INTEL CORPORATION
v.
PACT XPP SCHWEIZ AG

IPR2020-00541

Patent 9,075,605

Patent Owner's Presentation

PACT EX. 2028

Power Management involves "paradoxical goals"

8. Processor Control

This section describes the OS runtime aspects of managing the processor's power consumption and other controls while the system is in the *working* state[12]. The major controls over the processors are:

?? Processor power states: C0, C1, C2, C3
?? Processor clock throttling
?? Cooling control

These controls are used in combination by the operating software to achieve the desired balance of the following, sometimes paradoxical, goals:

?? Performance
?? Power consumption and battery life
?? Thermal requirements
?? Noise level requirements

Because the goals interact with each other, the operating software needs to implement a policy as to when and where tradeoffs between the goals are to be made[13]. For example, the operating software would determine when the audible noise of the fan is undesirable and would trade off that requirement for lower thermal requirements, which can lead to lower processing performance. Each processor control is discussed in the following sections along with how the control interacts with the various goals.

Ex. 1020 (ACPI) at 213.

The ’605 Patent

U.S. Patent 9,075,605 B2
Independent Claim 1

1. A method for operating a multiprocessor system comprising a plurality of data processing cells, each of the plurality of data processing units (a) including at least one Arithmetic Logic Unit (ALU) and a register unit and (b) being adapted for sequentially processing data, the method comprising:

the multiprocessor system setting a clock frequency, of at least a part of the multiprocessor system to a minimum in accordance with a number of pending operations of a first processor;

the multiprocessor system subsequently increasing the clock frequency of the at least the part of the multiprocessor system to a maximum in accordance with a number of pending operations of a second processor; and

the multiprocessor system subsequently reducing the clock frequency of the at least the part of the multiprocessor system in accordance with (a) an operating temperature threshold preventing over-temperature and (b) a hysteresis characteristic.

Ex. 1003 (US 9,075,605) at 15:6-16:5.

U.S. Patent 9,075,605 B2
Dependent Claim 2

1. A method for operating a multiprocessor system comprising a plurality of data processing cells, each of the plurality of data processing units (a) including at least one Arithmetic Logic Unit (ALU) and a register unit and (b) being adapted for sequentially processing data, the method comprising:

the multiprocessor system setting a clock frequency, of at least a part of the multiprocessor system to a minimum in accordance with a number of pending operations of a first processor;

the multiprocessor system subsequently increasing the clock frequency of the at least the part of the multiprocessor system to a maximum in accordance with a number of pending operations of a second processor; and

the multiprocessor system subsequently reducing the clock frequency of the at least the part of the multiprocessor system in accordance with (a) an operating temperature threshold preventing over-temperature and (b) a hysteresis characteristic.

Ex. 1003 (US 9,075,605) at 15:6-16:5.

2. The method of claim **1** wherein:
said clock frequency is determined in part by processor voltage.

Ex. 1003 (US 9,075,605) at 16:6-8.

PACT EX. 2028

U.S. Patent 9,075,605 B2
Dependent Claim 3

1. A method for operating a multiprocessor system comprising a plurality of data processing cells, each of the plurality of data processing units (a) including at least one Arithmetic Logic Unit (ALU) and a register unit and (b) being adapted for sequentially processing data, the method comprising:

the multiprocessor system setting a clock frequency, of at least a part of the multiprocessor system to a minimum in accordance with a number of pending operations of a first processor;

the multiprocessor system subsequently increasing the clock frequency of the at least the part of the multiprocessor system to a maximum in accordance with a number of pending operations of a second processor; and

the multiprocessor system subsequently reducing the clock frequency of the at least the part of the multiprocessor system in accordance with (a) an operating temperature threshold preventing over-temperature and (b) a hysteresis characteristic.

Ex. 1003 (US 9,075,605) at 15:6-16:5.

3. The method of claim **1** wherein:

said clock frequency is determined in part by available power.

Ex. 1003 (US 9,075,605) at 16:9-11.

PACT EX. 2028

U.S. Patent 9,075,605 B2
Dependent Claim 4

1. A method for operating a multiprocessor system comprising a plurality of data processing cells, each of the plurality of data processing units (a) including at least one Arithmetic Logic Unit (ALU) and a register unit and (b) being adapted for sequentially processing data, the method comprising:

the multiprocessor system setting a clock frequency, of at least a part of the multiprocessor system to a minimum in accordance with a number of pending operations of a first processor;

the multiprocessor system subsequently increasing the clock frequency of the at least the part of the multiprocessor system to a maximum in accordance with a number of pending operations of a second processor; and

the multiprocessor system subsequently reducing the clock frequency of the at least the part of the multiprocessor system in accordance with (a) an operating temperature threshold preventing over-temperature and (b) a hysteresis characteristic.

Ex. 1003 (US 9,075,605) at 15:6-16:5.

4. The method of claim **1** wherein:
said clock frequency is determined in part by a user setting.

Ex. 1003 (US 9,075,605) at 16:12-13.

PACT EX. 2028

U.S. Patent 9,075,605 B2
Dependent Claim 5

1. A method for operating a multiprocessor system comprising a plurality of data processing cells, each of the plurality of data processing units (a) including at least one Arithmetic Logic Unit (ALU) and a register unit and (b) being adapted for sequentially processing data, the method comprising:

the multiprocessor system setting a clock frequency, of at least a part of the multiprocessor system to a minimum in accordance with a number of pending operations of a first processor;

the multiprocessor system subsequently increasing the clock frequency of the at least the part of the multiprocessor system to a maximum in accordance with a number of pending operations of a second processor; and

the multiprocessor system subsequently reducing the clock frequency of the at least the part of the multiprocessor system in accordance with (a) an operating temperature threshold preventing over-temperature and (b) a hysteresis characteristic.

Ex. 1003 (US 9,075,605) at 15:6-16:5.

5. The method of claim **1** wherein:

said clock frequency is determined in part by available power and a user setting.

Ex. 1003 (US 9,075,605) at 16:14-16.

PACT EX. 2028

U.S. Patent 9,075,605 B2
Dependent Claim 6

1. A method for operating a multiprocessor system comprising a plurality of data processing cells, each of the plurality of data processing units (a) including at least one Arithmetic Logic Unit (ALU) and a register unit and (b) being adapted for sequentially processing data, the method comprising:

the multiprocessor system setting a clock frequency, of at least a part of the multiprocessor system to a minimum in accordance with a number of pending operations of a first processor;

the multiprocessor system subsequently increasing the clock frequency of the at least the part of the multiprocessor system to a maximum in accordance with a number of pending operations of a second processor; and

the multiprocessor system subsequently reducing the clock frequency of the at least the part of the multiprocessor system in accordance with (a) an operating temperature threshold preventing over-temperature and (b) a hysteresis characteristic.

Ex. 1003 (US 9,075,605) at 15:6-16:5.

6. The method of claim **1** wherein:

said number of pending operations is determined at by the fill level of one or more buffers.

Ex. 1003 (US 9,075,605) at 16:17-19.

U.S. Patent 9,075,605 B2
Independent Claim 1

1. A method for operating a multiprocessor system comprising a plurality of data processing cells, each of the plurality of data processing units (a) including at least one Arithmetic Logic Unit (ALU) and a register unit and (b) being adapted for sequentially processing data, the method comprising:

the multiprocessor system setting a clock frequency, of at least a part of the multiprocessor system to a minimum in accordance with a number of pending operations of a first processor;

the multiprocessor system subsequently increasing the clock frequency of the at least the part of the multiprocessor system to a maximum in accordance with a number of pending operations of a second processor; and

the multiprocessor system subsequently reducing the clock frequency of the at least the part of the multiprocessor system in accordance with (a) an operating temperature threshold preventing over-temperature and (b) a hysteresis characteristic.

Ex. 1003 (US 9,075,605) at 15:6-16:5.

3. The method of claim **1** wherein:
said clock frequency is determined in part by available power.

Ex. 1003 (US 9,075,605) at 16:8-11.

4. The method of claim **1** wherein:
said clock frequency is determined in part by a user setting.

Ex. 1003 (US 9,075,605) at 16:12-13.

5. The method of claim **1** wherein:
said clock frequency is determined in part by available power and a user setting.

Ex. 1003 (US 9,075,605) at 16:14-16.

PACT EX. 2028

A POSITA would not be motivated to combine different embodiments of Nicol

PACT EX. 2028

U.S. Patent 9,075,605 B2
Independent Claim 1

1. A method for operating a multiprocessor system comprising a plurality of data processing cells, each of the plurality of data processing units (a) including at least one Arithmetic Logic Unit (ALU) and a register unit and (b) being adapted for sequentially processing data, the method comprising:

the multiprocessor system setting a clock frequency, of at least a part of the multiprocessor system to a minimum in accordance with a number of pending operations of a first processor;

the multiprocessor system subsequently increasing the clock frequency of the at least the part of the multiprocessor system to a maximum in accordance with a number of pending operations of a second processor; and

the multiprocessor system subsequently reducing the clock frequency of the at least the part of the multiprocessor system in accordance with (a) an operating temperature threshold preventing over-temperature and (b) a hysteresis characteristic.

Ex. 1003 (US 9,075,605) at 15:6-16:5.

PACT EX. 2028

Nicol

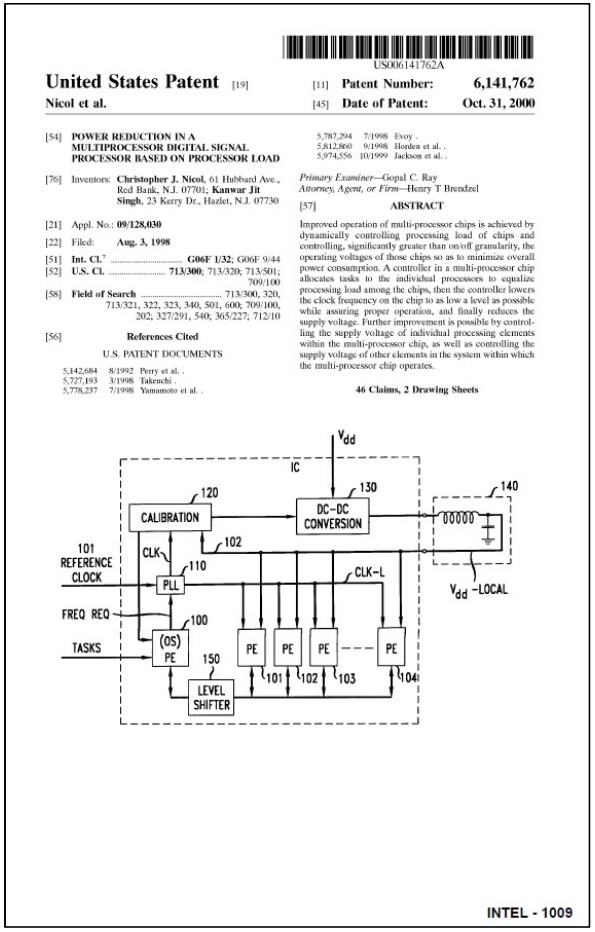

US006141762A

United States Patent [19]
Nicol et al.

[11] **Patent Number: 6,141,762**
[45] **Date of Patent: Oct. 31, 2000**

[54] **POWER REDUCTION IN A MULTIPROCESSOR DIGITAL SIGNAL PROCESSOR BASED ON PROCESSOR LOAD**

[76] Inventors: **Christopher J. Nicol**, 61 Hubbard Ave., Red Bank, N.J. 07701; **Kanwar Jit Singh**, 23 Kerry Dr., Hazlet, N.J. 07730

[21] Appl. No.: **09/128,030**

[22] Filed: **Aug. 3, 1998**

[51] **Int. Cl.**[7] **G06F 1/32**; G06F 9/44

[52] **U.S. Cl.** **713/300**; 713/320; 713/501; 709/100

[58] **Field of Search** 713/300, 320, 713/321, 322, 323, 340, 501, 600; 709/100, 202; 327/291, 540; 365/227; 712/10

[56] **References Cited**

U.S. PATENT DOCUMENTS

5,142,684 8/1992 Perry et al. .
5,727,193 3/1998 Takeuchi .
5,778,237 7/1998 Yamamoto et al. .
5,787,294 7/1998 Evoy .
5,812,860 9/1998 Horden et al. .
5,974,556 10/1999 Jackson et al. .

Primary Examiner—Gopal C. Ray
Attorney, Agent, or Firm—Henry T Brendzel

[57] **ABSTRACT**

Improved operation of multi-processor chips is achieved by dynamically controlling processing load of chips and controlling, significantly greater than on/off granularity, the operating voltages of those chips so as to minimize overall power consumption. A controller in a multi-processor chip allocates tasks to the individual processors to equalize processing load among the chips, then the controller lowers the clock frequency on the chip to as low a level as possible while assuring proper operation, and finally reduces the supply voltage. Further improvement is possible by controlling the supply voltage of individual processing elements within the multi-processor chip, as well as controlling the supply voltage of other elements in the system within which the multi-processor chip operates.

46 Claims, 2 Drawing Sheets

INTEL - 1009

[54] **POWER REDUCTION IN A MULTIPROCESSOR DIGITAL SIGNAL PROCESSOR BASED ON PROCESSOR LOAD**

[57] **ABSTRACT**

Improved operation of multi-processor chips is achieved by dynamically controlling processing load of chips and controlling, significantly greater than on/off granularity, the operating voltages of those chips so as to minimize overall power consumption. A controller in a multi-processor chip allocates tasks to the individual processors to equalize processing load among the chips, then the controller lowers the clock frequency on the chip to as low a level as possible while assuring proper operation, and finally reduces the supply voltage. Further improvement is possible by controlling the supply voltage of individual processing elements within the multi-processor chip, as well as controlling the supply voltage of other elements in the system within which the multi-processor chip operates.

Ex. 1009 (Nicol) Title and Abstract

PACT EX. 2028

Nicol

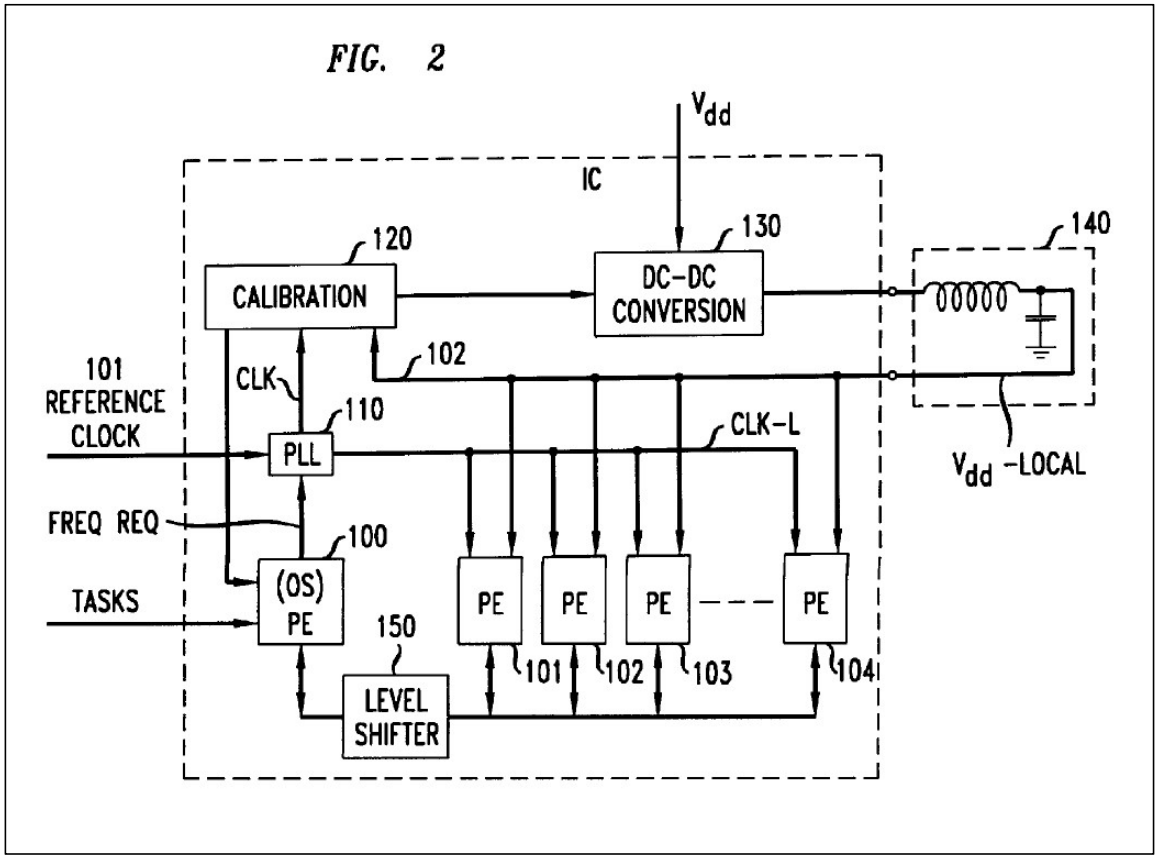

Ex. 1009 (Nicol) at Fig. 2.

In the FIG. **2** arrangement, in accordance with the principles disclosed herein, the applications that need to be processed are mapped to the N PEs under control of real time operating system (RTOS) executed on PE **100**. If the number of instructions that need to be executed for each task is known and made available to the operating system, a scheduler within the operating system can use this information to determine the best way to allocate the tasks to the available processors in order to balance the computation. The intermediate goal, of course, is to maximize the parallelism and to evenly distribute the load presented to the FIG. **2** system among all of the PE's.

Ex. 1009 (Nicol) at 3:9-21.

Hence, the operating system of PE **100** needs to ascertain the required completion time, divide the collection of tasks as evenly as possible (in terms of needed processing time), consider the PE with the tasks that require the most time to carry out, and adjust the clock frequency to insure that the most heavily loaded PE carries out its assigned tasks within the required completion time. Once the frequency is thus determined, a minimum supply voltage can be determined.

Ex. 1009 (Nicol) at 3:66-4:6.

PACT EX. 2028

Nicol

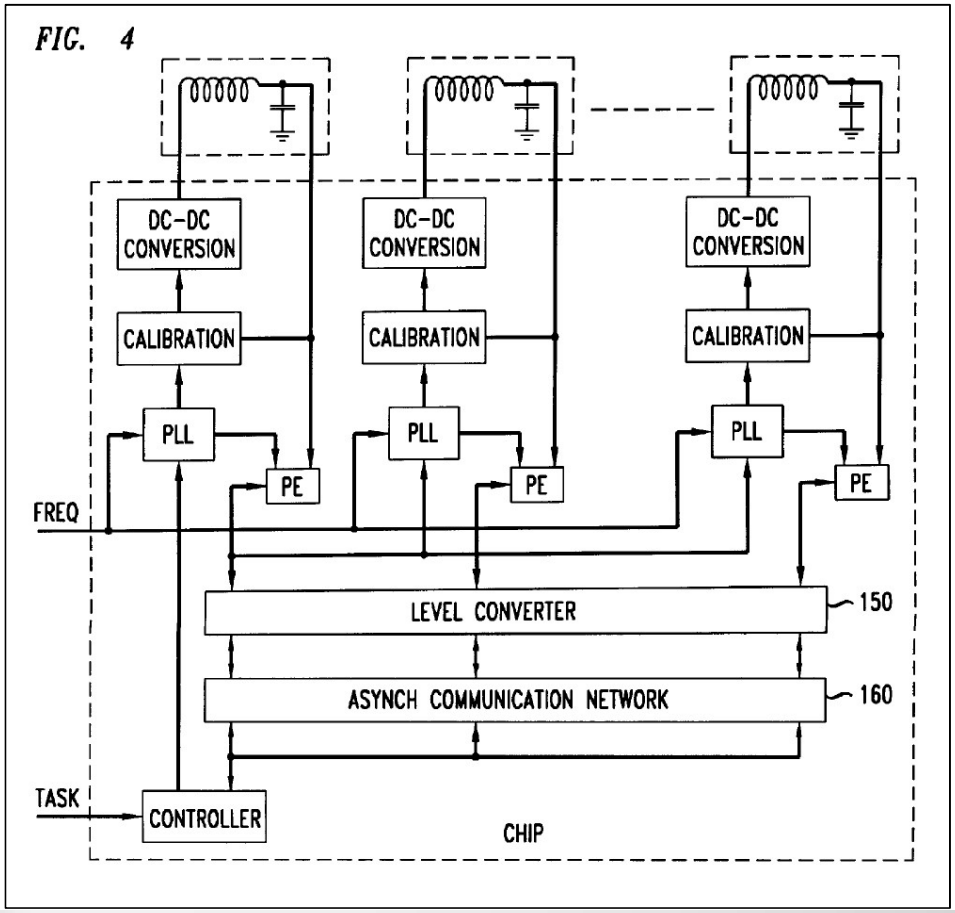

The notion of adjusting operating frequency to load and adjusting supply voltage to track the operating frequency can be extended to allow each PE to have its own supply voltage. The benefit of this approach for some applications becomes apparent when it is realized that the chip-wise voltage scaling is most effective when the load of the computation can be evenly distributed across all of the PEs. In some applications, however, one may encounter tasks that cannot be partitioned into concurrent evenly-loaded threads and, therefore, some PE within the multiprocessor would require a higher operating frequency and a higher operating voltage. This would require raising the frequency and voltage of the entire multiprocessor chip.

Ex. 1009 (Nicol) at 5:53-65.

A separate power supply for each PE in a chip overcomes this limitation by allowing the operating system to independently program the lowest operating frequency and corresponding lowest supply voltage for each PE. The architecture of such an arrangement is shown in FIG. **4**. Each PE in FIG. **4** needs an independent controller that performs the functions of PE **100** (except it does not divide tasks among PEs). As shown in FIG. **4**, all of the controllers are embodied in a single controller **200**, which may be just another processing element of the integrated circuit that contains the other processing elements. Each processing element also

Ex. 1009 (Nicol) at 5:66-6:9.

PACT EX. 2028

Petitioner identifies ACN in Figure 4 for "sequentially processing data"

> ***d) "and (b) being adapted for sequentially processing data, the method comprising:"***
>
> Nicol renders obvious this limitation under Patent Owner's construction of the term "sequentially processing data" to mean "passing results onto one or more other data processing units which are subsequently processing data." Ex. 1001 ¶147; Ex. 1016 (JCC) at 11 and 17. Specifically, Nicol teaches the need for "***synchronizing the exchange of data*** among the PEs [green] of a multiprocessor chip" through an "asynch[ronous] communication network" 160 (highlighted in blue in Figure 4). Ex. 1009, 6:26–28, 6:14–18, Figure 4. For example, Nicol's annotated Figure 4 below depicts the passing of data (red) from PE1 (green) through the asynchronous communication network (blue) to PE2 (green) for further processing. Ex. 1001 ¶148; Ex. 1009, Figure 4.

Paper 2 (Petition) at 30-31.

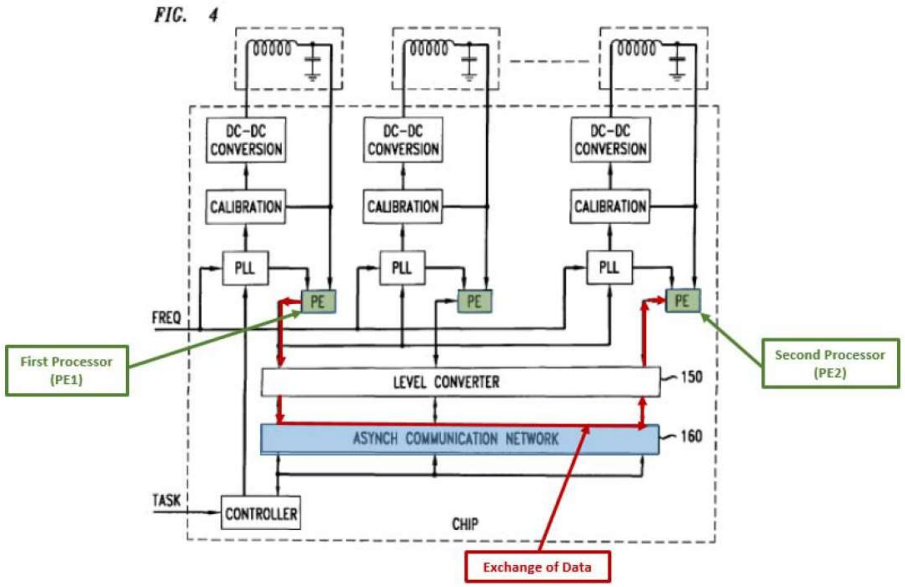

Id., Figure 4 (annotated).

