well as another interface for Multibus. I/O devices can be attached either to SCSI or to Multibus, as the customer desires. (Although all interfaces are labeled "Bus adapter," each is a unique design.)

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 615

In 1986, Sequent began the design of the Symmetry multiprocessor, assuming a microprocessor 300% to 400% faster than the 32032. The goal was to support as many processors as possible using the I/O controllers developed for the Balance system. This meant the bus had to remain compatible, though the new memory and bus system had to deliver roughly 300% to 400% higher bandwidth than the older system.

The goal of higher memory-system bandwidth with a similar bus was attacked on four levels. First, the cache was increased to 64 KB, increasing the hit rate and therefore the effective memory bandwidth as seen by the processor. Second, the cache policy was changed from write through to write back to reduce the number of write operations on the shared bus. To maintain cache coherency with write back, Symmetry uses a write-invalidate scheme (see pages 468–469). The third change was to double the bus width to 64 bits, thereby doubling the bus bandwidth to 53 MB/sec. The final change was to have each memory controller interleave memory as two banks (see Section 8.8), allowing the memory system to match the bandwidth of the wider bus. The memory system can have up to six controllers with up to 240-MB total main memory.

The use of high-level languages and the portability of the UNIX operating system allowed changing instruction sets to the faster Intel 80386. Running at a higher clock rate, with the faster Weitek 1167 floating-point accelerator, **and** with the improved memory system, a single 80386 ran from 214% to 776% faster for floating-point benchmarks and about 375% faster for integer benchmarks. Figure 10.5 (page 583) shows the organization of the Symmetry.

One of the other design constraints was that the new Symmetry boards had to work properly when put into the old Balance systems. Since the new system was to use write back and the old system used write through, the hardware team solved the problem by designing the new caches to support either write through or write back. Lovett and Thakkar [1988] took advantage of that feature to run parallel programs with both policies. Figure 10.6 shows bus utilization versus the number of processors for four parallel programs.

As mentioned above, bus utilization directly corresponds to the number of processors that can be used in such single-bus systems. Write-through caches should have higher bus utilization for the same number of processors since every write must go over the bus; or from a different perspective, the same bus should be able to support more processors if they use write-back caches. Figure 10.6 fulfills our expectations; the buses saturate with fewer than 16 processors with write through, but write back appears to scale to the full size.

There are two components to the bus traffic: normal misses and coherency support. Uniprocessor misses (compulsory, capacity, and conflict) can be reduced by larger caches and by better write policies, but the coherency traffic is a function of the parallel program. The primary benefit of write back for the programs in Figure 10.6 was simply reducing the number of writes on the bus due to the write-back policy, for there were few writes to shared data in these programs.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 616

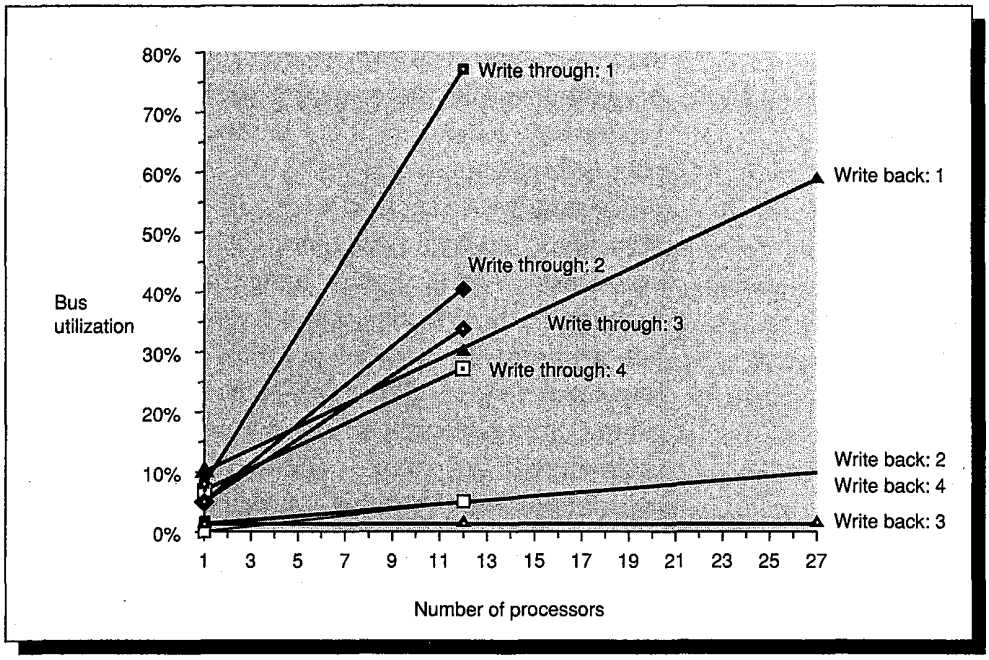

FIGURE 10.6 Comparing the impact of write-through versus write-back cache coherency on bus utilization of the Sequent Symmetry multiprocessor for four parallel benchmarks: (1) Butterfly Switch Simulator, (2) 2D Monte Carlo Simulation, (3) Ray Tracing , and (4) Parallel Linpack Benchmark. Lovett and Thakkar [1988] collected these data with a hardware performance monitor.

Another experiment evaluated the Symmetry as a timeshared (multiprogrammed) multiprocessor running ten independent programs. The experiment ran n copies of the program on n processors. This study found about half the programs started to stray from linearly increasing throughput at 6 to 8 processors with write through, yet with write back it stayed near linear for all but one of the ten programs for up to 28 processors. (The single dud was due to hot spots in the operating system rather than write-back coherency protocol.)

10.9 Fallacies and Pitfalls

Given the speculative nature of this chapter, it would seem that this section would not be needed. In good conscience, however, we submit two warnings.

Pitfall: Measuring performance of multiprocessors by linear speedup versus execution time.

"Mortar shot" graphs—plotting performance versus number of processors showing linear speedup, a plateau, and then a falling off—have long been used to judge the success of parallel processors. While scalability is one facet of a parallel program, it is not a direct measure of performance. The first question is

the power of the processors being scaled: A program that linearly improves performance to equal 100 Intel 8080s may be slower than the sequential version on a workstation. Be especially careful of floating-point–intensive programs, as processing elements without hardware assist may scale wonderfully but have poor collective performance.

Comparing execution times is only fair if you are comparing the best algorithms on each machine. (Of course, you can't subtract time for idle processors when evaluating a multiprocessor, so CPU time is inappropriate for multiprocessors.) Comparing the identical code on two machines may seem fair, but it is not; the parallel program may be slower on a uniprocessor than a sequential version. Sometimes, developing a parallel program will lead to algorithmic improvements, so that comparing the previously best-known sequential program with the parallel code—which seems fair—will not compare equivalent algorithms. To reflect this issue, sometimes the terms *relative speedup* (same program) and *true speedup* (best programs) are used. Results that suggest *superlinear* performance, when a program on n processors is more than n times faster than the equivalent uniprocessor, give a clue to unfair comparisons.

Fallacy: Amdahl's Law doesn't apply to parallel computers.

In 1987, the head of a research organization claimed that Amdahl's Law (see Section 1.3) had been broken by a MIMD machine. This hardly meant, however, that the law has been overturned for parallel computers; the neglected portion of the program will still limit performance. To try to understand the basis of the media reports, let's see what Amdahl [1967] originally said:

A fairly obvious conclusion which can be drawn at this point is that the effort expended on achieving high parallel processing rates is wasted unless it is accompanied by achievements in sequential processing rates of very nearly the same magnitude. [page 483]

One interpretation of the law was that since portions of every program must be sequential, there is a limit to the useful economic number of processors—say 100. By showing linear speedup with 1000 processors, this interpretation of Amdahl's Law was disproved.

The approach of the researchers was to change the input to the benchmark, so that rather than going 1000 times faster, they essentially computed 1000 times more work in comparable time. For their algorithm the sequential portion of the program was constant independent of the size of the input, and the rest was fully parallel—hence, linear speedup with 1000 processors.

Chapter 2 (see Section 2.2) describes the dangers of letting each experimenter select his own input for benchmarks. We see no reason why varying input is safe for evaluating performance of multiprocessors, nor why Amdahl's Law doesn't apply. What this research does point out is the importance of having benchmarks that are large enough to demonstrate performance of large-scale parallel processors.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 618

10.10 Concluding Remarks—Evolution Versus Revolution in Computer Architecture

Reading conference and journal articles from the last 20 years can leave one discouraged; so much effort has been expended with so little impact. Optimistically speaking, these papers act as gravel and, when placed logically together, form the foundation for the next generation of computers. From a more pessimistic point of view, if 90% of the ideas disappeared no one would notice.

One reason for this could be called the "von Neumann syndrome." By hoping to invent a new model of computation that will revolutionize computing, researchers are striving to become known as the von Neumann of the 21st century. Another reason is taste: researchers often select problems that no one else cares about. Even if important problems are selected, there is frequently a lack of experimental evidence to convincingly demonstrate the value of the solution. Moreover, when important problems are selected and the solutions are demonstrated, the proposed solutions may be too expensive relative to their

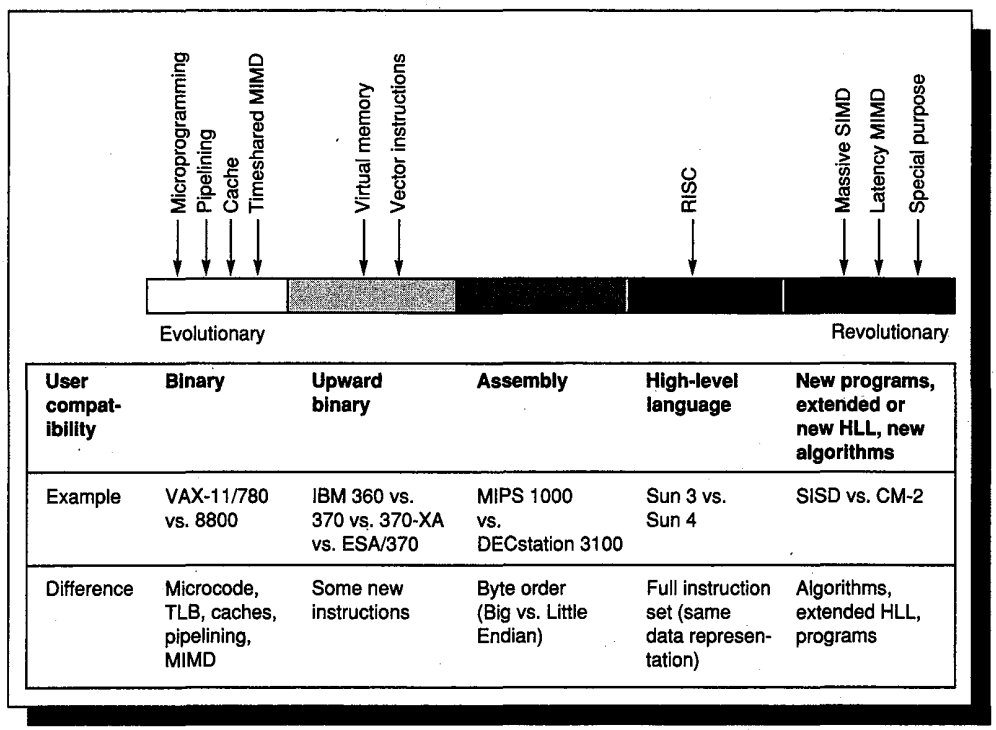

User compat-ibility	Binary	Upward binary	Assembly	High-level language	New programs, extended or new HLL, new algorithms
Example	VAX-11/780 vs. 8800	IBM 360 vs. 370 vs. 370-XA vs. ESA/370	MIPS 1000 vs. DECstation 3100	Sun 3 vs. Sun 4	SISD vs. CM-2
Difference	Microcode, TLB, caches, pipelining, MIMD	Some new instructions	Byte order (Big vs. Little Endian)	Full instruction set (same data represen-tation)	Algorithms, extended HLL, programs

FIGURE 10.7 The evolution-revolution spectrum of computer architecture. The first four columns are distinguished from the last column in that applications and operating systems can be ported from other computers rather than written from scratch. For example, RISC is listed in the middle of the spectrum because user compatibility is only at the level of high-level languages, while microprogramming allows binary compatibility, and latency-oriented MIMDs require changes to algorithms and extending HLLs. Time-shared MIMD means MIMDs justified by running many independent programs at once, while latency MIMD means MIMDs intended to run a single program faster.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 619

benefit. Sometimes this expense is measured as straightforward cost/performance—the performance enhancement does not merit the added cost. More often the expense of innovation is that it is too disruptive to computer users. Figure 10.7 shows what we mean by the *evolution-revolution spectrum* of computer architecture innovation. To the left are ideas that are invisible to the user (presumably excepting better cost, better performance, or both). This is the evolutionary end of the spectrum. At the other end are revolutionary architecture ideas. Those are the ideas that require new applications from programmers who must learn new programming languages and models of computation, and must invent new data structures and algorithms.

Revolutionary ideas are easier to publish than evolutionary ideas, but to be adopted they must have a much higher payoff. Caches are an example of an evolutionary improvement. Within five years after the first publication about caches almost every computer company was designing a machine with a cache. The RISC ideas were nearer to the middle of the spectrum, for it took closer to ten years for most companies to have a RISC product. An example of a revolutionary computer architecture is the Connection Machine. Every program that runs efficiently on that machine was either substantially modified or written especially for it, and programmers need to learn a new style of programming for it. Thinking Machines was founded in 1983, but only a few companies offer that style of machine.

There is value in projects that do not affect the computer industry because of lessons that they document for future efforts. The sin is not in having a novel architecture that is not a commercial success; the sin is in not quantitatively evaluating the strengths and weaknesses of the novel ideas. The next section mentions several machines whose primary contribution is documentation of the machine and experience using it.

When contemplating the future—and when inventing your own contributions to the field—remember the evolution-revolution spectrum. Also keep in mind the laws and principles of computer architecture found in the early chapters; these will surely guide computers of the future, just as they have guided computers of the past.

10.11 Historical Perspective and References

For over a decade prophets have voiced the contention that the organization of a single computer has reached its limits and that truly significant advances can be made only by interconnection of a multiplicity of computers in such a manner as to permit cooperative solution.... Demonstration is made of the continued validity of the single processor approach...

Amdahl [1967, 483]

The quotes at the chapter opening give the classic arguments for abandoning the current form of computing, and Amdahl [1967] gives the classic reply.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 620

Arguments for the advantages of parallel execution can be traced back to 19th century [Menabrea 1842]! Yet the effectiveness of the multiprocessor for reducing latency of individual important programs is still being determined.

The earliest ideas on SIMD-style computers are from Unger [1958] and Slotnick, Borck, and McReynolds [1962]. Slotnick's Solomon design formed the basis of the Illiac IV, perhaps the most infamous of the supercomputer projects. While successful in pushing several technologies useful in later projects, it failed as a computer. Costs escalated from the $8 million estimate in 1966 to $31 million by 1972, despite constructing only a quarter of the planned machine. Actual performance was at best 15 MFLOPS versus initial predictions of 1000 MFLOPS for the full system (see Hord [1982]). Delivered to NASA Ames Research 1972, the computer took three more years of engineering before it was usable. These events slowed investigation of SIMD, with Danny Hillis [1985] resuscitating this style in the Connection Machine: The cost of a program memory for each of 65,636 1-bit processors was prohibitive, and SIMD was the solution.

It is difficult to distinguish the first multiprocessor. The first computer from the Eckert-Mauchly Corporation, for example, had duplicate units to improve availability. Holland [1959] gave early arguments for multiple processors. After several laboratory attempts at multiprocessors, the 1980s first saw successful commercial multiprocessors. Bell [1985] suggests the key was that the smaller size of the microprocessor allowed the memory bus to replace the interconnection network hardware, and that portable operating systems meant multiprocessor projects no longer required the invention of a new operating system. This is the paper in which he defines the terms "multiprocessor" and "multicomputer." Two of the best-documented multiprocessor projects are the C.mmp [Wulf and Bell 1972 and Wulf and Habrison 1978] and Cm* [Swan et al. 1977 and Gehringer, Siewiorek, and Segall 1987]. Recent commercial multiprocessors include the Encore Multimax [Wilson 1987] and the Sequent Symmetry [Lovett and Thakkar 1988]. The Cosmic Cube is an early multicomputer [Seitz 1985]. Recent commercial multicomputers are the Intel Hypercube and the Transputer-based machines [Whitby-Strevens 1985]. Attempts at building a scalable shared-memory multiprocessor include the IBM RP3 [Pfister, Brantley, George, Harvey, Kleinfekder, McAuliffe, Melton, Norton, and Weiss 1985], the NYU Ultracomputer [Schwartz 1980 and Elder, Gottlieb, Kruskal, McAuliffe, Randolph, Snir, Teller, and Wilson 1985], and the University of Illinois Cedar project [Gajksi, Kuck, Lawrie, and Sameh 1983].

There is unbounded information on multiprocessors and multicomputers: Conferences, journal papers, and even books seem to be appearing faster than any single person can absorb the ideas. One good source is the International Conference on Parallel Processing, which has met annually since 1972. Two recent books on parallel computing have been written by Almasi and Gottlieb [1989] and Hockney and Jesshope [1988]. Eugene Miya of NASA Ames has collected an on-line bibliography of parallel-processing papers that contains more than 10,000 entries. To highlight a few papers, he sends out electronic

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 621

requests every January to ask which papers every serious student in the field should read. After collecting the ballots, he picks the ten papers most frequently recommended and publishes that list. Here is an alphabetical list of the winners: Andrews and Schneider [1983]; Batcher [1974]; Dewitt, Finkel, and Solomon [1984]; Kuhn and Padua [1981]; Lipovski and Tripathi [1977]; Russell [1978]; Seitz [1985]; Swan, Fuller, and Siewiorek [1977]; Treleaven, Brownbridge, and Hopkins [1982]; and Wulf and Bell [1972].

Special-purpose computers predate the stored-program computer. Brodersen [1989] gives a history of signal processing and its evolution to programmable devices. H. T. Kung [1982] coined the term "systolic array" and has been one of the leading proponents of this style of computer design. Recent research has been in the direction of making programmable systolic-array elements and providing a programming environment to simplify the programming task.

Its hard to predict the future, yet Gordon Bell has made two predictions for 1995. The first is that a computer capable of sustaining a TeraFLOPS—one million MFLOPS—will be constructed by 1995, either using a multicomputer with 4K to 32K nodes or a Connection Machine with several million processing elements [Bell 1989]. To put this prediction in perspective, each year the Gordon Bell Prize acknowledges advances in parallelism, including the fastest real program (highest MFLOPS). In 1988, the winner achieved 400 MFLOPS using a CRAY X-MP with four processors and 16 megawords and in 1989 the winner used an eight-processor CRAY Y-MP to run at 1680 MFLOPS. Machines and programs will have to improve by a factor of three each year for the fastest program to achieve 1 TFLOPS in 1995.

The second Bell prediction concerns the number of data streams in supercomputers shipped in 1995. Danny Hillis believes that while supercomputers with a small number of data streams may be best sellers, the biggest machines will be machines with many data streams, and these will perform the bulk of the computations. Bell bet Hillis that in the last quarter of calendar year 1995 more sustained MFLOPS will be shipped in machines using few data streams ($\leq$100) rather than many data streams ($\geq$1000). This bet concerns only supercomputers, defined as machines costing more than $1,000,000 and used for scientific applications. Sustained MFLOPS is defined for this bet as the number of floating-point operations per **month**, so availability of machines affects their rating. The loser must write and publish an article explaining why his prediction failed; your authors will act as judge and jury.

References

ALMASI, G. S. AND A. GOTTLIEB [1989]. *Highly Parallel Computing,* Benjamin/Cummings, Redwood City, Calif.

AMDAHL, G. M. [1967]. "Validity of the single processor approach to achieving large scale computing capabilities," *Proc. AFIPS Spring Joint Computer Conf.* 30, Atlantic City, N. J. (April) 483–485.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 622

ANDREWS, G. R. AND F. B. SCHNEIDER [1983]. "Concept and notations for concurrent programming," *Computing Surveys* 15:1 (March) 3–43.

BATCHER, K. E. [1974]. "STARAN parallel processor system hardware," *Proc. AFIPS National Computer Conference*, 405–410.

BELL, C. G. [1985]. "Multis: A new class of multiprocessor computers," *Science* 228 (April 26) 462–467.

BELL, C. G. [1989]. "The future of high performance computers in science and engineering," *Comm. ACM* 32:9 (September) 1091–1101.

BOUKNIGHT, W. J, S. A. DENEBERG, D. E. MCINTYRE, J. M. RANDALL, A. H. SAMEH, AND D. L. SLOTNICK [1972]. "The ILLIAC IV system," *Proc. IEEE* 60:4, 369–379. Also appears in D. P. Siewiorek, C. G. Bell, and A. Newell, *Computer Structures: Principles and Examples* (1982), 306–316.

BRODERSEN, R. W. [1989]. "Evolution of VLSI signal-processing circuits," *Proc. Decennial Caltech Conf. on VLSI* (March) 43–46, The MIT Press, Pasadena, Calif.

DEWITT, D. J., R. FINKEL, AND M. SOLOMON [1984]. "The CRYSTAL multicomputer: Design and implementation experience, Computer Sciences Tech. Rep. No. 553, University of Wisconsin-Madison, September.

ELDER, J., A. GOTTLIEB, C. K. KRUSKAL, K. P. MCAULIFFE, L. RANDOLPH, M. SNIR, P. TELLER, AND J. WILSON [1985]. "Issues related to MIMD shared-memory computers: The NYU Ultracomputer approach," *Proc. 12th Int'l Symposium on Computer Architecture* (June), Boston, Mass., 126–135.

FLYNN, M. J. [1966]. "Very high-speed computing systems," *Proc. IEEE* 54:12 (December) 1901–1909.

GAJSKI, D., D. KUCK, D. LAWRIE, AND A. SAMEH [1983]. "CEDAR—A large scale multiprocessor," *Proc. Int'l Conf. on Parallel Processing* (August) 524–529.

GEHRINGER, E. F., D. P. SIEWIOREK, AND Z. SEGALL [1987]. *Parallel Processing: The Cm* Experience*, Digital Press, Bedford, Mass.

HILLIS, W. D. [1985]. *The Connection Machine,* The MIT Press, Cambridge, Mass.

HOCKNEY, R. W. AND C. R. JESSHOPE [1988]. *Parallel Computers-2, Architectures, Programming and Algorithms,* Adam Hilger Ltd., Bristol, England and Philadelphia.

HOLLAND, J. H. [1959]. "A universal computer capable of executing an arbitrary number of subprograms simultaneously," *Proc. East Joint Computer Conf.* 16, 108–113.

HORD, R. M. [1982]. *The Illiac-IV, The First Supercomputer,* Computer Science Press, Rockville, Md.

KUHN, R. H. AND D. A. PADUA, EDS. [1981]. *Tutorial on Parallel Processing*, IEEE.

KUNG, H. T. [1982]. "Why systolic architectures?," *IEEE Computer* 15:1, 37–46.

LIPOVSKI , A. G. AND A. TRIPATHI [1977]. "A reconfigurable varistructure array processor," *Proc. 1977 Int'l Conf. of Parallel Processing* (August), 165–174.

LOVETT, T. AND S. THAKKAR [1988]. "The Symmetry multiprocessor system," *Proc. 1988 Int'l Conf. of Parallel Processing*, University Park, Pennsylvania, 303–310.

MENABREA, L. F. [1842]. "Sketch of the analytical engine invented by Charles Babbage," Bibiothèque Universelle de Genève (October).

MITCHELL, D. [1989]. "The Transputer: The time is now," *Computer Design*, RISC supplement, 40–41 (November).

PFISTER, G. F., W. C. BRANTLEY, D. A. GEORGE, S. L. HARVEY, W. J. KLEINFEKDER, K. P. MCAULIFFE, E. A. MELTON, V. A. NORTON, AND J. WEISS [1985]. "The IBM research parallel processor prototype (RP3): Introduction and architecture," *Proc. 12th Int'l Symposium on Computer Architecture* (June), Boston, Mass., 764–771.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 623

RUSSELL, R. M. [1978]. "The Cray-1 computer system," *Comm. ACM* 21:1 (January) 63–72.

SEITZ, C. [1985]. "The Cosmic Cube," *Comm. ACM* 28:1 (January) 22–31.

SLOTNICK, D. L., W. C. BORCK, AND R. C. MCREYNOLDS [1962]. "The Solomon computer," *Proc. Fall Joint Computer Conf.* (December), Philadelphia, 97–107.

SWAN, R. J., A. BECHTOLSHEIM, K. W. LAI, AND J. K. OUSTERHOUT [1977]. "The implementation of the Cm* multi-microprocessor," *Proc. AFIPS National Computing Conf.*, 645–654.

SWAN, R. J., S. H. FULLER, AND D. P. SIEWIOREK [1977]. "Cm*—A modular, multi-microprocessor," *Proc. AFIPS National Computer Conf.* 46, 637–644.

SWARTZ, J. T. [1980]. "Ultracomputers," *ACM Transactions on Programming Languages and Systems* 4:2, 484–521

TRELEAVEN, P. C., D. R. BROWNBRIDGE, and R. P. HOPKINS [1982]. "Data-driven and demand-driven computer architectures," *Computing Surveys*, 14:1 (March) 93–143.

UNGER, S. H. [1958]. "A computer oriented towards spatial problems," *Proc. Institute of Radio Engineers* 46:10 (October) 1744–1750.

VON NEUMANN, J. [1945]. "First draft of a report on the EDVAC." Reprinted in W. Aspray and A. Burks, eds., *Papers of John von Neumann on Computing and Computer Theory* (1987), 17–82, The MIT Press, Cambridge, Mass.

WHITBY-STREVENS C. [1985]. "The transputer," *Proc. 12th Int'l Symposium on Computer Architecture*, Boston, Mass. (June) 292–300.

WILSON, A. W., JR. [1987]. "Hierarchical cache/bus architecture for shared memory multiprocessors," *Proc. 14th Int'l Symposium on Computer Architecture* (June), Pittsburg, Penn., 244–252.

WULF, W. AND C. G. BELL [1972]. "C.mmp—A multi-mini-processor," *Proc. AFIPS Fall Joint Computing Conf.* 41, part 2, 765–777.

WULF, W. AND S. P. HARBISON [1978]. "Reflections in a pool of processors—An experience report on C.mmp/Hydra," *Proc. AFIPS 1978 National Computing Conf.* 48 (June), Anaheim, Calif. 939–951.

EXERCISES

10.1 [Discussion] <10.4> The weakness of SIMD for case statements, as well as the failure of the first machine to popularize SIMD, prevented exploration of SIMD designs while MIMD was still an open frontier. MIMD also has the advantage of riding the wave of improvements in SISD processors. Now that MIMD programming has not succumbed easily to assaults of computer scientists, the issue arises whether the simpler programming model of SIMD might lead it to victory over MIMD for large numbers of processors. It looks as if MIMD programs for thousands of processors will consist of thousands of copies of one program rather than thousands of different programs. Thus, the direction is toward a single **program** with multiple data streams, independent of whether the machine itself is SIMD or MIMD. What trends favor MIMD over SIMD, and vice versa? Be sure to consider utilization of memory and processors (including communication and synchronization).

10.2 [Discussion] <10.3–10.5> It might take approximately 100 clocks to communicate in a massively parallel SIMD or MIMD machine. What hardware techniques might

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 624

this time? How can you change the architecture or the programming model to make a computer more immune to such delays?

10.3 [Discussion] <10.4,10.8> What must happen before latency-oriented MIMD machines become commonplace?

10.4 [Discussion] <10.6> When do special-purpose processors make sense economically?

10.5 [Discussion] <10.8> Construct a scenario whereby a truly revolutionary architecture—pick your favorite candidate—will play a significant role. Significant is defined as 10% of the computers sold, 10% of the users, 10% of the money spent on computers, or 10% of some other figure of merit.

10.6 [30] <10.2> The CM-2 uses 64K 1-bit processors in SIMD mode. Bit-serial operations can easily be simulated 32 bits one step by a 32-bit-wide SISD, at least for logical operations. The CM-2 takes about 500 ns for such operations. If you have access to a fast SISD, calculate how long add and logical AND take on 64K 1-bit numbers.

10.7 [30] <10.2> Similar to the question above, a popular use of the CM-2 is to operate on 32-bit data using multiple steps with the 64K 1-bit processors. The CM-2 takes about 16 microseconds for a 32-bit AND or add. Simulate this activity on a fast SISD; calculate how long it takes to add and logical AND 64K 32-bit numbers.

10.8–10.12 <2.2,10.4> **If you have access to a few different multiprocessors or multicomputers, performance comparison is the basis of some projects.**

10.8 [50] <2.2,10.4> One argument for super-linear speedup (pages 585–586) is that time spent servicing interrupts or switching contexts is reduced when you have many processors, since only one need service interrupts and there are more processors to be shared by users. Measure the time spent on a workload in handling interrupts or context switching on a uniprocessor versus a multiprocessor. This workload may be a mix of independent jobs for a multiprogramming environment or a single large job. Does the argument hold?

10.9 [50] <2.2,10.4> A multiprocessor or multicomputer is typically marketed using programs that can scale performance linearly with the number of processors. The project would be to port programs written for one machine to the others and measure their absolute performance and how it changes as you change the number of processors. What changes need to be made to improve performance of the ported programs on each machine? What is the ratio of processor performance according to each program?

10.10 [50] <2.2,10.4> Instead of trying to create fair benchmarks, invent programs that make one multiprocessor or multicomputer look terrible compared to the others, and also programs that always make one look better than the others. It would be an interesting result if you couldn't find a program that made one multiprocessor or multicomputer look worse than the others. What are the key performance characteristics of each organization?

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 625

10.11 [50] <2.2,10.4> Multiprocessors and multicomputers usually show performance increases as you increase the number of processors, with the ideal being n times speedup for n processors. The goal of this biased benchmark is to make a program that gets worse performance as you add processors. For example, this means that 1 processor on the multiprocessor or multicomputer runs the program fastest, 2 is slower, 4 is slower than 2, and so on. What are the key performance characteristics for each organization that give inverse linear speedup?

10.12 [50] <10.4> Networked workstations can be considered multicomputers, albeit with slow communication relative to computation. Port multicomputer benchmarks to a network using remote procedure calls for communication. How well do the benchmarks scale on the network versus the multicomputer? What are the practical differences between networked workstations and a commercial multicomputer?

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 626

The Fast drives out the Slow even if the Fast is wrong.

W. Kahan

by David Goldberg
(Xerox Palo Alto Research Center)

A.1	**Introduction**	**A-1**
A.2	**Basic Techniques of Integer Arithmetic**	**A-2**
A.3	**Floating Point**	**A-12**
A.4	**Floating-Point Addition**	**A-16**
A.5	**Floating-Point Multiplication**	**A-20**
A.6	**Division and Remainder**	**A-23**
A.7	**Precisions and Exception Handling**	**A-28**
A.8	**Speeding Up Integer Addition**	**A-31**
A.9	**Speeding Up Integer Multiplication and Division**	**A-39**
A.10	**Putting It All Together**	**A-53**
A.11	**Fallacies and Pitfalls**	**A-57**
A.12	**Historical Perspective and References**	**A-58**
	Exercises	**A-63**

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.

A Computer Arithmetic

A.1 Introduction

A tremendous variety of algorithms have been proposed for use in floating-point accelerators. However, actual floating-point chips are usually based on refinements and variations of just a few basic algorithms. In this appendix, we focus on those algorithms. In addition to choosing algorithms for addition, subtraction, multiplication and division, the computer architect must decide whether to go beyond the basics. Should square root be implemented in hardware or software? Should extended precision be implemented? This appendix will give you the background for making these and other decisions.

Our discussion of floating point will focus almost exclusively on the IEEE floating-point standard (IEEE 754) because of its rapidly increasing acceptance. Although floating-point arithmetic involves manipulating exponents and shifting fractions, the bulk of the time in floating-point operations is spent operating on fractions using integer algorithms (but not necessarily using the integer hardware). Thus, after our discussion of floating point, we will take a more detailed look at integer algorithms.

Some good references on computer arithmetic, in order from least to most detailed, are Chapter 7 of Hamacher, Vranesic, and Zaky [1984], Gosling [1980], and Scott [1985].

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 629

A.2 Basic Techniques of Integer Arithmetic

Readers who have studied computer arithmetic before will find most of this section to be review.

Ripple-Carry Addition

The building blocks of an adder that can compute the sum of the n-bit numbers $a_{n-1}\cdots a_1a_0$ and $b_{n-1}\cdots b_1b_0$ are *half adders* and *full adders*. The half adder takes two bits a_i and b_i as input and produces a sum bit s_i and a carry bit c_{i+1} as output. Mathematically, $s_i = (a_i + b_i) \bmod 2$, and $c_{i+1} = \lfloor (a_i + b_i)/2 \rfloor$, where $\lfloor\ \rfloor$ is the floor function. As logic equations, $s_i = a_i \bar{b}_i + \bar{a}_i b_i$, and $c_{i+1} = a_i b_i$, where $a_i b_i$ means $a_i \wedge b_i$ and $a_i + b_i$ means $a_i \vee b_i$. The half adder is also called a (2,2) adder, since it takes two inputs and produces two outputs. The full adder is a (3,2) adder and is defined by the logic equations

A.2.1
$$s_i = a_i\bar{b}_i\bar{c}_i + \bar{a}_i b_i\bar{c}_i + \bar{a}_i\bar{b}_i c_i + a_i b_i c_i$$

A.2.2
$$c_{i+1} = a_i b_i + a_i c_i + b_i c_i$$

The input c_i is called the *carry in*, while c_{i+1} is the *carry out*. The principle problem in building an adder for n-bit numbers is propagating the carries. The most obvious way to solve this is with a *ripple-carry adder*, consisting of n full adders, as illustrated in Figure A.1. (In the figures in this appendix the least significant bit is always on the right.) The c_{i+1} output of the ith adder is fed into the c_{i+1} input of the next adder (the $(i + 1)$-th adder) with the lower order carry in c_0 set to 0. Since the low-order carry in is zero, the low-order adder could be a half adder. Later, however, we will see that setting the low-order carry-in bit to 1 is useful for performing subtraction.

From Equation A.2.2, there are two levels of logic involved in computing c_{i+1} from c_i. Thus, if the least significant bit generates a carry, and that carry gets propagated all the way to the last adder, the a_0 signal will pass through $2n$ levels of logic before the final gate can determine whether there is a carry out of the most significant place. In general, the time a circuit takes to produce an output is proportional to the maximum number of logic levels through which a signal travels. However, determining the exact relationship between logic levels and timings is highly technology dependent. Therefore, when comparing adders we will simply compare the number of logic levels in each one. For a ripple-carry adder that operates on n bits, there are $2n$ logic levels. Typical values of n are 32 for integer arithmetic and 53 for double-precision floating point. The ripple-carry adder is the slowest adder, but also the cheapest. It can be built with only n simple cells, connected in a simple, regular way.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 630

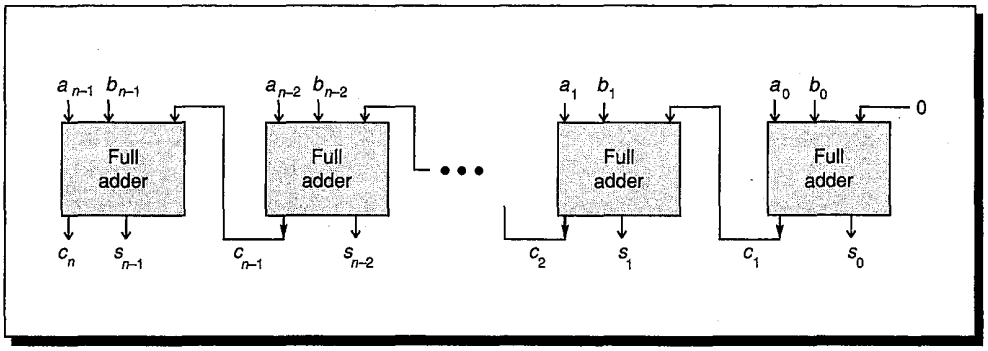

FIGURE A.1 Ripple-carry adder, consisting of *n* full adders. The carry out of one full adder is connected to the carry in of the adder for the next most significant bit. The carries ripple from the least significant bit (on the right) to the most significant bit (on the left).

Because the ripple-carry adder is relatively slow compared to the designs discussed in Section A.8, one might wonder why it is used at all. In technologies like CMOS, even though ripple adders take time O(n), the constant factor is very small. In such cases short ripple adders are often used as building blocks in larger adders.

Radix-2 Multiplication and Division

The simplest multiplier operates on two unsigned numbers, one bit at a time, as illustrated in Figure A.2(a) (page A-4). The numbers to be multiplied are $a_{n-1}a_{n-2}\cdots a_0$ and $b_{n-1}b_{n-2}\cdots b_0$, and they are placed in registers A and B, respectively. Register P is initially zero. There are two parts in each multiply step.

1. If the least significant bit of A is 1, then register B, containing $b_{n-1}b_{n-2}\cdots b_0$, is added to P; otherwise 00···00 is added to P. The sum is placed back into P.
2. Registers P and A are shifted right, with the low-order bit of P being moved into register A and the rightmost bit of A, which is not used in the rest of the algorithm, being shifted out.

After n steps, the product appears in registers P and A, with A holding the lower-order bits.

The simplest divider also operates on unsigned numbers and produces a bit at a time. A hardware divider is shown in Figure A.2(b). To compute a/b, put a in the A register, b in the B register, 0 in the P register, and then proceed as follows:

1. Shift the register pair (P,A) one bit left.
2. Subtract the content of register B (which is $b_{n-1}b_{n-2}\cdots b_0$) from register P.
3. If the result of step 2 is negative, set the low-order bit of A to 0, otherwise to 1.

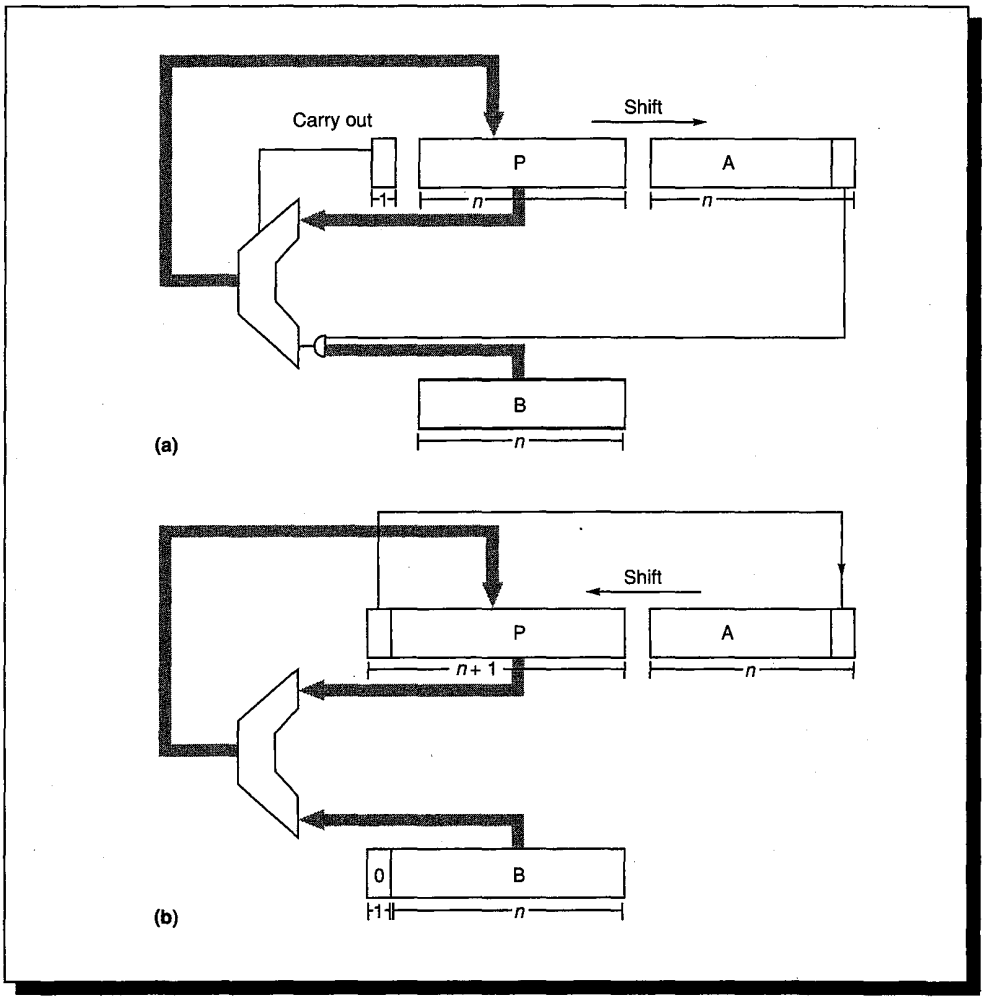

FIGURE A.2 Block diagram of simple multiplier (a) and divider (b) for *n*-bit unsigned integers. Each multiplication step consists of adding the contents of P to either B or 0 (depending on the low-order bit of A), replacing P with the sum, and then shifting both P and A one bit right. Each division step involves first shifting P and A one bit left, subtracting B from P, and if the difference is nonnegative, putting it into P. If the difference is nonnegative, the low-order bit of A is set to 1.

4. If the result of step 2 is negative, restore the old value of P by adding the contents of register B back into P.

After repeating this n times, the A register will contain the quotient, and the P register will contain the remainder. This algorithm is the binary version of the paper-and-pencil method; a numerical example is illustrated in Figure A.3(a) (page A-6).

Notice that the two block diagrams in Figure A.2 are very similar. The main difference is that the register pair (P,A) shifts right when multiplying and left when dividing. By allowing these registers to shift bidirectionally, the same hardware can be shared between multiplication and division.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 632

The division algorithm illustrated in Figure A.3(a) (page A-6) is called *restoring*, because if subtraction by b yields a negative result, the P register is restored by adding b back in. The restoration step (4 above) can be easily eliminated. To see why, let r be the contents of the (P,A) register pair, with a binary point between the low-order bit of P and the high-order bit of A. Then each step of the algorithm computes $2r - b$, putting the high-order word of this difference in P, and the low-order word in A. Suppose the result of a step is negative. Normally, we would add b back in (giving $2r$), shift (giving $4r$), and then subtract (obtaining $4r - b$). Suppose we didn't restore, but continued with the algorithm. First, shift the unrestored $2r - b$, yielding $4r - 2b$, then add b, giving $4r - b$. This is exactly what we would have obtained if we had restored! Thus, the *nonrestoring* algorithm is

If P is negative,

1a. Shift the register pair (P,A) one bit left.

2a. Add the contents of register B to P.

Else,

1b. Shift the register pair (P,A) one bit left.

2b. Subtract the contents of register B from P.

Finally,

3. If P is negative, set the low-order bit of A to 0, otherwise set it to 1.

After repeating this n times, the quotient is in A. If P is nonnegative, it is the remainder. Otherwise, it needs to be restored (i.e., add b), and then it will be the remainder. A numerical example is given in Figure A.3(b). Note that the sign of P must be tested before shifting, since the sign bit can be lost when shifting. However, because of two's complement arithmetic (discussed in the next section), the net result of shifting followed by the appropriate add/subtract operation will be the correct value. This comes about because the result of each step is a number r with $|r| \leq b$.

If a and b are unsigned numbers in the range $0 \leq a,b \leq 2^n - 1$, then the multiplier in Figure A.2 will work if register P is n bits long. However, for division, P must be extended to $n + 1$ bits in order to detect the sign of P. Thus the adder must also have $n + 1$ bits.

Why would anyone implement restoring division, which uses the same hardware as nonrestoring division (the control is slightly different) but involves an extra addition? In fact, the usual implementation for restoring division doesn't literally perform an add in step 4. Rather, the sign resulting from the subtraction is tested, and only if the sum is nonnegative is it loaded back into the P register.

As a final point, before beginning to divide, the hardware must check to see if the divisor is zero.

(a)

P	A	
00000	1110	Divide 14 = 1110 by 3 = 11. B always contains 0011
00001	110	step (1): shift
−00011		step (2): subtract
−00010	110**0**	step (3): result is negative, set quotient bit to 0
00001	110**0**	step (4): restore
00011	10**0**	step (1): shift
−00011		step (2): subtract
00000	10**01**	step (3): result is nonnegative, set quotient bit to 1
00001	**001**	step (1): shift
−00011		step (2): subtract
−00010	**0010**	step (3): result is negative, set quotient bit 0
00001	**0010**	step (4): restore
00010	**010**	step (1): shift
−00011		step (2): subtract
−00001	**0100**	step (3): result is negative, set quotient bit to 0
00010	**0100**	step (4): restore. The quotient is 0100 and the remainder is 00010.

(b)

00000	1110	Divide 14 = 1110 by 3 = 11. B always contains 0011
00001	110	step (1b): shift
+11101		step (2b): subtract *b* (add 2's complement)
11110	110**0**	step (3): P is negative, so set quotient bit to 0
11101	10**0**	step (1a): shift
+00011		step (2a): add *b*
00000	10**01**	step (3): P is nonnegative, so set quotient bit to 1
00001	**001**	step (1b): shift
+11101		step (2b): subtract *b*
11110	**0010**	step (3): P is negative, so set quotient bit to 0
11100	**010**	step (1a): shift
+00011		step (2a): add *b*
11111	**0100**	step (3): P is negative, so set quotient bit to 0
+00011		remainder is negative, so do final restore step
00010		The quotient is 0100 and the remainder is 00010

FIGURE A.3 Numerical example of (a) restoring division and (b) nonrestoring division.

Signed Numbers

There are four methods commonly used to represent signed n-bit numbers: *sign magnitude, two's complement, one's complement,* and *biased.* In the sign-magnitude system, the high-order bit is the sign bit, and the low-order $n-1$ bits are the magnitude of the number. In the two's complement system, a number and its negative add up to 2^n. In one's complement, the negative of a number is obtained by complementing each bit. In a biased system, a fixed bias is picked so that the sum of the bias and the number being represented will always be nonnegative. A number is represented by first adding it to the bias, and then encoding the sum as an ordinary unsigned number.

Example: How is –3 expressed in each of these formats?

Answer: The binary representation of 3 is 0011_2. In signed magnitude, –0011 = 1011. In two's complement $0011_2 + 1101_2 = 8$, so –0011 = 1101. In one's complement, –0011 = 1100. Using a bias of 8, 3 is represented by 1011, and –3 by 0101.

The most widely used system for representing integers, two's complement, is the system we will use here; one's complement is discussed in the Exercises. One reason for the popularity of two's complement is that addition is extremely simple: Simply discard the carry out from the high-order bit. To add 5 + –2, for example, add 0101 and 1110 to obtain 0011, resulting in the correct value of 3. A useful formula for the value of a two's complement number $a_{n-1}a_{n-2}\cdots a_1a_0$ is

A.2.3
$$-a_{n-1}2^{n-1} + a_{n-2}2^{n-2} + \cdots + a_1 2^1 + a_0$$

Overflow occurs when the result of the operation does not fit in the representation being used. For example, if unsigned numbers are being represented using four bits, then $6 = 0110_2$, and $11 = 1011_2$. Their sum (17) overflows because its binary equivalent (10001_2) doesn't fit into four bits. For unsigned numbers, detecting overflow is easy; it occurs exactly when there is a carry out of the most significant bit. For two's complement, things are trickier: Overflow occurs exactly when the carry into the high-order bit is different from the (to be discarded) carry out of the high-order bit. In the example of 5 + –2 above, a 1 is carried both into and out of the leftmost bit, avoiding overflow.

Negating a two's complement number involves complementing each bit and then adding 1. For instance, to negate 0011, complement it to get 1100 and then add 1 to get 1101. Thus, to implement $a-b$ using an adder, simply feed a and $\overline{b}$ (where $\overline{b}$ is the number obtained by complementing each bit of b) into the adder, and set the low-order, carry-in bit to 1. This explains why the rightmost adder in Figure A.1 is a full adder.

Multiplying two's complement numbers is not quite as simple as adding them. The obvious approach is to convert both operands to be nonnegative, do

an unsigned multiplication, and then (if the original operands were of opposite signs) negate the result. Although this is conceptually simple, it requires extra time and hardware. Here is a better approach: Suppose that we are multiplying a times b using the hardware shown in Figure A.2(a) (page A-4). Register A is loaded with the number a; B is loaded with b. Since the contents of register B is always b, we will use B and b interchangeably. The first thing to do when multiplying two's complement numbers is to ensure that when P is shifted, it is shifted arithmetically; that is, the bit shifted into the high-order bit of P should be the sign bit of P. Note that our n-bit–wide adder will now be adding n-bit two's complement numbers between -2^{n-1} and $2^{n-1} - 1$.

Next, suppose a is negative. The method for handling this case is called *Booth recoding*. Booth recoding is a very basic technique in computer arithmetic and will play a key role in Section A.9. Observe that multiplying by 0111_2 is the same as multiplying by $1000_2 - 1$. To perform this multiplication, subtract b from register P in the first multiplication cycle. Add zero in the second and third cycles. In the fourth cycle, add b. To apply this technique to a negative multiplier like $-4 = 1100_2$, think of it as an unsigned number and write it as $10000_2 - 0100_2$. If the multiplication algorithm only involves n steps ($n = 4$ in this case), the 10000_2 term is ignored, and we end up subtracting $0100_2 = 4$ times the multiplier—exactly the right answer. The advantage of Booth recoding is that it works equally well for positive and negative multipliers. To deal with negative values of a, then, all that is required is to sometimes subtract b from P, instead of either adding b or 0 to P. Here are the precise rules: If the initial content of A is $a_{n-1} \cdots a_0$, then at the ith multiply step, the low-order bit of register A is a_i, and

1. If $a_i = 0$ and $a_{i-1} = 0$ then add 0.
2. If $a_i = 0$ and $a_{i-1} = 1$ then add B.
3. If $a_i = 1$ and $a_{i-1} = 0$ then subtract B.
4. If $a_i = 1$ and $a_{i-1} = 1$ then add 0.

For the first step, when $i = 0$, take a_{i-1} to be 0.

Example: When multiplying −6 times −5, what is the sequence of values in the (P,A) register pair?

Answer: Initially, P is zero and A holds $-6 = 1010_2$. From Figure A.4, in the first step 0 is added to P giving (P,A) = 0000 1010. After shifting (P,A) = 0000 0101. In the next step, Figure A.4 shows that 0101 is added to P giving (P,A) = 0101 0101. Continuing, (P,A) = 0010 1010, 1101 1010, 1110 1101, 0011 1101, and finally 0001 1110.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 636

The four cases above can be restated as saying that in the ith step you should add $(a_{i-1} - a_i)$B to P. With this observation, it is easy to verify that these rules work, because the result of all the additions is

$$\sum_{i=0}^{n-1} b(a_{i-1} - a_i)2^i = b(-a_{n-1}2^{n-1} + a_{n-2}2^{n-2} + \cdots + a_1 2 + a_0)$$

From Equation A.2.3 (page A-7), the quantity in parenthesis is the value of A as a two's complement number.

The simplest way to implement the rules for Booth recoding is to extend the A register one bit to the right so that this new bit will contain a_{i-1}. Unlike the naive method of inverting any negative operands, this technique doesn't require extra steps or any special casing for negative operands. It has only a slightly more complicated control logic. If the multiplier is being shared with a divider, there will already be the capability for subtracting b, rather than adding it. To summarize, a simple method for handling two's complement multiplication is to pay attention to the sign of P when shifting it right, and to save the most recently shifted off bit of A to use in deciding whether to add or subtract b from P.

The reason for the term "recoding" is as follows. Consider representing numbers using 1, 0, and $\overline{1}$ where $\overline{1}$ represents -1; as an example, this allows us to also represent (recode) 0111 as $100\overline{1}$. Imagine a multiplication algorithm that worked as follows: Put a recoded number into the A register. If the low-order bit of A is 1, then add B. If it is $\overline{1}$, then subtract B. If the low-order bit is 0, then add 0. This imaginary algorithm has exactly the same effect as the Booth recoding method given above.

Booth recoding is usually the best method for designing hardware that operates on signed numbers. For hardware that doesn't directly implement it, however, performing Booth recoding in software or microcode is usually too slow, due to the conditional tests and branches. If the hardware supports arithmetic shifts (so that negative b is handled correctly), then the following

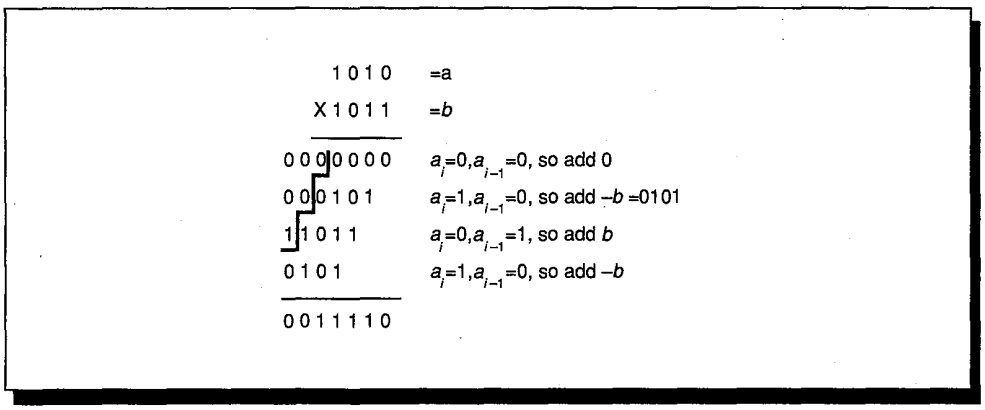

FIGURE A.4. Multiplication of $a = -6$ by $b = -5$ to get 30 using Booth recoding. The digits to the left of the jagged line are the sign-extended digits.

method can be used. Treat the multiplier a as if it were an unsigned number, and perform $n-1$ multiply steps. If $a < 0$ (in which case there will be a 1 in the low-order bit of the A register at this point), then subtract b from P; otherwise ($a \geq 0$) neither add nor subtract. In either case, do a final shift (for a total of n shifts) to get the low-order bit of the product into the low-order position of A. This works because it amounts to multiplying b by $-a_{n-1}2^{n-1} + \cdots + a_1 2 + a_0$, which is the value of $a_{n-1}\cdots a_0$ as a two's complement number by Equation A.2.3. If the hardware doesn't support arithmetic shift, then converting the operands to be nonnegative is probably the best approach.

Two final remarks: A good way to test a signed-multiply routine is to try $-2^{n-1} \times -2^{n-1}$, since this is the only case that produces a $2n-1$ bit result. Unlike multiplication, division is usually performed in hardware by converting the operands to be nonnegative and then doing an unsigned divide; because division is substantially slower (and less frequent) than multiplication, the extra time used to manipulate the signs has less impact than it does on multiplication.

Systems Issues

When designing an instruction set, there are a number of issues related to integer arithmetic that need to be resolved. Several of them are discussed here.

First, what should be done about integer overflow? This situation is complicated by the fact that detecting overflow is different depending on whether the operands are signed or unsigned integers. Consider signed arithmetic first. There are three approaches: Set a bit on overflow, trap on overflow, or do nothing on overflow. In the last case, software has to check whether or not an overflow occurred. The most convenient solution for the programmer is to have an enable bit. If this bit is turned on, then overflow causes a trap. If it is turned off, then overflow sets a bit. The advantage of this approach is that both trapping and nontrapping operations require only one instruction. Furthermore, as we will see in Section A.7, this is analogous to how the IEEE floating-point standard handles floating-point overflow. Figure A.5 shows how some common machines treat overflow.

What about unsigned addition? Notice that none of the architectures in Figure A.5 trap on unsigned overflow. The reason for this is that the primary use of unsigned arithmetic is in manipulating addresses. It is convenient to be able to subtract from an unsigned address by adding. For example, when $n = 4$, we can subtract 2 from the unsigned address $10 = 1010_2$ by adding $14 = 1110_2$. Even though $1010_2 + 1110_2$ sums to the answer we wanted ($1000_2 = 8$), this operation has an unsigned overflow. In other words, addresses are treated as both signed and unsigned numbers, making an overflow trap useless for address calculations.

A second issue concerns multiplication. Should the result of multiplying two n-bit numbers be a $2n$-bit result, or should multiplication just return the low-order n bits, signaling overflow if the result doesn't fit in n bits? The argument in favor of an n-bit result is that in virtually all high-level languages,

multiplication is an operation whose arguments are integer variables and whose result is an integer variable of the same type. Therefore, there is no way to generate code that utilizes a double-precision result. The argument in favor of a $2n$-bit result is that it can be used by an assembly language routine to speed up multiplication of multiple-precision integers substantially (by about a factor of 3).

A third issue concerns machines that want to execute one instruction every cycle. It is rarely practical to perform a multiplication or division in the same amount of time that an addition or register–register move takes. There are three possible approaches to this problem. The first is to have a single-cycle *multiply-step* instruction. This might do one step of the Booth algorithm. The second approach is to do integer multiplication in the floating-point unit and have it be part of the floating-point instruction set. (This is what DLX does.) The third approach is to have an autonomous unit in the CPU do the multiplication. In this case, the result can either be guaranteed to be delivered in a fixed number of cycles—and the compiler charged with waiting the proper amount of time—or there can be an interlock. The same comments apply to division as well. As examples, the SPARC has a multiply-step instruction but no divide-step instruction, and the MIPS R3000 has an autonomous unit that does multiplication and division (see Section E-6 for new extensions to SPARC for arithmetic). The designers of the HP Precision Architecture did an especially thorough job of analyzing the frequency of the operands for multiplication and division, and based their multiply and divide steps accordingly. (See Magenheimer et al. [1988] for details.)

A potential pitfall worth mentioning concerns multiple-precision addition. Many instruction sets offer a variant of the `ADD` instruction that adds three operands: two n-bit numbers together with a third single-bit number. This third number is the carry from the previous addition. Since the multiple-precision number will typically be stored in an array, it is important to be able to increment the array pointer without destroying the carry bit.

Machine	Trap on signed overflow?	Trap on unsigned overflow?	Set bit on signed overflow?	Set bit on unsigned overflow?
VAX	If enable is on	No	Yes. `ADD` sets V bit.	Yes. `ADD` sets C bit.
IBM 370	If enable is on	No	Yes. `ADD` sets cond code.	Yes. Logical `ADD` sets cond code.
Intel 8086	No	No	Yes. `ADD` sets V bit.	Yes. `ADD` sets C bit.
MIPS R3000	There are 2 `ADD` instructions: one always traps, the other never does.	No	No. Software must deduce it from sign of operands and result.	
SPARC	No	No	`ADDCC` sets V bit. `ADD` does not.	`ADDCC` sets C bit. `ADD` does not.

FIGURE A.5 Summary of how various machines handle integer overflow. Both the 8086 and SPARC have an instruction that traps if the V bit is set, so the cost of trapping on overflow is one extra instruction.

A.3 Floating Point

Introduction

Many applications require numbers that aren't integers. There are a number of ways that nonintegers can be represented. One is to use *fixed point*; that is, use integer arithmetic and simply imagine the binary point somewhere other than just to the right of the least significant digit. Adding two such numbers can be done with an integer add, whereas multiplication requires some extra shifting. Other representations that have been proposed involve storing the logarithm of a number and doing multiplication by adding the logarithms, or using a pair of integers (a,b) to represent the fraction a/b. However, there is only one noninteger representation that has gained widespread use, and that is the floating-point representation. In this system, a computer word is divided into two parts, an exponent and a significand. As an example, an exponent of -2 and significand of 1.5 might represent the number $1.5 \times 2^{-2} = 0.375$. The advantages of standardizing a particular representation are obvious. Numerical analysts can build up high-quality software libraries, computer designers can develop techniques for implementing high-performance hardware, and hardware vendors can build standard accelerators. Given the predominance of the floating-point representation, it appears unlikely that any other representation will come into widespread use.

A key fact about floating-point instructions is that their semantics are not as clear cut as the semantics of the rest of the instruction set, and in the past the behavior of floating-point operations varied considerably from one computer family to the next. The variations involved such things as the number of bits allocated to the exponent and significand, the range of exponents, how rounding was carried out, and the actions taken on exceptional conditions like underflow and overflow. Computer architecture books used to dispense advice on how to deal with all these details, but fortunately this is no longer necessary. That's because the computer industry is rapidly converging on the format specified by IEEE standard 754-1985. The advantages of using a standard variant of floating point are similar to those for using floating point over other noninteger representations. In this chapter we will discuss only the IEEE version of floating point. For further reading see IEEE [1985], Cody et al. [1984], Cody [1988], and Goldberg [1989].

Overview of the IEEE Standard

Probably the most notable feature of the standard is that it requires computation to continue in the face of exceptional conditions, such as dividing by zero or taking the square root of a negative number. The result of taking the square root of a negative number is a *NaN* (*N*ot *a N*umber), a bit pattern that does not represent an ordinary number. As an example of how NaNs might be useful, consider the

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 640

code for a zero finder that takes a function F as an argument and evaluates F at various points to determine a zero for it. If the zero finder accidentally probes outside the valid values for F, F may well cause an exception. Writing a zero finder that deals with this case is highly language and operating-system dependent, because it relies on how the operating system reacts to exceptions and how this reaction is mapped back into the programming language. In IEEE arithmetic it is easy to write a zero finder that handles this situation and runs on many different systems. After each evaluation of F, it simply checks to see if F has returned a NaN; if so, it knows it has probed outside the domain of F.

Because of the rules for performing arithmetic with NaNs, writing floating-point subroutines that can accept NaN as an argument rarely requires any special case checks. Suppose that arccos is computed in terms of arctan, using the formula $\arccos x = 2 \arctan(\sqrt{(1-x)/(1+x)})$. If arctan handles an argument of NaN properly, arccos will automatically do so too. That's because the IEEE standard specifies that when an argument of an operation is a NaN, the result should be a NaN. Therefore if x is a NaN, $1+x$, $1-x$, $(1+x)/(1-x)$ and $\sqrt{(1-x)/(1+x)}$ will also be NaNs. No checking for NaNs is required.

While the result of $\sqrt{-1}$ is a NaN, the result of 1/0 is not a NaN, but $+\infty$, which is another special value. The standard defines arithmetic on infinities (including $-\infty$) using rules such as $1/\infty = 0$. The formula $\arccos x = 2 \arctan(\sqrt{(1-x)/(1+x)})$ illustrates how infinity arithmetic can be used. Since $\arctan x$ asymptotically approaches $\pi/2$ as x approaches ∞, it is natural to define $\arctan(\infty) = \pi/2$, in which case $\arccos(-1)$ will automatically be computed correctly as $2 \arctan(\infty) = \pi$.

Another feature of the IEEE standard with implications for hardware is the rounding rule. When operating on two floating-point numbers, the result is usually a number that cannot be exactly represented as another floating-point number. For example, in a floating-point system using base 10 and two significant digits, $2.1 \times 0.5 = 1.05$. This needs to be rounded to two digits. Should it be rounded to 1.0 or 1.1? In the IEEE standard, such halfway cases are rounded to the number whose low-order digit is even. That is, 1.05 rounds to 1.0, not 1.1. The standard actually has four *rounding modes*. The default is *round to nearest*, which rounds to an even number in the case of ties. The other modes are round toward 0, round toward $+\infty$ and round toward $-\infty$.

The standard specifies four precisions: *single*, *single extended*, *double*, and *double extended*. The properties of these precisions are summarized in Figure A.6 (page A-14). Implementations are not required to have all four precisions, but are encouraged to support either the combination of single and single extended or all of single, double, and double extended. Let us consider single precision in more detail. Single-precision numbers are represented using 32 bits: 1 for the sign, 8 for the exponent, and 23 for the fraction. The exponent is a signed number represented using the bias method (as explained in Section A.2 above) with a bias of 127. We will always use the term *exponent field* to mean the unsigned number contained in bits one through nine and *exponent* to mean

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 641

the power to which two is to be raised. (In the standard these are called the "biased exponent" and the "unbiased exponent," respectively.) The fraction represents a number less than one, but the *significand* of the floating-point number is one plus the fraction part. In other words, if e is the value of the exponent field and f is the value of the fraction field, the number being represented is $1.f \times 2^{e-127}$.

Example: What single-precision number does the following 32-bit word represent?

```
1 10000001 01000000000000000000000
```

Answer: Considered as an unsigned number, the exponent field is 129, making the value of the exponent 129 − 127 = 2. The fraction part is $.01_2 = .25$, making the significand 1.25. Thus, this bit pattern represents the number $-1.25 \times 2^2 = -5$.

The fractional part of a floating-point number (.25 in the example above) must not be confused with the significand, which is one plus the fractional part. The leading 1 in the significand $1.f$ does not appear in the representation; that is, the leading bit is implicit. When performing arithmetic on IEEE format numbers, the fraction part normally needs to be *unpacked,* which is to say the implicit one needs to be made explicit.

In Figure A.6, the range of exponents for single precision is −126 to 127; accordingly, the exponent field ranges from 1 to 254. The exponent fields of 0 and 255 are used to represent special values. When the exponent field is 255, a zero fraction field represents infinity, and a nonzero fraction field represents a NaN. Thus, there is an entire family of NaNs. When the exponent and fraction fields are zero, then the number represented is zero. Because ordinary numbers always have a significand greater than or equal to 1—and are thus never zero—a special convention such as this is required to represent zero.

A zero exponent field and nonzero fraction part represent a *denormal* number, also sometimes called a *subnormal* number. These numbers make up the most controversial part of the standard. Later, in the discussion of multiplication, we will see why they are difficult to implement in hardware. In many floating-point systems if $E_{\min}$ is the smallest exponent, a number less than $1.0 \times 2^{E_{\min}}$

	Single	Single extended	Double	Double extended
p (bits of precision)	24	≥ 32	53	≥ 64
$E_{\max}$	127	≥ 1023	1023	≥ 16383
$E_{\min}$	−126	≤ −1022	−1022	≤ −16382
Exponent bias	127		1023	

FIGURE A.6 Format parameters for the IEEE 754 floating-point standard. The first row gives the number of bits in the significand. The blank boxes are unspecified parameters.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 642

cannot be represented, and a floating-point operation that results in a number less than this is simply flushed to zero. In the IEEE standard, on the other hand, numbers less than $1.0 \times 2^{E\min}$ are represented by shifting their fraction part to the right. This is called *gradual underflow*. Thus, as numbers decrease in magnitude below $2^{E\min}$, they gradually lose their significance and are only represented by zero when all their significance has been shifted out. For example, in base 10 with 4 significant figures, let $x = 1.234 \times 2^{E\min}$. Then $x/10 = 0.123 \times 10^{E\min}$, having lost a digit of precision; $x/100$ and $x/1000$ have even less precision, while $x/10000$ is finally small enough to be rounded to zero. Denormalized numbers are implemented by having a word with a zero exponent field represent the number $0.f \times 2^{E\min}$. One of the advantages of gradual underflow is that when it is used, if $x \neq y$, then $x - y \neq 0$. In a flush-to-zero system, this is not always true.

