

JOHN WATKINSON

# THE MPEG HANDBOOK

MPEG-1 MPEG-2 MPEG-4

THE MPEG HANDBOOK  
INCLUDES:

- THE MPEG-4 STANDARD
- MP3 AND NETWORKS
- EXTENSIVE EXAMPLES OF  
APPLICATIONS

NEW! AN ENTIRE NEW EDITION... UPDATED EVERYTHING YOU  
NEED TO KNOW ABOUT DIGITAL VIDEO



Focal Press

# **The MPEG Handbook**

**MPEG-1, MPEG-2, MPEG-4**

**John Watkinson**



**Focal Press**

OXFORD AUCKLAND BOSTON JOHANNESBURG MELBOURNE NEW DELHI

Focal Press  
An imprint of Butterworth-Heinemann  
Linacre House, Jordan Hill, Oxford OX2 8DP  
225 Wildwood Avenue, Woburn, MA 01801-2041  
A division of Reed Educational and Professional Publishing Ltd

 A member of the Reed Elsevier plc group

First published 2001

© John Watkinson 2001

All rights reserved. No part of this publication may be reproduced in any material form (including photocopying or storing in any medium by electronic means and whether or not transiently or incidentally to some other use of this publication) without the written permission of the copyright holder except in accordance with the provisions of the Copyright, Designs and Patents Act 1988 or under the terms of a licence issued by the Copyright Licensing Agency Ltd, 90 Tottenham Court Road, London, England W1P 0LP. Applications for the copyright holder's written permission to reproduce any part of this publication should be addressed to the publishers

**British Library Cataloguing in Publication Data**

A catalogue record for this book is available from the British Library

**Library of Congress Cataloguing in Publication Data**

A catalogue record for this book is available from the Library of Congress

For information on all Focal Press publications visit  
our website at [www.focalpress.com](http://www.focalpress.com)

ISBN 0 240 51656 7

Composition by Genesis Typesetting, Rochester, Kent  
Printed and bound in Great Britain



# Contents

---

<i>Preface</i>	xi
----------------	----

<i>Acknowledgements</i>	xiii
-------------------------	------

<b>Chapter 1 Introduction to compression</b>	<b>1</b>
--	----------

1.1 What is MPEG?	1
1.2 Why compression is necessary	3
1.3 MPEG-1, 2 and 4 contrasted	4
1.4 Some applications of compression	5
1.5 Lossless and perceptive coding	6
1.6 Compression principles	8
1.7 Video compression	13
1.7.1 Intra-coded compression	14
1.7.2 Inter-coded compression	15
1.7.3 Introduction to motion compensation	16
1.7.4 Film-originated video compression	18
1.8 Introduction to MPEG-1	19
1.9 MPEG-2: Profiles and Levels	20
1.10 Introduction to MPEG-4	22
1.11 Audio compression	26
1.11.1 Sub-band coding	26
1.11.2 Transform coding	26
1.11.3 Predictive coding	26
1.12 MPEG bitstreams	27
1.13 Drawbacks of compression	28



vi *Contents*

1.14	Compression pre-processing	29
1.15	Some guidelines	30
	References	31
 <b>Chapter 2 Fundamentals</b>		<b>32</b>
2.1	What is an audio signal?	32
2.2	What is a video signal?	32
2.3	Types of video	33
2.4	What is a digital signal?	35
2.5	Sampling	38
2.6	Reconstruction	42
2.7	Aperture effect	46
2.8	Choice of audio sampling rate	49
2.9	Video sampling structures	51
2.10	The phase-locked loop	54
2.11	Quantizing	56
2.12	Quantizing error	58
2.13	Dither	61
2.14	Introduction to digital processing	63
2.15	Logic elements	65
2.16	Storage elements	67
2.17	Binary coding	69
2.18	Gain control	78
2.19	Floating-point coding	81
2.20	Multiplexing principles	82
2.21	Packets	83
2.22	Statistical multiplexing	84
2.23	Timebase correction	84
	References	86
 <b>Chapter 3 Processing for compression</b>		<b>87</b>
3.1	Introduction	87
3.2	Transforms	90
3.3	Convolution	92
3.4	FIR and IIR filters	95
3.5	FIR filters	95
3.6	Interpolation	102
3.7	Downsampling filters	113
3.8	The quadrature mirror filter	114
3.9	Filtering for video noise reduction	118
3.10	Warping	119

29	3.11 Transforms and duality	126
30	3.12 The Fourier transform	128
31	3.13 The discrete cosine transform (DCT)	137
	3.14 The wavelet transform	138
32	3.15 The importance of motion compensation	144
	3.16 Motion-estimation techniques	146
32	3.16.1 Block matching	146
32	3.16.2 Gradient matching	148
33	3.16.3 Phase correlation	149
35	3.17 Motion-compensated displays	153
38	3.18 Camera-shake compensation	155
42	3.19 Motion-compensated de-interlacing	157
46	3.20 Compression and requantizing	158
49	References	163
51		
54	<b>Chapter 4 Audio compression</b>	<b>165</b>
56		
58	4.1 Introduction	165
61	4.2 The deciBel	166
63	4.3 Audio level metering	171
65	4.4 The ear	174
67	4.5 The cochlea	176
69	4.6 Level and loudness	178
78	4.7 Frequency discrimination	179
81	4.8 Critical bands	180
82	4.9 Beats	182
83	4.10 Codec level calibration	184
84	4.11 Quality measurement	185
84	4.12 The limits	187
86	4.13 Compression applications	187
	4.14 Audio compression tools	188
87	4.15 Sub-band coding	193
	4.16 Audio compression formats	194
87	4.17 MPEG audio compression	194
90	4.18 MPEG Layer I audio coding	197
92	4.19 MPEG Layer II audio coding	202
95	4.20 MPEG Layer III audio coding	204
95	4.21 MPEG-2 AAC – advanced audio coding	207
102	4.22 Dolby AC-3	213
113	4.23 MPEG-4 Audio	215
114	4.24 MPEG-4 AAC	215
118	4.25 Compression in stereo and surround sound	216
119	References	222