Paper 2 (Petition) at 31.

PACT EX. 2028

Petitioner identifies ACN in Figure 4 for "sequentially processing data"

1. "being adapted for sequentially processing data"

The Petition provides three independent reasons for why a POSITA would understand Nicol to disclose, or at least render obvious, "being adapted for sequentially processing data." *See* Pet., 30-35. While the POR mischaracterizes these independent grounds as being part of a "three step analysis" (POR, 20-21), no such effort is necessary: as explained in the Petition, Nicol discloses that data can be processed by one PE and then sequentially passed to another PE for further processing. Pet., 30-32. Nicol discloses this functionality via its PE-to-PE asynchronous communication network ("ACN") which passes data from one PE to another PE for sequential processing. *See id.*; Ex. 1009, 6:18-21 ("a synchronization schema must be implemented when there is a need to ***communicate data between PEs*** (or with other system elements) that operate at different frequencies."), 6:25-33 (discussing the need and the asynchronous communication network solution for "synchronizing ***the exchange of data among the PEs*** of a multiprocessor chip"); Ex. 1055, ¶16.

Paper 27 (Reply) at 8-9.

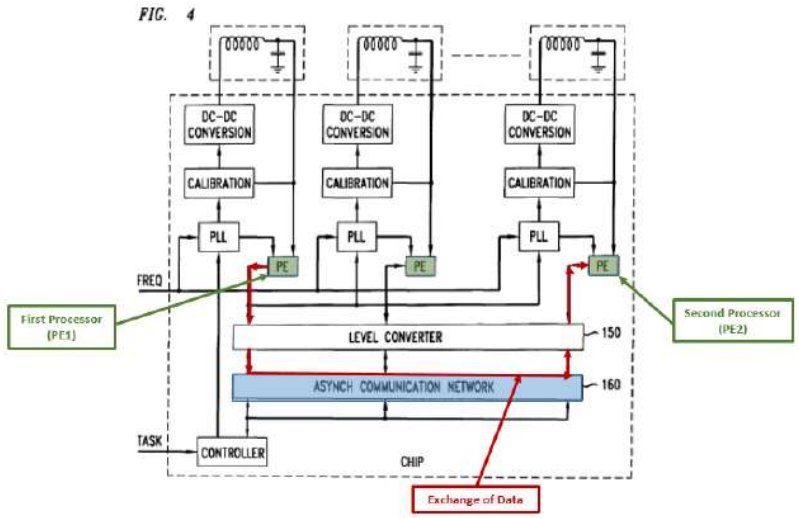

Pet., 30-32; Ex. 1009, Figure 4 (annotated). Despite this clear disclosure in Nicol,

Paper 27 (Reply) at 9.

PACT EX. 2028

Nicol

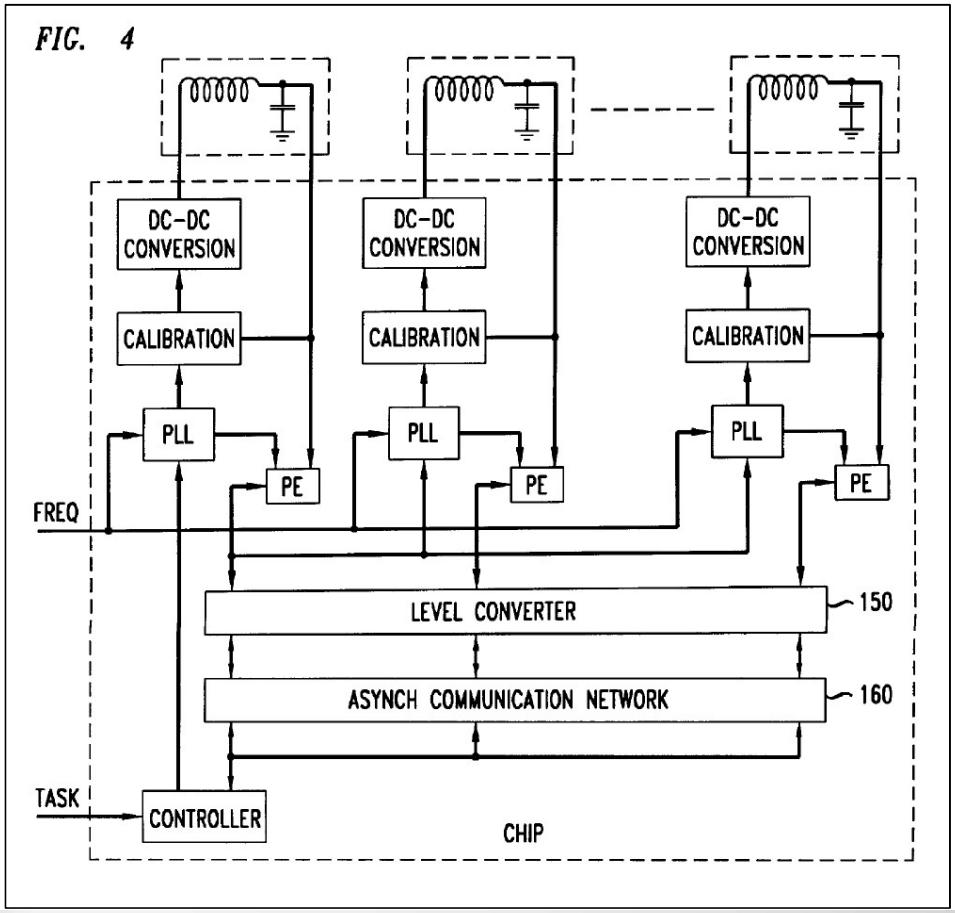

It may be noted that if the frequencies at which the individual PEs operate differ from one another and from other elements within the system where the multiprocessor chip is employed, there is an issue of synchronization that must be addressed. That is, a synchronization schema must be implemented when there is a need to communicate data between PEs (or with other system elements) that operate at different frequencies. It is possible to arrange the frequencies so that the collection of tasks that are assigned to the multiprocessor is completed at a predetermined time. In such a case, the synchronization problem of the multiprocessor vis-a-vis other elements within the system where the multiprocessor is employed is minimized. However, that leaves the issue of synchronizing the exchange of data among the PEs of a multiprocessor chip.

To effect such synchronization, each PE within the FIG. **4** arrangement is connection to an arrangement comprising elements **150** and **160**. Level converter **150** converts the variable voltage swings of the PEs to a fixed level swing, and network **160** resolves the issue of different clock domains.

Ex. 1009 (Nicol) at 6:14-33.

Petitioner relies on methodology of Figure 2 for setting clock frequency

> ***e)*** ***"the multiprocessor system setting a clock frequency, of at least a part of the multiprocessor system to a minimum in accordance with a number of pending operations of a first processor;"***
>
> Nicol discloses this claim element. Ex. 1001 ¶161. Nicol explains that the controller (red) in its multiprocessor system sets a clock frequency to a minimum according to a number of pending operations (instructions) of each of the PEs of the
>
> * * * * *
>
> Indeed, Nicol explains that the controller "allocates tasks to the individual processors [PEs, including PE1] to equalize processing load among the [PEs, including PE1], then the controller ***lowers the clock frequency*** on the chip to ***as low a level as possible*** while assuring proper operation, and finally reduces the supply voltage." Ex. 1009, 2:23–27. Nicol further explains that the controller takes into account the number of pending operations (instructions) of each PE—including of PE1 ("first processor")—in its determination of processing load and task allocation. Ex. 1009, 3:13–18 ("If ***the number of instructions*** that need to be executed for each task is known and made available to the operating system, a scheduler within the operating system can use this information to determine the best way to allocate the

Paper 2 (Petition) at 35-36.

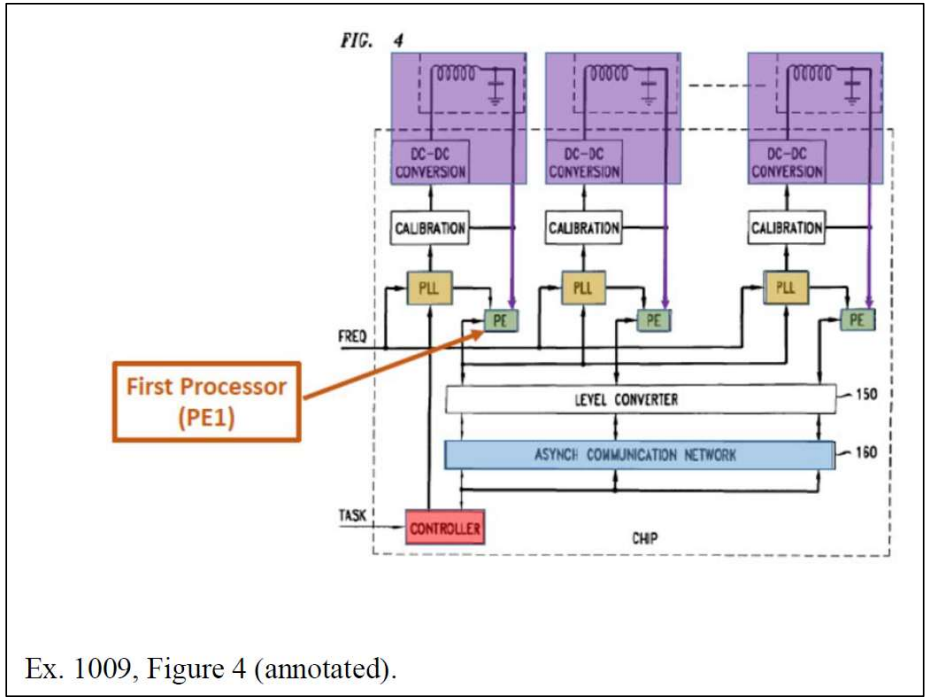

Ex. 1009, Figure 4 (annotated).

Paper 2 (Petition) at 36.

PACT EX. 2028

Petitioner relies on methodology of Figure 2 for setting clock frequency

f) ***"the multiprocessor system subsequently increasing the clock frequency of the at least the part of the multiprocessor system to a maximum in accordance with a number of pending operations of a second processor;"***

Nicol discloses this element. Ex. 1001 ¶173. As discussed above in Section X.A (Overview of Nicol) and in the preceding section XI.A.2.e, the controller in Nicol's multiprocessor system adjusts the clock frequency of the system according to a number of pending operations of each of the PEs—including the number of pending operations of PE2 ("second processor")—identified in Figure 4 (annotated) below. *Id.*; Ex. 1009, 2:23–27 (discussing allocating tasks to the PEs and setting an appropriate clock frequency for task execution) and 3:13–18 (discussing using "***the number of instructions*** that need to be executed for each task" for task allocation).

* * * * *

Notably, the controller takes into account the number of pending instructions for each PE—including PE1 ("first processor") and PE2 ("second processor")—because it sets the clock frequency based on identifying the processing requirement of "the most heavily loaded PE." Ex. 1001, ¶176; Ex. 1009, 3:66–4:5 (explaining that the operating system "needs to ascertain the required completion time [of the tasks], divide the collection of tasks as evenly as possible (in terms of needed processing time), consider the PE with the tasks that require the most time to carry out, and ***adjust the clock frequency*** [common to all PEs] to insure that ***the most heavily loaded PE*** carries out its assigned task within the required completion time.").

Paper 2 (Petition) at 39-40.

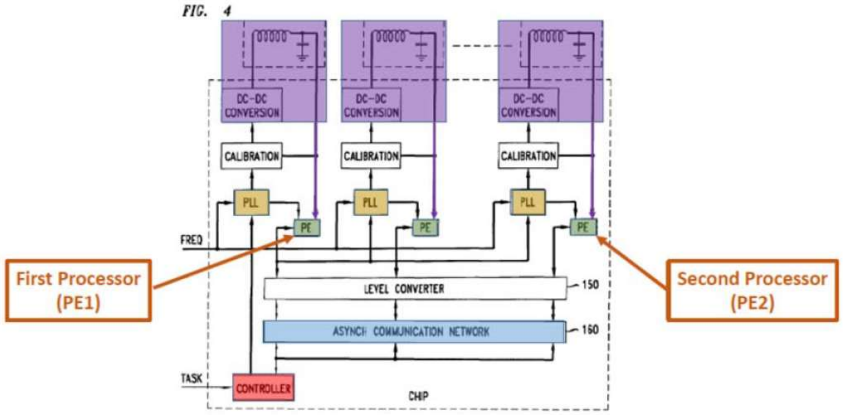

Ex. 1009, Figure 4 (annotated).

Paper 2 (Petition) at 40.

PACT EX. 2028

Nicol's controllers

Figure 2 controller PE 100

Hence, the operating system of PE **100** needs to ascertain the required completion time, divide the collection of tasks as evenly as possible (in terms of needed processing time), consider the PE with the tasks that require the most time to carry out, and adjust the clock frequency to insure that the most heavily loaded PE carries out its assigned tasks within the required completion time. Once the frequency is thus determined, a minimum supply voltage can be determined.

Ex. 1009 (Nicol) at 3:66-4:6.

Figure 4 controllers PE 100 and PE 200

ture of such an arrangement is shown in FIG. **4**. Each PE in FIG. **4** needs an independent controller that performs the functions of PE **100** (except it does not divide tasks among PEs). As shown in FIG. **4**, all of the controllers are embodied in a single controller **200**, which may be just another processing element of the integrated circuit that contains the other processing elements. Each processing element also requires a calibration circuit like circuit **120**, and a voltage converter circuit like circuits **130** and **140**. It also has a PE **200** that assigns the tasks given to the multi-processor chip of FIG. **4** among the PEs.

Ex. 1009 (Nicol) at 6:3-13.

PACT EX. 2028

A POSITA would not be motivated to combine Nicol and Kling

PACT EX. 2028

Nicol

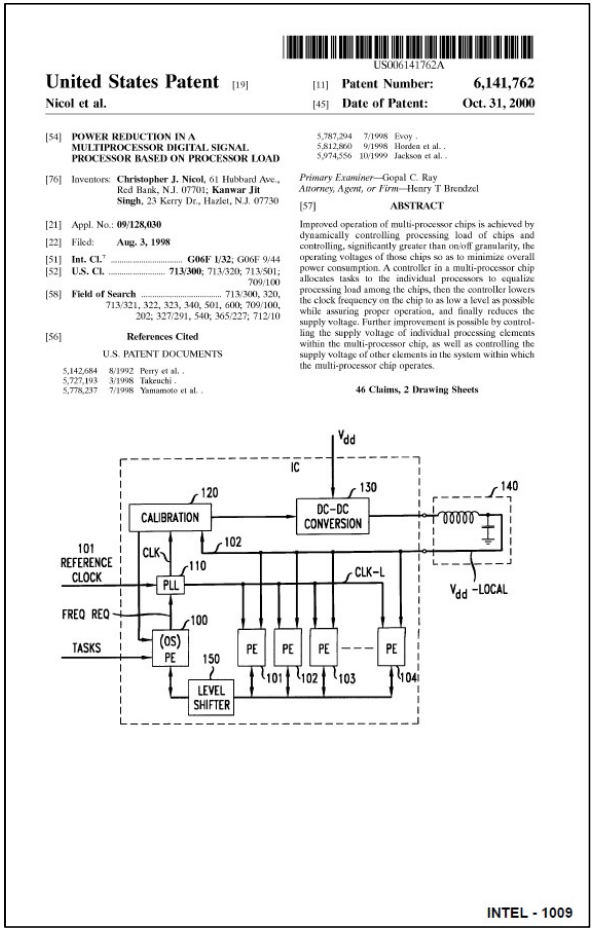

US006141762A

United States Patent [19]
Nicol et al.

[11] **Patent Number: 6,141,762**
[45] **Date of Patent: Oct. 31, 2000**

[54] **POWER REDUCTION IN A MULTIPROCESSOR DIGITAL SIGNAL PROCESSOR BASED ON PROCESSOR LOAD**

[76] Inventors: **Christopher J. Nicol**, 61 Hubbard Ave., Red Bank, N.J. 07701; **Kanwar Jit Singh**, 23 Kerry Dr., Hazlet, N.J. 07730

[21] Appl. No.: **09/128,030**

[22] Filed: **Aug. 3, 1998**

[51] **Int. Cl.**[7] **G06F 1/32**; G06F 9/44

[52] **U.S. Cl.** **713/300**; 713/320; 713/501; 709/100

[58] **Field of Search** 713/300, 320, 713/321, 322, 323, 340, 501, 600; 709/100, 202; 327/291, 540; 365/227; 712/10

[56] **References Cited**

U.S. PATENT DOCUMENTS

5,142,684 8/1992 Perry et al. .
5,727,193 3/1998 Takeuchi .
5,778,237 7/1998 Yamamoto et al. .
5,787,294 7/1998 Evoy .
5,812,860 9/1998 Horden et al. .
5,974,556 10/1999 Jackson et al. .

Primary Examiner—Gopal C. Ray
Attorney, Agent, or Firm—Henry T Brendzel

[57] **ABSTRACT**

Improved operation of multi-processor chips is achieved by dynamically controlling processing load of chips and controlling, significantly greater than on/off granularity, the operating voltages of those chips so as to minimize overall power consumption. A controller in a multi-processor chip allocates tasks to the individual processors to equalize processing load among the chips, then the controller lowers the clock frequency on the chip to as low a level as possible while assuring proper operation, and finally reduces the supply voltage. Further improvement is possible by controlling the supply voltage of individual processing elements within the multi-processor chip, as well as controlling the supply voltage of other elements in the system within which the multi-processor chip operates.

46 Claims, 2 Drawing Sheets

INTEL - 1009

[54] **POWER REDUCTION IN A MULTIPROCESSOR DIGITAL SIGNAL PROCESSOR BASED ON PROCESSOR LOAD**

[57] **ABSTRACT**

Improved operation of multi-processor chips is achieved by dynamically controlling processing load of chips and controlling, significantly greater than on/off granularity, the operating voltages of those chips so as to minimize overall power consumption. A controller in a multi-processor chip allocates tasks to the individual processors to equalize processing load among the chips, then the controller lowers the clock frequency on the chip to as low a level as possible while assuring proper operation, and finally reduces the supply voltage. Further improvement is possible by controlling the supply voltage of individual processing elements within the multi-processor chip, as well as controlling the supply voltage of other elements in the system within which the multi-processor chip operates.

Ex. 1009 (Nicol), Title and Abstract.

PACT EX. 2028

Nicol

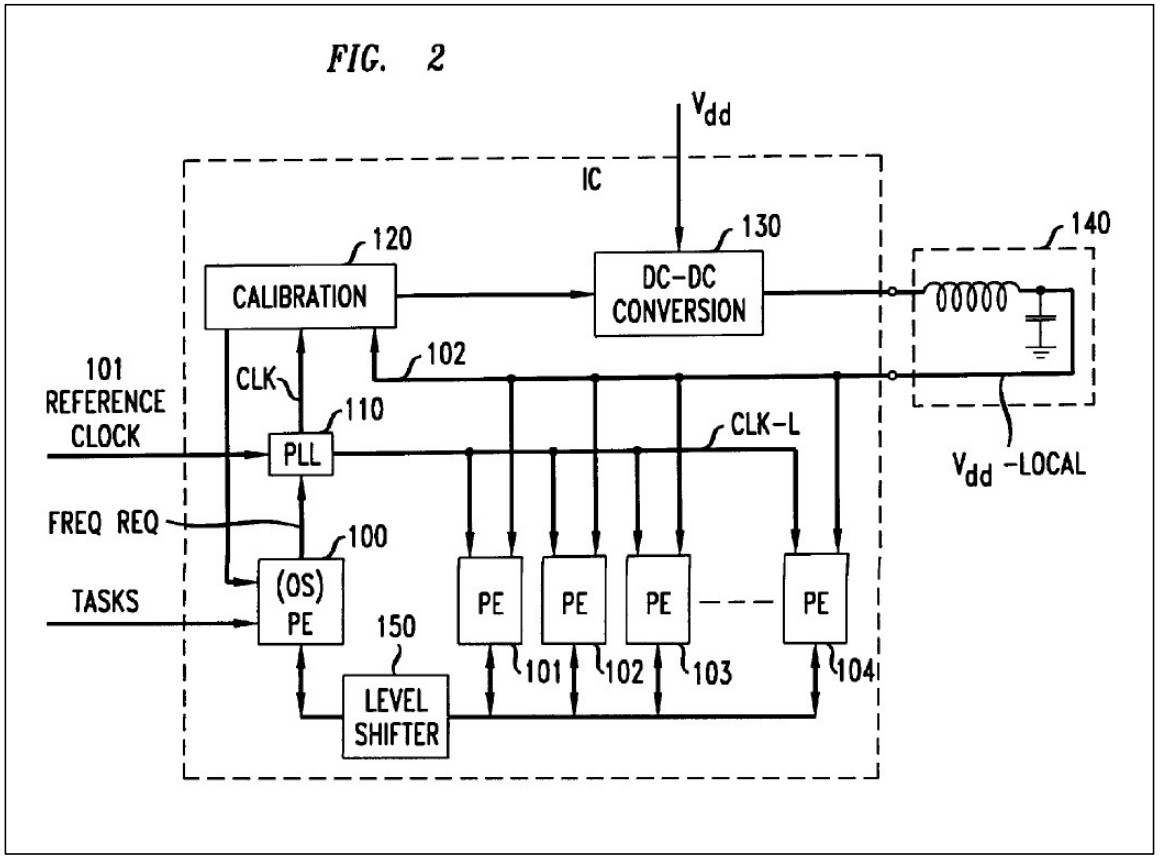

Ex. 1009 (Nicol) at Fig. 2.

In the FIG. **2** arrangement, in accordance with the principles disclosed herein, the applications that need to be processed are mapped to the N PEs under control of real time operating system (RTOS) executed on PE **100**. If the number of instructions that need to be executed for each task is known and made available to the operating system, a scheduler within the operating system can use this information to determine the best way to allocate the tasks to the available processors in order to balance the computation. The intermediate goal, of course, is to maximize the parallelism and to evenly distribute the load presented to the FIG. **2** system among all of the PE's.

Ex. 1009 (Nicol) at 3:9-21.

Hence, the operating system of PE **100** needs to ascertain the required completion time, divide the collection of tasks as evenly as possible (in terms of needed processing time), consider the PE with the tasks that require the most time to carry out, and adjust the clock frequency to insure that the most heavily loaded PE carries out its assigned tasks within the required completion time. Once the frequency is thus determined, a minimum supply voltage can be determined.

Ex. 1009 (Nicol) at 3:66-4:6.

PACT EX. 2028

Kling

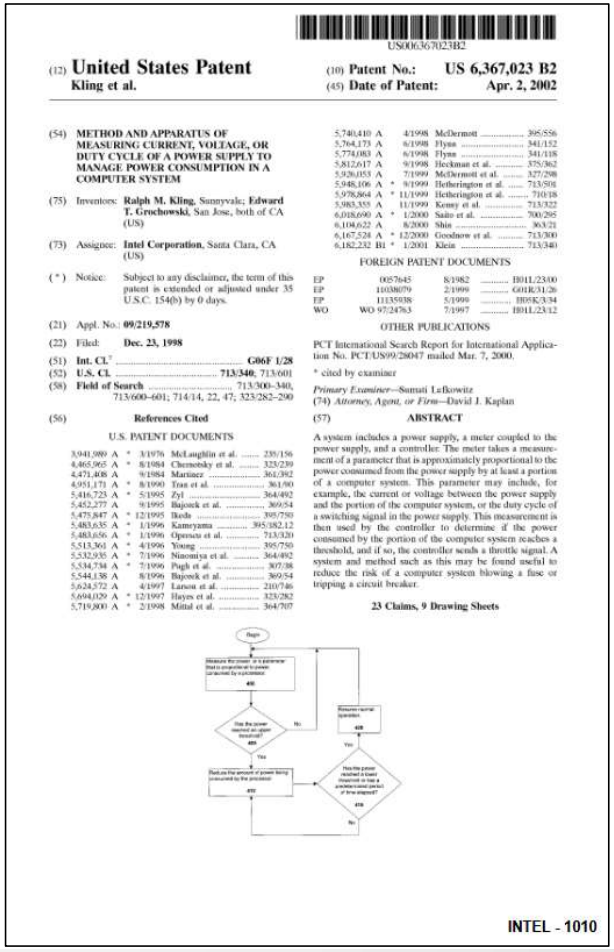

US006367023B2

(12) **United States Patent**
Kling et al.