The primary reason why the IEEE standard, like most other floating-point formats, uses biased exponents is that it means nonnegative numbers are ordered in the same way as integers. That is, the magnitude of floating-point numbers can be compared using an integer comparator. Another (related) advantage is that zero is represented by a word of all zeros. The down side of biased exponents is that adding them is slightly awkward, because it requires that the bias be subtracted from their sum.

As the IEEE standard becomes more widespread, it will become easier to port software and to write portable libraries that deal with floating-point exceptions. But the standard also has some drawbacks:

1. It was originally intended for microprocessors, so the requirements of high-performance implementations were not given high priority.

2. The standard contains optional parts. This results in difficult decisions for implementors—which parts should they implement?—and for portable software writers—should they avoid using any of the optional parts of the standard?

3. Gradual underflow has usually been implemented in a way that is orders of magnitude slower than flush to zero, so users often disable it.

4. There is as yet no industrial-strength, public-domain, IEEE floating-point test suite.

Although the standard may ultimately improve the quality of floating-point libraries, this has yet to happen because of the large base of VAXes, IBM/370s, and Crays, as well as the fact that there is no corresponding standard for how to access its features in software. On the other hand, both DEC and IBM have recently introduced machines that use IEEE arithmetic.

Some final comments on the standard:

1. Unlike most standards, IEEE 754 did not ratify or refine any existing system. Although most of the features of the standard appeared in at least one previous computer system, it is substantially different from what was current practice at the time.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 643

2. The standard says nothing about integer arithmetic or about transcendental functions (sin, cos, exp, and so forth). In particular, it says nothing about the accuracy that transcendentals should have, and it says nothing about the exceptional values of transcendentals, such as 0^0.

3. It is intended that a computer **system**—that is, some combination of hardware and software—will implement the standard. Thus, there is nothing wrong with designing hardware that does not completely implement the standard, as long as there is some way for software to provide what the hardware does not. In fact, the best design may well involve having rare cases handled by software.

A.4 Floating-Point Addition

There are two differences between floating-point arithmetic and integer arithmetic: An exponent field must be manipulated, in addition to the fraction field, and the result of a floating-point operation usually has to be rounded in order to be represented by another floating-point number of the same precision.

Rounding

The IEEE standard specifies that the result of an arithmetic operation should be the same as it would be if computed exactly and then rounded using the current rounding mode. The most difficult mode to implement is the default mode—round toward the nearest value (and round halfway cases to even). The naive approach to complying with the IEEE standard is to compute the sum exactly and then round. This would be quite expensive, since it would require a very long adder. To see how to satisfy the standard with less hardware, we will consider some examples.

There are two ways that rounding can occur during addition. For purposes of illustration we will use base 10, which is more natural for humans, and three significant digits. The first case requires rounding due to carry out on the left, as illustrated in Figure A.7(a). The second case requires rounding due to unequal exponents, as in Figure A.7(b). Figure A.7(c) shows that it is possible for both situations to occur simultaneously. In each of these cases, the sum must be computed to more than three places in order to perform rounding. In one case—when subtracting nearby numbers, as in Figure A.7(d)—the sum must be computed to more than three places, even though no rounding occurs. By temporarily ignoring the round-to-even requirement, each of these examples can be implemented with a four-digit-wide adder (that is, using one additional digit). Thus, in Figure A.7(b) the rightmost 6 of 2.56 can simply be dropped before adding. But there is one case, shown in Figure A.7(e), in which four digits are not enough. If the low-order digit of .0376 were shifted off, the answer would have been .973 instead of .972. However, it is easy to check (disregarding round to even) that

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 644

two extra digits are always enough. These extra digits are called the *guard* and *round* digits.

The round-to-even rule introduces an extra complication. Figure A.7(f) shows an example with five significant digits. It might appear at first that one needs to keep double the number of digits to perform round to even, as the rightmost 1 in 2.5001 determines whether the result will be 4.5676 or 4.5677.

Upon a little reflection one can see that it is only necessary to know whether or not there are any nonzero digits past the guard and round positions. This information can be stored in a single bit, usually called the *sticky bit*, which is implemented by examining each digit as it is shifted off. As soon as a nonzero digit appears, the sticky bit is set on and remains stuck on. To implement round to even, simply append the sticky bit to the right of the round digit just before rounding.

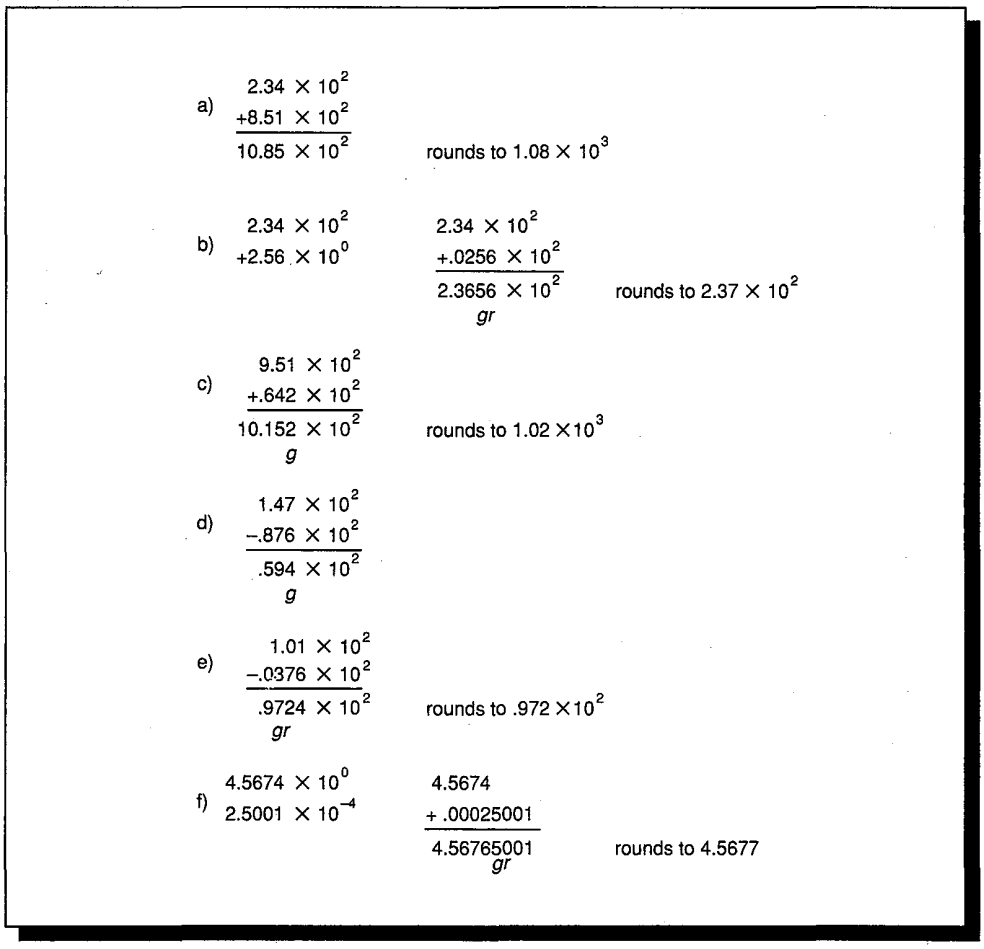

FIGURE A.7 Examples of rounding. In (a) there is rounding because of carry out on the left and in (b) because of unequal exponents, whereas in (c) both occur. Example (d) shows that one extra place must be kept even if there is no rounding, while (e) shows the situation in which two extra digits are needed. Finally (f), where $p = 5$, illustrates why a sticky bit is necessary to perform round to even. The letters g and r are placed under the guard and round digits.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 645

The Addition Algorithm

The notations e_i and s_i are used here for the exponent and significand fields of the floating-point number a_i. This means that the floating-point number has been unpacked and that s_i has an explicit leading bit. The basic procedure for adding two floating-point numbers a_1 and a_2 is straightforward and involves five steps.

1. If $e_1 < e_2$, swap the operands so that the difference of the exponents satisfies $d = e_1 - e_2 \geq 0$. Tentatively set the exponent of the result to e_1.

2. Shift s_2 by $d = e_1 - e_2$ places to the right. More precisely, put s_2 into a p-bit register and then extend that register MIN(2,d) bits to the right. Shift s_2 d places to the right. If $d > 2$, set the sticky bit to the logical OR of the $d-2$ bits that are shifted out of the extended register. Of the two extended bits, the most significant is the guard bit; the least significant is the round bit.

3. Append the sticky bit to s_2, and then add the two signed-magnitude fraction fields in a $p + 3$ bit adder. Call this preliminary sum S.

4. If there was a carry out from the most significant place in the previous step, shift the magnitude of S right by one. Otherwise, shift it left until it is normalized. Adjust the exponent of the result accordingly. The round bit is now set to the $(p + 1)$-st bit of the magnitude of S, and the sticky bit to the logical OR of all the bits to the right of the round bit.

5. Round the result using Figure A.8. If a table entry is nonempty, add 1 to the magnitude of S. Thus, if $S \geq 0$, you will be computing $S + 1$, otherwise $S - 1$.

The guard and round bits before shifting are marked in each of the examples of Figure A.7 (page A-17).

Example: Show how the addition algorithm proceeds on the operands of Figure A.7(f) when round to nearest is in effect.

Answer: In step 1, $e_1 = 0 > e_2 = -3$, so $d = 3$ and no swapping is necessary. In step 2, $g = 5$, $r = 0$, and sticky is the OR of 0, 0, and 1; hence, sticky is 1. In step 3 the numbers to be added are 4.5674 and 0.0002501, so the preliminary sum is $S = 4.5676501$. In step 4 there is no carry out, so d is still 3. The round bit is 5, and the sticky bit is $1 = 0 \vee 1$. In step 5, consulting the table tells us that because round and sticky are both nonzero, we must add 1 to the fifth digit of S, changing S from 45676 to 45676 + 1 = 45677.

Step 3 involves adding sign-magnitude numbers, and itself has three steps:

3a. Convert any negative numbers to two's complement.

3b. Perform a $(p + 4)$-bit two's complement addition ($p + 3$ bits of magnitude, 1 bit for the sign).

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 646

3c. If the result is negative, perform another two's complementation to put the result back into sign-magnitude form.

As is apparent from this, addition is quite a complicated operation. Here is one trick that can speed it up. A pair of numbers will only need to be variably shifted once, in either step 2 or step 4, but not in both. The reason is simple: If $|e_1 - e_2| > 1$, then step 4 can require a shift of at most one place. And if $|e_1 - e_2| \leq 1$, then step 2 obviously requires a shift of at most one step. A non-pipelined adder can exploit this and reduce the number of steps from five to four. An adder that uses each of the above steps as a pipeline stage can also use this reduction, though it requires duplicating the shifter and adder.

Step 3 can be time consuming, because it can involve as many as four additions: two to negate both operands (two's complementation done by performing a bitwise complementation followed by adding 1), a third for the addition itself, and then a fourth to negate the result. There are a number of ways to speed up this step. We have already seen that 1 can be added to a sum essentially for free by setting the low-order, carry-in bit of the adder to 1. If both operands are negative, we can set their sign bits to zero, remembering to negate the result. The add required when negating the result can be combined with the rounding step (which must be prepared to do an add anyway).

The rounding step requires a second full-precision add in addition to the one in step 3. It is possible to combine these into a single add. Observe that at the end of step 2, the g, r, and s bits are known; thus it is also known whether or not to round up, adding 1 to the pth most significant bit. What is not known is the position of the pth most significant bit, since its location depends on the result of the add in step 3; when adding numbers of the same sign, that position is determined by whether there is a carry out of the most significant bit. Therefore, the way to eliminate step 5 is to add in the round-up bit (if necessary) as part of step 3. Because the position is unknown, two versions of step 3 must be performed using two adders in parallel. Each adder assumes one of the two possibilities for the position where the round-up bit goes. This technique for reducing the number of addition steps is used on the Intel 860 [Kohn 1989]. When rounding, there is one complication that can arise: The addition of 1 could cause a carry out of the high-order bit. This case occurs only when the value of S is $11\cdots11$.

Rounding mode	$S \geq 0$	$S < 0$
$-\infty$		+1 if $r \vee s$
$+\infty$	+1 if $r \vee s$	
0		
Nearest	+1 if $r \wedge \overline{s} \wedge p_0$ or $r \wedge s$	+1 if $r \wedge \overline{s} \wedge p_0$ or $r \wedge s$

FIGURE A.8 Rules for implementing the IEEE rounding modes. Blank boxes mean that the p most significant bits of the preliminary sum S are the actual sum bits. If the condition in the box is true, add 1 to the pth most significant bit of S. The symbols r and s represent the round and sticky bits, while p_0 is the pth most significant bit of S.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 647

Denormalized Numbers

Very little changes in the above description if one of the inputs is a denormal number. There must be a test to see if the exponent field is 0. If it is, then when unpacking the significand there will not be a leading 1. By setting the exponent field to 1 when unpacking a denormal, the shifting rules in steps 1–5 are still correct.

In order to deal with denormalized outputs, step 4 must be modified slightly. The value in the P register is shifted left until P is normalized, or until the exponent becomes $E_{\min}$ (that is, the exponent field becomes 1). If the exponent is $E_{\min}$, and if after rounding, the high-order bit of P is 1, then the result is a normalized number and should be packed in the usual way, by omitting the 1. If, on the other hand, the high-order bit is 0, the result is denormal, and when the result is unpacked the exponent field must be set to 0.

Incidentally, detecting overflow is very easy. It can only happen if step 4 involves a shift right, and if the exponent field at that point is bumped up to 255 in single precision (or 2047 for double precision), or if this occurs after rounding.

Detecting underflow is complicated by the fact that it depends on whether there is a user trap handler. The IEEE standard specifies that if user trap handlers are enabled, the system must trap if the result is denormal. On the other hand, if trap handlers are disabled, then the underflow flag is set only if there is a loss of accuracy—that is, if the result must be rounded. The rationale for this is that if no accuracy is lost on an underflow, there is no point in setting a warning flag. But if a trap handler is enabled, the user might be trying to simulate flush-to-zero and should therefore be notified whenever a result dips below $1.0 \times 2^{E\min}$. This discussion is relevant for addition in that an addition or subtraction resulting in a denormal number will always be exact; because no accuracy can be lost to underflow, there is no need to set the underflow flag.

A.5 Floating-Point Multiplication

Floating-point multiplication is much like integer multiplication. Because floating-point numbers are stored in sign-magnitude form, the multiplier need only deal with unsigned numbers (although we have seen that Booth recoding handles signed two's complement numbers painlessly). If the fractions are unsigned p-bit numbers, then the product can have as many as $2p$ bits and must be rounded to a p-bit number. Besides multiplying the fraction parts, the exponent fields must be added, and the bias then subtracted from their sum.

Here is a straightforward method of handling rounding using the multiplier of Figure A.2 (page A-4): Multiply the two fractions to obtain a $2p$-bit product in the (P,A) registers. During the multiplication, the first $p - 2$ times a bit is shifted into the A register, OR it into the sticky bit. After the end of all the multiply

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 648

steps, the high-order bit of A is the guard bit, and the second high-order bit is the round bit. There are two cases:

1. The high-order bit of P is 0. Shift P left 1 bit, shifting in the g bit from A. Shifting the rest of A is not necessary.
2. The high-order bit of P is 1. Set $s := s \vee r$ and $r := g$, and add 1 to the exponent.

Now use the rules in Figure A.8 (page A-19) to round the result, adding the 1 (if necessary) into the low-order bit of P. The fraction (in unpacked form) is in the P register. Recall that the rounding operation can cause a carry out of the most significant bit. A good discussion of more efficient ways to implement rounding is in Santoro, Bewick, and Horowitz [1989].

Detecting overflow and underflow is slightly tricky. Consider the case of single precision. The exponent fields must be added together with −127. If the addition is done in a 10-bit adder, $-127 = 1110000001_2$, and overflow occurs when the high-order bits of the sum are 01 or if the sum is 0011111111. Underflow occurs when the high-order bits are 11 or the sum is 0000000000. Alternatively, the addition can be done using only an 8-bit adder. Simply add both exponents and $-127 = 10000001_2$. If the high-order bits of the exponent fields are different, no over/underflow is possible. If the high-order bits are both 1, the result has overflowed if it has 0 in the high-order bit or if it is 1111111. If both the exponents have high-order bits of zero, underflow has occurred if the sum has a high-order bit of 1, or if the sum is 00000000.

Denormals

From the description of the multiplication algorithm, one can see that after doing an integer multiplication on the fractions, the final result is obtained with at most one shift. With denormals, the situation changes completely. Suppose the input is normalized, but the output is denormal, so that in single precision the product has an exponent e with $e < -126$. Then the result must be shifted right by $-e - 126$ places. This requires extra hardware (a barrel shifter that wouldn't otherwise be needed) and extra time. The situation with denormal inputs isn't any better, because even if the final result is a normalized number, a variable shift is still required. Thus, high-performance, floating-point multipliers often do not handle denormalized numbers, but instead trap, letting software handle them. There are a few practical codes that generate many underflows, even when working properly, and these programs usually run quite a bit slower on systems that require denormals to be processed by a trap handler.

One procedure followed by some floating-point units is to have the multiplier deliver denormalized outputs in *wrapped* form. That is, the fraction part is normalized, and the exponent is wrapped around to a large positive number. This is exactly the result when following the multiplication algorithm for normalized numbers given above. Since the addition unit must have a barrel shifter, it is

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 649

usually straightforward to provide a way to convert wrapped numbers into their correct denormalized form by passing them through the adder. However, if a trap handler has to intervene in order to send wrapped numbers into the adder, multiplication will still be slowed down substantially.

There are some fine points that occur when a multiplication results in a denormal number. Consider the simple case of a base 2 floating-point system with 3-bit significands (hence two bits of fraction). The exact result of 1.11×2^{-2} multiplied by $1.11 \times 2^{E_{\min}}$ is $0.110001 \times 2^{E_{\min}}$. If the rounding mode is round toward plus infinity, the rounded result is the normal number $1.00 \times 2^{E_{\min}}$. Should underflow be signaled? Signaling underflow means that one is using the *before rounding* rule, because the result was denormal before rounding. Not signaling underflow means that one is using the *after rounding* rule, because the result is normalized after rounding. The IEEE standard provides for choosing either rule; however, the one chosen must be used consistently for all operations.

As mentioned in the addition section, the trap handler, if there is one, should be called whenever the result is denormal. If there is no trap handler, the underflow exception is signaled only when the result is denormal and inexact. Normally, inexact means there was a result that couldn't be represented exactly and had to be rounded. Consider again the example of $(1.11 \times 2^{-2}) \times (1.11 \times 2^{E_{\min}}) = 0.110001 \times 2^{E_{\min}}$, with round to nearest in effect. The delivered result is $0.11 \times 2^{E_{\min}}$, which had to be rounded, causing inexact to be signaled. But is it correct to also signal underflow? Gradual underflow loses significance because the exponent range is bounded. If the exponent range were unbounded, the delivered result would be $1.10 \times 2^{E_{\min}-1}$, exactly the same answer obtained with gradual underflow. The fact that denormalized numbers have fewer bits in their significand than normalized numbers therefore doesn't make any difference in this case. The commentary to the standard [Cody et al. 1984] encourages this as the criterion for setting the underflow flag. That is, it should be set whenever the delivered result is different from what would be delivered in a system with the same fraction size, but with a very large exponent range. However, owing to the difficulty of implementing this scheme, the standard allows setting the underflow flag whenever the result is denormal and different from the infinitely precise result.

Precision of Multiplication

In the discussion of integer multiplication, we mentioned that designers must decide whether to deliver the low-order word of the product or the entire product. A similar issue arises in floating-point multiplication, where the exact product can be rounded to the precision of the operands or to the next higher precision. In the case of integer multiplication, none of the standard high-level languages contains a construct that would generate a "single times single gets double" instruction. The situation is different for floating point. Not only do

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 650

many languages allow assigning the product of two single-precision variables to a double-precision one, but the construction can also be exploited by numerical algorithms. The best-known case is using iterative refinement to solve linear systems of equations.

A.6 Division and Remainder

Iterative Division

We earlier discussed an algorithm for integer division. Converting it into a floating-point division algorithm is similar to converting the integer multiplication algorithm into floating point. If the numbers to be divided are $s_1 2^{e1}$ and $s_2 2^{e2}$ then the divider will compute s_1/s_2, and the final answer will be this quotient multiplied by 2^{e1-e2}. Referring to Figure A.2(b) (page A-4), the alignment of operands is slightly different from integer division. Load s_2 into b and $s_1/2$ into P so that s_1 is shifted right one bit. Then the integer algorithm for division can be used, and the result will be of the form $q_0.q_1\cdots$. For floating-point division, the A register is not needed to hold the operands. To round, simply compute two additional quotient bits (guard and round) and use the remainder as the sticky bit. The guard digit is necessary because the first quotient bit might be zero. However, since the numerator and denominator are both normalized, it is not possible for the two most significant quotient bits to be zero.

There is a different approach to division, based on iteration. An actual machine that uses this algorithm will be discussed in Section A.10. First, we will describe the two main iterative algorithms and then discuss the pros and cons of iteration compared to the direct algorithms. There is a general technique for constructing iterative algorithms, called *Newton's iteration*, shown in Figure A.9.

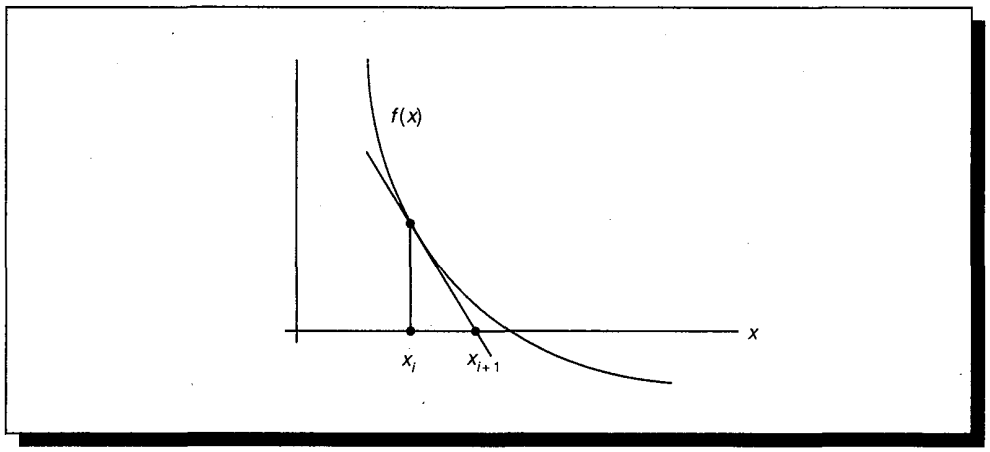

FIGURE A.9 Newton's iteration for zero finding. If x_i is an estimate for a zero of f, then x_{i+1} is a better estimate. To compute x_{i+1}, find the intersection of the x axis with the tangent line to f at x_i.

First, cast the problem in the form of finding the zero of a function. Then, starting from a guess for the zero, approximate the function by its tangent at that guess and form a new guess based on where the tangent has a zero. If x_i is a guess at a zero, then the tangent line has the equation

$$y - f(x_i) = f'(x_i)(x - x_i)$$

This equation has a zero at

A.6.1
$$x = x_{i+1} = x_i - \frac{f(x_i)}{f'(x_i)}$$

To recast division as finding the zero of a function, consider $f(x) = 1/x - b$. Since the zero of this function is at $1/b$, applying Newton's iteration to it will give an iterative method of computing $1/b$ from b. Using $f'(x) = -1/x^2$, Equation A.6.1 becomes

A.6.2
$$x_{i+1} = x_i - \frac{1/x_i - b}{-1/x_i^2} = x_i + x_i - x_i^2 b = x_i(2 - x_i b)$$

Thus, we could implement computation of a/b using the following method:

1. Scale b to lie in the range $1 \le b < 2$ and get an approximate value of $1/b$ (call it x_0) using a table lookup.
2. Iterate $x_{i+1} = x_i(2 - x_i b)$ until reaching an x_n that is accurate enough.
3. Compute ax_n and reverse the scaling done in step 1.

Here are some more details. How many times will step 2 have to be iterated? To say that x_i is accurate to p bits means that $(x_i - 1/b)/(1/b) = 2^{-p}$, and a simple algebraic manipulation shows $(x_{i+1} - 1/b)/(1/b) = 2^{-2p}$. Thus the number of correct bits doubles at each step. Newton's iteration is **self-correcting** in the sense that making an error in x_i doesn't really matter. That is, it treats x_i as a guess at $1/b$ and returns x_{i+1} as an improvement on it (roughly doubling the digits). One thing that would cause x_i to be in error is rounding error. More importantly, however, in the early iterations we can take advantage of the fact that we don't expect many correct bits by performing the multiplication in reduced precision, thus gaining speed without sacrificing accuracy. Some other applications of Newton's iteration are discussed in the Exercises.

The second iterative division method is sometimes called *Goldschmidt's algorithm*. It is based on the idea that to compute a/b, you should multiply the numerator and denominator by a number r with $rb \approx 1$. In more detail, let $x_0 = a$ and $y_0 = b$. At each step compute $x_{i+1} = r_i x_i$ and $y_{i+1} = r_i y_i$. Then the quotient $x_{i+1}/y_{i+1} = x_i/y_i = a/b$ is constant. If we pick r_i so that $y_i \to 1$, then $x_i \to a/b$, so the x_i converge to the answer we want. This same idea can be used to compute

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 652

other functions. For example, to compute the square root of a, let $x_0 = a$ and $y_0 = a$, and at each step compute $x_{i+1} = r_i^2 x_i$, $y_{i+1} = r_i y_i$. Then $x_{i+1}/y_{i+1}^2 = x_i/y_i^2 = 1/a$, so if the r_i are chosen to drive $x_i \to 1$, then $y_i \to \sqrt{a}$. This technique is used to compute square roots on the TI 8847.

Returning to Goldschmidt's division algorithm, set $x_0 = a$ and $y_0 = b$, and write $b = 1 - \delta$, where $|\delta| < 1$. If we pick $r_0 = 1 + \delta$, then $y_1 = r_0 y_0 = 1 - \delta^2$. We next pick $r_1 = 1 + \delta^2$, so that $y_2 = r_1 y_1 = 1 - \delta^4$, and so on. Since $|\delta| < 1$, $y_i \to 1$. With this choice of r_i, the x_i will be computed as $x_{i+1} = r_i x_i = (1 + \delta^{2^i}) x_i = (1 + (1 - b)^{2^i}) x_i$, or

A.6.3

$$x_{i+1} = a[1 + (1 - b)][1 + (1 - b)^2][1 + (1 - b)^4] \cdots [1 + (1 - b)^{2^i}]$$

There appear to be two problems with this algorithm. First, convergence is slow when b is not near 1 (that is, δ is not near 0); and second, the formula isn't self-correcting—since the quotient is being computed as a product of independent terms, an error in one of them won't get corrected. To deal with slow convergence, if you want to compute a/b, look up an approximate inverse to b (call it b'), and run the algorithm on ab'/bb'. This will converge rapidly since $bb' \approx 1$.

To deal with the self-correction problem, the computation should be run with a few bits of extra precision to compensate for rounding errors. However, Goldschmidt's algorithm does have a weak form of self-correction, in that the precise value of the r_i does not matter. Thus, in the first few iterations, you can choose r_i to be a truncation of $1 + \delta^{2^i}$ which may make these iterations run faster without affecting the speed of convergence. If r_i is truncated, then y_i is no longer exactly $1 - \delta^{2^i}$, so Equation A.6.3 can no longer be used, but it is easy to organize the computation so that it does not depend on the precise value of r_i. With these changes, Goldschmidt's algorithm is as follows (the notes in brackets show the connection with our earlier formulas).

1. Scale a and b so that $1 \le b < 2$.
2. Look up an approximation to $1/b$ (call it b') in a table.
3. Set $x_0 = ab'$ and $y_0 = bb'$.
4. Iterate until x_i is close enough to a/b :

$r \approx 2 - y$	[if $y_i = 1 + \delta_i$, then $r \approx 1 - \delta_i$]
$y = y \times r$	[$y_{i+1} = y_i \times r \approx 1 - \delta_i^2$]
$x = x \times r$	[$x_{i+1} = x_i \times r$]

The two iteration methods are related. Suppose in Newton's method that we unroll the iteration and compute each term x_{i+1} directly in terms of b, instead of recursively in terms of x_i. By carrying out this calculation, we discover that

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 653

$$x_{i+1} = x_0(2 - x_0 b)(1 + (x_0 b - 1)^2)(1 + (x_0 b - 1)^4)\cdots(1 + (x_0 b - 1)^{2^i})$$

This formula is of a very similar form to Equation A.6.3 when $a = 1$. In fact, if the iterations were done to infinite precision, the two methods would yield exactly the same sequence x_i.

The advantage of iteration is that it doesn't require special divide hardware, but can instead use the multiplier (which, however, requires extra control). Further, on each step, it delivers twice as many digits as in the previous step—unlike ordinary division, which produces a fixed number of digits at every step. There are two disadvantages with inverting by iteration. The first is that the IEEE standard requires division to be correctly rounded, but iteration only delivers a result that is close to the correctly rounded answer. In the case of Newton's iteration, which computes $1/b$ instead of a/b directly, there is an additional problem. Even if $1/b$ was correctly rounded, there is no guarantee that a/b will be. Take 5/7 as an example: To two digits of accuracy 1/7 is 0.14, and 5×0.14 is 0.70, but 5/7 is 0.71. The second disadvantage is that iteration does not give a remainder. This is especially troublesome if the floating-point divide hardware is being used to perform integer division, since a remainder operation is present in almost every high-level language.

Traditional folklore has held that the way to get a correctly rounded result from iteration is to compute $1/b$ to slightly more than $2p$ bits, compute a/b to slightly more than $2p$ bits, and then round to p bits. However, there is a faster way, which apparently was first implemented on the TI 8847. In this method, a/b is computed to about six extra bits of precision, giving a preliminary quotient q. By comparing qb with a (again with only six extra bits), it is possible to quickly decide whether q is correctly rounded or whether it needs to be bumped up or down by 1 in the least significant place. This algorithm is explored further in the Exercises.

One factor to take into account when deciding on division algorithms is the relative speed of division and multiplication. Since division is more complex than multiplication, it will run more slowly. As a general rule of thumb, division algorithms should try to achieve a speed that is about one-third that of multiplication. One argument in favor of this rule is that there are real programs (such as some versions of Spice) where the ratio of division to multiplication is 1:3. Another place where a factor of three arises is in the standard iterative method for computing square root. This method involves one division per iteration, but can be replaced by one using three multiplications. This is discussed in the Exercises.

Floating-Point Remainder

For nonnegative integers, integer division and remainder satisfy

$$a = (a \text{ DIV } b)b + a \text{ REM } b, \quad 0 \le a \text{ REM } b < b$$

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 654

A floating-point remainder x REM y can be similarly defined as $x = \text{INT}(x/y)y + x$ REM y. How should x/y be converted to an integer? The IEEE remainder function uses the round-to-even rule. That is, pick $n = \text{INT}(x/y)$ so that $|x/y - n| \leq 1/2$. If two different n satisfy this relation, pick the even one. Then REM is defined to be $x - yn$. Unlike integers where $0 \leq a$ REM $b < b$, for floating-point numbers $|x \text{ REM } y| \leq y/2$. Although this defines REM precisely, it is not a practical operational definition, because n can be huge. In single precision, n could be as large as $2^{127}/2^{-126} = 2^{253} \approx 10^{76}$.

There is a natural way to compute REM if a direct division algorithm is used. Proceed as if you were computing x/y. If $x = s_1 2^{e_1}$ and $y = s_2 2^{e_2}$ and the divider is as in Figure A.2(b) (page A-4), then load s_1 into P and s_2 into B. After $e_1 - e_2$ division steps, the P register will hold a number r of the form $x - yn$ satisfying $0 \leq r < y$. The IEEE remainder is then either r or $r - y$. It is only necessary to keep track of the last quotient bit produced, which is needed in order to resolve halfway cases. Unfortunately, $e_1 - e_2$ can be a lot of steps, and floating-point units typically have a maximum amount of time they are allowed to spend on one instruction. Thus, it is usually not possible to implement REM directly. None of the chips discussed in Section A.10 implement REM, but they could by providing a remainder-step instruction—this is what is done on the Intel 8087 family. A remainder step takes as arguments two numbers x and y, and performs divide steps until either the remainder is in P, or else n steps have been performed, where n is a small number, such as the number of steps required for division in the highest supported precision. The REM driver calls the REM-step instruction $\lfloor (e_1 - e_2)/n \rfloor$ times, initially using x as the numerator, but then replacing it with the remainder from the previous REM step. It is useful if the REM-step instruction returns the low-order three bits of the quotient, since when doing trigonometric argument reduction to the interval $(0, \pi/4)$, you need to know the value of n mod 8 in order to know what quadrant you are in.

Currently, most of the fastest floating-point chips don't implement remainder, even though it is a required part of the IEEE standard. Since the standard allows implementations to be a combination of hardware and software, the REM operation could be implemented entirely in software. However, availability of the REM-step instruction would make computing REM much simpler. Is a REM-step instruction worth it? For two reasons this situation is difficult to decide on the basis of frequency data. First, because REM is peculiar to the IEEE standard, few people are currently using it. Testing the demand for REM is somewhat like trying to estimate the demand for a new product. Second, the main benefit from REM is not an increase in performance, but rather an increase in accuracy, and it is not easy to quantify the value of accuracy. What we will do here is simply present the primary application of REM, which is argument reduction for periodic functions, like sin and cos.

There are some subtle issues involved in argument reduction. To simplify things, imagine that we are working in base 10 with 5 significant figures, and consider computing $\sin x$. Suppose that $x = 7$. Then we reduce by $\pi = 3.1416$

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 655

and compute $\sin(7) = \sin(7 - 2\times3.1416) = \sin(0.7168)$ instead. But suppose we want to compute $\sin(2.0 \times 10^5)$. Then $2 \times 10^5/3.1416 = 63661.8$, which in our 5-place system comes out to be 63662. Since multiplying 3.1416 times 63662 gives 200000.5392, which rounds to 2.0000×10^5, argument reduction reduces 2×10^5 to 0, which is not even close to being correct. The problem is that our 5-place system does not have the precision to do correct argument reduction. Suppose we had the REM operator. Then we could compute 2×10^5 REM 3.1416 and get −.5392. However, this is still not correct because we used 3.1416, which is an approximation for π. The value of 2×10^5 REM π is −.071513. The difficulty is that we subtracted two nearby numbers, 2×10^5 and 63662×3.1416, where 63662×3.1416 was slightly in error due to approximating π. Even though REM has the effect of performing the subtraction exactly, all the significant figures in 63662×3.1416 canceled, leaving behind only rounding error.

Traditionally, there have been two approaches to computing periodic functions with large arguments. The first is to return an error for their value when x is large. The second is to store π to a very large number of places and do exact argument reduction. The REM operator is not much help in either of these situations. There is a third approach that has been used in some math libraries, such as the Berkeley UNIX 4.3bsd release. In these libraries, π is computed to the nearest floating-point number. Let's call this machine π, and denote it by π'. Then when computing $\sin x$, reduce x using x REM π'. As we saw in the above example, x REM π' is quite different from x REM π, so that computing $\sin x$ as $\sin(x \text{ REM } \pi')$ will not give the exact value of $\sin x$. However, computing trigonometric functions in this fashion has the property that all familiar identities (such as $\sin^2 x + \cos^2 x = 1$) are true to within a few rounding errors. Thus, using REM together with machine π provides a simple method of computing trigonometric functions that is accurate for small arguments and still useful for large arguments in most applications.

A.7 Precisions and Exception Handling

Precisions

Implementations of the IEEE standard are only required to support single precision. Thus, the computer designer must make a choice about what other precisions to support. Because of the widespread use of double precision in scientific computing, double precision is almost always implemented.

Double-extended precision is more problematic. Although the Motorola 68882 and Intel 387 coprocessors implement extended precision, most of the more recently designed, high-performance floating-point chips do not implement extended precision. Among the reasons are that the 80-bit width of extended precision is awkward for 64-bit buses and registers, and that many high-level languages do not give the user access to extended precision. However, extended

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 656

precision is very useful to writers of mathematical software. As an example, consider writing a library routine to compute the length of a vector in the plane $\sqrt{x^2 + y^2}$. If x is larger than $2^{E_{max}/2}$, then computing this in the obvious way will overflow. This means that either the allowable exponent range for this subroutine will be cut in half, or a more complex algorithm using scaling will have to be employed. But if extended precision is available, then the simple algorithm will work. Computing the length of a vector is a simple task, and it is not difficult to come up with an algorithm that doesn't overflow. However, there are more complex problems for which extended precision means the difference between a simple, fast algorithm and a much more complex one. One of the best examples of this is binary/decimal conversion. An efficient algorithm for binary-to-decimal conversion that makes essential use of extended precision is very readably presented in Coonen [1984]. This algorithm is also briefly sketched in Goldberg [1989]. Computing accurate values for transcendental functions is another example of a problem that is made much easier if extended precision is present.

One very important fact about precision concerns *double rounding*. To illustrate in decimal, suppose that we want to compute 1.9×0.66, and that single precision is two digits, while extended precision is three digits. The exact result of the product is 1.254. Rounded to extended precision, the result is 1.25. When further rounded to single precision, we get 1.2. However, the result of 1.9×0.66 correctly rounded to single precision is 1.3. Thus, rounding twice may not produce the same result as rounding once. Suppose you want to build hardware that only does double-precision arithmetic. Can you simulate single precision by computing first in double precision and then rounding to single? The above example suggests that you can't. However, double rounding is not always dangerous. In fact, the following rule is true (although it is not easy to prove).

> *If x and y have p-bit significands, and $x + y$ is computed exactly and then rounded to q places, a second rounding to p places will not change the answer if $p \le (q-1)/2$. This is true not only for addition, but also for multiplication, division, and square root.*

In our example above, $q = 3$, and $p = 2$, so $2 \le (3 - 1)/2$ is not true. On the other hand, for IEEE arithmetic, double precison has $p = 53$, and single precision is $p = 24 \le (q - 1)/2 = 26$. Thus, single precision can be implemented by computing in double precision (that is, computing the answer exactly and then rounding to double) and then rounding to single precision.

The standard requires implementations to provide versions of addition, subtraction, multiplication, division, and remainder that take two operands of the same precision and produce a result of that precision. It also recommends that implementations allow operations that take operands of two different precisions and return a result whose precision is at least as wide as the widest operand. The standard allows implementations to combine two operands and return a result in a higher precision. Remember that the result of an operation is the exact result

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 657

rounded to the destination precision. What the standard does not allow is combining two operands and returning a result in a lower precision. Although at first this may seem like a minor restriction, consider again the problem of computing $\sqrt{x^2+y^2}$. If x and y are double, then you might like to compute $x^2 + y^2$ in extended precision and then compute a square root that takes an extended-precision argument and returns a double-precision answer. But this is not allowed by the standard.

There is a related issue. The standard permits combining two extended variables to produce a result that is stored in extended format, but rounded to double precision. However, this doesn't help in the square root example, because the result of the square root must still be explicitly converted from an extended format to a double-precision format.

Exceptions

The IEEE standard defines five exceptions: underflow, overflow, divide by zero, inexact, and invalid. By default, when these exceptions occur, they merely set a flag and the computation continues. The flags are *sticky*, meaning that once set they remain set until explicitly cleared. The standard strongly encourages implementations to provide a trap-enable bit for each exception. When an exception with an enabled trap handler occurs, a user trap handler is called, and the value of the associated exception flag is undefined.

The underflow, overflow, and divide-by-zero exceptions are found in most other systems. The *inexact exception* is peculiar to IEEE arithmetic and occurs when either the result of an operation must be rounded or when it overflows. In fact, since 1/0 and an operation that overflows both deliver ∞, the exception flags must be consulted to distinguish between them. The inexact exception is an unusual "exception," in that it is not really an exceptional condition because it occurs so frequently. Thus, enabling a trap handler for inexact will most likely have a severe impact on performance. The *invalid exception* is for things like $\sqrt{-1}$, 0/0 or $\infty - \infty$, which don't have any natural value as a floating-point number or as $\pm\infty$. Thus, 1/0 causes a divide by zero exception and delivers ∞, whereas 0/0 causes an invalid exception and delivers a NaN. There is a twist in IEEE underflow, because it is not always signaled when numbers fall below $1.0 \times 2^{E_{\min}}$. If a user trap handler is not installed, then underflow is signaled only if the result of an operation is below $2^{E_{\min}}$ and is inexact.

The IEEE standard assumes that when a trap occurs, it is possible to identify the operation that trapped and its operands. On machines with pipelining, or machines with multiple arithmetic units, when an exception occurs, it may not be enough to simply have the trap handler examine the program counter. Hardware support may be necessary in order to identify exactly which operation trapped. Another problem is illustrated by the following program fragment.

```
X = Y * Z;
Z = A + B;
```

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 658

These two instructions might well be executed in parallel. If the multiply traps, its argument Z could already have been overwritten by the addition, especially since addition is usually faster than multiplication. Computer systems that support trapping in the IEEE standard must provide some way to save the value of Z, either in hardware or by having the compiler avoid such a situation in the first place.

One approach to this problem, used in the MIPS R3010, is to treat floating-point exceptions similarly to page-fault exceptions. If an instruction that assigns a memory location to a register causes a page fault, the execution of the instruction must stall before it clobbers the register because (for example) that very register might be used to reference the memory that faulted. The key to making this work is that the memory address is computed early in the instruction cycle, before the instruction actually writes anything. A similar trick can be done with floating-point operations. An instruction that may cause an exception can be identified early in the instruction cycle. For example, an addition can overflow only if one of the operands has an exponent of E_{max}, and so on. This early check is conservative: It might flag an operation that doesn't actually cause an exception. However, if such false positives are rare, then this technique will have excellent performance. When an instruction is tagged as being possibly exceptional, special code in a trap handler can compute it without destroying any state. Remember that all these problems occur only when trap handlers are enabled. Otherwise, setting the exception flags during normal processing is straightforward.

There is a subtlety that should be mentioned that involves the underflow trap. When there is no underflow trap handler, the result of an operation that involves an underflow is a denormal number. When there is a trap handler, it is provided with the result of the operation with the exponent wrapped around. Now there is a potential double-rounding problem. If the rounding mode is round toward nearest, when there is a trap handler the result is correctly rounded to p significant bits. If there is no trap handler, the result is rounded to less than p bits, depending on how many leading zeros the denormal number has. If the trap handler wants to return the denormal result, it can't just round its argument, because that might lead to a double-rounding error. Thus, the trap handler must be passed at least one extra bit of information if it is to be able to deliver the correctly rounded result.

A.8 Speeding Up Integer Addition

The previous section showed that there are many steps that go into implementing floating-point operations. However, each floating-point operation eventually reduces to an integer operation. Thus, increasing the speed of integer operations will also lead to faster floating point.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 659

Integer addition is the simplest operation and the most important. Even for programs that don't do explicit arithmetic, addition must be performed to increment the program counter and to do address calculations. Despite the simplicity of addition, there isn't a single best way to perform high-speed addition. We will discuss three techniques that are in current use: carry lookahead, carry skip, and carry select.

Carry Lookahead

An n-bit adder is just a combinational circuit. It can therefore be written by a logic formula whose form is a sum of products and can be computed by a circuit with two levels of logic. How does one figure out what this circuit looks like? Recall from Equation A.2.1 that the formula for the ith sum bit is

A.8.1 $$s_i = a_i\bar{b}_i\bar{c}_i + \bar{a}_i b_i \bar{c}_i + \bar{a}_i\bar{b}_i c_i + a_i b_i c_i$$

The problem with this formula is that although we know the values of a_i and b_i—they are inputs to the circuit—we don't know c_i. So our goal is to write c_i in terms of a_i and b_i. To accomplish this, we first rewrite Equation A.2.2 (page A-2) as

A.8.2 $$c_{i+1} = g_i + p_i c_i, \quad g_i = a_i b_i, \quad p_i = a_i + b_i$$

Here is the reason for the symbols p and g: If g_i is true, then c_{i+1} is certainly true, so a carry is *generated.* Thus, g is for generate. If p_i is true, then if c_i is true, it is *propagated* to c_{i+1}. Start with Equation A.8.1 and use Equation A.8.2 to replace c_i with $g_{i-1} + p_{i-1}c_{i-1}$. Then, use Equation A.8.2 with $i - 1$ in place of i, to replace c_{i-1} with c_{i-2}, and so on. This gives the result

A.8.3 $$c_{i+1} = g_i + p_i g_{i-1} + p_i p_{i-1} g_{i-2} + \cdots + p_i p_{i-1} \cdots p_1 g_0 + p_i p_{i-1} \cdots p_1 p_0 c_0$$

An adder that computes carries using Equation A.8.3 is called a *carry-lookahead adder,* or CLA adder. A CLA adder requires one logic level to form p and g, two levels to form the carries, and two for the sum, for a grand total of five logic levels. This is a vast improvement over the $2n$ levels required for the ripple-carry adder.

Unfortunately, as is evident from Equation A.8.3 or from Figure A.10, a carry-lookahead adder on n bits requires a fan-in of $n+1$ at the OR gate as well as at the rightmost AND gate. Also, the p_{n-1} signal must drive n AND gates. In addition, the rather irregular structure and many long wires of Figure A.10 make it impractical to build a full carry-lookahead adder when n is large.

However, we can use the carry-lookahead idea to build an adder that has about $\log_2 n$ logic levels (substantially less than the $2n$ required by a ripple-carry adder), and yet has a simple, regular structure. The idea is to build up the p's and g's in steps. We have already seen that

$$c_1 = g_0 + c_0 p_0$$

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 660

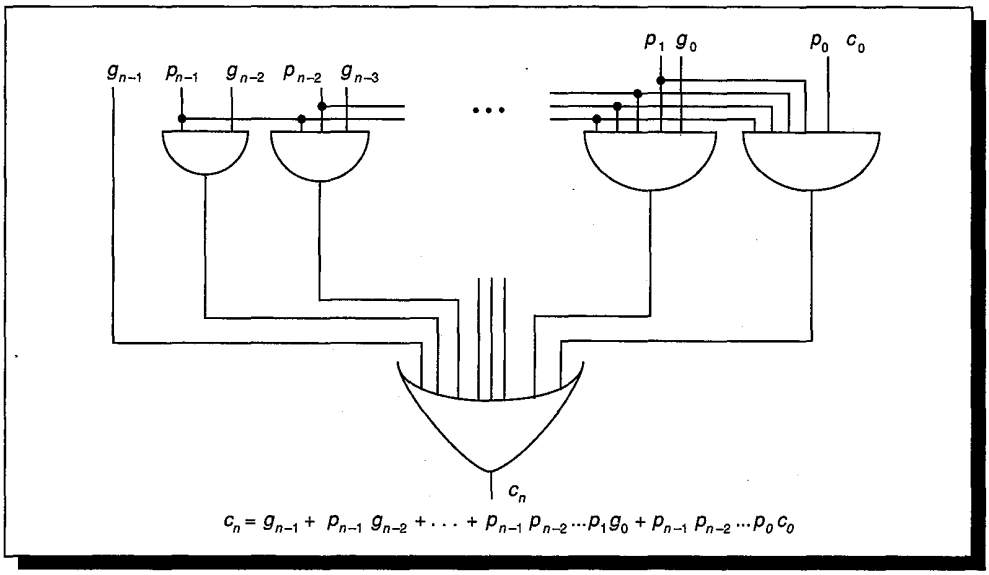

FIGURE A.10 Pure carry-lookahead circuit for computing the carry out c_n of an n-bit adder.

This says there is a carry out of the 0th position (c_1) if there is either a carry generated in the 0th position, or if there is a carry into the 0th position and the carry propagates. Similarly,

$$c_2 = G_{01} + P_{01}c_0$$

G_{01} means there is a carry generated out of the block consisting of the first two bits. P_{01} means that a carry propagates through this block. P and G have the following logic equations:

$$\begin{aligned} G_{01} &= g_1 + p_1 g_0 \\ P_{01} &= p_1 p_0 \end{aligned}$$

More generally, for any j with $i < j, j + 1 < k$, we have the recursive relations

A.8.4
$$c_{k+1} = G_{ik} + P_{ik}c_i$$

A.8.5
$$G_{ik} = G_{j+1,k} + P_{j+1,k}G_{ij}$$

A.8.6
$$P_{ik} = P_{ij}P_{j+1,k}$$

Equation A.8.5 says that a carry is generated out of the block consisting of bits i through k inclusive if it is generated in the high-order part of the block $(j + 1, k)$ or if it is generated in the low-order (i,j) part of the block and then propagated through the high part. These equations will also hold for $i \le j < k$ if we set $G_{ii} = g_i$ and $P_{ii} = p_i$.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 661

Example: Express P_{03} and G_{03} in terms of p's and g's.

Answer: Using A.8.6, $P_{03} = P_{01}P_{23} = P_{00}P_{11}P_{22}P_{33}$. Since $P_{ii} = p_i$, $P_{03} = p_0p_1p_2p_3$. For G_{03}, Equation A.8.5 says $G_{03} = G_{23} + P_{23}G_{01} = (G_{33} + P_{33}G_{22}) + (P_{22}P_{33})(G_{11} + P_{11}G_{00}) = g_3 + p_3g_2 + p_3p_2g_1 + p_3p_2p_1g_0$.

With these preliminaries out of the way, we can now show the design of a practical CLA adder. The adder consists of two parts. The first part computes various values of P and G from p_i and g_i, using Equations A.8.5 and A.8.6; the second part uses these P and G values to compute all the carries via Equation A.8.4. The first part of the design is in Figure A.11. At the top of the diagram, input numbers $a_7 \cdots a_0$ and $b_7 \cdots b_0$ are converted to p's and g's using cells of type 1. Then various P's and G's are generated by combining cells of type 2 in a binary-tree structure. The second part of the design is shown in Figure A.12. By feeding c_0 in at the bottom of this tree, all the carry bits come out the top. Each cell must know a pair of (P,G) values in order to do the conversion, and the value it needs is written inside the cells. Now compare Figure A.11 and Figure A.12. There is a one-to-one correspondence between cells, and the value of (P,G) needed by the carry-generating cells is exactly the value known by the

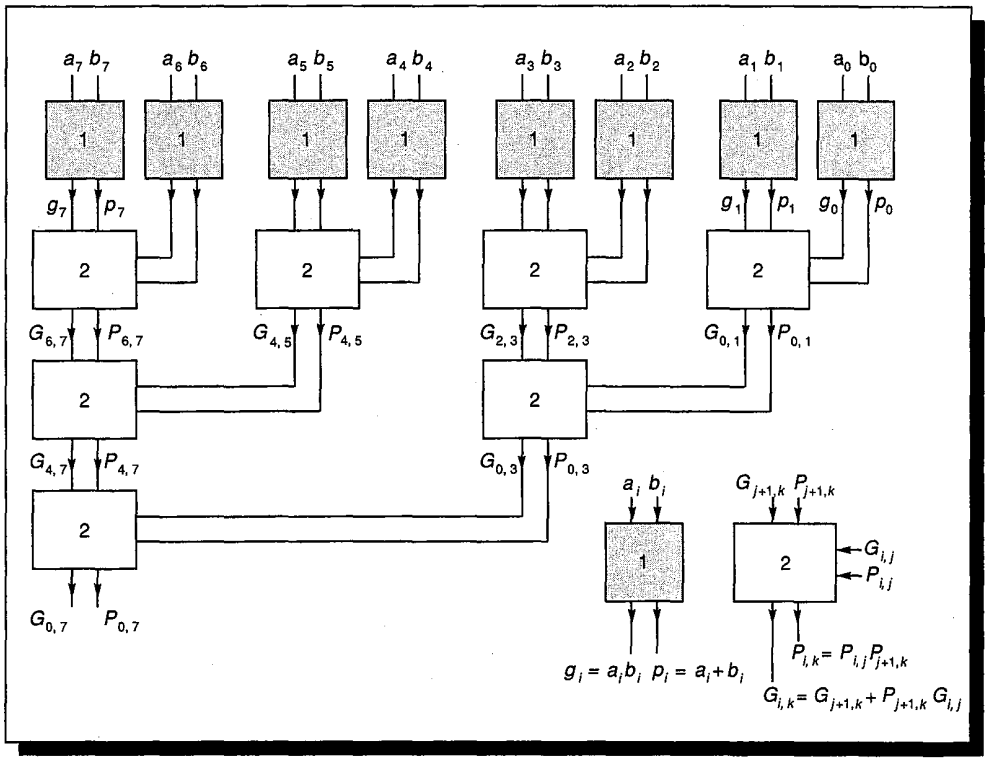

FIGURE A.11 First part of carry-lookahead tree. As signals flow from the top to the bottom, various values of P and G are computed.

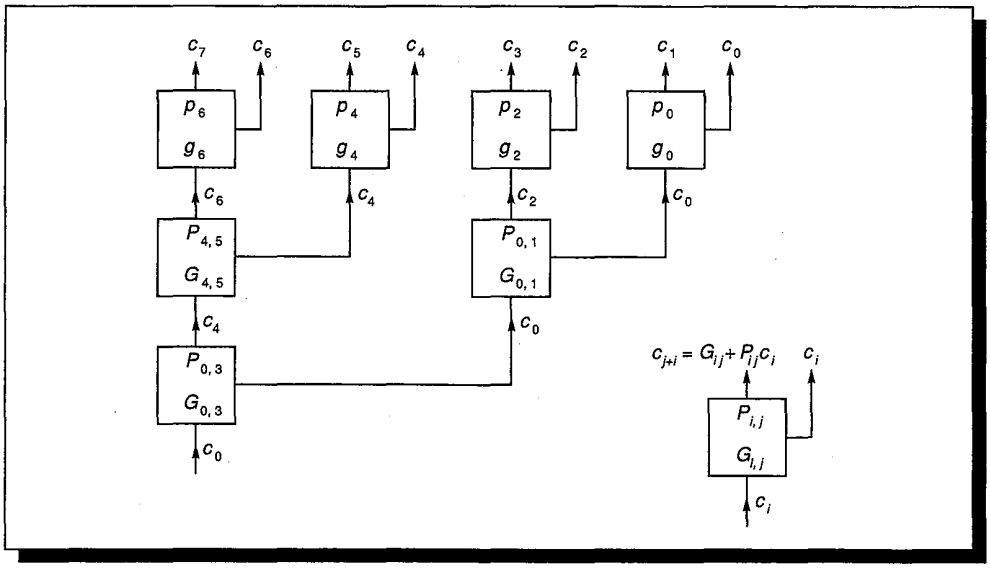

FIGURE A.12 Second part of carry-lookahead tree. Signals flow from the bottom to the top, combining with P and G to form the carries.

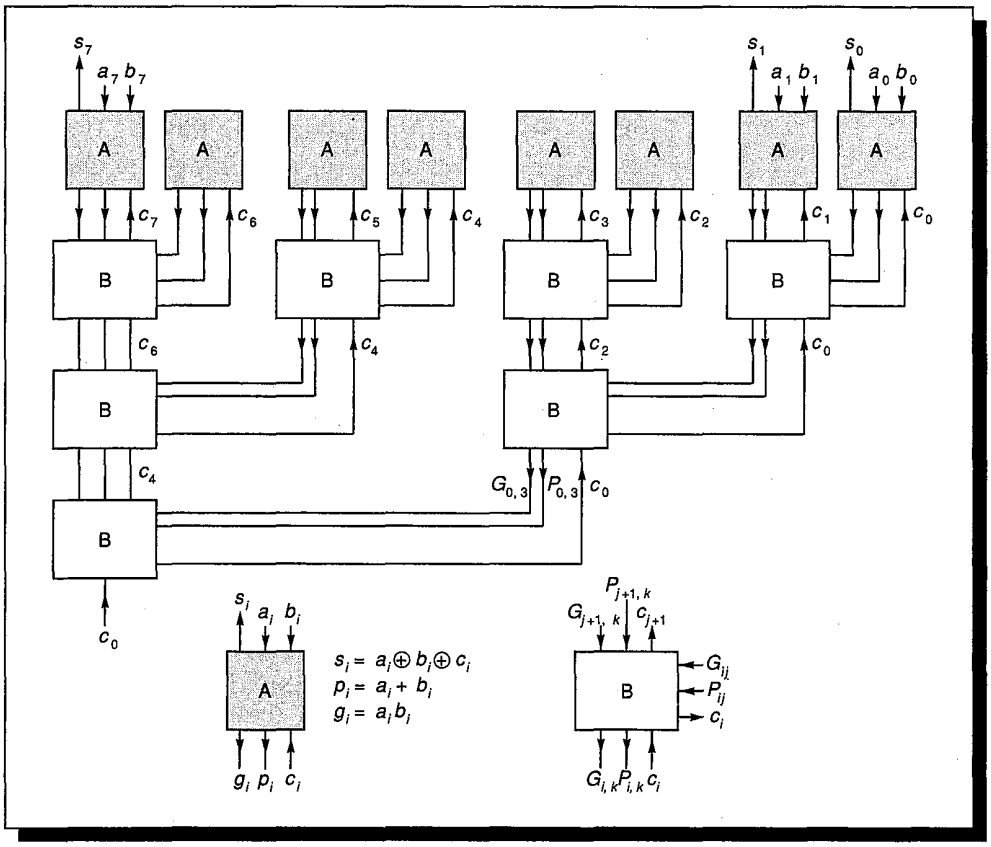

FIGURE A.13 Complete carry-lookahead tree adder. This is the combination of Figures A.11 and A.12. The numbers to be added enter at the top, flow to the bottom to combine with c_0, and then flow back up to compute the sum bits.

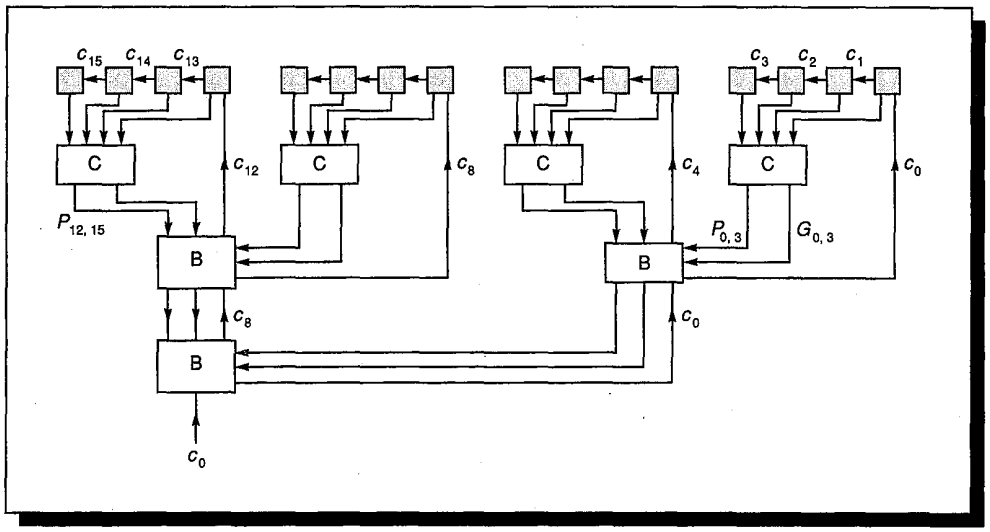

FIGURE A.14 Combination of CLA adder and ripple-carry adder. In the top row, carries ripple within each group of four boxes.

corresponding (P,G) generating cells. The combined cell is shown in Figure A.13. The numbers to be added flow into the top and downward through the tree, combining with c_0 at the bottom and flowing back up the tree to form the carries. Note that there is one thing missing from Figure A.13: a small piece of extra logic to compute c_8 for the carry out of the adder.

The bits in a CLA must pass through about $\log_2 n$ logic levels, compared with $2n$ for a ripple-carry adder. This is a substantial speed improvement, especially for a large n. Whereas the ripple-carry adder had n cells, however, the CLA adder has $2n$ cells, although in our layout they will take $n \log n$ space. The point is that a small investment in size pays off in a dramatic improvement in speed.

There are a number of technology-dependent modifications that can improve CLA adders. For example, if each node of the tree has three inputs instead of two, then the height of the tree will decrease from $\log_2 n$ to $\log_3 n$. Of course, the cells will be more complex and thus might operate more slowly, negating the advantage of the decreased height. For technologies where rippling works well, a hybrid design might be better. This is illustrated in Figure A.14. Carries ripple between adders at the top level, while the "B" boxes are the same as in Figure A.13. This design will be faster if the time to ripple between four adders is faster than the time it takes to traverse a level of "B" boxes.

Carry-Skip Adders

A *carry-skip adder* sits midway between a ripple-carry adder and a carry-lookahead adder, both in terms of speed and cost. (A carry-skip adder is not called a CSA, as that name is reserved for carry-save adders.) The motivation for this adder comes from examining the equations for P and G. For example,

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 664

$$P_{03} = p_0 p_1 p_2 p_3$$

$$G_{03} = g_3 + p_3 g_2 + p_3 p_2 g_1 + p_3 p_2 p_1 g_0$$

Computing P is much simpler than computing G, and a carry-skip adder only computes the P's. Such an adder is illustrated in Figure A.15. Carries begin rippling simultaneously through each block. If any block generates a carry, then the carry out of a block will be true, even though the carry in to the block may not be correct yet. If at the start of each add operation the carry in to each block is zero, then no spurious carry outs will be generated. Thus, the carry out of each block can thus be thought of as if it were the G signal. Once the carry out from the least significant block is generated, it not only feeds into the next block, but is also fed through the AND gate with the P signal from that next block. If the carry out and P signals are both true, then the carry **skips** the second block and is ready to feed into the third block, and so on. The carry-skip adder is only practical if the carry in signals can be easily cleared at the start of each operation—for example by precharging in CMOS.

To analyze the speed of a carry-skip adder, let's assume that it takes one time unit for a signal to pass through two logic levels. Then it will take k time units for a carry to ripple across a block of size k, and it will take one time unit for a carry to skip a block. The longest signal path in the carry-skip adder starts with a carry being generated at the 0th position. Then it takes k time units to ripple through the first block, $n/k - 2$ time units to skip blocks, and k more to ripple through the last block. To be specific: If we have a 20-bit adder broken into groups of 4 bits, it will take 11 time units to perform an add. Suppose we keep the least significant block at 4 bits, but combine the next two blocks into a single 8-bit block. Then the time of the adder drops to 10 time units. However, if we had combined three blocks instead of two, then the time to ripple through this 3-block unit (12 bits in all) would dominate the time to add. However, the general principle is important: For a carry-skip adder, making the interior blocks larger will speed up the adder. In fact, the same idea of varying the block sizes can sometimes speed up other adder designs as well. Because of the large amount of rippling, a carry-skip adder is most appropriate for technologies where rippling is fast.

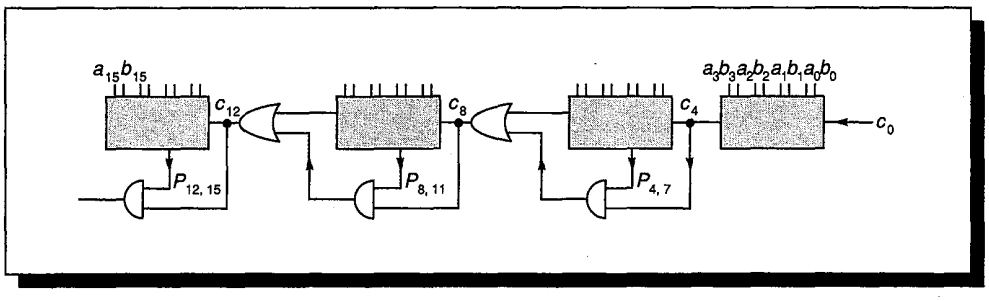

FIGURE A.15 Carry-skip adder.

Carry-Select Adder

A *carry-select adder* works on the following principle: Two additions are performed in parallel, one assuming the carry in is zero and the other assuming the carry in is one. When the carry in is finally known, the correct sum (which has been precomputed) is simply selected. An example of such a design is shown in Figure A.16. An 8-bit adder is divided into two halves, and the carry out from the lower half is used to select the upper half. If each block is computing its sum using rippling (a linear-time algorithm), then the design in Figure A.16 is twice

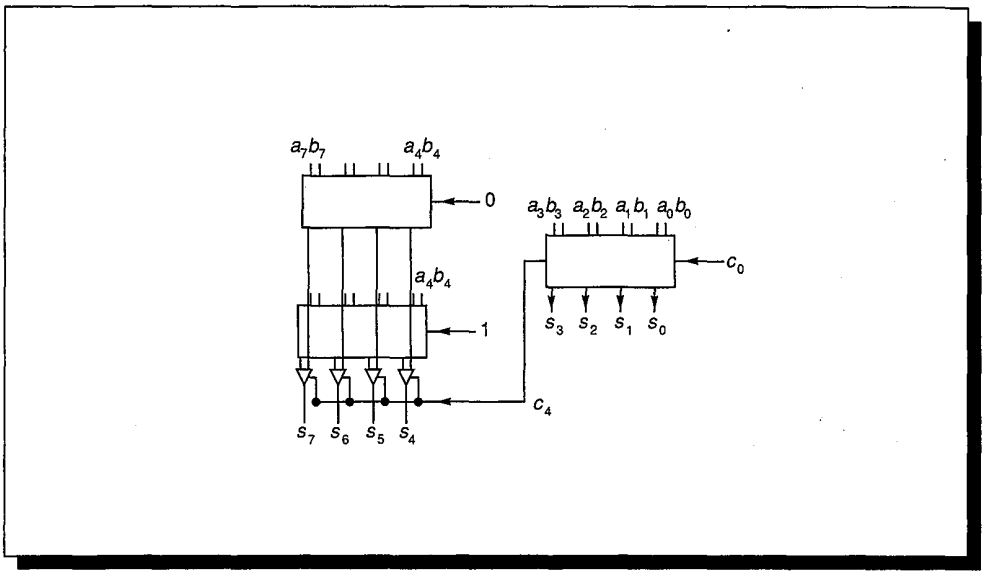

FIGURE A.16 Simple carry-select adder. At the same time that the sum of the low-order four bits are being computed, the high-order bits are being computed twice in parallel: once assuming that $c_4 = 0$, and once assuming $c_4 = 1$.