<b>Chapter 5</b>	<b>MPEG video compression</b>	<b>224</b>
5.1	The eye	224
5.2	Dynamic resolution	228
5.3	Contrast	232
5.4	Colour vision	233
5.5	Colour difference signals	235
5.6	Progressive or interlaced scan?	238
5.7	Spatial and temporal redundancy in MPEG	243
5.8	<i>I</i> and <i>P</i> coding	248
5.9	Bidirectional coding	249
5.10	Coding applications	252
5.11	Intra-coding	253
5.12	Intra-coding in MPEG-1 and MPEG-2	257
5.13	A bidirectional coder	261
5.14	Slices	265
5.15	Handling interlaced pictures	266
5.16	MPEG-1 and MPEG-2 coders	271
5.17	The elementary stream	272
5.18	An MPEG-2 decoder	273
5.19	MPEG-4	276
5.20	Video objects	278
5.21	Texture coding	281
5.22	Shape coding	287
5.23	Padding	289
5.24	Video object coding	290
5.25	Two-dimensional mesh coding	293
5.26	Sprites	298
5.27	Wavelet-based compression	300
5.28	Three-dimensional mesh coding	304
5.29	Animation	313
5.30	Scaleability	315
5.31	Coding artifacts	318
5.32	MPEG and concatenation	322
	References	328
<b>Chapter 6</b>	<b>Program and transport streams</b>	<b>329</b>
6.1	Introduction	329
6.2	Packets and time stamps	329
6.3	Transport streams	332
6.4	Clock references	333
6.5	Program Specific Information (PSI)	335

224	6.6 Multiplexing	336
224	6.7 Remultiplexing	338
228	Reference	339
232		
233		
235	<b>Chapter 7 MPEG applications</b>	<b>340</b>
238	7.1 Introduction	340
243	7.2 Video phones	342
248	7.3 Digital television broadcasting	342
249	7.4 The DVB receiver	355
252	7.5 CD-Video and DVD	356
253	7.6 Personal video recorders	361
257	7.7 Networks	363
261	7.8 FireWire	371
265	7.9 Broadband networks and ATM	373
266	7.10 ATM AALs	377
271	References	380
272		
273	<b>Index</b>	<b>381</b>
276		
278		
281		
287		
289		
290		
293		
298		
300		
304		
313		
315		
318		
322		
328		
329		
329		
329		
332		
333		
335		

# Preface

---

This book completely revises the earlier book entitled *MPEG-2*. It is an interesting insight into the rate at which this technology progresses that this book was in preparation only a year after *MPEG-2* was first published. The impetus for the revision is, of course, *MPEG-4* which is comprehensively covered here. The opportunity has also been taken to improve a number of explanations and to add a chapter on applications of *MPEG*.

The approach of the book has not changed in the slightest. Compression is a specialist subject with its own library of specialist terminology which is generally accompanied by a substantial amount of mathematics. I have always argued that mathematics is only a form of shorthand, itself a compression technique! Mathematics describes but does not explain, whereas this book explains and then describes.

A chapter of fundamentals is included to make the main chapters easier to follow. Also included are some guidelines which have been found practically useful in getting the best out of compression systems.

The reader who has endured this book will be in a good position to tackle the *MPEG* standards documents themselves, although these are not for the faint-hearted, especially the *MPEG-4* documents which are huge and impenetrable. One wonders what they will come up with next!

# Acknowledgements

---

Information for this book has come from a number of sources to whom I am indebted. The publications of the ISO, AES and SMPTE provided essential reference material. Thanks also to the following for lengthy discussions and debates: Peter de With, Steve Lyman, Bruce Devlin, Mike Knee, Peter Kraniuskas and Tom MacMahon. The assistance of MicroSoft Corp. and Tektronix Inc. is also appreciated. Special thanks to Mikael Reichel.

# 1

## Introduction to compression

---

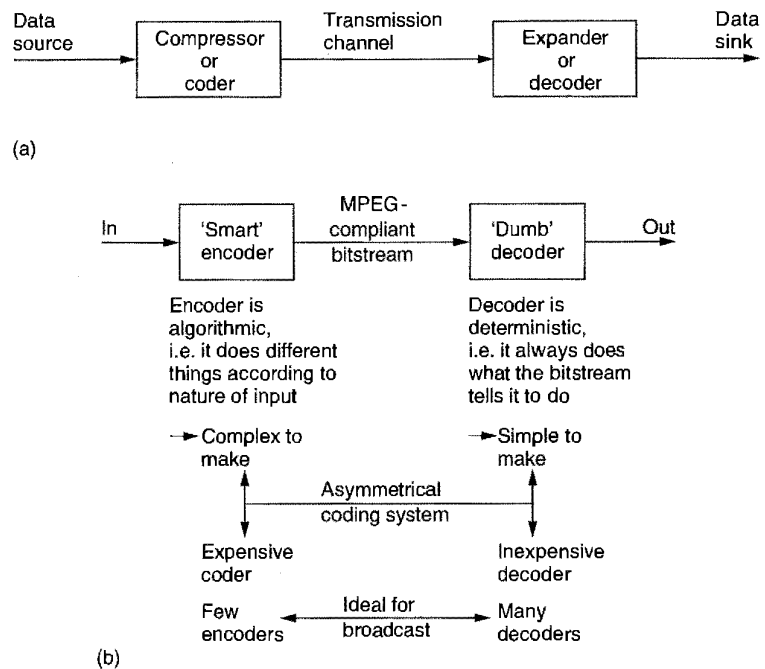
### 1.1 What is MPEG?

MPEG is actually an acronym for the Moving Pictures Experts Group which was formed by the ISO (International Standards Organization) to set standards for audio and video compression and transmission.

Compression is summarized in Figure 1.1. It will be seen in (a) that the data rate is reduced at source by the *compressor*. The compressed data are then passed through a communication channel and returned to the original rate by the *expander*. The ratio between the source data rate and the channel data rate is called the *compression factor*. The term *coding gain* is also used. Sometimes a compressor and expander in series are referred to as a *combander*. The compressor may equally well be referred to as a *coder* and the expander a *decoder* in which case the tandem pair may be called a *codec*.

Where the encoder is more complex than the decoder, the system is said to be asymmetrical. Figure 1.1(b) shows that MPEG works in this way. The encoder needs to be algorithmic or adaptive whereas the decoder is 'dumb' and carries out fixed actions. This is advantageous in applications such as broadcasting where the number of expensive complex encoders is small but the number of simple inexpensive decoders is large. In point-to-point applications the advantage of asymmetrical coding is not so great.