(10) **Patent No.: US 6,367,023 B2**
(45) **Date of Patent: Apr. 2, 2002**

(54) **METHOD AND APPARATUS OF MEASURING CURRENT, VOLTAGE, OR DUTY CYCLE OF A POWER SUPPLY TO MANAGE POWER CONSUMPTION IN A COMPUTER SYSTEM**

(75) Inventors: **Ralph M. Kling**, Sunnyvale; **Edward T. Grochowski**, San Jose, both of CA (US)

(73) Assignee: **Intel Corporation**, Santa Clara, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **09/219,578**

(22) Filed: **Dec. 23, 1998**

(51) **Int. Cl.**[7] **G06F 1/28**
(52) **U.S. Cl.** **713/340**; 713/601
(58) **Field of Search** 713/300–340, 713/600–601; 714/14, 22, 47; 323/282–290

(56) **References Cited**

U.S. PATENT DOCUMENTS

3,941,989	A *	3/1976	McLaughlin et al.	235/156
4,465,965	A *	8/1984	Chemosky et al.	323/239
4,471,408	A	9/1984	Martinez	361/392
4,951,171	A *	8/1990	Tran et al.	361/90
5,416,723	A *	5/1995	Zyl	364/492
5,452,277	A	9/1995	Bajorek et al.	369/54
5,475,847	A *	12/1995	Ikeda	395/750
5,483,635	A *	1/1996	Kameyama	395/182.12
5,483,656	A *	1/1996	Oprescu et al.	713/320
5,513,361	A *	4/1996	Young	395/750
5,532,935	A *	7/1996	Ninomiya et al.	364/492
5,534,734	A *	7/1996	Pugh et al.	307/38
5,544,138	A	8/1996	Bajorek et al.	369/54
5,624,572	A	4/1997	Larson et al.	210/746
5,694,029	A *	12/1997	Hayes et al.	323/282
5,719,800	A *	2/1998	Mittal et al.	364/707
5,740,410	A	4/1998	McDermott	395/556
5,764,173	A	6/1998	Flynn	341/132
5,774,083	A	6/1998	Flynn	341/118
5,812,617	A	9/1998	Heckman et al.	375/362
5,926,053	A	7/1999	McDermott et al.	327/298
5,948,106	A *	9/1999	Hetherington et al.	713/501
5,978,864	A *	11/1999	Hetherington et al.	710/18
5,983,355	A	11/1999	Kenny et al.	713/322
6,018,690	A *	1/2000	Saito et al.	700/295
6,104,622	A	8/2000	Shin	363/21
6,167,524	A *	12/2000	Goodnow et al.	713/300
6,182,232	B1 *	1/2001	Klein	713/340

FOREIGN PATENT DOCUMENTS

EP	0057645	8/1982	H01L/23/00
EP	11038079	2/1999	G01R/31/26
EP	11135938	5/1999	H05K/3/34
WO	WO 97/24763	7/1997	H01L/23/12

OTHER PUBLICATIONS

PCT International Search Report for International Application No. PCT/US99/28047 mailed Mar. 7, 2000.

* cited by examiner

Primary Examiner—Suresh Lefkowitz
(74) *Attorney, Agent, or Firm*—David J. Kaplan

(57) **ABSTRACT**

A system includes a power supply, a meter coupled to the power supply, and a controller. The meter takes a measurement of a parameter that is approximately proportional to the power consumed from the power supply by at least a portion of a computer system. This parameter may include, for example, the current or voltage between the power supply and the portion of the computer system, or the duty cycle of a switching signal in the power supply. This measurement is then used by the controller to determine if the power consumed by the portion of the computer system reaches a threshold, and if so, the controller sends a throttle signal. A system and method such as this may be found useful to reduce the risk of a computer system blowing a fuse or tripping a circuit breaker.

23 Claims, 9 Drawing Sheets

INTEL - 1010

(54) **METHOD AND APPARATUS OF MEASURING CURRENT, VOLTAGE, OR DUTY CYCLE OF A POWER SUPPLY TO MANAGE POWER CONSUMPTION IN A COMPUTER SYSTEM**

(57) **ABSTRACT**

A system includes a power supply, a meter coupled to the power supply, and a controller. The meter takes a measurement of a parameter that is approximately proportional to the power consumed from the power supply by at least a portion of a computer system. This parameter may include, for example, the current or voltage between the power supply and the portion of the computer system, or the duty cycle of a switching signal in the power supply. This measurement is then used by the controller to determine if the power consumed by the portion of the computer system reaches a threshold, and if so, the controller sends a throttle signal. A system and method such as this may be found useful to reduce the risk of a computer system blowing a fuse or tripping a circuit breaker.

Ex. 1010 (Kling), Title and Abstract.

PACT EX. 2028

Kling

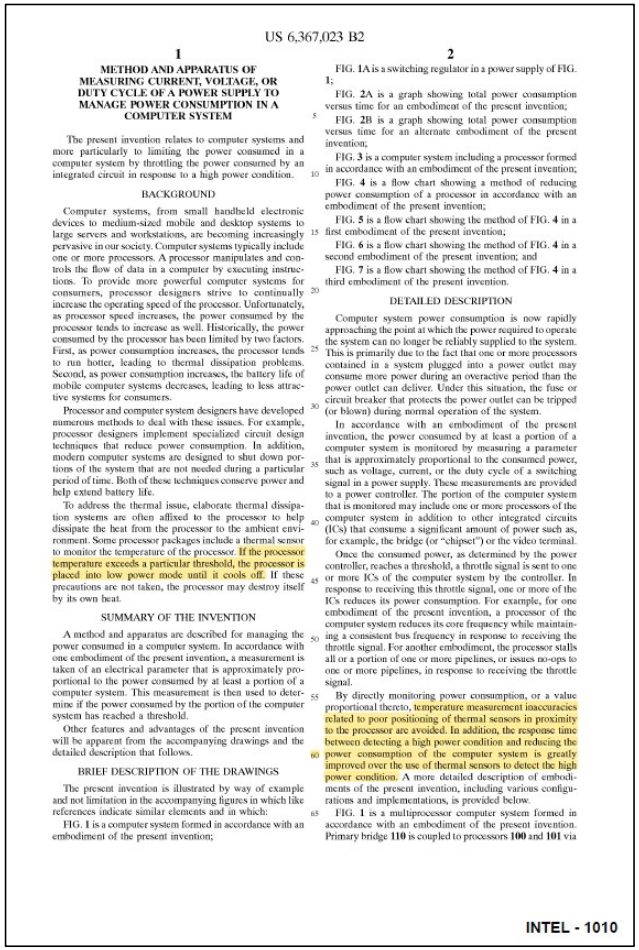

US 6,367,023 B2

1

METHOD AND APPARATUS OF MEASURING CURRENT, VOLTAGE, OR DUTY CYCLE OF A POWER SUPPLY TO MANAGE POWER CONSUMPTION IN A COMPUTER SYSTEM

The present invention relates to computer systems and more particularly to limiting the power consumed in a computer system by throttling the power consumed by an integrated circuit in response to a high power condition.

BACKGROUND

Computer systems, from small handheld electronic devices to medium-sized mobile and desktop systems to large servers and workstations, are becoming increasingly pervasive in our society. Computer systems typically include one or more processors. A processor manipulates and controls the flow of data in a computer by executing instructions. To provide more powerful computer systems for consumers, processor designers strive to continually increase the operating speed of the processor. Unfortunately, as processor speed increases, the power consumed by the processor tends to increase as well. Historically, the power consumed by the processor has been limited by two factors. First, as power consumption increases, the processor tends to run hotter, leading to thermal dissipation problems. Second, as power consumption increases, the battery life of mobile computer systems decreases, leading to less attractive systems for consumers.

Processor and computer system designers have developed numerous methods to deal with these issues. For example, processor designers implement specialized circuit design techniques that reduce power consumption. In addition, modern computer systems are designed to shut down portions of the system that are not needed during a particular period of time. Both of these techniques conserve power and help extend battery life.

To address the thermal issue, elaborate thermal dissipation systems are often affixed to the processor to help dissipate the heat from the processor to the ambient environment. Some processor packages include a thermal sensor to monitor the temperature of the processor. If the processor temperature exceeds a particular threshold, the processor is placed into low power mode until it cools off. If these precautions are not taken, the processor may destroy itself by its own heat.

SUMMARY OF THE INVENTION

A method and apparatus are described for managing the power consumed in a computer system. In accordance with one embodiment of the present invention, a measurement is taken of an electrical parameter that is approximately proportional to the power consumed by at least a portion of a computer system. This measurement is then used to determine if the power consumed by the portion of the computer system has reached a threshold.

Other features and advantages of the present invention will be apparent from the accompanying drawings and the detailed description that follows.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example and not limitation in the accompanying figures in which like references indicate similar elements and in which:

FIG. **1** is a computer system formed in accordance with an embodiment of the present invention;

2

FIG. 1A is a switching regulator in a power supply of FIG. 1;

FIG. **2A** is a graph showing total power consumption versus time for an embodiment of the present invention;

FIG. **2B** is a graph showing total power consumption versus time for an alternate embodiment of the present invention;

FIG. **3** is a computer system including a processor formed in accordance with an embodiment of the present invention;

FIG. **4** is a flow chart showing a method of reducing power consumption of a processor in accordance with an embodiment of the present invention;

FIG. **5** is a flow chart showing the method of FIG. **4** in a first embodiment of the present invention;

FIG. **6** is a flow chart showing the method of FIG. **4** in a second embodiment of the present invention; and

FIG. **7** is a flow chart showing the method of FIG. **4** in a third embodiment of the present invention.

DETAILED DESCRIPTION

Computer system power consumption is now rapidly approaching the point at which the power required to operate the system can no longer be reliably supplied to the system. This is primarily due to the fact that one or more processors contained in a system plugged into a power outlet may consume more power during an overactive period than the power outlet can deliver. Under this situation, the fuse or circuit breaker that protects the power outlet can be tripped (or blown) during normal operation of the system.

In accordance with an embodiment of the present invention, the power consumed by at least a portion of a computer system is monitored by measuring a parameter that is approximately proportional to the consumed power, such as voltage, current, or the duty cycle of a switching signal in a power supply. These measurements are provided to a power controller. The portion of the computer system that is monitored may include one or more processors of the computer system in addition to other integrated circuits (ICs) that consume a significant amount of power such as, for example, the bridge (or "chipset") or the video terminal.

Once the consumed power, as determined by the power controller, reaches a threshold, a throttle signal is sent to one or more ICs of the computer system by the controller. In response to receiving this throttle signal, one or more of the ICs reduces its power consumption. For example, for one embodiment of the present invention, a processor of the computer system reduces its core frequency while maintaining a consistent bus frequency in response to receiving the throttle signal. For another embodiment, the processor stalls all or a portion of one or more pipelines, or issues no-ops to one or more pipelines, in response to receiving the throttle signal.

By directly monitoring power consumption, or a value proportional thereto, temperature measurement inaccuracies related to poor positioning of thermal sensors in proximity to the processor are avoided. In addition, the response time between detecting a high power condition and reducing the power consumption of the computer system is greatly improved over the use of thermal sensors to detect the high power condition. A more detailed description of embodiments of the present invention, including various configurations and implementations, is provided below.

FIG. **1** is a multiprocessor computer system formed in accordance with an embodiment of the present invention. Primary bridge **110** is coupled to processors **100** and **101** via

INTEL - 1010

BACKGROUND

To address the thermal issue, elaborate thermal dissipation systems are often affixed to the processor to help dissipate the heat from the processor to the ambient environment. Some processor packages include a thermal sensor to monitor the temperature of the processor. If the processor temperature exceeds a particular threshold, the processor is placed into low power mode until it cools off. If these precautions are not taken, the processor may destroy itself by its own heat.

Ex. 1010 (Kling; US 6,367,023) at 1:38-46.

DETAILED DESCRIPTION

By directly monitoring power consumption, or a value proportional thereto, temperature measurement inaccuracies related to poor positioning of thermal sensors in proximity to the processor are avoided. In addition, the response time between detecting a high power condition and reducing the power consumption of the computer system is greatly improved over the use of thermal sensors to detect the high power condition. A more detailed description of embodiments of the present invention, including various configurations and implementations, is provided below.

Ex. 1010 (Kling) at 2:55-64; Paper 20 (POR) at 14, 56-57.

PACT EX. 2028

Kling

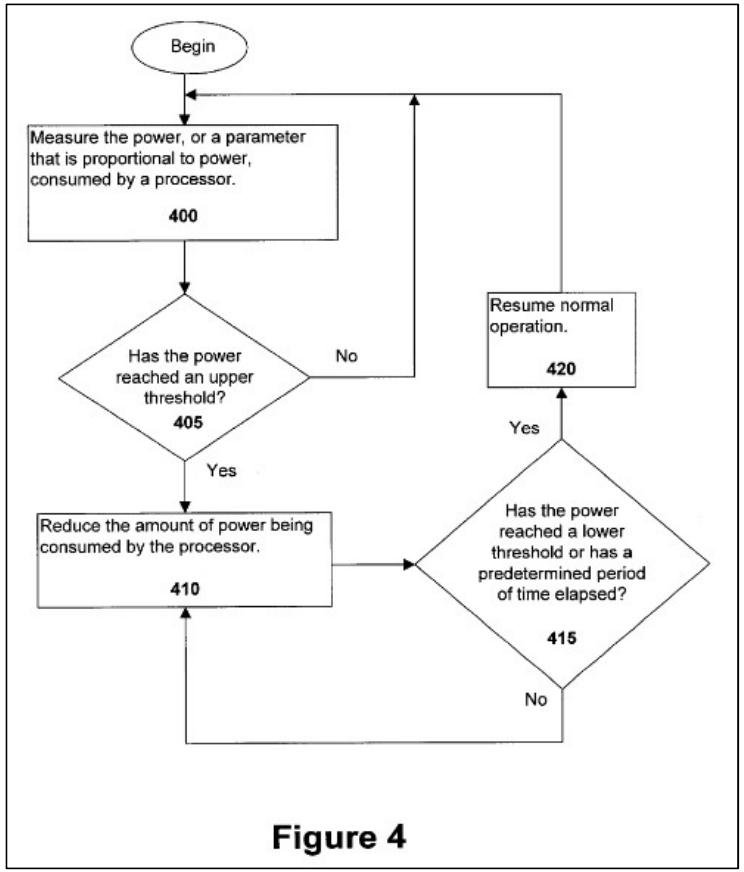

Figure 4

Ex. 1010 (Kling) at Fig. 4.

FIG. **4** is a flow chart showing a method of reducing power consumption of a processor in accordance with an embodiment of the present invention. At step **400**, the power, or a parameter that is proportional to power, consumed by the processor is measured. This power may be measured by a meter and the value provided in either digital or analog form to a power controller. For example, an ammeter may provide a measurement of the current consumed by the processor to the power controller. Next, at step **405**, the power controller determines if the power has reached an upper threshold. If the power has not reached the threshold, normal operation continues at step **400**. If, however, the power has reached the threshold, then, at step **410**, the amount of power being consumed by the processor is reduced. This reduction may be in response to a throttle signal sent from the power controller to the processor.

Ex. 1010 (Kling) at 7:39-54.

PACT EX. 2028

Kling

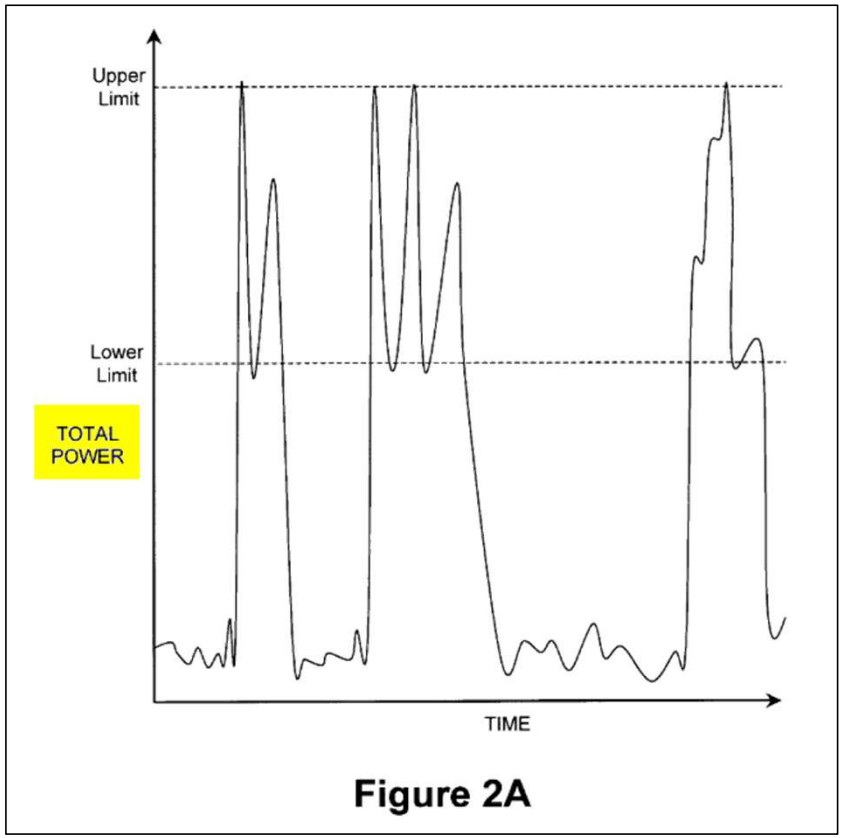

Figure 2A

Ex. 1010 (Kling) at Fig. 2A.

Once the consumed power, as determined by the power controller, reaches a threshold, a throttle signal is sent to one or more ICs of the computer system by the controller. In response to receiving this throttle signal, one or more of the ICs reduces its power consumption. For example, for one embodiment of the present invention, a processor of the computer system reduces its core frequency while maintaining a consistent bus frequency in response to receiving the throttle signal. For another embodiment, the processor stalls all or a portion of one or more pipelines, or issues no-ops to one or more pipelines, in response to receiving the throttle signal.

Ex. 1010 (Kling) at 2:42-55.

PACT EX. 2028

Nicol and Kling take different approaches

US006141762A

United States Patent [19]

Nicol et al.

[11] **Patent Number: 6,141,762**

[45] **Date of Patent: Oct. 31, 2000**

Hence, the operating system of PE **100** needs to ascertain the required completion time, divide the collection of tasks

as evenly as possible (in terms of needed processing time), consider the PE with the tasks that require the most time to carry out, and adjust the clock frequency to insure that the most heavily loaded PE carries out its assigned tasks within the required completion time. Once the frequency is thus determined, a minimum supply voltage can be determined.

Ex. 1009 (Nicol) at 3:66-4:5; Paper 20 (POR) at 10, 44.

US006367023B2

(12) **United States Patent**

Kling et al.

(10) **Patent No.: US 6,367,023 B2**

(45) **Date of Patent: Apr. 2, 2002**

power supply, and a controller. The meter takes a measurement of a parameter that is approximately proportional to the power consumed from the power supply by at least a portion of a computer system. This parameter may include, for example, the current or voltage between the power supply and the portion of the computer system, or the duty cycle of a switching signal in the power supply. This measurement is then used by the controller to determine if the power consumed by the portion of the computer system reaches a threshold, and if so, the controller sends a throttle signal. A

Ex. 1010 (Kling), Abstract; Paper 20 (POR) at 45.

PACT EX. 2028

Petitioner's motivation to combine Nicol and Kling

> *Third*, Nicol is focused on minimizing power consumption and does not discuss over-temperature protection. Because over-temperature protection is critical for high-performance systems such as Nicol's multiprocessor system and SoC, a POSITA seeking to improve Nicol would have been motivated to consult Kling, which discusses well-known techniques of providing over-temperature protection, as well as additional improvements. Ex. 1001 ¶127. One such improvement

Paper 2 (Petition) at 24.

Petitioner's motivation to combine Nicol and Kling

as well as additional improvements. Ex. 1001 ¶127. One such improvement described in Kling uses simple power measurement, such as measuring current, to detect excessive power consumption (when it is above a threshold) and use that information to "throttle" back processors to reduce power consumption. Ex. 1010, 5:38–59. A POSITA would have been motivated to implement this functionality in Nicol's system for over-temperature protection and, more generally, to detect excessive power consumption and reduce it, which is consistent with Nicol's objective of minimizing power consumption. Ex. 1009, 2:18–22; Ex. 1001 ¶129.

Paper 2 (Petition) at 24-25.

Petitioner's motivation to combine Nicol and Kling

For example, the Petition explains that Nicol does not discuss over-temperature protection and thus "a POSITA seeking to improve Nicol would have been motivated to consult Kling, which discusses well-known techniques of providing over-temperature protection, as well as additional improvements." Pet., 24-25. In particular, a POSITA would have been motivated to use temperature measurements both to improve voltage scaling (taught in Nicol) and for over-temperature protection (discussed in Kling). Pet., 42-43; Ex. 1055, ¶7. Likewise, a POSITA would have been motivated to use power measurements both for over-temperature protection (detailed in Kling) and to detect excessive power consumption consistent with Nicol's objective of minimizing power consumption. Pet., 25; Ex. 1055, ¶7.

Paper 27 (Reply) at 2-3.

Reasonable Expectation of Success

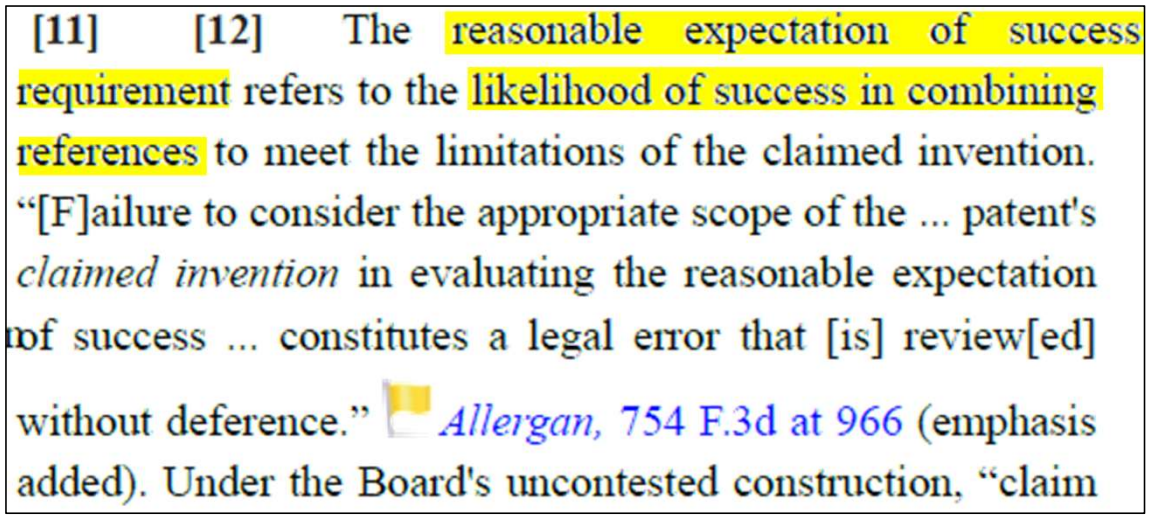

[11] [12] The reasonable expectation of success requirement refers to the likelihood of success in combining references to meet the limitations of the claimed invention. "[F]ailure to consider the appropriate scope of the ... patent's *claimed invention* in evaluating the reasonable expectation of success ... constitutes a legal error that [is] review[ed] without deference." *Allergan*, 754 F.3d at 966 (emphasis added). Under the Board's uncontested construction, "claim

Intelligent Bio-Systems, Inc. v. Illumina Cambridge Ltd., 821 F.3d 1359, 1367 (Fed. Cir. 2016).
Paper 31 (Sur-Reply) at 11.