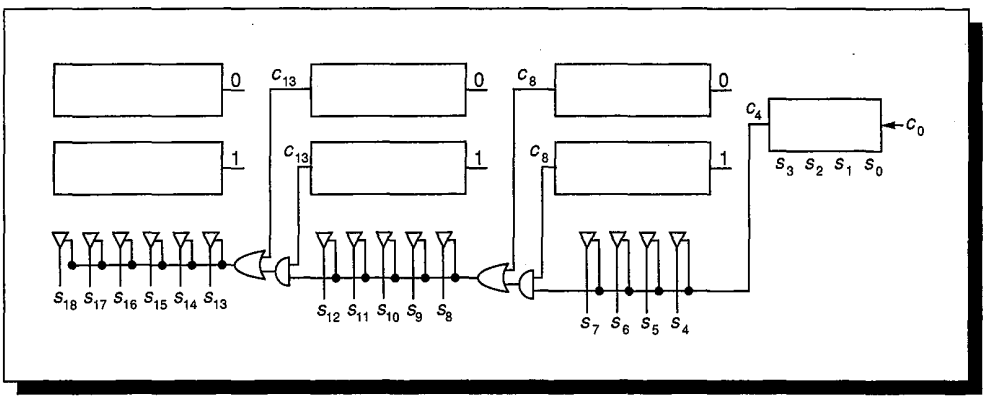

FIGURE A.17 Carry-select adder. As soon as the carry out of the rightmost block is known, it is used to select the other sum bits.

as fast at 50% more cost. However, note that the c_4 signal must drive many muxes, which may be very slow in some technologies. Instead of dividing the adder into halves, it could be divided into quarters for a still further speedup. This is illustrated in Figure A.17. If it takes k time units for a block to add k-bit numbers, and if it takes one time unit to compute the mux input from the two carry-out signals, then for optimal operation each block should be one bit wider than the next, as shown in Figure A.17. Therefore, as in the carry-skip adder, the best design involves variable-sized blocks.

As a summary of this section, the asymptotic time and space requirements for the different adders are given in Figure A.18. These different adders shouldn't be thought of as disjoint choices, but rather as building blocks to be used in constructing an adder. The utility of these different building blocks is highly dependent on the technology used. For example, the carry-select adder works well when a signal can drive many muxes, and the carry-skip adder is attractive in technologies where signals can be cleared at the start of each operation. Knowing the asymptotic behavior of adders is useful in understanding them, but relying too much on that behavior is a pitfall. The reason is that asymptotic behavior is only important as n grows very large. But n for an adder is the bits of precision, and double precision today is the same as it was twenty years ago—about 53 bits. Although it is true that as computers get faster, computations get longer—and thus have more rounding error, which in turn requires more precision—this effect grows very slowly with time.

	Time	Space
Ripple	$O(n)$	$O(n)$
CLA	$O(\log n)$	$O(n \log n)$
Carry skip	$O(\sqrt{n})$	$O(n)$
Carry select	$O(\sqrt{n})$	$O(n)$

FIGURE A.18 Asymptotic time and space requirements for four different types of adders.

A.9 Speeding Up Integer Multiplication and Division

The multiplication and division algorithms presented in Section A.2 are fairly slow, producing one bit per cycle (although that cycle might be a fraction of the CPU instruction cycle time). In this section we discuss various techniques for higher performance multiplication and division.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 667

Shifting Over Zeros

Shifting over zeros is a technique that is not currently used much, but is instructive to consider. It is distinguished by the fact that its execution time is operand dependent. Its lack of use is primarily attributable to its failure to offer enough speedup over bit-at-a-time algorithms. In addition, pipelining, synchronization with the CPU, and good compiler optimization are difficult with algorithms that run in variable time. In multiplication, the idea behind shifting over zeros is to add logic that detects when the low-order bit of the A register is zero (see Figure A.2(a)) and, if so, skip the addition step and proceed directly to the shift step—hence the term *shifting over zeros*. This technique becomes more useful if the number of zeros in the A operand can be increased. The Exercises discuss how well Booth recoding does in increasing zeros.

What about shifting for division? In nonrestoring division, an ALU operation (either an addition or subtraction) is performed at every step, so that there appears to be no opportunity for skipping an operation. But think about division this way: To compute a/b, subtract multiples of b from a, and then report how many subtractions were done. At each stage of the subtraction process the remainder must fit into the P register of Figure A.2(b) (page A-4). In the case when the remainder is a small positive number, you normally subtract b; but suppose instead you only shifted the remainder and subtracted b the next time. As long as the remainder was sufficiently small (its high-order bit 0), after shifting it still would fit into the P register, and no information would be lost. However, this method does require changing the way we keep track of the number of times b has been subtracted from a. This idea usually goes under the name of *SRT division*, for Sweeney, Robertson, and Tocher, who independently proposed algorithms of this nature. The main extra complication of SRT division is that the quotient bits cannot be determined immediately from the sign of P at each step, as it can be in ordinary nonrestoring division.

More precisely, to divide a by b where a and b are n-bit numbers, load a and b into the A and B registers, respectively, of Figure A.2 (page A-4).

1. If B has k leading zeros when expressed using n bits, shift all the registers left k bits. After this shift, since b has $n + 1$ bits, its most significant bit will be 0, and its second-most-significant bit will be 1.

2. For $i = 0, n - 1$ do

   If the top three bits of P are equal, set $q_i = 0$ and shift (P,A) one bit left.

   If the top three bits of P are not all equal and P is negative, set $q_i = \overline{1}$, shift (P,A) one bit left, and add B.

   Otherwise set $q_i = 1$, shift (P,A) one bit left, and subtract B

   Endloop

3. If the final remainder is negative, correct the remainder by adding B, and correct the quotient by subtracting 1 from q_0. Finally, the remainder must be shifted k bits right, where k is the initial shift.

A numerical example is given in Figure A.19. Although we are discussing integer division, it helps in explaining the algorithm to move the binary point from the right of the least significant bit to the left of the most significant bit. Thus if $n = 4$ and the operation is 9/4, the A register holds 0.1001 and (remembering that the B register has $n+1$ bits), the B register holds 0.0100.

Since this changes the binary point in both the numerator and denominator, the quotient is not affected. The remainder being a two's complement number, a P register of 1.1110_2 represents $-1/8$. With this convention, the P register holds numbers satisfying $-1 \le P < 1$. The first step of the algorithm shifts b so that $b \ge 1/2$. As before, let r be the value of the (P,A) pair. Our rule for which ALU operation to perform is this: If $-1/4 \le r < 1/4$ (true whenever the top three bits of P are equal), then compute $2r$ by shifting (P,A) left one bit; else if $r < 0$ (and hence $r < -1/4$, since otherwise it would have been eliminated by the first condition), then compute $2r + b$ by shifting and then adding, else $r \ge 1/4$ and subtract b from $2r$. Using $b \ge 1/2$, it is easy to check that these rules keep $-1/2 \le r < 1/2$. For nonrestoring division, we only have $|r| \le b$, and we need P to be $n+1$ bits wide. But for SRT division, the bound on r is tighter, namely $-1/2 \le r < 1/2$. Thus, we can save a bit by eliminating the high-order bit of P (and b and the adder). In particular, the test for equality of the top three bits of P becomes a test on just two bits.

P	A	
00000	1000	B contains 0011, so shift all registers left two places
00010	0000	B now contains 1100. Top bits of P are equal, so shift and set $q_0 = 0$
00100	0000	Top bits are not equal, so set $q_1 = 1$
01000	0000	shift and
+10100		subtract B
11100	0000	Top bits equal, so shift and set $q_2 = 0$
11000	0000	Top bits are unequal, so set $q_3 = -1$
10000	0000	shift and
+01100		add B
11100		Remainder is negative, so restore it and subtract 1 from q_0
+01100		
01000		This must be shifted right two places to give remainder
		Remainder = 10, $q = 0101\bar{} - 1 = 0010$

FIGURE A.19 SRT division of 1000/0011.

The algorithm might change slightly in an implementation of SRT division. After each ALU operation, the P register can be shifted as many places as necessary to make either $P \geq 1/4$ or $P < -1/4$. By shifting k places, k quotient bits are set equal to zero all at once. For this reason SRT division is sometimes described as one that keeps the remainder normalized to $|r| \geq 1/4$.

Notice that the value of the quotient bit computed in a given step is based on which operation is performed in that step (which in turn depends on the result of the operation from the previous step). This is in contrast to nonrestoring division, where the quotient bit computed in ith step depends on the result of the operation in the same step. This difference is reflected in the fact that when the final remainder is negative, the last quotient bit must be adjusted in SRT division, but not in nonrestoring division. However, the key fact about the quotient bits in SRT division is that they can include $\overline{1}$. Therefore the quotient bits can't be stored in the low-order bits of the A register; furthermore, the quotient must be converted to ordinary two's complement in a full adder. A common way to do this is to accumulate the positive quotient bits in one register and the negative quotient bits in another, and then subtract the two registers after all the bits are known. Because there is more than one way to write a number in terms of the digits -1, 0, 1, SRT division is said to use a *redundant* quotient representation.

The differences between SRT division and ordinary nonrestoring division can be summarized as follows:

1. ALU decision rule: In nonrestoring division, it is determined by the sign of P; in SRT, it is determined by the two most significant bits of P.

2. Quotient determination: In nonrestoring division, it is immediate from the signs of P; in SRT, it must be computed in a full n-bit adder.

3. Speed: SRT division will be faster on operands that produce zero quotient bits.

Speeding Up Multiplication with a Single Adder

As mentioned before, shifting-over techniques are not used much in current hardware. We now discuss some methods that are in more widespread use. Methods that increase the speed of multiplication can be divided into two classes: those that use a single adder and those that use multiple adders. Let's first discuss techniques that use a single adder.

In the discussion of addition we noted that, because of carry propagation, it is not practical to perform addition with two levels of logic. Using the cells of Figure A.13, adding two 64-bit numbers will require a trip through seven cells to compute the P's and G's, and seven more to compute the carry bits, which will require at least 28 logic levels. Each multiplication step will require a trip through this adder. A way to avoid this computation in each step is to use *carry-save adders* (CSA). A carry-save adder is simply n independent full adders. A

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 670

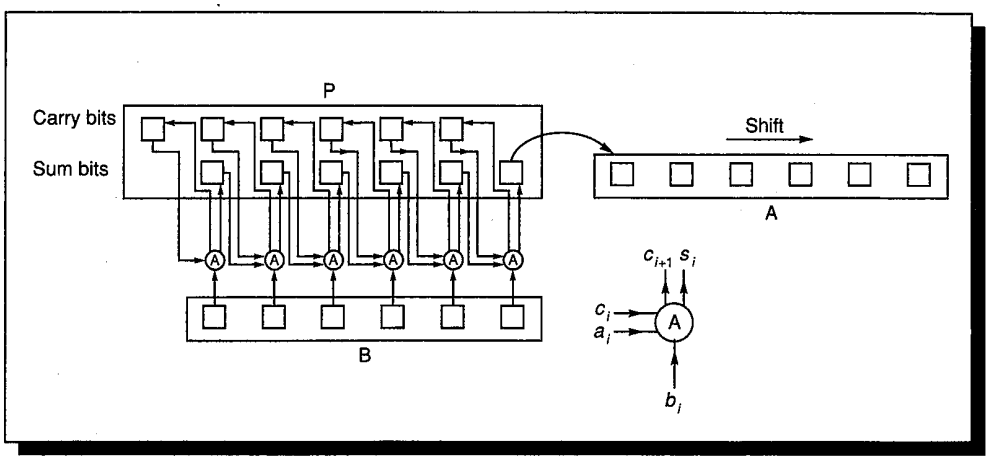

FIGURE A.20 Carry-save multiplier. Each circle represents a (3,2) adder working independently. At each step, the only bit of P that needs to be shifted is the low-order sum bit.

multiplier using such an adder is illustrated in Figure A.20. Each circle marked "A" is a single-bit full adder, and each box represents one bit of a register. Each addition operation results in a pair of bits, stored in the sum and carry parts of P. Since each add is independent, only two logic levels are involved in the add—a vast improvement over 28.

To operate the multiplier in Figure A.20, load the sum and carry bits of P with zero and perform the first ALU operation. (If Booth recoding is used, it might be a subtraction rather than an addition.) Then shift the low-order sum bit of P into A, as well as shifting A itself. The $n - 1$ high-order bits of P don't need to be shifted because on the next cycle the sum bits are fed into the next lower order adder. Each addition step is dramatically increased in speed, since each add cell is working independently of the others, and no carry is propagated. There are two drawbacks to carry-save adders. First, they require more hardware because there must be a copy of register P to hold the carry outputs of the adder. Second, after the last step, the high-order word of the result must be fed into an ordinary adder to combine the sum and carry parts. This could be accomplished by feeding the output of P into the adder used to perform the addition operation. Multiplying with a carry-save adder is sometimes called redundant multiplication because P is represented using two registers. Since there are many ways to represent P as the sum of two registers, this representation is redundant. The term *carry-propagate adder* (CPA) is used to denote an adder that is not a CSA. A propagate adder may propagate its carries using ripples, carry lookahead, or some other method.

Another way to speed up multiplication without using extra adders is to examine k low-order bits of A at each step, rather than just one bit. This is often called *higher-radix multiplication*. As an example, suppose that $k = 2$. If the pair of bits is 00, add 0 to P, and if it is 01, add B. If it is 10, simply shift b one bit left before adding it to P. Unfortunately, if the pair is 11, it appears we would

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 671

have to compute $b + 2b$. But this can be avoided by using a higher-radix version of Booth recoding. Imagine A as a base 4 number: When the digit 3 appears, change it to $\overline{1}$ and add 1 to the next higher digit to compensate. The name for this technique, *overlapping triplets*, comes from the fact that it looks at 3 bits to determine what multiple of b to use, whereas ordinary Booth recoding looks at 2 bits.

The precise rules for overlapping triplets are given in Figure A.21. Besides having more complex control logic, this technique also requires that the P register be one bit wider to accommodate the possibility of $2b$ or $-2b$ being added to it. It is also possible to use a radix-8 (or even higher) version of Booth recoding. In that case, however, it will be necessary to use the multiple 3B as a potential summand. Radix-8 multipliers normally compute 3B once and for all at the beginning of a multiplication operation.

Current pair		Previous	Multiple
$i+1$	i	$i-1$	
0	0	0	0
0	0	1	$+b$
0	1	0	$+b$
0	1	1	$+2b$
1	0	0	$-2b$
1	0	1	$-b$
1	1	0	$-b$
1	1	1	0

FIGURE A.21 Multiples of b to use for radix-4 Booth recoding. For example, if the two low-order bits of the A register are both 1, and the last bit to be shifted out of the A register was 0, then the correct multiple is $-b$, obtained from the second to last row of the table.

Faster Multiplication with Many Adders

If the space for many adders is available, then multiplication speed can be improved. Figure A.22 shows a block diagram of a simple *array multiplier* for multiplying two 8-bit numbers using seven CSAs and one propagate adder. As it still takes eight additions to compute the product, the latency of computing a product is not dramatically different from using a single carry-save adder. However, with the hardware in Figure A.22, multiplication can be pipelined, increasing the total throughput. On the other hand, although this level of pipelining is sometimes used in array processors, it is not used in any of the single-chip, floating-point accelerators discussed in Section A.10. Pipelining is discussed in general in Chapter 6 and by Kogge [1981] in the context of multipliers.

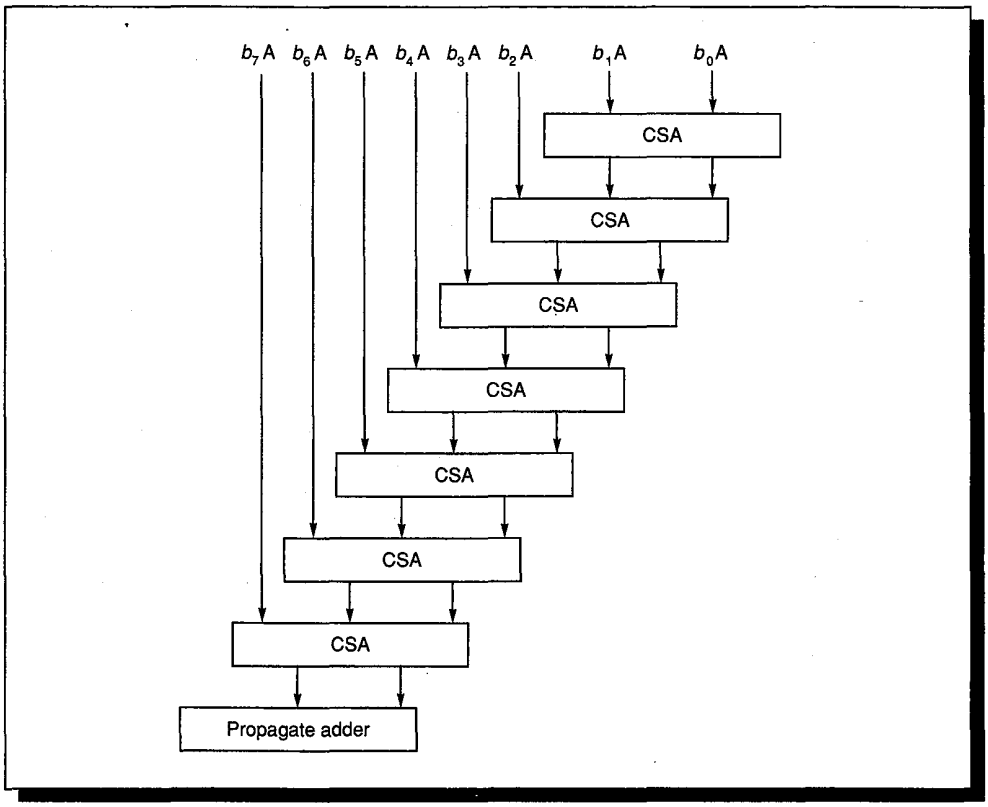

FIGURE A.22 Block diagram of an array multiplier. The 8-bit number in A is multiplied by $b_7b_6\cdots b_0$. Each box marked "CSA" is a carry-save adder.

With the technology of 1990, it is not possible to fit an array large enough to multiply two double-precision numbers on a single chip and have space left over for the other arithmetic operations. Thus, a popular design is to use a two-pass arrangement such as the one shown in Figure A.23 (page A-46). The first pass through the array "retires" four bits of B. Then the result of this first pass is fed back into the top to be combined with the next four summands. The result of this second pass is then fed into a CPA. This design, however, loses the ability to be pipelined.

If arrays require as many addition steps as the much cheaper arrangement in Figure A.2, why are they so popular? First of all, using an array has a smaller latency than using a single adder—because the array is a combinational circuit, the signals flow through it directly without being clocked. Although the two-pass adder of Figure A.23 would normally still use a clock, the cycle time for passing through k arrays can be less than k times the clock that would be needed for a design like the one in Figure A.2. Secondly, the array is amenable to various schemes for further speedup. One of them is shown in Figure A.24 (page A-47). The idea of this design is that two adds proceed in parallel or, to put it another way, each stream passes through only half the adders. Thus, it runs at almost twice the speed of the multiplier in Figure A.22. This *even/odd* multiplier

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 673

is popular in VLSI because of its regular structure. Arrays can also be speeded up using asynchronous logic. One of the reasons why the multiplier of Figure A.2 (page A-4) needs a clock is to keep the output of the adder from feeding back into the input of the adder before the output has fully stabilized. Thus, if the array in Figure A.23 is long enough so that no signal can propagate from the top through the bottom in the time it takes for the first adder to stabilize, it may be possible to avoid clocks altogether. Williams et al. [1987] discusses a design using this idea, although it is for dividers instead of for multipliers.

The techniques of the previous paragraph still have a multiply time of $O(n)$, but the time can be reduced to $\log n$ using a tree. The simplest tree would combine pairs of summands $b_0A \cdots b_{n-1}A$, cutting the number of summands from n to $n/2$. Then these $n/2$ numbers would be added in pairs again, reducing to $n/4$, and so on, and resulting in a single sum after $\log n$ steps. However, this simple binary-tree idea doesn't map into full (3,2) adders, which reduce three inputs to two rather than reducing two inputs to one. A tree that does use full adders, known as a *Wallace tree*, is shown in Figure A.25. When computer arithmetic units were built out of MSI parts, a Wallace tree was the design of choice for high-speed multipliers. There is, however, a problem with implementing them in VLSI.

Figures A.22–A.24 are sufficiently concise that it may be hard to visualize all the adders involved in an array multiplier. Figure A.26 (page A-49) shows each individual adder in a 4-bit array multiplier. Figure A.26(b) shows the inputs to the circuit, and Figure A.26(c) shows how those inputs are connected by adders.

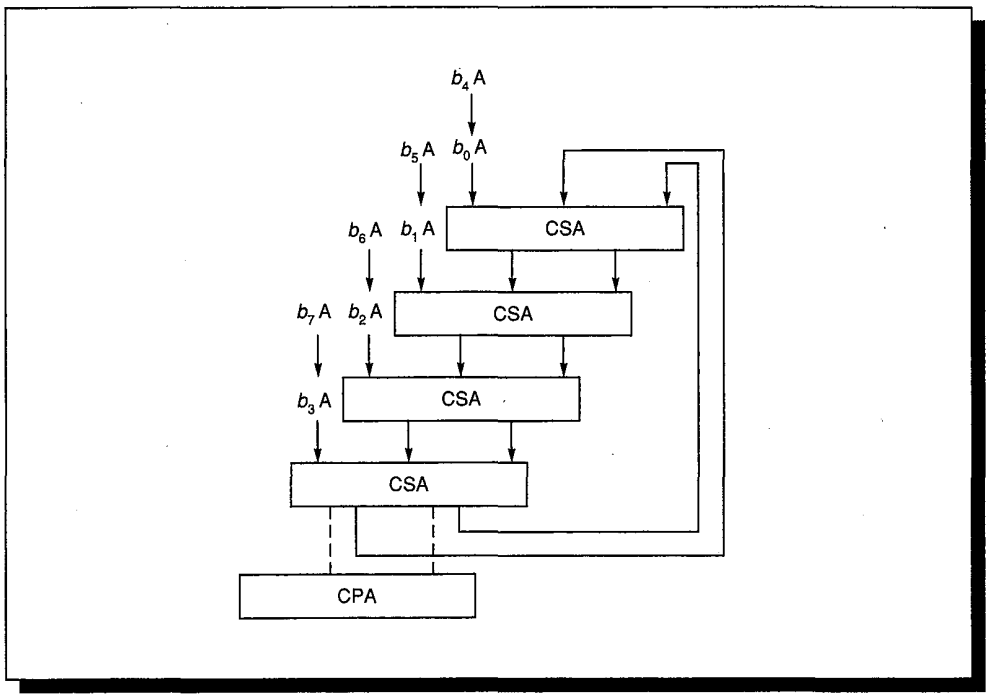

FIGURE A.23 Multipass array multiplier. Multiplies two 8-bit numbers with about half the hardware of that in Figure A.22. At the end of the second pass, the bits flow into the CPA.

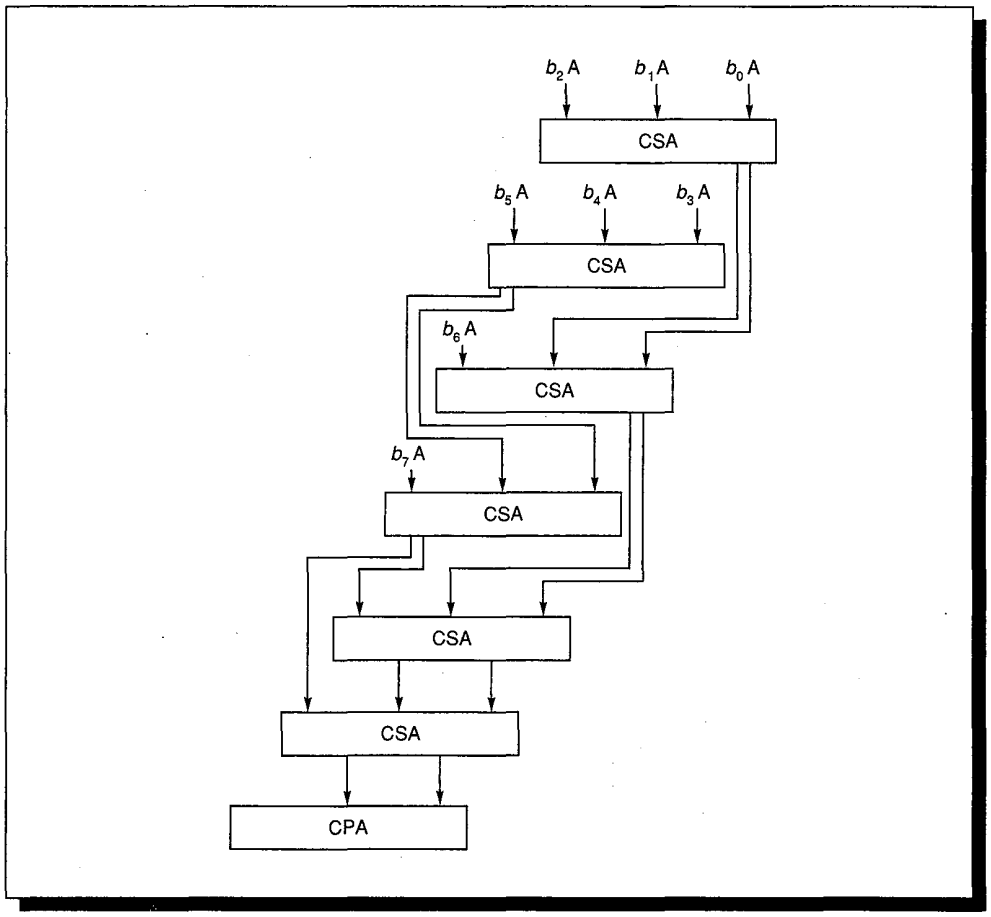

FIGURE A.24 Even/odd array. The first two adders work in parallel. Their results are fed into the third and fourth adders, which also work in parallel, and so on.

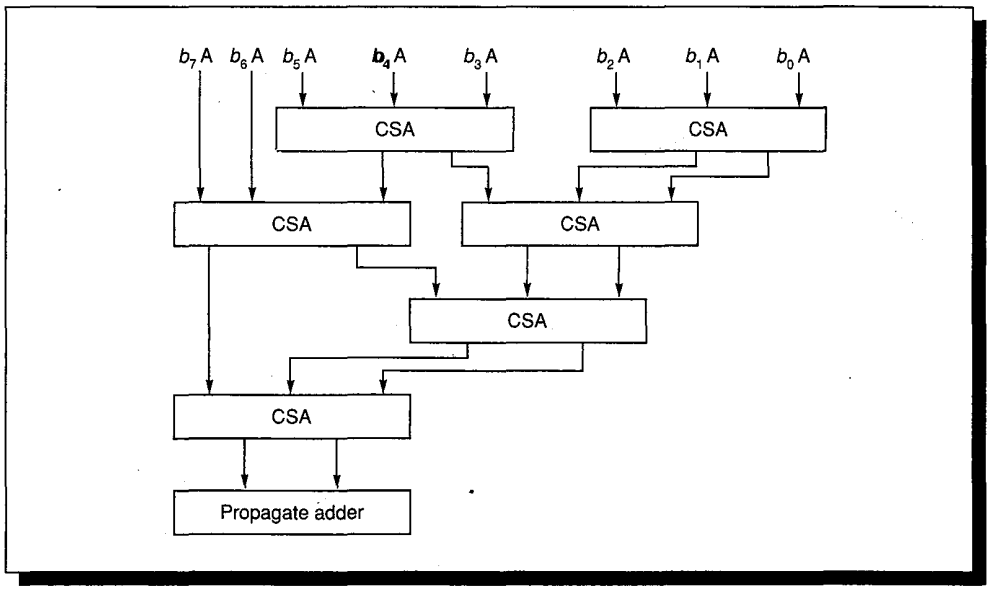

FIGURE A.25 Wallace-tree multiplier.

Each row of adders in A.26(c) corresponds to a single box in A.26(a). In actual implementation the array would be laid out as a square, not "twisted" as shown in the picture. (Lining up bits of the same significance in the same column makes the picture easier to understand.) If you try to fill in all the adders and paths for the Wallace tree of Figure A.25 (page A-47), you will discover that it does not have the nice, regular structure of Figure A.26. This is why VLSI designers have often chosen to use other $\log n$ designs such as the *binary-tree multiplier*, which is discussed next.

The problem with adding summands in a binary tree is that of coming up with a (2,1) adder that combines two digits and produces a single-sum digit. Because of carries, this isn't possible using binary notation, but it can be done with some other representation. We will use the *signed-digit representation* 1, $\bar{1}$, and 0, which we used previously to understand Booth's algorithm. This representation has two costs. First, it takes two bits to represent each signed digit. Second, the algorithm for adding two signed-digit numbers a_i and b_i is complex and requires examining $a_i a_{i-1} a_{i-2}$ and $b_i b_{i-1} b_{i-2}$. Although this means you must look two bits back, in binary addition you might have to look an arbitrary number of bits back (because of carries).

We can describe the algorithm for adding two signed-digit numbers as follows. First, compute sum and carry bits s_i and c_{i+1} using the table in Figure A.27. Then compute the final sum as $s_i + c_i$. The tables are set up so that this final sum does not generate a carry.

Example: What is the sum of the signed-digit numbers $1\bar{1}0$ and 001 ?

Answer: The two low-order bits sum to $0 + 1 = 1\bar{1}$, the next pair sums to $\bar{1} + 0 = 0\bar{1}$, and the high-order pair sums to $1 + 0 = 01$, so the sum is $1\bar{1} + 0\bar{1}0 + 0100 = 10\bar{1}$.

This, then, defines a (2,1) adder. With this in hand, we can use a straightforward binary tree to perform multiplication. In the first step it adds b_0A + b_1A in parallel with b_2A + b_3A, ···, b_{n-2}A + b_{n-1}A. The next step adds the results of these sums in pairs, and so on. Although the final sum must be run through a carry-propagate adder to convert it from signed-digit form to two's complement, this final add step is necessary in any multiplier using CSAs.

To summarize, both Wallace trees and signed-digit trees are $\log n$ multipliers. The Wallace tree uses the fewer gates but is harder to lay out. The signed-digit tree has a more regular structure, but requires two bits to represent each digit and has more complicated add logic. As with adders, it is possible to combine different multiply techniques. For example, Booth recoding and arrays can be combined. In Figure A.22 (page A-45) instead of having each input be b_iA, we could have it be $b_i b_{i-1}$A, and in order to avoid having to compute the multiple $3b$, we can use Booth recoding.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 676

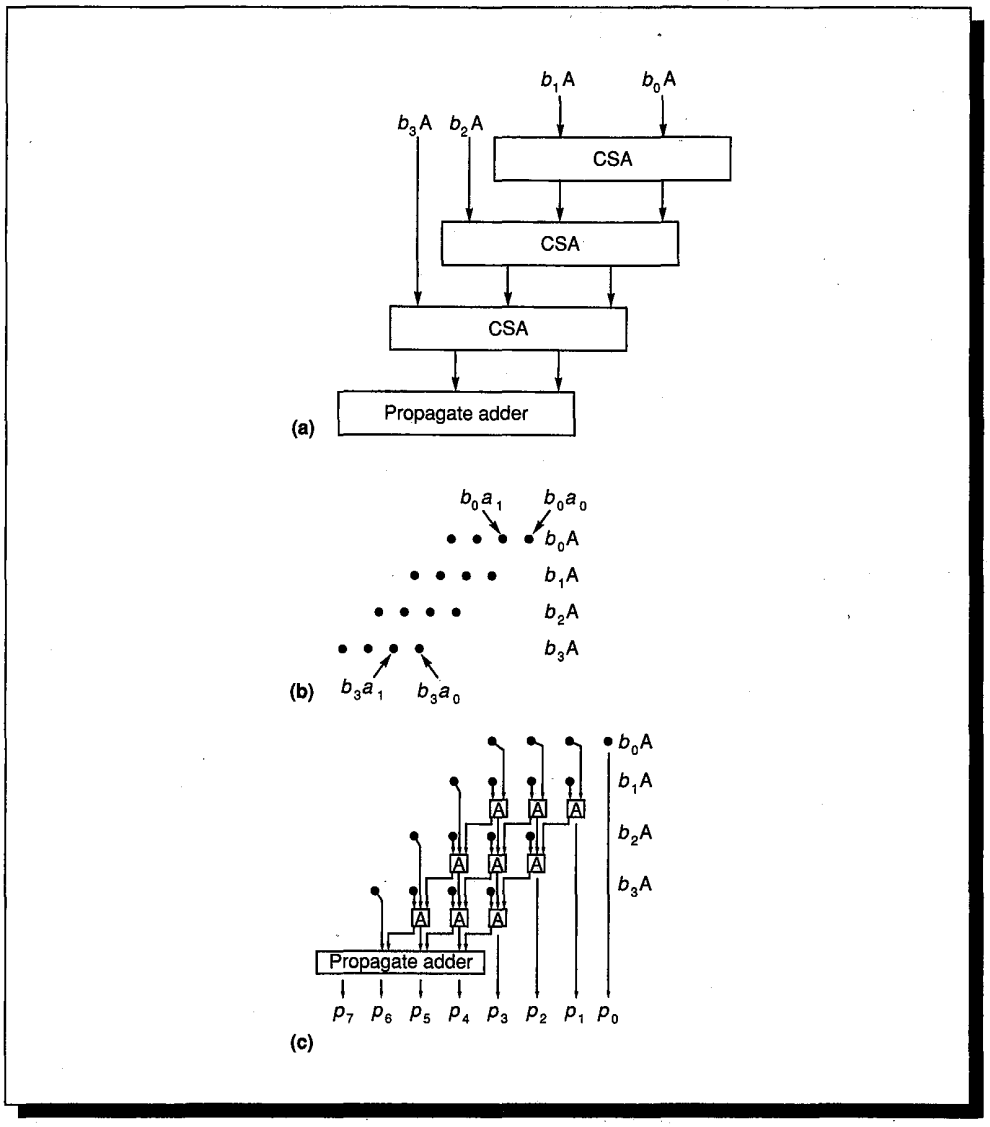

FIGURE A.26 Block diagram of an array multiplier (a); the inputs to the array (b); the array expanded to show all the adders (c).

1	1	$\bar{1}$	0	1 x		$\bar{1}$ x	
$+1$	$+\bar{1}$	$+\bar{1}$	$+0$	$+0$ y		$+0$ y	
1 0	0 0	$\bar{1}$ 0	0 0	1 $\bar{1}$	if $x \geq 0$ and $y \geq 0$	0 $\bar{1}$	if $x \geq 0$ and $y \geq 0$
				0 1	otherwise	$\bar{1}$ 1	otherwise

FIGURE A.27 Signed-digit addition table. The leftmost sum shows that when computing 1 + 1, the sum bit is 0 and the carry bit is 1.

Faster Division with One Adder

The two techniques for speeding up multiplication with a single adder were carry-save adders and higher-radix multiplication. There is a difficulty when trying to utilize these approaches to speed up nonrestoring division. The problem with CSAs is that at the end of each cycle the value of P, since it is in carry-save form, is not known exactly. In particular, the sign of P is uncertain, yet it is the sign of P that is used to compute the quotient digit and decide on the next ALU operation. When a higher radix is used, the problem is deciding what value to subtract from P. In the paper-and-pencil method, you have to guess the quotient digit. In binary division there are only two possibilities; we were able to finesse the problem by initially guessing one and then adjusting the guess based on the sign of P. This doesn't work in higher radices because there are more than two possible quotient digits, rendering quotient selection potentially quite complicated: You would have to compute all the multiples of b and compare them to P.

Both the carry-save technique and higher-radix division can be made to work if we use a redundant quotient representation. Recall from our discussion of SRT division that by allowing the quotient digits to be −1, 0, or 1, there is often a choice of which one to pick. The idea in the previous algorithm was to choose zero whenever possible because that meant an ALU operation could be skipped. In carry-save division, the idea is that because the remainder (P register) is not known exactly (being stored in carry-save form), the exact quotient digit is also not known. But thanks to the redundant representation, the remainder doesn't have to be known precisely in order to pick a quotient digit. This is illustrated in Figure A.28, where the x axis represents r_i, the contents of the (P,A) register pair after i steps. The line labeled $q_i = 1$ shows the value that r_{i+1} would be if we choose $q_i = 1$, and similarly for the lines $q_i = 0$ and $q_i = -1$. We can choose any value for q_i, as long as $r_{i+1} = rP_i - q_iB$ satisfies $|r_{i+1}| \leq B$. The allowable ranges are shown in the right half of Figure A.28. Thus we only need to know r precisely enough to decide in which range in Figure A.28 it lies.

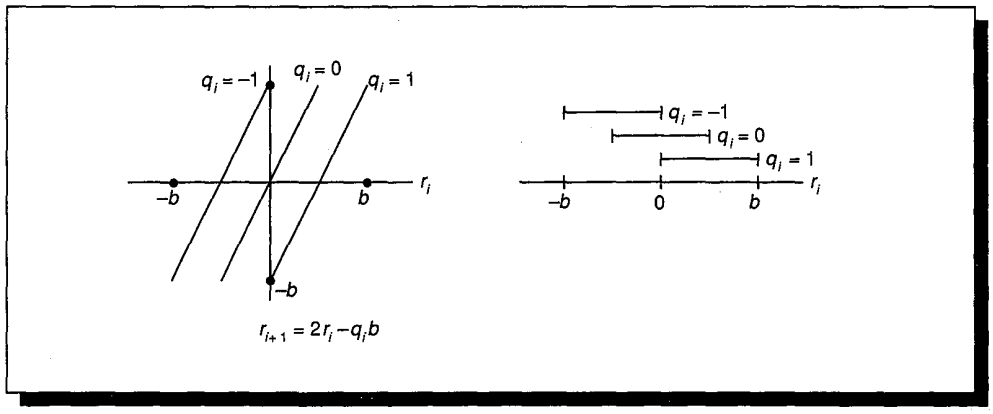

FIGURE A.28 Quotient selection for radix-2 division. The x axis represents the i th remainder, which is the quantity in the (P,A) register pair. The y axis shows the value of the remainder after one additional divide step. Each bar on the right-hand graph gives the range of r_i values for which it is permissible to select the associated value of q_i.

This is the basis for using carry-save adders. Look at the high-order bits of the carry-save adder and sum them in a propagate adder. Then use this approximation of r to compute q_i, usually by means of a lookup table. The same technique works for higher-radix division (whether or not a carry-save adder is used). The high-order bits P can be used to index a table that gives one of the allowable quotient digits.

The design challenge when building a high-speed SRT divider is figuring out how many bits of P and B need to be examined. For example, suppose that we take a radix of 4, use quotient digits of 2, 1, 0, $\overline{1}$, $\overline{2}$, but have a propagate adder. How many bits of P and B need to be examined? Deciding this involves two steps. For ordinary radix-2 nonrestoring division, because at each stage $|r| \le b$, the P buffer won't overflow. But for radix 4, $r_{i+1} = 4r_i - q_i b$ is computed at each stage, and if r_i is near b, then $4r_i$ will be near $4b$, and even the largest quotient digit will not bring r back to the range $|r_{i+1}| \le b$. In other words, the remainder might grow without bound. However, restricting $|r_i| \le 2b/3$ makes it easy to check that r_i will stay bounded.

After figuring out the bound that r_i must satisfy, we can draw the diagram in Figure A.29, which is analogous to Figure A.28. If r_i is between $(1/12)b$ and $(5/12)b$, we can pick $q = 1$, and so on. Or to put it another way, if r/b is between 1/12 and 5/12, we can pick $q = 1$. Suppose we look at 4 bits of P and 4 bits of b, and the high bits of P (not counting the $(n + 1)$-st sign bit) are $0011xxx\cdots$, while the high bits of b are $1001xxx\cdots$. To simplify calculation, imagine the binary point at the left end of each register. Since we truncated, r (the value of P concatenated with A) could have a value from .0011 to .0100, and b could have a value from .1001 to .1010. Thus r/b could be as small as .0011/.1010 or as large as .0100/.1001. But $.0011_2/.1010_2 = 3/10 < 1/3$ would require a quotient bit of 1, while $.0100_2/.1001_2 = 4/9 > 5/12$ would require a quotient bit of 2. In other words, 4 bits of P and 4 bits of b aren't enough to pick a quotient bit. It turns out that 5 bits of P and 4 bits of b are enough. This can be verified by writing a simple program that checks all the cases.

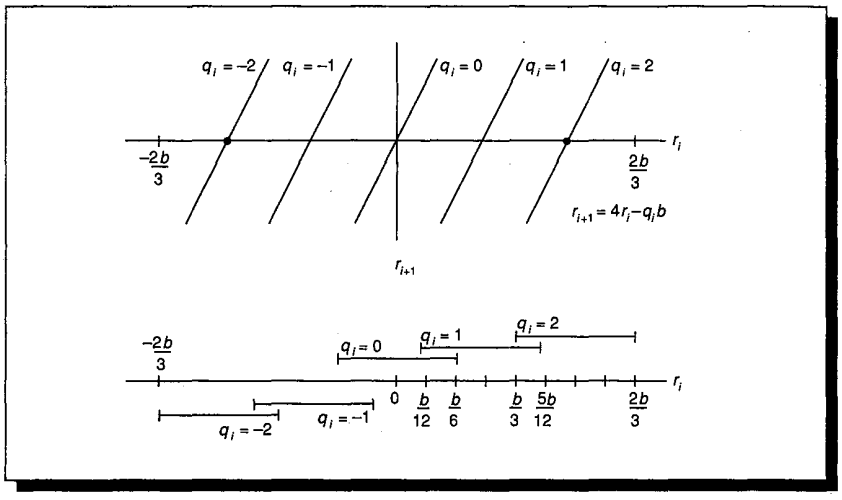

FIGURE A.29 Quotient selection for radix-4 division.

Example: Suppose that the radix is 4 and the quotient digits are 2, 1, 0, $\overline{1}$, $\overline{2}$, but this time a CSA is used instead of a propagate adder. How many bits of the P and B registers need to be examined?

Answer: Once again $|r_i| \leq 2b/3$, and the ranges of the q_i are still as in Figure A.29. If the top 4 bits of the sum part and the carry part of P are respectively 0010 and 0001, then the sum part ranges from 0010 to 0011 and the carry part from 0001 to 0010. Accordingly, the true value of r ranges from 0010 + 0001 = 0011 to 0011 + 0010 = 0101. Given, therefore, a CPA that adds the top 4 bits of the carry and sum parts of P, and a sum of 0011, the true sum will be anywhere from 0011 to 0101. A program that checks all the cases will show that 6 bits of P and 4 bits of b are needed to predict a quotient digit. The result of such a program is shown in Figure A.30. For example, if b is 1001$xxx\cdots$ and r is 001101$xxx\cdots$, then the top 4 bits of b are 9 and the top 6 bits of r are 13, making the quotient digit 1. But if r were $001110_2 = 14$, the quotient digit would have to be 2.

b	Range of P		q	b	Range of P		q
8	−21	−14	−2	12	−32	−20	−2
8	−13	−5	−1	12	−20	−7	−1
8	−5	3	0	12	−8	6	0
8	3	11	1	12	5	18	1
8	12	21	2	12	18	32	2
9	−24	−16	−2	13	−34	−21	−2
9	−15	−6	−1	13	−21	−7	−1
9	−6	4	0	13	−8	6	0
9	4	13	1	13	5	19	1
9	14	24	2	13	19	34	2
10	−26	−17	−2	14	−37	−22	−2
10	−16	−6	−1	14	−23	−7	−1
10	−6	4	0	14	−9	7	0
10	4	14	1	14	5	21	2
10	15	26	2	14	20	37	2
11	−29	−18	−2	15	−40	−24	−2
11	−18	−6	−1	15	−25	−8	−1
11	−7	5	0	15	−10	8	0
11	4	16	1	15	6	23	1
11	16	29	2	15	22	40	2

FIGURE A.30 Quotient digits for radix-4 SRT division with a CSA. The top row says that if the high-order 4 bits of b are $1000_2 = 8$, and if the top 6 bits of P are between $110010_2 = -14$ and $10101_2 = -21$, then the quotient digit is −2.

Although these are simple cases, all SRT analyses proceed in the same way. First compute the range of r_i, then plot r_i against r_{i+1} to find the quotient ranges, and finally write a program to compute how many bits are necessary. (It is sometimes also possible to compute the required number of bits analytically.) Two final comments about high-radix SRT division are in order. First, Figure A.30 is not symmetrical. Thus, for a radix-4 CSA divider, the lookup table needs not only 6 bits of P, but also the sign of P. Second, the quotient lookup table has a fairly regular structure. This means it is usually cheaper to encode it as a PLA rather than in ROM.

A.10 Putting It All Together

In this section, we will compare the Weitek 3364, the MIPS R3010, and the Texas Instruments 8847 (see Figures A.31 and A.32, pages A-54–A-55). In many ways, these are ideal chips to compare. They each implement the IEEE standard for addition, subtraction, multiplication, and division on a single chip. All were introduced in 1988 and run with a cycle time of about 40 nanoseconds. However, as we will see, they use quite different algorithms. The Weitek chip is well described in Birman et al. [1988], the MIPS chip is described in less detail in Rowen, Johnson, and Ries [1988], and the details of the TI chip have yet to be published.

There are a number of things that these three chips have in common. They perform addition and multiplication in parallel, and they implement neither extended precision nor the IEEE remainder operation. We discussed earlier how an efficient REM could be provided in software if only chips would implement a remainder-step function. The designers of these chips probably decided not to

	MIPS R3010	Weitek 3364	TI 8847
Clock cycle time (ns)	40	50	30
Size (mil^2)	114,857	147,600	156,180
Transistors	75,000	165,000	180,000
Pins	84	168	207
Power (watts)	3.5	1.5	1.5
Cycles/add	2	2	2
Cycles/mult	5	2	3
Cycles/divide	19	17	11
Cycles/sq root	–	30	14

FIGURE A.31 Summary of the three floating-point chips discussed in this section. The cycle times are for production parts available in June 1989. The cycle counts are for double-precision operations.

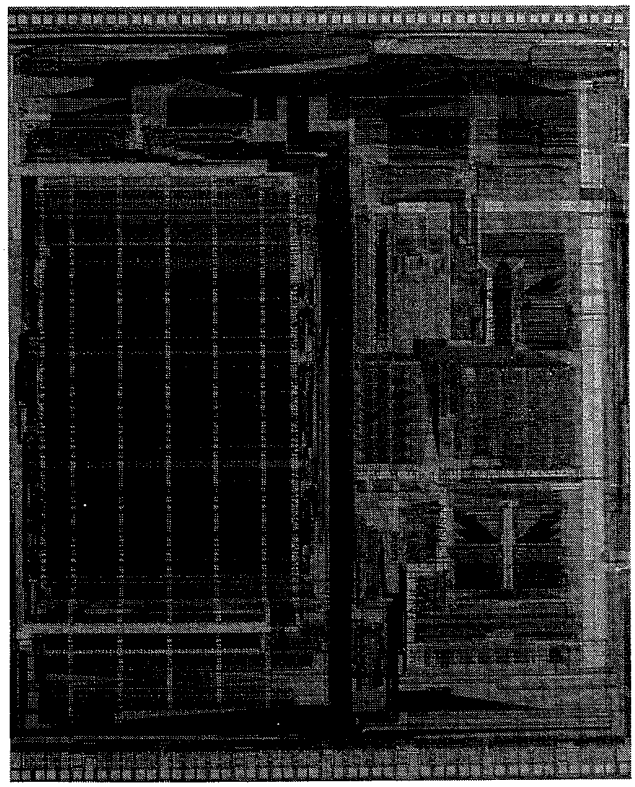

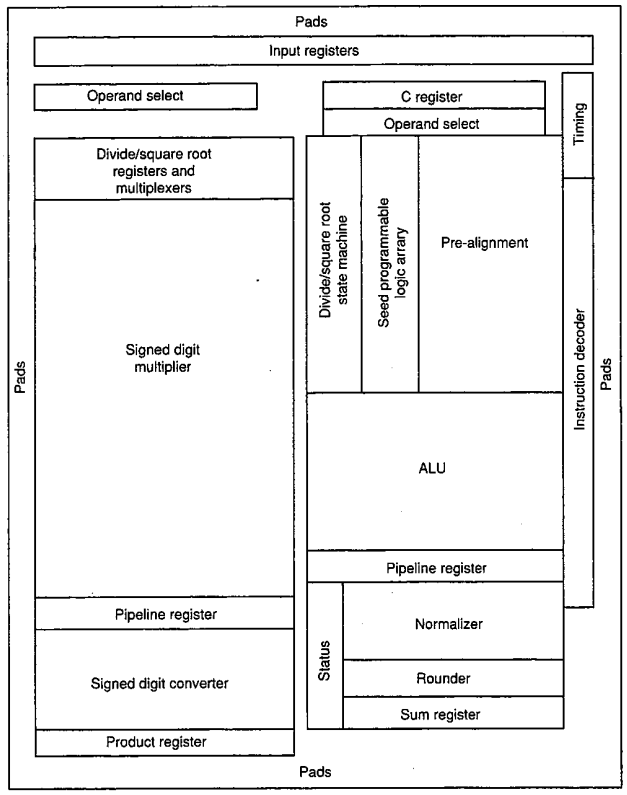
Pads
Input registers
Operand select
C register
Operand select
Timing
Divide/square root registers and multiplexers
Divide/square root state machine
Seed programmable logic arrary
Pre-alignment
Signed digit multiplier
Instruction decoder
Pads
Pads
ALU
Pipeline register
Pipeline register
Normalizer
Status
Signed digit converter
Rounder
Sum register
Product register
Pads

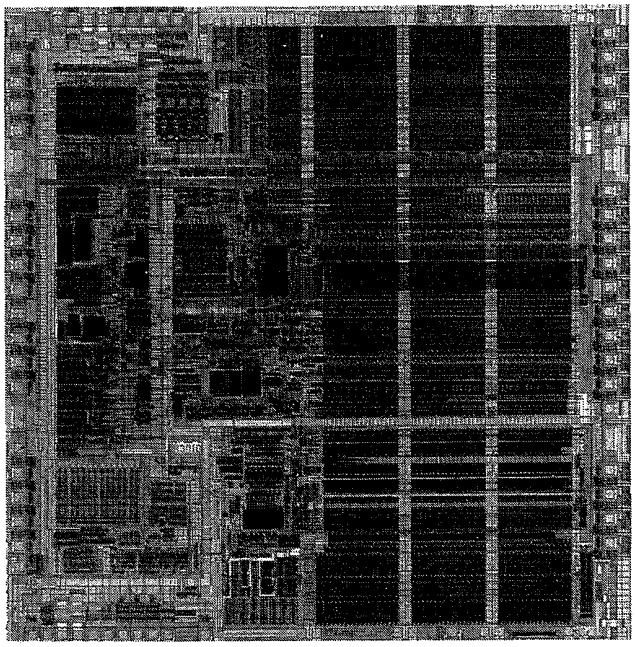

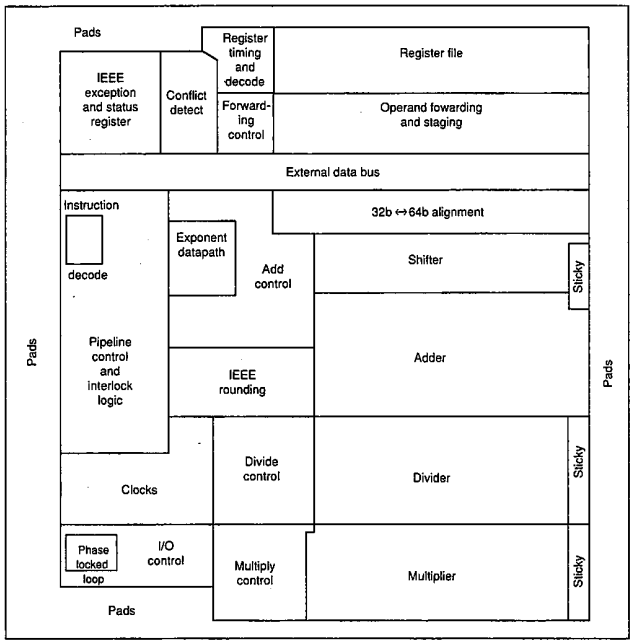
Pads
Register timing and decode
Register file
IEEE exception and status register
Conflict detect
Forward-ing control
Operand fowarding and staging
External data bus
Instruction
32b ↔ 64b alignment
Exponent datapath
Add control
Shifter
Sticky
decode
Pads
Pipeline control and interlock logic
IEEE rounding
Adder
Pads
Divide control
Divider
Sticky
Clocks
Phase locked loop
I/O control
Multiply control
Multiplier
Sticky
Pads

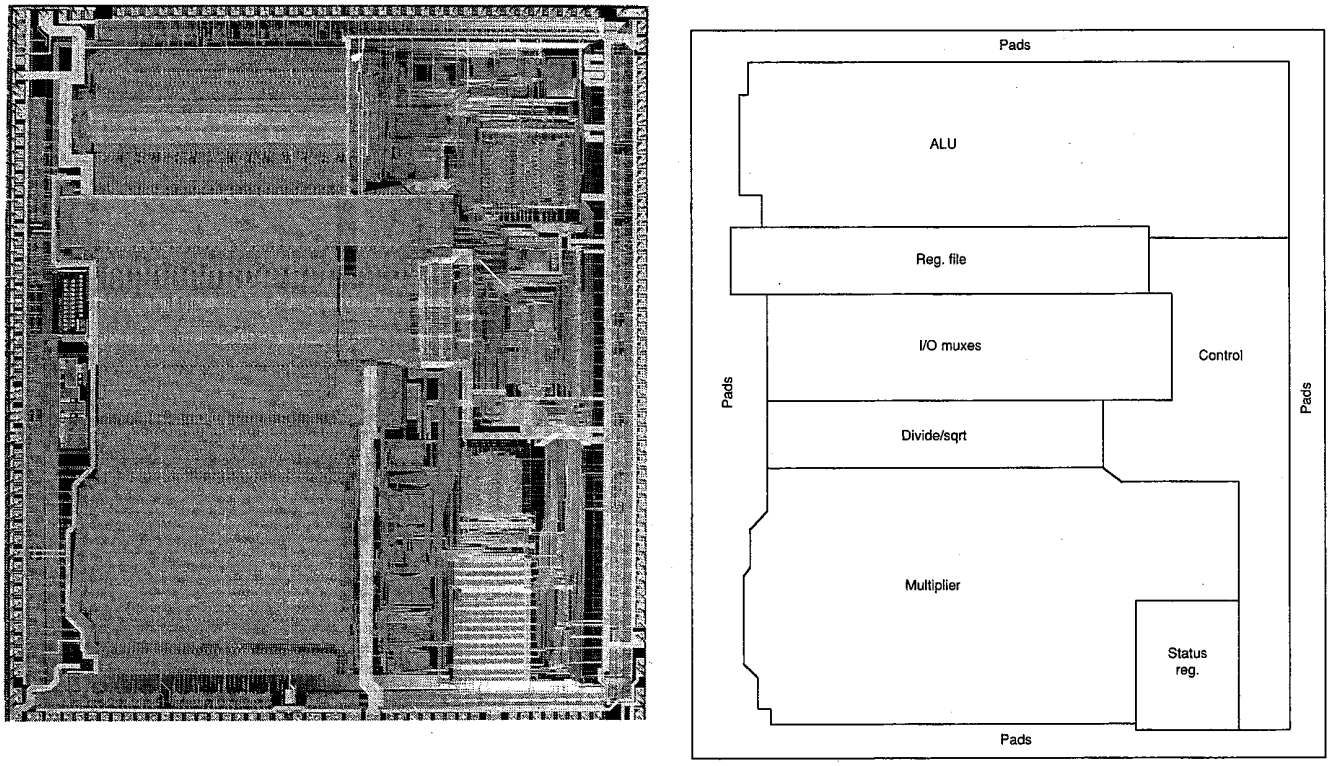

FIGURE A.32 Chip layout. In the left-hand column are the photomicrographs; the right-hand column shows the corresponding floor plans. Top left is the TI 8847, bottom left is the MIPS R3010, and above is the Weitek 3364.

provide extended precision because the most influential users are those who run portable codes, which can't rely on extended precision. However, as we have seen, extended precision can make for faster and simpler math libraries.

A summary of the three chips is given in Figures A.31 (page A-53) and A.32. Note that a higher transistor count generally leads to smaller cycle counts. Comparing the cycles/op numbers needs to be done carefully because the figures for the MIPS chip are those for a complete system (R3000/3010 pair), while the Weitek and TI numbers are for standalone chips, and are usually larger when used in a complete system.

The MIPS chip has the fewest transistors of the three. This is reflected in the fact that it is the only chip of the three that does not have any pipelining or hardware square root. Further, the multiplication and addition operations are not completely independent because they share the carry-propagate adder that performs the final rounding (as well as the rounding logic). Addition on the R3010 uses a mixture of ripple, CLA, and carry select. A carry-select adder is used in the fashion of Figure A.16 (page A-38). Within each half, carries are propagated using a hybrid ripple-CLA scheme of the type indicated in Figure A.14. However, this is further tuned by varying the size of each block, rather than having each fixed at four bits (as they are in Figure A.14 on page A-36). The multiplier is midway between the designs of Figures A.2 (page A-4) and A.22 (page A-45). It has an array just large enough so that output can be fed back into the input without having to be clocked. Also, it uses radix-4 Booth recoding and the even-odd technique of Figure A.24 (page A-47). The R3010 can do a divide and multiply in parallel (like the Weitek chip but unlike the TI chip). The divider is a radix-4 SRT method with quotient digits −2, −1, 0, 1, and 2, and is similar to that described in Taylor [1985]. Double-precision division is about four times slower than multiplication. The R3010 shows that for chips using an $O(n)$ multiplier, an SRT divider can operate fast enough to keep a reasonable ratio between multiply and divide.

The Weitek 3364 has independent add, multiply, and divide units, and also uses radix-4 SRT division. However, the add and multiply operations on the Weitek chip are pipelined. The three addition stages are (1) exponent compare, (2) add followed by shift (or vice versa), and (3) final rounding. Stages (1) and (3) take only a half-cycle, allowing the whole operation to be done in two cycles, even though there are three pipline stages. The multiplier uses an array of the style of Figure A.23 but uses radix-8 Booth recoding, which means it must compute 3 times the multiplier. The three multiplier pipeline stages are (1) compute $3b$, (2) pass through array, and (3) final carry-propagation add and round. Single precision passes through the array once, double precision twice. Like addition, the latency is two cycles. The Weitek chip uses an interesting addition algorithm. It is a variant on the carry-skip adder pictured in Figure A.15 (page A-37). However P_{ij}, which is the logical AND of many terms, is computed by rippling, performing one AND per ripple. Thus, while the carries propagate left within a block, the value of P_{ij} is propagating right within the next block, and the block sizes are chosen so that both waves complete at the same time. Unlike

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 684

the MIPS chip, the 3364 has hardware square root, which shares the divide hardware. The ratio of double-precision multiply to divide is 2:17. The large disparity between multiply and divide is due to the fact that multiplication uses radix-8 Booth recoding, while division uses a radix-4 method. In the MIPS R3010, multiplication and division use the same radix.

The notable feature of the TI 8847 is that it does division by iteration (using the Goldschmidt algorithm discussed in Section A.6). This improves the speed of division (the ratio of multiply to divide is 3:11), but means that multiplication and division cannot be done in parallel as on the other two chips. Addition has a two-stage pipeline. Exponent compare, fraction shift, and fraction addition are done in the first stage, normalization and rounding in the second stage. Multiplication uses a binary tree of signed-digit adders and has a three-stage pipeline. The first stage passes through the array retiring half the bits, the second stage passes through the array a second time, and the third stage converts from signed-digit form to two's complement. Since there is only one array, a new multiply operation can only be initiated in every other cycle. However, by slowing down the clock, two passes through the array can be made in a single cycle. In this case, a new multiplication can be initiated in each cycle. The 8847 adder uses a carry-select algorithm rather than carry lookahead. As mentioned in Section A.6, the TI carries 60 bits of precision in order to do correctly rounded division.

These three chips illustrate the different tradeoffs made by designers with similar constraints. One of the most interesting things about these chips is the diversity of their algorithms. Each uses a different add algorithm, as well as a different multiply algorithm. In fact, Booth recoding is the only technique that is universally used by all the chips.

A.11 Fallacies and Pitfalls

Fallacy: Underflows rarely occur in actual floating-point application code.

Although most codes rarely underflow, there are actual codes that underflow frequently. SDRWAVE [Kahaner 1988], which solves a one-dimensional wave equation, is one such example. This program underflows quite frequently, even when functioning properly. Measurements on one machine show that adding hardware support for gradual underflow would cause SDRWAVE to run about 50% faster.

Fallacy: Conversions between integer and floating point are rare.

In fact, in Spice they are as frequent as divides. The assumption that conversions are rare leads to a mistake in the SPARC instruction set, which does not provide an instruction to move from integer registers to floating-point registers.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 685

Pitfall: Don't increase the speed of a floating-point unit without increasing its memory bandwidth.

A typical use of a floating-point unit is to add two vectors to produce a third vector. If these vectors consist of double-precision numbers, then each floating-point add will use three operands of 64 bits each, or 24 bytes of memory. The memory bandwidth requirements are even greater if the floating-point unit can perform addition and multiplication in parallel (as most do).

Pitfall: $-x$ is not the same as $0 - x$.

This is a fine point in the IEEE standard that has tripped up some designers. Because floating-point numbers use the sign/magnitude system, there are two zeros, +0 and –0. The standard says that $0 - 0 = +0$, whereas $-(0) = -0$. Thus $-x$ is not the same as $0 - x$ when $x = 0$.

A.12 Historical Perspective and References

The earliest computers used fixed point rather than floating point. In "Preliminary Discussion of the Logical Design of an Electronic Computing Instrument," Burks, Goldstine, and von Neumann put it like this:

There appear to be two major purposes in a "floating" decimal point system both of which arise from the fact that the number of digits in a word is a constant fixed by design considerations for each particular machine. The first of these purposes is to retain in a sum or product as many significant digits as possible and the second of these is to free the human operator from the burden of estimating and inserting into a problem "scale factors" — multiplicative constants which serve to keep numbers within the limits of the machine.

There is, of course, no denying the fact that human time is consumed in arranging for the introduction of suitable scale factors. We only argue that the time so consumed is a very small percentage of the total time we will spend in preparing an interesting problem for our machine. The first advantage of the floating point is, we feel, somewhat illusory. In order to have such a floating point, one must waste memory capacity which could otherwise be used for carrying more digits per word. It would therefore seem to us not at all clear whether the modest advantages of a floating binary point offset the loss of memory capacity and the increased complexity of the arithmetic and control circuits. [Bell and Newell 1971, 97]

This enables us to see things from the perspective of early computer designers, who believed that saving computer time and memory were more important than saving programmer time.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 686

The original papers introducing the Wallace tree, Booth recoding, SRT division, overlapped triplets, and so on, are reprinted in Swartzlander [1980]. A good explanation of an early machine (the IBM 360/91) that used a pipelined Wallace tree, Booth recoding, and iterative division is in Anderson et al. [1967]. A discussion of the average time for single-bit SRT division is in Freiman [1961]; this is one of the few interesting historical papers that does not appear in Swartzlander.

The standard book of Mead and Conway [1980] discouraged the use of CLAs as not being cost effective in VLSI. Brent and Kung [1982] was an important paper that helped combat that view. An example of a detailed layout for CLAs can be found in Ngai and Irwin [1985] or in Weste and Eshraghian [1985]. Takagi, Yasuura, and Yajima [1985] provides a detailed description of a signed-digit–tree multiplier.

Although the IEEE standard is being widely adopted, there are still three other important floating-point systems in use: the IBM/370, the DEC VAX, and the Cray. We will briefly discuss these older formats. The VAX format is closest to the IEEE standard. Its single-precision format (F format) is like IEEE single precision in that it has a hidden bit, 8 bits of exponent, and 23 bits of fraction. However, it does not have a sticky bit, which causes it to round halfway cases up instead of to even. The VAX has a slightly different exponent range than IEEE single: $E_{\min}$ is -128 rather than -126 as in IEEE, and $E_{\max}$ is 126 instead of 127. The main differences between VAX and IEEE are the lack of special values and gradual underflow. The VAX has a reserved operand, but it works like a signaling NaN: it traps whenever it is referenced. Originally, the VAX's double precision (D format) also had 8 bits of exponent. However, as this is too small for many applications, a G format was added; like the IEEE standard, this format has 11 bits of exponent. The VAX also has an H format, which is 128 bits long.

The IBM/370 floating-point format uses base 16 rather than base 2. This means it cannot use a hidden bit. In single precision, it has 7 bits of exponent and 24 bits (6 hex digits) of fraction. Thus, the largest representable number is $16^{2^7} = 2^{4\times 2^7} = 2^{2^9}$, compared with 2^{2^8} for IEEE. However, a number that is normalized in the hexadecimal sense only needs to have a nonzero leading digit. When interpreted in binary, the three most significant bits could be zero. Thus, there are potentially fewer than 24 bits of significance. The reason for using the higher base was to minimize the amount of shifting required when adding floating-point numbers. However, this is less significant in current machines, where the floating-point add time is usually fixed independent of the operands. Another difference between 370 arithmetic and IEEE arithmetic is that the 370 has neither a round digit nor a sticky digit, which effectively means that it truncates rather than rounds. Thus, in many computations, the result will systematically be too small. Unlike the VAX and IEEE arithmetic, every bit pattern is a valid number. Thus, library routines must establish conventions for what to return in case of errors. In the IBM FORTRAN library, for example, $\sqrt{-4}$ returns 2!

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 687

Arithmetic on Cray computers is interesting because it is driven by a motivation for the highest possible floating-point performance. It has a 15-bit exponent field and a 48-bit fraction field. Addition on Cray computers does not have a guard digit, and multiplication is even less accurate than addition. Thinking of multiplication as a sum of p numbers, each $2p$ bits long, what Cray computers do is to drop the low-order bits of each summand. Thus, analyzing the exact error characteristics of the multiply operation is not easy. Reciprocals are computed using iteration, and division of a by b is done by multiplying a times $1/b$. The errors in multiplication and reciprocation combine to make the last three bits of a divide operation unreliable. At least Cray computers serve to keep numerical analysts on their toes!