The approach of the ISO to standardization in MPEG is novel because it is not the encoder which is standardized. Figure 1.2(a) shows that instead the way in which a decoder shall interpret the bitstream is defined. A decoder which can successfully interpret the bitstream is said to be *compliant*. Figure 1.2(b) shows that the advantage of standardizing the decoder is that over time encoding algorithms can improve yet compliant decoders will continue to function with them.



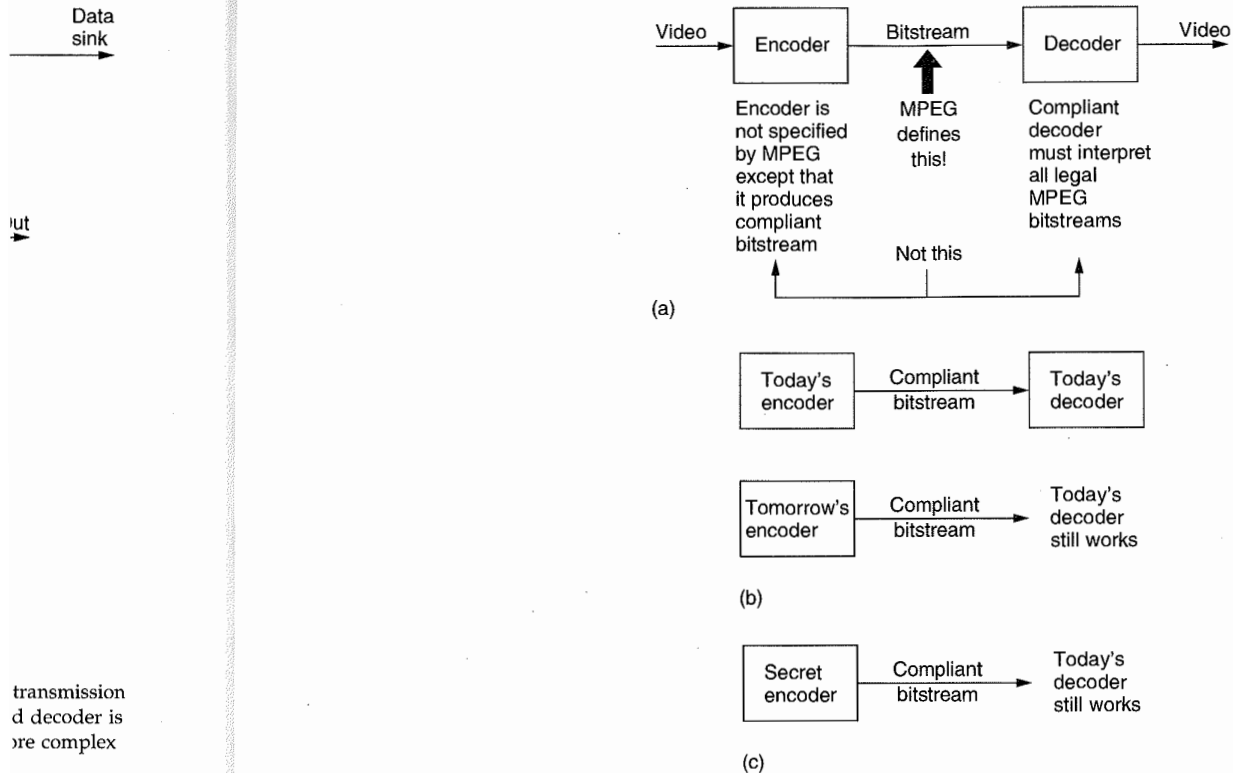
**Figure 1.1** In (a) a compression system consists of compressor or coder, a transmission channel and a matching expander or decoder. The combination of coder and decoder is known as a codec. (b) MPEG is asymmetrical since the encoder is much more complex than the decoder.

It should be noted that a compliant decoder must correctly be able to interpret every allowable bitstream, whereas an encoder which produces a restricted subset of the possible codes can still be compliant.

The MPEG standards give very little information regarding the structure and operation of the encoder. Provided the bitstream is compliant, any coder construction will meet the standard, although some designs will give better picture quality than others. Encoder construction is not revealed in the bitstream and manufacturers can supply encoders using algorithms which are proprietary and their details do not need to be published. A useful result is that there can be competition between different encoder designs which means that better designs can evolve. The user will have greater choice because different levels of cost and complexity can exist in a range of coders yet a compliant decoder will operate with them all.

MPEG is, however, much more than a compression scheme as it also standardizes the protocol and syntax under which it is possible to combine or multiplex audio data with video data to produce a digital equivalent of a television program. Many such programs can be





**Figure 1.2** (a) MPEG defines the protocol of the bitstream between encoder and decoder. The decoder is defined by implication, the encoder is left very much to the designer. (b) This approach allows future encoders of better performance to remain compatible with existing decoders. (c) This approach also allows an encoder to produce a standard bitstream while its technical operation remains a commercial secret.

combined in a single multiplex and MPEG defines the way in which such multiplexes can be created and transported. The definitions include the metadata which decoders require to demultiplex correctly and which users will need to locate programs of interest.

As with all video systems there is a requirement for synchronizing or genlocking and this is particularly complex when a multiplex is assembled from many signals which are not necessarily synchronized to one another.

## 1.2 Why compression is necessary

Compression, bit rate reduction, data reduction and source coding are all terms which mean basically the same thing in this context. In essence the

same (or nearly the same) information is carried using a smaller quantity or rate of data. It should be pointed out that in audio *compression* traditionally means a process in which the dynamic range of the sound is reduced. In the context of MPEG the same word means that the bit rate is reduced, ideally leaving the dynamics of the signal unchanged. Provided the context is clear, the two meanings can co-exist without a great deal of confusion.

There are several reasons why compression techniques are popular:

- (a) Compression extends the playing time of a given storage device.
- (b) Compression allows miniaturization. With fewer data to store, the same playing time is obtained with smaller hardware. This is useful in ENG (electronic news gathering) and consumer devices.
- (c) Tolerances can be relaxed. With fewer data to record, storage density can be reduced making equipment which is more resistant to adverse environments and which requires less maintenance.
- (d) In transmission systems, compression allows a reduction in bandwidth which will generally result in a reduction in cost. This may make possible a service which would be impracticable without it.
- (e) If a given bandwidth is available to an uncompressed signal, compression allows faster than real-time transmission in the same bandwidth.
- (f) If a given bandwidth is available, compression allows a better-quality signal in the same bandwidth.