Prior Art Needs to be Enabled

To render a claim obvious, the prior art, taken as a whole, must enable a skilled artisan to make and use the claimed invention. *See* *In re Kumar*, 418 F.3d 1361, 1368 (Fed. Cir. 2005). In general, a prior art reference asserted under § 103 does not necessarily have to enable its own disclosure, i.e., be "self-enabling," to be relevant to the obviousness inquiry. *See* *Symbol Techs., Inc. v. Opticon, Inc.*, 935 F.2d 1569, 1578 (Fed. Cir. 1991) ("While a reference must enable someone to practice the invention in order to anticipate under § 102(b), a non-enabling reference may qualify as prior art for the purpose of determining obviousness under § 103.");

But even though a non-enabling reference can play a role in an obviousness analysis, the evidence of record must still establish that a skilled artisan could have made the claimed invention. As the Sixth Circuit aptly explained:

Raytheon Techs. Corp. v. General Electric Co., 2021 WL 1432964, at *4 (Fed. Cir. 2021); Paper 31 (Sur-Reply) at 14.

PACT EX. 2028

Combining Nicol with Kling is difficult and could occupy an entire career

(c) Incorporating Kling's upper power limit would not be a simple matter as Petitioner asserts

75. In my opinion, the idea that a POSITA would easily incorporate the invention of Kling into Nicol via simple add-on software is untrue. Balancing upper limits with Nicol's power management system would have been a complex task in which a POSITA would have to balance the power needs of individual processors to accomplish tasks within a designated time along with temperature consideration which is also an element of the claim. A POSITA may spend their entire career developing such a system that responds appropriately and efficiently to these variables.

Ex. 2024 (Mangione-Smith POR Decl.) at 59, ¶75.

Nicol does not disclose “sequentially processing data”

PACT EX. 2028

U.S. Patent 9,075,605 B2
Independent Claim 1

1. A method for operating a multiprocessor system comprising a plurality of data processing cells, each of the plurality of data processing units (a) including at least one Arithmetic Logic Unit (ALU) and a register unit and (b) being adapted for sequentially processing data, the method comprising:

the multiprocessor system setting a clock frequency, of at least a part of the multiprocessor system to a minimum in accordance with a number of pending operations of a first processor;

the multiprocessor system subsequently increasing the clock frequency of the at least the part of the multiprocessor system to a maximum in accordance with a number of pending operations of a second processor; and

the multiprocessor system subsequently reducing the clock frequency of the at least the part of the multiprocessor system in accordance with (a) an operating temperature threshold preventing over-temperature and (b) a hysteresis characteristic.

Ex. 1003 (US 9,075,605) at 15:6-16:5.

Claim Construction

Data processing unit and/or **data processing units . . . adapted for sequentially processing data** '047 Patent, Claims 1-5, 8, 10, 15, 19, 22, 24 '605 Patent, Claim 1	Sequential data processors that pass results to one or more additional data processing units to perform additional computations

Ex. 1054 (Claim Construction Order) at 2.

Nicol

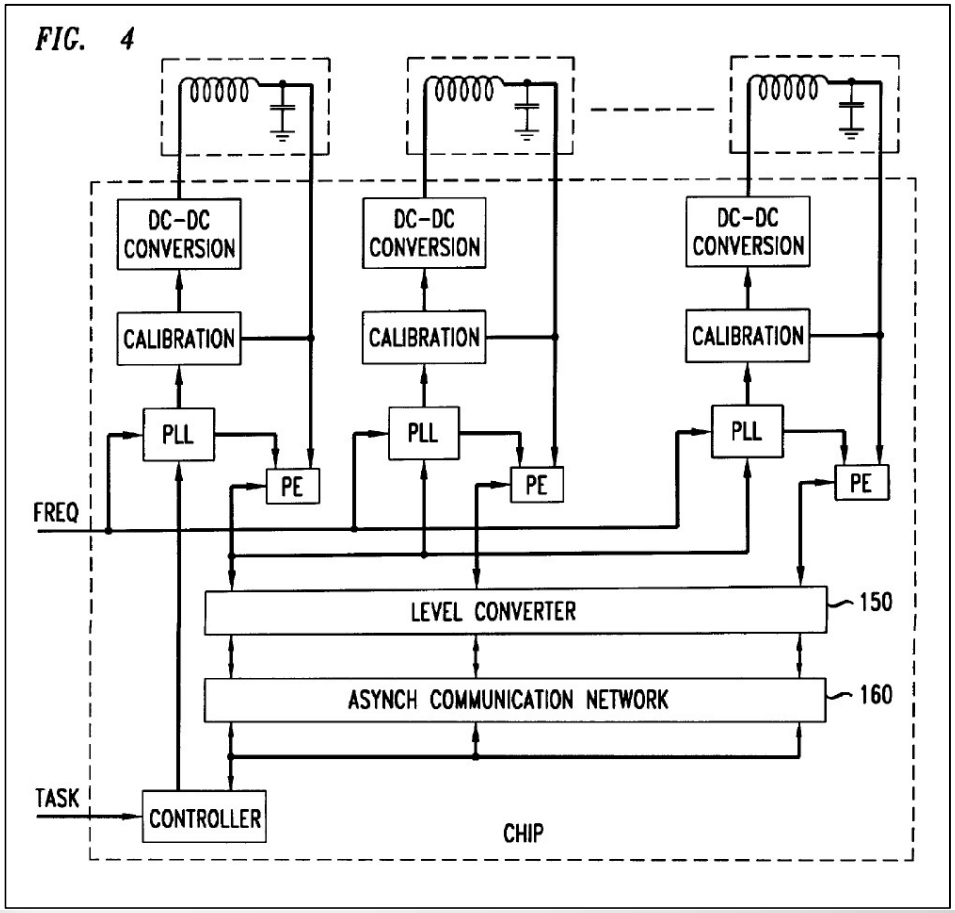

It may be noted that if the frequencies at which the individual PEs operate differ from one another and from other elements within the system where the multiprocessor chip is employed, there is an issue of synchronization that must be addressed. That is, a synchronization schema must be implemented when there is a need to communicate data between PEs (or with other system elements) that operate at different frequencies. It is possible to arrange the frequencies so that the collection of tasks that are assigned to the multiprocessor is completed at a predetermined time. In such a case, the synchronization problem of the multiprocessor vis-a-vis other elements within the system where the multiprocessor is employed is minimized. However, that leaves the issue of synchronizing the exchange of data among the PEs of a multiprocessor chip.

To effect such synchronization, each PE within the FIG. **4** arrangement is connection to an arrangement comprising elements **150** and **160**. Level converter **150** converts the variable voltage swings of the PEs to a fixed level swing, and network **160** resolves the issue of different clock domains.

Ex. 1009 (Nicol) at 6:14-33.

PACT EX. 2028

Mangione-Smith Depo. Tr. (1/15/21)

A. Yeah, I don't recall if I said that they were all communicating amongst each other. The level shifter is there to do voltage equalization between the PEs 101 through 104 and the PE 100. There is -- it's clear to me from the written description that the application for figure 2 anticipates communications back to the PE 100 from the various computation PEs.

Ex. 1052 (Mangione-Smith 1/15/21 Depo.) at 46:20-47:3.

In the case of communications among the PEs, there is a whole world of scheduling aspects, scheduling issues, that would have to be addressed because one task is now dependent upon another. None of which are discussed in Nicol.

Ex. 1052 (Mangione-Smith 1/15/21 Depo.) at 47:15-20.

PACT EX. 2028

Mangione-Smith Depo. Tr. (1/15/21)

A. Yeah. It's not clear to me that figure 4 discloses direct communication among the PEs. It may be intended to do that. It does disclose -- direct communication among PEs of that sort was generally known. One of the interesting things here is the differences in clock speed, and the synchronizer would address that.

Ex. 1052 (Mangione-Smith 1/15/21 Depo.) at 109:22-110:5.

So a person of ordinary skill looking at figure 4 would understand -- well, figure 4 and the spec, would understand that, you know, Nicol is intending for the PEs to talk to one another but that there might be additional sort of changes required in order to completely allow that; right?

A. I wouldn't --

Q. Did I get that right?

A. Yeah, I wouldn't quite phrase it that way. I think a person of ordinary skill in the art would see that Nicol anticipates that figure 4 could be used to build a system where the PEs communicate and envisions that as a possibility. I don't know that it actually discloses that on its own.

Ex. 1052 (Mangione-Smith 1/15/21 Depo.) at 110:10-111:2.

PACT EX. 2028

Petitioner has not Established the Disclosure of "Temperature Threshold" (Claim 1)

U.S. Patent 9,075,605 B2
Independent Claim 1 ("temperature threshold")

1. A method for operating a multiprocessor system comprising a plurality of data processing cells, each of the plurality of data processing units (a) including at least one Arithmetic Logic Unit (ALU) and a register unit and (b) being adapted for sequentially processing data, the method comprising:

the multiprocessor system setting a clock frequency, of at least a part of the multiprocessor system to a minimum in accordance with a number of pending operations of a first processor;

the multiprocessor system subsequently increasing the clock frequency of the at least the part of the multiprocessor system to a maximum in accordance with a number of pending operations of a second processor; and

the multiprocessor system subsequently reducing the clock frequency of the at least the part of the multiprocessor system in accordance with (a) an operating temperature threshold preventing over-temperature and (b) a hysteresis characteristic.

Ex. 1003 (US 9,075,605) at 15:7-16:5.

PACT EX. 2028

Petitioner relies on Kling for the “temperature threshold”

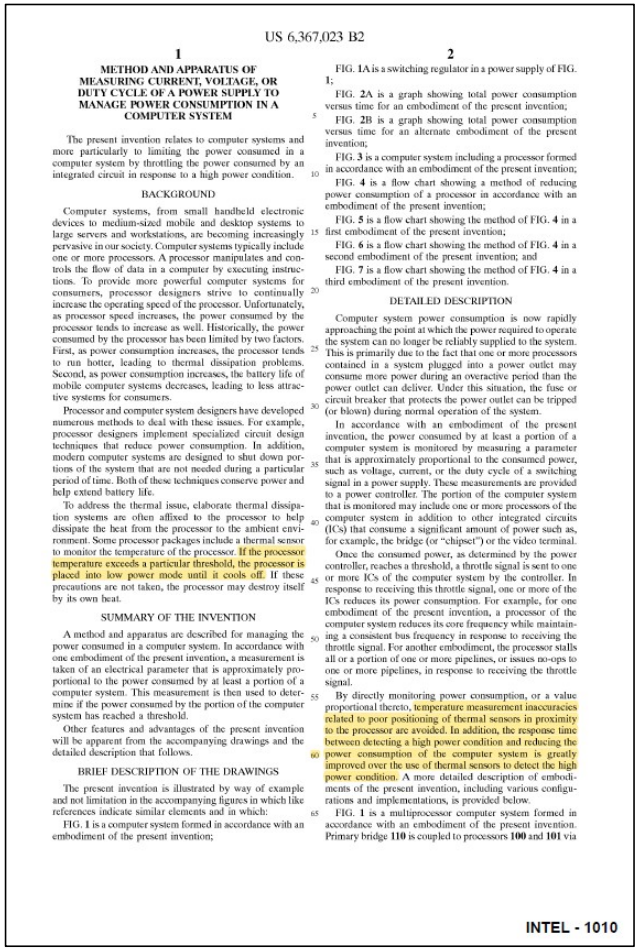

US 6,367,023 B2

1

METHOD AND APPARATUS OF MEASURING CURRENT, VOLTAGE, OR DUTY CYCLE OF A POWER SUPPLY TO MANAGE POWER CONSUMPTION IN A COMPUTER SYSTEM

The present invention relates to computer systems and more particularly to limiting the power consumed in a computer system by throttling the power consumed by an integrated circuit in response to a high power condition.

BACKGROUND

Computer systems, from small handheld electronic devices to medium-sized mobile and desktop systems to large servers and workstations, are becoming increasingly pervasive in our society. Computer systems typically include one or more processors. A processor manipulates and controls the flow of data in a computer by executing instructions. To provide more powerful computer systems for consumers, processor designers strive to continually increase the operating speed of the processor. Unfortunately, as processor speed increases, the power consumed by the processor tends to increase as well. Historically, the power consumed by the processor has been limited by two factors. First, as power consumption increases, the processor tends to run hotter, leading to thermal dissipation problems. Second, as power consumption increases, the battery life of mobile computer systems decreases, leading to less attractive systems for consumers.

Processor and computer system designers have developed numerous methods to deal with these issues. For example, processor designers implement specialized circuit design techniques that reduce power consumption. In addition, modern computer systems are designed to shut down portions of the system that are not needed during a particular period of time. Both of these techniques conserve power and help extend battery life.

To address the thermal issue, elaborate thermal dissipation systems are often affixed to the processor to help dissipate the heat from the processor to the ambient environment. Some processor packages include a thermal sensor to monitor the temperature of the processor. If the processor temperature exceeds a particular threshold, the processor is placed into low power mode until it cools off. If these precautions are not taken, the processor may destroy itself by its own heat.

SUMMARY OF THE INVENTION

A method and apparatus are described for managing the power consumed in a computer system. In accordance with one embodiment of the present invention, a measurement is taken of an electrical parameter that is approximately proportional to the power consumed by at least a portion of a computer system. This measurement is then used to determine if the power consumed by the portion of the computer system has reached a threshold.

Other features and advantages of the present invention will be apparent from the accompanying drawings and the detailed description that follows.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example and not limitation in the accompanying figures in which like references indicate similar elements and in which:

FIG. **1** is a computer system formed in accordance with an embodiment of the present invention;

2

FIG. 1A is a switching regulator in a power supply of FIG. 1;

FIG. **2A** is a graph showing total power consumption versus time for an embodiment of the present invention;

FIG. **2B** is a graph showing total power consumption versus time for an alternate embodiment of the present invention;

FIG. **3** is a computer system including a processor formed in accordance with an embodiment of the present invention;

FIG. **4** is a flow chart showing a method of reducing power consumption of a processor in accordance with an embodiment of the present invention;

FIG. **5** is a flow chart showing the method of FIG. **4** in a first embodiment of the present invention;

FIG. **6** is a flow chart showing the method of FIG. **4** in a second embodiment of the present invention; and

FIG. **7** is a flow chart showing the method of FIG. **4** in a third embodiment of the present invention.

DETAILED DESCRIPTION

Computer system power consumption is now rapidly approaching the point at which the power required to operate the system can no longer be reliably supplied to the system. This is primarily due to the fact that one or more processors contained in a system plugged into a power outlet may consume more power during an overactive period than the power outlet can deliver. Under this situation, the fuse or circuit breaker that protects the power outlet can be tripped (or blown) during normal operation of the system.

In accordance with an embodiment of the present invention, the power consumed by at least a portion of a computer system is monitored by measuring a parameter that is approximately proportional to the consumed power, such as voltage, current, or the duty cycle of a switching signal in a power supply. These measurements are provided to a power controller. The portion of the computer system that is monitored may include one or more processors of the computer system in addition to other integrated circuits (ICs) that consume a significant amount of power such as, for example, the bridge (or “chipset”) or the video terminal.

Once the consumed power, as determined by the power controller, reaches a threshold, a throttle signal is sent to one or more ICs of the computer system by the controller. In response to receiving this throttle signal, one or more of the ICs reduces its power consumption. For example, for one embodiment of the present invention, a processor of the computer system reduces its core frequency while maintaining a consistent bus frequency in response to receiving the throttle signal. For another embodiment, the processor stalls all or a portion of one or more pipelines, or issues no-ops to one or more pipelines, in response to receiving the throttle signal.

By directly monitoring power consumption, or a value proportional thereto, temperature measurement inaccuracies related to poor positioning of thermal sensors in proximity to the processor are avoided. In addition, the response time between detecting a high power condition and reducing the power consumption of the computer system is greatly improved over the use of thermal sensors to detect the high power condition. A more detailed description of embodiments of the present invention, including various configurations and implementations, is provided below.

FIG. **1** is a multiprocessor computer system formed in accordance with an embodiment of the present invention. Primary bridge **110** is coupled to processors **100** and **101** via

INTEL - 1010

BACKGROUND

To address the thermal issue, elaborate thermal dissipation systems are often affixed to the processor to help dissipate the heat from the processor to the ambient environment. Some processor packages include a thermal sensor to monitor the temperature of the processor. If the processor temperature exceeds a particular threshold, the processor is placed into low power mode until it cools off. If these precautions are not taken, the processor may destroy itself by its own heat.

Ex. 1010 (Kling) at 1:38-46.

DETAILED DESCRIPTION

By directly monitoring power consumption, or a value proportional thereto, temperature measurement inaccuracies related to poor positioning of thermal sensors in proximity to the processor are avoided. In addition, the response time between detecting a high power condition and reducing the power consumption of the computer system is greatly improved over the use of thermal sensors to detect the high power condition. A more detailed description of embodiments of the present invention, including various configurations and implementations, is provided below.

Ex. 1010 (Kling) at 2:55-64; Paper 20 (POR) at 14, 56-57.

PACT EX. 2028

Petitioner relies on Kling for the "temperature threshold"

Kling explains that processing systems must be able to rapidly detect and mitigate over-temperature situations. Ex. 1010, 1:38–46 ("To address the thermal issue [in modern computer systems] . . . processor packages include a thermal sensor to monitor the temperature of the processor. *If the processor temperature exceeds a particular threshold*, the processor is placed into low power mode until it cools off. If these precautions are not taken, the processor may destroy itself by its own heat.").

Paper 2 (Petition) at 42.

Petitioner has not Established the Disclosure of "Hysteresis" (Claim 1)

U.S. Patent 9,075,605 B2
Independent Claim 1 ("hysteresis")

1. A method for operating a multiprocessor system comprising a plurality of data processing cells, each of the plurality of data processing units (a) including at least one Arithmetic Logic Unit (ALU) and a register unit and (b) being adapted for sequentially processing data, the method comprising:

the multiprocessor system setting a clock frequency, of at least a part of the multiprocessor system to a minimum in accordance with a number of pending operations of a first processor;

the multiprocessor system subsequently increasing the clock frequency of the at least the part of the multiprocessor system to a maximum in accordance with a number of pending operations of a second processor; and

the multiprocessor system subsequently reducing the clock frequency of the at least the part of the multiprocessor system in accordance with (a) an operating temperature threshold preventing over-temperature and (b) a hysteresis characteristic.

Ex. 1003 (US 9,075,605) at 15:7-16:5.

Petitioner's interpretation of "hysteresis"

value. Another example is the manner in which a thermostat utilizes the right amount of hysteresis to maintain the desired temperature. Too little hysteresis would result in constant switching between cycles, and too much would allow too large a temperature interval between cycles. 4. In an os-

Ex. 1053 (Wiley) at 355; Paper 27 (Reply) at 22.

PACT EX. 2028

Petitioner's argument on "hysteresis"

Ex. 1053 (Wiley), 355 ("[A] thermostat utilizes the right amount of hysteresis to maintain the desired temperature. Too little hysteresis would result in constant switching between cycles, and too much would allow too large a temperature interval between cycles.")). This admission is fatal to PO's strained position. Just like thermostat control and as taught in Kling, an over-temperature protection system requires the use of hysteresis to avoid constantly switching between normal operation mode (*e.g.*, high frequency/voltage for maximum performance) and an over-temperature protection mode (*e.g.*, reduced frequency/voltage to "cool off" an over-heating PE). Pet., 44-45; Ex. 1055, ¶37; Ex. 1010, 5:45-54; Ex. 1021, 471.

Paper 27 (Reply) at 22.

Petitioner points to Kling's power oscillation

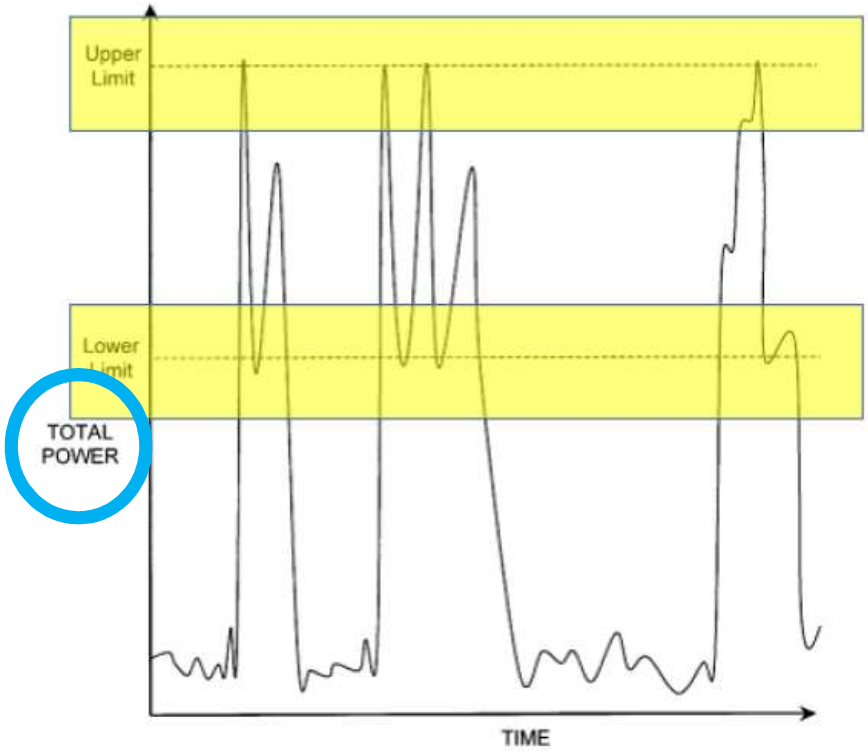

Figure 2A

Paper 2 (Petition) at 45.

tion. In accordance with this embodiment, the lower threshold may be set to a predetermined value selected to provide hysteresis to reduce the occurrence of power oscillation between the upper and lower thresholds. This lower thresh-

Ex. 1010 (Kling) at 5:51-54; Paper 2 (Petition) at 44.

Kling also discusses the use of hysteresis to improve robustness and "reduce the occurrence of power oscillation," as would also have been well-known in the art.

Paper 2 (Petition) at 44.

PACT EX. 2028

Nicol does not have power oscillation

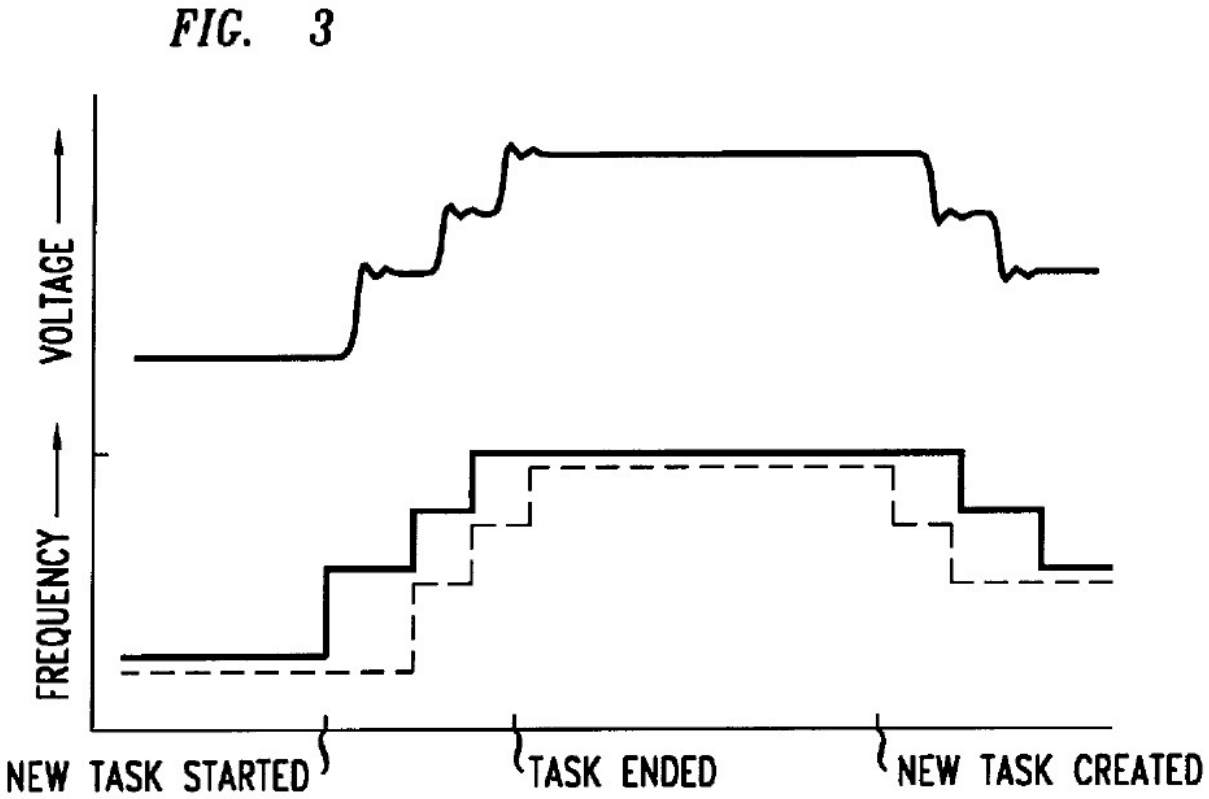

Ex. 1009 (Nicol) at Fig. 3; Paper 20 (POR) at 11, 41-42.