The IEEE standardization process began in 1977, inspired mainly by W. Kahan, and is based partly on Kahan's work with the IBM 7094 at the University of Toronto [Kahan 1968]. The standardization process was a lengthy affair, with gradual underflow causing the most controversy. (According to Cleve Moler, visitors to the U.S. were advised that the sights not to be missed were Las Vegas, the Grand Canyon, and the IEEE standards committee meeting.) The standard was finally approved in 1985. The Intel 8087 was the first major commercial IEEE implementation and appeared in 1981, before the standard was finalized. It contains features that were eliminated in the final standard, such as projective bits. According to Kahan, the length of double-extended precision was based on what could be implemented in the 8087. Although the IEEE standard was not based on any existing floating-point system, most of its features were present in some other system. For example the CDC 6600 reserved special bit patterns for INDEFINITE and INFINITY, while the idea of denormal numbers appears in Goldberg [1967] as well as in Kahan [1968]. Kahan was awarded the 1989 Turing prize in recognition of his work on floating point.

References

ANDERSON, S. F., J. G. EARLE, R. E. GOLDSCHMIDT, AND D. M. POWERS [1967]. "The IBM System/360 Model 91: Floating-point execution unit," *IBM J. Research and Development* 11, 34–53. Reprinted in [Swartzlander 1980].

Good description of an early high-performance floating-point unit that used a pipelined Wallace-tree multiplier and iterative division.

ATKINS, D. E. [1968]. "Higher-radix division using estimates of the divisor and partial remainders," *IEEE Trans. on Computers* C-17:10, 925–934. Reprinted in [Swartzlander 1980].

This is the standard reference for high-radix SRT division.

BELL, C. G. AND A. NEWELL, [1971]. *Computer Structures: Readings and Examples,* McGraw-Hill, New York.

BIRMAN, M., G. CHU, L. HU, J. MCLEOD, N. BEDARD, F. WARE, L. TORBAN, AND C. M. LIM [1988]. "Design of a high-speed arithmetic datapath," *Proc. ICCD: VLSI Computers and Processors,* 214–216.

Fairly detailed description of the Weitek 3364 floating-point chip.

BRENT, R. P. AND H. T. KUNG [1982] "A regular layout for parallel adders," *IEEE Trans. on Computers* C-31, 260–264.

This is the paper that popularized CLA adders in VLSI.

BURKS, A. W., H. H. GOLDSTINE, AND J. VON NEUMANN, [1946]. *Preliminary Discussion of the Logical Design of an Electronic Computing Instrument.*

CODY, W. J. [1988]. "Floating point standards: Theory and practice," in *Reliability in Computing: The Role of Interval Methods in Scientific Computing,* R. E. Moore, (ed.), Academic Press, Boston, Mass., 99–107.

Presents a status of hardware and software implementations of the standard.

CODY, W. J., J. T. COONEN, D. M. GAY, K. HANSON, D. HOUGH, W. KAHAN, R. KARPINSKI, J. PALMER, F. N. RIS, AND D. STEVENSON [1984]. "A proposed radix- and word-length-independent standard for floating-point arithmetic," *IEEE Micro* 4:4, 86–100.

Contains a draft of the 854 standard, which is more general than 754. The significance of this article is that it contains commentary on the standard, most of which is equally relevant to 754.

COONEN, J. [1984]. *Contributions to a Proposed Standard for Binary Floating-Point Arithmetic,* Ph.D. Thesis, Univ. of Calif., Berkeley.

The only detailed discussion of how rounding modes can be used to implement efficient binary decimal conversion.

FREIMAN, C. V. [1961]. "Statistical analysis of certain binary division algorithms," *Proc. IRE* 49:1, 91–103.

Contains an analysis of the performance of shifting-over-zeros SRT division algorithm.

GOLDBERG, D. [1989]. "Floating-point and computer systems," *Xerox Tech. Rep.* CSL-89-9. A version of this paper will appear in *Computing Surveys.*

Contains an in-depth tutorial on the IEEE standard from the software point of view.

GOLDBERG, I. B. [1967]. "27 bits are not enough for 8-digit accuracy," *Comm. ACM* 10:2, 105–106.

This paper proposes using hidden bits and gradual underflow.

GOSLING, J. B. [1980]. *Design of Arithmetic Units for Digital Computers,* Springer-Verlag NewYork, Inc., New York.

A concise, well-written book, although it focuses on MSI designs.

HAMACHER, V. C., Z. G. VRANESIC, AND S. G. ZAKY [1984]. *Computer Organization,* 2nd ed., McGraw-Hill, New York.

Introductory computer architecture book with a good chapter on computer arithmetic.

HWANG, K. [1979]. *Computer Arithmetic: Principles, Architecture, and Design,* Wiley, New York.

This book contains the widest range of topics of the computer arithmetic books.

IEEE [1985]. "IEEE standard for binary floating-point arithmetic," *SIGPLAN Notices* 22:2, 9–25.

IEEE 754 is reprinted here.

KAHAN, W. [1968]. "7094-II system support for numerical analysis," *SHARE Secretarial Distribution* SSD-159.

This system had many features that were incorporated into the IEEE floating-point standard.

KAHANER, D. K. [1988]. "Benchmarks for 'real' programs," *SIAM News* (November).

The benchmark presented in this article turns out to cause many underflows.

KNUTH, D. [1981]. *The Art of Computer Programming,* vol II, 2nd ed., Addison-Wesley, Reading, Mass.

Has a section on the distribution of floating-point numbers.

KOGGE, P. [1981]. *The Architecture of Pipelined Computers,* McGraw-Hill, New York.

Has brief discussion of pipelined multipliers.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 689

KOHN, L. AND S.-W. FU, [1989]. "A 1,000,000 transistor microprocessor," *IEEE Int'l Solid-State Circuits Conf.*, 54–55.

A brief overview of the Intel 860, whose floating-point addition algorithm is discussed in Section A.4.

MAGENHEIMER, D. J., L. PETERS, K. W. PETTIS, AND D. ZURAS, [1988]. "Integer multiplication and division on the HP Precision Architecture," *IEEE Trans. on Computers* 37:8, 980–990.

Rationale for the integer- and divide-step instructions in the Precision architecture.

MEAD, C. AND L. CONWAY [1980]. *Introduction to VLSI Systems*, Addison-Wesley, Reading, Mass.

NGAI, T-F. AND M. J. IRWIN [1985]. "Regular, area-time efficient carry-lookahead adders," *Proc. Seventh IEEE Symposium on Computer Arithmetic*, 9–15.

Describes a CLA adder like that of Figure A.13, where the bits flow up and then come back down.

PENG, V., S. SAMUDRALA, AND M. GAVRIELOV [1987]. "On the implementation of shifters, multipliers, and dividers in VLSI floating point units," *Proc. Eighth IEEE Symposium on Computer Arithmetic*, 95–102.

Highly recommended survey of different techniques actually used in VLSI designs.

ROWEN, C., M. JOHNSON, and P. RIES [1988]. "The MIPS R3010 floating-point coprocessor," *IEEE Micro* 53–62 (June).

SANTORO, M. R., G. BEWICK, and M. A. HOROWITZ [1989]. "Rounding algorithms for IEEE multipliers," *Proc. Ninth IEEE Symposium on Computer Arithmetic*, 176–183.

A very readable discussion of how to efficiently implement rounding for floating-point multiplication.

SCOTT, N. R. [1985]. *Computer Number Systems and Arithmetic*, Prentice-Hall, Englewood Cliffs, N.J.

SWARTZLANDER, E., ED. [1980]. *Computer Arithmetic*, Dowden, Hutchison and Ross (distributed by Van Nostrand, New York).

A collection of historical papers.

TAKAGI, N., H. YASUURA, AND S. YAJIMA [1985]."High-speed VLSI multiplication algorithm with a redundant binary addition tree," *IEEE Trans. on Computers* C-34:9, 789–796.

A discussion of the binary-tree signed multiplier that was the basis for the design used in the TI 8847.

TAYLOR, G. S. [1981]. "Compatible hardware for division and square root," *Proc. Fifth IEEE Symposium on Computer Arithmetic*, 127–134.

Good discussion of a radix-4 SRT division algorithm.

TAYLOR, G. S. [1985]. "Radix 16 SRT dividers with overlapped quotient selection stages," *Proc. Seventh IEEE Symposium on Computer Arithmetic*, 64–71.

Describes a very sophisticated high-radix division algorithm.

WESTE, N. AND K. ESHRAGHIAN [1985]. *Principles of CMOS VLSI Design*, Addison-Wesley, Reading, Mass.

This textbook has a section on the layouts of various kinds of adders.

WILLIAMS, T. E., M. HOROWITZ, R. L. ALVERSON, AND T. S. YANG [1987]. "A self-timed chip for division," *Advanced Research in VLSI, Proc. 1987 Stanford Conf.*, The MIT Press, Cambridge, Mass.

Describes a divider that tries to get the speed of a combinational design without using the area that would be required by one.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 690

EXERCISES

A.1 [15/15/20] <A.3> Represent the following numbers as single-precision and double-precision IEEE floating-point numbers.

a. [15] 10

b. [15] 10.5

c. [20] 0.1

A.2 [10/15/20] <A.8> Complete the details of the block diagrams for the following adders.

a. [10] In Figure A.11, show how to implement the "1" and "2" boxes in terms of AND and OR gates.

b. [15] In Figure A.14, what signals need to flow from the adder cells in the top row into the "C" cells? Write the logic equations for the "C" box.

c. [20] Show how to extend the block diagram in A.13 so it will produce the carry-out bit c_8.

A.3 [15/15] <A.4> Floating-point addition.

a. [15] In a decimal system with $p = 5$, compute $-4.5673 + 4.9999 \times 10^{-5}$ assuming round to nearest. Give the value of the guard and round digits, and the sticky bit.

b. [15] What is the value of the sum for the other three rounding modes?

A.4 [15] <A.3> Show that if gradual underflow is not used, then it is no longer true that $x \neq y$ if and only if $x - y \neq 0$.

A.5 [25] <A.9> Write out the analogue of Figure A.21 for radix-8 Booth recoding.

A.6 [15] <A.3> Is the ordering of nonnegative floating-point numbers the same as integers when denormalized numbers are also considered? What if the denormalized numbers are represented using the wrapped representation mentioned in Section A.5?

A.7 [25/10] <A.2> One's complement.

a. [25] When adding two's complement numbers, you discard the carry out from the most significant bit. Show that in one's complement, you must add the carry back into the low end.

b. [10] Find the rule for detecting overflow in one's complement.

A.8 [15] <A.2> Equations A.2.1 and A.2.2 are for adding two n-bit numbers. Derive similar equations for subtraction, where there will be a borrow instead of a carry.

A.9 [15/20] <A.2> More one's complement.

a. [15] A complication that arises with one's complement arithmetic is that zero has two representations. Show that even if the negative form of zero is never an input, the adder in Equation A.2.1 (with c_0 the end around carry) can still produce a negative zero.

b. [20] Use the fact that $a + b = a - (-b)$ together with the subtractor circuit of the previous problem to derive a different one's complement adder. Can this adder ever produce negative zero?

A.10 [20] <A.2> On a machine that doesn't detect integer overflow in hardware, show how you would detect overflow on a signed addition operation in software.

A.11 [25] <A.9> In the array of Figure A.23, the fact that an array can be pipelined is not exploited. Can you come up with a design that feeds the output of the bottom CSA into the bottom CSAs instead of the top one, and that will run faster than the arrangement of Figure A.23?

A.12 [15] <A.9> For ordinary Booth recoding, the multiple of b used in the ith step is simply $a_{i-1} - a_i$. Can you find a similar formula for radix-4 Booth recoding (overlapped triplets)?

A.13 [25/15/30] <A.9> Shifting-over-zeros multiplication.

a. [25] Does Booth recoding always increase the number of zeros in a number? Can it ever decrease the number of zeros?

b. [15] Given the number $a_{n-1}\cdots a_0$, define $c_0 = 0$, and define c_i to be the carry out from adding a_i, a_{i-1}, and c_{i-1}. Then *modified Booth recoding* gives a number with digits $A_i = a_i + c_i - 2c_{i+1}$. What is the recoding of 01101?

c. [30] Show that modified Booth recoding never decreases the number of zeros.

A.14 [20/15/20/15/20/15] <A.6> Iterative square root.

a. [20] Use Newton's method to derive an iterative algorithm for square root. The formula will involve a division.

b. [15] What is the fastest way you can think of to divide a floating-point number by 2?

c. [20] If division is slow, then the iterative square root routine will also be slow. Use Newton's method on $f(x) = 1/x^2 - a$ to derive a method that doesn't use any divisions.

d. [15] Assume that the ratio division by 2 : floating-point add : floating-point multiply is 1:2:4. What ratios of multiplication time to divide time makes each iteration step in the method of Part c faster than each iteration in the method of Part a?

e. [20] When using the method of Part a, how many bits need to be in the initial guess in order to get double-precision accuracy after 3 iterations? (You may ignore rounding error.)

f. [15] Suppose that when Spice runs on the TI 8847, it spends 16.7% of its time in the square root routine (this percentage has been measured on other machines). Using the

values in Figure A.31 and assuming 3 iterations, how much slower would Spice run if square root was implemented in software using the method of Part a?

A.15 [30/10] <A.2> This problem presents an algorithm for adding signed-magnitude numbers. If A and B are integers of opposite signs, let a and b be their magnitudes.

a. [30] Show that the following rules for manipulating the unsigned numbers a and b gives $A + B$

1. Complement one of the operands.
2. Using end around carry (as in the one's complement adder of problem A.7) add the complemented operand and the other (uncomplemented) one.
3. If there was a carry out, the sign of the result is the sign associated with the uncomplemented operand.
4. Otherwise, if there was no carry out, complement the result, and give it the sign of the complemented operand.

b. [10] <A.4> In our discussion of floating-point add, we suggested that when the result is negative the +1 needed to do two's complement be done in the rounding unit. Use the result of Part A to devise a floating-point adder that doesn't require this.

A.16 [15] <A.7> Our example that showed that double rounding can give a different answer from rounding once used the round-to-even rule. If halfway cases are rounded up, is double rounding still dangerous?

A.17 [15/30] <A.9> The text discussed radix-4 SRT division with quotient digits of –2, –1, 0, 1, 2. Suppose that 3 and –3 are also allowed as quotient digits.

a. [15] What relation replaces $|r_i| \le 2b/3$?

b. [30] How many bits of b and P do you need to examine ?

A.18 [25] <A.6,A.9> The discussion of the remainder-step instruction assumed that division was done using a bit-at-a-time algorithm. What would have to change if division was implemented using a higher-radix method?

A.19 [20/20/25/25/20] <A.3> Signed-logarithm representation.

a. [20] Suppose you want to represent a number x by its sign and $\log|x|$. Then if $\log|x|$ is to be nonnegative, x must be ≥ 1. You can allow smaller x if you represent x by $\log k|x|$ for some constant k. Use 0 if $k|x| < 1$. Now $\log k|x|$ will not be an integer, but it can be represented as a fixed-point number. If we put the binary point m bits to the left of the least significant bit, write down formulas for converting x to signed-logarithm form and back.

b. [20] Give the rules for multiplication and division.

c. [25] Show that no matter what base of logs is used, this system cannot exactly represent all of 1, 2, and 3.

d. [25] Show how to implement addition using a table containing 2^{p-1} entries of $p-1$ bits each, where the signed logarithm number is stored in a p-bit register.

e. [20] Show that for numbers which are exactly representable in this system, multiplication is exact, addition is not, but $a(b+c) = ab + ac$ exactly (when there is no over/underflow).

A.20 [20/10] <A.8> Carry-skip adders.

a. [20] Assuming that time is proportional to logic levels, what (fixed) block size gives the fastest addition for an adder of some fixed total length?

b. [10] Explain why the carry-skip adder takes time $\sqrt{n}$.

A.21 [Discussion] In the MIPS approach to exception handling, you need a test for determining whether two floating-point operands could cause an exception. This should be fast and also not have too many false positives. Can you come up with a practical test? The performance cost of your design will depend on the distribution of floating-point numbers. This is discussed in Knuth [1981] and Swartzlander [1980].

A.22 [35] <A.8> The simplest carry-select adder replaces an n-bit adder with $n/2$ bit adders and a mux. A more complex carry-select adder would use $n/4$-bit adders and more muxes. Can you design an adder that uses muxes and 1-bit adders and runs in O(log n) time? Such an adder is called a *conditional-sum adder*.

A.23 [10/15/20/15/15] <A.6> Correctly rounded iterative division. Let a and b be floating-point numbers with p-bit significands ($p = 53$ in double precision). Let q be the exact quotient $q = a/b$. Suppose that $\bar{q}$ is the result of an iteration process, that $\bar{q}$ has a few extra bits of precision, and that $0 < q - \bar{q} < 2^{-p}$.

a. [10] If x is a floating-point number, and $1 \le x < 2$, what is the next representable number after x?

b. [15] Show how to compute q' from $\bar{q}$, where q' has $p + 1$ bits of precision and $|q - q'| < 2^{-p}$.

c. [20] Assuming round to nearest, show that the correctly rounded quotient is either q', $q' - 2^{-p}$, or $q' + 2^{-p}$.

d. [15] Give rules for computing the correctly rounded quotient from q' based on the low- order bit of q' and the sign of $a - bq'$.

e. [15] Solve Part c for the other three rounding modes.

B Complete Instruction Set Tables

B.1 VAX User Instruction Set	**B-2**
B.2 System/360 Instruction Set	**B-6**
B.3 8086 Instruction Set	**B-9**

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 695

B.1 VAX User Instruction Set

The following tables include all the VAX user instructions; the system instructions are not included.

The underscore following the instruction name implies that the instruction will operate upon any data type contained in the parentheses following that instruction. The data type abbreviations are:

B = byte (8 bits)
W = word (16 bits)
L = longword (32 bits)
Q = quadword (64 bits)
O = octaword (128 bits)
F = F_floating (32 bits)
D = D_floating (64 bits)
G = G_floating (64 bits)
H = H_floating (128 bits)

Integer and Floating-Point Logical and Arithmetic Instructions

Instruction	Description
ADAWI	Add aligned word interlocked
ADD_2	Add (B,W,L,F,D,G,H) 2 operand
ADD_3	Add (B,W,L,F,D,G,H) 3 operand
ADWC	Add with carry
ASH_	Arithmetic shift (L,Q)
BIC_2	Bit clear (B,W,L) 2 operand
BIC_3	Bit clear (B,W,L) 3 operand
BICPSW	Bit clear processor status word
BIS_2	Bit set (B,W,L) 2 operand
BIS_3	Bit set (B,W,L) 3 operand
BISPSW	Bit set processor status word
BIT_	Bit test (B,W,L)
CLR_	Clear (B,W,L=F,Q=D=G,O=H)
CVT_	Convert (B,W,L,F,D,G,H)(B,W,L,F,D,G,H) except BB, WW, LL, FF, DD, GG, HH, DG, and GD
CVTR_L	Convert rounded (F,D,G,H) to longword
CMP_	Compare (B,W,L,F,D,G,H)
DEC_	Decrement (B,W,L)
DIV_2	Divide (B,W,L,F,D,G,H) 2 operand
DIV_3	Divide (B,W,L,F,D,G,H) 3 operand
EDIV	Extended divide
EMOD_	Extended modulus (F,D,G,H)
EMUL	Extended multiply

Instruction	Description
INC_	Increment (B,W,L)
INDEX	Compute index
MCOM_	Move complemented (B,W,L)
MNEG_	Move negated (B,W,L,F,D,G,H)
MOVA_	Move address (B,W,L=F,Q=D=G,O=H)
MOV_*	Move (B,W,L,F,D,G,H,Q,O)**—general move between two operands
MOVPSL	Move from processor status longword
MOVZ_	Move zero-extended (BW,BL,WL)
MUL_2	Multiply (B,W,L,F,D,G,H) 2 operand
MUL_3	Multiply (B,W,L,F,D,G,H) 3 operand
POLY_	Polynomial evaluation (F,D,G,H)
POPR	Pop registers from stack
PUSHA_	Push address (B,W,L=F,Q=D=G,O=H) on stack
PUSHL	Push longword on stack
PUSHR	Push registers on stack
ROTL	Rotate longword
SBWC	Subtract with carry
SUB_2	Subtract (B,W,L,F,D,G,H) 2 operand
SUB_3	Subtract (B,W,L,F,D,G,H) 3 operand
TST_	Test (B,W,L,F,D,G,H)
XOR_2	Exclusive or (B,W,L) 2 operand
XOR_3	Exclusive or (B,W,L) 3 operand

Branch, Jump, and Procedure Call Instructions

Instruction	Description
ACB_	Add, compare and branch (B,W,L.F,D,G,H)
AOBLEQ	Add one and branch less than or equal
AOBLSS	Add one and branch less than
BB_	Branch on bit (set, clear)
BBS_	Branch on bit (set, clear) and (set, clear) bit
BB_I	Branch on bit set (clear) and set (clear) bit interlocked
BCC	Branch carry cleared
BCS	Branch carry set
BEQL	Branch equal
BEQLU	Branch equal unsigned
BGEQ	Branch greater than or equal
BGEQU	Branch greater than or equal unsigned
BGTR	Branch greater than

Instruction	Description
BGTRU	Branch greater than unsigned
BLB_	Branch on low bit (set, clear)
BLEQ	Branch less than or equal
BLEQU	Branch less than or equal unsigned
BLSS	Btranch less than
BLSSU	Branch less than unsigned
BNEQ	Branch not equal
BNEQU	Branch not equal unsigned
BR_	Jump with (B,W) displacement
BSB_	Branch to subroutine with (B,W) displacement
BV_	Branch overflow (set,clear)
CALLG	Call procedure with general argument list
CALLS	Call procedure with stack argument list
CASE_	Case on (B,W,L)
JMP	Jump
JSB	Jump to subroutine
RET	Return from procedure
RSB	Return from subroutine
SOBGEQ	Subtract one and branch greater than or equal
SOBGTR	Subtract one and branch greater than

Decimal and String Instructions

Instruction	Description
ADDP4	Add packed 4 operand
ADDP6	Add packed 6 operand
ASHP	Arithmetic shift packed and round
CMPC3	Compare characters 3 operand
CMPC5	Compare characters 5 operand
CMPP3	Compare packed 3 operand
CMPP4	Compare packed 4 operand
CRC	Calculate cyclic redundancy check
CVTLP	Convert long to packed
CVTPL	Convert packed to long
CVTPT	Convert packed to trailing
CVTTP	Convert trailing to packed
CVTPS	Convert packed to separate
CVTSP	Convert separate to packed
DIVP	Divide packed
EDITPC	Edit packed to character string

Instruction	Description
LOCC	Locate character
MATCHC	Match characters
MOVC3	Move character 3 operand
MOVC5	Move character 5 operand
MOVP	Move packed
MOVTC	Move translated characters
MOVTUC	Move translated until character
MULP	Multiply packed
SCANC	Scan characters
SKPC	Skip character
SPANC	Span characters
SUBP4	Subtract packed 4 operand
SUBP6	Subtract packed 6 operand

Variable-Length Bit Field Instructions

Instruction	Description
CMPV	Compare field
CMPZV	Compare zero-extended field
EXTV	Extract field
EXTZV	Extract zero-extended field
INSV	Insert field
FFS	Find first set
FFC	Find first clear

Queue Instructions

Instruction	Description
INSQHI	Insert entry into queue at head, interlocked
INSQTI	Insert entry into queue at tail, interlocked
INSQUE	Insert entry in queue
REMQHI	Remove entry from queue at head, interlocked
REMQTI	Remove entry from queue at tail, interlocked
REMQUE	Remove entry from queue

B.2 System/360 Instruction Set

The 360 instruction set is shown in the following tables, organized by instruction type and format. System/370 contains 15 additional user instructions.

Integer/Logical and Floating-Point R–R Instructions

The * indicates the instruction is floating point, and may be either D (double precision) or E (single precision).

Instruction	Description
ALR	Add logical register
AR	Add register
A*R	FP addition
CLR	Compare logical register
CR	Compare register
C*R	FP compare
DR	Divide register
D*R	FP divide
H*R	FP halve
LCR	Load complement register
LC*R	Load complement
LNR	Load negative register
LN*R	Load negative
LPR	Load positive register
LP*R	Load positive
LR	Load register
L*R	Load FP register
LTR	Load and test register
LT*R	Load and test FP register
MR	Multiply register
M*R	FP multiply
NR	And register
OR	Or register
SLR	Subtract logical register
SR	Subtract register
S*R	FP subtraction
XR	Exclusive or register

Branches and Status Setting R–R Instructions

These are R–R format instructions that either branch or set some system status; several of them are privileged and legal only in supervisor mode.

Instruction	Description
BALR	Branch and link
BCTR	Branch on count
BCR	Branch/condition
ISK	Insert key
SPM	Set program mask
SSK	Set storage key
SVC	Supervisor call

Integer/Logical and Floating-Point Instructions— RX Format

These are all RX format instructions. The symbol "+" means either a word operation (and then stands for nothing) or H (meaning halfword); for example, A+ stands for the two opcodes A and AH. The symbol "*" is D or E standing for double- or single-precision floating point.

Instruction	Description
A+	Add
A*	FP add
AL	Add logical
C+	Compare
C*	FP compare
CL	Compare logical
D	Divide
D*	FP divide
L+	Load
L*	Load FP register
M+	Multiply
M*	FP multiply
N	And
O	Or
S+	Subtract
S*	FP subtract
SL	Subtract logical
ST+	Store
ST*	Store FP register
X	Exclusive or

Branches and Special Loads and Stores—RX format

Instruction	Description
BAL	Branch and link
BC	Branch condition
BCT	Branch on count
CVB	Convert-binary
CVD	Convert-decimal
EX	Execute
IC	Insert character
LA	Load address
STC	Store character

RS and SI Format Instructions

These are the RS and SI format instructions. The symbol "*" may be A (arithmetic) or L (logical).

Instruction	Description
BXH	Branch/high
BXLE	Branch/low-equal
CLI	Compare logical immediate
HIO	Halt I/O
LPSW	Load PSW
LM	Load multiple
MVI	Move immediate
NI	And immediate
OI	Or immediate
RDD	Read direct
SIO	Start I/O
SL*	Shift left A/L
SLD*	Shift left double A/L
SR*	Shift right A/L
SRD*	Shift right double A/L
SSM	Set system mask
STM	Store multiple
TCH	Test channel
TIO	Test I/O
TM	Test under mask
TS	Test and set
WRD	Write direct
XI	Exclusive or immediate

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 702

SS Format Instructions

These are all decimal or string instructions.

Instruction	Description
AP	Add packed
CLC	Compare logical chars
CP	Compare packed
DP	Divide packed
ED	Edit
EDMK	Edit and mark
MP	Multiply packed
MVC	Move character
MVN	Move numeric
MVO	Move with offset
MVZ	Move zone
NC	And characters
OC	Or characters
PACK	Pack (Character → decimal)
SP	Subtract packed
TR	Translate
TRT	Translate and test
UNPK	Unpack
XC	Exclusive or characters
ZAP	Zero and add packed

B.3 8086 Instruction Set

These charts contain the instruction set of the 8086; floating-point instructions that are neither included nor used by the 8086 benchmarks are not included.

Arithmetic and Logical Instructions

Instruction	Description
AAA	ASCII adjust after addition
AAD	ASCII adjust before division
AAM	ASCII adjust after multiplication
AAS	ASCII adjust after subtraction
ADC	Add with carry
ADD	Integer addition
AND	Logical and
CBW/CWD/CDQ	Convert byte to word/word to dword/dword to quad
CLC	Clear the carry flag
CLD	Clear the direction flag
CLI	Clear the interrupt flag
CMC	Complement the carry flag
CMP	Compare
DAA	Decimal adjust after addition
DAS	Decimal adjust after subtraction
DEC	Decrement
DIV	Unsigned divide
IDIV	Signed divide
IMUL	Signed multiplication
INC	Increment
MUL	Unsigned multiplication
NEG	Negate
NOT	Not
OR	Inclusive or
RCL	Rotate through carry left
RCR	Rotate through carry right
ROL	Rotate left
ROR	Rotate right
SAL/SHL	Shift arithmetic left
SAR	Shift arithmetic right
SBB	Subtract with borrow
SHR	Shift logical right
STC	Set carry flag
STD	Set direction flag
STI	Set interrupt flag
SUB	Subtract
TEST	Logical compare
XOR	Exclusive or

Control Instructions

Instruction	Description
CALL	Call procedure (intrasegment)
CALL	Call procedure (intersegment)
HLT	Halt
INT	Call to interrupt procedure
INTO	On overflow call interrupt procedure
IRET	Interrupt return
JB/JNAE/JC	Jump below
JBE/JNA	Jump below or equal
JCXZ/JECXZ	Jump CX/ECX zero
JE/JZ	Jump equal
JL/JNGE	Jump less
JLE/JNG	Jump less or equal
JMP	Jump (intrasegment)
JMPF	Jump (intersegment)
JNB/JAE/JNC	Jump not below
JNBE/JA	Jump not below or equal
JNE/JNZ	Jump not equal
JNL/JCE	Jump not less
JNLE/JG	Jump not less or equal
JNO	Jump no overflow
JNP/JPO	Jump not parity
JNS	Jump not sign
JO	Jump overflow
JP/JPE	Jump parity
JS	Jump sign
LOCK	Bus lock
RET	Return (intrasegment)
RETF	Return (intersegrnent)

Data Transfer Instructions

Instruction	Description
IN	Input from a port
LAHF	Load flags into AH register
LDS	Load pointer to DS
LEA	Load effective address
LES	Load pointer to ES
LOCK	Bus lock
MOV	Move
OUT	Output to a port
POP	Pop off stack
POPF/POPFD	Pop from stack into flags
PUSH	Push onto stack
PUSH	Push segment register onto the stack
PUSHF/PUSHFD	Push flags onto stack
SAHF	Store AH register into flags
XCHC	Exchange
XLAT/XLATB	Table lookup translation

String Instructions

Instruction	Description
CMPS/CMPSB/CMPSW/CMPSD	Compare string
LODS/LODSB/LODSW/LODSD	Load string
MOVS/MOVSB/MOVSW/MOVSD	Move string
REP	Repeat
REPE/REPZ	Repeat while equal
REPNE/REPNZ	Repeat while not equal
SCAS/SCASB/SCASW/SCASD	Scan string
STOS/STOSB/STOSW/STOSD	Store string

C Detailed Instruction Set Measurements

C.1 VAX Detailed Measurements C-2
C.2 360 Detailed Measurements C-3
C.3 Intel 8086 Detailed Measurements C-4
C.4 DLX Detailed Instruction Set Measurements C-5

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.

C.1 VAX Detailed Measurements

Instruction	GCC	Spice	TeX	COBOLX	Average
Control	**30%**	**18%**	**30%**	**25%**	**26%**
Conditional Branch	20%	13%	19%	18%	17%
BRB,BRW	6%	3%	4%	5%	5%
CALLS,CALLG	2%	1%	4%	0%	2%
RET	2%	1%	4%	0%	2%
JMP				2%	1%
Arithmetic, logical	**40%**	**23%**	**33%**	**24%**	**30%**
CMP*	12%	5%	11%	9%	9%
ADDL_	5%	12%	4%		5%
INCL	3%		3%	5%	3%
MOVA*	1%	3%	4%	2%	3%
TSTL	4%	2%	3%		2%
CLRL	3%	1%	2%	3%	2%
SUB*_	3%	1%	3%		2%
CVT*L	6%			0%	2%
ASHL	3%		3%	0%	2%
MULL_	0%			5%	1%
Data transfer	**19%**	**15%**	**28%**	**4%**	**16%**
MOVL	15%	9%	17%	4%	11%
PUSHL	3%		7%		2%
MOVQ		6%			1%
MOVZ*L	1%		4%		1%
Floating point	**0%**	**23%**	**0%**	**0%**	**6%**
MULD_		9%			2%
SUBD_		6%			1%
ADDD_		6%			1%
DIVD_		3%			1%
CMPD		2%			
Decimal, string	**0%**	**0%**	**1%**	**38%**	**10%**
CVTTP,CVTPT				19%	5%
MOVC3,MOVC5			1%	9%	2%
ADDP4				6%	1%
CMPP_				2%	1%
CMPC3				2%	1%
Totals	**88%**	**79%**	**92%**	**88%**	**87%**

FIGURE C.1 Instructions responsible for more than 1.5% of the dynamic executions in any benchmark. The instructions are broken into five classes, printed in boldface. The data in those rows give the total frequency for the operations in that class. Cells representing a contribution of 1% or less are empty, except the average column can have an entry of 1%. Because of rounding, the average can differ from what might appear to be correct if based on the figures in the individual columns.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 708

C.2 360 Detailed Measurements

Instruction	PLIC	FORTGO	PLIGO	COBOLGO	Average
Control	**32%**	**13%**	**5%**	**16%**	**16%**
BC,BCR	28%	13%	5%	14%	15%
BAL,BALR	3%			2%	1%
Arithmetic, logical	**29%**	**35%**	**29%**	**9%**	**26%**
A,AR	3%	17%	21%		10%
SR	3%	7%			3%
SLL		6%	3%		2%
LA	8%	1%	1%		2%
CLI	7%				2%
NI				7%	2%
C	5%	4%	4%	0%	3%
TM	3%	1%		3%	2%
MH			2%		1%
Data transfer	**17%**	**40%**	**56%**	**20%**	**33%**
L,LR	7%	23%	28%	19%	19%
MVI	2%		16%	1%	5%
ST	3%		7%		3%
LD		7%	2%		2%
STD		7%	2%		2%
LPDR		3%			1%
LH	3%				1%
IC	2%				1%
LTR		1%			0%
Floating point		**7%**			**2%**
AD		3%			1%
MDR		3%			1%
Decimal, string	**4%**			**40%**	**11%**
MVC	4%			7%	3%
AP				11%	3%
ZAP				9%	2%
CVD				5%	1%
MP				3%	1%
CLC				3%	1%
CP				2%	1%
ED				1%	0%
Total	**82%**	**95%**	**90%**	**85%**	**88%**

FIGURE C.2 (See previous page.) Distribution of instruction execution frequencies for the four 360 programs. All instructions with a frequency of execution greater than 1.5% are included. Immediate instructions, which operate on only a single byte, are included in the section that characterizes their operation, rather than with the long character-string versions of the same operation. By comparison, the average frequencies for the major instruction classes of the VAX are 23% (control), 28% (arithmetic), 29% (data transfer), 7% (floating point), and 9% (decimal). Once again, a 1% entry in the average column can occur because of entries in the constituent columns.

C.3 Intel 8086 Detailed Measurements

Instruction	Turbo C	MASM	Lotus	Average
Control	**21%**	**20%**	**32%**	**24%**
Conditional jumps	10%	12%	9%	10%
CALL, CALLF	4%	3%	5%	4%
RET, RETF	4%	3%	5%	4%
LOOP			12%	4%
JMP	3%	2%	2%	2%
Arithmetic, logical	**23%**	**24%**	**26%**	**25%**
CMP	8%	9%	5%	7%
SAL, SHR, RCR	2%	1%	11%	5%
ADD	3%	2%	3%	3%
OR, XOR	4%	2%	2%	3%
INC, DEC	3%	4%	3%	3%
SUB	2%	3%		2%
CBW	1%	1%		1%
TEST		2%	2%	1%
Data transfer	**49%**	**46%**	**30%**	**42%**
MOV	29%	31%	21%	27%
LES	6%	2%		3%
PUSH	10%	8%	4%	7%
POP	5%	6%	5%	5%
Totals	**93%**	**90%**	**88%**	**90%**

FIGURE C.3 The instructions responsible for more than 1.5% of the executions on any of the three benchmarks Some very similar instructions were combined for simplicity. Although MASM makes some use of string operations, the frequency is too low to make the table.

C.4 DLX Detailed Instruction Set Measurements

Instruction	GCC	Spice	TeX	US Steel	Average
Control	**21%**	**5%**	**7%**	**23%**	**14%**
B--Z	19%	2%	7%	16%	11%
J	2%	3%		3%	2%
JAL				2%	0%
JR				2%	0%
Arithmetic, logical	**37%**	**28%**	**41%**	**49%**	**39%**
ADDU,ADDUI	17%	16%	20%	27%	20%
LHI	2%	7%	10%	3%	5%
SLL	5%	5%	5%	4%	5%
LI	4%		4%	6%	4%
S--,S--I	5%		3%	3%	3%
AND,ANDI	2%			3%	1%
SRA	2%			2%	1%
OR,ORI				2%	1%
Data transfer	**28%**	**35%**	**33%**	**10%**	**26%**
LW	18%	8%	19%	5%	13%
SW	10%	2%	12%	5%	7%
LBU			2%		1%
LD		14%			4%
SD		6%			1%
MOVFP2I, MOVI2FP		5%			1%
Floating point	**0%**	**15%**	**0%**	**0%**	**4%**
FMUL		5%			1%
FADD		4%			1%
FSUB		3%			1%
FDIV		3%			1%
Totals	**85%**	**83%**	**82%**	**82%**	**83%**

FIGURE C.4 Instruction mixes for GCC, Spice, TeX, and the U.S. Steel COBOL benchmark. Some instructions were combined, both in the interest of space and because the combined class more correctly reflects what the processor is doing. The instruction class "B--Z" includes all conditional branches (which are all compares to zero). The class "S--,S--I" includes all set conditional instructions, both immediate and register–register. Immediate operations have been combined with the non-immediate class for all operations except loads, where they are distinctly different. Again, a blank space means that the instruction is not responsible for more than 1.5% of the executions, and the average may appear at 1% or less because the instruction is not used by all benchmarks.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 711

D Time Versus Frequency Measurements

D.1 Time Distribution on the VAX-11/780 D-2
D.2 Time Distribution on the IBM 370/168 D-4
D.3 Time Distribution on an 8086 in an IBM PC D-6
D.4 Time Distribution on a DLX Relative D-8

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.

D.1 Time Distribution on the VAX-11/780

We know from Chapters 2 and 3 that measuring instruction counts alone can be misleading. In this appendix we will examine the time distributions for some programs running on these four machines. For the 360, the 8086, and DLX, we will show the time distribution averaged over the three programs in the graph format used earlier. For the VAX, we will use measurements reported in Clark and Levy [1982] (see References in Chapter 4).

Figure D.1 shows the distribution of instruction executions, both by time and by frequency of occurrence. These data were measured by Emer and reported by Clark and Levy for a VAX-11/780 running VMS with multiple users doing three primary tasks:

1. Updating indexed files
2. Executing a matrix multiplication routine
3. Doing program development, including editing, compiling, and debugging

Figure D.1 includes any user instruction that accounts for more than 1% of the instruction executions or more than 1% of execution time. There are 26 instructions that fit this description, and together they account for 59% of the executions and 58% of the time. The measured data include the operating system and file system overhead.

Time distributions are particularly important on architectures like the VAX, where the number of cycles for an instruction may vary from one or two up to tens or hundreds.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 714

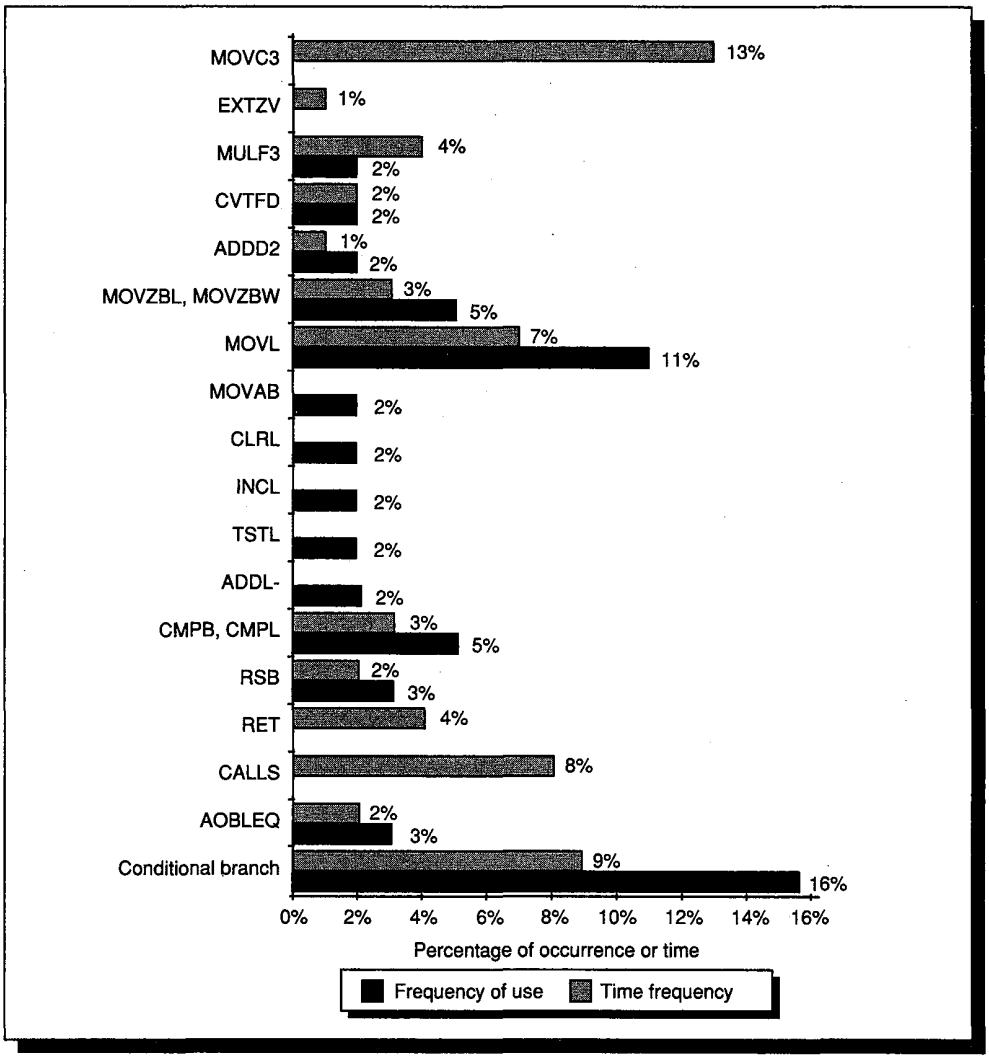

FIGURE D.1 Time and frequency distribution for a multiuser workload on a VAX-11/780 running VMS. This data includes all user instructions that are responsible for more than 1% of either the instruction executions or the execution time. (Two operating system instructions (`REI` and `MTPR`), each of which accounts for about 1% of the execution time, are not included.) The absence of an execution-frequency bar or time-frequency bar for an entry (such as `MOVC3` or `TSTL`) means that the time frequency or execution-time frequency is below 1% (not that it is 0!). Clark and Levy [1982] commented that the large percentage of time consumed by the `MOVC3` in the time distribution is somewhat abnormal for a nonbusiness workload and has not been observed in other measurements on the 11/780.

D.2 Time Distribution on the IBM 370/168

Figure D.2 shows the time distribution on an IBM 370/168 for the same programs we discussed in Chapter 4 and included in Figure 4.28 (page 175). All instructions that are responsible for more than 1.5% of the execution frequency and the execution time for at least one program are included. Several

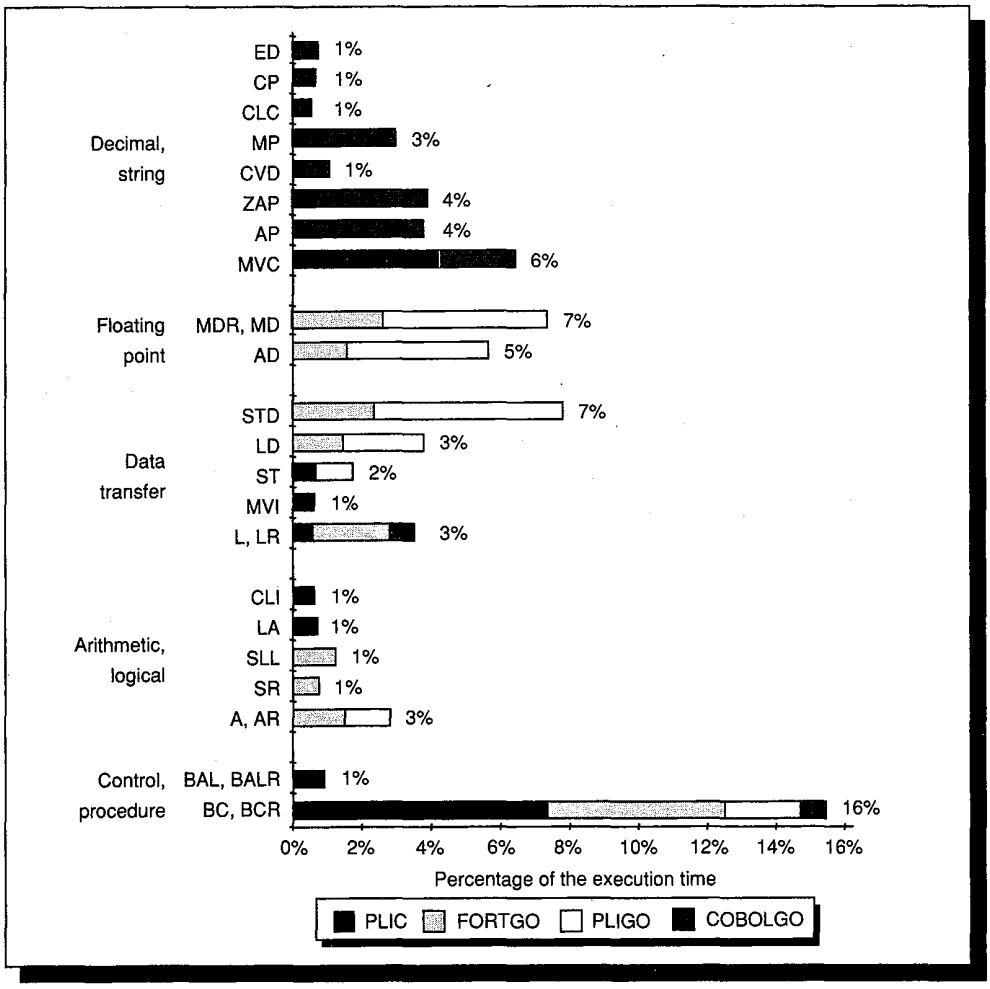

FIGURE D.2 Time distribution for the four programs discussed in Chapter 4 running on an IBM 370/168. The corresponding data on execution frequency appears in Figure 4.28 (page 175), or in table form in Figure C.2. Any instruction with greater than 1.5% frequency in the time distribution and in the execution-count distribution is included in this chart. Shustek [1978] (see References in Chapter 4) computed these numbers using a model of the 370/168 CPU. The model predicts the execution time for the programs and has an overall accuracy for each program of about 99% except on PLIGO, where it has an 8% error.

instructions appeared in the time distribution that were not in the frequency distribution, where their occurrence was too low. These instructions, which are not in Figure 4.28, are

TRT—Translate and test, a string instruction used by the PL/I compiler, most likely to scan the input source; takes 5.4% of the time in that program.

DP—Divide packed, a low frequency but long-running instruction that takes 18.7% of the time in COBOLGO.

DDR—Divide double register, a floating-point divide, infrequent but long running at 5.2% of the FORTGO execution time.

LM and STM—Load multiple and store multiple, with frequencies just below 1%, are somewhat slower than the average instruction; thus, they take 3% to 4% of the cycles in PLIGO.

BCT,BXLE—Loop branches that involve incrementing counts or doing other compares; BCT consumes about 2% of the time in PLIC, and BXLE consumes 3.5% in FORTGO.

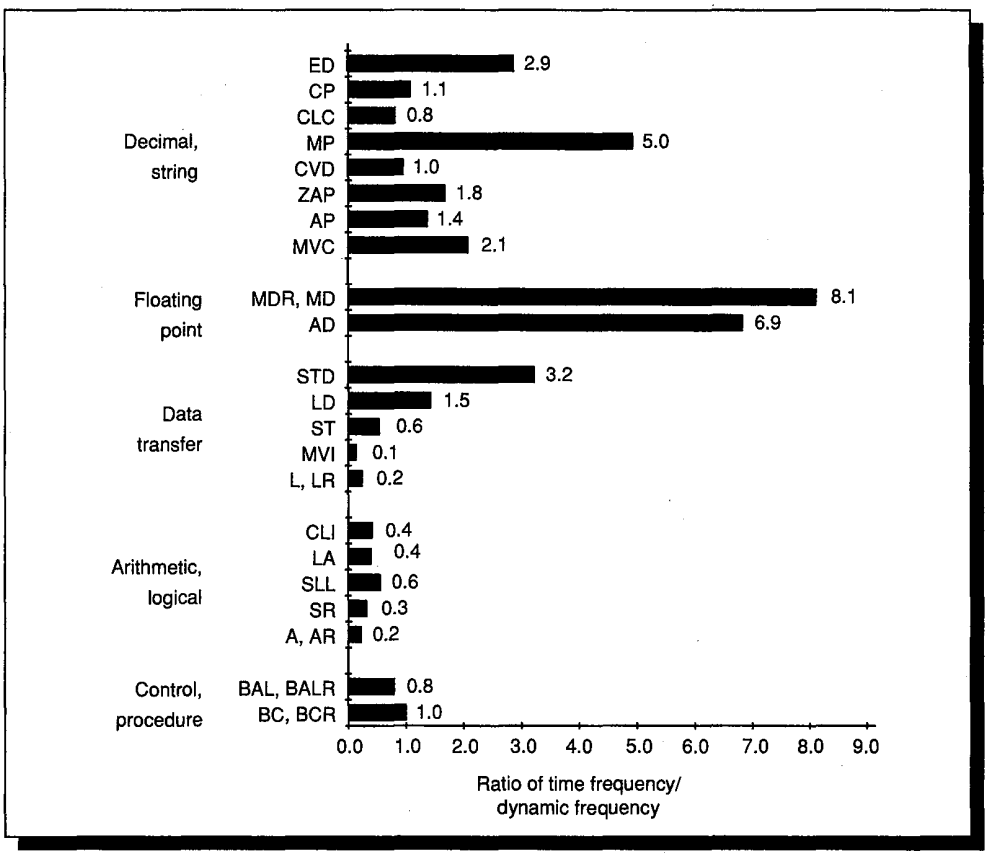

FIGURE D.3 Time frequency (percent of cycles doing this instruction as measured on an IBM 370/168) divided by dynamic frequency (percent of executions for this instruction). The programs are those in Chapter 4. This data is obtained directly from Figures 4.28 (page 175) and Figure D.2. This clearly shows that the floating-point instructions are the most expensive.

Several of the simpler but lower-frequency data transfer and ALU instructions that appeared in the frequency distribution do not appear in the time distribution because they constitute a very small percentage of the execution time. In total, the instructions shown in Figure D.2 account for 89% of the instruction executions and 72% of the execution time.

Figure D.3 gives the average execution time divided by the average frequency for those instructions that appear in both distributions. This measurement is a ratio that indicates the relative cost of an instruction. For example, an instruction that is responsible for 10% of the executions and 10% of the execution time will have a ratio of 1:1, or a cost factor of 1, and a CPI equal to the average CPI on the machine.

D.3 Time Distribution on an 8086 in an IBM PC

Figure D.4 continues our examination of time distribution by looking at the top time-consuming instructions on the 8086 for the same programs as measured in Chapter 4. These curves look very similar to those in Figure 4.32 (page 178), the frequency distribution for the 8086 (shown in table form in Figure C.3, page C-4). Two arithmetic and logical instructions, `CBW` and `SUB`, that appeared in the frequency distribution do not appear in the top of the execution-time distribution. Additionally, there are four instructions that have a significant contribution to the time frequency but are not in the execution-frequency distribution:

- String instructions `SCAS` (a string search) and `MOVS` (a string move). Both instructions are used in MASM, where they account for 8% and 7% of the execution time, respectively. `MOVS` is also used in Lotus, where it accounts for 6.6% of the program's execution time.
- Integer multiply and divide `ML16` and `DV16`. These are used in Lotus, where they respectively account for 10% and 4% of the program's execution time.

Together, the instructions in Figure D.4 are responsible for 87% of the instruction executions and 85% of the execution time.

Figure D.5 shows the ratio of execution time to execution frequency in the same fashion used for the IBM 360. Calls, returns, and loading a segment register consume a larger percentage of the execution time relative to their dynamic occurrence. However, the overall execution time profile of the 8086 is much closer to the execution frequency profile—the correspondence is often 1:1, and never as high as 1:2. This is primarily because the variation in CPI among instructions is small compared to an overall average CPI of 14.1. The long-running instructions that do not even appear in the frequency counts but are major consumers of execution time (and would have a high CPI) are the string instructions and integer multiply and divide.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 718

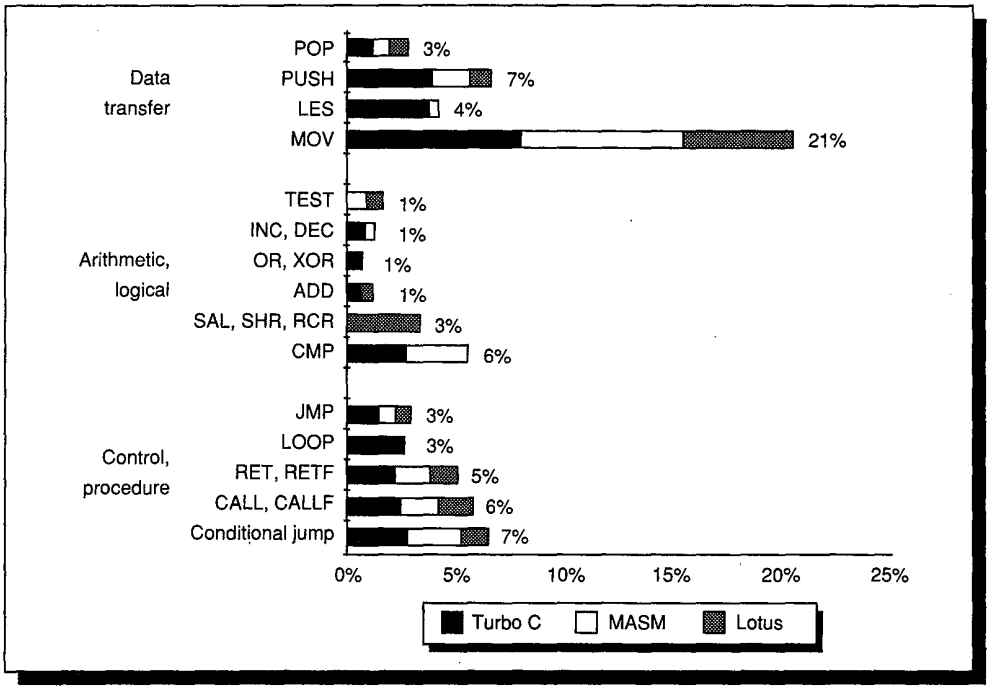

FIGURE D.4 The 8086 time distribution as measured on an IBM PC running MS-DOS. The format and data are the same as in Figure 4.32 (page 178).

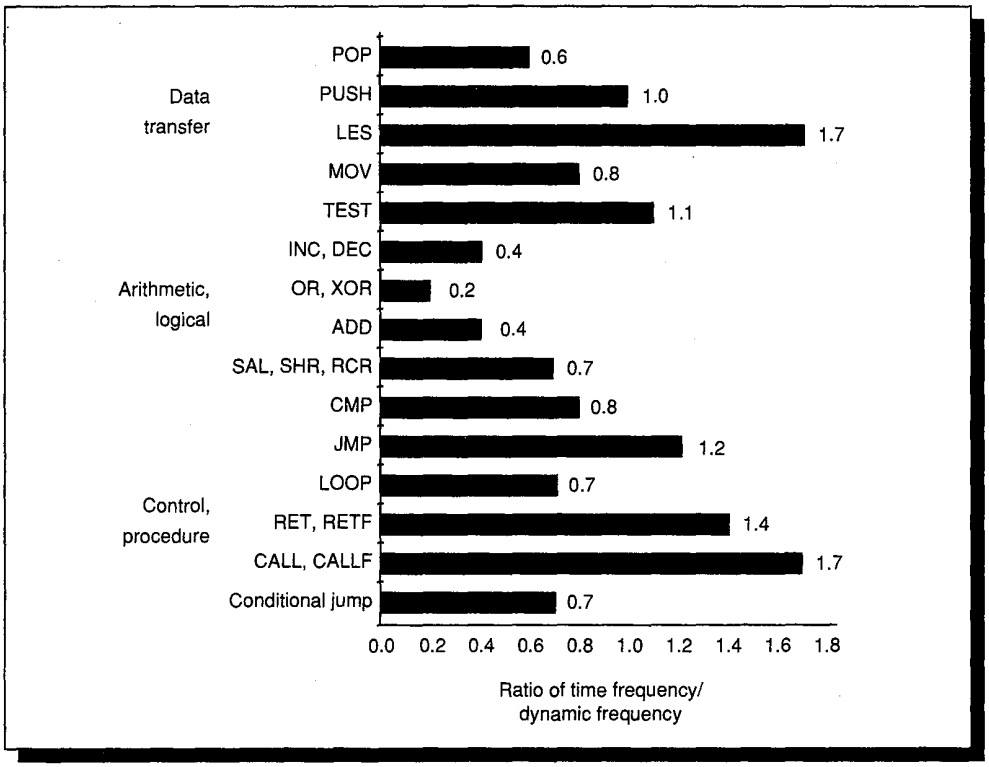

FIGURE D.5 Time distribution divided by frequency distribution for the 8086. This data is directly derived from Figures 4.32 (page 178) and D.4. The distribution is remarkably flatter than that for the IBM 360 or the VAX.

D.4 Time Distribution on a DLX Relative

To obtain a time distribution for DLX, we turn to the DECstation 3100, which has an instruction set architecture very similar to DLX (see Appendix E). The time distribution on the DECstation 3100 for the same programs measured in Chapter 4 (Figure 4.34 on page 181 and in table form in Figure C.4 is shown in Figure D.6. Figure D.6 includes all instructions that contribute more than 1% to the execution time. In total, these instructions account for 81% of all instruction executions and 97% of the execution time.

This time distribution is by far the closest to the frequency distribution. This is because under ideal conditions almost all instructions in DLX can take one cycle; only the `LD` and `SD` instructions must take two cycles. Of course, these perfect conditions never arise. The average CPI using the DECstation 3100 as a base is about 1.6 for GCC, TeX, and COBOLX, and about 2.1 for Spice.

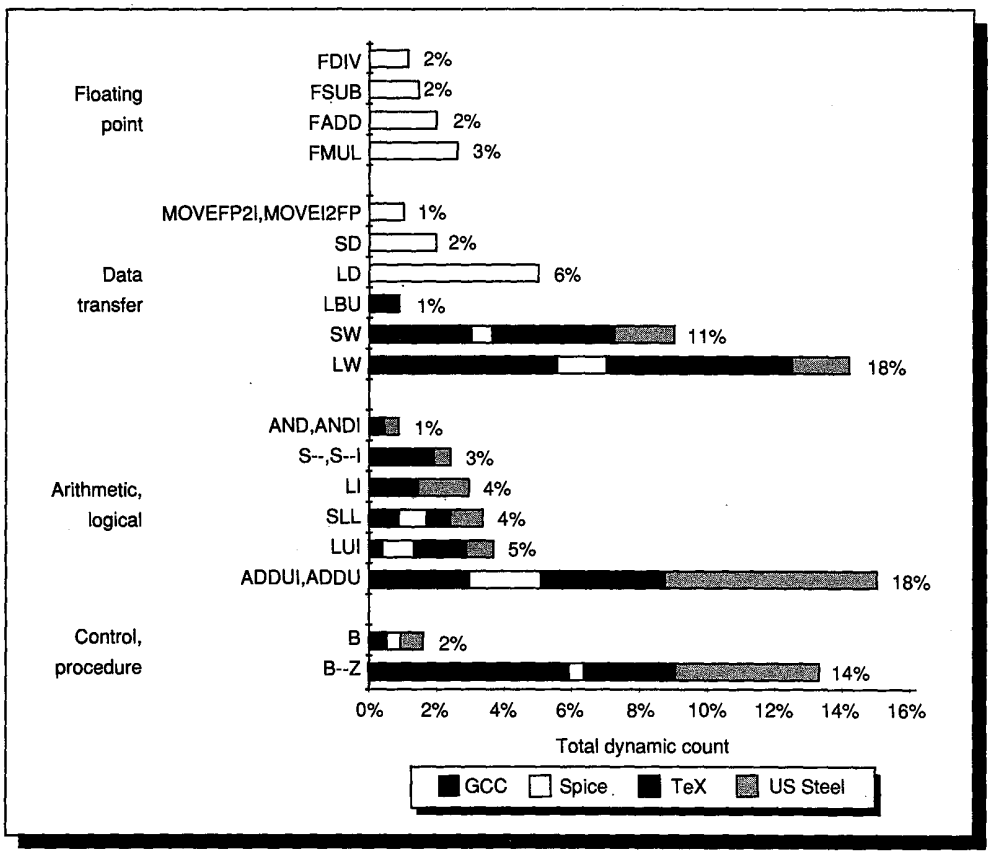

FIGURE D.6 The time distribution for our three benchmarks plus the US Steel COBOL benchmark as they would run on DLX using the CPI measurements from a DECstation 3100.

Figure D.7 shows contribution to execution time over contribution to execution frequency for the top instructions. Like the 360 and 8086 charts, a value above 1 indicates that this instruction has a higher CPI than the average instruction. Remember, though, that the ratio does not indicate the CPI for the instruction. However, we can use this figure to find the CPI for an instruction, given the base CPI for a specific program.

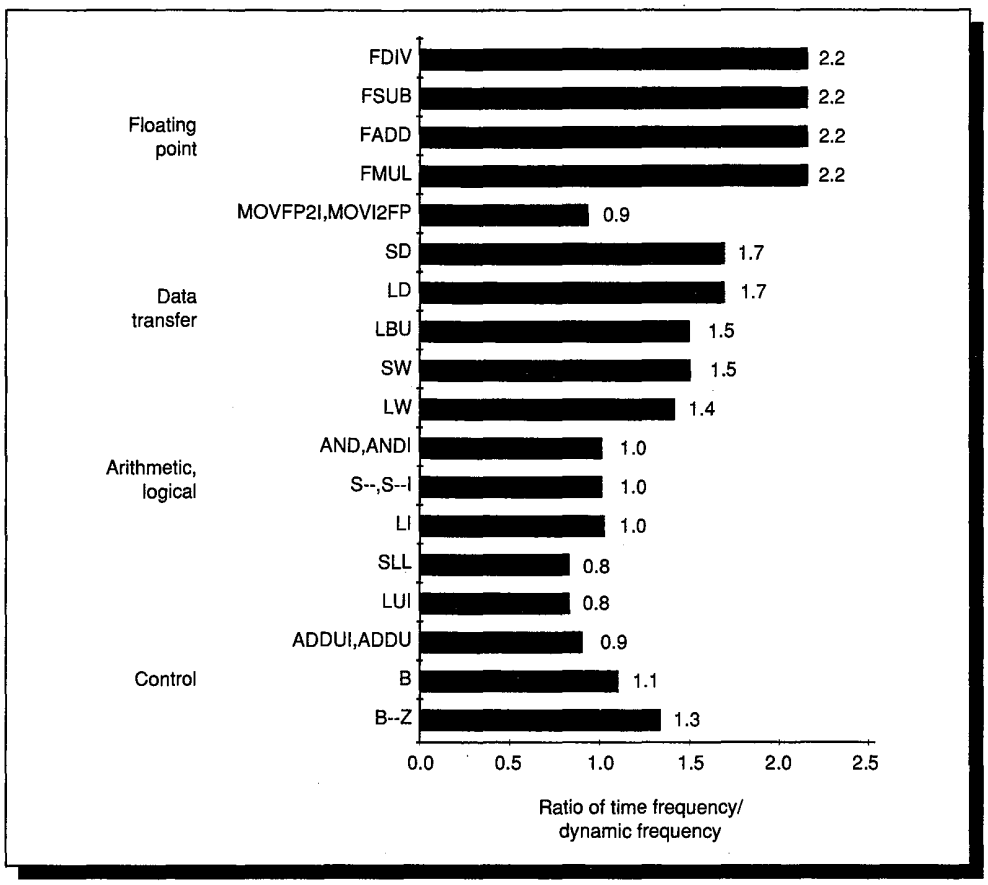

FIGURE D.7 Time frequency divided by execution frequency for DLX as measured using the time data from Figure D.6 and the frequency data from Figure 4.34 (page 181). The integer register–floating-point register moves are inexpensive, since they are really register–register operations. Surprisingly, the double-precision memory references are **not** twice as expensive as the 32-bit loads and stores. Can you hypothesize why based on the discussions of pipelining and cache design?

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 721

RISC: any computer announced after 1985.

Steven Przybylski (a designer of the Stanford MIPS)

E.1	**Introduction**	**E-1**
E.2	**Addressing Modes and Instruction Formats**	**E-2**
E.3	**Instructions: The DLX Subset**	**E-4**
E.4	**Instructions: Common Extensions to DLX**	**E-9**
E.5	**Instructions Unique to MIPS**	**E-12**
E.6	**Instructions Unique to SPARC**	**E-15**
E.7	**Instructions Unique to M88000**	**E-17**
E.8	**Instructions Unique to i860**	**E-19**
E.9	**Concluding Remarks**	**E-23**
E.10	**References**	**E-24**

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.

E Survey of RISC Architectures

E.1 Introduction

We cover four examples of Reduced Instruction Set Computer (RISC) architectures in this appendix:

- Intel 860;
- MIPS R3000/R3010 (plus a section on MIPS II, used in the R6000);
- Motorola M88000; and
- SPARC, developed originally by Sun Microsystems.

We also include DLX, the instruction set architecture invented for this book. (A review of DLX can be found in the back inside cover or in pages 160–167 of Chapter 4.) Characteristics of these architectures are found in Figure E.1.

There has never been another class of computers that were so similar. This similarity allows the presentation of four architectures at once, with DLX thrown in for good measure! After presenting the addressing modes and instruction formats, the instructions are presented in three steps:

- Instructions found in DLX;
- Instructions not found in DLX but found in two or more architectures; and
- The unique instructions and characteristics of each architecture.

We conclude with a speculation about the future directions for RISCs.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 723

	DLX	i860	MIPS	M88000	SPARC
Date announced	1990	1989	1986	1988	1987
Instruction size (bits)	32	32	32	32	32
Address space (size, model)	32 bits, flat	32 bits, flat	32 bits, flat	32 bits, flat	32 bits, flat
Data alignment	Aligned	Aligned	Aligned	Aligned	Aligned
Data addressing modes	1	2	1	3	2
Protection	Page	Page	Page	Page	Page
Page size	4 KB	4 KB	4 KB	4 KB	4–64 KB
I/O	Memory mapped	Memory mapped	Memory mapped	Memory mapped	Memory mapped
Integer registers (size, model, number)	31 GPR x 32 bits	31 GPR x 32 bits	31 GPR x 32 bits	31 GPR x 32 bits	31 GPR x 32 bits
Separate floating-point registers	32 x 32 or 16 x 64 bits	30 x 32 or 15 x 64 bits	16 x 32 or 16 x 64 bits	0	32 x 32 or 16 x 64
Floating-point format	IEEE 754 single, double	IEEE 754 single, double	IEEE 754 single, double	IEEE 754 single, double	IEEE 754 single,double

FIGURE E.1 Summary of five recent architectures. Except for number of data address modes and some instruction set details, the integer instruction sets of these architectures of the late 1980s are identical. Contrast this to Figure E.13, page E-23.