### 1.3 MPEG-1, 2 and 4 contrasted

The first compression standard for audio and video was MPEG-1.<sup>1,2</sup> Although many applications have been found, MPEG-1 was basically designed to allow moving pictures and sound to be encoded into the bit rate of an audio Compact Disc. The resultant Video-CD was quite successful but has now been superseded by DVD. In order to meet the low bit requirement, MPEG-1 downsampled the images heavily as well as using picture rates of only 24–30 Hz and the resulting quality was moderate.

The subsequent MPEG-2 standard was considerably broader in scope and of wider appeal.<sup>3</sup> For example, MPEG-2 supports interlace and HD whereas MPEG-1 did not. MPEG-2 has become very important because it has been chosen as the compression scheme for both DVB (digital video broadcasting) and DVD (digital video disk). Developments in standardizing scalable and multi-resolution compression which would have become MPEG-3 were ready by the time MPEG-2 was ready to be

standardized and so this work was incorporated into MPEG-2, and as a result there is no MPEG-3 standard.

MPEG-4<sup>4</sup> uses further coding tools with additional complexity to achieve higher compression factors than MPEG-2. In addition to more efficient coding of video, MPEG-4 moves closer to computer graphics applications. In the more complex Profiles, the MPEG-4 decoder effectively becomes a rendering processor and the compressed bitstream describes three-dimensional shapes and surface texture. It is to be expected that MPEG-4 will become as important to Internet and wireless delivery as MPEG-2 has become in DVD and DVB.


## 1.4 Some applications of compression

The applications of audio and video compression are limitless and the ISO has done well to provide standards which are appropriate to the wide range of possible compression products.

MPEG coding embraces video pictures from the tiny screen of a videophone to the high-definition images needed for electronic cinema. Audio coding stretches from speech-grade mono to multichannel surround sound.

Figure 1.3 shows the use of a codec with a recorder. The playing time of the medium is extended in proportion to the compression factor. In the case of tapes, the access time is improved because the length of tape needed for a given recording is reduced and so it can be rewound more quickly. In the case of DVD (digital video disk aka digital versatile disk) the challenge was to store an entire movie on one 12 cm disk. The storage density available with today's optical disk technology is such that consumer recording of conventional uncompressed video would be out of the question.

In communications, the cost of data links is often roughly proportional to the data rate and so there is simple economic pressure to use a high compression factor. However, it should be borne in mind that implement-



**Figure 1.3** Compression can be used around a recording medium. The storage capacity may be increased or the access time reduced according to the application.

ing the codec also has a cost which rises with compression factor and so a degree of compromise will be inevitable.

In the case of video-on-demand, technology exists to convey full bandwidth video to the home, but to do so for a single individual at the moment would be prohibitively expensive. Without compression, HDTV (high-definition television) requires too much bandwidth. With compression, HDTV can be transmitted to the home in a similar bandwidth to an existing analog SDTV channel. Compression does not make video-on-demand or HDTV possible; it makes them economically viable.

In workstations designed for the editing of audio and/or video, the source material is stored on hard disks for rapid access. Whilst top-grade systems may function without compression, many systems use compression to offset the high cost of disk storage. In some systems a compressed version of the top-grade material may also be stored for browsing purposes.

When a workstation is used for *off-line* editing, a high compression factor can be used and artifacts will be visible in the picture. This is of no consequence as the picture is only seen by the editor who uses it to make an EDL (edit decision list) which is no more than a list of actions and the timecodes at which they occur. The original uncompressed material is then *conformed* to the EDL to obtain a high-quality edited work. When *on-line* editing is being performed, the output of the workstation is the finished product and clearly a lower compression factor will have to be used. Perhaps it is in broadcasting where the use of compression will have its greatest impact. There is only one electromagnetic spectrum and pressure from other services such as cellular telephones makes efficient use of bandwidth mandatory. Analog television broadcasting is an old technology and makes very inefficient use of bandwidth. Its replacement by a compressed digital transmission is inevitable for the practical reason that the bandwidth is needed elsewhere.

Fortunately in broadcasting there is a mass market for decoders and these can be implemented as low-cost integrated circuits. Fewer encoders are needed and so it is less important if these are expensive. Whilst the cost of digital storage goes down year on year, the cost of the electromagnetic spectrum goes up. Consequently in the future the pressure to use compression in recording will ease whereas the pressure to use it in radio communications will increase.

## 1.5 Lossless and perceptive coding

Although there are many different coding techniques, all of them fall into one or other of these categories. In *lossless* coding, the data from the expander are identical bit-for-bit with the original source data. The

actor and so

convey full  
individual at the  
sion, HDTV  
th compres-  
width to an  
e video-on-  
ole.

r video, the  
st top-grade  
se compres-  
compressed  
r browsing

ompression  
This is of no  
s it to make  
ons and the  
material is  
k. When on-  
ation is the  
I have to be  
ression will  
pectrum and  
ces efficient  
g is an old  
eplacement  
tical reason

coders and  
er encoders  
Whilst the  
ost of the  
future the  
ie pressure

m fall into  
from the  
data. The

so-called 'stacker' programs which increase the apparent capacity of disk drives in personal computers use lossless codecs. Clearly with computer programs the corruption of a single bit can be catastrophic. Lossless coding is generally restricted to compression factors of around 2:1.

It is important to appreciate that a lossless coder cannot guarantee a particular compression factor and the communications link or recorder used with it must be able to function with the variable output data rate. Source data which result in poor compression factors on a given codec are described as *difficult*. It should be pointed out that the difficulty is often a function of the codec. In other words data which one codec finds difficult may not be found difficult by another. Lossless codecs can be included in bit-error-rate testing schemes. It is also possible to cascade or *concatenate* lossless codecs without any special precautions.

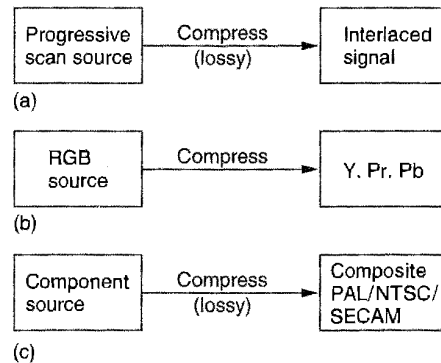
Higher compression factors are only possible with *lossy* coding in which data from the expander are not identical bit-for-bit with the source data and as a result comparing the input with the output is bound to reveal differences. Lossy codecs are not suitable for computer data, but are used in MPEG as they allow greater compression factors than lossless codecs. Successful lossy codecs are those in which the errors are arranged so that a human viewer or listener finds them subjectively difficult to detect. Thus lossy codecs must be based on an understanding of psycho-acoustic and psycho-visual perception and are often called *perceptive* codes.