Kennedy's "hysteresis" is about "noise rejection and oscillator circuits"

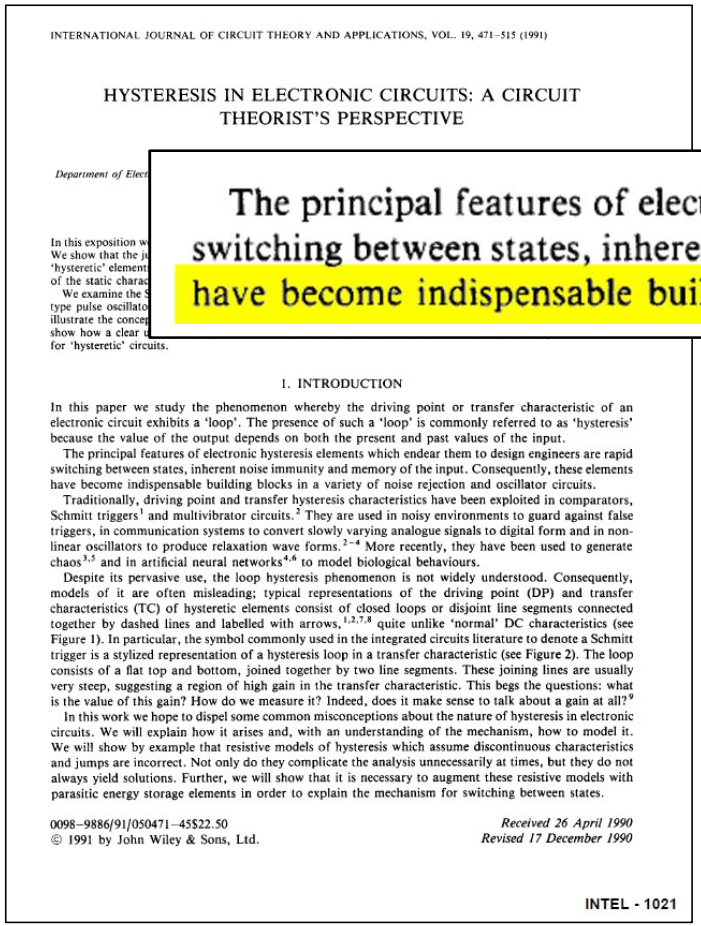

INTERNATIONAL JOURNAL OF CIRCUIT THEORY AND APPLICATIONS, VOL. 19, 471–515 (1991)

HYSTERESIS IN ELECTRONIC CIRCUITS: A CIRCUIT THEORIST'S PERSPECTIVE

Department of Elect

In this exposition w
We show that the ju
'hysteretic' element
of the static charac
We examine the S
type pulse oscillato
illustrate the concep
show how a clear u
for 'hysteretic' circuits.

1. INTRODUCTION

In this paper we study the phenomenon whereby the driving point or transfer characteristic of an electronic circuit exhibits a 'loop'. The presence of such a 'loop' is commonly referred to as 'hysteresis' because the value of the output depends on both the present and past values of the input.

The principal features of electronic hysteresis elements which endear them to design engineers are rapid switching between states, inherent noise immunity and memory of the input. Consequently, these elements have become indispensable building blocks in a variety of noise rejection and oscillator circuits.

Traditionally, driving point and transfer hysteresis characteristics have been exploited in comparators, Schmitt triggers[1] and multivibrator circuits.[2] They are used in noisy environments to guard against false triggers, in communication systems to convert slowly varying analogue signals to digital form and in non-linear oscillators to produce relaxation wave forms.[2–4] More recently, they have been used to generate chaos[3,5] and in artificial neural networks[4,6] to model biological behaviours.

Despite its pervasive use, the loop hysteresis phenomenon is not widely understood. Consequently, models of it are often misleading; typical representations of the driving point (DP) and transfer characteristics (TC) of hysteretic elements consist of closed loops or disjoint line segments connected together by dashed lines and labelled with arrows,[1,2,7,8] quite unlike 'normal' DC characteristics (see Figure 1). In particular, the symbol commonly used in the integrated circuits literature to denote a Schmitt trigger is a stylized representation of a hysteresis loop in a transfer characteristic (see Figure 2). The loop consists of a flat top and bottom, joined together by two line segments. These joining lines are usually very steep, suggesting a region of high gain in the transfer characteristic. This begs the questions: what is the value of this gain? How do we measure it? Indeed, does it make sense to talk about a gain at all?[9]

In this work we hope to dispel some common misconceptions about the nature of hysteresis in electronic circuits. We will explain how it arises and, with an understanding of the mechanism, how to model it. We will show by example that resistive models of hysteresis which assume discontinuous characteristics and jumps are incorrect. Not only do they complicate the analysis unnecessarily at times, but they do not always yield solutions. Further, we will show that it is necessary to augment these resistive models with parasitic energy storage elements in order to explain the mechanism for switching between states.

0098–9886/91/050471–45$22.50
© 1991 by John Wiley & Sons, Ltd.

Received 26 April 1990
Revised 17 December 1990

INTEL - 1021

The principal features of electronic hysteresis elements which endear them to design engineers are rapid switching between states, inherent noise immunity and memory of the input. Consequently, these elements have become indispensable building blocks in a variety of noise rejection and oscillator circuits.

Ex. 1021 (Kennedy) at 471; Paper 20 (POR) at 43.

1021 (Kennedy, M. & Chua, L., *Hysteresis in Electronic Circuits: A Circuit Theorist's Perspective,* International Journal of Circuit Theory and Applications, Vol. 19, (1991)) at 471 (explaining that circuits with hysteresis characteristics were "indispensable" and commonly used "in noisy environments and to guard against false triggers"),.

Paper 20 (POR) at 43.

PACT EX. 2028

Nicol uses an alternate method to prevent noise

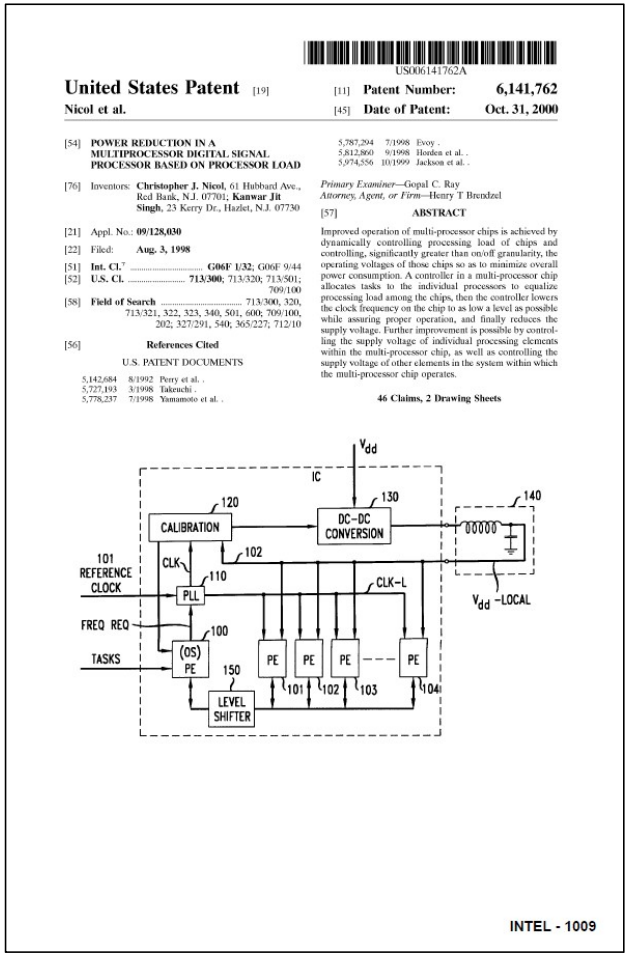

US006141762A

United States Patent [19]
Nicol et al.

[11] **Patent Number: 6,141,762**
[45] **Date of Patent: Oct. 31, 2000**

[54] **POWER REDUCTION IN A MULTIPROCESSOR DIGITAL SIGNAL PROCESSOR BASED ON PROCESSOR LOAD**

[76] Inventors: **Christopher J. Nicol**, 61 Hubbard Ave., Red Bank, N.J. 07701; **Kanwar Jit Singh**, 23 Kerry Dr., Hazlet, N.J. 07730

[21] Appl. No.: **09/128,030**

[22] Filed: **Aug. 3, 1998**

[51] **Int. Cl.**[7] **G06F 1/32**; G06F 9/44

[52] **U.S. Cl.** **713/300**; 713/320; 713/501; 709/100

[58] **Field of Search** 713/300, 320, 713/321, 322, 323, 340, 501, 600; 709/100, 202; 327/291, 540; 365/227; 712/10

[56] **References Cited**

U.S. PATENT DOCUMENTS

5,142,684 8/1992 Perry et al. .
5,727,193 3/1998 Takeuchi .
5,778,237 7/1998 Yamamoto et al. .
5,787,294 7/1998 Evoy .
5,812,860 9/1998 Horden et al. .
5,974,556 10/1999 Jackson et al. .

Primary Examiner—Gopal C. Ray
Attorney, Agent, or Firm—Henry T Brendzel

[57] **ABSTRACT**

Improved operation of multi-processor chips is achieved by dynamically controlling processing load of chips and controlling, significantly greater than on/off granularity, the operating voltages of those chips so as to minimize overall power consumption. A controller in a multi-processor chip allocates tasks to the individual processors to equalize processing load among the chips, then the controller lowers the clock frequency on the chip to as low a level as possible while assuring proper operation, and finally reduces the supply voltage. Further improvement is possible by controlling the supply voltage of individual processing elements within the multi-processor chip, as well as controlling the supply voltage of other elements in the system within which the multi-processor chip operates.

46 Claims, 2 Drawing Sheets

INTEL - 1009

> It is noted that the frequency is gradually programmed into the system (as shown by the stepped changes in FIG. **3**). This prevents excessive noise on the V_{dd}-local supply voltage and possible circuit failure. For example, if the system

Ex. 1009 (Nicol) at 5:29-31, Fig. 3; Paper 20 (POR) at 43.

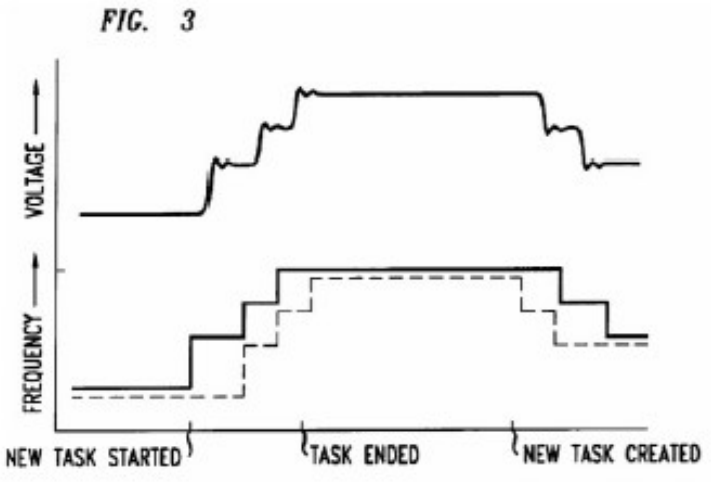

PACT EX. 2028

Petitioner has not Established the Motivation to Combine for the Method Claim

Claim 1(c)

1. A method for operating a multiprocessor system comprising a plurality of data processing cells, each of the plurality of data processing units (a) including at least one Arithmetic Logic Unit (ALU) and a register unit and (b) being adapted for sequentially processing data, the method comprising:

the multiprocessor system setting a clock frequency, of at least a part of the multiprocessor system to a minimum in accordance with a number of pending operations of a first processor;

the multiprocessor system subsequently increasing the clock frequency of the at least the part of the multiprocessor system to a maximum in accordance with a number of pending operations of a second processor; and

the multiprocessor system subsequently reducing the clock frequency of the at least the part of the multiprocessor system in accordance with (a) an operating temperature threshold preventing over-temperature and (b) a hysteresis characteristic.

Ex. 1003 (US 9,075,605) at 15:6-16:5.

Kling is about power threshold, not temperature threshold

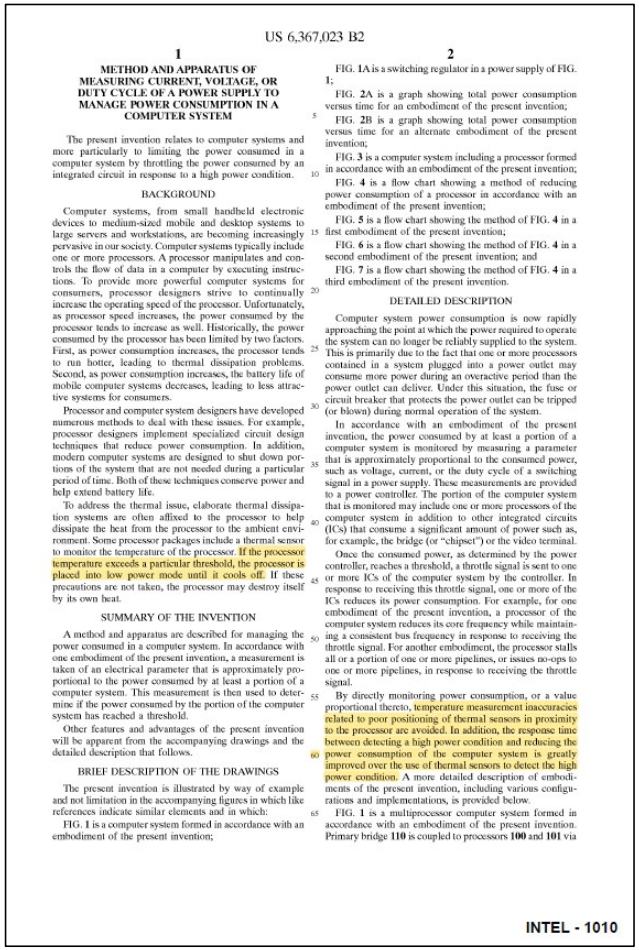

US 6,367,023 B2

1

METHOD AND APPARATUS OF MEASURING CURRENT, VOLTAGE, OR DUTY CYCLE OF A POWER SUPPLY TO MANAGE POWER CONSUMPTION IN A COMPUTER SYSTEM

The present invention relates to computer systems and more particularly to limiting the power consumed in a computer system by throttling the power consumed by an integrated circuit in response to a high power condition.

BACKGROUND

Computer systems, from small handheld electronic devices to medium-sized mobile and desktop systems to large servers and workstations, are becoming increasingly pervasive in our society. Computer systems typically include one or more processors. A processor manipulates and controls the flow of data in a computer by executing instructions. To provide more powerful computer systems for consumers, processor designers strive to continually increase the operating speed of the processor. Unfortunately, as processor speed increases, the power consumed by the processor tends to increase as well. Historically, the power consumed by the processor has been limited by two factors. First, as power consumption increases, the processor tends to run hotter, leading to thermal dissipation problems. Second, as power consumption increases, the battery life of mobile computer systems decreases, leading to less attractive systems for consumers.

Processor and computer system designers have developed numerous methods to deal with these issues. For example, processor designers implement specialized circuit design techniques that reduce power consumption. In addition, modern computer systems are designed to shut down portions of the system that are not needed during a particular period of time. Both of these techniques conserve power and help extend battery life.

To address the thermal issue, elaborate thermal dissipation systems are often affixed to the processor to help dissipate the heat from the processor to the ambient environment. Some processor packages include a thermal sensor to monitor the temperature of the processor. If the processor temperature exceeds a particular threshold, the processor is placed into low power mode until it cools off. If these precautions are not taken, the processor may destroy itself by its own heat.

SUMMARY OF THE INVENTION

A method and apparatus are described for managing the power consumed in a computer system. In accordance with one embodiment of the present invention, a measurement is taken of an electrical parameter that is approximately proportional to the power consumed by at least a portion of a computer system. This measurement is then used to determine if the power consumed by the portion of the computer system has reached a threshold.

Other features and advantages of the present invention will be apparent from the accompanying drawings and the detailed description that follows.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example and not limitation in the accompanying figures in which like references indicate similar elements and in which:

FIG. **1** is a computer system formed in accordance with an embodiment of the present invention;

2

FIG. 1A is a switching regulator in a power supply of FIG. 1;

FIG. **2A** is a graph showing total power consumption versus time for an embodiment of the present invention;

FIG. **2B** is a graph showing total power consumption versus time for an alternate embodiment of the present invention;

FIG. **3** is a computer system including a processor formed in accordance with an embodiment of the present invention;

FIG. **4** is a flow chart showing a method of reducing power consumption of a processor in accordance with an embodiment of the present invention;

FIG. **5** is a flow chart showing the method of FIG. **4** in a first embodiment of the present invention;

FIG. **6** is a flow chart showing the method of FIG. **4** in a second embodiment of the present invention; and

FIG. **7** is a flow chart showing the method of FIG. **4** in a third embodiment of the present invention.

DETAILED DESCRIPTION

Computer system power consumption is now rapidly approaching the point at which the power required to operate the system can no longer be reliably supplied to the system. This is primarily due to the fact that one or more processors contained in a system plugged into a power outlet may consume more power during an overactive period than the power outlet can deliver. Under this situation, the fuse or circuit breaker that protects the power outlet can be tripped (or blown) during normal operation of the system.

In accordance with an embodiment of the present invention, the power consumed by at least a portion of a computer system is monitored by measuring a parameter that is approximately proportional to the consumed power, such as voltage, current, or the duty cycle of a switching signal in a power supply. These measurements are provided to a power controller. The portion of the computer system that is monitored may include one or more processors of the computer system in addition to other integrated circuits (ICs) that consume a significant amount of power such as, for example, the bridge (or "chipset") or the video terminal.

Once the consumed power, as determined by the power controller, reaches a threshold, a throttle signal is sent to one or more ICs of the computer system by the controller. In response to receiving this throttle signal, one or more of the ICs reduces its power consumption. For example, for one embodiment of the present invention, a processor of the computer system reduces its core frequency while maintaining a consistent bus frequency in response to receiving the throttle signal. For another embodiment, the processor stalls all or a portion of one or more pipelines, or issues no-ops to one or more pipelines, in response to receiving the throttle signal.

By directly monitoring power consumption, or a value proportional thereto, temperature measurement inaccuracies related to poor positioning of thermal sensors in proximity to the processor are avoided. In addition, the response time between detecting a high power condition and reducing the power consumption of the computer system is greatly improved over the use of thermal sensors to detect the high power condition. A more detailed description of embodiments of the present invention, including various configurations and implementations, is provided below.

FIG. **1** is a multiprocessor computer system formed in accordance with an embodiment of the present invention. Primary bridge **110** is coupled to processors **100** and **101** via

INTEL - 1010

> By directly monitoring power consumption, or a value proportional thereto, temperature measurement inaccuracies related to poor positioning of thermal sensors in proximity to the processor are avoided. In addition, the response time between detecting a high power condition and reducing the power consumption of the computer system is greatly improved over the use of thermal sensors to detect the high power condition. A more detailed description of embodiments of the present invention, including various configurations and implementations, is provided below.

Ex. 1010 (Kling) at 2:55-64; Paper 20 (POR) at 14, 56-57.

PACT EX. 2028

Claim 1(c) - hysteresis

Kling also discusses the use of hysteresis to improve robustness and "reduce the occurrence of power oscillation." as would also have been well-known in the art.

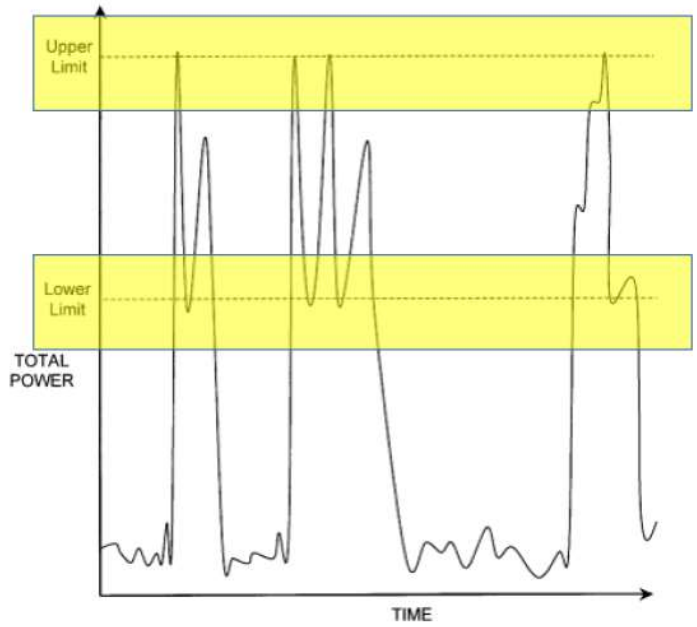

Figure 2A

Paper 2 (Petition) at 44-45.

PACT EX. 2028

Petitioner argues temperature threshold reached before power threshold

spots" that power measurements alone could miss. Ex. 1055, ¶10. This is because the same amount of power consumed by the system (*e.g.*, a chip) may not be high enough to trigger a power measurement trip point, but still could cause a local temperature to rise to an unsafe level if most of the power is consumed by one small circuit of the chip, creating a dangerous "hot spot." *Id.*, ¶11.

Paper 27 (Reply) at 5.

Petitioner argues temperature threshold reached before power threshold

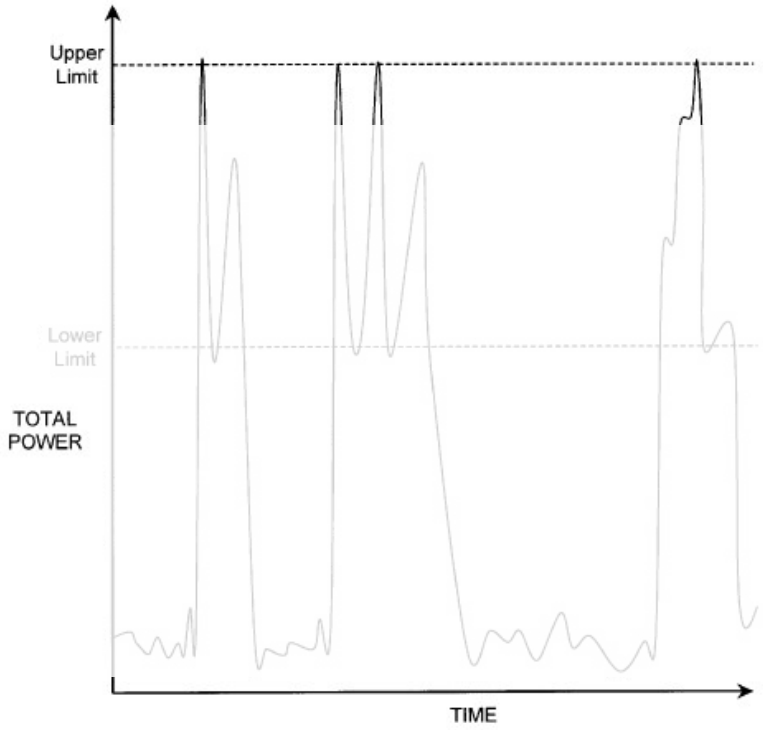

Figure 2A

Ex. 1010 (Kling) at Fig. 2A.