E.2 Addressing Modes and Instruction Formats

Figure E.2 shows the data addressing modes supported by each architecture. Since all have one register that always has the value 0—in fact, it is `r0` in every architecture—the absolute address mode with limited range can be synthesized using `r0` as the base in displacement addressing. Similarly, register-indirect addressing is synthesized by using displacement addressing with an offset of 0. Simplified addressing modes is one distinguishing feature between these and prior architectures.

Addressing mode	DLX	i860	MIPS	M88000	SPARC
Register + offset (displacement or based)	√	√	√	√	√
Register + register (indexed)	--	√	--	√	√
Register + scaled register (scaled)	--	--	--	√	--

FIGURE E.2 Summary of data addressing modes. (These addressing modes are explained in Section 3.4, pages 94–103) While the i860 does have indexed data addressing for all loads and floating-point stores, it is not available for integer stores.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 724

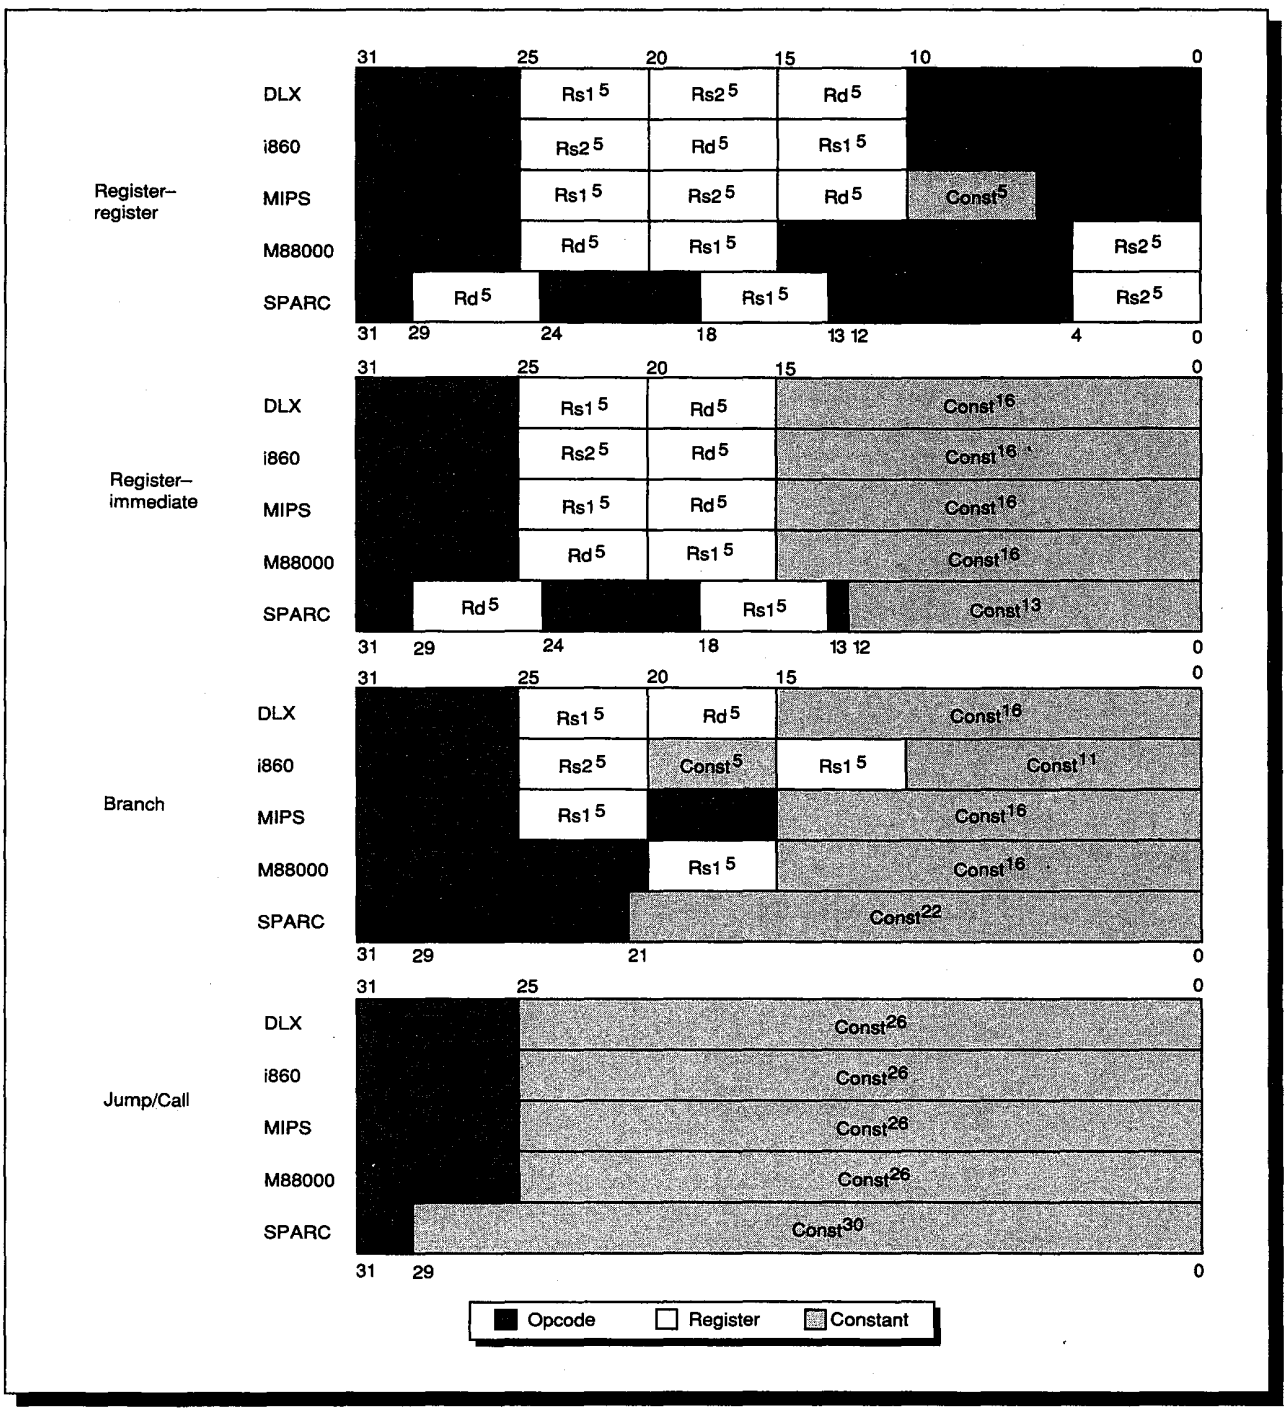

FIGURE E.3 Instruction formats for five architectures. These four formats are found in all five architectures. (The superscript notation in this figure means something different from our standard notation; it shows the width of a field in bits.) While the register fields are located in similar pieces of the instruction, beware that the destination and two source fields are scrambled. Here are the meanings of the abbreviations: Op = the main opcode, Opx =an opcode extension, Rd = the destination register, Rs1 = source register 1, Rs2 = source register 2, and Const = a constant (used as an immediate or as an address). The main variation for the M88000 is register–immediate format when the operation doesn't need a full 16-bit immediate: an opcode extension field is placed in the upper bits of the constant field. The variation for the i860 is using Rs1 in the Branch format to specify a 5-bit constant as well as a register.

References to code are normally PC-relative, although register indirect is supported for returning from procedures and for case statements. One variation is that PC-relative branch addresses in everything but DLX are shifted left 2 bits before being added to the PC, thereby increasing the branch distance. This works because the length of all instructions is one word and instructions must be word aligned in memory.

Figure E.3 (page E-3) shows the format of instructions, which includes the size of the address in the instructions. Each instruction set architecture uses these four primary instruction formats. The primary differences are subtle, concerning how to extend constant fields to 32 bits. Figure E.4 shows the variations.

Format: instruction category	DLX	i860	MIPS	M88000	SPARC
Branch: all	Sign	Sign	Sign	Sign	Sign
Jump/Call: all	Sign	Sign	--	Sign	Sign
Register–immediate: data transfer	Sign	Sign	Sign	Zero	Sign
Register–immediate: arithmetic	Sign	Sign	Sign	Zero	Sign
Register–immediate: logical	Sign	Zero	Zero	Zero	Sign

FIGURE E.4 Summary of constant extension. The constant in the Jump and Call instructions of MIPS are not sign extended since they only replace the lower 28 bits of the PC, leaving the upper 4 bits unchanged.

E.3 Instructions: The DLX Subset

The similarities of each architecture allow simultaneous descriptions of the architectures, starting with the operations equivalent to DLX.

DLX Instructions

Almost every instruction found in DLX instructions is found in the other architectures, as Figure E.5 shows. (For reference, definitions of the DLX instructions are found on pages 160 to 167 of Chapter 4 and the back inside cover.) Instructions are listed under four categories: "Data transfer," "Arithmetic, logical," "Control," and "Floating point." A fifth category in the figure shows conventions for register usage and pseudoinstructions on each architecture. If a DLX instruction requires a short sequence of instructions, these instructions are separated by semicolons in Figure E.5. (To avoid confusion, the destination register will **always** be the leftmost operand in this appendix, independent of the notation normally used with each architecture.)

Every architecture must have a scheme for compare and conditional branch, but even with all the similarities, each of these architectures has found a different way to perform the operation. The advantages and disadvantages of the general options are found on pages 105–109 of Chapter 3.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 726

Instruction name	DLX	i860	MIPS	M88000	SPARC
Data transfer (Instruction formats)	**R–I**	**R–I, R–R**	**R–I**	**R–I, R–R**	**R–I, R–R**
Load byte signed	LB	LD.B	LB	LD.B	LDSB
Load byte unsigned	LBU	LD.B; AND ...,x00FF,...	LBU	LD.BU	LDUB
Load halfword signed	LH	LD.S	LH	LD.H	LDSH
Load halfword unsigned	LHU	LD.S; AND ...,xFFFF...	LHU	LD.HU	LDUH
Load word	LW	LD.L	LW	LD	LD
Load SP float	LF	FLD.L	LWC1	LD	LDF
Load DP float (see E.5 for MIPS)	LD	FLD.D	LWC1 Rd; LWC1 Rd+1	LD.D	LDDF
Store byte	SB	ST.B	SB	ST.B	STB
Store halfword	SH	ST.S	SH	ST.H	STH
Store word	SW	ST.L	SW	ST	ST
Store SP float	SF	FST.L	SWC1	ST	STF
Store DP float (see E.5 for MIPS)	SD	FST.D	SWC1 Rd; SWC1 Rd+1	ST.D	STDF
Read, write special registers	MOVS2I, MOVI2S	LD.C, ST.C	MF_, MT_	LDCR,FLDCR STCR,FSTCR	RD, LDFSR, WR, STFSR
Move int. to FP reg.	MOVI2FP	IXFR	MFC1	not applicable	ST;LDF,
Move FP to int. reg.	MOVFP2I	FXFR	MTC1	not applicable	STF;LD
Arithmetic, logical (Instruction formats)	**R–R, R–I**	**R–R, R–I**	**R–R, R–I**	**R–R, R–I**	**R–R, R–I**
Add	ADDU,ADDUI	ADD,ADDU	ADDU,ADDIU	ADDU	ADD
Add (trap if overflow)	ADD,ADDI	ADD; INTOVR	ADD,ADDI	ADD	ADDcc; TVS
Sub	SUBU,SUBUI	SUB,SUBU	SUBU	SUBU	SUB
Sub (trap if overflow)	SUB,SUBI	SUB; INTOVR	SUB	SUB	SUBcc; TVS
Multiply (see E.6 for SPARC)	MULTU, MULTUI	FMLOW	MULT, MULTU	MUL	MULScc;...; MULScc
Multiply (trap if ovf)	MULT,MULTI	--	--	--	-- (see E.6)
Divide *	DIVU,DIVUI	--	DIV,DIVU	DIV,DIVU	-- (see E.6)
Divide (trap if ovf)	DIV,DIVI	--	--	--	-- (see E.6)
And	AND,ANDI	AND	AND,ANDI	AND	AND
Or	OR,ORI	OR	OR,ORI	OR	OR
Xor	XOR,XORI	XOR	XOR,XORI	XOR	XOR
Load high part reg.	LHI	OR.H ...,r0,...	LUI	OR.U ...,r0,...	SETHI *(B fmt.)*
Shift left logical	SLL,SLLI	SHL	SLLV,SLL	MAK	SLL
Shift right logical	SRL,SRLI	SHR	SRLV,SRL	EXTU	SRL
Shift right arithmetic	SRA,SRAI	SHRA	SRAV,SRA	EXT	SRA
Compare	S-(<,>,≤,≥,=,≠)	SUB r0,...	SLT,SLTU, SLTI,SLTIU	CMP	SUBcc r0,...

Instruction Name	DLX	i860	MIPS	M88000	SPARC
Control (Instruction formats)	**B, J/C**	**B, J/C**	**B, J/C**	**B, J/C**	**B, J/C**
Branch on integer compare	BEQ,BNE	BC.T,BNC.T, BTE,BTNE	BEQ,BNE,B_Z (<,>,≤,≥)	BB1.N,BB0.N, BCND.N	Bicc (<,>,≤,≥,=,≠)
Branch on floating-point compare	BFPT,BFPF	BC.T,BNC.T	BC1T,BC1F	BB1.N,BB0.N, BCND.N	FBfcc (<,>,≤,≥,=,...)
Jump, jump register	J,JR	BR, BRI	J,JR	BR.N,JMP.N	B, JMPL r0,...
Call, call register	JAL,JALR	CALL, CALLI	JAL,JALR	BSR.N,JSR.N	CALL,JMPL
Trap	TRAP	TRAP	BREAK	TCND, TB0	Ticc
Return from interrupt	RFE	BRI *(trap bits≠0)*	JR; RFE	RTE	RETT
Floating point (Instruction formats)	**R–R**	**R–R**	**R–R**	**R–R**	**R–R**
Add single, double	ADDF, ADDD	FADD.SS, FADD.DD	ADD.S, ADD.D	FADD.SSS, FADD.DDD	FADDS, FADDD
Sub single, double	SUBF, SUBD	FSUB.SS, FSUB.DD	SUB.S, SUB.D	FSUB.SSS, FSUB.DDD	FSUBS, FSUBD
Mult single, double	MULF, MULD	FMUL.SS, FMUL.DD	MUL.S, MUL.D	FMUL.SSS, FMUL.DDD	FMULS, FMULD
Div single, double	DIVF, DIVD	--, --	DIV.S, DIV.D	FDIV.SSS, FDIV.DDD	FDIVS, FDIVD
Compare	_F, _D (<,>,≤,≥,=,...)	PF_.SS, PF_.DD (>,≤,=)	C_.S, C_.D (<,>,≤,≥,=,...)	FCMP.SS, FCMP.DD	FCMPS, FCMPD
Move R–R	MOVF	FIADD.SS ...,f0,	MOV.S	ADD ...,r0,...	FMOVS
Convert (single,double,integer) to (single,double,integer)	CVTF2D, CVTD2F, CVTF2I, CVTD2I, CVTI2F, CVTI2D	FADD.SD ..f0.., FADD.DS ..f0.., FIX.SS, FIX.DS, --, --	CVT.S.D, CVT.D.S, CVT.S.W, CVT.D.W, CVT.W.S, CVT.W.D	FADD.SSD r0, --, INT.SS, INT.SD, FLT.SS, FLT.DS	FSTOD, FDTOS, FSTOI, FDTOI, FITOS, FITOD
Conventions					
Register with value 0	r0	r0	r0	r0	r0
Return address reg.	r31	r1	r31	r1	r31
Noop	ADD r0,r0,r0	SHL r0,r0,r0	SLL r0,r0,r0	OR r0,r0,r0	SETHI r0,0
Move R–R integer	ADD ...,r0,...	SHL ...,r0,...	ADD ...,r0,...	OR ...,r0,...	OR ...,r0,...
Operand order	OP Rd,Rs1,Rs2	OP Rs1,Rs2,Rd	OP Rd,Rs1,Rs2	OP Rd,Rs1,Rs2	OP Rs1,Rs2,Rd

FIGURE E.5 Instructions equivalent to DLX. Dashes mean the operation is not available in that architecture, or not synthesized in a few instructions. Such a sequence of instructions is shown separated by semicolons. If there are several choices of instructions equivalent to DLX, they are separated by commas. Finally, "not applicable" means that while this operation is not directly available, other changes in the architecture means it wouldn't make sense. This later category is for the M88000, since integer and floating-point instructions sharing the same registers means separate floating-point move instructions are unnecessary. Note that in the "Arithmetic, logical" category DLX and MIPS use separate instruction mnemonics to indicate an immediate operand, while the i860, M88000, and SPARC offer immediate versions of these instructions but use a single mnemonic. (Of course these are separate opcodes!) Both MIPS and SPARC have new instructions that were not implemented in the first machine and that apply to some of these cases: see Sections E.5 and E.6.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 728

Compare and Conditional Branch

SPARC uses the traditional four condition code bits stored in the program status word: Negative, Zero, Carry, and Overflow. They can be set on any arithmetic or logical instruction, but unlike earlier architectures this setting is optional on each instruction. This leads to fewer problems in pipelined implementation (page 334 in Chapter 6). While condition codes can be set as a side effect of an operation, explicit compares are synthesized with a subtract using `r0` as the destination. Floating point uses separate condition codes to encode the IEEE 754 conditions, requiring a floating-point compare instruction. SPARC conditional branches test condition codes to determine all possible unsigned and signed relations.

MIPS uses the contents of registers to evaluate conditional branches. Any two registers can be compared for equality (`BEQ`) or inequality (`BNE`) and then the branch is taken if the condition holds. The set-on-less-than instructions (`SLT,SLTI, SLTU,SLTIU`) compare two operands and then set the destination register to 1 if less and to 0 otherwise. These instructions are enough to synthesize the full set of relations. Because of the popularity of comparisons to 0, MIPS includes special compare-and-branch instructions for all such comparisons: greater than or equal to zero (`BGEZ`), greater than zero (`BGTZ`), less than or equal to zero (`BLEZ`), and less than zero (`BLTZ`). Of course, equal and not equal to zero can be synthesized using `r0` with `BEQ` and `BNE`. Like SPARC, MIPS uses a condition code for floating point with separate floating-point compare and branch instructions.

The M88000 also uses registers to evaluate conditions and optimizes compare to 0 with a separate set of compare-and-branch instructions (`BCND.N`). Comparison of arbitrary operands differs. MIPS offers several compare instructions to set the register to 0 or 1 depending on the selected condition, but the M88000 uses a single instruction (`CMP`) and sets 10 bits of the destination register showing the relationship of the two operands. These bits represent equality ($=, \neq$) plus all relations for signed ($<, \leq, >, \geq$) and unsigned ($<, \leq, >, \geq$) operands. Instructions that branch if a bit in a register is 1 (`BB1.N`) or 0 (`BB0.N`) complete the conditional branch set. (Another option is using `EXTU` with `CMP` to set a register to 0 or 1 and then using `BCND.N`. Using `EXT` instead of `EXTU` sets a register to 0 or –1, if so desired.) Since there is a common register set for integer and floating point, floating-point compare uses the same scheme: set bits of a register and branch based on the result using `BB1.N` or `BB0.N`.

The Intel i860 uses condition codes for branches like SPARC, except that the i860 condition codes are set implicitly as part of every integer arithmetic or logical instruction. Also unlike SPARC, the i860 uses just two bits of conditions: `OF` and `CC`. `OF` is set only by the integer add and subtract instructions, and is used to indicate overflow. There is no conditional branch instruction to test this bit, but the `INTOVR` instruction will cause a trap if the bit is set. The CC bit is set or cleared depending on the operation. The logical instructions (`AND,OR,XOR`) set CC if the result is 0. The unsigned arithmetic instructions (`ADDU,SUBU`) set CC

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 729

if there is a carry out of the most significant bit. Signed subtract (SUBS) sets CC if Rs2 > Rs1, while signed add (ADDS) sets CC if Rs2 is less than the two's complement of Rs1. Floating-point comparison instructions set CC if the condition tested is true: greater than (PFGT), less than or equal (PFLE), or equal (PFEQ).

The i860 conditional branch instructions (BC.T and BNC.T) test CC and branch depending on whether CC is 1 or 0. The i860 also has conditional branch instructions based on equality of two operands: BTE jumps if they are equal and BTNE jumps if they are not.

Figure E.6 summarizes the four schemes used for conditional branches.

	DLX	i860	MIPS	M88000	SPARC
Number of condition code bits (integer and FP)	1 FP	1 both, 1 integer	1 FP	--	4 integer, 2 FP
Basic compare instructions (integer and FP)	1 integer, 1 FP	1 FP	1 integer, 1 FP	1 integer, 1 FP	1 FP
Basic branch instructions (integer and FP)	1 integer, 1 FP	1 both, 1 integer	2 integer, 1 FP	1 both, 1 integer	1 integer, 1 FP
Compare register with register/const and branch	=,≠	=,≠	=,≠	--	--
Compare register to zero and branch	=,≠	=,≠	=,≠,<,≤,>,≥	=,≠,<,≤,>,≥	--

FIGURE E.6 Summary of five approaches to conditional branches. Integer compare on the i860 and SPARC is synthesized with an arithmetic instruction that sets the condition codes using r0 as the destination.

Integer Multiply and Divide

Multiply and divide are usually implemented as multicycle instructions and are thus not a good match for the single-cycle execution goal of the rest of the integer instructions, requiring separate integration into the pipeline. Each architecture takes a different approach to integer multiply and divide as well as conditional branch. The i860 uses the same scheme as DLX: there is a floating-point instruction (FMLOW) that treats the contents of two floating-point registers as integers, leaving a 32-bit result in the lower 32 bits of a double-precision pair of floating-point registers. Programs do integer divide using i860 floating-point instructions. (Floating-point divide uses Newton-Raphson iteration; see pages E-19–E-20.)

The combined integer and floating-point register file allows the M88000 to use the floating-point unit to perform integer multiply and divide, as the operands do not have to be moved to and from the floating-point registers. The one complication in the first version of the architecture, the MC88100, is a negative dividend or negative divisor results in a trap. Software then makes the operands positive, uses the divide instruction, and then complements the quotient (if necessary). A zero divisor traps as well, as we would hope.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 730

In the MIPS architecture the 64-bit product of an integer multiply or the quotient/remainder of an integer divide is placed in a special registers HI and LO. This computation is treated as an independent unit executing in parallel with the integer and floating-point units. The appropriate result is transferred to the correct register with a `MFHI` or `MFLO` instruction. Attempts to read the registers before the computation is complete stalls the processor. There is no trap for overflow or divide by zero. These are typically checked by explicit integer instructions that execute in parallel with the divide. (See Section E.5 for architectural extensions not implemented in the first MIPS machines.)

SPARC provides a multiply step instruction. When used in a loop it calculates a full 64-bit product using the special register Y. It is loaded with the multiplier and receives the least significant word of the product. Magenheimer, Peters, Pettis, and Zuras [1988] measured the size of operands in multiplies and divides to show how well the multiply step would work. Using this data for C programs, Muchnick [1988] found that by making special cases the average multiply by a constant takes 6 clock cycles and multiply of variables takes 24 clock cycles. There is no divide step in the SPARC. (See Section E.6 for architectural extensions not implemented in the first SPARC machines.)

E.4 Instructions: Common Extensions to DLX

Figure E.7 (pages E-10–E-11) lists instructions not found in Figure E.5 (pages E-5–E-6) in the same four categories. Instructions are put in this list if they appear in more than one of the four architectures. The instructions are defined using the hardware description language, which is described on the page facing the inside back cover and on pages 160–167 of Chapter 4.

While most of the categories are self-explanatory, a few bear comment:

- The "Atomic swap" row means a primitive that can exchange a register with memory without interruption. This is useful for operating system semaphores in uniprocessors as well as for multiprocessor synchronization (see pages 471–473 of Chapter 8.)
- In the "Endian" row, "Big or Little" means there is a bit in the program status register that allows the processor to act either as Big Endian or Little Endian. This can be accomplished by simply complementing some of the least significant bits of the address in data transfer instructions.
- The "Coprocessor operations" row lists several categories that allow for the processor to be extended with special-purpose hardware.
- The "Implicit conversions" row under "Floating point" means that floating-point operands in these architectures do not have to all be the same size, and the floating-point unit performs a conversion as part of the operation. The i860 allows for two single-precision operands to produce a double-precision

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 731

result while the M88000 allows for any combination of single and double precisions for each of the three operands.

One difference that needs a longer explanation is the optimized branches. Figure E.8 (page E-12) shows the options. The i860 and M88000 offer branches that take effect immediately, like branches on earlier architectures. This avoids executing NOPs when there is no instruction to fill the delay slot. SPARC provides a version of delayed branch that makes it easier to fill the delay slot. The "annulling" branch executes the instruction in the delay slot only if the branch is taken; otherwise the instruction is annulled. This means the instruction at the target of the branch can safely be copied into the delay slot since it will only be executed if the branch is taken. The restrictions are that the target is not another branch and that the target is known at compile time. SPARC also offers a nondelayed jump because an unconditional branch with the annul bit set does **not** execute the following instruction.

After covering the similarities, we will cover the unique features of each architecture, ordering them by length of description of the unique features from shortest to longest.

Name	Definition	i860	MIPS	M88000	SPARC
Data transfer					
Atomic swap R/M (for semaphores)	Temp←Rd; Rd← Mem[x]; Mem[x]←Temp	LOCK;LD.L; UNLOCK; ST.L;	-- (see E.5)	XMEM, XMEMBU	SWAP
Load double integer	Rd←Mem[x]; Rd+1←Mem[x+4]	--	--	LD.D	LDD
Store double integer	Mem[x]←Rd; Mem[x+4]←Rd+1	--	--	ST.D	STD
Load coprocessor	Coprocessor←Mem[x]	--	LWCi	--	LDC
Store coprocessor	Mem[x]←Coprocessor	--	SWCi	--	STC
Endian	(Big/Little Endian?)	Big or Little	Big or Little	Big or Little	Big
Cache flush	(Flush cache block at this address)	FLUSH	-- (see E.5)	--	FLUSH
Arithmetic, logical					
Support for multi-word integer add	CarryOut,Rd ← Rs1 + Rs2 + OldCarryOut	ADDU;BNC; ADDU ...,...,#1	ADDU;SLTU; ADDU	ADDU.CIO	ADDXcc
Support for multi-word integer sub	CarryOut,Rd ← Rs1 – Rs2 + OldCarryOut	SUBU;BNC; ADDU ....,...,#1	SUBU;SLTU; SUBU	SUBU.CIO	SUBXcc
And not	Rd ← Rs1 & !(Rs2)	ANDNOT	--	AND.C *(R–R)*	ANDN
Or not	Rd ← Rs1 | !(Rs2)	--	--	OR.C *(R–R)*	ORN
Xor not	Rd ← Rs1 ^ !(Rs2)	--	--	XOR.C *(R–R)*	XNOR

	Definition	i860	MIPS	M88000	SPARC
Arithmetic, logical *(continued)*					
And high immediate	$Rd_{0..15} \leftarrow Rs1_{0..15}$ & (Const<<16); $Rd_{16..31} \leftarrow 0$	ANDH *(R–I)*	--	AND.U *(R–I)*	--
Or high immediate	$Rd_{0..15} \leftarrow Rs1_{0..15}$ | (Const<<16); $Rd_{16..31} \leftarrow 0$	ORH *(R–I)*	--	OR.U *(R–I)*	--
Xor high immediate	$Rd_{0..15} \leftarrow Rs1_{0..15}$ ^ (Const<<16); $Rd_{16..31} \leftarrow 0$	XORH *(R–I)*	--	XOR.U *(R–I)*	--
Coprocessor operations	(Defined by coprocessor)	--	COPi	--	CPop
Control					
Optimized delayed branches	(Branch not always delayed)	BC,BNC	--	BB1,BB0, BCND	Bicc,A
Optimized floating-point branches	(Branch not always delayed)	BC,BNC	--	BB1,BB0, BCND	Bfcc,A
Conditional trap	if (COND) {R31←PC; PC ←0..0#i}	--	-- (see E.5)	TB1, TB0, TCND	Ticc
Branch on coprocessor	if (CoProc COND) {PC ←PC+Const}	--	BCiT,BCiF	--	Bccc
No. control regs.	Misc. regs (virtual memory, interrupts,...)	6	12	32	7
Floating point					
Negate	Fd ← Fs ^ x80000000	--	NEG.S, NEG.D	XOR.U 8000	NEGS
Absolute value	Fd ← Fs & x7FFFFFFF	--	ABS.S, ABS.D	AND.U 7FFF	ABSS
Truncate to integer	Fd ← unrounded integer part of Fs	FTRUNC.SS, FTRUNC.DS	--	TRNC.SS, TRNC.SD	--
Implicit conversions	(Convert as part of operation)	_.SD (2 single operands, 1 double result)	--	_.SSD,_.SDS, _.SDD,_.DSS, _.DSD._.DDS (all combinations)	--

FIGURE E.7 Instructions not found in DLX but found in two or more of the four architectures. Both MIPS and SPARC have new instructions that were not implemented in the first machine and that apply to some of these cases: see Sections E.5 and E.6.

	Delayed branch	(Plain) Branch	Annulling delayed branch
Found in architectures	All 5 RISCs	i860, M88000	SPARC
Execute following instruction	Always	Only if branch **not** taken	Only if branch taken

FIGURE E.8 When the instruction following the branch is executed for three types of branches.

E.5 Instructions Unique to MIPS

Starting with data transfer instructions, MIPS is unlike the others since the architecture requires that the instruction following a load does not refer to the value being loaded. The MIPS Assembler inserts a NOOP instruction if this situation occurs.

Nonaligned Data Transfers

The other unique feature of MIPS data transfer is special instructions to handle misaligned words in memory. A rare event in most programs, it is included for COBOL programs where the programmer can force misalignment by declarations. While all these architectures trap if you try to load a word or store a word to a misaligned address, on all architectures misaligned words can be accessed without traps by using 4 load byte instructions and then assembling the result using shifts and logical ORs. The MIPS load and store word left and right instructions (`LWL`, `LWR`, `SWL`, `SWR`) allow this to be done in just 2 instructions: `LWL` loads the left portion of the register and `LWR` loads the right portion of the register. `SWL` and `SWR` do the corresponding stores. Figure E.9 shows how they work. Unlike other loads, a `LWL` followed by a `LWR` does not require a NOOP even though both will specify the same register since fields do not overlap.

TLB Instructions

TLB misses are handled in software in the MIPS R2000, so the instruction set also has instructions for manipulating the registers of the TLB (see pages 437–438 and 443–445 in Chapter 8 for more on TLBs.) These registers are considered part of the "system coprocessor" and thus can be accessed by the instructions that move between coprocessor registers and integer registers. The contents of a TLB entry are read by loading via Read Indexed TLB Entry (`TLBR`) and written using either Write Indexed TLB Entry (`TLBWI`) or Write Random TLB Entry (`TLBWR`). The TLB contents are searched using Probe TLB for Matching Entry (`TLBP`).

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 734

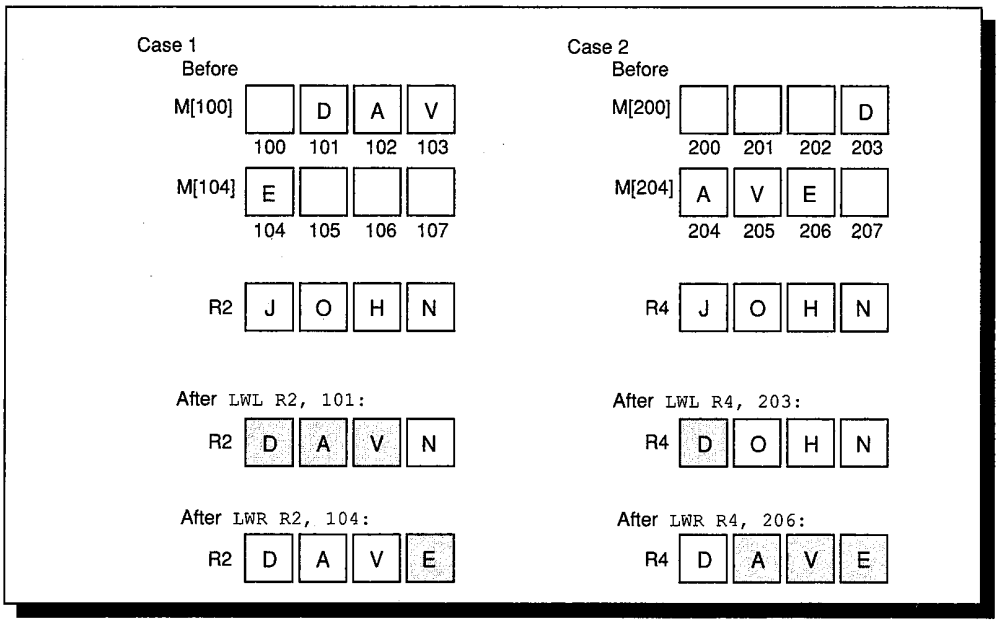

FIGURE E.9 MIPS instructions for unaligned word reads. This figure assumes operating in Big Endian mode. Case (1) first loads the 3 bytes 101,102, and 103 into the left of R2 leaving the least significant byte undisturbed. The following `LWR` simply loads byte 104 into the least significant byte of R2 leaving the other bytes of the register unchanged using `LWL`. Case (2) first loads byte 203 into the most significant byte of R4 and the following `LWR` loads the other 3 bytes of R4 from memory bytes 204, 205, and 206. `LWL` reads the word with the first byte from memory, shifts to the left to discard the unneeded byte(s), and changes only those bytes in Rd. The byte(s) transferred are from the first byte until the lowest-order byte of the word. The following `LWR` addresses the last byte, right shifts to discard the unneeded byte(s), and finally changes only those bytes of Rd. The byte(s) transferred are from the last byte up to the highest-order byte of the word. Store word left (`SWL`) is simply the inverse of `LWL`, and store word right (`SWR`) is the inverse of `LWR`. Changing to Little Endian mode flips which bytes are selected and discarded. (If big/little–left/right–load/store seems confusing, don't worry, it works!)

Remaining Instructions

Below is a list of the remaining unique details of the MIPS architecture:

- *NOR:* This logical instruction calculates !(Rs1 | Rs2).
- *Constant shift amount:* Nonvariable shifts use the 5-bit constant field shown in the register–register format in Figure E.3.
- *SYSCALL:* This special trap instruction is used to invoke the operating system.
- *Move to/from control registers:* `CTCi` and `CFCi` move between the integer registers and control registers.

- *Limited single-precision registers:* Although the 32 floating-point registers can be addressed individually for loads and stores, single-precision operands for floating-point operations can use only the 16 even floating-point registers.
- *Jump/Call not PC-relative:* The 26-bit address of jumps and calls is not added to the PC. It is shifted left 2 bits and replaces the lower 28 bits of the PC. This would only make a difference if the program was located near a 256-MB boundary.
- *Conditional procedure call instructions*: `BGEZAL` saves the return address and branches if the contents of Rs1 is greater than or equal to zero, and `BLTZAL` does the same for less than zero. The purpose of these instructions is to get a PC-relative call.

There is no specific provision in the MIPS architecture for floating-point execution to proceed in parallel with integer execution, but the MIPS implementations of floating point allow this to happen by checking to see if arithmetic interrupts are possible early in the cycle; normally interrupts are not possible and integer and floating point operate in parallel (see page A-31 in Appendix A).

MIPS II

With the announcement of the R6000 came a set of extensions to the original MIPS architecture described above. Here are the additions of MIPS II:

- *Interlocked loads:* The MIPS II Assembler need not insert a NOP after a load if there is a dependency on the following instruction, as the hardware will automatically stall.
- *Branch likely:* Equivalent to the SPARC annulled branches, this instruction executes the instruction in the delay slot only if the branch is taken.
- *Load double floating point and store double floating point:* MIPS II takes a single instruction to load or store double-precision floating-point numbers.
- *`SQRT`:* Single- and double-precision floating-point square root are added to the floating-point operations.
- *Conditional trap instructions:* These match the conditional branch instructions, except they are not delayed: When the trap is taken, the following instruction is **not** executed. These instructions are useful for range checking, popular in Ada.

E.6 Instructions Unique to SPARC

Register Windows

The primary unique feature of SPARC is register windows (pages 450–453 of Chapter 8), used to reduce the register save/restore overhead of procedure calls and returns. SPARC can have between 2 and 32 windows, using 8 registers each for the globals, locals, incoming parameters, and outgoing parameters (see Figure 8.34 page 452.) (Given each window has 16 unique registers, an implementation of SPARC can have as few as 40 physical registers and as many as 520, although most have 128 to 136, so far.) Rather than tie window changes with call and return instructions, SPARC has the separate instructions `SAVE` and `RESTORE`. `SAVE` is used to "save" the caller's window by pointing to the next window of registers in addition to performing an add instruction. The trick is that the source registers are from the caller's window of the addition operation while the destination register is in the callee's window. SPARC compilers typically use this instruction for changing the stack pointer to allocate local variables in a new stack frame. `RESTORE` is the inverse of `SAVE`, bringing back the caller's window while acting as an add instruction, with the source registers from the callee's window and the destination register in the caller's window. This automatically deallocates the stack frame. Compilers can also make use of it for generating the callee's final return value. Unlike earlier register window architectures, SPARC uses a Window Invalid Mask, which is used in real-time applications, that allows the windows to be partitioned between different processes.

Another data transfer feature is alternate space option for loads and stores. This simply allows the memory system to identify memory accesses to input/output devices, or to control registers for devices such as the cache and memory-management unit.

Support for LISP and Smalltalk

The primary remaining arithmetic feature is tagged addition and subtraction. The designers of SPARC spent some time thinking about languages like LISP and Smalltalk, and this influenced some of the features of SPARC already discussed: register windows, conditional trap instructions, calls with 32-bit instruction addresses, and multiword arithmetic (see Taylor [1986] and Ungar [1984]). A small amount of support is offered for tagged data types with operations for addition, subtraction, and hence comparison. The two least significant bits indicate whether the operand is an integer (coded as 00), so `TADDcc` and `TSUBcc` set the overflow bit if either operand is not tagged as integer or if the result is too large. A subsequent conditional branch or trap instruction can decide what to do. (If the operands are not integers, software recovers the operands, checks the

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 737

types of the operands, and invokes the correct operation based on those types.) Two other versions of these instructions make the conditional trap unnecessary, as TADDccTV and TSUBccTV trap if the overflow is set. It turns out that the misaligned memory access trap can also be put to use for tagged data, since loading from a pointer with the wrong tag can be an invalid access. Figure E.10 shows both types of tag support.

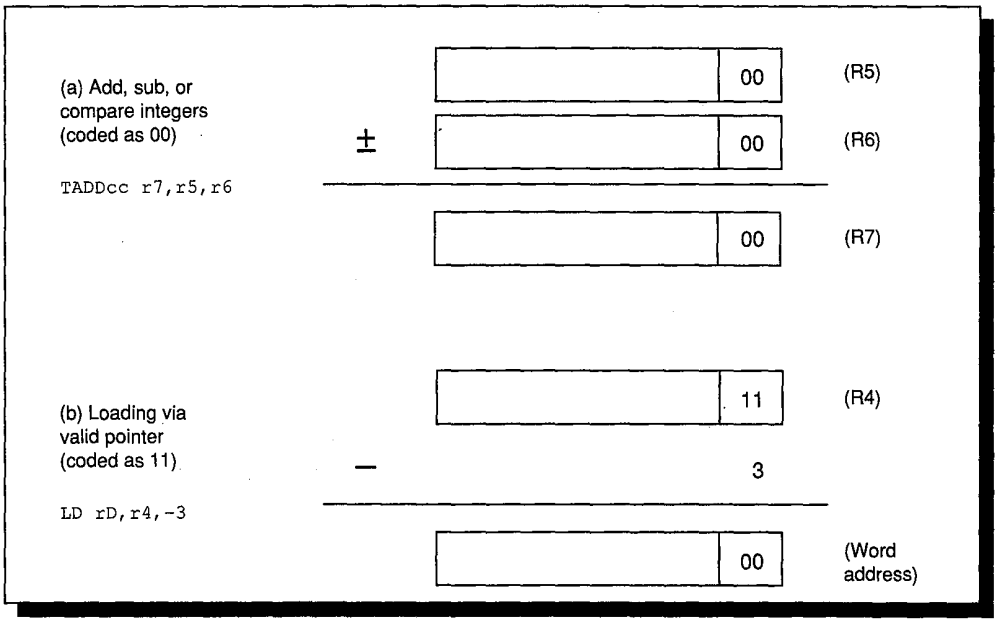

FIGURE E.10 SPARC uses the two least significant bits to encode different data types for the tagged arithmetic instructions. (a) shows integer arithmetic, which takes a single cycle as long as the operands and the result are integers. (b) shows that the misaligned trap can be used to catch invalid memory accesses, such as trying to use an integer as a pointer. For languages with paired data like LISP, an offset of –3 can be used to access the even word of a pair (CAR) and +1 can be used for the odd word of a pair (CDR).

Overlapped Integer and Floating-Point Operations

SPARC allows floating-point instructions to overlap execution with integer instructions. To recover from an interrupt during such a situation, SPARC has a queue of pending floating-point instructions and their addresses. STDFQ allows the processor to empty the queue. The second floating-point feature is the inclusion of floating-point square root instructions FSQRTS and FSQRTD.

Remaining Instructions

The remaining unique features of SPARC are:

- *JMPL* uses Rd to specify the return address register, so specifying r31 makes it similar to JALR in DLX and specifying r0 makes it like JR.

- *LDSTUB* loads the value of the byte into Rd and then stores FF_{16} into the addressed byte. This instruction can be used to implement a semaphore.
- *LDDC and STDC* provide load double and store double for the coprocessor.
- *UNIMP* causes an unimplemented instruction interrupt. Muchnick [1988] explains how this is used for proper execution of aggregate returning procedures in C.

Finally, SPARC includes opcodes for instructions that are emulated in software on early implementations. SPARC application programs generally call dynamically linked library routines to perform these operations, but the opcodes would result in a trap if executed. The instructions are:

- *Signed and unsigned integer multiply and divide*, with both operands and the results being integer registers. The extra 32 bits of a product and the 32-bit remainder of a divide are placed in the Y register.
- *Quadruple precision floating-point arithmetic*, allowing the floating-point registers to act as eight 128-bit registers.
- *Multiple precision floating-point results for multiply*, meaning two single-precision operands can result in a double-precision product and two double-precision operands can result in a quadruple-precision product. These instructions can be useful in complex arithmetic and some models of floating-point calculations.

E.7 Instructions Unique to M88000

The most distinguishing feature of the M88000 is the single set of 32 registers for both integer and floating-point operations. This simplifies the instruction set at the cost of fewer registers for floating-point programs.

Bit Instructions

The next feature unique to the M88000 is a full set of bit-field instructions, shown in Figure E.11 (page E-18). (While we usually number the most significant bit 0, in this table we follow Motorola's notation, which numbers the most significant bit 31 and the least significant bit 0.) Bit-field instructions need an extra operand to specify the width of the field in addition to the destination register, source register, and beginning of the bit field. This 5-bit width field is located next to the bit field in source 2. The M88000 encodes a width of 0 to mean the full 32-bit value, hence the traditional shift instructions (`SLL`, `SRL`, `SRA`) are simply the corresponding bit-field instructions (`MAK`, `EXTU`, `EXT`) with 0 in the width field.

Name	Instruction	Notation
CLR	Clear bit field	$Rd_{(o+w)..(o+1)} \leftarrow 0^w$
SET	Set bit field	$Rd_{(o+w)..(o+1)} \leftarrow 1^w$
EXT	Extract signed bit field	if (w==0) {Rd $\leftarrow Rs1_{31}{}^{o}$ ## (Rs1 >> o) } else {Rd $\leftarrow (Rs1_{(o+w)})^{o}$ ## $(Rs1_{(o+w)..(o+1)}$ >> o) }
EXTU	Extract unsigned bit field	if (w==0) {Rd $\leftarrow 0^{o}$ ## (Rs1 >> o) } else {Rd $\leftarrow 0^{o}$ ## $(Rs1_{(o+w)..(o+1)}$ >> o) }
MAK	Make bit field	if (w==0) {Rd $\leftarrow$ Rs1 << o} else {$Rd_{(o+w)..(o+1)} \leftarrow Rs1_{(w-1)..0}$ }
ROT	Rotate right	Rd $\leftarrow Rs1_{(o-1)..0}$ ## $Rs1_{31..o}$
FF0	Find first bit clear	for (i=31;$Rs2_i$==0||i<0;i $\leftarrow$i-1); /* loop until = 0*/ if(i<0) {Rd $\leftarrow$ 32} else {Rd $\leftarrow$ i}
FF1	Find first bit set	for (i=31;$Rs2_i$==1||i<0;i $\leftarrow$i-1); /* loop until =1 */ if(i<0) {Rd $\leftarrow$ 32} else {Rd $\leftarrow$ i}

FIGURE E.11 The M88000 bit-field instructions. The bit offset, o, is the least significant five bits of the second operand and the bit-field width, w, is the five bits next to the offset. The subscript notation specifies a bit field while the superscript notation means replicate the bit that many times. Note that in this table, bit 31 refers to the most significant bit, and 0 refers to the least significant bit.

Remaining Instructions

The final unique instructions are load address (LDA), MASK, round to nearest integer (NINT), trap on bounds (TBND), and exchange control register (XCR):

- LDA loads Rd with the effective address rather than the data in memory. The only time this is different from ADDU is for scaled addressing of nonbyte data.
- MASK is simply another case of logical AND immediate: This instruction clears the other half of the word while AND immediate leaves it undisturbed. Thus, ANDI in DLX is arguably closer to MASK than to AND immediate in the M88000.
- NINT differs from INT in that it rounds to the nearest integer no matter how the rounding modes are set (see Appendix A, pages A-16 to A-17).
- TBND traps if Rs1 > Rs2, treating them as unsigned numbers (see page 239 in Chapter 5 for an explanation of how an unsigned comparison can check two signed bounds at once).
- XCR exchanges a control register with an integer register.

In addition to instructions, here are a few features that distinguish the M88000:

- Double-length operations use Rn and Rn+1 rather than an even-odd register pair. This gives the M88000 more flexibility in register allocation, which is important given the lack of floating-point registers.
- The first implementation, the MC88100, allows all multicycle instructions to overlap execution with following instructions unless there is a data hazard (see pages 264–265 in Chapter 6). Also, all floating-point instructions except divide are pipelined, taking just one cycle to issue single-precision operations and two cycles to issue double-precision operations. The 88000 provides a set of shadow registers (see Section 5.6) for floating-point operands to help software handle both precise and imprecise interrupts (see Motorola [1988]).
- There are special data transfers, identified by appending .USR to the instructions, that allow access to the user's data while in supervisor mode.

E.8 Instructions Unique to i860

The i860 has many unique features. Before covering the special extensions for graphics and high-performance floating point, let's cover the traditional areas.

The unique data transfers are for floating point only. The i860 provides 128-bit loads (`FLD.Q`) and stores (`FST.Q`) of pairs of 64-bit floating-point registers. It also provides an optional addressing mode on all floating-point loads and stores: the effective address (sum of Rs1/Const and Rs2) is stored back into Rs2. One unique characteristic is that the i860 seems to run out of opcode bits for load instructions because it uses the least significant bit to distinguish load halfword from load word. This works fine for the register–register format since bit 0 is an opcode extension field in this format, but in register–immediate format this is the least significant bit of the constant field. To avoid crazy addressing problems, this bit is cleared when used as an address. This prevents having an odd value in an index register that is corrected by an odd byte address in the constant field for halfword and word data transfers (see E.10(b) on page E-16 for a reason this is useful.)

The only unique arithmetic logical instruction is a double-length shift-right logical (`SHRD`). Rs1 and Rs2 are shifted right as a pair and then the 32 least significant bits are placed into Rd. Since there is no room in the instruction to specify the shift amount, `SHRD` uses the shift amount from the last `SHR` instruction. This value is saved in the 5-bit SC field of the program status word. By the way, `SHRD` can be used to perform a 32-bit rotate by having Rs1 and Rs2 specify the same register.

The i860 control instructions include a loop instruction called `BLA`. This instruction both performs an add and a conditional branch. Since it is likely that another instruction in the loop would change the condition code, the i860 has a special loop condition code (`LCC`) just for this instruction. `BLA` performs `Rd` ← `Rs1+Rs2` and branches if `LCC` equals 1. In addition, `BLA` sets the `LCC` for the

next time through the loop if Rs2 ≥ –Rs1 and clears it otherwise. (LCC is set just the opposite of how ADDS sets CC.)

While i860 does not have floating-point divide, it does have a floating-point reciprocal instruction (FRCP). Used with Newton-Raphson iteration (pages A-23–A-24 of Appendix A), this calculates divide that disagrees with the IEEE floating-point standard (IEEE 754) in the 2 least significant bits. Intel offers software to produce the correctly rounded result at twice the cycle count. A similar instruction, FRSQR, calculates a reciprocal step for square root. The floating-point instructions also include 64-bit integer addition and subtraction (FIADD.DD and FISUB.DD) using the floating-point registers.

This covers the unique features in the traditional categories, so let's describe the new categories of the i860.

Graphics Instructions

The graphics or *pixel* instructions of i860 operate on 64 bits of data at a time, with each word representing several pixels. Pixel instructions are intended to be useful in graphics operations such as hidden surface elimination (see page 525 in Chapter 9), distance interpolation, and three-dimensional shading using intensity interpolation. These special-purpose instructions are not simple to understand, so interested readers should refer to the manual for details.

The overview of the operations is that two bits in the program status word determine the size of the pixels in a 64-bit word. Pixels can be 8-, 16-, or 32-bits wide, with each size containing fields representing intensity of the primary colors red, blue, and green. Some pixel instructions work with a 64-bit accumulator called the MERGE register, useful in collecting the results of a series of calculations on pixels. In addition to "merge" instructions (FADDP and FADDZ), the i860 has instructions for z buffers (page 525) that compare two sets of four 16-bit (FZCHKS) or two 32-bit (FZCHKL) values, storing the smaller values in the 64-bit destination register and setting bits indicating which was smaller in the program status word. Pixel-store instructions (PST) then use those bits to selectively store only those pixels that were smaller. Finally, the FORM instruction is used to move the MERGE register into a floating-point register and then clear MERGE.

Pipelined Mode

For higher performance, the i860 offers pipelined versions of all the floating-point and pixel instructions. One model for these instructions is to use them to build vector primitives, allowing procedures to be written to implement vector operations (see Chapter 7). The hope is that existing vectorizing compilers could invoke these more efficient procedures. Another model, used by compilers currently under development at Intel's behest, tries to compile directly into these instructions for both vector and nonvector codes.

In *pipelined mode,* an instruction is launched every cycle, but unlike other pipelined machines, there is no hardware to remember where the results are to be stored. Basically, the instruction issuing at the stage the pipeline completes specifies the destination! There are four independent pipelines in the i860, and each pipeline advances only when the next instruction of that type is executed. Figure E.12 shows the i860 pipelines, the number of pipeline stages, and instructions that advance each pipeline. Thus, the source fields and opcode specify the operation to be launched while the destination field specifies the register to be loaded by an instruction of the same type that is in the final stage at this cycle.

Pipeline	No. of Stages	Instructions using pipeline
FP multiplier	3 (single operands) 2 (double operands)	PFMUL
FP adder	3	PFADD, PFSUB, PFGT, PFLE, PFEQ, PFIX, PFTRUNC
FP load	3	PFLD
Graphics	1	PFIADD, PFISUB, PFZCHKS, PFZCHKL, PFADDP, PFADDZ, PFORM

FIGURE E.12 i860 pipelines, including the number of pipeline stages and instructions. All adder and multiplier instructions allow single-precision operands with single-precision results (.SS), single operands with double results (.SD), and double-precision operands with double-precision results (.DD). Since the number of stages differs for multiply depending on single or double, Intel recommends not mixing precisions involving multiplication.

For example, look at the sequence below for the floating-point adder pipeline (assume the operands are specified with the result on the left):

```
PFADD.SS    F4, F2, F3        ;Single Prec. Add
PFSUB.DD    F10, F8, F6       ;Double Prec. Sub
PFMUL.DD    F16, F12, F14     ;Double Prec. Mul
PFADD.SS    F19, F17, F18     ;Single Prec. Add
PFADD.SS    F22, F20, F21     ;Single Prec. Add
```

The floating-point adder pipeline is three stages, so the first instruction launches a floating-point add of `F2` and `F3`, but `F4` is loaded from the operation in the adder pipeline launched three instructions earlier. The multiply in this sequence does not advance the adder pipeline, so the third adder pipeline instruction following the first instruction (one subtract and two adds) is the final instruction in the sequence, meaning that `F22` ← `F2+F3`.

The load pipeline has an interesting interaction with the data cache. As long as the data is in the cache, it is fetched from the cache. On a miss the data is fetched from memory, but the cache is not updated with the new data. This

policy prevents operations on large data structures from filling the cache with data that will not be reused and throwing out data that would be reused. The programmer must decide on whether to use scalar loads (FLD) or pipelined loads (PFLD), depending on whether the data is likely to be reused or not.

Scalar instructions will normally empty the pipeline. (The exception is the load pipeline because FLD or LD don't empty it.) Thus, before executing a scalar floating-point instruction there must be a sequence of dummy pipelined instructions that store the results away. For example, there is no pipelined version of the floating-point instruction used for integer multiply (FMLOW), so the pipeline must be drained if an integer multiply is needed during a floating-point calculation.

Summarizing pipelined mode on the i860, the advantages are

- Pipeline control is simple (basically it is done in software).
- It doesn't need many registers, since they are not reserved during the operation.

The disadvantages are:

- Operations must be performed to empty the pipeline.
- The interrupt mechanism is complicated, taking longer to recover the state.
- Sometimes the pipeline is hard to use.
- Code size may mushroom (this has not yet been quantified).

Add/Sub and Multiply

To squeeze even more performance from the floating-point unit, the i860 has pipelined instructions that simultaneously perform an add and multiply (PFAM and PFMAM) or a subtract and multiply (PFSM and PFMSM), advancing the pipelines of both the add and multiply units. Since each instruction needs 4 sources and 2 destinations, the i860 has three registers that can also be used in addition to the three floating-point registers specified in the instruction. The registers KI and KR, optionally loaded from Rs1, can be sources for the multiplier, and register T can be a destination of the multiplier or a source for the adder. The final stage of adder pipeline and multiplier pipeline can also be sources. Four bits in each instruction specify a variety of combinations of the operands and the operations.

Dual Instruction Mode

Finally, the i860 allows an integer and a floating-point instruction to be fetched and executed simultaneously. This long instruction word or superscalar form of operation (pages 318–322 in Chapter 6) is called *dual-instruction mode* in the i860. Simultaneous execution occurs in this mode when the upper instruction of

an aligned doubleword is an integer instruction and the lower is a floating-point instruction with the "D" bit set (bit 9 = 1). Entering or exiting the mode is delayed: When the i860 finds an instruction with the D bit set, it executes one more instruction before entering dual-instruction mode; and, similarly, when the i860 is in dual-instruction mode and finds a D bit not set, it executes one more pair before going to sequential execution.

Clearly, highest performance comes when the i860 is in both dual-instruction and pipelined modes.

E.9 Concluding Remarks

This appendix covers the addressing modes, instruction formats, and all instructions found in four recent architectures. While the later sections concentrate on the differences, it would not be possible to cover four architectures in these few pages if there were not so many similarities. In fact, we would guess that more than 90% of the instructions executed for any of these architectures would be found in Figure E.3 (page E-3). To illustrate this homogeneity, Figure E.13 gives a summary for four architectures from the 1970s similar to Figure E.1 (page E-2). (Imagine trying to write a single appendix in this style for those architectures.) In the history of computing, there has never been such widespread agreement on computer architecture.

	IBM 360/370	Intel 8086	Motorola 68000	DEC VAX
Date announced	1964/1970	1978	1980	1977
Instruction size(s) (bits)	16, 32, 48	8,16,24,32, 40,48	16,32,48, 64,80	8,16,24,32, ..., 432
Addressing (size, model)	24 bits, flat	4+16 bits, segmented	24 bits, flat	32 bits, flat
Data aligned?	Yes 360/ No 370	No	16-bit aligned	No
Data addressing modes	4	5	9	≥ 14
Protection	Page	None	Optional	Page
Page size	4 KB	--	0.25 to 32 KB	0.5 KB
I/O	Opcode	Opcode	Memory mapped	Memory mapped
Integer registers (size, model, number)	16 GPR x 32 bits	8 dedicated data x 16 bits	8 data & 8 address x 32 bits	15 GPR x 32 bits
Separate floating-point registers	4 x 64 bits	Optional: 8 x 80 bits	Optional: 8 x 80 bits	0
Floating-point format	IBM	IEEE 754 single, double, extended	IEEE 754 single, double, extended	DEC

FIGURE E.13 Summary of four 1970s architectures. Unlike the architectures in Figure E.1 (page E-2), there is little agreement between these architectures in any category. (See Chapter 4 for more details on the 370, 8086, and VAX.)

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 745

This style of architectures cannot remain static, however. One hard lesson is that address space must grow, so the 32-bit size of all these architectures must expand for them to survive. In terms of their implementation, we expect all to offer superscalar execution of 2 to 4 instructions per cycle. The hardware technology will go beyond the current CMOS VLSI and ECL to BiCMOS, and possibly even Gallium Arsenide. Our guess is that all of them will grow beyond the current market of workstations and peripheral controllers to minicomputers, mainframes, and even supercomputers, with increasing numbers of processors per computer class.

E.10 References

INTEL [1989]. *i860 64-Bit Microprocessor Programmer's Reference Manual.*

KANE, G. [1988]. *MIPS RISC Architecture,* Prentice-Hall, Englewood Cliffs, N. J.

MOTOROLA [1988]. *MC88100 RISC Microprocessor User's Manual.*

MAGENHEIMER, D. J., L. PETERS, K. W. PETTIS AND D. ZURAS [1988]. "Integer multiplication and division on the HP Precision Architecture," *IEEE Trans. on Computers,* 37:8, 980–990.

MUCHNICK, S. S. [1988]. "Optimizing compilers for SPARC," *Sun Technology* (Summer) 1:3, 64–77.

SUN MICROSYSTEMS [1989]. *The SPARC Architectural Manual,* Version 8, Part No. 800-1399-09, August 25, 1989.

TAYLOR, G., P. HILFINGER, J. LARUS, D. PATTERSON, AND B. ZORN [1986]. "Evaluation of the SPUR LISP architecture," *Proc. 13th Symposium on Computer Architecture (*June), Tokyo.

UNGAR, D., R. BLAU, P. FOLEY, D. SAMPLES, AND D. PATTERSON [1984]. "Architecture of SOAR: Smalltalk on a RISC," *Proc. 11th Symposium on Computer Architecture* (June), Ann Arbor, Mich., 188–197.

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 746

References

The following is a compilation of all the references listed in the reference section of each chapter. The page number of where each reference appears in the book is in parentheses after the reference.