In perceptive coding, the greater the compression factor required, the more accurately must the human senses be modelled. Perceptive coders can be forced to operate at a fixed compression factor. This is convenient for practical transmission applications where a fixed data rate is easier to handle than a variable rate. The result of a fixed compression factor is that the subjective quality can vary with the 'difficulty' of the input material. Perceptive codecs should not be concatenated indiscriminately especially if they use different algorithms. As the reconstructed signal from a perceptive codec is not bit-for-bit accurate, clearly such a codec cannot be included in any bit error rate testing system as the coding differences would be indistinguishable from real errors.

Although the adoption of digital techniques is recent, compression itself is as old as television. Figure 1.4 shows some of the compression techniques used in traditional television systems.

Most video signals employ a non-linear relationship between brightness and the signal voltage which is known as gamma. Gamma is a perceptive coding technique which depends on the human sensitivity to video noise being a function of the brightness. The use of gamma allows the same subjective noise level with an eight-bit system as would be achieved with a fourteen-bit linear system.

One of the oldest techniques is interlace, which has been used in analog television from the very beginning as a primitive way of reducing



**Figure 1.4** Compression is as old as television. (a) Interlace is a primitive way of halving the bandwidth. (b) Colour difference working invisibly reduces colour resolution. (c) Composite video transmits colour in the same bandwidth as monochrome.

bandwidth. As will be seen in Chapter 5, interlace is not without its problems, particularly in motion rendering. MPEG-2 supports interlace simply because legacy interlaced signals exist and there is a requirement to compress them. This should not be taken to mean that it is a good idea.

The generation of colour difference signals from *RGB* in video represents an application of perceptive coding. The human visual system (HVS) sees no change in quality although the bandwidth of the colour difference signals is reduced. This is because human perception of detail in colour changes is much less than in brightness changes. This approach is sensibly retained in MPEG.

Composite video systems such as PAL, NTSC and SECAM are all analog compression schemes which embed a subcarrier in the luminance signal so that colour pictures are available in the same bandwidth as monochrome. In comparison with a linear-light progressive scan *RGB* picture, gamma-coded interlaced composite video has a compression factor of about 10:1.

In a sense MPEG-2 can be considered to be a modern digital equivalent of analog composite video as it has most of the same attributes. For example, the eight-field sequence of the PAL subcarrier which makes editing difficult has its equivalent in the GOP (group of pictures) of MPEG.

## 1.6 Compression principles

In a PCM digital system the bit rate is the product of the sampling rate and the number of bits in each sample and this is generally constant.

Nevertheless the *information* rate of a real signal varies. In all real signals, part of the signal is obvious from what has gone before or what may come later and a suitable receiver can predict that part so that only the true information actually has to be sent. If the characteristics of a predicting receiver are known, the transmitter can omit parts of the message in the knowledge that the receiver has the ability to re-create it. Thus all encoders must contain a model of the decoder.

One definition of information is that it is the unpredictable or surprising element of data. Newspapers are a good example of information because they only mention items which are surprising. Newspapers never carry items about individuals who have *not* been involved in an accident as this is the normal case. Consequently the phrase 'no news is good news' is remarkably true because if an information channel exists but nothing has been sent then it is most likely that nothing remarkable has happened.

The unpredictability of the punch line is a useful measure of how funny a joke is. Often the build-up paints a certain picture in the listener's imagination, which the punch line destroys utterly. One of the author's favourites is the one about the newly married couple who didn't know the difference between putty and petroleum jelly – their windows fell out.

The difference between the information rate and the overall bit rate is known as the redundancy. Compression systems are designed to eliminate as much of that redundancy as practicable or perhaps affordable. One way in which this can be done is to exploit statistical predictability in signals. The information content or *entropy* of a sample is a function of how different it is from the predicted value. Most signals have some degree of predictability. A sine wave is highly predictable, because all cycles look the same. According to Shannon's theory, any signal which is totally predictable carries no information. In the case of the sine wave this is clear because it represents a single frequency and so has no bandwidth.

At the opposite extreme a signal such as noise is completely unpredictable and as a result all codecs find noise *difficult*. The most efficient way of coding noise is PCM. A codec which is designed using the statistics of real material should not be tested with random noise because it is not a representative test. Second, a codec which performs well with clean source material may perform badly with source material containing superimposed noise. Most practical compression units require some form of pre-processing before the compression stage proper and appropriate noise reduction should be incorporated into the pre-processing if noisy signals are anticipated. It will also be necessary to restrict the degree of compression applied to noisy signals.

All real signals fall part-way between the extremes of total predictability and total unpredictability or noisiness. If the bandwidth (set by the

way of  
our  
monochrome.

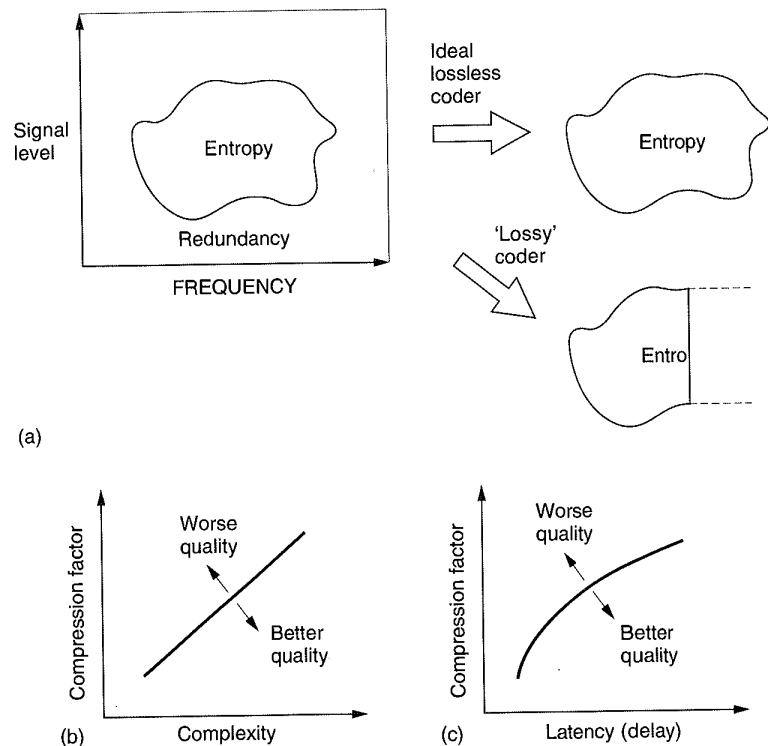
without its  
rts interlace  
requirement  
it is a good

B in video  
sual system  
f the colour  
ion of detail  
is approach

AM are all  
e luminance  
ndwidth as  
e scan RGB  
ompression

l equivalent  
ributes. For  
hich makes  
pictures) of

mpling rate  
ly constant.