"according to a hysteresis characteristic"

"operating temperature threshold preventing over-temperature"

PACT EX. 2028

Method Claims Require More Than "Being Capable"

United States Court of Appeals for the Federal Circuit

"While [the petitioner] was only required to identify a prior art reference that discloses an apparatus '**capable of**' performing the recited functions to prove that the **apparatus claims** would have been obvious, **more was required with respect to the method claims**. Specifically, [the petitioner] needed to present evidence and argument that a person of ordinary skill **would have been motivated to operate [the prior art] in a manner that satisfied the [claim] limitation**."

ParkerVision, Inc. v. Qualcomm Incorporated, 903 F.3d 1354, 1363 (Fed. Cir. 2018).

PACT EX. 2028

Method Claims Require More Than "Being Capable"

"[T]o show the obviousness of the **claimed method**, Petitioner **cannot merely show an apparatus capable of performing the method** would have been obvious but **must present** evidence and argument that a person of ordinary skill would have been **motivated to operate the prior art apparatus in a manner that satisfied the limitations of the challenged claims**."

United Industries Corporation et al v. Susan McKnight, Inc. et al., Paper 31 (Case IPR2017-01687, PTAB, Jan. 22, 2019) at 10, 27 (citing *ParkerVision, Inc. v. Qualcomm Incorporated*, 903 F.3d 1354 (Fed. Cir. 2018)).

PACT EX. 2028

Petitioner Fails to Establish a Reasonable Expectation of Achieving What is Claimed

Claim 1

1. A method for operating a multiprocessor system comprising a plurality of data processing cells, each of the plurality of data processing units (a) including at least one Arithmetic Logic Unit (ALU) and a register unit and (b) being adapted for sequentially processing data, the method comprising:

the multiprocessor system setting a clock frequency, of at least a part of the multiprocessor system to a minimum in accordance with a number of pending operations of a first processor;

the multiprocessor system subsequently increasing the clock frequency of the at least the part of the multiprocessor system to a maximum in accordance with a number of pending operations of a second processor; and

Ex. 1003 (US 9,075,605) at 15:6-20.

the multiprocessor system subsequently reducing the clock frequency of the at least the part of the multiprocessor system in accordance with (a) an operating temperature threshold preventing over-temperature and (b) a hysteresis characteristic.

Ex. 1003 (US 9,075,605) at 16:1-5.

PACT EX. 2028

Nicol

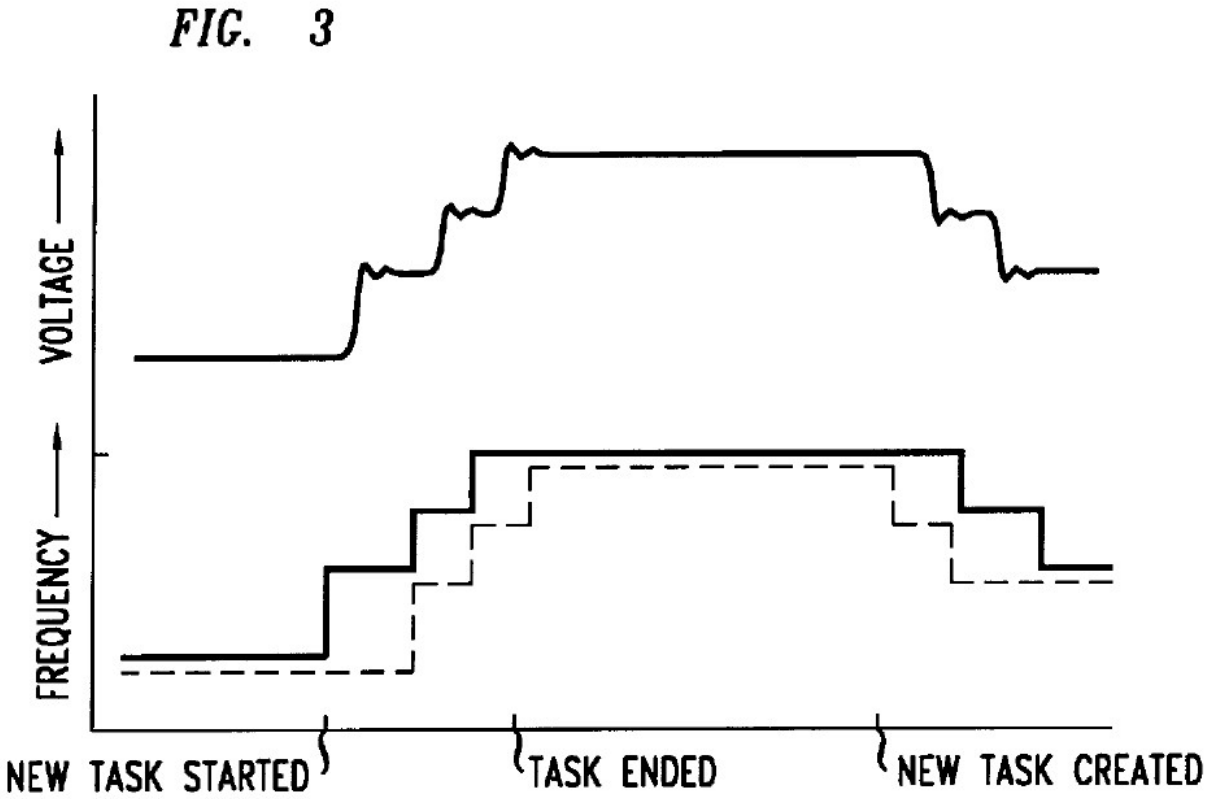

Ex. 1009 (Nicol) at Fig. 3.

Kling's voltage threshold

Pet., 45-46. Kling teaches reducing clock frequency when a voltage representing the power consumed by a processor reaches a threshold. *Id.*; Ex. 1010, 2:31-51 ("the

Paper 27 (Reply) at 22.

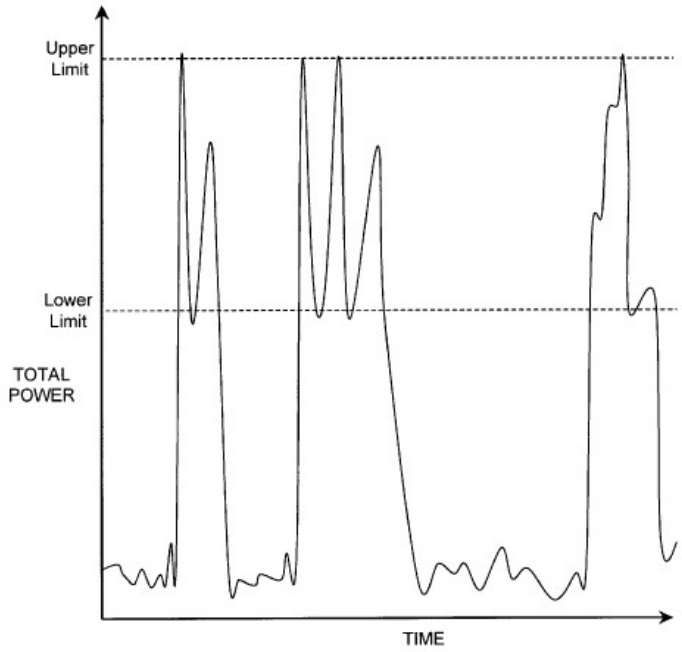

Kling's voltage threshold

Ex. 1010 (Kling) at Fig. 2A.

Case 1: Nicol's voltage < Kling's voltage threshold

The alleged "maximum" will not trigger Kling's threshold

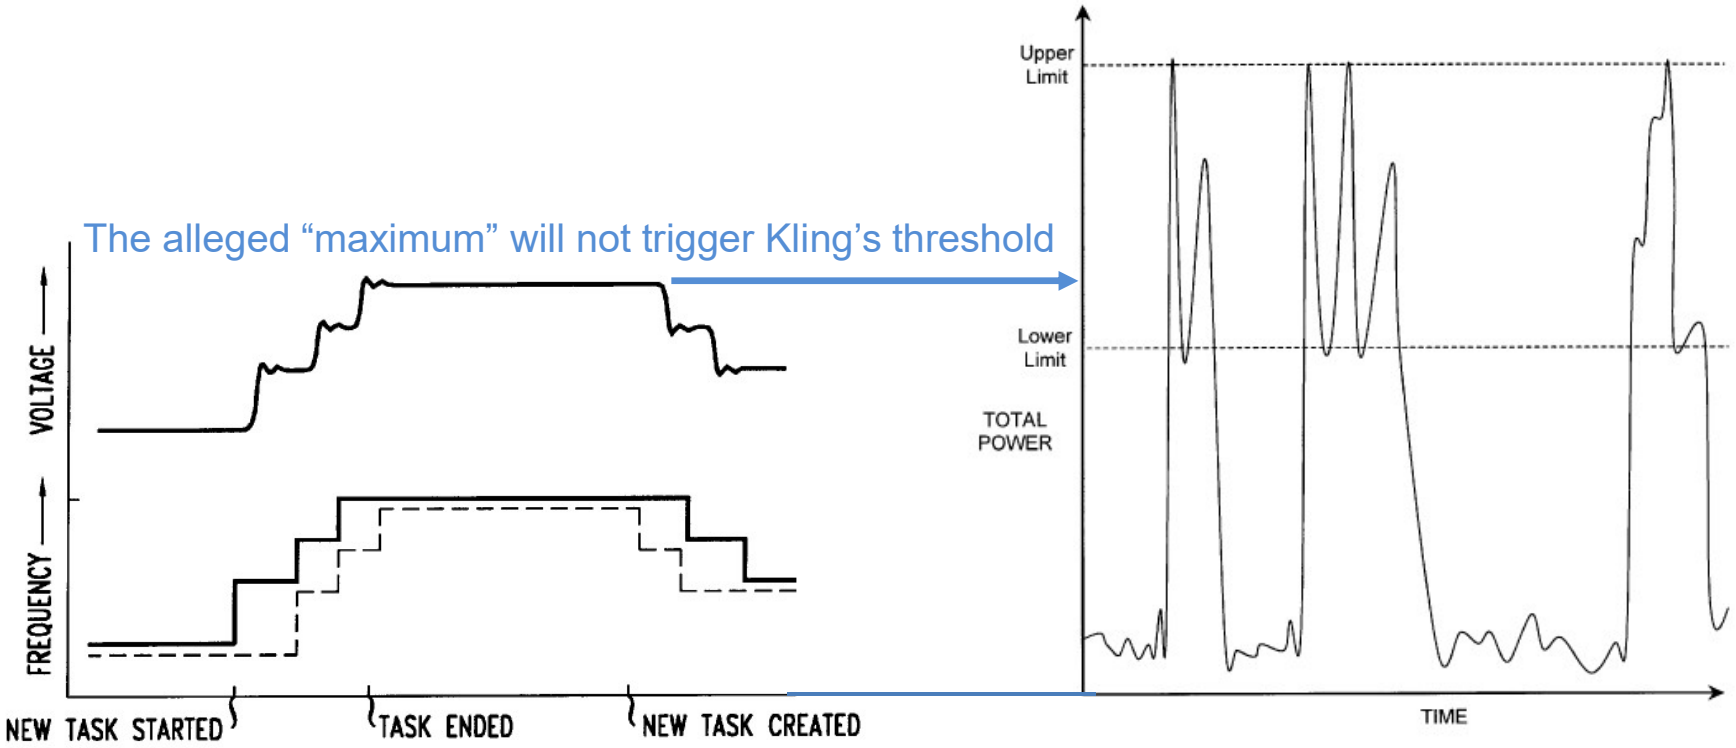

Ex. 1009 (Nicol) at Fig. 3.

Ex. 1010 (Kling) at Fig. 2A.

Nicol

Kling

PACT EX. 2028

Case 2: Nicol's voltage > Kling's voltage threshold

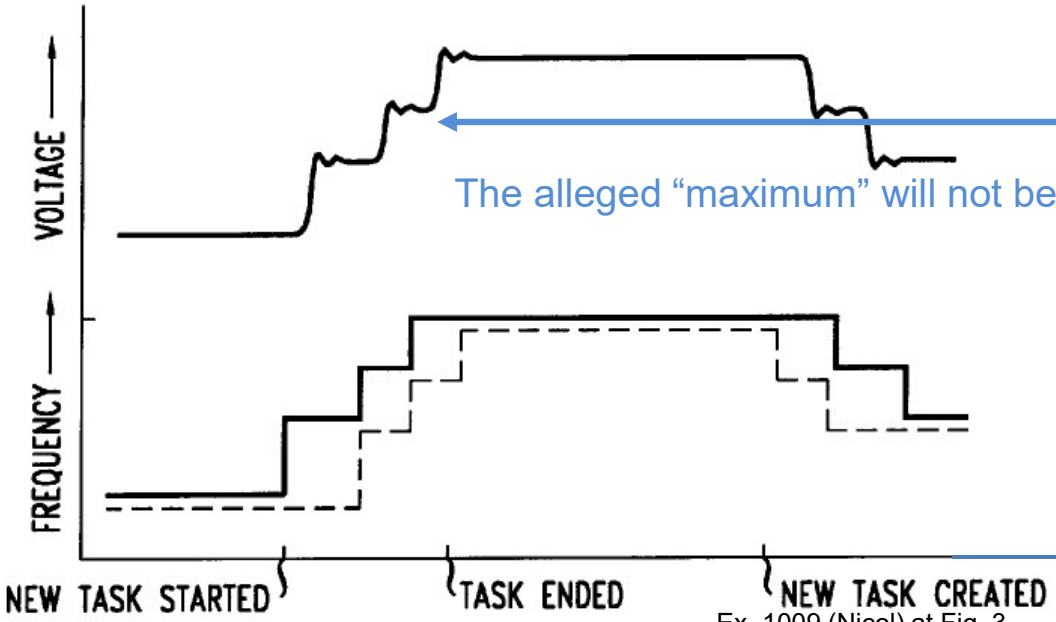

The alleged "maximum" will not be achieved

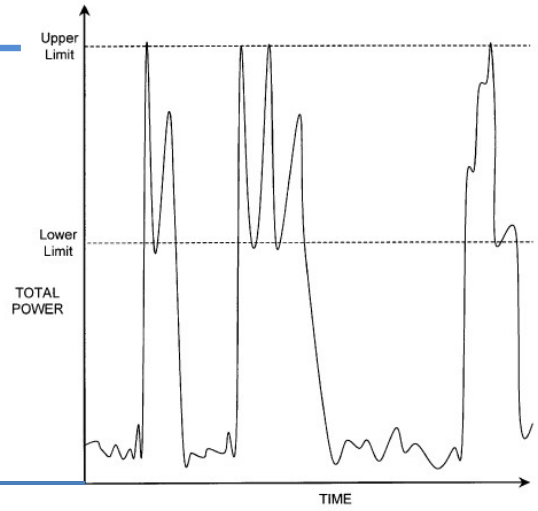

Ex. 1009 (Nicol) at Fig. 3.

Ex. 1010 (Kling) at Fig. 2A.

Nicol

Kling

PACT EX. 2028

Case 3: Nicol's voltage = Kling's voltage threshold

The voltage will be reduced immediately after it reaches Kling's upper limit; thus, the temperature threshold will not be reached.

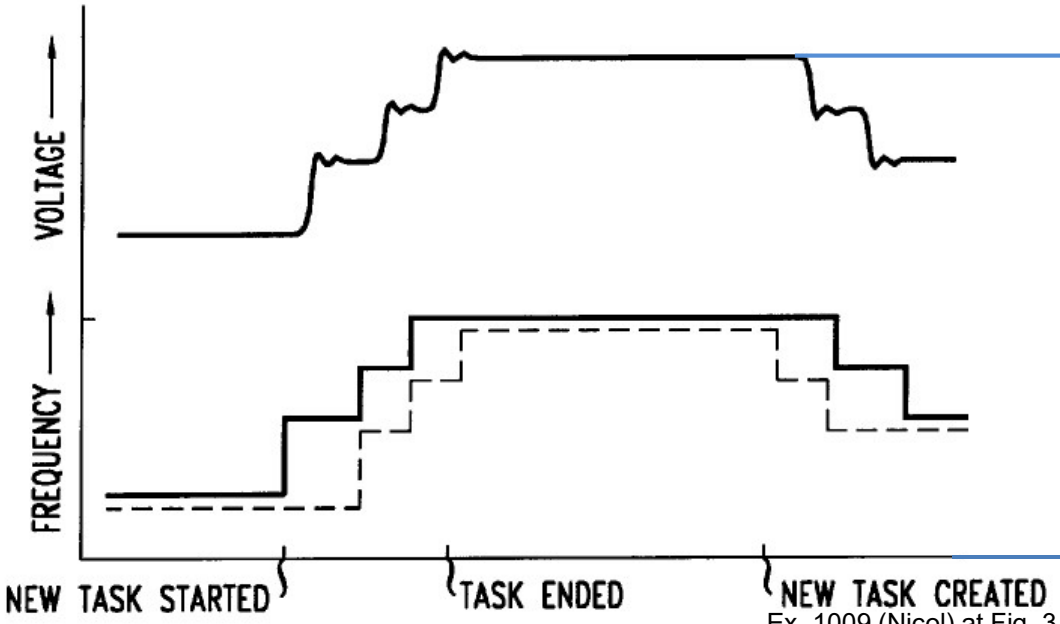

Ex. 1009 (Nicol) at Fig. 3.

Nicol

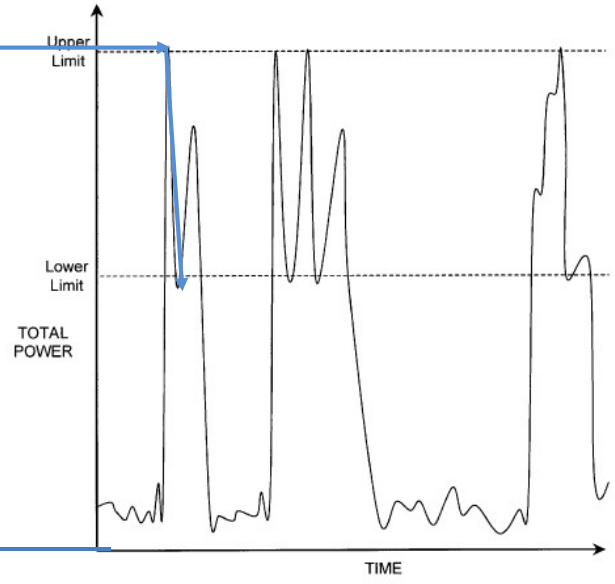

Ex. 1010 (Kling) at Fig. 2A.

Kling

PACT EX. 2028

Kling's response time

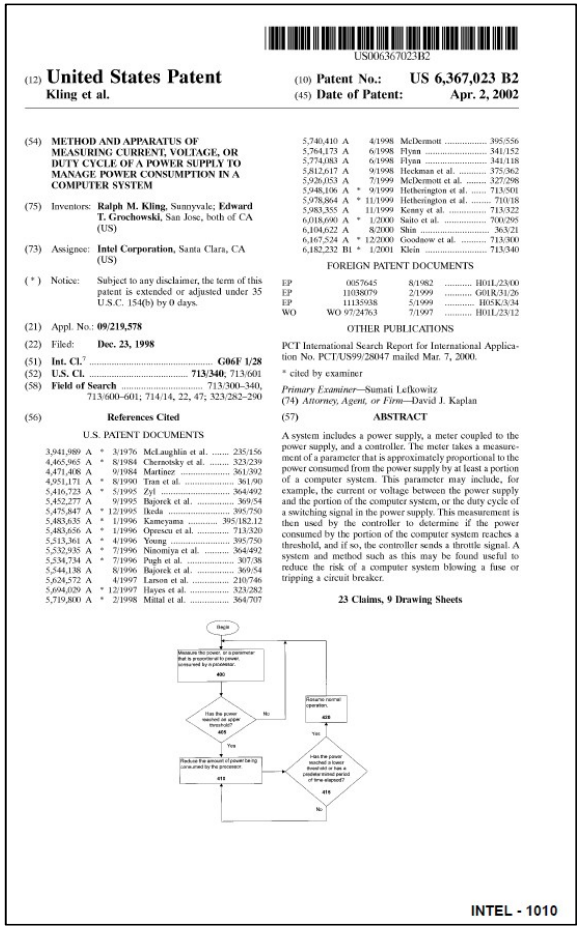

US006367023B2

(12) **United States Patent**
Kling et al.

(10) **Patent No.: US 6,367,023 B2**
(45) **Date of Patent: Apr. 2, 2002**

(54) **METHOD AND APPARATUS OF MEASURING CURRENT, VOLTAGE, OR DUTY CYCLE OF A POWER SUPPLY TO MANAGE POWER CONSUMPTION IN A COMPUTER SYSTEM**

(75) Inventors: **Ralph M. Kling**, Sunnyvale; **Edward T. Grochowski**, San Jose, both of CA (US)

(73) Assignee: **Intel Corporation**, Santa Clara, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **09/219,578**

(22) Filed: **Dec. 23, 1998**

(51) **Int. Cl.**[7] **G06F 1/28**

(52) **U.S. Cl.** **713/340**; 713/601

(58) **Field of Search** 713/300–340, 713/600–601; 714/14, 22, 47; 323/282–290

(56) **References Cited**

U.S. PATENT DOCUMENTS

3,941,989	A *	3/1976	McLaughlin et al.	235/156
4,465,965	A *	8/1984	Chernotsky et al.	323/239
4,471,408	A	9/1984	Martinez	361/392
4,951,171	A *	8/1990	Tran et al.	361/90
5,416,723	A *	5/1995	Zyl	364/492
5,452,277	A	9/1995	Bajorek et al.	369/54
5,475,847	A *	12/1995	Ikeda	395/750
5,483,635	A *	1/1996	Kameyama	395/182.12
5,483,656	A *	1/1996	Oprescu et al.	713/320
5,513,361	A *	4/1996	Young	395/750
5,532,935	A *	7/1996	Ninomiya et al.	364/492
5,534,734	A *	7/1996	Pugh et al.	307/38
5,544,138	A	8/1996	Bajorek et al.	369/54
5,624,572	A	4/1997	Larson et al.	210/746
5,694,029	A *	12/1997	Hayes et al.	323/282
5,719,800	A *	2/1998	Mittal et al.	364/707
5,740,410	A	4/1998	McDermott	395/556
5,764,173	A	6/1998	Flynn	341/152
5,774,083	A	6/1998	Flynn	341/118
5,812,617	A	9/1998	Heckman et al.	375/362
5,926,053	A	7/1999	McDermott et al.	327/298
5,948,106	A *	9/1999	Hetherington et al.	713/501
5,978,864	A *	11/1999	Hetherington et al.	710/18
5,983,355	A	11/1999	Kenny et al.	713/322
6,018,690	A *	1/2000	Saito et al.	700/295
6,104,622	A	8/2000	Shin	363/21
6,167,524	A *	12/2000	Goodnow et al.	713/300
6,182,232	B1 *	1/2001	Klein	713/340

FOREIGN PATENT DOCUMENTS

EP	0057645	8/1982	H01L/23/00
EP	11038079	2/1999	G01R/31/26
EP	11135938	5/1999	H05K/3/34
WO	WO 97/24763	7/1997	H01L/23/12

OTHER PUBLICATIONS

PCT International Search Report for International Application No. PCT/US99/28047 mailed Mar. 7, 2000.