ADAMS, T. AND R. ZIMMERMAN [1989]. "An analysis of 8086 instruction set usage in MS DOS programs," *Proc. Third Symposium on Architectural Support for Programming Languages and Systems* (April) Boston, 152–161. (p. 188)

AGARWAL, A. [1987]. *Analysis of Cache Performance for Operating Systems and Multiprogramming*, Ph.D. Thesis, Stanford Univ., Tech. Rep. No. CSL-TR-87-332 (May). (p. 487)

AGARWAL, A., R. L. SITES, AND M. HOROWITZ [1986]. "ATUM: A new technique for capturing address traces using microcode," *Proc. 13th Annual Symposium on Computer Architecture* (June 2–5), Tokyo, Japan, 119–127. (p. 486)

AGERWALA, T. AND J. COCKE [1987]. "High performance reduced instruction set processors," IBM Tech. Rep. (March). (p. 340)

ALEXANDER, W. G. AND D. B. WORTMAN [1975]. "Static and dynamic characteristics of XPL programs," *Computer* 8:11 (November) 41–46. (pp. 130, 187)

ALLIANT COMPUTER SYSTEMS CORP. [1987]. *Alliant FX/Series: Product Summary* (June), Acton, Mass. (p. 395)

ALMASI, G. S. AND A. GOTTLIEB [1989]. *Highly Parallel Computing,* Benjamin/Cummings, Redwood City, Calif. (p. 589)

AMDAHL, G. M. [1967]. "Validity of the single processor approach to achieving large scale computing capabilities," *Proc. AFIPS Spring Joint Computer Conf.* 30, Atlantic City, N. J. (April) 483–485. (pp. 26, 588)

AMDAHL, G. M., G. A. BLAAUW, AND F. P. BROOKS, JR. [1964]. "Architecture of the IBM System/360," *IBM J. Research and Development* 8:2 (April) 87–101. (pp. 127, 186)

ANDERSON, D. W., F. J. SPARACIO, AND R. M. TOMASULO [1967]. "The IBM 360 Model 91: Machine philosophy and instruction handling," *IBM J. of Research and Development* 11:1 (January) 8–24. (p. 339)

ANDERSON, S. F., J. G. EARLE, R. E. GOLDSCHMIDT, AND D. M. POWERS [1967]. "The IBM System/360 Model 91: Floating-point execution unit," *IBM J. Research and Development* 11, 34–53. Reprinted in [Swartzlander 1980]. (p. A-59)

ANDREWS, G. R. AND F. B. SCHNEIDER [1983]. "Concept and notations for concurrent programming," *Computing Surveys* 15:1 (March) 3–43. (p. 590)

ANON ET AL. [1985]. "A measure of transaction processing power," Tandem Tech. Rep. TR 85.2. Also appeared in *Datamation*, April 1, 1985. (p. 511)

ARCHIBALD, J. AND J.-L. BAER [1986]. "Cache coherence protocols: Evaluation using a multiprocessor simulation model," *ACM Trans. on Computer Systems* 4:4 (November) 273–298. (p. 487)

ATANASOFF, J. V. [1940]. "Computing machine for the solution of large systems of linear equations," Internal Report, Iowa State University. (p. 24)

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 747

ATKINS, D. E. [1968]. "Higher-radix division using estimates of the divisor and partial remainders," *IEEE Trans. on Computers* C-17:10, 925–934. Reprinted in [Swartzlander 1980]. (p. A-60)

BAER, J.-L. AND E.-H. WANG [1988]. "On the inclusion property for multi-level cache hierarchies," *Proc. 15th Annual Symposium on Computer Architecture* (May–June), Honolulu, 73–80. (p. 487)

BAKOGLU, H. B., G. F. GROHOSKI, L. E. THATCHER, J. A. KAHLE, C. R. MOORE, D. P. TUTTLE, W. E. MAULE, W. R. HARDELL, D. A. HICKS, M. NGUYEN PHU, R. K. MONTOYE, W. T. GLOVER , AND S. DHAWAN [1989]. "IBM second-generation RISC machine organization," *Proc. Int'l Conf. on Computer Design, IEEE* (October) Rye, N.Y., 138–142. (p. 340)

BANERJEE, U. [1979]. *Speedup of Ordinary Programs*, Ph.D. Thesis, Dept. of Computer Science, Univ. of Illinois at Urbana-Champaign (October). (p. 395)

BARTON, R. S. [1961]. "A new approach to the functional design of a computer," *Proc. Western Joint Computer Conf.*, 393–396. (p. 127)

BASHE, C. J., L. R. JOHNSON, J. H. PALMER, AND E. W. PUGH [1986]. *IBM's Early Computers*, MIT Press, Cambridge, Mass. (p. 561)

BASHE, C. J., W. BUCHHOLZ, G .V. HAWKINS, J .L. INGRAM, AND N. ROCHESTER [1981]. "The architecture of IBM's early computers," *IBM J. of Research and Development* 25:5 (September) 363–375. (p. 561)

BATCHER, K. E. [1974]. "STARAN parallel processor system hardware," *Proc. AFIPS National Computer Conf.*, 405–410. (p. 590)

BELL , C. G. AND W. D. STRECKER [1976]. "Computer structures: What have we learned from the PDP-11?," *Proc. Third Annual Symposium on Computer Architecture* (January), Pittsburgh, Penn., 1–14. (p. 485)

BELL, C. G. [1984]. "The mini and micro industries," *IEEE Computer* 17:10 (October) 14–30. (p. 27)

BELL, C. G. [1985]. "Multis: A new class of multiprocessor computers," *Science* 228 (April 26) 462–467. (p. 589)

BELL, C. G. [1989]. "The future of high performance computers in science and engineering," *Comm. ACM* 32:9 (September) 1091–1101. (p. 590)

BELL, C. G. AND A. NEWELL, [1971]. *Computer Structures: Readings and Examples*, McGraw-Hill, New York. (p. A-58)

BELL, C. G., J. C. MUDGE, AND J. E. MCNAMARA [1978]. *A DEC View of Computer Engineering*, Digital Press, Bedford, Mass. (p. 80)

BELL, C. G., R. CADY, H. MCFARLAND, B. DELAGI, J. O'LAUGHLIN, R. NOONAN, AND W. WULF [1970]. "A new architecture for mini-computers: The DEC PDP-11," *Proc. AFIPS SJCC*, 657–675. (p. 127)

BERRY, M., D. CHEN, P. KOSS, D. KUCK [1988]. "The Perfect Club benchmarks: Effective performance evaluation of supercomputers," CSRD Report No. 827 (November), Center for Supercomputing Research and Development, University of Illinois at Urbana-Champaign. (p. 80)

BIRMAN, M., G. CHU, L. HU, J. MCLEOD, N. BEDARD, F. WARE, L. TORBAN, AND C. M. LIM [1988]. "Design of a high-speed arithmetic datapath," *Proc. ICCD: VLSI Computers and Processors*, 214–216. (p. A-53)

BLAKKEN, J. [1983]. "Register windows for SOAR," in *Smalltalk On A RISC: Architectural Investigations*, Proc. of CS 292R (April) 126–140. (p. 451)

BLOCH, E. [1959]. "The engineering design of the Stretch computer," *Proc. Fall Joint Computer Conf.*, 48–59. (p. 338)

BORRILL, P. L. [1986]. "32-bit buses—An objective comparison," *Proc. Buscon 1986 West*, San Jose, Calif., 138–145. (p. 533)

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 748

BOUKNIGHT, W. J, S. A. DENEBERG, D. E. MCINTYRE, J. M. RANDALL, A. H. SAMEH, AND D. L. SLOTNICK [1972]. "The Illiac IV system," *Proc. IEEE* 60:4, 369–379. Also appears in D. P. Siewiorek, C. G. Bell, and A. Newell, *Computer Structures: Principles and Examples* (1982), 306–316. (p. 570)

BRADY, J. T. [1986]. "A theory of productivity in the creative process," *IEEE CG&A* (May) 25–34. (p. 560)

BRENT, R. P. AND H. T. KUNG [1982] "A regular layout for parallel adders," *IEEE Trans. on Computers* C-31, 260–264. (p. A-59)

BRODERSEN, R. W. [1989]. "Evolution of VLSI signal-processing circuits," *Proc. Decennial Caltech Conf. on VLSI* (March) 43–46, The MIT Press, Pasadena, Calif. (p. 590)

BUCHER, I. Y. [1983]. "The computational speed of supercomputers," *Proc. SIGMETRICS Conf. on Measuring and Modeling of Computer Systems,* ACM (August) 151–165. (p. 395)

BUCHER, I. Y. AND A. H. HAYES [1980]. "I/O Performance measurement on Cray-1 and CDC 7000 computers," *Proc. Computer Performance Evaluation Users Group, 16th Meeting*, NBS 500-65, 245–254. (p. 562)

BUCHOLTZ, W. [1962]. *Planning a Computer System: Project Stretch,* McGraw-Hill, New York. (p. 338)

BURKS, A. W., H. H. GOLDSTINE, AND J. VON NEUMANN [1946]. "Preliminary discussion of the logical design of an electronic computing instrument," Report to the U.S. Army Ordnance Department, p. 1; also appears in *Papers of John von Neumann,* W. Aspray and A. Burks, eds., The MIT Press, Cambridge, Mass. and Tomash Publishers, Los Angeles, Calif., 1987, 97–146. (p. 24)

CALLAHAN, D., J. DONGARRA, AND D. LEVINE [1988]. "Vectorizing compilers: A test suite and results," *Supercomputing '88,* ACM/IEEE (November), Orlando, Fla., 98–105. (p. 377)

CASE, R. P. AND A. PADEGS [1978]. "The architecture of the IBM System/370," *Comm. ACM* 21:1, 73–96. Also appears in D. P. Siewiorek, C. G. Bell, and A. Newell, *Computer Structures: Principles and Examples* (1982), McGraw-Hill, New York, 830–855. (pp. 186, 485)

CENSIER, L. M. AND P. FEAUTRIER [1978]. "A new solution to the coherence problem in multicache systems," *IEEE Trans. on Computers* C-27:12 (December) 1112–1118. (p. 487)

CHAITIN, G. J., M. A. AUSLANDER, A. K. CHANDRA, J. COCKE, M. E. HOPKINS, AND P. W. MARKSTEIN [1982]. "Register allocation via coloring," *Computer Languages* 6, 47–57. (p. 130)

CHARLESWORTH, A. E. [1981]. "An approach to scientific array processing: The architecture design of the AP-120B/FPS-164 family," *Computer* 14:12 (December) 12–30. (p. 340)

CHEN, P. [1989]. *An Evaluation of Redundant Arrays of Inexpensive Disks Using an Amdahl 5890,* M. S. Thesis, Computer Science Division, Tech. Rep. UCB/CSD 89/506. (p. 507)

CHEN, S. [1983]. "Large-scale and high-speed multiprocessor system for scientific applications," *Proc. NATO Advanced Research Work on High Speed Computing* (June); also in K. Hwang, ed., "Supercomputers: Design and applications," *IEEE* (August) 1984. (p. 394)

CHEN, T. C. [1980]. "Overlap and parallel processing" in *Introduction to Computer Architecture,* H. Stone, ed., Science Research Associates, Chicago, 427–486. (p. 339)

CHOW, F. C. [1983]. *A Portable Machine-Independent Global Optimizer—Design and Measurements,* Ph. D. Thesis, Stanford Univ. (December). (p. 130)

CHOW, F. C. AND J. L. HENNESSY [1984]. "Register allocation by priority-based coloring," *Proc. SIGPLAN '84 Compiler Construction (ACM SIGPLAN Notices* 19:6, June) 222–232. (p. 130)

CHOW, F., M. HIMELSTEIN, E. KILLIAN, AND L. WEBER [1986]. "Engineering a RISC compiler system," *Proc. COMPCON* (March), San Francisco, 132–137. (p. 197)

CLARK, D. W. [1983]. "Cache performance of the VAX-11/780," *ACM Trans. on Computer Systems* 1:1, 24–37. (p. 486)

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 749

CLARK, D. W. [1987]. "Pipelining and performance in the VAX 8800 processor," *Proc. Second Conf. on Architectural Support for Programming Languages and Operating Systems,* IEEE/ACM (March), Palo Alto, Calif., 173–177. (p. 272)

CLARK, D. W. AND H. LEVY [1982]. "Measurement and analysis of instruction set use in the VAX-11/780," *Proc. Ninth Symposium on Computer Architecture* (April), Austin, Tex., 9–17. (p. 188)

CLARK, D. W. AND J. S. EMER [1985]. "Performance of the VAX-11/780 translation buffer: Simulation and measurement," *ACM Trans. on Computer Systems* 3:1, 31–62. (p. 486)

CLARK, D. W. AND W. D. STRECKER [1980]. "Comments on 'the case for the reduced instruction set computer', " *Computer Architecture News* 8:6 (October) 34–38. (p. 130)

CLARK, D. W., P. J. BANNON, AND J. B. KELLER [1988]. "Measuring VAX 8800 performance with a histogram hardware monitor," *Proc. 15th Annual Symposium on Computer Architecture* (May–June), Honolulu, Hawaii, 176–185. (pp. 213, 486)

COCKE, J. AND J. T. SCHWARTZ [1970]. *Programming Languages and Their Compilers,* Courant Institute, New York Univ., New York City. (p. 130)

COCKE, J., AND J. MARKSTEIN [1980]. "Measurement of code improvement algorithms," *Information Processing* 80, 221–228. (p. 130)

CODD, E. F. [1962]. "Multiprogramming," in F.L. Alt and M. Rubinoff, *Advances in Computers,* vol. 3, Academic Press, New York, 82. (p. 241)

CODY, W. J. [1988]. "Floating point standards: Theory and practice," in *Reliability in Computing: The Role of Interval Methods in Scientific Computing,* R. E. Moore, (ed.), Academic Press, Boston, Mass., 99–107. (p. A-12)

CODY, W. J., J. T. COONEN, D. M. GAY, K. HANSON, D. HOUGH, W. KAHAN, R. KARPINSKI, J. PALMER, F. N. RIS, AND D. STEVENSON [1984]. "A proposed radix- and word-length-independent standard for floating-point arithmetic," *IEEE Micro* 4:4, 86–100. (p. A-12)

COHEN, D. [1981]. "On holy wars and a plea for peace," *Computer* 14:10 (October) 48–54. (p. 95)

COLWELL, R. P, C. Y. HITCHCOCK, III, E. D. JENSEN, H. M. B. SPRUNT, AND C. P. KOLLAR, [1985]. "Computers, complexity, and controversy," *Computer* 18:9 (September) 8–19. (p. 125)

COLWELL, R. P., R. P. NIX, J. J. O'DONNELL, D. B. PAPWORTH, AND B. K. RODMAN [1987]. "A VLIW architecture for a trace scheduling compiler," *Proc. Second Conf. on Architectural Support for Programming Languages and Operating Systems,* IEEE/ACM (March), Palo Alto, Calif., 180–192. (p. 340)

CONTI, C., D. H. GIBSON, AND S. H. PITKOWSKY [1968]. "Structural aspects of the System/360 Model 85, part I: General organization," *IBM Systems J.* 7:1, 2–14. (pp. 77, 486)

COONEN, J. [1984]. *Contributions to a Proposed Standard for Binary Floating-Point Arithmetic,* Ph.D. Thesis, Univ. of Calif., Berkeley. (p. A-29)

CRAWFORD, J. H AND P. P. GELSINGER [1987]. *Programming the 80386,* Sybex, Alameda, Calif. (pp. 188, 446)

CURNOW, H. J. AND B. A. WICHMANN [1976]. "A synthetic benchmark," *The Computer J.* 19:1. (p.77)

DAVIDSON, E. S. [1971]. "The design and control of pipelined function generators," *Proc. Conf. on Systems, Networks, and Computers,* IEEE (January), Oaxtepec, Mexico, 19–21. (p. 339)

DAVIDSON, E. S., A. T. THOMAS, L. E. SHAR, AND J. H. PATEL [1975]. "Effective control for pipelined processors," *COMPCON, IEEE* (March), San Francisco, 181–184. (p. 339)

DEHNERT, J. C., P. Y.-T. HSU, AND J. P. BRATT [1989]. "Overlapped loop support on the Cydra 5," *Proc. Third Conf. on Architectural Support for Programming Languages and Operating Systems* (April), IEEE/ACM, Boston, 26–39. (p. 340)

DEROSA, J., R. GLACKEMEYER, AND T. KNIGHT [1985]. "Design and implementation of the VAX 8600 pipeline," *Computer* 18:5 (May) 38–48. (p. 328)

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 750

DEWITT, D. J., R. FINKEL, AND M. SOLOMON [1984]. "The CRYSTAL multicomputer: Design and implementation experience, Computer Sciences Tech. Rep. No. 553, University of Wisconsin-Madison, September. (p. 590)

DIGITAL EQUIPMENT CORPORATION [1987]. *Digital Technical J.* 4 (March), Hudson, Mass. (This entire issue is devoted to the VAX 8800 processor.) (p. 341)

DITZEL, D. R. [1981]. "Reflections on the high-level language Symbol computer system," *Computer* 14:7 (July) 55–66. (p. 129)

DITZEL, D. R. AND D. A. PATTERSON [1980]. "Retrospective on high-level language computer architecture," in *Proc. Seventh Annual Symposium on Computer Architecture,* La Baule, France (June) 97–104. (p. 130)

DITZEL, D. R. AND H. R. MCLELLAN [1987]. "Branch folding in the CRISP microprocessor: Reducing the branch delay to zero," *Proc. 14th Symposium on Computer Architecture* (June), Pittsburgh, 2–7. (p. 339)

DITZEL, D. R., AND H. R. MCLELLAN [1982]. "Register allocation for free: The C machine stack cache," *Symposium on Architectural Support for Programming Languages and Operating Systems* (March 1–3), Palo Alto, Calif., 48–56. (p. 487)

DOHERTY, W. J. AND R. P. KELISKY [1979]. "Managing VM/CMS systems for user effectiveness," *IBM Systems J.* 18:1, 143–166. (p. 560)

DONGARRA, J. J. [1986]. "A survey of high performance computers," *COMPCON, IEEE* (March) 8–11. (p. 394)

EARLE, J. G. [1965]. "Latched carry-save adder," *IBM Technical Disclosure Bull.* 7 (March) 909–910. (p. 254)

EGGERS, S. [1989]. *Simulation Analysis of Data Sharing in Shared Memory Multiprocessors* , Ph.D. Thesis, Univ. of California, Berkeley, Computer Science Division Tech. Rep. UCB/CSD 89/501 (April). (p. 487)

ELDER, J., A. GOTTLIEB, C. K. KRUSKAL, K. P. MCAULIFFE, L. RANDOLPH, M. SNIR, P. TELLER, AND J. WILSON [1985]. "Issues related to MIMD shared-memory computers: The NYU Ultracomputer approach," *Proc. 12th Int'l Symposium on Computer Architecture* (June), Boston, Mass., 126–135. (p. 589)

ELLIS, J. R., J. A. FISHER, J. C. RUTTENBERG, AND A. NICHOLAU [1984]. "Parallel processing: A smart compiler and a dumb machine," *Proc. SIGPLAN Conf. on Compiler Construction* (June), Montreal, Canada, 37–47. (p. 340)

ELSHOFF, J. L. [1976]. "An analysis of some commercial PL/I programs," *IEEE Trans. on Software Engineering* SE-2 2 (June) 113–120. (p. 130)

EMER, J. S. AND D. W CLARK [1984]. "A characterization of processor performance in the VAX-11/780," *Proc. 11th Symposium on Computer Architecture* (June), Ann Arbor, Mich., 301–310. (pp. 189, 213, 342, 486)

E•SUN MICROSYSTEMS [1989]. *The SPARC Architectural Manual,* Version 8, Part No. 800-1399-09, August 25, 1989.

FABRY, R. S. [1974]. "Capability based addressing," *Comm. ACM* 17:7 (July) 403–412. (p. 485)

FAZIO, D. [1987]. "It's really much more fun building a supercomputer than it is simply inventing one," *COMPCON, IEEE* (February) 102-105. (p. 394)

FEIERBACK, G AND D. STEVENSON [1979]. "The Illiac-IV," in *Infotech State of the Art Report on Supercomptuers*, Maidenhead, England. This data also appears in D. P. Siewiorek, C. G. Bell, and A. Newell, *Computer Structures: Principles and Examples* (1982), McGraw-Hill, New York, 268–269. (p. 556)

FISHER, J. A. [1983]. "Very long instruction word architectures and ELI-512," *Proc. Temth Symposium on Computer Architecture* (June), Stockholm, Sweden. (p. 340)

FLEMMING, P. J. AND J. J. WALLACE [1986]. "How not to lie with statistics: The correct way to summarize benchmarks results," *Comm. ACM* 29:3 (March) 218–221. (p. 78)

FLYNN, M. J. [1966]. "Very high-speed computing systems," *Proc. IEEE* 54:12 (December) 1901–1909. (pp. 351, 591)

FOLEY, J. D. AND A. VAN DAM [1982]. *Fundamentals of Interactive Computer Graphics*, Addison-Wesley, Reading, Mass. (p. 561)

FOSTER, C. C. AND E. M. RISEMAN [1972]. "Percolation of code to enhance parallel dispatching and execution," *IEEE Trans. on Computers* C-21:12 (December) 1411–1415. (p. 340)

FOSTER, C. C., R. H. GONTER, AND E. M. RISEMAN [1971]. "Measures of opcode utilization," *IEEE Trans. on Computers* 13:5 (May) 582–584. (p. 129)

FRANK, P. D. [1987]. "Advances in Head Technology," presentation at *Challenges in Winchester Technology* (December 15), Santa Clara Univ. (p. 561)

FRANK, S. J. [1984]. "Tightly coupled multiprocessor systems speed memory access times," *Electronics* 57:1 (January) 164–169. (p. 487)

FREIMAN, C. V. [1961]. "Statistical analysis of certain binary division algorithms," *Proc. IRE* 49:1, 91–103. (p. A-59)

FRIESENBORG, S. E. AND R. J. WICKS [1985]. "DASD expectations: The 3380, 3380-23, and MVS/XA," Tech. Bulletin GG22-9363-02 (July 10), Washington Systems Center. (p. 554)

FULLER, S. H. [1976]. "Price/performance comparison of C.mmp and the PDP-11," *Proc. Third Annual Symposium on Computer Architecture* (Texas, January 19–21), 197–202. (p. 80)

FULLER, S. H. AND W. E. BURR [1977]. "Measurement and evaluation of alternative computer architectures," *Computer* 10:10 (October) 24–35. (p. 78)

GAGLIARDI, U. O. [1973]. "Report of workshop 4–software-related advances in computer hardware," *Proc. Symposium on the High Cost of Software*, Menlo Park, Calif., 99–120. (p. 129)

GAJSKI, D., D. KUCK, D. LAWRIE, AND A. SAMEH [1983]. "CEDAR—A large scale multiprocessor," *Proc. Int'l Conf. on Parallel Processing* (August) 524–529. (p. 589)

GARNER, R., A. AGARWAL, F. BRIGGS, E. BROWN, D. HOUGH, B. JOY, S. KLEIMAN, S. MUNCHNIK, M. NAMJOO, D. PATTERSON, J. PENDLETON, AND R. TUCK [1988]. "Scaleable processor architecture (SPARC)," *COMPCON, IEEE* (March), San Francisco, 278–283. (p. 190)

GEHRINGER, E. F., D. P. SIEWIOREK, AND Z. SEGALL [1987]. *Parallel Processing: The Cm* Experience*, Digital Press, Bedford, Mass. (p. 587)

GIBSON, D. H. [1967]. "Considerations in block–oriented systems design," *AFIPS Conf. Proc.* 30, SJCC, 75–80. (p. 486)

GIBSON, J. C. [1970]. "The Gibson mix," Rep. TR. 00.2043, IBM Systems Development Division, Poughkeepsie, N.Y. (Research done in 1959.) (p. 77)

GOLDBERG, D. [1989]. "Floating-point and computer systems," *Xerox Tech. Rep.* CSL-89-9. A version of this paper will appear in *Computing Surveys*. (p. A-29)

GOLDBERG, I. B. [1967]. "27 bits are not enough for 8-digit accuracy," *Comm. ACM* 10:2, 105–106. (p. A-60)

GOLDSTEIN, S. [1987]. "Storage performance—an eight year outlook," Tech. Rep. TR 03.308-1 (October), Santa Teresa Laboratory, IBM, San Jose, Calif. (p. 561)

GOLDSTINE, H. H. [1972]. *The Computer: From Pascal to von Neumann*, Princeton University Press, Princeton, N.J. (p. 25)

GOODMAN, J. R. [1983]. "Using cache memory to reduce processor memory traffic," *Proc. Tenth Annual Symposium on Computer Architecture* (June 5–7), Stockholm, Sweden, 124–131. (p. 487)

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 752

GOODMAN, J. R. and M.-C. Chiang [1984]. "The use of static column RAM as a memory hierarchy," *Proc. 11th Annual Symposium on Computer Architecture* (June 5–7), Ann Arbor, Mich., 167–174. (p. 488)

GOSLING, J. B. [1980]. *Design of Arithmetic Units for Digital Computers,* Springer-Verlag NewYork, Inc., New York. (p. A-61)

GRAY, W. P. [1989]. Memorandum of Decision, No. C-84-20799-WPG, U.S. District Court for the Northern District of California (February 7, 1989). (p. 244)

GROSS, T. R. [1983]. *Code Optimization of Pipeline Constraints,* Ph.D. Thesis (December), Computer Systems Lab., Stanford Univ. (p. 342)

HALBERT, D. C. AND P. B. KESSLER [1980]. "Windows of overlapping register frames," *CS 292R Final Reports* (June) 82–100. (p. 489)

HAMACHER, V. C., Z. G. VRANESIC, AND S. G. ZAKY [1984]. *Computer Organization,* 2nd ed., McGraw-Hill, New York. (p. A-61)

HAUCK, E. A., AND B. A. DENT [1968]. "Burroughs' B6500/B7500 stack mechanism," *Proc. AFIPS SJCC*, 245–251. (p. 131)

HENLY, M. AND B. MCNUTT [1989]. "DASD I/O characteristics: A comparison of MVS to VM," Tech. Rep. TR 02.1550 (May), IBM, General Products Division, San Jose, Calif. (pp. 80, 562)

HENNESSY, J. [1984]. "VLSI processor architecture," *IEEE Trans. on Computers* C-33:11 (December) 1221–1246. (p. 190)

HENNESSY, J. [1985]. "VLSI RISC processors," *VLSI Systems Design* VI:10 (October) 22–32. (p. 191)

HENNESSY, J. L. AND T. R. GROSS [1983]. "Postpass code optimization of pipeline constraints," *ACM Trans. on Programming Languages and Systems* 5:3 (July) 422–448. (p. 342)

HENNESSY, J., N. JOUPPI, F. BASKETT, AND J. GILL [1981]. "MIPS: A VLSI processor architecture," *Proc. CMU Conf. on VLSI Systems and Computations* (October), Computer Science Press, Rockville, Md. (p. 191)

HENNESSY, J. L., N. JOUPPI, F. BASKETT, T. R. GROSS, AND J. GILL [1982]. "Hardware/software tradeoffs for increased performance," *Proc. Symposium on Architectural Support for Programming Languages and Operating Systems* (March), 2–11. (p. 131)

HENNESSY, J. [1984]. "VLSI processor architecture," *IEEE Trans. on Computers* C-33:11 (December) 1221–1246. (p. 189)

HILL, M. D. [1987]. *Aspects of Cache Memory and Instruction Buffer Performance,* Ph.D. Thesis, Univ. of California at Berkeley Computer Science Division, Tech. Rep. UCB/CSD 87/381 (November). (p. 489)

HILL, M. D. [1988]. "A case for direct mapped caches," *Computer* 21:12 (December) 25–40. (p. 489)

HILLIS, W. D. [1985]. *The Connection Machine,* The MIT Press, Cambridge, Mass. (p. 591)

HINTZ, R. G. AND D. P. TATE [1972]. "Control data STAR-100 processor design," *COMPCON, IEEE* (September) 1–4. (p. 396)

HOCKNEY, R. W. AND C. R. JESSHOPE [1988]. *Parallel Computers-2, Architectures, Programming and Algorithms,* Adam Hilger Ltd., Bristol, England and Philadelphia. (p. 591)

HOLLAND, J. H. [1959]. "A universal computer capable of executing an arbitrary number of subprograms simultaneously," *Proc. East Joint Computer Conf.* 16, 108–113. (p. 591)

HOLLINGSWORTH, W., H. SACHS AND A. J. SMITH [1989]. "The Clipper processor: Instruction set architecture and implementation," *Comm. ACM* 32:2 (February), 200–219. (p. 80)

HORD, R. M. [1982]. *The Illiac-IV, The First Supercomputer,* Computer Science Press, Rockville, Md. (p. 591)

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 753

HOWARD, J. H. ET AL. [1988]. "Scale and performance in a distributed file system," *ACM Trans. on Computer Systems* 6:1, 51–81. (p. 512)

HUGUET, M. AND T. LANG [1985]. "A reduced register file for RISC architectures," *Computer Architecture News* 13:4 (September) 22–31. (p. 489)

HWANG, K. [1979]. *Computer Arithmetic: Principles, Architecture, and Design,* Wiley, New York. (p. A-61)

HWU, W.-M. AND Y. PATT [1986]. "HPSm, a high performance restricted data flow architecture having minimum functionality," *Proc. 13th Symposium on Computer Architecture* (June), Tokyo, 297–307. (p. 339)

IBM [1982]. *The Economic Value of Rapid Response Time*, GE20-0752-0 White Plains, N.Y., 11–82. (p. 560)

IEEE [1985]. "IEEE standard for binary floating-point arithmetic," *SIGPLAN Notices* 22:2, 9–25. (p. A-12)

IMPRIMIS [1989]. "Imprimis Product Specification, 97209 Sabre Disk Drive IPI-2 Interface 1.2 GB," Document No. 64402302 (May). (p. 558)

INTEL [1989]. *i860 64-Bit Microprocessor Programmer's Reference Manual.* (E-24)

JORDAN, K. E. [1987]. "Performance comparison of large-scale scientific computers: Scalar mainframes, mainframes with vector facilities, and supercomputers," *Computer* 20:3 (March) 10–23. (p. 395)

JOUPPI N. P. AND D. W. WALL [1989]. "Available instruction-level parallelism for superscalar and superpipelined machines," *Proc. Third Conf. on Architectural Support for Programming Languages and Operating Systems,* IEEE/ACM (April), Boston, 272–282. (p. 340)

KAHAN, W. [1968]. "7094-II system support for numerical analysis," *SHARE Secretarial Distribution* SSD-159. (p. A-60)

KAHANER, D. K. [1988]. "Benchmarks for 'real' programs," *SIAM News* (November). (p. A-57)

KAHN, R. E. [1972]. "Resource-sharing computer communication networks," *Proc. IEEE* 60:11 (November) 1397–1407. (p. 561)

KANE, G. [1986]. *MIPS R2000 RISC Architecture,* Prentice Hall, Englewood Cliffs, N.J. (p. 190)

KANE, G. [1988]. *MIPS RISC Architecture,* Prentice-Hall, Englewood Cliffs, N. J. (E-24)

KATZ, R. H., D. A. PATTERSON, AND G. A. GIBSON [1990]. "Disk system architectures for high performance computing," *Proc. IEEE* 78:2 (February). (p. 561)

KATZ, R. H., S. EGGERS, D. A. WOOD, C. PERKINS, AND R. G. SHELDON [1985]. "Implementing a cache consistency protocol," *Proc. 12th Annual Symposium on Computer Architecture*, 276–283. (p. 487)

KELLER R. M. [1975]. "Look-ahead processors," *ACM Computing Surveys* 7:4 (December) 177–195. (p. 339)

KELLY, E. [1988]. "'SCRAM Cache' in Sun-4/110 beats traditional caches," *Sun Technology* 1:3 (Summer) 19–21. (p. 487)

KILBURN, T., D. B. G. EDWARDS, M. J. LANIGAN, F. H. SUMNER [1962]. "One-level storage system," *IRE Transactions on Electronic Computers* EC-11 (April) 223–235. Also appears in D. P. Siewiorek, C. G. Bell, and A. Newell, *Computer Structures: Principles and Examples* (1982), McGraw-Hill, New York, 135–148. (pp. 26, 487)

KIM, M. Y. [1986]. "Synchronized disk interleaving," *IEEE Trans. on Computers* C-35:11 (November). (p. 561)

KNUTH, D. [1981]. *The Art of Computer Programming,* vol II, 2nd ed., Addison-Wesley, Reading, Mass. (p. A-61)

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 754

KNUTH, D. E. [1971]. "An empirical study of FORTRAN programs," *Software Practice and Experience,* Vol. 1, 105–133. (p. 27)

KOGGE, P. M. [1981]. *The Architecture of Pipelined Computers,* McGraw-Hill, New York. (pp. 339, A-44)

KOHN, L. AND S.-W. FU, [1989]. "A 1,000,000 transistor microprocessor," *IEEE Int'l Solid-State Circuits Conf.,* 54–55. (p. A-19)

KROFT, D. [1981]. "Lockup-free instruction fetch/prefetch cache organization," *Proc. Eighth Annual Symposium on Computer Architecture* (May 12–14), Minneapolis, Minn., 81–87. (p. 487)

KUCK, D., P. P. BUDNIK, S.-C. CHEN, D. H. LAWRIE, R. A. TOWLE, R. E. STREBENDT, E. W. DAVIS, JR., J. HAN, P. W. KRASKA, Y. MURAOKA [1974]. "Measurements of parallelism in ordinary FORTRAN programs," *Computer* 7:1 (January) 37–46. (p. 395)

KUHN, R. H. AND D. A. PADUA, EDS. [1981]. *Tutorial on Parallel Processing*, IEEE. (p. 590)

KUNG, H. T. [1982]. "Why systolic architectures?," *IEEE Computer* 15:1, 37–46. (p. 590)

KUNKEL, S. R. AND J. E. SMITH [1986]. "Optimal pipelining in supercomputers," *Proc. 13th Symposium on Computer Architecture* (June), Tokyo, 404–414. (p. 339)

LAM, M. [1988]. "Software pipelining: An effective scheduling technique for VLIW machines," *SIGPLAN Conf. on Programming Language Design and Implementation,* ACM (June), Atlanta, Ga., 318–328. (p. 340)

LAMPSON, B. W. [1982]. "Fast procedure calls," *Symposium on Architectural Support for Programming Languages and Operating Systems* (March 1–3), Palo Alto, Calif., 66–75. (p. 487)

LARSON, JUDGE E. R. [1973]. "Findings of Fact, Conclusions of Law, and Order for Judgment," File No. 4–67, Civ. 138, *Honeywell v. Sperry Rand and Illinois Scientific Development,* U.S. District Court for the District of Minnesota, Fourth Division (October 19). (p. 24)

LEE, R. [1989]. "Precision architecture," *Computer* 22:1 (January) 78–91. (p. 190)

LEINER, A. L. [1954]. "System specifications for the DYSEAC," *J. ACM* 1:2 (April) 57–81. (p. 561)

LEINER, A. L. AND S. N. ALEXANDER [1954]. "System organization of the DYSEAC," *IRE Trans. of Electronic Computers* EC-3:1 (March) 1–10. (p. 561)

LEVY, H. M. AND R. H. ECKHOUSE, JR. [1989]. *Computer Programming and Architecture: The VAX*, 2nd ed., Digital Press, Bedford, Mass. 358–372. (pp. 188, 243)

LEVY, J. V. [1978]. "Buses: The skeleton of computer structures," in *Computer Engineering: A DEC View of Hardware Systems Design*, C. G. Bell, J. C. Mudge, and J. E. McNamara, eds., Digital Press, Bedford, Mass. (p. 561)

LINCOLN, N. R. [1982]. "Technology and design tradeoffs in the creation of a modern supercomputer," *IEEE Trans. on Computers* C-31:5 (May) 363–376. (p. 393)

LIPOVSKI, A. G. AND A. TRIPATHI [1977]. "A reconfigurable varistructure array processor," *Proc. 1977 Int'l Conf. of Parallel Processing* (August), 165–174. (p. 590)

LIPTAY, J. S. [1968]. "Structural aspects of the System/360 Model 85, part II: The cache," *IBM Systems J.* 7:1, 15–21. (p. 486)

LOVETT, T. AND S. THAKKAR [1988]. "The Symmetry multiprocessor system," *Proc. 1988 Int'l Conf. of Parallel Processing*, University Park, Pennsylvania, 303–310. (p. 589)

LUBECK, O., J. MOORE, AND R. MENDEZ [1985]. "A benchmark comparison of three supercomputers: Fujitsu VP-200, Hitachi S810/20, and CRAY X-MP/2," *Computer* 18:12 (December) 10–24. (pp. 75, 395)

LUNDE, A. [1977]. "Empirical evaluation of some features of instruction set processor architecture," *Comm. ACM* 20:3 (March) 143–152. (p. 129)

MABERLY, N. C. [1966]. *Mastering Speed Reading,* New American Library, Inc., New York. (p. 513)

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 755

MAGENHEIMER, D. J., L. PETERS, K. W. PETTIS AND D. ZURAS [1988]. "Integer multiplication and division on the HP Precision Architecture," *IEEE Trans. on Computers,* 37:8, 980–990. (p. E-9)

MAGENHEIMER, D. J., L. PETERS, K. W. PETTIS, AND D. ZURAS, [1988]. "Integer multiplication and division on the HP Precision Architecture," *IEEE Trans. on Computers* 37:8, 980–990. (p. A-11)

MCCALL, K. [1983]. "The Smalltalk-80 benchmarks," *Smalltalk 80: Bits of History, Words of Advice,* G. Krasner, ed., Addison-Wesley, Reading, Mass., 153–174. (p. 451)

MCCREIGHT, E. [1984]. "The Dragon computer system: An early overview," Tech. Rep. Xerox Corp. (September). (p. 487)

MCFARLING, S. [1989]. "Program optimization for instruction caches," *Proc. Third Int'l Conf. on Architectural Support for Programming Languages and Operating Systems* (April 3–6), Boston, Mass., 183–191. (p. 496)

MCFARLING, S. AND J. HENNESSY [1986]. "Reducing the cost of branches," *Proc. 13th Symposium on Computer Architecture* (June), Tokyo, 396–403. (p. 340)

MCKEEMAN, W. M. [1967]. "Language directed computer design," *Proc. 1967 Fall Joint Computer Conf.,* Washington, D.C., 413–417. (p. 128)

MCKEVITT, J., ET AL. [1977]. *8086 Design Report*, internal memorandum. (p. 229)

MCMAHON, F. M. [1986]. "The Livermore FORTRAN kernels: A computer test of numerical performance range," Tech. Rep. UCRL-55745, Lawrence Livermore National Laboratory, Univ. of California, Livermore, Calif. (December). (p. 78)

MEAD, C. AND L. CONWAY [1980]. *Introduction to VLSI Systems,* Addison-Wesley, Reading, Mass. (p. A-59)

MENABREA, L. F. [1842]. "Sketch of the analytical engine invented by Charles Babbage," Bibiothèque Universelle de Genève (October). (p. 589)

METCALFE, R. M. AND D. R. BOGGS [1976]. "Ethernet: Distributed packet switching for local computer networks," *Comm. ACM* 19:7 (July) 395–404. (p. 560)

MEYERS, G. J. [1978]. "The evaluation of expressions in a storage-to-storage architecture," *Computer Architecture News* 7:3 (October), 20–23. (p. 127)

MEYERS, G. J. [1982]. *Advances in Computer Architecture*, 2nd ed., Wiley, N.Y. (p. 129)

MIRANKER, G. S., J. RUBENSTEIN, AND J. SANGUINETTI [1988]. "Squeezing a Cray-class supercomputer into a single-user package," *COMPCON, IEEE* (March) 452–456. (p. 395)

MITCHELL, D. [1989]. "The Transputer: The time is now," *Computer Design*, RISC supplement, 40–41 (November). (p. 570)

MIURA, K. AND K. UCHIDA [1983]. "FACOM vector processing system: VP100/200," *Proc. NATO Advanced Research Work on High Speed Computing* (June); also in K. Hwang, ed., "Supercomputers: Design and applications," *IEEE* (August 1984) 59–73. (p. 394)

MOORE, B., A. PADEGS, R. SMITH, AND W. BUCHOLZ [1987]. "Concepts of the System/370 vector architecture," *Proc. 14th Symposium on Computer Architecture* (June), ACM/IEEE, Pittsburgh, Pa., 282–292. (p. 394)

MORSE, S., B. RAVENAL, S. MAZOR, AND W. POHLMAN [1980]. "Intel Microprocessors—8008 to 8086," *Computer* 13:10 (October). (p. 188)

MOTOROLA [1988]. *MC88100 RISC Microprocessor User's Manual.* (E-19)

MOUSSOURIS, J., L. CRUDELE, D. FREITAS, C. HANSEN, E. HUDSON, S. PRZYBYLSKI, T. RIORDAN, AND C. ROWEN [1986]. "A CMOS RISC processor with integrated system functions," *Proc. COMPCON, IEEE* (March), San Francisco. (p. 189)

MUCHNICK, S. S. [1988]. "Optimizing compilers for SPARC," *Sun Technology* (Summer) 1:3, 64–77. (p. E-9)

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 756

NEWMAN, W. N. AND R. F. SPROULL [1979]. *Principles of Interactive Computer Graphics*, 2nd ed., McGraw-Hill, New York. (p. 561)

NGAI, T-F. AND M. J. IRWIN [1985]. "Regular, area-time efficient carry-lookahead adders," *Proc. Seventh IEEE Symposium on Computer Arithmetic*, 9–15. (p. A-59)

NICHOLAU, A. AND J. A. FISHER [1984]. "Measuring the parallelism available for very long instruction word architectures," *IEEE Trans. on Computers* C-33:11 (November) 968–976. (p. 340)

OUSTERHOUT, J. K. ET AL. [1985]. "A trace-driven analysis of the UNIX 4.2 BSD file system," *Proc. Tenth ACM Symposium on Operating Systems Principles*, Orcas Island, Wash., 15–24. (p. 538)

PADUA, D. AND M. WOLFE [1986]. "Advanced compiler optimizations for supercomputers," *Comm. ACM* 29:12 (December) 1184–1201. (p. 395)

PAPAMARCOS, M. AND J. PATEL [1984]. "A low coherence solution for multiprocessors with private cache memories," *Proc. of the 11th Annual Symposium on Computer Architecture* (June), Ann Arbor, Mich., 348–354. (p. 487)

PATTERSON, D. A. [1983]. "Microprogramming," *Scientific American* 248:3 (March), 36–43. (p. 244)

PATTERSON, D. A. [1985]. "Reduced Instruction Set Computers," *Comm. ACM* 28:1 (January) 8–21. (p. 189)

PATTERSON, D. A. AND C. H. SEQUIN [1981]. "Lockup-free instruction fetch/prefetch cache organization," *Proc. Eighth Annual Symposium on Computer Architecture* (May 12–14), Minneapolis, Minn., 443–458. (p. 487)

PATTERSON, D. A. AND D. R. DITZEL [1980]. "The case for the reduced instruction set computer," *Computer Architecture News* 8:6 (October), 25–33. (pp. 130, 189)

PATTERSON, D. A., G. A. GIBSON, AND R. H. KATZ [1987]. "A case for redundant arrays of inexpensive disks (RAID)," Tech. Rep. UCB/CSD 87/391, Univ. of Calif. Also appeared in *ACM SIGMOD Conf. Proc.*, Chicago, Illinois, June 1–3, 1988, 109–116. (p. 561)

PENG, V., S. SAMUDRALA, AND M. GAVRIELOV [1987]. "On the implementation of shifters, multipliers, and dividers in VLSI floating point units," *Proc. Eighth IEEE Symposium on Computer Arithmetic*, 95–102. (p. A-62)

PFISTER, G. F., W. C. BRANTLEY, D. A. GEORGE, S. L. HARVEY, W. J. KLEINFEKDER, K. P. MCAULIFFE, E. A. MELTON, V. A. NORTON, AND J. WEISS [1985]. "The IBM research parallel processor prototype (RP3): Introduction and architecture," *Proc. 12th Int'l Symposium on Computer Architecture* (June), Boston, Mass., 764–771. (p. 589)

PHISTER, M., JR. [1979]. *Data Processing Technology and Economics*, 2nd ed., Digital Press and Santa Monica Publishing Company. (p. 80)

PRZYBYLSKI, S. A. [1990]. *Cache Design: A Performance-Directed Approach*, Morgan Kaufmann Publishers, San Mateo, Calif. (p. 487)

PRZYBYLSKI, S. A., M. HOROWITZ, AND J. L. HENNESSY [1988]. "Performance tradeoffs in cache design," *Proc. 15th Annual Symposium on Computer Architecture* (May–June), Honolulu, Hawaii, 290–298. (p. 481)

RADIN, G. [1982]. "The 801 minicomputer," *Proc. Symposium Architectural Support for Programming Languages and Operating Systems* (March), Palo Alto, Calif. 39–47. (p. 189)

RAMAMOORTHY, C. V. AND H. F. LI [1977]. "Pipeline architecture," *ACM Computing Surveys* 9:1 (March) 61–102. (p. 339)

REDMOND, K. C. AND T. M. SMITH [1980]. *Project Whirlwind—The History of a Pioneer Computer*, Digital Press, Boston, Mass. (p. 25)

REIGEL, E. W., U. FABER, AND D. A. FISCHER, [1972]. "The Interpreter—a microprogrammable building block system," *Proc. AFIPS 1972 Spring Joint Computer Conf.* 40, 705–723. (p. 244)

ROBERTS, D., G. TAYLOR, AND T. LAYMAN [1990]. "An ECL RISC microprocessor designed for two-level cache," *IEEE COMPCON* (February). (p. 487)

ROBINSON, B. AND L. BLOUNT [1986]. "The VM/HPO 3880-23 performance results," IBM Tech. Bulletin, GG66-0247-00 (April), Washington Systems Center, Gathersburg, Md. (p. 553)

ROWEN, C., M. JOHNSON, and P. RIES [1988]. "The MIPS R3010 floating-point coprocessor," *IEEE Micro* 53–62 (June). (p. A-53)

RUSSELL, R. M. [1978]. "The CRAY-1 computer system," *Comm. ACM* 21:1 (January) 63–72. (pp. 393, 590)

RYMARCZYK, J. [1982]. "Coding guidelines for pipelined processors," *Proc. Symposium on Architectural Support for Programming Languages and Operating Systems,* IEEE/ACM (March), Palo Alto, Calif., 12–19. (p. 339)

SALEM, K. AND H. GARCIA-MOLINA [1986]. "Disk striping," *IEEE 1986 Int'l Conf. on Data Engineering.* (p. 561)

SAMPLES, A. D. AND P. N. HILFINGER [1988]. "Code reorganization for instruction caches," Tech. Rep. UCB/CSD 88/447 (October), Univ. of Calif., Berkeley. (p. 496)

SANTORO, M. R., G. BEWICK, and M. A. HOROWITZ [1989]. "Rounding algorithms for IEEE multipliers," *Proc. Ninth IEEE Symposium on Computer Arithmetic,* 176–183. (p. A-21)

SCHNECK, P. B. [1987]. *Supercomputer Architecture,* Kluwer Academic Publishers, Norwell, Mass. (p. 394)

SCOTT, N. R. [1985]. *Computer Number Systems and Arithmetic,* Prentice-Hall, Englewood Cliffs, N.J. (p. A-1)

SCRANTON, R. A., D. A. THOMPSON, AND D. W. HUNTER [1983]. "The access time myth," Tech. Rep. RC 10197 (45223) (September 21), IBM, Yorktown Heights, N.Y. (p. 561)

SEITZ, C. [1985]. "The Cosmic Cube," *Comm. ACM* 28:1 (January) 22–31. (p. 590)

SHURKIN, J. [1984]. *Engines of the Mind: A History of the Computer,* W. W. Norton, New York. (p. 25)

SHUSTEK, L. J. [1978]. "Analysis and performance of computer instruction sets," Ph.D. Thesis (May), Stanford Univ., Stanford, Calif. (p. 187)

SITES, R. [1979]. *Instruction Ordering for the CRAY-1 Computer,* Tech. Rep. 78-CS-023 (July), Dept. of Computer Science, Univ. of Calif., San Diego. (p. 339)

SITES, R. L., [1979]. "How to use 1000 registers," *Caltech Conf. on VLSI* (January). (p. 487)

SLATER, R. [1987]. *Portraits in Silicon,* The MIT Press, Cambridge, Mass. (p. 25)

SLOTNICK, D. L., W. C. BORCK, AND R. C. MCREYNOLDS [1962]. "The Solomon computer," *Proc. Fall Joint Computer Conf.* (December), Philadelphia, 97–107. (p. 589)

SMITH, A. AND J. LEE [1984]. "Branch prediction strategies and branch target buffer design," *Computer* 17:1 (January) 6–22. (p. 339)

SMITH, A. J. [1982]. "Cache memories," *Computing Surveys* 14:3 (September) 473–530. (p. 486)

SMITH, A. J. [1985]. "Disk cache—miss ratio analysis and design considerations," *ACM Trans. on Computer Systems* 3:3 (August) 161–203. (p. 538)

SMITH, A. J. [1986]. "Bibliography and readings on CPU cache memories and related topics," *Computer Architecture News* (January) 22–42. (p. 486)

SMITH, B. J. [1981]. "Architecture and applications of the HEP multiprocessor system," *Real-Time Signal Processing IV* 298 (August) 241–248. (p. 395)

SMITH, J. E. [1981]. "A study of branch prediction strategies," *Proc. Eighth Symposium on Computer Architecture* (May), Minneapolis, 135–148. (p. 339)

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 758

SMITH, J. E. [1984]. "Decoupled access/execute computer architectures," *ACM Trans. on Computer Systems* 2:4 (November), 289–308. (p. 340)

SMITH, J. E. [1988]. "Characterizing computer performance with a single number," *Comm. ACM* 31:10 (October) 1202–1206. (p. 78)

SMITH, J. E. [1989]. "Dynamic instruction scheduling and the Astronautics ZS-1," *Computer* 22:7 (July) 21–35. (p. 340)

SMITH, J. E. AND A. R. PLEZKUN [1988]. "Implementing precise interrupts in pipelined processors," *IEEE Trans. on Computers* 37:5 (May) 562–573. (p. 339)

SMITH, J. E. AND J. R. GOODMAN [1983]. "A study of instruction cache organizations and replacement policies," *Proc. Tenth Annual Symposium on Computer Architecture* (June 5–7), Stockholm, Sweden,, 132–137. (p. 490)

SMITH, J. E., G. E. DERMER, B. D. VANDERWARN, S. D. KLINGER, C. M. ROZEWSKI, D. L. FOWLER, K. R. SCIDMORE, J. P. LAUDON [1987]. "The ZS-1 central processor," *Proc. Second Conf. on Architectural Support for Programming Languages and Operating Systems,* IEEE/ACM (March), Palo Alto, Calif., 199–204. (p. 340)

SMITH, M. D., M. JOHNSON, AND M. A. HOROWITZ [1989]. "Limits on multiple instruction issue," *Proc. Third Conf. on Architectural Support for Programming Languages and Operating Systems,* IEEE/ACM (April), Boston, Mass., 290–302. (p. 341)

SMITH, W. R., R. R. RICE, G. D. CHESLEY, T. A. LALIOTIS, S. F. LUNDSTROM, M. A. CHALHOUN, L. D. GEROULD, AND T. C. COOK [1971]. "SYMBOL: A large experimental system exploring major hardware replacement of software," *Proc. AFIPS Spring Joint Computer Conf.,* 601–616. (p. 129)

SMOTHERMAN , M. [1989]. "A sequencing-based taxonomy of I/O systems and review of historical machines," *Computer Architecture News* 17:5 (September) 5–15. (pp. 241, 561)

SOHI, G. S., AND S. VAJAPEYAM [1989]. "Tradeoffs in instruction format design for horizontal architectures," *Proc. Third Conf. on Architectural Support for Programming Languages and Operating Systems,* IEEE/ACM (April), Boston, Mass. 15–25. (p. 341)

SPEC [1989]. "SPEC Benchmark Suite Release 1.0," October 2, 1989. (p. 48)

SPORER, M., F. H. MOSS AND C. J. MATHAIS [1988]. "An introduction to the architecture of the Stellar Graphics supercomputer," *COMPCON, IEEE* (March) 464–467. (p. 395)

STERN, N. [1980]. "Who invented the first electronic digital computer," *Annals of the History of Computing* 2:4 (October) 375–376. (p. 24)

STRAPPER, C. H. [1989]. "Fact and fiction in yield modelling," Special Issue of the *Microelectronics Journal* entitled *Microelectronics into the Nineties*, Oxford, UK; Elsevier (May). (p. 80)

STRAPPER, C. H., F. H. ARMSTRONG, AND K. SAJI [1983]. "Integrated circuit yield statistics," *Proc. IEEE* 71:4 (April) 453–470. (p. 80)

STRECKER, W. D. [1976]. "Cache memories for the PDP-11?," *Proc. Third Annual Symposium on Computer Architecture* (January), Pittsburgh, Penn., 155–158. (pp. 187, 486)

STRECKER, W. D. [1978]. "VAX-11/780: A virtual address extension to the PDP-11 family," *Proc. AFIPS National Computer Conf.* 47, 967-980. (128, 187)

STRECKER, W. D. AND C. G. BELL [1976]. "Computer structures: What have we learned from the PDP-11?," *Proc. Third Symposium on Computer Architecture.* (p. 187)

SUTHERLAND, I. E. [1963]. "Sketchpad: A man-machine graphical communication system," *Spring Joint Computer Conf.* 329. (p. 561)

SWAN, R. J., A. BECHTOLSHEIM, K. W. LAI, AND J. K. OUSTERHOUT [1977]. "The implementation of the Cm* multi-microprocessor," *Proc. AFIPS National Computing Conf.,* 645–654. (p. 589)

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 759

SWAN, R. J., S. H. FULLER, AND D. P. SIEWIOREK [1977]. "Cm*—A modular, multi-microprocessor," *Proc. AFIPS National Computer Conf.* 46, 637–644. (p. 590)

SWARTZ, J. T. [1980]. "Ultracomputers," *ACM Transactions on Programming Languages and Systems* 4:2, 484–521 (p. 592)

SWARTZLANDER, E., ED. [1980]. *Computer Arithmetic,* Dowden, Hutchison and Ross (distributed by Van Nostrand, New York). (p. A-59)

TAKAGI, N., H. YASUURA, AND S. YAJIMA [1985]."High-speed VLSI multiplication algorithm with a redundant binary addition tree," *IEEE Trans. on Computers* C-34:9, 789–796. (p. A-59)

TANENBAUM, A. S. [1978]. "Implications of structured programming for machine architecture," *Comm. ACM* 21:3 (March) 237–246. (p. 128)

TANG, C. K. [1976]. "Cache system design in the tightly coupled multiprocessor system," *Proc. 1976 AFIPS National Computer Conf.*, 749–753. (p. 487)

TAYLOR, G. S. [1981]. "Compatible hardware for division and square root," *Proc. Fifth IEEE Symposium on Computer Arithmetic,* 127–134. (p. A-62)

TAYLOR, G. S. [1985]. "Radix 16 SRT dividers with overlapped quotient selection stages," *Proc. Seventh IEEE Symposium on Computer Arithmetic,* 64–71. (p. A-56)

TAYLOR, G. S., P. N. HILFINGER, J. R. LARUS, D. A. PATTERSON, AND B. G. ZORN [1986]. "Evaluation of the SPUR Lisp architecture," *Proc. 13th Annual Symposium on Computer Architecture* (June 2–5), Tokyo, Japan, 444–452. (pp. 189, 451)

TAYLOR, G., P. HILFINGER, J. LARUS, D. PATTERSON, AND B. ZORN [1986]. "Evaluation of the SPUR LISP architecture," *Proc. 13th Symposium on Computer Architecture (*June), Tokyo. (p. E-15)

THACKER, C. P. AND L. C. STEWART [1987]. "Firefly: a multiprocessor workstation," *Proc. Second Int'l Conf. on Architectural Support for Programming Languages and Operating Systems,* Palo Alto, Calif., 164–172. (p. 487)

THACKER, C. P., E. M. MCCREIGHT, B. W. LAMPSON, R. F. SPROULL, AND D. R. BOGGS [1982]. "Alto: A personal computer," in *Computer Structures: Principles and Examples*, D. P. Siewiorek, C. G. Bell, and A. Newell, eds., McGraw-Hill, New York, 549–572. (p. 560)

THADHANI, A. J. [1981]. "Interactive user productivity," *IBM Systems J.* 20:4, 407–423. (p. 560)

THISQUEN, J. [1988]. "Seek time measurements," *Amdahl Peripheral Products Division Tech. Rep.* (May). (p. 558)

THORLIN, J. F. [1967]. "Code generation for PIE (parallel instruction execution) computers," *Spring Joint Computer Conf.* (April), Atlantic City, N.J. (p. 339)

THORNTON, J. E. [1964]. "Parallel operation in Control Data 6600," *Proc. AFIPS Fall Joint Computer Conf.* 26, part 2, 33–40. (pp. 128, 339)

THORTON, J. E. [1970]. *Design of a Computer, the Control Data 6600,* Scott, Foresman, Glenview, Ill. (p. 339)

TJADEN, G. S. AND M. J. FLYNN [1970]. "Detection and parallel execution of independent instructions," *IEEE Trans. on Computers* C-19:10 (October) 889–895. (p. 340)

TOMASULO, R. M. [1967]. "An efficient algorithm for exploring multiple arithmetic units," *IBM J. of Research and Development* 11:1 (January) 25–33. (p. 339)

TRELEAVEN, P. C., D. R. BROWNBRIDGE, and R. P. HOPKINS [1982]. "Data-driven and demand-driven computer architectures," *Computing Surveys*, 14:1 (March) 93–143. (p. 590)

TROIANI, M., S. S. CHING, N. N. QUAYNOR, J. E. BLOEM, AND F. C. COLON OSORIO [1985]. "The VAX 8600 I Box, a pipelined implementation of the VAX architecture," *Digital Technical J.* 1 (August) 4–19. (p. 328)

TUCKER, S. G. [1967]. "Microprogram control for the System/360," *IBM Systems Journal* 6:4, 222–241. (p. 242)

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 760

UNGAR, D. M. [1987]. *The Design of a High Performance Smalltalk System*, The MIT Press Distinguished Dissertation Series, Cambridge, Mass. (p. 451)

UNGAR, D., R. BLAU, P. FOLEY, D. SAMPLES, AND D. PATTERSON [1984]. "Architecture of SOAR: Smalltalk on a RISC," *Proc. 11th Symposium on Computer Architecture* (June), Ann Arbor, Mich., 188–197. (p. 189)

UNGAR, D., R. BLAU, P. FOLEY, D. SAMPLES, AND D. PATTERSON [1984]. "Architecture of SOAR: Smalltalk on a RISC," *Proc. 11th Symposium on Computer Architecture* (June), Ann Arbor, Mich., 188–197. (p. E-15)

UNGER, S. H. [1958]. "A computer oriented towards spatial problems," *Proc. Institute of Radio Engineers* 46:10 (October) 1744–1750. (p. 589)

VON NEUMANN, J. [1945]. "First draft of a report on the EDVAC." Reprinted in W. Aspray and A. Burks, eds., *Papers of John von Neumann on Computing and Computer Theory* (1987), 17–82, The MIT Press, Cambridge, Mass. (p. 592)

WAKERLY, J. [1989]. *Microcomputer Architecture and Programming,* J. Wiley, New York. (p. 188)

WANG, E.-H., J.-L. BAER, AND H. M. LEVY [1989]. "Organization and performance of a two-level virtual-real cache hierarchy," *Proc. 16th Annual Symposium on Computer Architecture* (May 28–June 1), Jerusalem, Israel, 140–148. (p. 487)

WATANABE, T. [1987]. "Architecture and performance of the NEC supercomputer SX system," *Parallel Computing* 5, 247–255. (p. 394)

WATERS, F., ED. [1986]. *IBM RT Personal Computer Technology,* IBM, Austin, Tex., SA 23-1057. (p. 190)

WATSON, W. J. [1972]. "The TI ASC–A highly modular and flexible super computer architecture," *Proc. AFIPS Fall Joint Computer Conf.,* 221–228. (p. 393)

WEICKER, R. P. [1984]. "Dhrystone: A synthetic systems programming benchmark," *Comm. ACM* 27:10 (October) 1013–1030. (p. 47)

WEISS, S. AND J. E. SMITH [1984]. "Instruction issue logic for pipelined supercomputers," *Proc. 11th Symposium on Computer Architecture* (June), Ann Arbor, Mich., 110–118. (p. 339)

WEISS, S. AND J. E. SMITH [1987]. "A study of scalar compilation techniques for pipelined supercomputers," *Proc. Second Conf. on Architectural Support for Programming Languages and Operating Systems* (March), IEEE/ACM, Palo Alto, Calif., 105–109. (p. 340)

WESTE, N. AND K. ESHRAGHIAN [1985]. *Principles of CMOS VLSI Design,* Addison-Wesley, Reading, Mass. (p. A-59)

WHITBY-STREVENS C. [1985]. "The transputer," *Proc. 12th Int'l Symposium on Computer Architecture*, Boston, Mass. (June) 292–300. (p. 589)

WICHMANN, B. A. [1973]. *Algol 60 Compilation and Assessment*, Academic Press, New York. (p. 46)

WIECEK, C. [1982]. "A case study of the VAX 11 instruction set usage for compiler execution," *Proc. Symposium on Architectural Support for Programming Languages and Operating Systems* (March), IEEE/ACM, Palo Alto, Calif., 177–184. (p. 188)

WILKES, M. [1965]. "Slave memories and dynamic storage allocation," *IEEE Trans. Electronic Computers* EC-14:2 (April) 270–271. (p. 486)

WILKES, M. V. [1953]. "The best way to design an automatic calculating machine," in *Manchester University Computer Inaugural Conf.*, 1951, Ferranti, Ltd., London. (Not published until 1953.) Reprinted in "The Genesis of Microprogramming" in *Annals of the History of Computing* 8:116. (p. 241)

WILKES, M. V. [1982]. "Hardware support for memory protection: Capability implementations," *Proc. Symposium on Architectural Support for Programming Languages and Operating Systems* (March 1–3), Palo Alto, Calif., 107–116. (pp. 107, 486)

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 761

WILKES, M. V. [1985]. *Memoirs of a Computer Pioneer*, The MIT Press, Cambridge, Mass. (pp. 25, 241)

WILKES, M. V. AND J. B. STRINGER [1953]. "Microprogramming and the design of the control circuits in an electronic digital computer," *Proc. Cambridge Philosophical Society* 49:230–238. Also reprinted in D. P. Siewiorek, C. G. Bell, and A. Newell, *Computer Structures: Principles and Examples* (1982), McGraw-Hill, New York, 158–163, and in "The Genesis of Microprogramming" in *Annals of the History of Computing* 8:116. (p. 248)

WILKES, M. V. AND W. RENWICK [1949]. *Report of a Conf. on High Speed Automatic Calculating Machines,* Cambridge, England. (p. 88)

WILKES, M. V., D. J. WHEELER, AND S. GILL [1951]. *The Preparation of Programs for an Electronic Digital Computer*, Addison-Wesley Press, Cambridge, Mass. (p. 24)

WILLIAMS, T. E., M. HOROWITZ, R. L. ALVERSON, AND T. S. YANG [1987]. "A self-timed chip for division," *Advanced Research in VLSI, Proc. 1987 Stanford Conf.,* The MIT Press, Cambridge, Mass. (p. A-46)

WILSON, A. W., JR. [1987]. "Hierarchical cache/bus architecture for shared memory multiprocessors," *Proc. 14th Int'l Symposium on Computer Architecture* (June), Pittsburg, Penn., 244–252. (p. 589)

WULF, W. [1981]. "Compilers and computer architecture," *Computer* 14:7 (July) 41–47. (p. 130)

WULF, W. A., R. LEVIN AND S. P. HAREISON [1981]. *Hydra/C.mmp: An Experimental Computer System,* McGraw-Hill, New York. (p. 485)

WULF, W. AND C. G. BELL [1972]. "C.mmp—A multi-mini-processor," *Proc. AFIPS Fall Joint Computing Conf.* 41, part 2, 765–777. (p. 590)

WULF, W. AND S. P. HAREISON [1978]. "Reflections in a pool of processors—An experience report on C.mmp/Hydra," *Proc. AFIPS 1978 National Computing Conf.* 48 (June), Anaheim, Calif. 939–951. (p. 589)

Index

Bold page numbers indicate term definitions.