**Figure 1.5** (a) A perfect coder removes only the redundancy from the input signal and results in subjectively lossless coding. If the remaining entropy is beyond the capacity of the channel some of it must be lost and the codec will then be lossy. An imperfect coder will also be lossy as it fails to keep all entropy. (b) As the compression factor rises, the complexity must also rise to maintain quality. (c) High compression factors also tend to increase latency or delay through the system.

sampling rate) and the dynamic range (set by the wordlength) of the transmission system are used to delineate an area, this sets a limit on the information capacity of the system. Figure 1.5(a) shows that most real signals only occupy part of that area. The signal may not contain all frequencies, or it may not have full dynamics at certain frequencies.

Entropy can be thought of as a measure of the actual area occupied by the signal. This is the area that *must* be transmitted if there are to be no subjective differences or *artifacts* in the received signal. The remaining area is called the *redundancy* because it adds nothing to the information conveyed. Thus an ideal coder could be imagined which miraculously sorts out the entropy from the redundancy and only sends the former. An ideal decoder would then re-create the original impression of the information quite perfectly. As the ideal is approached, the coder complexity and the latency or delay both rise. Figure 1.5(b) shows how





complexity increases with compression factor. The additional complexity of MPEG-4 over MPEG-2 is obvious from this. Figure 1.5(c) shows how increasing the codec latency can improve the compression factor.

Obviously we would have to provide a channel which could accept whatever entropy the coder extracts in order to have transparent quality. As a result moderate coding gains which only remove redundancy need not cause artifacts and result in systems which are described as *subjectively lossless*. If the channel capacity is not sufficient for that, then the coder will have to discard some of the entropy and with it useful information. Larger coding gains which remove some of the entropy must result in artifacts. It will also be seen from Figure 1.5 that an imperfect coder will fail to separate the redundancy and may discard entropy instead, resulting in artifacts at a sub-optimal compression factor.

A single variable-rate transmission or recording channel is traditionally unpopular with channel providers, although newer systems such as ATM support variable rate. Digital transmitters used in DVB have a fixed bit rate. The variable rate requirement can be overcome by combining several compressed channels into one constant rate transmission in a way which flexibly allocates data rate between the channels. Provided the material is unrelated, the probability of all channels reaching peak entropy at once is very small and so those channels which are at one instant passing easy material will make available transmission capacity for those channels which are handling difficult material. This is the principle of statistical multiplexing.

Where the same type of source material is used consistently, e.g. English text, then it is possible to perform a statistical analysis on the frequency with which particular letters are used. Variable-length coding is used in which frequently used letters are allocated short codes and letters which occur infrequently are allocated long codes. This results in a lossless code. The well-known Morse code used for telegraphy is an example of this approach. The letter e is the most frequent in English and is sent with a single dot. An infrequent letter such as z is allocated a long complex pattern. It should be clear that codes of this kind which rely on a prior knowledge of the statistics of the signal are only effective with signals actually having those statistics. If Morse code is used with another language, the transmission becomes significantly less efficient because the statistics are quite different; the letter z, for example, is quite common in Czech.

The Huffman code<sup>5</sup> is also one which is designed for use with a data source having known statistics. The probability of the different code values to be transmitted is studied, and the most frequent codes are arranged to be transmitted with short wordlength symbols. As the probability of a code value falls, it will be allocated longer wordlength.



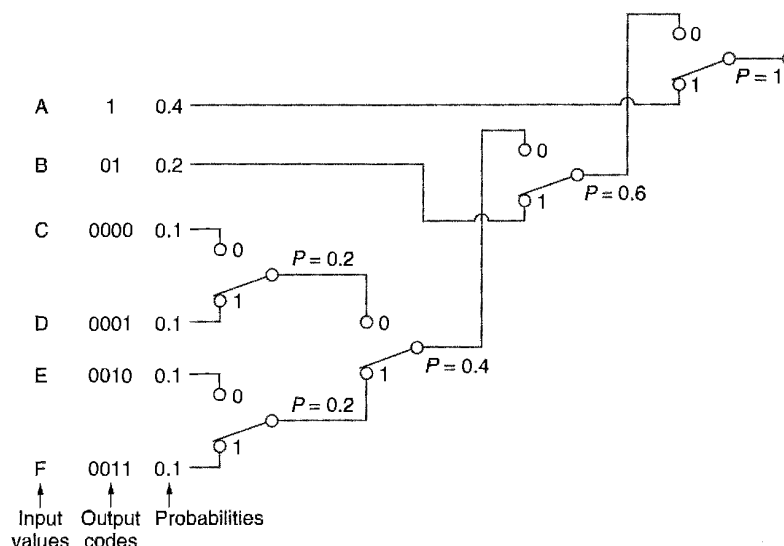
er  
ty



out signal and  
the capacity of  
nperfect coder  
tor rises, the  
s also tend to

ngth) of the  
limit on the  
at most real  
t contain all  
quencies.

occupied by  
are to be no  
e remaining  
information  
miraculously  
e former. An  
sion of the  
, the coder  
shows how




**Figure 1.6** The Huffman code achieves compression by allocating short codes to frequent values. To aid deserializing the short codes are not prefixes of longer codes.

The Huffman code is used in conjunction with a number of compression techniques and is shown in Figure 1.6.

The input or *source* codes are assembled in order of descending probability. The two lowest probabilities are distinguished by a single code bit and their probabilities are combined. The process of combining probabilities is continued until unity is reached and at each stage a bit is used to distinguish the path. The bit will be a zero for the most probable path and one for the least. The compressed output is obtained by reading the bits which describe which path to take going from right to left.