* cited by examiner

Primary Examiner—Sumati Lefkowitz
(74) *Attorney, Agent, or Firm*—David J. Kaplan

(57) **ABSTRACT**

A system includes a power supply, a meter coupled to the power supply, and a controller. The meter takes a measurement of a parameter that is approximately proportional to the power consumed from the power supply by at least a portion of a computer system. This parameter may include, for example, the current or voltage between the power supply and the portion of the computer system, or the duty cycle of a switching signal in the power supply. This measurement is then used by the controller to determine if the power consumed by the portion of the computer system reaches a threshold, and if so, the controller sends a throttle signal. A system and method such as this may be found useful to reduce the risk of a computer system blowing a fuse or tripping a circuit breaker.

23 Claims, 9 Drawing Sheets

INTEL - 1010

By directly monitoring power consumption, or a value proportional thereto, temperature measurement inaccuracies related to poor positioning of thermal sensors in proximity to the processor are avoided. In addition, the response time between detecting a high power condition and reducing the power consumption of the computer system is greatly improved over the use of thermal sensors to detect the high power condition. A more detailed description of embodiments of the present invention, including various configurations and implementations, is provided below.

Ex. 1010 (Kling) at 2:55-64; Paper 20 (POR) at 14, 56-57.

PACT EX. 2028

Nicol and Kling do not render claim 2 obvious

U.S. Patent 9,075,605 B2
Dependent Claim 2

1. A method for operating a multiprocessor system comprising a plurality of data processing cells, each of the plurality of data processing units (a) including at least one Arithmetic Logic Unit (ALU) and a register unit and (b) being adapted for sequentially processing data, the method comprising:

the multiprocessor system setting a clock frequency, of at least a part of the multiprocessor system to a minimum in accordance with a number of pending operations of a first processor;

the multiprocessor system subsequently increasing the clock frequency of the at least the part of the multiprocessor system to a maximum in accordance with a number of pending operations of a second processor; and

the multiprocessor system subsequently reducing the clock frequency of the at least the part of the multiprocessor system in accordance with (a) an operating temperature threshold preventing over-temperature and (b) a hysteresis characteristic.

Ex. 1003 (US 9,075,605) at 15:6-16:5.

2. The method of claim **1** wherein:
said clock frequency is determined in part by processor voltage.

Ex. 1003 (US 9,075,605) at 16:6-8.

PACT EX. 2028

Kling's voltage threshold

Pet., 45-46. Kling teaches reducing clock frequency when a voltage representing the power consumed by a processor reaches a threshold. *Id.*; Ex. 1010, 2:31-51 ("the

Paper 27 (Reply) at 22.

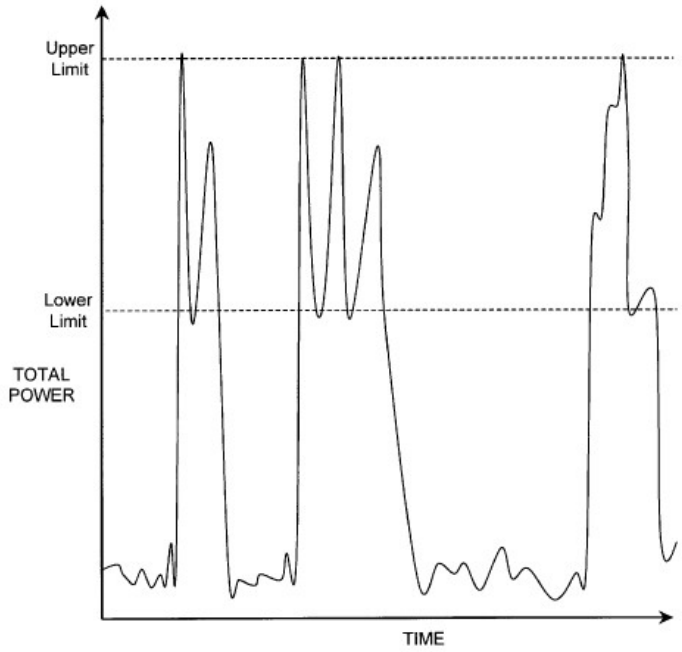

Kling's voltage threshold

Ex. 1010 (Kling) at Fig. 2A.

PACT EX. 2028

Nicol

US006141762A

United States Patent [19]

Nicol et al.

[11] **Patent Number:** 6,141,762

[45] **Date of Patent:** Oct. 31, 2000

Hence, the operating system of PE **100** needs to ascertain the required completion time, divide the collection of tasks

as evenly as possible (in terms of needed processing time), consider the PE with the tasks that require the most time to carry out, and adjust the clock frequency to insure that the most heavily loaded PE carries out its assigned tasks within the required completion time. Once the frequency is thus determined, a minimum supply voltage can be determined.

Ex. 1009 (Nicol) at 3:66-4:5; Paper 20 (POR) at 10, 44.

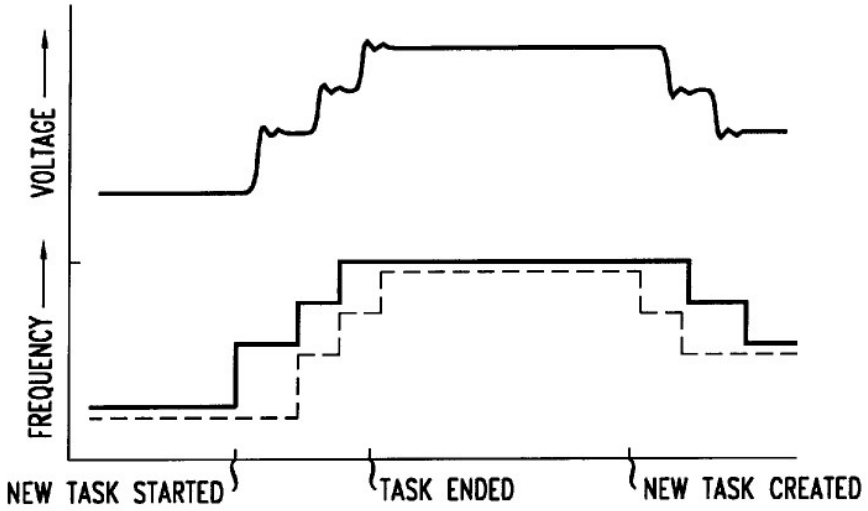

Ex. 1009 (Nicol) at Fig. 3.

PACT EX. 2028

Nicol and Kling do not render claim 3 obvious

PACT EX. 2028

U.S. Patent 9,075,605 B2
Dependent Claim 3

1. A method for operating a multiprocessor system comprising a plurality of data processing cells, each of the plurality of data processing units (a) including at least one Arithmetic Logic Unit (ALU) and a register unit and (b) being adapted for sequentially processing data, the method comprising:

the multiprocessor system setting a clock frequency, of at least a part of the multiprocessor system to a minimum in accordance with a number of pending operations of a first processor;

the multiprocessor system subsequently increasing the clock frequency of the at least the part of the multiprocessor system to a maximum in accordance with a number of pending operations of a second processor; and

the multiprocessor system subsequently reducing the clock frequency of the at least the part of the multiprocessor system in accordance with (a) an operating temperature threshold preventing over-temperature and (b) a hysteresis characteristic.

Ex. 1003 (US 9,075,605) at 15:6-16:5.

3. The method of claim 1 wherein:
said clock frequency is determined in part by available power.

Ex. 1003 (US 9,075,605) at 16:9-11.

Kling's throttle signal

Once the consumed power, as determined by the power controller, reaches a threshold, a throttle signal is sent to one or more ICs of the computer system by the controller. In response to receiving this throttle signal, one or more of the ICs reduces its power consumption. For example, for one embodiment of the present invention, a processor of the computer system reduces its core frequency while maintaining a consistent bus frequency in response to receiving the throttle signal. For another embodiment, the processor stalls all or a portion of one or more pipelines, or issues no-ops to one or more pipelines, in response to receiving the throttle signal.

Ex. 1010 (Kling) at 2:42-55.

discussed above in Section XI.A.1 and XI.A.3, a POSITA would have found it obvious to add Kling's power control functionality to Nicol's system. Ex. 1001 ¶199. Kling teaches measuring the power (based on the current or voltage across an appropriate resistor) of a processor and comparing the measured power against an "upper limit" threshold, as shown in Figure 2A (annotated) below. Ex. 1001, ¶200;

The upper limit in Figure 2A is the available power to the processor. Ex. 1010, 5:21–25 (explaining that "the upper limit is a constant value … associated with an approximately maximum power that can be reliably consumed by the computer system …"); Ex. 1001 ¶¶202–203. If the measured power reaches the upper limit, the controller sends a "throttle signal" to the processor to instruct it to reduce its clock frequency to reduce power consumption. Ex. 1010, 5:38–40, 2:48–

Paper 2 (Petition) at 47-48
citing Ex. 1010 (Kling) at 5:38-40, 2:48-51.

PACT EX. 2028

Case 1: Nicol's voltage < Kling's voltage threshold

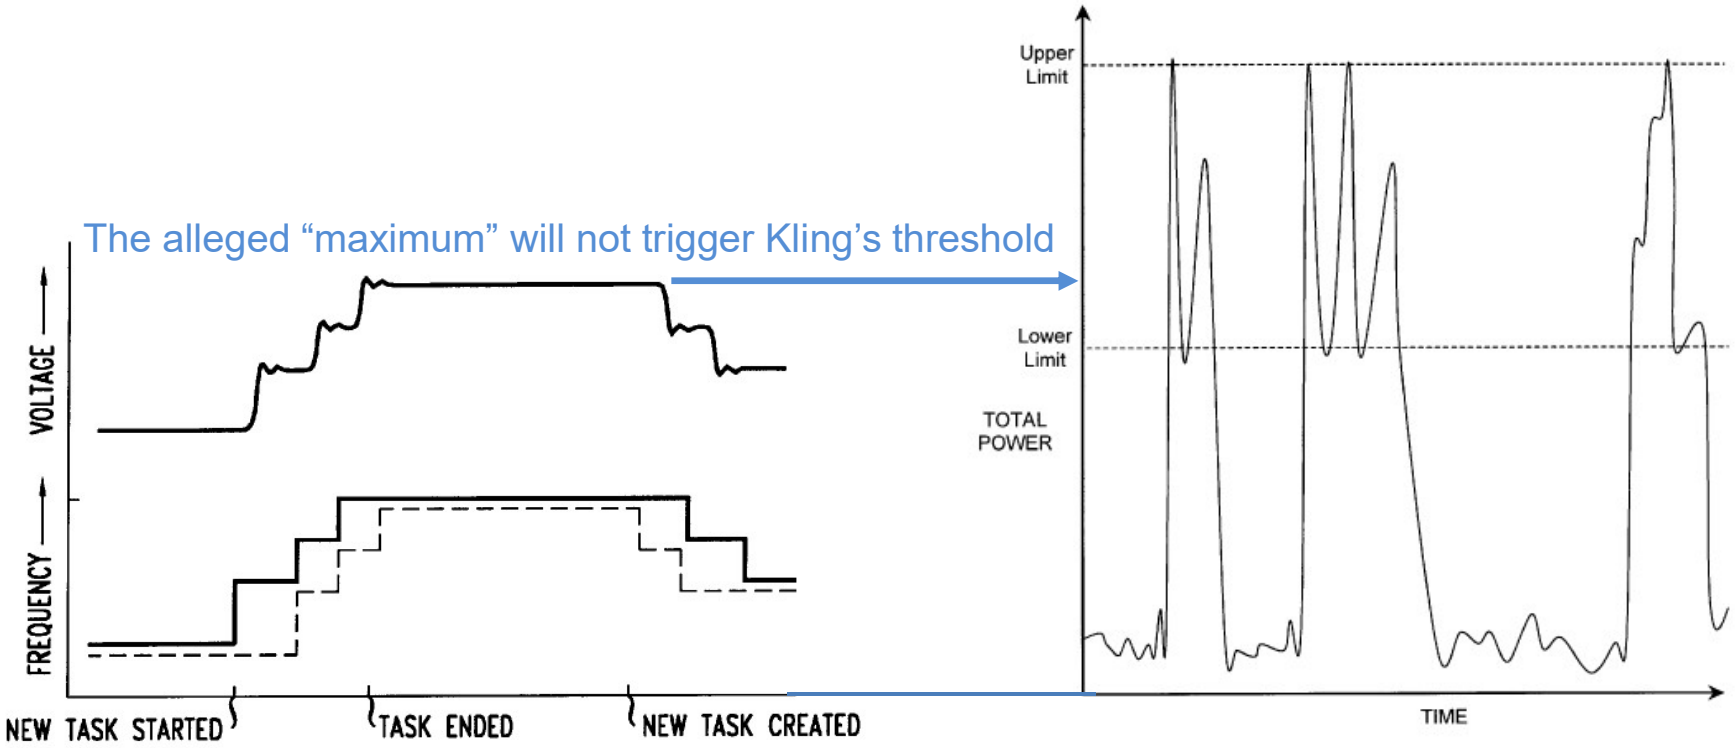

Ex. 1009 (Nicol) at Fig. 3.

Ex. 1010 (Kling) at Fig. 2A.

Nicol

Kling

PACT EX. 2028

Case 2: Nicol's voltage > Kling's voltage threshold

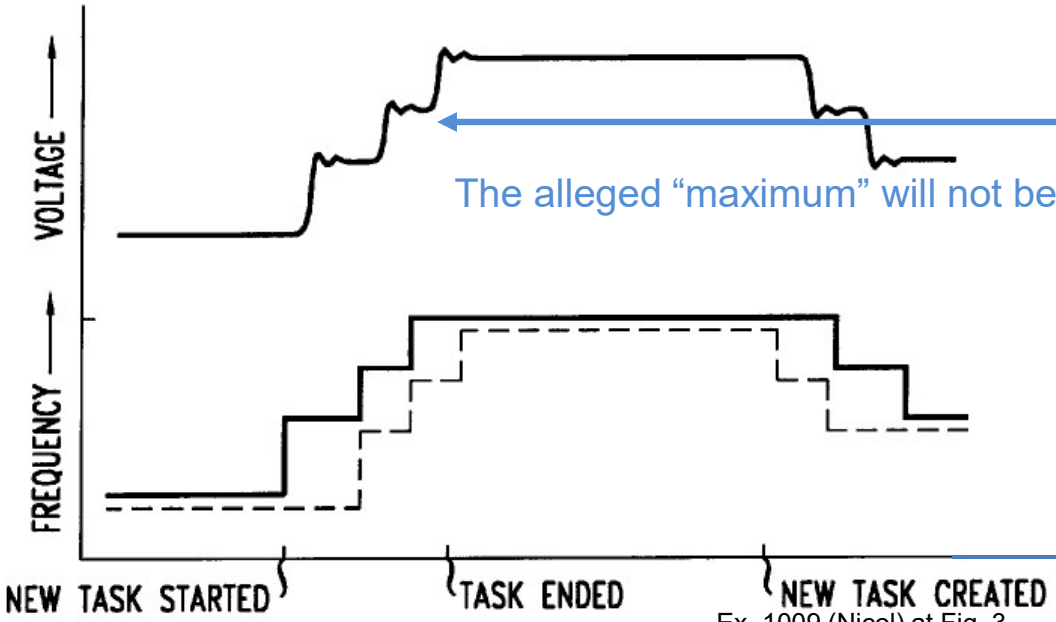

The alleged "maximum" will not be achieved

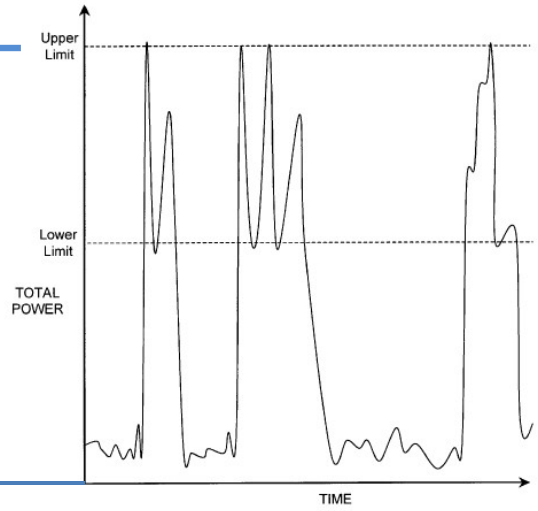

Ex. 1009 (Nicol) at Fig. 3.

Ex. 1010 (Kling) at Fig. 2A.

Nicol

Kling

PACT EX. 2028

Nicol does not teach that clock frequency is determined in part by available power

Hence, the operating system of PE **100** needs to ascertain the required completion time, divide the collection of tasks as evenly as possible (in terms of needed processing time), consider the PE with the tasks that require the most time to carry out, and adjust the clock frequency to insure that the most heavily loaded PE carries out its assigned tasks within the required completion time. Once the frequency is thus determined, a minimum supply voltage can be determined.

Ex. 1009 (Nicol) at 3:66-4:6.

C. Petitioner has not Carried its Burden to Show That Nicol in view of Kling Renders Obvious Dependent Claim 3

Petitioner has not shown that Nicol in view of Kling discloses "said clock frequency is determined in part by available power." Petitioner relies on Kling for this element. *Id*. at ¶72. Petitioner quotes Kling as disclosing "the upper limit is a constant value … associated with an approximately maximum power that can be reliably consumed by the computer system." Pet., 48. Petitioner argues that "[a]

Paper 20 (POR) at 47.

PACT EX. 2028

Petitioner does not explain how the combination of Nicol and Kling determines clock frequency

51. A POSITA would have understood that this functionality can be readily implemented in Nicol's system with reasonable expectation of success. Ex. 1001 ¶204. The functionality involves relatively simple add-on software to Nicol's more complex task scheduler software residing in the controller, and relies on well-known measurement techniques, such as measuring a voltage across an appropriately placed resistor, and would have yielded predictable results. Ex. 1001, ¶¶205–206. In addition, a POSITA would have been motivated to add the functionality to Nicol's system to implement the design incentives of protecting against over-temperature by controlling power consumption, consistent with Nicol's objective. Ex. 1001, ¶206.

Paper 2 (Petition) at 48.

E. Nicol in View of Kling renders Obvious Dependent Claim 3

PO's argument that Nicol in view of Kling does not render obvious "said clock frequency is determined in part by available power" is baseless. POR, 47. PO does not deny that Kling teaches an available power limit that, when exceeded, would trigger frequency reduction. POR, 49; *see also* Ex. 1055, ¶46. Nothing more is required to satisfy this claim element.[6]

Paper 27 (Reply) at 24.

A POSITA would not have a reasonable expectation of success when combining Nicol and Kling

(c) **Incorporating Kling's upper power limit would not be a simple matter as Petitioner asserts**

75. In my opinion, the idea that a POSITA would easily incorporate the invention of Kling into Nicol via simple add-on software is untrue. Balancing upper limits with Nicol's power management system would have been a complex task in which a POSITA would have to balance the power needs of individual processors to accomplish tasks within a designated time along with temperature consideration which is also an element of the claim. A POSITA may spend their entire career developing such a system that responds appropriately and efficiently to these variables.

Ex. 2024 (Mangione-Smith POR Decl.) at 59, ¶75.

PACT EX. 2028

Nicol and Kling do not render claim 4 obvious

U.S. Patent 9,075,605 B2
Dependent Claim 4

1. A method for operating a multiprocessor system comprising a plurality of data processing cells, each of the plurality of data processing units (a) including at least one Arithmetic Logic Unit (ALU) and a register unit and (b) being adapted for sequentially processing data, the method comprising:

the multiprocessor system setting a clock frequency, of at least a part of the multiprocessor system to a minimum in accordance with a number of pending operations of a first processor;

the multiprocessor system subsequently increasing the clock frequency of the at least the part of the multiprocessor system to a maximum in accordance with a number of pending operations of a second processor; and

the multiprocessor system subsequently reducing the clock frequency of the at least the part of the multiprocessor system in accordance with (a) an operating temperature threshold preventing over-temperature and (b) a hysteresis characteristic.

Ex. 1003 (US 9,075,605) at 15:6-16:5.

4. The method of claim **1** wherein:
said clock frequency is determined in part by a user setting.

Ex. 1003 (US 9,075,605) at 16:12-13.

PACT EX. 2028

Kling's modifiable upper threshold

system, typically the processors. The upper threshold set in the power controller is shown as the upper limit line in the graph of FIG. **2A**. This threshold may be permanently set within the power controller or may be modifiable by the system designer or system user via software or hardware control.

For one embodiment of the present invention, the upper limit is a constant value as shown in FIG. **2A**. This constant value may be associated with an approximately maximum power that can be reliably consumed by the computer system from a power outlet into which the system is plugged before the outlet's fuse or circuit breaker trips. This value may be changed by, for example, modifying the settings stored in the basic input/output system (BIOS) of the computer. For another embodiment of the present invention, the

Ex. 1010 (Kling) at 5:14-28.

Kling's modifiable lower threshold

Nicol in view of Kling renders obvious this claim element. Ex. 1001 ¶207. As discussed above in Section XI.A.1, XI.A.3 and XI.A.4, a POSITA would have found it obvious to add Kling's power control functionality to Nicol's system. Ex. 1001 ¶208. Kling teaches measuring the power (current or voltage across an appropriate resistor) of a processor and comparing the measured power against an "upper limit" threshold as well as a "lower limit" threshold, as shown in Figure 2A (annotated) below. The thresholds are user settings. Ex. 1010, 5:57–59 ("Alternatively, ***the lower threshold may be modifiable by a user*** or automatically adjusted within the computer system to, for example, reduce power oscillation.").

Paper 2 (Petition) at 49.

Kling's upper and lower limit power thresholds

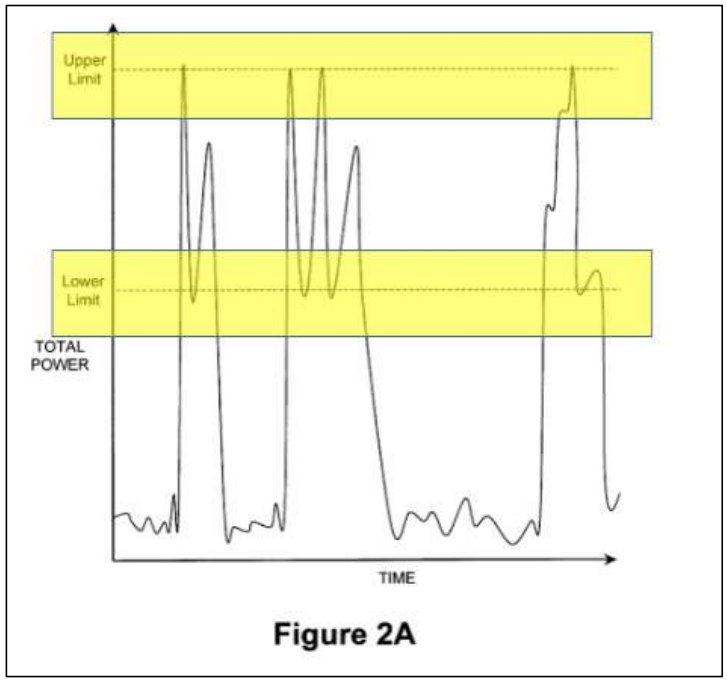

Figure 2A

Paper 2 (Petition) at 49
citing Ex. 1010 (Kling) at Fig. 2A.

Nicol in view of Kling renders obvious this claim element. Ex. 1001 ¶207. As discussed above in Section XI.A.1, XI.A.3 and XI.A.4, a POSITA would have found it obvious to add Kling's power control functionality to Nicol's system. Ex. 1001 ¶208. Kling teaches measuring the power (current or voltage across an appropriate resistor) of a processor and comparing the measured power against an "upper limit" threshold as well as a "lower limit" threshold, as shown in Figure 2A (annotated) below. The thresholds are user settings. Ex. 1010, 5:57–59 ("Alternatively, ***the lower threshold may be modifiable by a user*** or automatically adjusted within the computer system to, for example, reduce power oscillation.").

Paper 2 (Petition) at 49.