∞ (*see* infinity)
-∞ (*see* infinity)
+∞ (*see* infinity)
10000 (*see* Apollo DN 10000)
11/780 (*see* Digital Equipment Corporation, VAX-11/780)
11/785 (see Digital Equipment Corporation, VAX-11/785)
2000 (*see* MIPS Computer Corporation, 2000; Digital Equipment Corporation, VAXstation)
2:1 cache rule, front endsheet
2100 (*see* Sequent Corporation)
29000 (*see* AMD 29000)
3000 (*see* MIPS Computer Corporation, 3000)
3010 (*see* MIPS Computer Corporation, 3010)
3090 (*see* International Business Machines Corp.; disk, magnetic, IBM 3990 storage subsystem and)
3090-600S (*see* International Business Machines Corp., IBM 3090-600S)
3100 (*see* Digital Equipment Corporation, DECstation; Digital Equipment Corporation, VAXstation)
3364 (*see* Weitek 3364)
360 (*see* International Business Machines Corp., IBM 360)
360/85 (*see* International Business Machines Corp., IBM 360/85)
360/91 (*see* International Business Machines Corp., IBM 360/91)
370 (*see* International Business Machines Corp., IBM 370)
370/158 (*see* International Business Machines Corp., IBM 370/158)
370-XA (*see* International Business Machines Corp., IBM 370-XA)
3990 (*see* International Business Machines Corp.; disk, magnetic, IBM 3990 storage subsystem and)
68000 (*see* Motorola Corporation, 68000)
6809 (*see* Motorola Corporation, 6809)
701 (*see* International Business Machines Corp., IBM 701)
7030 (*see* International Business Machines Corp., IBM 7030)
704 (*see* International Business Machines Corp., IBM 704)
7090 (*see* International Business Machines Corp., IBM 7090)
8000 (*see* Sequent Corporation)
801 (*see* International Business Machines Corp., IBM 801)
8012 (*see* International Business Machines Corp., IBM 8012)
80186 (*see* Intel Corporation, 80x86, 80186)
80286 (*see* Intel Corporation, 80x86, 80286)
80386 (*see* Intel Corporation, 80x86, 80386)
80486 (*see* Intel Corporation, 80x86, 80486)
80x86 (*see* Intel Corporation, 80x86)
8080 (*see* Intel Corporation, 8080)
8086 (*see* Intel Corporation, 80x86, 8086)
8088 (*see* Intel Corporation, 8088)
8550 (*see* Digital Equipment Corporation, VAX)
860 (*see* Intel Corporation, 860; Intel Corporation, i860)
8600 (*see* Digital Equipment Corporation, VAX)
8700 (*see* Digital Equipment Corporation, VAX)
88000 (*see* Motorola Corporation, 88000)
88100 (*see* Motorola Corporation, 88100)
88200 (*see* Motorola Corporation, 88200)
8847 (*see* Texas Instruments, 8847)
90/10 rule, front endsheet (*see also* locality, principle of)
90/50 branch-taken rule, front endsheet) (see also branch, taken)

A

aborts, **216** (*see also* interrupts)
absolute addressing (*see* addressing mode, direct)
access alignment (*see* data alignment)
access authorization (*see* virtual memory, Intel 80286/80386 and; virtual memory, protection schemes of)
access bit, **446** (*see also* virtual memory, page table)
access latency, **405** (*see also* memory hierarchy, access time; cache, access time)
access time 19, 20, **405**, 420 **425** (*see also* memory hierarchy, access time; cache access time)
access time gap, 518 (fig.), **519**
accumulator architecture (*see* architecture, accumulator)
accumulator-based architecture (*see* architecture, accumulator)
Adams, T., 188
adders, A-39 (fig.) (*see also* arithmetic, integer, ripple-carry addition; arithmetic, integer, speeding up addition)
addition (*see* arithmetic, addition, floating-point; arithmetic, integer, addition; arithmetic, integer, speeding up addition)
address (*see also* addressing mode)
  consumption of, front endsheet, 16
  effective, **97**–98
  fault, **433** (*see also* virtual memory, page fault)
  memory, 12, 18, 21, 93, 94-103, 115-117, 134
    shared versus multiple, 578-579
  space, 16, 19 (*see also* cache; memory; memory hierarchy; virtual memory; virtual memory, processes and)
    consequences of too small an, 480-481
    extensions of, 483
    on the Intel 80286, 445-446
    on the VAX-11/780, 441
  specifier, **102** (*see also* addressing mode)
  translation, **433** (*see also* virtual memory, address translation)
address-consumption rate, front endsheet, 16, 480-481

addressing mode, 97-103, 126, 134, 136 (*see also* Digital Equipment Corporation, VAX; DLX; Intel Corporation, Intel 80x86, 8086; International Business Machines Corp., IBM 360)
autodecrement, **98**
autoincrement, **98**
encoding of, 102-103
direct (absolute), **98**
displacement (based), 98-100, 105-106, 114, 133
field, 100, 102-103, 106
size, 100
value, 100
immediate (literal), 98-102
field, 102
value, 101
indexed, **98**, 136
memory indirect (memory deferred), **98**-99
operand specifiers, 145 (fig.), 169 (fig.), 173 (fig.), 177 (fig.), 180 (fig.)
register deferred (indirect), **98**, 136
of RISC architectures, E-2
scaled (index), **98**, 126
after rounding, A-20, **A-22** (*see also* arithmetic, rounding and)
Aiken, 24
algorithm, 14
Agarwal, A., 190
Alexander, W. G., 130, 187
aliased variables, **116**-117
aliases, **460**
alignment (*see also* data alignment; stack, alignment of)
interrupts and, 215
on the DLX, 221, 231
alignment network, 96-97, 135
Alto, 560
ALU (*see* arithmetic logic unit)
AMD 29000, 167, 190
Amdahl, G. M., 17, 26, 127, 186, 242, 588 (*see also* Amdahl's Law)
Amdahl/Case rule (*see* Case/Amdahl rule of thumb)
Amdahl's Law, 8-11, 22, 26, 29, 575-576, 586 (*see also* Case/Amdahl rule of thumb)
CPU-DRAM performance gap and, 426, 427 (fig.), 432
I/O and, 500, 555, 559
Amdahl's rule of thumb, 426 (*see also* Case/Amdahl rule of thumb)
Annual International Symposium on computer architecture (*see* architecture, Annual International Symposium on)
anti-aliasing, **460** (*see also* virtual cache)
antidependence, **374** (*see also* vector processor, antidependence)
AP-120B (*see* Floating-Point Systems)
Apollo DN 10000, 340
Archibald, J., 471, 487, 488
architecture, 3, 4, 5, **13**, 128 (*see also* Digital Equipment Corporation, VAX; DLX; HLLCA; Intel Corporation, 860; Intel Corporation, 80x86; International Business Machines Corp., IBM 360; MIPS Computer Corporation, R3000; Motorola Corporation; SPARC)
accumulator, 24, 90-92, 127
Annual International Symposium on, 80
decoupled (*see* decoupled architecture)
definitions, front endsheet
evolution-revolution spectrum of, 587-**588**
evolution versus revolution, 587-588
formualas, front endsheet
general-purpose register (*see* general-purpose register architecture)
architecture (*continued*)
Harvard, 25
instruction set (*see* instruction set, architecture)
load/store (*see* load/store architecture)
memory-memory (*see* memory-memory architecture)
performance evaluation of, 78-80
register-memory (*see* register-memory architecture)
register-register (*see* register-register architecture)
revolutionary, 593
rules of thumb (see rules of thumb)
simulator, 48
stack (*see* stack architecture)
systolic (*see* systolic architecture)
trends of, 16
trivia, front endsheet
vector (*see* vector processor, architecture)
areal density of disk, **518** (*see also* maximum areal density; disk, magnetic)
arithmetic, 15, 201, A-1
addition, floating-point, A-16–A-20
algorithm for, A-18–A-19
denormals and, A-20
rounding in, A-16–A-17 (*see also* infinity)
addition, integer, A-2–A-3 (*see also* arithmetic, integer)
add, subtract and multiply instructions in Intel 860, E-22
Booth recoding, **A-8**–A-9, A-20, A-40, A-43–A-44, A-48, A-56, A-59
modified, A-64
decimal, 15, 103, 109-110 (*see also* arithmetic, integer; arithmetic, floating-point)
denormals, **A-14**–A-15, A-20, A-21–A-22, A-31, A-60
division, integer, A-3–A-7 (*see also* arithmetic, integer)
nonrestoring, **A-5**, A-6 (fig.), A-40–A-41, A-42
speeding up, A-50
restoring, **A-5**, A-6 (fig.)
division, floating-point, A-23–A-26
exceptions, A-30–A-31
overflow, **A-7**, A-10, A-11 (fig.), A-30–A-31
floating-point addition, A-20
floating-point multiplication, A-21
integer, A-10
underflow, A-20, A-21–A-22, A-30, A-57
gradual underflow, **A-15**, A-22, A-59–A-60, A-63
underflow trap, A-31
exponents and, A-1, A-12, **A-13**–A-14, A-15
exponent field, **A-13**–A-14, A-20
fallacies and pitfalls, A-57–A-58
floating-point, A-12–A-31, A-57, A-58, A-59 (*see also* arithmetic, IEEE standard and)
addition, A-16–A-20 (*see also* arithmetic, addition, floating-point)
division, A-23–A-26
exceptions, A-30–A-31
multiplication, A-20–A-23 (*see also* arithmetic, multiplication, floating-point)
precision, A-22–A-23, A-28–A-30
remainder, A-26–A-28
history of, A-58–A-60
IEEE standard and, 109, A-1, A-12–A-16, A-60, E-2, E-23
integer, A-2–A-11, A-57
basic techniques of, A-2–A-11
multiple-precision addition, A-11
radix-2 multiplication and division, A-3–A-6
ripple-carry addition, **A-2**–A-3, A-32, A-36 (fig.), A-39 (fig.)
signed numbers and, A-7–A-10
speeding up addition, A-31–A39

arithmetic, integer, speeding up addition (*continued*)
carry-lookahead adder (CLA), **A-32**, A-36, A-39 (fig.) (*see also* carry)
carry-select adder, **A-38**–A-39, A-39 (fig.), A-56, A-66 (*see also* carry)
carry-skip adder, **A-36**–A-37, A-39 (fig.) (*see also* carry)
speeding up division, A-39–A-42, A-50–A-53
shifting over zeros, **A-40**–A-42
with a single adder, A-50–A-53
speeding up multiplication, A-39–A-50
shifting over zeros, **A-40**
with a single adder, A-42–A-44
with many adders, A-44–A-49
systems issues of, A-10–A-11
multiplication and division, integer, A-3–A-7 (*see also* arithmetic, integer)
multiplication, floating-point, A-20–A-23
denormals and, A-21–A-2
precision, A-22–A-23
operations, 103
precision, A-22–A-23, A-28–A-30
double-extended, A-28, A-60
multiple-precision addition, A-11
remainder, A-4–A-5, A-40–A-42, A-50
floating-point, A-26–A-28
REM, A-26–A-28, A-53 (*see also* arithmetic, remainder)
rounding and, A-13, A-16–A-17, A-18, A-19, A-20–A-21, A-23, A-26 (*see also* infinity)
after rounding, A-20, **A-22**
before rounding, A-17, **A-22**
double rounding, **A-29**, A-64
rounding errors, A-24, A-25
rounding mode, **A-13**, A-22
signed, A-7–A-10, A-58
signed-digit representation, **A-48**
signed-logarithm representation, A-65
significand, A-12, **A-14**, A-18, A-22, A-29
square root, A-25, A-26, A-29–A-30, A-64
of a negative number, A-12–A-13
systems and, A-10–A-11 (*see also* not a number; infinity)
arithmetic and logical instructions, 92
coprocessor operations, **E-9**, E-11
in RISC architectures, E-5
arithmetic and logical operators, 103
arithmetic logic unit (ALU), 39-42, 201
clock cycles per instruction and, 224-226, 235
DLX states and, 222 (fig.), 225-226
effect on rain forest from papers about, 201
encoding and, 235-236
instructions and operations of, 91, 93, 101, 103, 106, 120, 123, 132, 133, 136, 202-203, 211, 213, 229-234 (figs.), 237 (fig.)
arithmetic mean (*see* mean, arithmetic)
arithmetic operations, 103 (*see also* arithmetic; instruction set)
arithmetic overflow (*see* interrupts, arithmetic overflow and)
arm, **516** (*see also* disk, magnetic)
Armstrong, F. H., 81
ARPANET, 527, 528 (fig.), 561 (*see also* networks)
array, A-45–A-46 (*see also* systolic arrays)
array multiplier, **A-44**, A-45–A-47 (figs.), A-49 (fig.), A-56 (*see also* arithmetic, integer, speeding up multiplication)
array of disks (*see* disk array)
array processor (*see* single instruction stream, multiple data stream computer)
ASCII, 109
ASP (*see* cost, average selling price)
ASPLOS (Architectural Support for Programming Languages and Operating Systems) conference, 130
associativity, 420 (*see also* cache, fully associative; cache, set associative)
asynchronous bus, **530** (*see also* bus)
Atanasoff, J. V., 24
Atlas computer, 26, 485
atomic, **471** (*see also* cache, coherency, synchronization)
atomic swap instruction (*see* data transfer)
atomic swap operation, 471 (*see also* cache, coherency, synchronization)
attributes field, **446** (*see also* virtual memory, page table; virtual memory; Intel 80286/80386 and)
Auslander, M. A., 130
autodecrement, **98** (*see also* addressing mode)
autoincrement, **98** (*see also* addressing mode)
availability, **520** (*see also* input/output, reliability)
average instruction execution time, 228
average memory-access time, **461** (*see also* memory hierarchy, access time; cache, access time; cache, two-level caches)

B

B5000 (*see* Burroughs)
B5500 (*see* Burroughs)
B6500 (*see* Burroughs)
Baer, J.-L., 471, 487, 488
Balance (*see* Sequent Corporation)
balance (tradeoffs), 121, 131, 135, 140-141, 220 (*see also* design, computer; Case/Amdahl rule of thumb)
pipelining
balance among stages, 252
balance in issue, 320
software and hardware, 14-16, 21, 28
bandwidth, 5, 18, 19 (fig.), 29, 124, 135
performance measures of main memory and, 425
I/O and (*see* input/output, performance, throughput)
bandwidth, I/O (*see* input/output, performance, throughput)
Banerjee test (*see* vector processor, data dependences, Banerjee test)
Barton, R. S., 127
base, **439** (*see also* virtual memory, protection schemes of; virtual memory, Intel 80286/80386 and)
based addressing mode (*see* addressing mode, displacement)
base field, **446** (*see also* virtual memory, page table)
basic architecture of vector processor (*see* vector processors, architecture)
basic block, **115**
Baskett, F., 130
BCD (*see* binary-coded decimal)
before rounding, A-17, **A-22** (*see also* arithmetic, rounding)
behavior, **512** (*see also* input/output, devices)
Bell, C. G., 81, 127, 590
bet with Hillis, 590
W. D. Strecker and, 485, 488
benchmark, 42, 43, 45-48, 53, 72, 75, 81, 82, 83, 85-86 (*see also* disk, magnetic, I/O benchmarks for; input/output, performance)
file system I/O, 512 (*see also* input/output, benchmarks)
historical perspective, 77-80
kernels, 45
Linpack (*see* vector processor, Linpack benchmark)
Perfect Club, 75, 79-80
vectorization and, 375
SPEC (System Performance Evaluation Cooperative), 48, 72-73, 79, 81, 83

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 765

benchmark (*continued*)
supercomputer I/O, 510-511 (*see also* input/output, benchmarks)
synthetic, 45-48, 73-74, 80, 86
Dhrystone, 45, 47, 73-74, 81, 85, 86
Whetstone, 45-46, 73-74, 77, 82, 83, 86
toy, 45
TP-1, 511, 511 (fig.), 565
transaction processing I/O, 511-512 (*see also* input/output, transactions and)
unfair, **490**
benchmark programs (*see* benchmark)
Berkeley RISC (*see* reduced instruction set computer, Berkeley)
Berry, M. D., 81
biased exponent, A14, A-15 (*see also* arithmetic, exponents and)
Big Endian, front endsheet, **95**
Bigelow, Julian, 24
binary-coded decimal (BCD), **109**
packed, **109**
unpacked, **109**
binary-tree multiplier, **A-48** (*see also* arithmetic)
bit block transfer, 521 (*see also* graphic displays)
bit blts, **521** (*see also* graphics displays)
bit-field instructions in Motorola 88000, E-17–E-18
bit map, **521** (*see also* graphics displays)
Blaauw, G. A., 127, 186
Blau, R., 189
block, **404** (*see also* memory hierarchy, blocks and; cache, blocks and; virtual memory, page; virtual memory, segment)
block-frame address, **405** (*see also* memory hierarchy, blocks and; cache, block-frame address of)
block identification, **407**, 484
caches and, 410-411
virtual caches and, 459-460
virtual memory and, 435-436
block-offset address, **405** (*see also* memory hierarchy, blocks and)
block-offset field, **410** (*see also* cache, blocks and)
block placement, **407**, 484 (*see also* conflict miss)
caches and, 408-409, 420 (*see also* cache, fully associative; cache, set associative; cache, direct mapped)
subblocks, 456-457
VAX-11/780 cache and, 419
virtual memory and, 434-435
block replacement, **407**, 484
caches and, 411-412, 420
early restart, 458
first-in–first-out (FIFO), **412**
least-recently used (LRU), **411**-412, 436
versus random, 412 (fig.)
on the VAX-11/780, 443
out-of-order fetch, 458
random, **411**
versus least-recently used, 412 (fig.)
VAX-11/780 cache and, 411
virtual memory, 436
block size (*see* cache, blocks and, size; virtual memory, paged, page size)
Boggs, D., 560, 562
Booth recoding, **A-8** (*see also* arithmetic, Booth recoding)
bound, **439** (*see also* virtual memory, protection schemes of)
bounds checking (*see* virtual memory, Intel 80286/80386 and)
Brady, J., 509, 560, 562
branch, 103, **104**-109, 133 (*see also* jump; branch-prediction schemes)
behavior, 272-273
clock cycles and, 224-225, 237
condition code (CC), 37, **106**, 201, 282-283, 335
conditional, 104-108, 203, 209 (*see also* branch instruction)
of RISC architectures, E-8
condition register, **106**
delay, **272**, 273-277, 282, 335 (*see also* hazard, branch-delay slots; branch-prediction schemes)
delayed, 274 (fig.), **275**, 276-279, 339
scheduling, 274-275
DLX and (*see* branch instruction, of DLX)
frequency, 272
hazard (*see* hazard, branch)
instruction, 37-38, 104
branch conditions of DLX, 203, 237
conditional, 37-38, 104
of DLX, 203, 224 (fig.), 230 (fig.), 234-237
loop, 108
not taken, **270**, 273 (*see also* branch-prediction schemes, predict-not-taken)
offset, 105-106
optimization, **114** (*see also* optimization)
penalty, 271-272, 277
determining, 313
reduction, 273-278, 307-314
on DLX, 276-277, 310
optimization and, 114-115, 119-120
PC (program-counter)-relative branches, **105**
pipelining and (*see* branch-prediction schemes)
prediction (*see* branch-prediction schemes)
scheduling, 274-275, 282
schemes, 277 (*see also* branch-prediction schemes; dynamic hardware branch prediction)
taken, 107-**108**, **270**, 273 (*see also* branch-prediction schemes, predict-taken)
90/50 branch-taken rule, front endsheet
target, 105-106
unconditional (*see* jump)
branch-delay slots, 274, **275**-276, 279, 335
empty, 276
filled, 276
scheduling, 274 (fig.), 276 (fig.), 345
branch likely instruction, **E-14** (*see also* delayed branch)
branch-prediction buffer, **308**-310
branch-prediction schemes 273-277, 308-314, 339-340 (*see also* dynamic hardware branch prediction; misprediction penalty)
prediction accuracy, 309-310, 313
predict-not-taken, 273-**274**, 277 (fig.), 309, 312-313
predict-taken, 274-275, 277 (fig.), 309, 312-313, 331-332
reducing branch penalties with dynamic hardware prediction, 307-314
branch-target buffer, **310**-312, 339-340
bridge, **527** (*see also* networks)
Briggs, F., 190
Brooks, F. P., 127, 186, 445
Brown, E., 190
bubble, **265** (*see also* pipeline stall)
Burks, A. W., 24
Burr, W. E., 79
Burroughs
B5000, 127
B5500, 71
B6500, 127, 131

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 766

bus, **528**-532, 560 (*see also* memory bus)
DLX, 200 (fig.), 201
FutureBus, 532, 532 (fig.)
IBM PC-AT bus, 531
instructions on the DLX, 211, 230
intelligent peripheral interface (IPI), **531**, 532 (fig.), 560
Multibus II, 532, 532 (fig.)
NuBus, 15, **561**
options for, 530 (fig.), 531 (fig.)
PDP-11 Unibus, 531, 560
small computer systems interface (SCSI), 532 (fig.), 560-**561**
standards for, 531-532
comparison of five bus standards, 532 (fig.)
transactions, 529
VME bus, 532, 532 (fig.)
bus error, 215 (fig.)
bus masters, **529** (*see also* bus)
bus transaction, **529** (*see also* bus)
bus width, 428 (*see also* memory, organization of)
busy-wait (*see* spin waiting)
bypass, **261**, 263, 292, 338 (*see also* forwarding)
bypass registers, 263
byte, 219
alignment of on DLX, 221, 231
load byte on DLX, 232, 235
byte addressed machine, 95
byte ordering, 95 (*see also* Little Endian and Big Endian)

C

cache, **19**-20, 25, 26, 224, 238, **408**-425, 454-474, 481, 483, 484 (fig.), 486-487 (*see also* memory; memory hierarchy; virtual memory; block identification; block placement; block replacement; write strategy)
2:1 cache rule, front end sheet
access time and, 420
average memory-access time, 418-419, 454
block size versus, 423-424 (figs.)
virtual cache, 459-460
blocks and, 408, 420, 425, 454
address tag, 410
identification (*see* block identification, caches and)
placement (*see* block placement, caches and)
replacement (*see* block replacement, caches and)
size, 420, 423, 454, 469
versus memory access time, 423-424 (figs.)
subblocks, 456-457
VAX-11/780, 414
block-frame address of, 410, 412 (fig.), 414
coherency, 466, 467-474, 487
block size and, 469
cache-coherency problem, **466** (fig.), 468
cache-coherency protocols, **467**-473
directory based, **467**-468
example of, 469-471
hits and, 471
misses and, 468-471
snooping, **467**-474, 487
summary of, 471 (fig.)
write broadcast, **469**-470 (fig.)
write invalidate, **468**-469, 470 (fig.)
example, 473 (fig.)
multilevel caches and, 468
read hits and, 471
read misses and, 469, 471
sequential consistency, 474
synchronization, 471-474
cache, coherency, synchronization (*continued*)
lock variable, **471**, 472 (fig.), 473
unlock, 472 (fig.)
weak consistency, 474
write hits and, 471
write misses and, 469, 471
data-only, 423-425
differences between virtual memory and, 438
direct mapped, 408, 409 (fig.), 410 (fig.), 418-422, 456, 481, 486
2:1 cache rule, front endsheet
address portions of, 410 (fig.)
conflict misses and, 420
disk, 537, **566** (*see also* input/output, interfacing to an operating system)
file, 537, 538 (fig.) (*see also* input/output, interfacing to an operating system)
fully associative, **408**, 409 (fig.), 410 (fig.), 418-422, 454
block placement and, 410 (fig.)
block replacement and, 411, 420
misses and, 420
hit, 412-413, 414, 460
rate, 411
read, 412
reducing hit times by making writes faster, 455-457
making cache hits faster with virtually addressed caches, 459-460
instruction-only, 423-425
I/O and, 466-467
least-recently used block replacement (*see* block replacement, least-recently used)
miss, **19**, 412, 414, 418, 419-422, 429, 459 (*see also* cache, miss rate; cache, write miss)
capacity, 419, 420, 421-422 (figs.)
compulsory, 419, 420, 421-422 (figs.)
conflict, 420, 421-422 (figs.)
reducing miss penalty, 457-458 (*see also* cache, two-level caches)
"three Cs" (capacity, compulsory, conflict), 420, 421-422 (figs.), 484
miss rate, 416, 418-419, 481 (*see also* cache, miss; cache, write miss)
2:1 cache rule, front endsheet
compared to misses per instruction, 417
data-only versus instruction-only miss rates, 424-425 (fig.)
for random vs. least-recently used block replacement, 412
multiprocessors and, 468 (see cache, coherency)
on DLX, 482 (fig.)
on the VAX-11/780, 482 (fig.)
reducing by reducing cache flushes, 466-467
versus cache size, 455 (fig.)
for two-level caches, 462 (fig.)
using a process-identifier tag (PID), 459 (fig.)
mixed cache (*see* cache, unified)
multilevel (*see* cache, two-level caches)
multiprocessors and (*see* cache, coherency)
n-way set associative (*see* cache, set associative, *n*-way)
parameters, typical, 408 (fig.) (*see also* parameters, typical ranges of)
performance, 416-419, 454-474, 481, 483
pipelined machines and, 334
random (*see* block replacement, random)
reads and, 412, 416 (*see also* cache, miss)
read miss rate, 416
register versus, speed of, 483
set associative, **409**, 409 (fig.), 410 (fig.), 454, 481

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 767

cache, set associative (*continued*)
block-offset field, **410**, 411 (fig.)
conflict misses and, 420
index field, **410**, 411 (fig.)
n-way, 409, 420-422
SRAM relationship, 426
stale data and (*see* stale data)
subblocks, **456**-457, 492
summary of, 484 (fig.)
synchronization (*see* cache, coherency)
tag field, **410**, 411 (fig.)
two-level caches, 460-465, 484 (fig.), 487
average memory-access time for, 461
coherency, 468 (*see also* cache, coherency)
parameters, typical for, 463 (fig.) (*see also* parameters, typical ranges of)
relative execution time, 463 (fig.), 465 (fig.)
size of, 464
summary of, 484 (fig.)
valid bit, **410**
VAX-11/780 and, 414-416 (*see also* memory hierarchy, VAX-11/780 and; virtual memory, VAX-11/780 and)
miss rates for, 482 (fig.)
unified, **423**
vectors as an alternative to caches, 352
virtual, **460**
virtual memory and, 434, 438
write back, **413**-414, 429, 469
clean, **413**
dirty, **413**
dirty bit, **413**
write buffer, **413**, 457, 482-483 (*see also* cache, writes and)
fallacy of, 482-483
VAX-11/780 and, 413, 416, 477, 483
write stalls and, 457-458
write miss, 413-414, 416, 457-458
no write allocate, **413**-414, 416
making faster, 457-458 (*see also* subblock placement)
multiprocessors and, 468
rate, 416
write allocate, **413**-414
writes and, 413-414, 416 (*see also* cache, write miss)
making writes faster, 455-457
multiprocessor, 468
write strategy (*see* write strategy, caches and)
write through, **413**-414, 416, 457, 477
cache-coherency problem, **466** (*see also* cache, coherency)
cache-coherency protocols, **467** (*see also* cache, coherency)
cache-coherency example, 473 (fig.) (*see also* cache, coherency)
cache machine, 334 (*see also* cache)
Cady, R., 127
call (*see* procedure call/return)
caller, 124-125
callee-saving, **108**-109, 124-125
caller-saving, **108**-109
call gate, **448** (*see also* virtual memory, protection schemes of; virtual memory, Intel 80286/80386 and)
`CallS` instruction, **122**, 124-125, 137, 213
capabilities, **441**, 485 (*see also* virtual memory, protection schemes of)
capacity miss, **420** (*see also* cache, miss, capacity)
capacity
of DRAMs (*see* dynamic random access memory)
of SRAMs (*see* static random access memory)
carry, A-2, A-11 (*see also* carry in; carry out; carry-lookahead adder; carry-propagate adder; carry-save adder; carry-select adder; carry-skip adder)
carry in, **A-2**, A-3 (fig.), A-7, A-37, A-38
carry-lookahead adder (CLA), **A-32** (*see also* arithmetic, integer, speeding up addition)
carry out, **A-2**, A-3 (fig.), A-7, A-15, A-16 (fig.), A-18–A-19, A-33, A-37
carry-propagate adder (CPA), **A-43**, A-48, A-51, A-56
carry-save adder (CSA), **A-42**–A-44, A-45 (fig.), A-51
carry-select adder, **A-38** (*see also* arithmetic, integer, speeding up integer addition)
carry-skip adder, **A-36** (*see also* arithmetic, integer, speeding up addition)
CAS (*see* column-access strobe)
Case, R., 17, 186
Case/Amdahl rule of thumb, front endsheet, **17**, 426 (*see also* balance, software and hardware; rules of thumb; performance)
CPU-DRAM performance gap and, 426, 427 (fig.), 432
cathode ray tube (CRT), **521** (*see also* graphics displays)
CC (*see* branch, condition code)
CD (*see* disk, optical)
CDB (*see* common data bus)
CDC (*see* Control Data Corporation)
CD-ROM, **519** (*see* disk, optical)
centralized memory (*see* memory, centralized)
central processing unit (CPU), 8, 13, 90-92, **199** (*see also* datapath; processor)
balance and, 17
memory hierarchy and, 18-19 (*see also* memory hierarchy)
system performance and, 11, 16
CPU-DRAM performance gap, 426, 427 (fig.), 432
idle time, 500-501
interfacing to I/O (*see* input/output, interfacing to the CPU)
CPU-execution clock cycles and caches, 416
CPU-memory buses, **529** (*see also* bus)
performance, **35**, 36-40, 71 (*see also* performance)
time, **35**-40, 41, 67-69, 122
caches and, 416, 418
I/O and, 499
system CPU time, **35**
user CPU time, **35**
chaining, **378** (*see also* vector processor, chaining and)
Chaitin, G. J., 130
Chandra, A. K., 130
character strings, 109
channel, **548** (*see also* disk, magnetic, IBM 3990 storage subsystem and)
channel controllers, **534** (*see also* input/output, interfacing to the CPU)
channel program, **549** (*see also* disk, magnetic, IBM 3990 storage subsystem and)
Chow, F. C., 114-115, 117, 130
CISC (complex instruction set computer) (*see* reduced instruction set computer; Digital Equipment Corporation, VAX)
CLA (*see* carry-lookahead adder)
Clark, D. W., 130, 171, 188, 189, 486, 488
clean, **413** (*see also* cache, write back)
clock, 36
period (*see* clock cycle)
tick (*see* clock cycle)
clock cycle, 29, **36**-38, 75, 77, 79, 81, 134, 201, 224, 228 (*see also* clock cycle time; clock cycles per instruction)
ALU and, 224-226, 235
branches and, 224-225, 237-238

clock cycle (*continued*)
caches and, 416
control and, 204
DLX and, 224, 235, 237-238
microinstructions and, 211
per instruction (CPI) (*see* clock cycles per instruction)
pipelines and, 278 (fig.), 351
reducing, 207
register file and, 201
stalls and, 213-214, 224
clock cycles per floating-point operation (CPF), 360-361, 378, 392
clock cycles per instruction (CPI), 36-41, 71-72, 77, 82, 94, 132, 134, 199, 224
CPF versus, 392
ALU and, 224-226, 235
DLX and, 224, 225 (fig.), 235, 238
performance and, 210
pipelining and, 252, 258, 351
reducing, 207
by adding hardwired control, 213-214
by parallelism, 214, 314-327
by pipelining, 252, 258
with special case microcode, 213
clock cycle time, 5, 36-41, 81, 199, 201, 228 (*see also* clock rate)
caches and, 416, 481
control and, 210, 227, 240
interrupts and, 214
pipelined machines and, 251-255
clock rate, 36-37, 41, 68, 71, 84, 135, 228 (*see also* clock cycle time)
clock skew, 253-254, 336
CM (*see* Connection Machine)
Cm* multiprocessor, 589 (*see also* multiprocessor)
C.mmp multiprocessor, 589 (*see also* multiprocessor)
COBOL, 15
Cocke, J., 130, 189, 340
code, 45
condition (*see* branch, condition code)
optimized, 41-42, 49, 73
size, 70-71, 73, 78-79, 92, 103, 121, 135, 324
source, 43, 48
system, 35
unoptimized, 41-42, 73
user, 35
code motion, **114** (*see also* optimization)
coherency (*see* cache, coherency)
cold start misses, **419** (*see also* cache, miss, compulsory)
collision misses, **420** (*see also* cache, miss, conflict)
coloring, graph, 113-114, 130
color map, **523** (*see also* graphics displays)
color table, **523** (*see also* graphics displays)
column-access strobe (CAS), **425**
column-major order, **366**, 367 (fig.)
committed instruction, **280**
common case
importance in design, 8
common data bus (CDB), **300**-307, 349
common subexpression elimination, 114 (*see also* optimization)
global, **112**, 114
communication, 573, 574, 592-593, 594
explicit, 579
implicit, 578-579
overhead, 575, 581
compare, 101, 103, 106-107
macrocode improvement of, on VAX-11/780, 239
in RISC architectures, E-7–E-8
Comparability, of instruction sets (*see* object-code compatibility)
compare and branch instruction, 106-**107**
comparison (*see* compare)
comparators
for hazard detection, 263, 269 (*see also* hazard, detection)
compiler, 5, 16, 17, 19, 21, 28, 92-94, 111-122
complexity of, 111, 120-121
future directions for, 581-582
optimizing, 41, 47, 67, 73-74, 81, 111-120, 126, 130, 131, 136
performance and, 37, 42-48, 71-72, 79
structure of, 111-115
vector processor and (*see* vector processor, compilers and)
completion, out-of-order (*see* out-of-order completion)
completion rate, **358** (*see also* vector processor)
complex instruction set computer (CISC) (*see* reduced instruction set computer)
compulsory miss, **419** (*see also* cache, miss, compulsory)
computer architecture (*see* architecture, computer)
Computer Museum, 25
computer program, **243**
condition code (*see* branch, condition code)
conditional branch (*see* branch, conditional)
conditionally executed statements (*see* vector processor, conditionally executed statements and)
conditional-sum adder, **A-66**
conflict miss, **420** (*see also* cache, miss, conflict)
connect/disconnect bus, **530** (*see also* bus)
Connection Machine
CM, 589-590, 591
CM-2, 573, 577, 593
constant bit density, **516** (*see also* disk, magnetic)
constant extension of RISC architectures (*see* reduced instruction set computer)
constant propagation, **114** (*see also* optimization)
context switch, **438** (*see also* virtual memory, processes and)
virtual caches and, 459-460 (fig.)
Conti, C. J., 78
control 199, 201 (fig.)
DLX and, 220-224, 228-234 (*see also* DLX, instruction set, control-flow instructions)
flow (*see also* control-flow instructions)
hardwired, 204-207, 210
reducing CPI by adding hardwired control, 213-214
improving DLX performance when control is hardwired, 226-228
performance of, 207, 224-225, 237
reducing hardware costs of hardwired control, 205-206, 213-214
interrupts and, 217-218
microprogrammed/microcoded 208-214, 238-243
ABCs of microprogramming, 209-210
microcoded control for DLX, 228-234
performance of, 238, 240-241
performance of microcoded control for DLX, 235
reducing cost and improving performance of DLX when control is microcoded, 235-238
reducing hardware costs by encoding control lines, 210-211
reducing hardware costs with multiple microinstruction formats, 211-212
special case microcode, 213
writable control store (WCS) and, 238-239

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 769

Control Data Corporation (CDC), 353, 394
CDC 6600, 71, 128, 132, 292, 295, 299-300, 338-339
ETA-10, 394
control dependences (see hazard, branch)
control-flow instructions, 104-109, 122 (*see also* DLX, instruction set, control-flow instructions)
in RISC architectures, E-6
control hazard (*see* hazard, control)
controller, disk (*see* disk, magnetic)
controller time, **516** (*see also* disk, magnetic)
control operators, **103**
control store, 209, **210**, 212-213, 235, 239 (*see also* writable control store)
coprocessor, **580**, A-28
coprocessor operations (*see* arithmetic and logical instructions)
copy back, **413** (*see also* cache, write back)
copy propagation, **114** (*see also* optimization)
Cosmic Cube multicomputer, 589
cost, 34, 53-54, 80 (*see also* die; integrated circuit; package; wafer; workstation)
average selling price (ASP), 64, 66, 85
average discount, 64-65, 84-85
comparing price of media versus price of packaged system, 556-557
direct, 64-66, 85
DRAM, 556-557
indirect, 64
list price, 64-66, 70, 84-85
magnetic disk, 556-557
versus access time for SRAM, DRAM, and magnetic disks, 518 (fig.)
versus price, 61-64, 65 (fig.), 66 (fig.), 84-85.
cost/performance, 11, 16, 21, 25, 76 (*see also* performance)
design, 34
fallacies, 70
optimizing, 14
price/performance, 47, 66-70, 80
CPA (*see* carry-propagate adder)
CPF (*see* clock cycles per floating-point operation)
CPI (*see* clock cycles per instruction)
CPU (*see* central processing unit)
Crawford, J., 188
Cray, Seymour, 71
Cray Research machines, 34, 353, 390, 391, 393
arithmetic on, A-60
CRAY-1, 353, 377, 391, 392, 393
CRAY-2, 43, 353, 377 (fig.)
CRAY X-MP, 74-75, 80, 353, 376-377, 391, 392, 394, 493
CRAY Y-MP, 353, 391-392, 394
critical section (*see* synchronization)
CRT (*see* cathode ray tube; graphics displays)
Crudele, L., 189
CSA, A-42 (*see also* carry-save adder)
Curnow, H. J., 78
cycle time (*see also* clock cycle time)
of DRAM (*see* dynamic random access memory)
of SRAM (*see* static random access memory)
Cydra 5 (*see* Cydrome Cydra 5)
Cydrome Cydra 5, 340
cylinder, **516** (*see also* disk, magnetic)
Cypress Corporation
Cypress CY7C601 microprocessor, 84, 493

D

DASD, **514** (*see also* direct-access storage device; disk, magnetic; input/output)
data alignment, 95-96
data antidependency (*see* antidependency)
data area, global, **116**
data dependences (*see* hazard, data)
vector processing and, 375 (*see also* vector processor, data dependences)
Data General Nova, 560
data hazard (*see* hazard, data)
data integrity, **520** (*see also* input/output, reliability)
data references, 123-124, 132-133
data transfer, 79, 135
atomic swap instruction, **E-9**–E-10
Endian option, **E-9**–E-10 (*see also* Big Endian; Little Endian)
non-aligned, **E-12**–E-13
in RISC architectures, E-5
data transfer operator, **103**
data trunks, **295**
data-only cache (*see* cache, data-only)
data parallelism, **573**
datapath, **201**
control and, 227
data from, 205, 206 (fig.)
design, 204, 207
DLX architecture and, 221
microinstructions and, 208-209, 211, 214
data rate, **511**, **512** (*see also* input/output, devices)
DAXPY (*see* vector processor, Linpack benchmark)
DEC (*see* Digital Equipment Corporation)
decimal arithmetic (*see* arithmetic, decimal)
decimal operations, 15, **103** (*see also* arithmetic, decimal)
decoupled architecture, 321
defects per unit area, 59-60
deferred addressing (*see* addressing mode, memory indirect)
deferred branching (*see* branch, delayed)
definitions, front endsheet
DeLagi, B., 127
delay slot, **268** (*see also* branch-delay slot; load delay slot)
delayed branch (*see* branch, delayed)
delayed load, **268**, 339
denormal, **A-14** (*see also* arithmetic, denormals and)
Dent, B. A., 127
dependences 264, 269 (fig.), 287 (*see also* hazard; vector processor, data dependences)
anti- (*see* vector processor, antidependence)
output (*see* vector processor, output dependence)
true data (*see* vector processor, data dependences)
vector processing and (*see* vector processor, data dependences)
depth (*see* pipelining, depth of a pipeline)
description language, 141-142, inside back cover
descriptor table, **446** (*see also* virtual memory, page table; virtual memory, Intel 80286/80386 and)
design, computer, 8, 13
complexity and time, 15-16
computer-aided, **580**
high-performance, 34
low-cost, 34
tradeoffs, 8, 14
trends and, 16-17
designer, computer (*see* architect, computer)
detailed measurements (*see* instruction set, measurements)
device level select (DLS), **553** (*see also* disk, magnetic, IBM 3990 storage subsystem and)

DG (*see* Data General)
Dhrystone, 28, 45 (*see also* benchmarks, synthetic)
die, 55-58 (*see also* wafer)
  area, 59-60, 61
  cost of, 55, 59-60, 62, 84, 85
  photographs of, 58
  testing, 60
    cost of, 55, 60, 62, 84
  yield, 59-61, 62, 80
difficulties in implementing pipelines (*see* pipelines, difficulties in implementation)
Digital Equipment Corporation (DEC), 15
  DECstation, 19
  DECstation 3100, 68, 167, 190, D-8–D-9
  PDP-8, 91
  PDP-10, 93
  PDP-11, 93, 104, 127-128, 131-132, 142, 187, 480-481, 531, 561
    bus of, 531
    Unibus and (*see* bus, Unibus)
  VAX, 25, 91, 93, 97, 101-102, 103-104, 123, 128-129, 140, 147, 169-172, 187-188
    addressing modes, 144-147, 145 (fig.), 169 (fig.)
      usage, 169-171, 170
    condition codes, 147
    data types, 143
    floating-point arithmetic on, A-59
    instruction mixes, 171-172
    instruction set, 142-144, 146 (fig.), 147 (fig.) (see also Digital Equipment Corporation, VAX, user instruction set)
      format, 141, 144-145, 147
      instruction length, 145 (fig.), 147 (fig.)
      usage measurement, 140, 168, 169-172, 186 (fig.), C-2
    interrupts, 215 (fig.), 218 (fig.), 219
    operand specifiers (*see* Digital Equipment Corporation, VAX, addressing modes)
    operations on, 147
    registers, 143-144
    summary of, E-23
    user instruction set, B-2–B-5
      branch, jump, and procedure call instructions, B-3–B-4
      decimal and string instructions, B-4–B-5
      integer and floating-point logical and arithmetic instructions, B-1–B-2
      queue instructions, B-5
      variable-length bit field instructions, B-5
  VAX-11/780, 13, 19, 28, 29, 142, 187-188
    address space, 441
    cache in, 414-416 (*see also* cache, VAX-11/780 and)
    instruction-prefetch buffer in, 450
    memory hierarchy of (*see* memory hierarchy, VAX-11/780 and)
    page-table entry of (*see* page-table entry)
    time distributions on, D-2–D-3
    translation-lookaside buffer (*see* virtual memory, translation-lookaside buffer)
    virtual memory in (*see* virtual memory, VAX-11/780 and)
    write buffers on, 413, 416, 477, 483
  VAX-11/785, 13
  VAX 8550, 28
  VAX 8600, 13, 28. 329. 337 (*see also* pipelining, VAX 8600 and)
    EBox, 328-332
    FBox, 328-332
    IBox, 328-332, 334
    IFetch, 230-238, 330
Digital Equipment Corp., VAX 8600 (*continued*)
    MBox, 328-331, 333
    Opfetch, 329-334
  VAX 8700
    frequency of process switches on, 439 (fig.)
  VAXstation 2000, 68
  VAXstation 3100, 68
digital signal processor (DSP), **580**
direct (absolute) addressing, **98** (*see also* addressing mode)
discount (*see* cost)
display (*see* cathode ray tube; graphics displays)
direct-access storage devices, **514** (*see also* disk, magnetic; input/output)
direct mapped, **408** (*see also* cache, direct-mapped)
direct memory access (DMA), **534**-535 (*see also* input/output, DMA and)
directory based, **467**-468 (*see* cache, coherency)
dirty, **413** (*see also* cache, write back)
dirty bit, **413** (*see also* cache, write back; virtual memory, dirty bits and)
disk, 6, 19, 20, 29 (*see also* disk, magnetic; disk, optical)
  growth rule, front endsheet, 17
  storage, 3, 19
  technology, 17
disk array, 520-521
  availability of, 520-521
  reliability of, 520-521
disk cache, 537 (*see also* input/output, interfacing to an operating system)
disk controller, **516** (*see also* disk, magnetic)
disk drive (*see* disk, magnetic)
disk-growth rule, front endsheet, 17
disk, magnetic, 514-520, 561
  access time gap and, 518 (fig.), 519
  array of (*see* disk array)
  capacity of, 517 (fig.), 518, 547 (*see also* maximum areal density)
  characteristics of, 515-516, 517 (fig.)
  comparison of four manufacturers, 517 (fig.)
  cost of, 556-557
  cost versus access time, 518 (fig.)
  data rate of, 514 (fig.), 517 (fig.)
  extended storage (ES), **519**
  future of, 518-519, 561
  IBM 3990 storage subsystem and, 546-554, 567
    changes in response time with improvements in 3380D, 553 (fig.)
    channels and, 548, 554
    channel program for, 549
    control hierarchy, 547-549
    data-transfer hierarchy, 547, 548 (fig.), 549
    DLS and, 552-553
    DPR and, 552-553
    head of string, 549, 552
    IOCB of, 549
    RPS, 551, 552, 554
    speed-matching buffers of, 549
    storage director of, 549
    summary of, 553-554
    tracing a disk read, 549-553
  I/O benchmarks for, 510-512 (*see also* input/output, performance)
    file system, 512
    supercomputer, 510-511
    transaction processing, 511-512
    TP-1, 511, 511 (fig.)
  organization of, 515 (fig.)

disk magnetic (*continued*)
seeks and, **516**, 557-558
average seek time, 516, 557, 563
formulas for, 557, 558 (fig.)
seek distance measurements, 559 (fig.)
versus seek distance, 558 (fig.)
solid state disks (SSDs), **519** (*see also* dynamic random access memory)
disk, optical, 519-520
write-once misperception, 519
displacement (based) addressing mode (*see* addressing mode)
Ditzel, D. R., 129, 130, 189
division (*see* arithmetic, division, floating-point; arithmetic, division, integer)
DLS (*see* device level select)
DLX, 117, 122-123, 160-167, 179-183, 188
addressing mode usage, 179-180
alignment, 221, 231
bus, 200 (fig.), 201, 211, 230
control (*see* control, DLX and)
datapath, 221
instruction mixes, 180-183
instruction set, 161-166, 165 (fig.). E-4–E-6
arithmetical logical instructions, 163
branch instructions, 203, 224 (fig.), 230 (fig.), 234-237
common extensions to, E-9–E-12
control-flow instructions, 163, 164 (fig.), 183
format, 166 (fig.)
jump instructions, 222-225
load and store instructions in, 161-163, 203
usage measurement, 179-183, 181 (fig.), 186 (fig.), C-5
load byte, 232, 235
machines related to, 166
miss rates for, 482 (fig.)
pipelining (*see* pipelining, DLX and)
registers, 161-162
register window benefits on, 453 (fig.)
states, 205, 221-224 (figs.), 225-226
summary of, E-2
superscalar (*see* superscalar)
time distribution on, D-8–D-9
vector processing and (*see* vector processor, DLXV and)
DLXV, 353 (*see also* vector processor, DLXV)
DMA (*see* direct memory access)
Doherty, W., 560, 562
done bit, **534** (*see also* input/output)
double precision (*see* arithmetic, precision)
double-extended precision (*see* arithmetic, precision)
double rounding, **A-29** (*see also* arithmetic, rounding)
doubleword, **95**
DPR (*see* dynamic path reconnection)
DRAM (*see* dynamic random access memory; memory, DRAM)
DRAM-growth rule, front endsheet
DRAM-specific interleaving for improving main memory performance, 431-432 (*see also* memory, interleaved)
drive (*see* disk, magnetic)
dual instruction mode
in Intel 860, E-22
dual-issue, 322, 340
dynamic address translation (*see* virtual memory, address translation)
dynamic branch prediction (*see* dynamic hardware branch prediction)
dynamic detection of memory hazard (*see* hazard, memory, dynamic detection of)
dynamic hardware branch prediction, 307-314, 339-340 (*see also* branch-prediction schemes)
dynamic measurements (*see* instruction set, measurements, dynamic)
dynamic path reconnection (DPR), **552** (*see also* disk, magnetic, IBM 3990 storage subsystem and)
dynamic random access memory (DRAM), 16, 17, 29, 425-427, 431-432 (*see also* static random access memory; memory, DRAM; virtual memory)
capacity of, 426, 431
cost of, 556-557
cost versus access time for, 518 (fig.)
cycle time of, 426, 432 (fig.)
growth rule, front endsheet, 17
interleaving and, 431-432
performance increase, 426 (fig.), 427 (fig.)
static column, **431**, 487
solid state disk and, **519**, 564
times of, 426 (fig.)
video, 523 (*see also* graphics displays)
dynamic scheduling **291**, 290-313, 321-322, 339-340
multiple instruction issue and, 321-322
reducing branch penalties with dynamic hardware prediction, 307-314 (*see also* branch-prediction schemes; dynamic hardware branch prediction)
scoreboard approach (*see* scoreboard)
Tomasulo algorithm (*see* Tomasulo algorithm)

E

Earle latch (*see* latches)
early restart, **458** (*see also* cache, miss)
EBCDIC, 109
EBox (*see* Digital Equipment Corporation, VAX 8600)
Eckhouse, R., 188
Eckert, J. P., 23-25, 241
Eckert-Mauchly Computer Corporation, 25
EDSAC (Electronic Delay Storage Automatic Calculator), 24, 241-242
EDVAC (Electronic Discrete Variable Automatic Computer), 23-24
Edwards, D. B. G., 26
Eggers, S., 471, 487, 488
elapsed time, 35-36, 67, 69, 72
Emer, J., 79
emulation, **242**
empty slots (*see* delay slots)
encoding, **210**-211, 235 (*see also* addressing mode)
Encore Multimax multiprocessor, 589 (*see also* multiprocessor)
Endian option (*see* data transfer)
Engelbart, D., 560
ENIAC (Electronic Numerical Integrator And Calculator), 23-24
entry time, **508** (*see also* input/output, transactions and)
error bit, **534** (*see also* input/output)
ES (*see* extended storage)
ESA/370 (*see* International Business Machines, IBM ESA/370)
ETA-10 (*see* Control Data Corporation, ETA-10)
Ethernet, **526** (*see also* networks)
evaluation of vector performance (*see* vector performance, analyzing)
even/odd multiplier, **A-45** (*see also* arithmetic)
exceptions, **216** (*see also* arithmetic, exceptions; interrupts)
execution, 252, 294, 301, 330
in a pipeline, 252, 285, 294, 301, 330
out-of-order (*see* out-of-order execution)

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 772

execution (*continued*)
simulation, 289
mode of, 8, 10, 29
execution time, 5-7, 27, **35**, 28, 29 (*see also* response time; performance; mean)
average instruction, 77
locality of reference and, 11-12 (*see also* locality)
normalized, 52-53, 83
performance and, 6, 35, 40-45, 48-49, 71-72, 81
speedup and, 10
total, 50, 83
weighted, 51, 84
executive process, **440** (*see also* virtual memory, processes and)
explicit communication (*see* communication, explicit)
exponent, **A-13** (*see also* arithmetic, exponents and)
exponent field, **A-13** (*see also* arithmetic, exponents and, exponent field)
extended storage (ES), **519**

F

Fabry, R., 485, 488
false sharing, **469** (*see also* cache, coherency)
fast page mode of DRAM, **432**
faults, **216** (*see also* interrupts)
FBox (*see* Digital Equipment Corporation, VAX 8600)
fetch on write, **413** (*see also* cache, write miss)
fields, **209**
FIFO (*see* block replacement, first-in–first-out)
file cache, 537 (*see also* input/output, interfacing to an operating system)
file server
versus workstation, 500
file systems, 512
file system I/O benchmark (*see* benchmark; input/output, performance; disk, magnetic, I/O benchmarks for)
filled slots (*see* branch-delay slots)
finite state diagram, **204**, 206
for the DLX, 205, 220
interrupts and, 217
firmware (*see* microprogramming)
first-in–first-out (FIFO), **412** (*see also* block replacement, first-in–first-out)
first part done (FPD), **219**-220
first reference misses, **419** (*see also* cache, miss, compulsory)
Fisher, J., 340
fixed-field decoding, **202**
fixed point, **A-12**, A-58 (*see also* arithmetic, integer)
Flemming, P. J., 79
floating point (FP), 15, 19 (*see also* arithmetic, floating-point)
arithmetic (*see* arithmetic, floating-point)
CDC 6600 and, 291-293
floating-point operations per second (FLOPS), 360-361 (*see also* vector processor, performance)
IBM 360/91 and, 299-300
millions of floating-point operations per second (MFLOPS), 43-44, 74-75, 78, 83, 86, 383, 386 (*see also* vector processor, performance)
native, 43-44, 81, 83
normalized, 43-44, 83
overflow (*see* arithmetic, exception, overflow)
floating-point arithmetic, quadruple precision, **E-17**
floating-point compares, 106-107
floating-point format (*see* arithmetic, IEEE standard and)
floating-point instructions (*see* floating-point operations)
floating-point operations, 14, 103, 284-290, 318-319 (*see also* pipelining, DLX and, floating-point)
implicit conversions, **E-9**, E-11
in RISC architectures, E-6
overlapped, in SPARC, E-16
floating-point operations per second (*see* floating point, floating-point operations per second)
floating-point operator, **103**
floating-point references, 119
floating-point pipeline (*see* pipelining, DLX and, floating-point)
floating-point register, 114, 118-119, 124
floating-point stalls, 290
floating-point standard, 109
Floating-Point Systems AP-120B, 340
FLOPS (*see* floating point, floating-point operations per second)
Flynn bottleneck, **351**, 352 (*see also* vector processor)
Foley, P., 189
format field, **211**
format of instructions (*see* instruction syntax)
FORTRAN, 119, 130
Absoft System V88 2.0a compiler 83
F77 compiler, 126
FORTRAN 8X, 581
FORTRAN 77, 581
forwarding, **261**-265, 269, 286, 339
Foster, C. C., 129
FP (*see* floating point)
FPD (*see* first part done) on VAX
fraction, computation time, 10
enhanced, 10
fragmentation and reassembly, **527** (*see also* networks)
frame address (*see* memory hierarchy, block)
frame buffer, **521** (*see also* graphics displays)
freezing the pipeline, 273, 334
Freitas, D., 189
frequency distributions (*see* instruction set, measurements)
full adders, **A-2** (*see also* arithmetic)
Fuller, S. F., 78, 80
fully associative, **408** (*see also* cache, fully associative)
functional requirements (*see* requirements, functional)
functional units, 255-258, 291-298, 300-305, 318-319, 323-324, 338
multiple, 284-285, 338, 346
vector processing and (*see* vector processor, functional units)
functional unit status, 295, 296-298 (figs.), 303-305 (figs.)
FutureBus, 532, 532 (fig.) (*see also* bus)
future file, **288** (*see also* out-of-order completion)

G

Gagliardi, U. O., 129
Gajksi, D., 589
Garner, R., 190
gateway, **527** (*see also* networks)
gather, **380** (*see also* vector processor, sparse matrices)
GCD (*see* greatest common divisor)
Gelsinger, P., 188
general-purpose register (GPR) architecture (*see* register, general-purpose register architecture)
generate, **A-32** (*see also* arithmetic; carry)
generation, computer, 26
geometric mean (*see* mean, geometric)
Gibson, D. H., 78
Gibson, J. C., 77, 78, 80
Gibson mix, 77, 78, 80

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 773

gigaflop (*see* floating point, millions of floating-point operations per second)
Gill, J., 130
Gill, S., 24
global address space, **446** (*see also* virtual memory, processes and)
global data area (*see* data area, global)
global miss rate, **461** (*see* cache, miss; cache, two-level caches)
Gnu C compiler, 67, 69-70, 79, 85
Goldschmidt's algorithm, **A-24**–A-25, A57
Goldstine, H. H., 23-25
Gonter, R. H., 129
Goodman, J., 487, 488
Gottlieb, A., 589
GPR (*see* register, general-purpose register architecture)
gradual underflow (*see* arithmetic, exceptions, underflow)
graphics instructions in Intel 860, E-20
graphics displays, 521-525, 560, 561
  color map, **523**, 524 (fig.)
  cost of, 523-524
  frame buffer, 521, 522 (fig.)
  future directions in, 525-526
  hidden surface elimination, **525**
    z-buffer approach to, **525**
  performance demands of, 524-525
  tasks and their performance requirements, 525 (fig.)
  video DRAMs, 524, 525
gray-scale displays, **521** (*see also* graphics displays)
greatest common divisor (GCD), **373** (*see also* vector processor, data dependences)
growth rules (*see* disk, growth rule; dynamic random access memory, growth rule)
Gross, T. R., 335, 339

H

half adders, **A-2** (*see also* arithmetic)
halfwords, **95**
Hansen, C., 189
hard disk (*see* disk, magnetic)
hard drive (*see* disk, magnetic)
hardware branch prediction **291** (*see also* dynamic hardware branch prediction)
hardware, **13** (*see also* balance, software and hardware)
  "smaller is faster," 18
  industry growth and, 21
hardwired control (*see* control, hardwired)
harmonic mean (*see* mean, harmonic)
Harvard University, 24-25
Hauck, E. A., 127
hazard, **257**-258, 278 (*see also* dependences; vector processor, data dependences)
  branch, **270**-272, 280, 307
    handling on VAX 8800, 331-332
    (*see also* branch, penalty)
  control, **257** (*see also* hazard, branch)
  data, **257**, 260-269, 282, 283-284, 286-290, 291-298, 300-306, 346 (*see also* vector processor, data dependences; pipelining)
    handling on VAX 8800, 331-332
  detection, 268-269, 334 (*see also* branch, penalty)
    VAX 8600 and, 328-329
    DLX and data hazard detection 268-269
    DLX and structural hazard detection, 292
    floating point and, 286
    overlapped integer and floating-point instructions and, 285
    scoreboard and, 293-298
hazard, detection (*continued*)
    Tomasulo algorithm and, 300, 302-306
  dynamic detection of memory hazards, 291-298, 300-306, 339
  RAW, **264**, 286, 294, 297, 301, 331
    vectors and (*see* vector processor, data dependences)
  memory, dynamic detection of, 339
  structural, **257**, 258-259, 284, 286, 294, 300
    CPI and, 260
    DLX and, 289, 291-292, 300
    superscalar machine and, 319
  true data dependences (*see* vector processor, data dependences)
  vector processing and, 375 (*see also* vector processor, data dependences)
  WAR, **264**, 286, 293-295, 304
  WAW, **264**, 287, 293-295, 304
hazards, reducing (*see* hazard, detection; branch, penalty, reduction)
head of string, **549** (*see also* disk, magnetic, IBM 3990 storage subsystem and)
Henly, M., 79
Hennessy, J. L., 130, 189
Hewlett-Packard
  Precision, 167, 190
hidden surface elimination, **525** (*see also* graphics displays)
higher-radix multiplication, **A-43**, A-50 (*see also* arithmetic, integer, speeding up multiplication)
high-level language, 16, **111**, 115-116, 121, 124, 127-129, 131, 135
High-Level Language Computer Architecture (HLLCA), **129**-130
high-performance design (*see* design, high-performance)
Hilfinger, P., 189
Hill, M., 421, 424, 481, 486-487, 489
Hillis, D., 577, 589, 590
  bet with Bell, 590
history, computer, 23-27
history file, **288** (*see also* out-of-order completion)
hit, **404** (*see also* memory hierarchy, hit; cache, hit)
Hitachi S810/20, 74
hit rate, **404** (*see also* memory hierarchy, hit rate; cache, hit)
hit time, **405** (*see also* memory hierarchy)
Hopkins, M. E., 130
horizontal microcode (*see* microcode, horizontal)
horizontal microinstruction (*see* microcode, horizontal)
Hough, D., 190
How is a block found? (*see* block identification)
HP (*see* Hewlett-Packard)
Hudson, E., 189

I

i860 (*see* Intel Corporation, i860)
IAS (Institute for Advanced Study) (*see* Princeton University)
IBox (*see* Digital Equipment Corporation, VAX 8600)
IBM (see International Business Machines Corporation)
IC (*see* integrated circuit)
ideal performance in pipelining, 258-259
identification, block (*see* block identification)
IEEE (*see* arithmetic, IEEE standard and)
Ifetch (*see* Digital Equipment Corporation, VAX 8600)
Illiac IV, 554, 555 (fig.), 573, 589, 591
immediate (literal) addressing mode (*see* addressing mode)
IMP (*see* interface message processor; networks)
implementation, 13
  hardware, 14, 21

implementation (*continued*)
performance evaluation and, 78-79
software, 14
technology of, 16
implicit communication (*see* communication, implicit)
implicit conversions (*see* floating-point instructions)
imprecise interrupt (*see* interrupts, imprecise)
improving performance of vector processors (*see* vector processor, improving performance)
in-order instruction issue, 291
index, 47, 98
index addressing mode (*see* addressing mode, scaled)
indexed addressing mode, **98** (*see also* addressing mode)
index field, **410** (*see also* cache, set associative)
index vector, **380** (*see also* vector processor, sparse matrices)
indirect addressing mode (*see also* addressing mode, register deferred; addressing mode, memory indirect)
induction variable elimination, **114** (*see also* optimization)
inexact exception, **A-30** (*see also* arithmetic, exceptions
infinity, A-13, A-14, A-19 (fig.), A-22, A-30, A-60 (*see also* arithmetic, rounding and; not a number)
infinite precision, A-22
initiation rate, **358** (*see also* vector processor, initiation rate)
input/output (I/O), 6, 11, 15 (fig.), 17, 22, 499-501, 554-561 (*see also* disk, magnetic; graphics displays; networks; bus)
bandwidth (*see* input/output, throughput)
benchmarks (*see* input/output, performance; disk, magnetic, I/O benchmarks for)
CPU time and, 499
DMA and, 534-535, 537, 561
virtual, 537, 538 (fig.)
IBM and, 546
IBM 3990 storage subsystem (*see* disk, magnetic, IBM 3990 storage subsystem and)
idle time and, 500-501 (*see also* input/output, people and)
designing a system for, 539-546
devices, 512-514, 560-561 (*see also* disk, magnetic; graphics displays; networks; bus)
categorized by behavior, partner, and data rate, 513 (fig.)
data rate, **511**, **512**, 514
examples of, 513 (fig.)
keyboards, 513
fallacies and pitfalls of, 554-559
history of, 560-561
importance of (*see* input/output, system performance and)
interfacing to the CPU, 533-535 (*see also* input/output, DMA and)
delegating I/O responsibility from the CPU, 534-535
fallacy of moving functions from CPU to I/O, 555-556
interfacing to an operating system, 535-538
caches causing problems with, 535-537
caches helping with, 537-538
disk cache, 537-538
effectiveness of, 538 (fig.)
stale data and, 535-536
virtual memory and, 537
latency (*see* input/output, response time)
operating systems and, 535 (*see also* input/output, interfacing to an operating system)
overlapping (*see* input/output, system performance and)
people and, 508-509, 513, 560
peak I/O rates for, 513 (fig.)
transactions per hour versus computer response time, 510 (fig.)
performance, 506-512, 539-546, 555-556 (*see also* input/output, response time; input/output, throughput; benchmark; disk, magnetic, I/O benchmarks for)
input/output, performance (*continued*)
cost/performance, 539, 555
producer-server model of response time and throughput, 506 (fig.), 508 (fig.)
reliability, 520-521
response time (latency), 506, 507, 509 (fig.), 560
disk array and, 520
graphics displays and, 522-524
IBM 3380D and, 553 (fig.)
magnetic disk and, 507 (fig.)
networks and, 528
transaction time and, 509
versus throughput, 507 (fig.), 507-509
versus transactions per hour, 510 (fig.)
supercomputers and, 529, 564
system performance and, 501-506, 555-556
Amdahl's Law and, 500, 555, 559
cost/performance, 555
time formulas for, 502-506
overlapped execution of I/O, 502 (fig.), 502-506
throughput (bandwidth), 506, 507, 544
bus and, 532
disk array and, 520
graphics displays and, 522-524
magnetic disk and, 507 (fig.)
networks and, 528
versus response time, 507 (fig.), 507-509
transactions and, **508**, 509 (fig.)
transaction processing (TP), **511**-512
transaction time, **508**
entry time, **508**, 509 (fig.), 560
system response time, **508** (*see also* input/output, performance, response time)
think time, **508**. 509 (fig.), 560
transactions per hour versus response time, 510 (fig.)
user transaction, 509 (fig.)
types (*see* input/output devices)
Institute for Advanced Study (IAS) (*see* Princeton University Institute for Advanced Study)
instruction (*see also* instruction set)
architecture (*see* instruction set, architecture)
average execution time, 77
control (*see* control; control-flow instructions)
count, 36-42, 72-73, 94, 99, 121, 123
optimization and, 119-120
density, **94**
encoding, 94, 102-103
fetch and decode rate, 351
format (*see* instruction set)
of RISC architectures, E-3
frequencies (*see* instruction set, measurements)
interruption and restart, 279-282, 287-289, 332
issue, **266**, 286-289, 292-296, 300-306, 339-340 (*see also* dual-issue; multiple instruction issue)
issue more than one instruction, 318-320
multiple instruction issue with dynamic scheduling, 321-325
scoreboard and, 292-296, 293 (fig.)
superscalar machines and, 318-320
stalls and, 284
measurements (*see* instruction set, measurements)
mix, 39, 45, 73, 77
path length, 36 (*see also* instruction count)
parallelism, 314-328, 340-341 (*see also* vector processor)
increasing with loop unrolling, 315-318
increasing with software pipelining and trace scheduling, 325-328

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 775

instruction (*continued*)
reference, 124
scheduling, **267**-268, 274-278, 339
set (*see* instruction set)
size, 103
static, 12
status, 295, 296-298 (figs.), 303, 305 (fig.), 308 (fig.)
syntax, 141 (*see also* instruction set, architecture)
instruction-level parallelism (*see* instruction, parallelism)
instruction-only cache (*see* cache, instruction-only)
instruction-prefetch buffer, **449**-450, 484 (fig.)
summary of, 484 (fig.)
VAX-11/780 and, 450 (fig.)
instruction set (*see also* instruction; DLX; Intel Corporation, 860; MIPS Computer Corporation, R3000; Motorola Corporation, 88000)
architecture, 13, 16, 17, 37, 90-94
comparison, 70
complications (*see* pipelining, difficulties in implementation, instruction set complications)
control (*see* control; control-flow instructions)
frequencies (*see* instruction set, instruction frequencies; instruction set, measurements)
instruction frequencies (*see also* instruction set, measurements)
DEC VAX, 172 (fig.)
DLX, 181 (fig.)
IBM 360, 175 (fig.)
Intel 8086, 178 (fig.)
measurements, 139-141, 142, 167-168, 184, 185 (fig.), 186 (fig.), D-2
DEC VAX, 169-172, 172 (fig.)
detailed measurements, C-2
DLX, 179-183, 181 (fig.)
detailed measurements, C-5
dynamic, 90, **139**, 140 (fig.)
comparisons of, by architecture, 186 (fig.)
frequency distributions, D-2, D-3 (fig.)
IBM 360, 173-176, 175 (fig.), 185 (fig.), 186 (fig.)
detailed measurements, C-3–C-4
Intel 8086, 176-178 (fig.), 186 (fig.)
detailed measurements, C-4
static, **139**
time distributions, 139, 171, 184-185, D-2–D-9
8086 in an IBM PC, D-6–D-8
DLX relative, D-8–D-9
IBM 370/168, D-4–D-6
VAX-11/780, D-2–D-3
performance and, 36-37, 39, 67
processor (ISP) (*see* instruction set, architecture
usage (*see* instruction set, measurements)
user (see Digital Equipment Corporation, VAX, user instruction set)
instruction set processor (ISP) (*see* instruction set, architecture)
integer arithmetic (*see* arithmetic, integer)
integer compares, 106-107
integer multiply and divide
in RISC architectures, E-8–E-9
signed and unsigned, in SPARC, **E-17**
integer operations, 15
integer overflow (*see* arithmetic, exception, overflow)
integer pipeline (*see* pipeline, DLX and, integer)
integer register, 114, 117-119, 124, 136
integer variables, **109**, 117
integrated circuit (IC), 3, 5, 13, 17, 26
cost of, 54-58
yield, 59, 81
Intel Corporation
Intel 4004 and 8008, 188
Intel 432, 125
Intel 8080, 153, 188
Intel 8088, 188
Intel 80x86, front endsheet, 153, 188, 449
Intel 8086, 91, 97, 104, 141, 153-160, 176-179, 188, 445
addressing modes, 155-156
usage, 177 (fig.)
address space, 154
compatibility mode, 153
flaws, 184
instruction mixes, 176-178
instruction set, 153-160, 158 (fig.), B-9–B-12
arithmetic and logical instructions, B-10
control instructions, B-11
data transfer instructions, B-12
formats, 141, 157, 159 (fig.), 160 (fig.)
string instructions, B-12
usage measurement, 156 (fig.), 168, 176-179, 186 (fig.), C-4
interrupts, 215 (fig.)
operations on, 156-160
postbyte encoding, 160 (fig.)
registers, 153-155, 154 (fig.)
summary of, E-23
time distribution on, D-6–D-7
Intel 80186, 153, 188
Intel 80286, 153, 188, 445-446, 448-449
call gates on (*see* call gate)
descriptor table, **446** (*see also* virtual memory, page table; virtual memory, Intel 80286/80386 and)
protection on (*see* virtual memory, Intel 80286/80386 and)
virtual memory on (*see* virtual memory, Intel 80286/80386 and)
Intel 80386, 153, 188
protection on (*see* virtual memory, Intel 80286/80386 and)
virtual memory on (*see* virtual memory, Intel 80286/80386 and)
Intel 80486, 56, 58, 84, 153, 188
Intel 860, 84, 167, 190, 340, 493, E-2
instruction set, E-5-E-6
common extensions to DLX instructions, E-10–E-11
unique, E-19–E-23
Intel i860, 493 (*see also* Intel Corporation, Intel 860)
Intel Hypercube multicomputer, 589
intelligent devices, **560** (*see also* bus)
intelligent peripheral interface (IPI), **531**, 532 (fig.), 560 (*see also* bus)
interface message processor (IMP), **527** (*see also* networks)
interference graph, **113**
interleaved memory, **429** (*see also* memory, interleaved)
interleaving factor, **429** (*see also* memory, interleaved)
interlocked loads instruction (*see* load interlock)
interlock (*see* pipeline interlock; hazard, data; load interlock)
internal fragmentation, **437** (*see* virtual memory, page size and)
internal storage, 90-92
International Business Machines Corp. (IBM), front end sheet, 15, 25, 80
IBM 3090, 547
storage (*see* disk, magnetic, IBM 3990 storage subsystem and)
IBM 3090-600S, 75
IBM 360, 16, 25, 77, 91, 93, 104, 127-128, 148-152, 172-176, 186-187, 242, 485, 557

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 776

International Business Machines, IBM 360 (*continued*)
addressing modes, 149-150
usage, 173-174
flaws, 183-184
IBM 360/85 (*see* International Business Machines Corp., IBM 360/85
instruction mixes, 175-176
instruction set, 148-150, 151 (fig.), 152 (fig.), B-6–B-9
formats, 149-151
usage, 174 (fig.)
register–indexed (RX), 149-150, 174
branches and special loads and stores, RX format, B-8
integer/logical and floating-point instructions, RX format, B-7
register–register (RR), 149-150, 174
branches and status setting R-R instructions, B-7
integer/logical and floating-point R-R instructions, B-6
register–storage (RS), 149-150, 174
RS and SI format instructions, B-8
storage–immediate (SI), 150, 174
RS and SI format instructions, B-8
storage–storage (SS), 150, 174, 177
SS format instructions, B-9
usage measurement, 172-176, 185 (fig.), 186 (fig.), C-3–C-4
Shustek's thesis on, 172-173, 185, 187
interrupts, 215 (fig.), 219-220
operations on, 151-152
registers, 141, 149-150, 174, 177
summary of, E-23
IBM 360/85, 26, 80, 486
IBM 360/91, 299-300, 339
IBM 360/IBM 370 (*see* International Business Machines Corp., IBM 360; International Business Machines Corp., IBM 370)
IBM 370, 148, 186-187, 394, 485 (*see also* International Business Machines Corp., IBM 360)
floating-point system on, A-59
IBM 370/158, 77, 78
IBM 370/168, D-4–D-6
IBM 370-XA, 148, 187
IBM 3990 storage subsystem (*see* disk, magnetic, IBM 3990 storage subsystem and)
IBM 701, 25, 26
IBM 704, 338
IBM 801, 189, 190
IBM 7030, 104, 338
IBM 7090, 129, 242
IBM ESA/370, 148
IBM PC, 34, 176, 184, 188, D-6–D-7
bus of, 531
IBM PC-AT, 531
IBM PL.8 compiler, 130
IBM RP3 multiprocessor, 589 (*see also* multiprocessor)
IBM RT-PC, 93, 190
IBM Stretch (7030), 77
IBM System/360 (*see* International Business Machines Corp., IBM 360)
IBM System/370 (*see* International Business Machines Corp., IBM 370)
MIPS definition and, 78
Stretch (*see* International Business Machines Corp., IBM 7030)
interprocedural register allocation, **453** (*see also* register windows)
interrupt-driven I/O, **534** (*see also* input/output, interfacing to the CPU)
interrupts, **214**-220
8600 and, 332-334
arithmetic overflow and, 214-215, 217 (fig.), 218 (fig.), 241
comparison on four computers, 215 (fig.)
DLX and, 229, 235, 237
history of, 241
how control checks for interrupts, 217-218
page faults and, 215, 217 (fig.), 218 (fig.)
pipelining and, 261, 276, 279-282 (*see also* interrupts, imprecise; interrupts, precise)
virtual memory and, 440
what's hard about interrupts, 218-220
interrupts, imprecise, 287-288
interrupts, precise, **280**, 334, 339
invalid exception, **A-30** (*see also* arithmetic, exceptions)
inverted page table, **435** (*see also* virtual memory)
I/O (*see* input/output)
I/O bandwidth (*see* input/output, performance, throughput)
I/O bus, **529** (*see also* bus)
IOCB (*see* I/O control block)
I/O control block, **534-535**, **549** (*see also* input/output, interfacing to the CPU; disk, magnetic, IBM 3990 storage subsystem and)
I/O controllers, **534** (*see also* input/output, interfacing to the CPU)
I/O latency (*see* input/output, performance, response time)
I/O processor, **534** (*see also* input/output, interfacing to the CPU)
I/O rate, **511** (*see also* input/output, transactions and)
I/O response time (*see* input/output, performance, response time)
I/O throughput (*see* input/output, performance, throughput)
Iowa State University, 24
IPI (*see* intelligent peripheral interface; bus)
ISP (instruction set processor) (*see* instruction set, architecture)
issue (*see* instruction issue)
issue more than one instruction (*see* instruction issue; superscalar)
issued, **266** (*see also* instruction issue)

J

Japanese supercomputers (*see* supercomputers, Japanese)
Jouppi, N., 130
Joy, B., 190
jump, **104**-105, 120 (*see also* branch)
conditional, 23
on the DLX, 222-225

K

Kahn, R., 561, 562
Kane, G., 190
Katz, R., 487, 488
Kelisky, R. P., 560
kernel process, **440** (*see also* virtual memory, processes and)
kernel programs, 43, 45-48, 77
Livermore FORTRAN, 43, 77, 80
Kilburn, T., 26, 432, 485, 489
Kleiman, S., 190
Knuth, D. E., 26-27
Kuck, D., 589
Kung, H. T., 590