In the case of computer data, there is no control over the data statistics. Data to be recorded could be instructions, images, tables, text files and so on; each having their own code value distributions. In this case a coder relying on fixed source statistics will be completely inadequate. Instead a system is used which can learn the statistics as it goes along. The Lempel-Ziv-Welch (LZW) lossless codes are in this category. These codes build up a conversion table between frequent long source data strings and short transmitted data codes at both coder and decoder and initially their compression factor is below unity as the contents of the conversion tables are transmitted along with the data. However, once the tables are established, the coding gain more than compensates for the initial loss. In some applications, a continuous analysis of the frequency of code selection is made and if a data string in the table is no longer being used with sufficient frequency it can be deselected and a more common string substituted.



Lossless codes are less common for audio and video coding where perceptive codes are permissible. The perceptive codes often obtain a coding gain by shortening the wordlength of the data representing the signal waveform. This must increase the noise level and the trick is to ensure that the resultant noise is placed at frequencies where human senses are least able to perceive it. As a result although the received signal is measurably different from the source data, it can *appear* the same to the human listener or viewer at moderate compressions factors. As these codes rely on the characteristics of human sight and hearing, they can only be fully tested subjectively.

The compression factor of such codes can be set at will by choosing the wordlength of the compressed data. Whilst mild compression will be undetectable, with greater compression factors, artifacts become noticeable. Figure 1.5 shows that this is inevitable from entropy considerations.

## 1.7 Video compression

Video signals exist in four dimensions: these are the attributes of the pixel, the horizontal and vertical spatial axes and the time axis. Compression can be applied in any or all of those four dimensions. MPEG assumes an eight-bit colour difference signal as the input, requiring rounding if the source is ten-bit. The sampling rate of the colour signals is less than that of the luminance. This is done by downsampling the colour samples horizontally and generally vertically as well. Essentially an MPEG system has three parallel simultaneous channels, one for luminance and two colour difference, which after coding are multiplexed into a single bitstream.

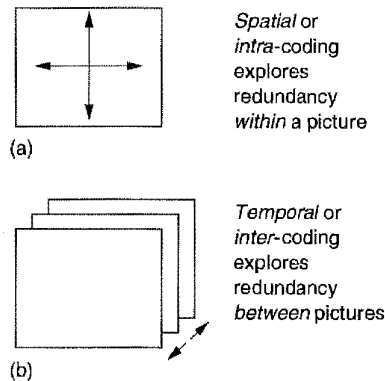


Figure 1.7 (a) Spatial or intra-coding works on individual images. (b) Temporal or inter-coding works on successive images.

Figure 1.7(a) shows that when individual pictures are compressed without reference to any other pictures, the time axis does not enter the process which is therefore described as *intra-coded* (intra = within) compression. The term *spatial coding* will also be found. It is an advantage of intra-coded video that there is no restriction to the editing which can be carried out on the picture sequence. As a result compressed VTRs such as Digital Betacam, DVC and D-9 use spatial coding. Cut editing may take place on the compressed data directly if necessary. As spatial coding treats each picture independently, it can employ certain techniques developed for the compression of still pictures. The ISO JPEG (Joint Photographic Experts Group) compression standards<sup>6,7</sup> are in this category. Where a succession of JPEG coded images are used for television, the term 'Motion JPEG' will be found.

Greater compression factors can be obtained by taking account of the redundancy from one picture to the next. This involves the time axis, as Figure 1.7(b) shows, and the process is known as *inter-coded* (inter = between) or *temporal* compression.

Temporal coding allows a higher compression factor, but has the disadvantage that an individual picture may exist only in terms of the differences from a previous picture. Clearly editing must be undertaken with caution and arbitrary cuts simply cannot be performed on the MPEG bitstream. If a previous picture is removed by an edit, the difference data will then be insufficient to re-create the current picture.

### 1.7.1 Intra-coded compression

Intra-coding works in three dimensions on the horizontal and vertical spatial axes and on the sample values. Analysis of typical television pictures reveals that whilst there is a high spatial frequency content due to detailed areas of the picture, there is a relatively small amount of energy at such frequencies. Often pictures contain sizeable areas in which the same or similar pixel values exist. This gives rise to low spatial frequencies. The average brightness of the picture results in a substantial zero frequency component. Simply omitting the high-frequency components is unacceptable as this causes an obvious softening of the picture.

A coding gain can be obtained by taking advantage of the fact that the amplitude of the spatial components falls with frequency. It is also possible to take advantage of the eye's reduced sensitivity to noise in high spatial frequencies. If the spatial frequency spectrum is divided into frequency bands the high-frequency bands can be described by fewer bits not only because their amplitudes are smaller but also because more noise can be tolerated. The wavelet transform (MPEG-4 only) and the discrete cosine transform used in JPEG and MPEG-1, MPEG-2 and MPEG-4 allow

two-dimensional pictures to be described in the frequency domain and these are discussed in Chapter 3.

### 1.7.2 Inter-coded compression


Inter-coding takes further advantage of the similarities between successive pictures in real material. Instead of sending information for each picture separately, inter-coders will send the difference between the previous picture and the current picture in a form of differential coding. Figure 1.8 shows the principle. A picture store is required at the coder to allow comparison to be made between successive pictures and a similar store is required at the decoder to make the previous picture available. The difference data may be treated as a picture itself and subjected to some form of transform-based spatial compression.

The simple system of Figure 1.8(a) is of limited use as in the case of a transmission error, every subsequent picture would be affected. Channel switching in a television set would also be impossible. In practical systems a modification is required. One approach is the so-called 'leaky predictor' in which the next picture is predicted from a limited number of previous pictures rather than from an indefinite number. As a result errors cannot propagate indefinitely. The approach used in MPEG is that periodically some absolute picture data are transmitted in place of difference data.

Figure 1.8(b) shows that absolute picture data, known as *I* or *intra* pictures are interleaved with pictures which are created using difference data, known as *P* or *predicted* pictures. The *I* pictures require a large amount of data, whereas the *P* pictures require fewer data. As a result the instantaneous data rate varies dramatically and buffering has to be used to allow a constant transmission rate. The leaky predictor needs less buffering as the compression factor does not change so much from picture to picture.

The *I* picture and all of the *P* pictures prior to the next *I* picture are called a group of pictures (GOP). For a high compression factor, a large number of *P* pictures should be present between *I* pictures, making a long GOP. However, a long GOP delays recovery from a transmission error. The compressed bitstream can only be edited at *I* pictures as shown.