Case 1: Nicol's voltage < Kling's voltage threshold

The alleged "maximum" will not trigger Kling's threshold

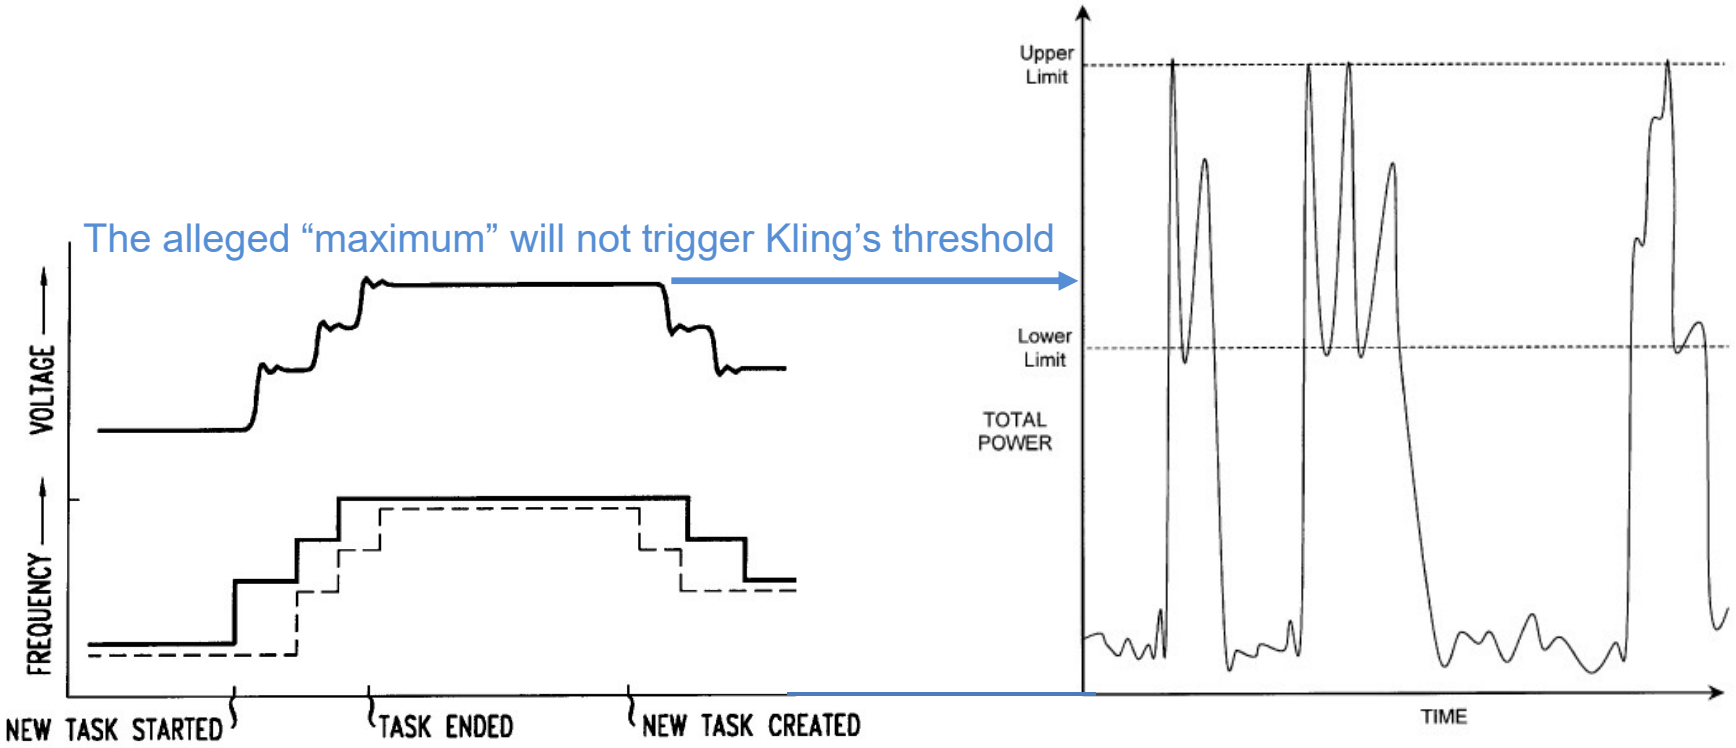

Ex. 1009 (Nicol) at Fig. 3.

Ex. 1010 (Kling) at Fig. 2A.

Nicol

Kling

PACT EX. 2028

Case 2: Nicol's voltage > Kling's voltage threshold

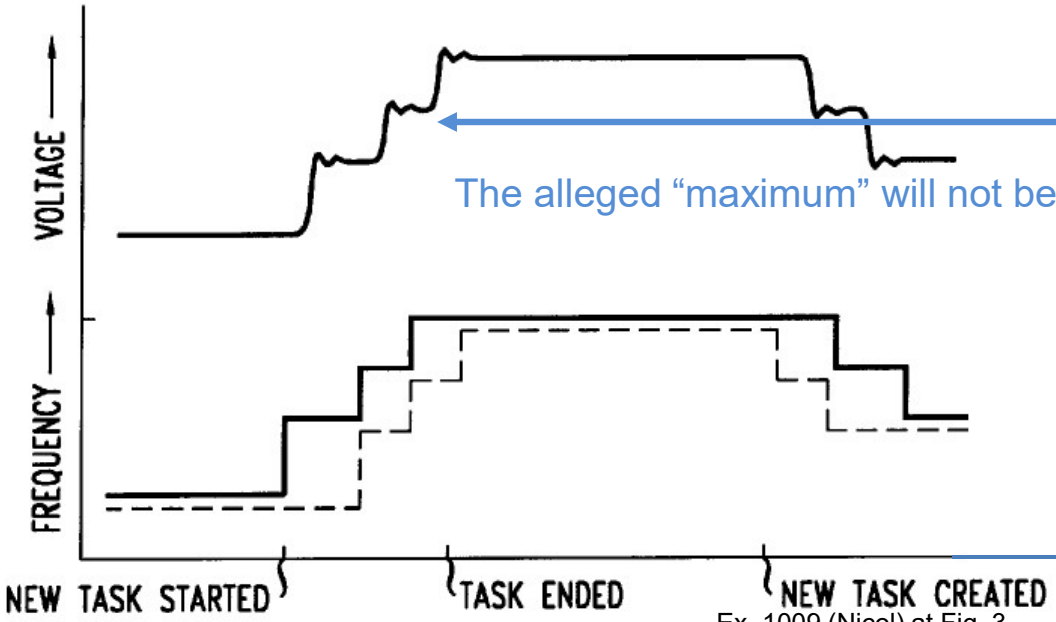

The alleged "maximum" will not be achieved

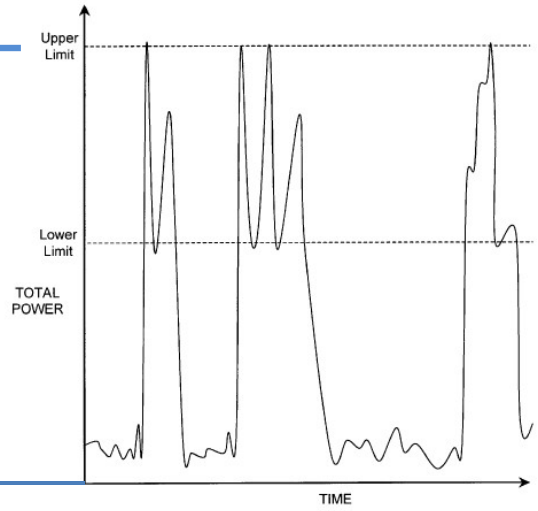

Ex. 1009 (Nicol) at Fig. 3.

Ex. 1010 (Kling) at Fig. 2A.

Nicol

Kling

PACT EX. 2028

Nicol and Kling do not render claim 5 obvious

Petitioner did not show when available power and user setting both determine the system's clock frequency

1. A method for operating a multiprocessor system comprising a plurality of data processing cells, each of the plurality of data processing units (a) including at least one Arithmetic Logic Unit (ALU) and a register unit and (b) being adapted for sequentially processing data, the method comprising:

the multiprocessor system setting a clock frequency, of at least a part of the multiprocessor system to a minimum in accordance with a number of pending operations of a first processor;

the multiprocessor system subsequently increasing the clock frequency of the at least the part of the multiprocessor system to a maximum in accordance with a number of pending operations of a second processor; and

the multiprocessor system subsequently reducing the clock frequency of the at least the part of the multiprocessor system in accordance with (a) an operating temperature threshold preventing over-temperature and (b) a hysteresis characteristic.

Ex. 1003 (US 9,075,605) at 15:6-16:5.

5. The method of claim **1** wherein:

said clock frequency is determined in part by available power and a user setting.

Ex. 1003 (US 9,075,605) at 16:14-16.

PACT EX. 2028

Teachings of Nicol and Kling

US006141762A

United States Patent [19]

Nicol et al.

[11] **Patent Number:** 6,141,762

[45] **Date of Patent:** Oct. 31, 2000

[57] **ABSTRACT**

Improved operation of multi-processor chips is achieved by dynamically controlling processing load of chips and controlling, significantly greater than on/off granularity, the operating voltages of those chips so as to minimize overall power consumption. A controller in a multi-processor chip allocates tasks to the individual processors to equalize processing load among the chips, then the controller lowers the clock frequency on the chip to as low a level as possible while assuring proper operation, and finally reduces the supply voltage. Further improvement is possible by controlling the supply voltage of individual processing elements within the multi-processor chip, as well as controlling the supply voltage of other elements in the system within which the multi-processor chip operates.

Ex. 1009 (Nicol), Abstract.

US006367023B2

(12) **United States Patent**

Kling et al.

(10) **Patent No.:** US 6,367,023 B2

(45) **Date of Patent:** Apr. 2, 2002

For one embodiment of the present invention, the upper limit is a constant value as shown in FIG. **2**A. This constant value may be associated with an approximately maximum power that can be reliably consumed by the computer system from a power outlet into which the system is plugged before the outlet's fuse or circuit breaker trips. This value

Ex. 1010 (Kling) at 5:20-26.

PACT EX. 2028

Case 1: Nicol's voltage < Kling's voltage threshold

The alleged "maximum" will not trigger Kling's threshold

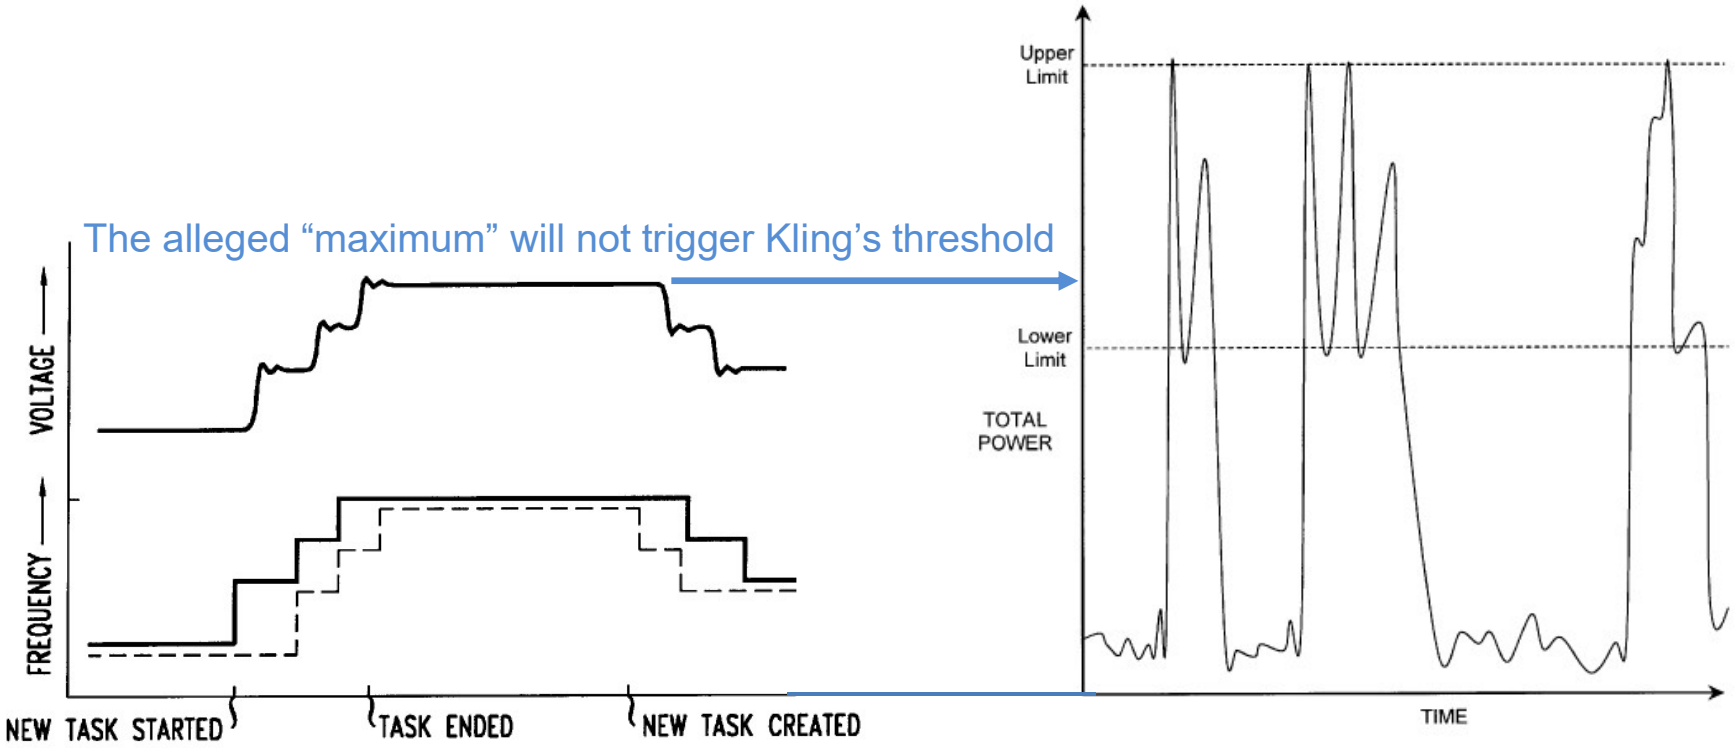

Ex. 1009 (Nicol) at Fig. 3.

Ex. 1010 (Kling) at Fig. 2A.

Nicol

Kling

PACT EX. 2028

Case 2: Nicol's voltage > Kling's voltage threshold

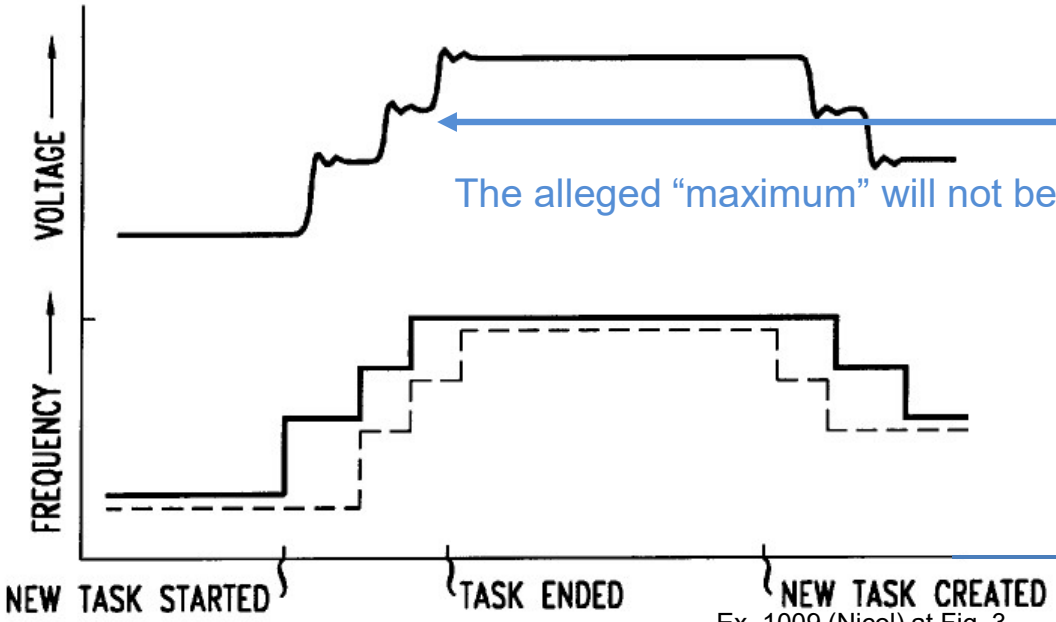

The alleged "maximum" will not be achieved

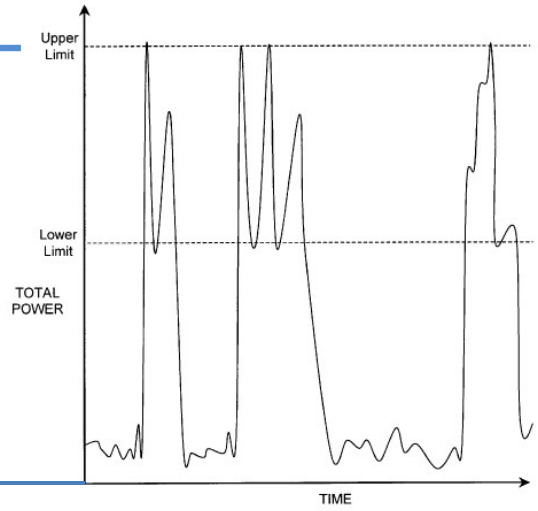

Ex. 1009 (Nicol) at Fig. 3.

Ex. 1010 (Kling), at Fig. 2A.

Nicol

Kling

PACT EX. 2028

Nicol and Kling do not render claim 6 obvious

U.S. Patent 9,075,605 B2
Dependent Claim 6

1. A method for operating a multiprocessor system comprising a plurality of data processing cells, each of the plurality of data processing units (a) including at least one Arithmetic Logic Unit (ALU) and a register unit and (b) being adapted for sequentially processing data, the method comprising:

the multiprocessor system setting a clock frequency, of at least a part of the multiprocessor system to a minimum in accordance with a number of pending operations of a first processor;

the multiprocessor system subsequently increasing the clock frequency of the at least the part of the multiprocessor system to a maximum in accordance with a number of pending operations of a second processor; and

the multiprocessor system subsequently reducing the clock frequency of the at least the part of the multiprocessor system in accordance with (a) an operating temperature threshold preventing over-temperature and (b) a hysteresis characteristic.

Ex. 1003 (US 9,075,605) at 15:6-16:5.

6. The method of claim **1** wherein:

said number of pending operations is determined at by the fill level of one or more buffers.

Ex. 1003 (US 9,075,605) at 16:17-19

PACT EX. 2028

Claim 6: "said number of pending operations is determined [] by the fill level of one or more buffers"

Nicol renders obvious this claim element. Ex. 1001 ¶217. Nicol teaches estimating processing loads based on the number of instructions of each of the "to do" tasks to be allocated to the PEs:

> If the ***number of instructions*** that need to be executed for each task is known and made available to the operating system, a scheduler within the operating system can use this information to determine the best way to allocate the tasks to the available processors in order to balance the computation.

Paper 2 (Petition) at 51.

Claim 1: Petitioner argues that "said number of pending operation" is the "number of pending operations of [each individual] processor."

2. Independent Claim 1

e) ***"the multiprocessor system setting a clock frequency, of at least a part of the multiprocessor system to a minimum in accordance with a number of pending operations of a first processor;"***

Nicol discloses this claim element. Ex. 1001 ¶161. Nicol explains that the controller (red) in its multiprocessor system sets a clock frequency to a minimum according to a number of pending operations (instructions) of each of the PEs of the system, including the number of pending operations of PE1 ("first processor") in the annotated Figure 4 below. Ex. 1001 ¶162; Ex. 1009, Figure 4.

Paper 2 (Petition) at 25, 35.

PACT EX. 2028

Claim 1: Petitioner argues that "said number of pending operation" is the "number of pending operations of [each individual] processor."

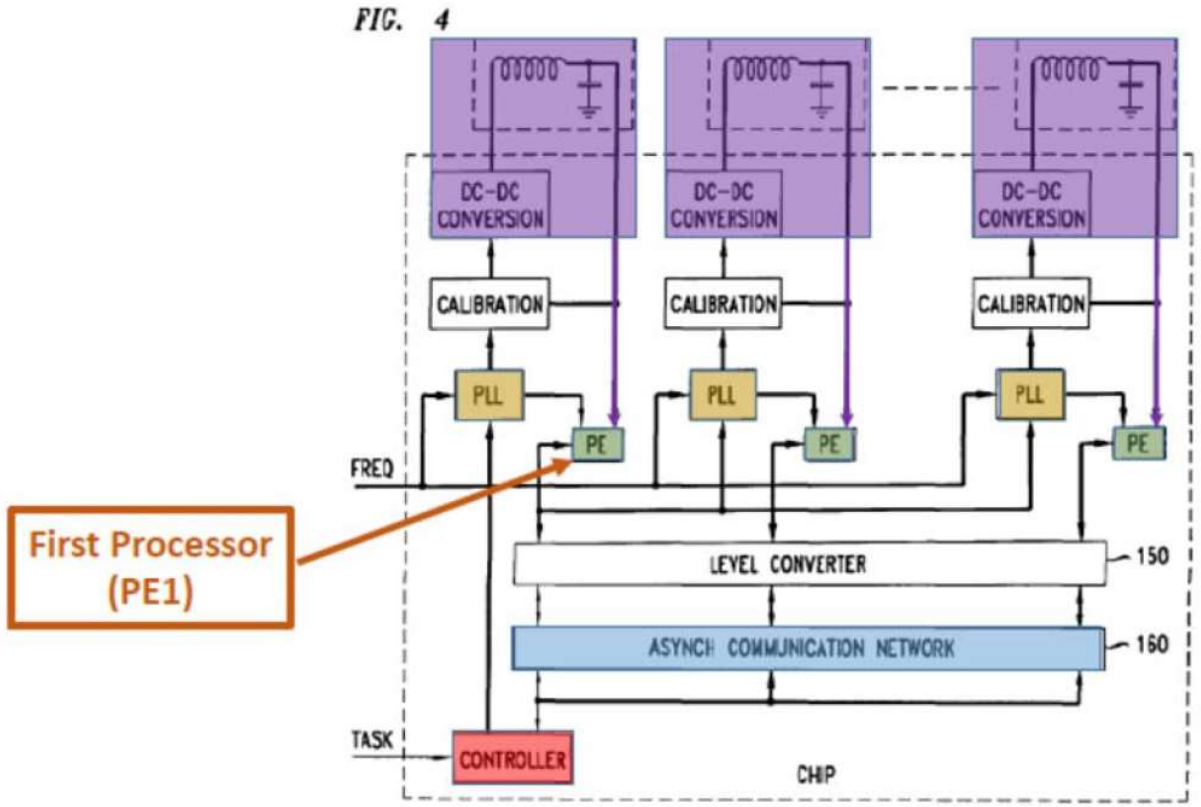

Paper 2 (Petition) at 36 (annotated by Petitioner).

PACT EX. 2028

Claim 1: Petitioner argues that "said number of pending operation" is the "number of pending operations of [each individual] processor."

> Nicol discloses this claim element. Ex. 1001 ¶161. Nicol explains that the controller (red) in its multiprocessor system sets a clock frequency to a minimum according to a number of pending operations (instructions) of each of the PEs of the system, including the number of pending operations of PE1 ("first processor") in the annotated Figure 4 below. Ex. 1001 ¶162; Ex. 1009, Figure 4.

Paper 2 (Petition) at 35.

PACT EX. 2028

The fill level of the data buffers may indicate the processor load, but the reverse is not necessarily true

1022 at 155 and Figure 3.11 (showing an "instruction queue"). Indeed, the applicant stated during the prosecution of the '605 patent that "one of ordinary skill in the art would understand that ***the fill level of a data buffer [is] representative of the number of waiting operations or processor load.***" Ex. 1004 at 377.

Paper 2 (Petition) at 51.

Cline and Hong proves no more than inspecting the buffer fill level was independently known in the prior art

An invention "composed of several elements is not proved obvious **merely by demonstrating that each of its elements was, independently, known** in the prior art."

KSR Int'l Co. v. Teleflex Inc., 550 U.S. 398, 418 (2007).

PACT EX. 2028