L

LAN (*see* local area network; networks)
language
  assembly, 16
  high-level (*see* high-level language)
  programming, 17
language-oriented architecture (*see* high-level language)
Lanigan, M. J., 26
Larus, J., 189
latch delay, 253
latches, 253-255, 339
  Earle latch, 254-255
  latch overhead, 336
latency, **5**, 18 (*see also* execution time, performance)
  access time, 425-426
  cycle time of, 425-426
  I/O latency (*see* input/output, performance, response time)
  performance measures of main memory and, 425
  throughput and, 8
latency, I/O (*see* input/output, performance, response time)
learning curve, 54, 55 (*see also* yield)
least-recently used (LRU), **411** (*see also* block replacement, least-recently used)
Lee, R., 190
length of vector (*see* vector processor, vector length)
Levy, H., 171, 188
limit field, **446** (*see also* virtual memory, page table; virtual memory, Intel 80286/80386 and)
line, **408** (*see also* cache, blocks and)
linear speedup, 576, 585-586, 593, 594
Linpack (*see* vector processor, Linpack benchmark; benchmarks), 28, 45
LISP, support for in SPARC, E-15–E-16
literal addressing mode (*see* addressing mode, immediate)
Little Endian, front endsheet, 95
LIW (*see* long instruction word)
live ranges, **113**
load and store buffers, 301-303, 308 (fig.)
load delay, **268**, 278, 290
load interlock, 267, 269
  in MIPS II architecture, **E-14**
load/store architecture, 39-42, **93, 94**, 124, 337 (*see also* reduced instruction set computer; DLX)
local address space, **446** (*see also* virtual memory, processes and)
local area networks (LAN), **526** (*see also* networks)
locality (*see also* memory hierarchy, principle of locality and)
  principle of (90/10 locality rule), front endsheet, 11-12, **403**
  program, 26
  of reference, 11-12, 18, 20
  spatial, 12 (fig.), 29, **403**
  temporal, 12, **403**
local miss rate, **461** (*see* cache, miss; cache, two-level caches)
lock/unlock operations (*see* synchronization)
lock variables, **471** (*see also* cache, coherency, synchronization)
logic
  operations, 15
  technology, 17 (fig.)
long-haul networks, **527** (*see also* networks)
long instruction word (LIW), **323**, 340
  in Intel 860, E-22–E-23
loop, 114-115 (*see also* loop unrolling)
  branch (*see* branch, loop)
  software-pipelined (*see* pipelining, software-pipelined loop)
loop-carried dependences, **372** (*see also* vector processor, data dependences)
loop unrolling, **316**, 325-326, 340
  increased instruction-level parallelism with, 315-318
  superscalar DLX and, 319-320
  unrolled loop, 316-318, 320, 326, 327 (fig.)
loosely-coupled MIMD (*see* multicomputer)
low-cost design (*see* design, low-cost)
lower level, **404** (*see* memory hierarchy; cache; memory; virtual memory)
Lunde, A., 129
LRU (*see* block replacement, least-recently used)

M

M680x0 (*see* Motorola Corporation)
M88000 (*see* Motorola Corporation, 88000)
macro-, 208
MAD (*see* maximum areal density)
magnetic disk (*see* disk, magnetic)
mainframe, 3-4
  versus minicomputer, 499
main memory (*see* memory, main)
Manchester, University of (*see* University of Manchester)
margin, gross, 64-66, 76, 85
Mark I (University of Manchester), 24
Mark-I, -II, -III, -IV (Harvard University), 24-25
market, computer
  effect on design, 4, 13, 14, 15
marketplace (*see* market, computer)
Markstein, J., 130
Markstein, P. W., 130
Mauchly, J., 23-25, 241
maximally encoded, **212** (*see also* microcode, vertical)
maximum areal density of disks (MAD), **518**-519, 561 (*see also* disk, magnetic)
  MAD formula, 518, 561
maximum vector length (MVL), **364** (*see also* vector processor, vector length)
Mazor, S., 188
MBox (*see* Digital Equipment Corporation, VAX 8600)
McFarland, H., 127
McKeeman, W. M., 128
McMahon, F. M., 78, 79
McNamara, J. E., 81
McNutt, B., 79
mean
  arithmetic, 50-53, 69-70, 78
    weighted, 51, 53, 84
  geometric, 52-53, 72, 78, 83-84
  harmonic, 50, 52, 75, 78, 81
    weighted, 51
mean time to failure (MTTF), **520** (*see also* input/output, reliability)
mean time to repair (MTTR), **520** (*see also* input/output, reliability)
measurements, dynamic, **139** (*see also* instruction set, measurements)
measurements of instruction set usage (*see* instruction set, measurements)
measurements, static, **139** (*see also* instruction set, measurements)
media price (*see* cost)
megahertz (*see* clock rate)
megaFLOPS (*see* MFLOPS)
memory, 5, 13, 14, 15 (*see also* bandwidth cache; dynamic random access memory; static random access memory; memory hierarchy; virtual memory; block identification; block placement; block replacement; write strategy)

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 778

memory (*continued*)
bandwidth, 257, 260, 324, 329 (*see also* memory, organization of)
in vector machines, 361-363, 392
banks, **361**-363 (*see also* memory, interleaved)
bus, 13, 15, 18 (fig.), 29
cell, 18
centralized, **578**
versus distributed, 578-579
consistency, 474 (*see also* cache, coherency)
sequential, 474
weak, 474
deferred addressing mode (*see* addressing mode, memory indirect)
DRAMs and, 16, 17, 29, 425-427 (*see also* dynamic random access memory)
interleaving, 431-432 (*see also* memory, interleaved)
refresh cost, 426
hazard (*see* hazard)
hierarchy (*see* memory hierarchy)
indirect (memory deferred) addressing, **98** (*see also* addressing mode)
interleaved, 429-431
disadvantage of, 430
DRAM-specific, 431-432
interleaving factor, **429**
latency, access time of, **425**-426 (*see also* latency; memory hierarchy, access time)
latency, cycle time, of, **425**-426 (*see also* latency)
magnetic core, 25, 425
main, 19-20, 25, 425-432
bandwidth, 425 (*see also* bandwidth)
latency, 425 (*see also* latency)
mapping, **433** (*see* virtual memory, address translation; virtual memory, Intel 80286/80386 and)
memory-mapped I/O, **533** (*see also* input/output, interfacing to the CPU)
memory-memory architecture (*see* memory-memory architecture)
memory-memory vector machine, **353** (*see also* vector processor, vector machines)
organization of, 427, 428 (fig.) (*see also* memory, interleaved; memory, wider)
performance, 485
CPU-DRAM performance gap, 426, 427 (fig.), 432
increasing with DRAM-specific interleaving, 431-432
pipeline, 259 (*see also* pipelining; load delay; load and store buffers)
read-only (ROM), **205**, 208, 239, 241-242
future of microprogramming and, 241
reference, 93-97, 110, 116-119, 123, 129, 134, 260, 264
CDC 6600 and, 293
computed, **117**
IBM 360/91, 301-303
save/restore, **116**, 118-119
register-memory architecture (*see* register-memory architecture)
software and, 16
stall clock cycles, **224**
stall cycle, **224**
caches and, 416
static random access (SRAM) (*see* static random access memory)
virtual (*see* virtual memory)
wider, 428-429
memory hierarchy, **19**, 18-20, 22, 29-30, 402, 403-407, 484 (fig.) (*see also* cache; cache, two-level caches; memory; virtual memory; virtual memory, translation-lookaside buffer; block identification; block placement; block replacement; write strategy; instruction pre-fetch buffer; register windows)
access time, **405**-406, 420, 425-426 (*see also* cache, access time)
average memory-access time, 405, 407 (*see also* cache, access time; cache, two-level caches)
blocks and, **404**-407 (*see also* cache, blocks and; memory, block; virtual memory, paged; block)
block-frame address, **405**
block-offset address, **405**
fixed block size, 404, 406
miss penalty and block size, 406 (fig.), 423 (fig.)
variable block size, 404, 434 (*see also* virtual memory, segmented)
cache's relationship to, 408 (*see also* cache)
fallacies and pitfalls of, 480-483
history of, 485-487
hit, **404** (*see also* cache, hit)
hit rate, **404** (*see also* cache, hit)
implications of, to CPU, 407
levels (*see* memory hierarchy, lower level; cache; memory; virtual memory)
lower level, **404** (*see also* cache; virtual memory; memory)
main memory's relationship to, 425 (*see also* memory)
miss, **404** (*see also* cache, miss; cache, write miss; virtual memory, page fault)
performance, 405-407, 485 (*see also* cache, performance; virtual memory, performance)
principle of locality and, **403**-404, 484 (*see also* loaclity)
address translation and, 437 (*see also* virtual memory, translation-lookaside buffer)
spatial locality, **403**, 406, 486 (*see also* locality)
cache block size and, 422, 458, 465 (*see also* cache, blocks and)
shared data and, 469
temporal locality, **403**, 406, 486 (*see also* locality)
least-recently used and, 411 (*see also* block replacement, least-recently used)
shared data and, 469
summary of examples of, 484 (fig.)
thrashes and, 420
upper level, **404**-407 (*see also* memory; cache; virtual memory)
VAX-11/780 and, 475-480 (*see also* cache, VAX-11/780 and; virtual memory, VAX-11/780 and)
average number of clock cycles per 780 instruction, 477
miss rates for, versus DLX, 482 (fig.)
miss rates for the VAX-11/780 TLB, 479 (fig.)
misses per hundred instructions for the VAX-11/780 TLB, 479 (fig.)
overall picture of, 476 (fig.)
physical-instruction-buffer address (PIBA), **475**
virtual-instruction-buffer address (VIBA), **475**
write buffers, 413, 416, 477, 483
virtual memory's relationship to, 433 (*see also* virtual memory)
memory deferred addressing mode (*see* addressing mode, memory indirect)
memory indirect (memory deferred) addressing, **98** (*see also* addressing mode)
memory interleaving (*see* memory, interleaved)
memory, main (*see* memory, main)

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 779

memory-mapped I/O, **533** (*see also* input/output, interfacing to the CPU)
memory-memory architecture, **93**-94, 122-124, 128-29, 134
memory-memory vector machine, **353** (*see also* vector processor, vector machines)
memory width (*see* memory, wider)
Metcalfe, R., 560, 562
metric, computer, 14, 18
MFLOPS (*see* floating point, millions of floating-point operations per second)
MHz (megahertz) (*see* clock rate)
micro-, 208
micro-architecture (*see* organization)
microcode, 208, 213 (*see also* control, microprogrammed/microcoded; microprogram)
  compared to macrocode, 238
  horizontal, **212**, 214, 244
  legal status of, as program, 243
  vertical, **212**, 244
microcoded control (*see* control, microprogrammed/microcoded)
microcomputer, 3-4
microinstruction, 208, 228 (*see also* microcode; control, microprogrammed/microcoded)
microinstruction format, 209
  on the DLX (*see* control, DLX and)
  reducing hardware costs with multiple microinstruction formats, 211-212
microprocessor, 4, 16, 75 (*see also* Cypress Corporation; Intel Corporation; MIPS Computer Corporation; Motorola Corporation; National 32032 microprocessor)
  comparison of 188-190
  Intel 80x86 (*see* Intel Corporation, 80x86)
  MIPS R2000 (*see* MIPS Computer Corporation)
  MIPS R3000 (*see* MIPS Computer Corporation)
  Motorola 680x0 (*see* Motorola Corporation)
  Motorola 88000 (*see* Motorola Corporation)
  SPARC (see SPARC)
  "super-", 500
microprogram, 208, 211 (*see also* microcode; control, microprogrammed/microcoded)
  counter, 228
  DLX microprogram (*see* control, DLX and)
  horizontal, **212**, 214, 244
  legal status of, as program, 243
  microprogram memory (*see* control store)
  structure of, 209
  vertical, **212**, 244
microprogrammed control (*see* control, microprogrammed/microcoded)
microprogramming, **208**, 209
  ABCs of microprogramming, 209-210
  (*see also* control, microprogrammed/microcoded)
millions of floating-point operations per second (MFLOPS) (*see* floating point, millions of floating-point operations per second)
millions of instructions per second (MIPS), 17, 40-42, 44, 67
  native, 42, 71, 78
  relative, 42, 72, 77-78
MIMD computer (*see* multiple instruction streams–multiple data streams computer)
minicomputer, 3-4
  PDP-11 (*see* Digital Equipment Corporation, PDP-11)
  VAX-11/780 (*see* Digital Equipment Corporation, VAX-11/780)
  VAX-8600 (*see* Digital Equipment Corporation, VAX 8600)
  VAX 8700 (*see* Digital Equipment Corporation, VAX 8700)
minicomputer (*COntinued*)
  versus mainframe, 499
  versus workstation, 499
minimally encoded, **212** (*see also* microcode, horizontal)
minus infinity (*see* infinity)
MIPS (*see* millions of instructions per second)
MIPS (*see also* Stanford MIPS)
MIPS Computer Systems, Inc., 41, 68, 93, 189, 339
  MIPS II architecture, E-14
  MIPS R2000, 104, 167, 179, 189-190, 289, 395
  MIPS R3000, 84, 167, 179, 189-190, 289, 492, E-2
    instruction set, E-5–E-6
      common extensions to DLX instructions, E-10–E-11
      unique, E-12–E-14
  MIPS R3010, A-31, A-53 (fig.), A-56, E-5–E-6 (*see also* MIPS Computer Systems, Inc., MIPS R3000)
mirroring, **521**
MISD computer (*see* multiple instruction streams–single data stream computer)
mispredicted branch (*see* misprediction penalty)
misprediction penalty, 277, 310, 311 (fig.), 312-313, 328 (*see also* branch-prediction schemes)
miss, **404** (*see also* cache, miss; virtual memory, page fault; memory hierarchy, miss)
misses per instruction, **417** (*see also* cache, miss)
miss penalty, **405** (*see also* memory hierarchy, miss; memory hierarchy, block; cache, miss; virtual memory, miss penalty)
miss rate, **404** (*see also* cache, miss)
MIT (Massachusetts Institute of Technology), 25
mixed cache, **423** (*see also* cache)
model for vector performance (*see* vector processor, performance, model for)
modify bit, **443** (*see also* virtual memory, page table; virtual memory, dirty bit)
Morse, S., 188
MOS, 59
Motorola Corporation
  C88000 1.8.4m14 C compiler, 83
  6809, 91
  68000, 93, 188
    architecture of, E-23
    interrupts on, 215 (fig.)
  88000, 167, 190, 495
    architecture of, E-2
    instruction set, E-5–E-6
      common extensions to DLX instructions, E-10–E-11
      unique, E-17–E-19
  88100, 84, 492
  88200, 492
Moussouris, J., 189
`MOVC3`, 219, 245-246
MTTF (*see* mean time to failure)
MTTR (*see* mean time to repair)
Muchnik, S., 190
Mudge, J. C., 81
Multibus II, 532, 532 (fig.) (*see also* bus)
multicomputer, 589, 593-594
multicycle operations, 283
  DLX and, 284-289
Multiflow machine, 340
multilevel cache (*see* cache, two-level caches)
multilevel inclusion property, **465** (*see also* cache, two-level caches)
multiple functional units (*see* functional units, multiple)
multiple instruction issue, 318-320, 321-325, 340
  dynamic scheduling and, 321-322

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 780

multiple instruction streams–multiple data streams (MIMD) computer, 574-576, 578, 587, 591, 592, 593
loosely-coupled MIMD (*see* multicomputer)
tightly-coupled MIMD (*see* multiprocessor)
multiple instruction streams–single data stream (MISD) computer, 573, 580
multiple operations per instruction, 323-325, 340
multiple-precision addition, A-11
multiple private address spaces, **578**
multiplication (*see* arithmetic, multiplication, floating-point; arithmetic, integer)
multiply-step instruction, **A-11** (*see also* arithmetic)
multiprocessor, 72-73, **574**-575, 581, 589, 593-594
caches on (*see* cache, coherency)
Cm* multiprocessor, 589
C.mmp multiprocessor, 589
Encore Multimax multiprocessor, 589
IBM RP3 multiprocessor, 589
measuring performance of, 585-586
miss rate, 468 (*see also* cache, coherency)
Symmetry multiprocessor, 582-585, 589
writes and, 468 (*see also* cache, coherency)
MVL (*see* maximum vector length; vector processor, vector length)

N

n-way set associative (*see* cache)
NaN (*see* not a number)
nano-, **244**
nanocode, 244-245
nanoinstruction, 244-245
Namjoo, M., 190
National 32032 microprocessor, 583
negative infinity (*see* infinity)
networks, 15 (fig.), 526-528
ARPANET, 527, 528 (fig.), 561
Ethernet, **526**, 528, 560
hierarchy of, 528 (fig.)
local area network (LAN), **526**-527, 528 (fig.)
range of characteristics, 526 (fig.)
RS232, 526, 528 (fig.)
Newton's iteration, **A-23**–A-24, A-25, A-26
New York University (NYU) Ultracomputer, 589
nibble mode, **431** (*see also* memory, DRAM)
ninety/ten rule (*see* locality, principle of)
nonrestoring division, **A-5** (*see also* arithmetic, nonrestoring)
nonunit strides, **367** (*see also* vector processor, stride)
Noonan, R., 127
no operation (NOP), 491
Spice miss rates with and without, 491 (fig.)
NOP (*see* no operation)
not a number (NaN), **A-12**–A-14, A-30 (*see also* arithmetic)
not taken, **270** (*see also* branch, not taken)
Nova (*see* Data General)
no write allocate, **413** (*see also* cache, write miss)
NuBus, 15, **561** (*see also* bus)
n-way set associative (*see* cache)

O

object-code compatibility, 4
offset address (*see* memory hierarchy, block)
O'Laughlin, J., 127
one level store (*see* virtual memory)
one's complement, **A-7** (*see also* arithmetic, signed)
operand specifier (*see* addressing mode)
operand
naming of, 90-92
type and size, 109-111
operand storage, 91–92
in memory, 92-94
operating system, 127-129
operations, **103**
operators (*see* operations)
operating system, 13, 15 (fig.), 19 (fig.)
operand specifier, 330-332
Opfetch (*see* Digital Equipment Corporation, Opfetch)
optical disk (*see* disk, optical)
optical compact disk, **519** (*see also* disk, optical)
optical write-once disk (*see* disk, optical)
optimization
global, 112, **114-115**, 131
high-level, 112, **114**
local, **114-115**
machine-dependent, **114-115**
organization, **13** (*see also* memory, organizations of; memory hierarchy)
organizations for improving main memory performance (*see* memory, organizations of)
effect on design time, 16
out of order
completion (*see* out-of-order completion)
execution (*see* out-of-order completion)
interrupts, 280-282
out-of-order completion, **287**-289, **291**-293, 304
out-of-order execution, **291**-292, 299-300, 339 (*see also* scoreboard; Tomasulo algorithm)
out-of-order fetch, **458** (*see also* cache, miss)
output dependence, **374** (*see also* vector processor, output dependence)
overflow, **A-7** (*see also* arithmetic, exception, overflow)
overflow, window (*see* register windows)
overlap (*see* pipelining)
overlapped integer and floating-point instructions, 285
overlapped loop iterations, 308
overlapping
I/O (*see* input/output, system performance and)
triplets, **A-44**, A-59 (*see also* arithmetic, integer, speeding up multiplication)
vector processing and, 360, 389-390
overlays, **433** (*see also* virtual memory, overlays)

P

P0, 441 (*see also* virtual memory, VAX-11/780 and)
P1, 441 (*see also* virtual memory, VAX-11/780 and)
package (*see also* cost)
cost of, 55, 60-62, 84
design and, 54
packaged system price (*see* cost)
packed (*see also* binary-coded decimal, packed)
packet switched approach, **527** (*see also* networks)
packets, **526** (*see also* networks)
packing operation, 110
Padegs, A., 186
page, 19, **433**, **434** (*see also* virtual memory, page; address, memory)
paged segments, **434** (*see also* virtual memory, page; virtual memory, segment)
page fault, 19, **433** (*see also* virtual memory, page fault; interrupts, page faults and)

page fault (*continued*)
pipelining and, 279-282
page mode for DRAMs, **431** (*see also* memory, DRAM)
page size (*see* virtual memory, paged, page size)
page table, **435** (*see also* virtual memory, page table)
page-table entry (PTE), **443** (*see also* virtual memory, page table)
on the Intel 80286/80386, 446
on the VAX-11/780, 443, 475
parallelism (*see also* instruction, parallelism)
in pipelining, 252, 314
instruction-level parallelism and pipelining, 314-328, 340-341 (*see also* instruction, parallelism)
parameters, typical ranges of
cache, 408 (fig.)
translation-lookaside buffers, 438 (fig.)
two-level cache, 463 (fig.)
VAX-11/780 TLB, 443 (fig.)
virtual memory, 433 (fig.)
partner, **512** (*see also* input/output, devices)
pass, 111, **112**, 114
Patterson, David, 130, 189, 190
PC (*see* program counter; branch)
PC (personal computer) (*see* Intel Corporation, 80x86; Intel Corporation, 8088; International Business Machines Corp., IBM PC)
PDP (*see* Digital Equipment Corporation)
peak performance (*see* vector processor, performance, peak performance)
Pegasus computer, 127
penalty for misprediction (*see* misprediction penalty)
Pendleton, J., 190
Perfect Club benchmark (*see* benchmark)
performance, **5-8**, **35**, 36-40, 71, 502 (*see also* input/output, system performance and; bandwidth; cost/performance; latency; response time; throughput)
Amdahl's Law and, 8-11
cache (*see* cache, performance)
cost and, 22, 26, 34
CPU and, 11, 16 (*see also* central processing unit, performance)
design requirements and, 13-17
"faster than," 6-7, 28
graphics display (*see* graphics displays, performance demands of)
growth of, 3, 4 (fig.), 5, 6, 21, 28
improving, 502-506 (*see also* input/output, system performance and)
input/output (*see* input/output, performance)
locality of reference and, 18, 20 (*see also* locality)
memory hierarchy (*see* memory hierarchy, performance; memory, performance; cache, performance; virtual memory, performance)
peak, 71, 74-75
pipelining performance improvement (*see* pipelining, DLX and, performance of)
RISC performance advantage (*see* reduced instruction set computer, performance advantage of)
"slower than," 7
system, **35** (*see also* input/output, system performance and)
vector processor (*see* vector processor, performance)
virtual memory (*see* virtual memory, performance)
peripheral, 499 (*see also* input/output, devices; disk, magnetic; graphics displays; networks; bus)
personal computer (PC), 560 (*see also* Intel Corporation, 80x86; Intel Corporation, 8088; International Business Machines Corp., IBM PC)
personal computer (*continued*)
versus workstation, 500
Pfister, G. F., 589
phase, **112** (*see also* pass)
phase-ordering problem, **111**-112
Phister, M., 81
physical addresses, **433** (*see also* virtual memory, address translation)
physical-instruction-buffer address (PIBA), **475** (*see also* memory hierarchy, VAX-11/780 and)
PIBA (*see* physical-instruction-buffer address)
PID, **460** (*see also* process-identifier tag)
pin grid array (PGA), 60, 84 (*see also* package)
pipeline, 8, 22, 25, **251** (*see also* pipelining)
pipelined bus, **530** (*see also* bus)
pipelined machines, 352
pipelined mode, **E-21**
in Intel 860, E-20–E-22
pipeline hazard (*see* hazard, data)
pipeline hazard detection (*see* hazard, detection)
pipeline interlock, **265**-267, 339 (*see also* load interlock)
DLX and, 267-268
pipeline reservation tables, 256, 339
pipeline scheduling, **114,** 119, **267**-268, 315-317, 339 (*see also* optimization; dynamic scheduling)
pipeline speedup, 258-259, 277
pipeline stall, 257-259, **265**-266, 278, 285, 290 (fig.)
branch delay and, 273-278
control hazard and, 269-271, 270 (fig.)
vector machines and, 352, 357-358
pipeline throughput (*see* pipelining, speedup)
pipelining, **251**-349
balance among stages, 252
balance in issue, 320
clock cycles and, 351
depth of a pipeline, 253, 258, 336, 339
difficulties in implementation, 278-284
dealing with interrupts, 279-282
instruction set complications, 282-284, 334-335
DLX and, 252-257, 270, 252-257, 278-282, 300, 301 (fig.)
floating-point, 260, 284-290, 299-300
integer, 252-278
performance of, 278, 290
superscalar DLX (*see* superscalar)
dynamic hardware prediction, 307-314 (*see also* branch-prediction schemes)
multiple instruction issue and, 321-322
dynamic scheduling, **291**, 290-307, 340
multiple instruction issue and, 321-322
scoreboard approach (*see* scoreboard)
Tomasulo algorithm (*see* Tomasulo algorithm)
hazards of (*see* hazard)
instruction-level parallelism, 314-328, 340-341
dynamic scheduling and, 321-322
loop unrolling and, 315-318
software pipelining and, 325-328
superscalar machines and, 318-320
trace scheduling and, 325-328
VLIW approach and, 322-325
Intel 860 and, E-21
making the pipeline work, 255-257
performance of, 278, 290
software for, **325**-328, 340
software-pipelined loop, 325, 327 (fig.)
speedup, 251-253, 289
superscalar DLX (*see* superscalar)
timing of instructions, 254, 260 (*see also* pipeline speedup)

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 782

pipelining (*continued*)
VAX 8600 and, 328-334
dealing with interrupts, 332-334
handling data dependences, 331
handling control dependences, 331-332
operand decode and fetch, 330-331
writes and (*see* write result in a pipeline)
pipe segment, **251**
pipe stage, **251-253**, 255-256, 285
Pitkowsky, S. H., 78
pixel instructions of Intel 860 (*see* graphics instructions)
pixels, **521** (*see also* graphics displays)
PLA (*see* programmed logic array)
placement, block (*see* block placement)
plastic quad flat pack (PQFP), 60 (*see also* package)
plus infinity (*see* infinity)
Pohlman, W., 188
polling, **534** (*see also* input/output, interfacing to the CPU)
pollution point, **406** (*see also* memory hierarchy, block; cache, blocks and)
position-independence, **105**
positive infinity (*see* infinity)
precise interrupts (*see* interrupts, precise)
precision (*see* arithmetic, precision)
Precision (*see* Hewlett-Packard, Precision)
predicting system performance (*see* input/output)
prediction of branching (*see* branch-prediction schemes)
prediction accuracy (*see* branch-prediction schemes, prediction accuracy)
predict-not-taken (*see* branch-prediction schemes, predict-not-taken)
predict-taken (*see* branch-prediction schemes, predict-taken)
present bit, **446** (*see also* virtual memory, page table; virtual memory, Intel 80286/80386 and)
price (*see* cost)
primitive, 121
Princeton University Institute for Advanced Study (IAS), 24
principle of locality, **403** (*see also* locality; memory hierarchy, principle of locality and)
procedure call/return, 73, 81, 103-105, 108-109, 114, 116, 137
fallacies and pitfalls, 124-125
procedure inlining (*see* procedure integration)
procedure integration, **112**, 114-115
process, **438** (*see also* virtual memory, processes and)
process-identifier tag (PID), **460** (*see also* cache)
processing
parallel, 22, 26 (*see also* parallelism)
sequential, 26
processor, **199**, 211 (*see also* central processing unit)
computation and, 201
control and, 201, 204, 214
datapath and, 201
special-purpose, 580
processor-memory-switch level (*see* organization)
process segments, 441 (*see also* virtual memory, VAX-11/780 and)
process switch, **438** (*see also* virtual memory, processes and)
producer-server model (*see* input/output, performance)
program
behavior (*see* instruction-prefetch buffer; register windows)
benchmarks (*see* benchmark)
of channel (*see* channel program)
program counter (PC), **105**
PC (program counter)-relative addressing, **97**-98, 104-106
PC (program-counter)-relative branches (*see* branch)
VAX 8600 and, 332
programmable read-only memory (PROM), 63
programmed logic array (PLA), **205**-206, 230, 232
PROM (*see* programmable read-only memory)
propagate, **A-32**–A-33 (*see also* carry-propagate adder; arithmetic)
protection, **432** (*see also* virtual memory, protection schemes of; virtual memory, Intel 80286/80386 and)
protocols
coherency (*see* cache, coherency)
networks and, **527** (*see also* networks)
multiprocessors and (*see* cache, coherency)
Przybylski, S., 189
PTE (*see* page-table entry)
Puzzle (*see* benchmarks, toy)

Q

Q1 (*see* block placement)
Q2 (*see* block identification)
Q3 (*see* block replacement)
Q4 (*see* write strategy)
questions for classifying memory hierarchies (*see* block identification; block placement; block replacement; write strategy)
Quicksort (*see* benchmarks, toy)
queueing delay, **516** (*see also* disk, magnetic)
queues, 321, 340

R

Radin, G., 189
RAID (redundant arrays of inexpensive disks) (*see* disk array)
random, **411** (*see also* block replacement, random)
ranges of parameters (*see* parameters, typical ranges of)
RAR (*see* read after read)
RAS (*see* row-access strobe)
raster, **521** (*see also* graphics displays)
raster cathode ray tube (CRT) display, **521** (*see also* graphics displays)
raster refresh buffer, **521** (*see also* graphics displays)
Ravenal, B., 188
RAW (*see* read after write)
RAW hazard (*see* hazard, RAW)
read after read (RAR), **265**
read after write (RAW), **264** (*see also* hazard, RAW)
read miss rate, 416 (*see also* cache, reads and; cache, miss)
read-only memory (ROM), **205**, 208, 239, 241-242
future of microprogramming and, 241
read-only protection, 440 (*see also* virtual memory, protection schemes of)
read-only storage (*see* read-only memory)
read-write head, **516** (*see also* disk, magnetic)
recurrence, **373** (*see also* vector processor, data dependences)
recursive doubling, **382** (*see also* vector processor, vector reduction)
reduced instruction set computer (RISC), 130, 131, 132, **188**-190, 337, 339-340 (*see also* International Business Machines Corp., IBM 801; MIPS Computer Corporation)
architecture, survey of, E-1–E-24
addressing mode, E-2
arithmetic and logical instructions, E-5
conditional branch of RISC, E-8
constant extension, E-4
control-flow instructions, E-6
data transfer, E-5
floating-point instructions, E-6
instruction format, E-3
integer multiply and divide, E-8–E-9

reduced instruction set, architecture (*continued*)
Berkeley, 189
performance advantage of, 189
reducing branch penalties (*see* branch-prediction schemes)
Redmond, K. C., 25
reduction **382** (*see also* vector processor, vector reduction and)
redundant, **A-42** (*see also* arithmetic, integer, speeding up division, shifting over zeros)
redundant arrays of inexpensive disks (*see* disk array)
reference bit, **436** (*see also* block replacement, least-recently used)
refresh, 426 (*see also* memory, DRAM)
refresh rate, **521** (*see also* graphics displays)
register, 19, 20, 22, 90-94
allocation, 108-109, 112-144, 115-119, 130
caches versus, speed of, 483
DEC VAX, 143-144
DLX, 161-162
file, 324
field, 102-103
general-purpose register (GPR) architecture, **91**-94, 127-128
comparison of, 93-94
hazard (*see* hazard, register)
IBM 360, 148-150
Intel 8086, 153-155, 154 (fig.)
machine (*see* register, general-purpose register architecture)
register-memory architecture (*see* register-memory architecture)
renaming, 307, 339, 340
antidependences and output dependences and, 374-375
register deferred (indirect) addressing, **98** (*see also* addressing mode)
register-indexed (RX) (*see* International Business Machines Corp., IBM 360, instruction set)
register-register (RR) (*see* International Business Machines Corp., IBM 360, instruction set)
register-register architecture (*see* register-register architecture)
register-storage (RS) (*see* International Business Machines Corp., IBM 360, instruction set)
result status, 295, 296 (fig.), 297 (fig.), 302-303
set, 91, 118-119
shadow (*see* shadow registers)
tags, 303-306
vector (*see* vector processor, registers)
vector-length (*see* vector processor, vector length)
vector-mask (*see* vector processor, vector-mask registers)
windows (*see* register windows)
register-memory architecture, **93**-94, 128
register-memory instruction, 39-40
register-register architecture, **93**-94 (*see also* load/store architecture)
register-storage architecture (*see* register-memory architecture; International Business Machines Corp., IBM 360, instruction set)
register windows, **450**-454, 484 (fig.), 487, E-15
benefits of, on DLX, 453 (fig.)
load and store benefits, 453 (fig.)
number of versus overflow rate, 451 (fig.)
pros and cons of, 453-454
summary of, 484 (fig.)
reliability, **520** (*see also* input/output, reliability)
relocation, **433** (*see also* virtual memory, relocation and)
REM, A-26–A-28, A-53 (*see also* arithmetic, remainder)
remainder (*see* arithmetic, remainder)
Remington-Rand Corporation, 25
replacement, block (*see* block replacement)
requested protection level, **448** (*see also* virtual memory, protection schemes of; virtual memory, Intel 80286/80386 and)
requirements, functional, 13-14, 15 (fig.)
reservation stations, **300**-308, 321
resources
allocation of, 8, 11
pipelines and, 255-257, 287
VLIW approach and, 323
response time, 6, 22, 506 (*see also* execution time; performance; input/output, performance, response time)
definition of, 5
restartable, **218**-220, 240, 279-282
restoring division (*see* arithmetic, division, integer, restoring)
result buffer, 263
result store, 330 (fig.), 331
return (*see* procedure call/return)
rings, **440** (*see also* virtual memory, protection schemes of)
ripple-carry adder, **A-2** (*see also* arithmetic, integer, ripple-carry addition)
Riordan, T., 189
RISC (*see* reduced instruction set computer)
RISC-I and RISC-II, 189, 190
Riseman, E. M., 129
ROM (*see* read only memory)
rotational positional sensing (RPS), 551 (*see also* disk, magnetic, IBM 3990 storage subsystem and)
rotation delay, **516** (*see also* disk, magnetic)
rotation latency, **516** (*see also* disk, magnetic)
rounding (*see* arithmetic, rounding and)
rounding modes, **A-13** (*see also* arithmetic, rounding and)
row-access strobe (RAS), **425**
Rowan, C., 189
row-major order, **366**, 367 (fig.)
RPS (*see* rotational positional sensing)
RPS miss, **552** (*see also* disk, magnetic, IBM 3990 storage subsystem and)
RR (*see* register-register)
RS (*see* register-storage)
RS232, 526 (*see also* networks)
rules of thumb, front endsheet (*see also* Case/Amdahl rule of thumb)
2:1 cache rule, front endsheet
90/10 locality rule, front endsheet
90/50 branch-taken rule, front endsheet
address-consumption rate, front endsheet
Amdahl/Case rule, front endsheet
disk-growth rate, front endsheet
DRAM-growth rule, front endsheet
RX (*see* register-indexed)

S

S810/20 (*see* Hitachi S810/20)
safe calls from user to OS gates, 448 (*see also* virtual memory, Intel 80286/80386 and)
Saji, K., 81
Samples, D., 189
SAXPY (*see* vector processor, Linpack benchmark)
scalability, **574**, 585
scalar expansion, **382** (*see also* vector processor, vector reduction and)
scalar variable, **116**
global, **116**, 119
scaled (index) addressing, **98** (*see also* addressing mode)
scatter, **380** (*see also* vector processor, sparse matrices and)

scatter-gather, **380** (*see also* vector processor, sparse matrices and)
scheduling, 268 (*see also* branch, scheduling, branch-delay scheduling; dynamic scheduling; instruction scheduling; pipeline scheduling)
scheduling the branch-delay slot (*see* branch-delay slot)
scheduling effectiveness, 268, 276, 278
schemes for branch-prediction (*see* branch-prediction schemes)
Schwartz, J. T., 130, 589
scoreboard, 291-299, 398-399, 346
  components of, 296 (fig.)
  dynamic scheduling around hazards with a scoreboard, 291-299
  hazard detection, 293 (*see also* hazard, detection)
  instruction issue, 293 (fig.)
  tables, 296-298 (figs.)
scoreboard approach (*see* scoreboard)
scoreboarding, **292** (*see also* scoreboard)
SCRAM (*see* static column DRAM)
SCSI (*see* small computer systems interface)
sectors, **515** (*see also* disk, magnetic)
seek, **516** (*see also* disk, magnetic, seeks and)
seek time, **516** (*see also* disk, magnetic, seeks and)
segment, **433**, **434** (*see also* virtual memory, segment)
segment descriptor, **446** (*see also* virtual memory, page-table)
self-modifying code, 335
semantic clash, **124**
semantic gap, 124, **129**
semaphore (*see* synchronization)
set associative (*see also* cache, set associative)
Sequent Corporation, 583
  Balance 8000, 583
  Balance 2100, 583
  Symmetry multiprocessor, 582-585, 589 (*see also* multiprocessor)
sequential consistency, 474 (*see also* cache, coherency)
sequential processing (*see* processing, sequential)
shadow registers, **246**
shadowing, **521**
shared caches (*see* cache, coherency)
shared memory (*see* virtual memory, shared; virtual memory, Intel 80286/80386 and)
shared-memory processor, **574**-575, 578-579, 589, 591, 592
shifting over zeros, **A-40** (*see also* arithmetic, integer, speeding up division, shifting over zeros)
short-circuiting, **261** (*see also* forwarding)
Shurkin, J., 25
Shustek, L. J., 138, 172-173, 185, 187
SI (*see* storage-immediate)
Sieve Of Erastosthenes (*see* benchmarks, toy)
signal
  delay, 18
  propagation, 18
sign-magnitude, **A-7** (*see also* arithmetic, signed)
signed-digit representation, **A-48** (*see also* arithmetic, signed)
signed-logarithm representation, A-65 (*see also* arithmetic, signed)
signed numbers (*see* arithmetic, signed)
SIMD computer (*see* single instruction stream–multiple data stream computer)
simulate the execution (*see* execution, simulation)
single instruction stream–multiple data stream (SIMD) computer, **572**-574, 578, 589, 592. 593
single level store, 432 (*see also* virtual memory)
Slater, R., 25
Slotnick, D. L., 589
slots (*see* branch-delay slots; load delay)
small computer systems interface (SCSI), 15, 532 (fig.), 560-**561** (*see also* bus)
Smalltalk, support for in SPARC, E-15–E-16
Smith, A., 486, 489
Smith, J. E., 79
Smith, T. M., 25
snoop, **467** (*see also* cache, coherency)
snooping, **467** (*see also* cache, coherency)
Snoopy cache (*see* cache, coherency)
software, 16-17 (*see also* balance, software and hardware)
software pipelining (*see* pipelining, software)
software-pipelined loop (*see* pipelining, software-pipelined loop)
solid state disks (SSDs), **519** (*see also* dynamic random access memory)
source code (*see* code, source)
SPARC, 167, 190
  architecture, 190
  instructions, E-5–E-6
    common extensions to DLX instructions, E-10–E-11
    unique, E-15–E-17
  summary of, E-2
SPARCstation 1 (*see* SPARC)
sparse matrices (*see* vector processor, sparse matrices and)
spatial locality, **403** (*see also* locality; memory hierarchy, principle of locality and; locality)
SPEC (System Performance Evaluation Cooperative) (see benchmark programs)
special-purpose processor (*see* processor, special-purpose)
speed-matching buffer, **540**, **549** (*see also* input/output)
speedup, 9-11, 20, 26, 28, 29
  definition of, 9
  enhanced, 10
  overall, 10
Spice program, 12, 44, 45, 67, 69, 70, 72, 79, 83, 86
spin lock, 473 (*see also* cache, coherency, synchronization)
spin waiting, 472 (*see also* cache, coherency, synchronization)
split transactions, **530** (*see also* bus)
square root (*see* arithmetic, square root)
SRAM (*see* static random access memory)
SRT division, **A-40**, A-41, A-42, A-51, A-53, A-56, A-59 (*see also* arithmetic, integer, speeding up division, shifting over zeros)
SS (*see* storage-storage)
SSD (*see* solid state disks)
stack, 98, 114, 116-118, 124-125, 127, 131, 134, 136 (*see also* stack architecture)
  alignment of, 124
  height reduction, **114** (*see also* optimization)
stack architecture, **90**-92, 127
stale data, **466**, 535-537 (*see also* cache; virtual memory; input/output)
stall, 213-**214** (*see also* memory stall cycles)
stall, pipeline (*see* pipeline stall)
standards
  bus (*see* bus, standards)
Stanford MIPS, 189 (*see also* MIPS Computer Systems Inc.)
start-up time **358** (*see also* vector processor, start-up time)
state-assignment problem, **206**
states, **201**, 204-206 (*see also* finite state diagram)
  clock cycles and, 224-225, 228
  DLX and, 205, 221-224 (figs.), 225
  interrupts and, 216, 218-219
  PLA and, 206
static column DRAM, **431** (*see also* dynamic random access memory, static column; memory, DRAM)
static measurements (*see* instruction set, measurements, static)

static random access memory (SRAM), 426, 431 (*see also* dynamic random access memory; memory)
capacity of, 426
cost versus access time of, 518 (fig.)
cycle time of, 426
static scheduling, 267, 274-275, 290-**291**, 315-317 (*see also* dynamic scheduling)
versus dynamic scheduling, 321, 340, 349
versus Tomasulo algorithm, 307
Stern, N., 24
sticky, **A-30**
sticky bit, **A-17**–A-18, A-23, A-59
storage (*see* memory; disk; disk, magnetic; input/output)
storage director, **549** (*see also* disk, magnetic, IBM 3990 storage subsystem and)
storage hierarchy (*see* memory hierarchy)
storage-immediate (SI) (*see* International Business Machines Corp., IBM 360, instruction set)
storage-storage (SS) (*see* International Business Machines Corp., IBM 360, instruction set)
storage-storage architecture (*see* memory-memory architecture)
storage subsystem, IBM (*see* disk, magnetic, IBM 3990 storage subsystem and)
stored-program computer, 23-25
store in, **413** (*see also* cache, write back)
store through, **413** (*see also* cache, write through)
Strapper, C. H., 81
strategy for writes (*see* write strategy)
Strecker, W. W., 130
Strecker (*see* Bell, C. G. and W. D. Strecker)
strength reduction, **114** (*see also* optimization)
Stretch (*see* International Business Machines Corp. IBM 7030)
stride, **367** (*see also* vector processor, stride)
string operations, 15
string operators, 103
strip mining, 364-365
subblock placement, **456** (*see also* cache, subblocks)
subblocks, **456** (*see also* cache, subblocks)
subexpression (*see* common subexpression elimination)
summary of memory hierarchy examples, 484 (fig.)
Sumner, F. H., 26
Sun Microsystems (*See also* SPARC)
1.2 FORTRAN compiler, 83
C compiler, 83
FORTRAN 77 compiler, 82
supercomputer, 3-4, 500
CRAY-1 (*see* Cray Research machines, CRAY-1)
CRAY-2 (*see* Cray Research machines, CRAY-2)
CRAY X-MP (*see* Cray Research machines, CRAY X-MP)
CRAY Y-MP (*see* Cray Research machines, CRAY Y-MP)
Fujitsu (*see* supercomputer, Japanese)
I/O and (*see* input/output, supercomputers and; disk, magnetic, I/O benchmarks for)
Japanese, 353, 390, 394
NEC SX-2 (*see* supercomputer, Japanese)
supercomputer I/O benchmarks (*see* input/output, supercomputers and; disk, magnetic, I/O benchmarks for)
"super-microprocessor", 500
superpipelined, **337**, 340-341
superscalar DLX (*see* superscalar, DLX)
superscalar
DLX, 318-320, 325
instruction issue, 318-320
instruction level parallelism, 318-320
in Intel 860, E-22–E-23
loop unrolling and, 319-320
machines, 318-320, 340-341, 573, 581
pipeline on, 319 (fig.)
processor, 337-338
structural hazards and, 319
superscalar machines, **318** (*see also* superscalar)
superscalar pipeline (*see* superscalar)
superscalar processor, **337** (*see also* superscalar)
supervisor process, **440** (*see also* virtual memory, processes and)
sustained performance (*see* vector processor, performance, sustained performance)
sustained rate (*see* vector processor, sustained rate)
Sutherland, I., 521, 561, 563
SYMBOL Project, 129, 132
Synapse N+1, 471, 487
synchronization, 471 (*see also* cache, coherency)
synchronous bus, **530** (*see also* bus)
synonyms, **460** (*see also* aliases)
synthetic benchmark (*see* benchmark, synthetic)
system CPU time (*see* central processing unit, CPU time, system)
system mode, **440** (*see also* virtual memory, protection schemes of)
system operators, **103**
system performance, **35** (*see also* performance; input/output, system performance and)
system response time, **508** (*see also* input/output, performance, response time)
system segments, 441 (*see also* virtual memory, VAX-11/780 and)
systolic architecture, **580**, 591
systolic array, **580**, 590 (*see also* array)

T

tag field, **410** (*see also* cache)
Tagged architecture (*see* SPARC)
tagging of data, 307, 339
taken branch, **270** (*see also* branch, taken)
Taylor, G., 189
technology (*see* design, computer; disk; implementation; logic)
temporal locality, **403** (*see also* locality; memory hierarchy, principle of locality and)
terminal network, **526** (*see also* networks)
test and set, **473** (*see* cache, coherency, synchronization)
TeX, 45, 67, 69, 70, 79, 80, 86
Texas Instruments
8847, A-26, A-53 (fig.), A-57
Thacker, C., 487, 490, 560, 563
Thadhani, A., 560, 563
think time, **508** (*see also* input/output, transactions and)
thrash, **420** (*see also* memory hierarchy; cache)
"three Cs" (*see* cache, miss)
three-operand format, 93-94
throughput, 5-6, 22
latency and, 8
I/O and, 500-501 (*see also* input/output, performance, throughput)
of pipeline (*see also* pipelining, speedup)
TI (*see* Texas Instruments)
ticks (*see* clock cycles)
tightly-coupled MIMD (*see* multiprocessor)
$Time_{best}$ of CPU and I/O overlapped, 503-505 (*see also* input/output, system performance and)
time distributions (*see* instruction set, measurements)
$Time_{scaled}$ of CPU and I/O overlapped, 503-504 (*see also* input/output, system performance and)
timesharing, 575

CAVIUM-1035
Cavium, Inc. v. Alacritech, Inc.
Page 786

Time$_{worst}$ of CPU and I/O overlapped, 503-505 (*see also* input/output, system performance and)
timing of instructions (*see* pipelining, timing of instructions)
TLB (*see* virtual memory, translation-lookaside buffer)
TLB instruction, **E-12**
Tomasulo algorithm, 299-307, 339 (*see also* dynamic scheduling)
  DLX and, 301 (fig.)-307
  hazard detection and, 300, 304 (*see also* hazard, detection)
  versus static scheduling, 307
toy benchmark (*see* benchmark, toy)
TP (*see* transaction processing)
TP-1, 510, 511 (fig.) (*see also* benchmark; disk, magnetic, I/O benchmarks for)
TPI (*see* clock cycles per instruction)
trace, **326**
trace compaction, **326**
trace scheduling, **323**, **326**, 325-328, 340
  VLIW and, 326
trace selection, **326**
tracks, **515** (*see also* disk, magnetic)
tradeoffs (*see* balance)
traffic ratio, **491**, **567**
transaction, **508** (*see also* input/output, transactions and)
transaction processing (TP), 14, 15, **511** (*see also* input/output, transaction; disk, magnetic, I/O benchmarks for)
transaction processing I/O benchmarks (*see also* disk, magnetic, I/O benchmarks for)
transaction time, **508** (*see also* input/output, transaction and)
transfer, **104** (*see* also branch)
transfer time, **516**, **405** (*see also* memory hierarchy, miss; disk, magnetic)
translation-lookaside buffer (TLB), **437** (*see also* virtual memory, translation-lookaside buffer)
Transputer-based multicomputer, 589
traps, **216** (*see also* interrupts)
trivia, front endsheet
Trojan horses, **445** (*see also* virtual memory, protection schemes of)
true data dependence, **374** (*see also* vector processor, data dependences)
Tuck, R., 190
Tucker, S., 242
two-bit prediction, **309**-310 (*see also* branch-prediction schemes)
two-level cache (*see* cache, two-level caches)
two-operand format, 93
two's complement, A-5, **A-7**–A-9, A-18–A-19
two-to-one cache rule, front end sheet
typical parameters (*see* parameters, typical ranges of)
typical program, 183

U

Ultrix C compiler, 68
unbiased exponents, A-14 (*see also* arithmetic, exponents and)
unconditional branches (*see* jump)
underflow (*see* arithmetic, exceptions, underflow)
underflow trap (*see* arithmetic, exceptions, underflow)
underflow, window (*see* window registers)
underpipelined, **337**, 344
unfair benchmarks, **490** (*see also* benchmark)
Ungar, D., 189
Unibus (*see* bus, Unibus)
unified, **423** (*see also* cache)
uniprocessor, 72-73
UNIVAC I, 241
UNIVAC I, 25, 26
University of Illinois Cedar project, 589
University of Manchester, 24, 485
University of Pennsylvania Moore School, 23-24
UNIX, 4, 15 (*see also* operating system)
unlock, **472** (fig.) (*see also* cache, coherency, synchronization)
unpacked, **A-14** (*see also* binary-coded decimal, unpacked)
unpacking operation, 110
unrolled loop (*see* loop unrolling)
untaken branch, **270** (*see also* branch, not taken)
upper level, **404** (*see also* memory hierarchy; cache; memory; virtual memory)
usage (*see* instruction set, measurements)
use bit, **436** (*see also* block replacement, least-recently used)
useful slots, 276 (*see also* branch-delay slots)
user code (*see* code, user)
user CPU time (*see* central processing unit, CPU time, user)

V

valid bit, **410**, **443** (*see also* cache, blocks and; virtual memory, page table)
VAX (*see* Digital Equipment Corporation, VAX)
VAXstation (*see* Digital Equipment Corporation, VAX)
VAX units of performance (VUP), 78
vector, **352** (*see also* vector processor)
  mode, 28
  operations, in Intel 860, E-20
  processor, 25
  rate, 28
vector architecture (*see* vector processor, architecture of)
vector functional units (*see* vector processor, functional units)
vectorization, percentage of, 28
vector length (*see* vector processor, vector length)
vector-length register (VLR), **364** (*see also* vector processor, vector length)
vector-mask control, **379** (*see also* vector processor, vector-mask control)
vector-mask register, **379** (*see also* vector processor, vector-mask register)
vector processor, 351-401
  advantage, 352
  antidependences, **374**-375
  architecture, 353-358
  chaining and, 377-**378**
  compilers and, 371-377 (fig.) (*see also* hazard)
  completion rate, **358**
  component, 353-354
  conditionally executed statements and, 379-382
  data dependences, 360, 371-377, 395 (*see also* antidependences; output dependences)
    Banerjee test, 374
    GCD test, 373
    loop-carried dependences, **372**-373
    RAW hazard, 374 (*see also* vector processor, true data dependence)
    recurrence, **373**
    sparse matrices and, 380-381, 382
    true data dependence, **374**
    WAR hazard, 374 (*see also* antidependence)
    WAW hazard, 374 (*see also* output dependence)
  DAXPY (*see* vector processor, Linpack benchmark)
  dependences (*see* vector processor, antidependences; vector processor, data dependences; vector processor, output dependences)
  DLXV, 353-363, 383-390
    initiation rate, 358

vector processor, DLXV (*continued*)
start-up time of, 358, 361 (fig.)
stride and, 368
vector instructions, 356 (fig.)
vector length and, 364-365
effectiveness (*see* vector processor, performance)
fallacies and pitfalls of, 390-392
functional units, 354
Flynn bottleneck and, **351**
history of, 393-395
improving performance, 377-382, 388-390
by chaining, 377-378
with conditionally executed statements and sparse matrices, 379-382
by vector reduction, 382
with multiple memory pipelines, 388-390
initiation rate, **358**-363
chaining and, 378
Linpack benchmark, 357
DAXPY loop, 357
SAXPY loop, 357, 360, 384, 388
in FORTRAN, 364
memory banks and, **361**-363
memory bank conflicts, 368
mod bank number, 362
output dependences, **374**
overlap, 360, 389-390
peak (*see* vector processor, performance, peak performance)
performance, 375-377
improving (*see* vector processor, improving performance)
analyzing, 369-371, 383-390
evaluating (*see* virtual processor, performance, analyzing)
model of, 369-371
length-related measures, 384
memory bandwidth and, 392
peak performance, 385-386, 390-391
SAXPY performance, 388-390
scalar performance comparison, 391-392
sustained performance, 386-388
reduction (*see* vector processor, vector reduction)
registers, 353, 354
renaming, 374-375
vector, 354 (fig.)
vector-length register (*see* vector processor, vector length)
vector-mask register (*see* vector processor, vector-mask register)
SAXPY (*see* vector processor, Linpack benchmark)
sparse matrices and, 380-382
gather, **380**, 393
index vector, **380**-381, 382
scatter, **380**, 393
scatter-gather, **380**, 381
start-up time, **358**-361, 390
early vector machines and, 390
start-up penalties on the DLXV, 361 (fig.)
stride, **367**, 366-369
nonunit strides, **367**, 393
sustained rate, **360**, 378, 385, 386-388
scalar machines and, 392
Japanese supercomputers and, 390
vector reduction and, **382**
vector length 364-366, 384
maximum vector length (MVL), **364**, 379
vector-length registers (VLR), **364**
vector machines, 22, **352**-353, 355 (fig.), 390-395, 581
memory-memory vector machine, **353**, 390-391, 393
start-up times and, 390
vector processor, vector machines (*continued*)
vector register machine, **353**, 364
vector-mask control, **379**
vector-mask register, **379**-380
vector reduction and, **382**-383
recursive doubling, **382**-383
scalar expansion, **382**
vector stride (*see* vector processor, stride)
vector reduction (*see* vector processor, vector reduction and)
vector register machine, **353** (*see also* vector processor, vector machines)
vector registers (*see* vector processor, registers)
vector stride (*see* vector processors, stride)
vertical microcode (*see* microcode, vertical)
vertical microinstruction (*see* microcode, vertical)
very long instruction word (VLIW), 318, **323**, 322-325, 337-338, 573, 580
instructions, 323
trace scheduling and, 326
VIBA (*see* virtual-instruction-address buffer)
video DRAM, **524** (*see also* graphics displays)
video look-up table, **523** (*see also* graphics displays)
virtual addresses, **433** (*see also* virtual memory, address translation)
virtual cache, **460**
anti-aliasing, **460**
virtual DMA, **537** (*see also* input/output, DMA and)
virtual-instruction-buffer address (VIBA), **475** (*see also* memory hierarchy, VAX-11/780 and)
virtual memory, 14, 19, 26, 103, 127, 129, **432**-449, 484 (fig.) (*see also* cache; memory; memory hierarchy; block identification; block placement; block replacement; write strategy)
address translation, **433**, 435, 436 (fig.), 440, 442-443, 460 (*see also* virtual memory, translation-lookaside buffer)
techniques for fast address translation, 437-438
on the VAX-11/780, 442-443
block (*see* virtual memory, page; virtual memory, segment)
block information (*see* block information, virtual memory)
block placement (*see* block placement, virtual memory)
block replacement (*see* block replacement, virtual memory)
caches and, 434, 438
differences between caches and, 434
dirty bits and, 436, 438
DMA and, 537
Intel 80286/80386 and, 445-449
attributes field, **446**
bounds checking on, 446
memory mapping on, 446
page-table entry of, 446
protection on, 446, 448-449
safe calls from user to OS gates, 448
segment descriptor of, 447 (fig.)
sharing on, 446-447
miss penalty, 434
overlays, **433**
paged, **433**, **434**
internal fragmentation and, **437**
page size, 437
versus segmentation, 434, 435 (fig.), 441
page fault, **433**, 434, 436 (*see also* cache, miss, "three Cs")
page table, **435**, 437
conserving memory with, 442-443
page-table entry (PTE) on the VAX-11/780, 443, 475
page-table entry/segment descriptor of the Intel 80286/80386, 446

virtual memory (*continued*)
parameters, typical, 433 (fig.) (*see also* parameters, typical ranges of)
processes and, 438-439 (*see also* virtual memory, protection schemes of)
address space, 432
user, kernel and supervisor processes, **440**
protection schemes of, **432**-433, 439-441, 443, 446-449 (*see also* virtual memory, Intel 80286/80386 and, protection on)
base register, **439**
bound register, **439**
read-only protection, 439
rings of security levels, **440**
Trojan horses and, **445**, 447
relocation and, **433**, 434
segmented, **433**, **434**
fallacy of, 483
versus paging, 434, 435 (fig.), 441
shared, 433, 445-446
stale data and (*see* stale data)
summary of, 484 (fig.)
translation-lookaside buffer (TLB), **437**-438, 484 (fig.)
parameters typical of, 438 (fig.) (*see also* parameters, typical ranges of)
miss rates for the VAX-11/780 TLB, 479
misses per hundred instructions on the VAX-11/780, 479
on the VAX-11/780, 443 (fig.), 444-445, 475
summary of, 484 (fig.)
TLB instruction-stream miss rate, 478
VAX-11/780 and, 441-445, 448-449 (*see also* cache, VAX-11/780 and; memory hierarchy, VAX-11/780 and)
area P0, 441
area P1, 441
miss rates for the VAX-11/780 TLB, 479
misses per hundred instructions on the VAX-11/780, 479
operation of the VAX-11/780 TLB, 444 (fig.)
page-table entry (PTE) on the VAX-11/780, 443, 475
parameters typical of, 443 (fig.) (*see also* parameters, typical ranges of)
process segments of, 441
system segments of, 441
writes and (*see* write strategy, virtual memory and)
VLIW (*see* very long instruction word)
VLR (*see* vector-length register; vector processors, vector-length)
VME bus, 532, 532 (fig.) (*see also* bus)
VMS
C compiler, 68
fort (FORTRAN compiler), 68
von Neumann, J., 23-24
"von Neumann syndrome", 587

W

wafer, 55-57, 59
chips per, 59, 84
cost of, 59-60, 62
dies per, 59, 61-62
photographs of, 56-57
yield, 59-60, 62, 84, 85
wait states, **224**
Wakerly, J., 188
Wallace, J. J., 79
Wallace tree, **A-46**, A-47, A-59 (*see also* array multiplier; arithmetic)
WAR (*see* write after read)
Ward, S., 561
Waters, F., 190
WAW (*see* write after write)
WCS (*see* writable control store)
weak consistency, 474 (*see also* cache, coherency)
weighted means (*see* mean)
Weitek 3364, A-53 (fig.), A-56–A-57
What happens on a write? (*see* write strategy)
Wheeler, D. J., 24
Where can a block be placed? (*see* block placement)
Whetstone (*see* benchmark programs, synthetic)
Which block should be replaced on a miss? (*see* block replacement)
Whirlwind, 25
Wichmann, B. A., 78
wider main memory (*see* memory, wider)
width of memory (*see* memory, wider)
width of bus (*see* memory, organization of)
Wiecek, 169, 171, 188
Wilkes, M., 24, 25, 425, 485, 486, 490
window overflow, **450** (*see also* register windows)
window underflow, **450** (*see also* register windows)
word, **95**
word reads, unaligned, **E-13**
workstation, 499-500, 560
cost of, 61, 63, 86
DECstation 3100 (*see* Digital Equipment Corporation, DECstation 3100)
file server versus, 500
minicomputer versus, 499
personal computer versus, 500
VAXstation 2000 (*see* Digital Equipment Corporation, VAXstation 2000)
VAXstation 3100 (*see* Digital Equipment Corporation, VAXstation, 3100)
SPARCstation I (*see* SPARC)
workload, 45
WORM (*see* write-once, read-many)
Wortman, D. B., 130, 187
wrapped form, **A-21**–A-22
wrapped fetch, **458** (*see* cache, miss)
writable control store (WCS), **239**-240, 248
write after read (WAR), **264** (*see also* hazard, WAR)
write after write (WAW), **264** (*see also* hazard, WAW)
write allocate, **413** (*see also* cache, write miss)
write around, **413** (*see also* cache, write miss)
write back, **413** (*see also* cache, write back)
virtual memory and, 436
write broadcast, **469** (*see also* cache, coherency)
write buffer (*see* cache, write buffer; cache, writes and)
write invalidate, **469** (*see also* cache, coherency)
write miss rate, 416 (*see also* cache, write miss)
write-once optical disk, **519** (*see also* disk, optical)
write-once, read-many (WORM), 497
write result in a pipeline, 294, 296-298 (figs.), 301, 303 (fig.), 305-306 (fig.), 308 (fig.), 333, 347
write stall, **413** (*see also* cache, writes and; cache, write buffer; write stalls and)
write strategy, **407**, 484
caches and, 412-414, 468 (*see also* cache, writes and)
virtual memory and, 436
write through, **413** (*see also* cache, write through)
Wulf, W., 127, 485, 490

X

X-MP (see Cray Research machines)

Y

yield, 54-55, 80, 81 (*see also* die; integrated circuit; wafer)
  final test, 55, 60-62
  scrap and, 64
Y-MP (see Cray Research machines)

Z

z buffer, **525** (*see also* graphics displays)
Zimmermann, R., 188
Zorn, B., 191
Zuse, 24

QA76.9/.A73/H392/1990
Computer architecture : a quantitative
approach / David A. Patterson, John L.

crdd c. 1 SVO

DLX Standard Instruction Set

Instruction type / opcode	Instruction meaning
Data transfers	**Move data between registers and memory, or between the integer and FP or special registers; only memory address mode is 16-bit displacement + contents of an integer register**
LB, LBU, SB	Load byte, load byte unsigned, store byte
LH, LHU, SH	Load halfword, load halfword unsigned, store halfword
LW, SW	Load word, store word (to/from integer registers)
LF, LD, SF, SD	Load SP float, load DP float, store SP float, store DP float
MOVI2S, MOVS2I	Move from/to integer register to/from a special register
MOVF, MOVD	Copy one floating-point register or a DP pair to another register or pair
MOVFP2I, MOVI2FP	Move 32 bits from/to FP registers to/from integer registers
Arithmetic, logical	**Operations on integer or logical data in integer registers; signed arithmetic instructions trap on overflow**
ADD, ADDI, ADDU, ADDUI	Add, add immediate (all immediates are 16 bits); signed and unsigned
SUB, SUBI, SUBU, SUBUI	Subtract, subtract immediate; signed and unsigned
MULT, MULTU, DIV, DIVU	Multiply and divide, signed and unsigned; operands must be floating-point registers; all operations take and yield 32-bit values
AND, ANDI	And, and immediate
OR, ORI, XOR, XORI	Or, or immediate, exclusive or, exclusive or immediate
LHI	Load high immediate—loads upper half of register with immediate
SLL, SRL, SRA, SLLI, SRLI, SRAI	Shifts: both immediate (S__I) and variable form (S__); shifts are shift left logical, right logical, right arithmetic
S__, S__I	Set conditional: "__" may be EQ, NE, LT, GT, LE, GE
Control	**Conditional branches and jumps; PC-relative or through register**
BEQZ, BNEZ	Branch integer register equal/not equal to zero; 16-bit offset from PC
BFPT, BFPF	Test comparison bit in the FP status register and branch; 16-bit offset from PC
J, JR	Jumps: 26-bit offset from PC (J) or target in register (JR)
JAL, JALR	Jump and Link: save PC+4 to R31, target is 26-bit offset from PC (JAL) or a register (JALR)
TRAP	Transfer to operating system at a vectored address (see Chapter 5)
RFE	Return to user code from an exception; restore user mode (see Chapter 5)
Floating point	**Floating-point operations on DP and SP formats**
ADDD, ADDF	Add DP,SP numbers
SUBD, SUBF	Subtract DP,SP numbers
MULTD, MULTF	Multiply DP,SP floating point
DIVD, DIVF	Divide DP, SP floating point
CVTF2D, CVTF2I, CVTD2F, CVTD2I, CVTI2F, CVTI2D	Convert instructions: CVTx2y converts from type x to type y, where x and y are one of I (integer), D (double precision), or F (single precision); both operands are in the FP registers
__D, __F	DP and SP compares: "__" may be EQ, NE, LT, GT, LE, GE; sets comparison bit in FP status register

Notation	Meaning	Example	Meaning
←	Data transfer. Length of the transfer is given by the destination's length; the length is specified when not clear.	R1←R2;	Transfer contents of R2 to R1. Registers have a fixed length, so transfers shorter than the register size must indicate which bits are used.
M	Array of memory accessed in bytes. The starting address for a transfer is indicated as the index to the memory array.	R1←M[x];	Place contents of memory location x into R1. If a transfer starts at M[i] and requires 4 bytes, the transferred bytes are M[i], M[i+1], M[i+2], and M[i+3].
$\leftarrow_n$	Transfer an *n*-bit field, used whenever length of transfer is not clear.	$M[y] \leftarrow_{16} M[x]$;	Transfer 16 bits starting at memory location x to memory location y. The length of the two sides should match.
X_n	Subscript selects a bit.	$R1_0 \leftarrow 0$;	Change sign bit of R1 to 0. (Bits are numbered from MSB starting at 0.)
$X_{m..n}$	Subscript selects a bit field.	$R3_{24..31} \leftarrow M[x]$;	Moves contents of memory location x into low-order byte of R3.
X^n	Superscript replicates a field.	$R3_{0..23} \leftarrow 0^{24}$;	Sets high-order three bytes of R3 to 0.
##	Concatenates two fields.	$R3 \leftarrow 0^{24}$ ## M[x]	Moves contents of location x into low byte of R3; clears upper three bytes.
		F2##F3←$_{64}$M[x];	Moves 64 bits from memory starting at location x; first 32 bits go into F2, second 32 into F3.
*, &	Dereference a pointer; get the address of a variable.	P*←&x;	Assign to object pointed to by p the address of the variable x.
<<, >>	C logical shifts (left,right)	R1 << 5	Shift R1 left 5 bits.
==, !=, >, <, >=, <=	C relational operators: equal, not equal, greater, less, greater or equal, less or equal	(R1==R2) & (R3!=R4)	True if the contents of R1 equal the contents of R2 and the contents of R3 do not equal the contents of R4.
&, |, ^, !	C bitwise logical operations: and, or, exclusive or, and complement.	(R1 & (R2 | R3))	Bitwise and of R1 and the bitwise or of R2 and R3.

DLX Pipeline Structure

Stage	ALU instruction	Load or store instruction	Branch instruction
IF	IR←Mem[PC]; PC←PC+4;	IR←Mem[PC]; PC←PC+4;	IR←Mem[PC]; PC←PC+4;
ID	A←Rs1; B←Rs2; PC1←PC IR1←IR	A←Rs1; B←Rs2; PC1←PC IR1←IR	A←Rs1; B←Rs2; PC1←PC IR1←IR
EX	ALUoutput←A *op* B; or ALUoutput←A *op* $((IR1_{16})^{16}\#\#IR1_{16..31})$;	DMAR←A+ $((IR1_{16})^{16}\#\#IR1_{16..31})$; SMDR←B;	ALUoutput←PC1 + $((IR1_{16})^{16}\#\#IR1_{16..31})$; cond←(Rs1 *op* 0);
MEM	ALUoutput1←ALUoutput	LMDR←Mem[DMAR]; or Mem[DMAR]←SMDR;	if (cond) PC←ALUoutput;
WB	Rd←ALUoutput1;	Rd←LMDR;	