In the case of moving objects, although their appearance may not change greatly from picture to picture, the data representing them on a fixed sampling grid will change and so large differences will be generated between successive pictures. It is a great advantage if the effect of motion can be removed from difference data so that they only reflect the changes in appearance of a moving object since a much greater coding gain can then be obtained. This is the objective of motion compensation.



**Figure 1.8** An inter-coded system (a) uses a delay to calculate the pixel differences between successive pictures. To prevent error propagation, intra-coded pictures (b) may be used periodically.

### 1.7.3 Introduction to motion compensation

In real television program material objects move around before a fixed camera or the camera itself moves. Motion compensation is a process which effectively measures motion of objects from one picture to the next so that it can allow for that motion when looking for redundancy between pictures. Figure 1.9 shows that moving pictures can be expressed in a three-dimensional space which results from the screen area moving along the time axis. In the case of still objects, the only motion is along the time




Figure 1.9 Objects travel in a three-dimensional space along the optic flow axis which is only parallel to the time axis if there is no movement.

axis. However, when an object moves, it does so along the *optic flow axis* which is not parallel to the time axis. The optic flow axis is the locus of a point on a moving object as it takes on various screen positions.

It will be clear that the data values representing a moving object change with respect to the time axis. However, looking along the optic flow axis the appearance of an object only changes if it deforms, moves into shadow or rotates. For simple translational motions the data representing an object are highly redundant with respect to the optic flow axis. Thus if the optic flow axis can be located, coding gain can be obtained in the presence of motion.

A motion-compensated coder works as follows. A reference picture is sent, but is also locally stored so that it can be compared with another picture to find motion vectors for various areas of the picture. The reference picture is then shifted according to these vectors to cancel inter-picture motion. The resultant *predicted* picture is compared with the actual picture to produce a *prediction error* also called a *residual*. The prediction error is transmitted with the motion vectors. At the receiver the reference picture is also held in a memory. It is shifted according to the transmitted motion vectors to re-create the predicted picture and then the prediction error is added to it to re-create the original.

In prior compression schemes the predicted picture followed the reference picture. In MPEG this is not the case. Information may be brought back from a later picture or forward from an earlier picture as appropriate.

1.7.4 *Film-originated video compression*

Film can be used as the source of video signals if a telecine machine is used. The most common frame rate for film is 24 Hz, whereas the field rates of television are 50 Hz and 60 Hz. This incompatibility is patched over in two different ways. In 50 Hz telecine, the film is simply played slightly too fast so that the frame rate becomes 25 Hz. Then each frame is converted into two television fields giving the correct 50 Hz field rate. In 60 Hz telecine, the film travels at the correct speed, but alternate frames are used to produce two fields then three fields. The technique is known as 3:2 pulldown. In this way two frames produce five fields and so the correct 60 Hz field rate results. The motion portrayal of telecine is not very good as moving objects judder, especially in 60 Hz systems. Figure 1.10 shows how the optic flow is portrayed in film-originated video.

When film-originated video is input to a compression system, the disturbed optic flow will play havoc with the motion-compensation system. In a 50 Hz system there appears to be no motion between the two fields which have originated from the same film frame, whereas between the next two fields large motions will exist. In 60 Hz systems, the motion will be zero for three fields out of five.

With such inputs, it is more efficient to adopt a different processing mode which is based upon the characteristics of the original film. Instead of attempting to manipulate fields of video, the system de-interlaces pairs of fields in order to reconstruct the original film frames. This can be done by a fairly simple motion detector. When substantial motion is measured




Figure 1.10 Telecine machines must use 3:2 pulldown to produce 60 Hz field rate video.



between successive fields in the output of a telecine, this is taken to mean that the fields have come from different film frames. When negligible motion is detected between fields, this is taken to indicate that the fields have come from the same film frame.

In 50 Hz video it is quite simple to find the sequence and produce de-interlaced frames at 25 Hz. In 60 Hz 3:2 pulldown video the problem is slightly more complex because it is necessary to locate the frames in which three fields are output so that the third field can be discarded, leaving, once more, de-interlaced frames at 25 Hz. Whilst it is relatively straightforward to lock-on to the 3:2 sequence with direct telecine output signals, if the telecine material has been edited on videotape the 3:2 sequence may contain discontinuities. In this case it is necessary to provide a number of field stores in the de-interlace unit so that a series of fields can be examined to locate the edits. Once telecine video has been de-interlaced back to frames, intra- and inter-coded compression can be employed using frame-based motion compensation.

MPEG transmissions include flags which tell the decoder the origin of the material. Material originating at 24 Hz but converted to interlaced video does not have the motion attributes of interlace because the lines in two fields have come from the same point on the time axis. Two fields can be combined to create a progressively scanned frame. In the case of 3:2 pulldown material, the third field need not be sent at all as the decoder can easily repeat a field from memory. As a result the same compressed film material can be output at 50 or 60 Hz as required.

Recently conventional telecine machines have been superseded by the *datacine* which scans each film frame into a pixel array which can be made directly available to the MPEG encoder without passing through an intermediate digital video standard. Datacines are used extensively for mastering DVDs from film stock.

## 1.8 Introduction to MPEG-1


As mentioned above, the intention of MPEG-1 is to deliver video and audio at the same bit rate as a conventional audio CD. As the bit rate was a given, this was achieved by subsampling to half the definition of conventional television. In order to have a constant input bit rate irrespective of the frame rate, 25 Hz systems have a picture size of  $352 \times 288$  pixels whereas 30 Hz systems have a picture size of  $352 \times 240$  pixels. This is known as *common intermediate format* (CIF). If the input is conventional interlaced video, CIF can be obtained by discarding alternate fields and downsampling the remaining active lines by a factor of two.

ie machine is  
reas the field  
ty is patched  
imply played  
each frame is  
field rate. In  
ernate frames  
que is known  
ds and so the  
elecine is not  
stems. Figure  
ted video.

system, the  
compensation  
ween the two  
reas between  
is, the motion

nt processing  
film. Instead  
terlaces pairs  
s can be done  
is measured

► Time



field rate video.

As interlaced systems have very poor vertical resolution, down-sampling to CIF actually does little damage to still images, although the very low picture rates damage motion portrayal.

Although MPEG-1 appeared rather rough on screen, this was due to the very low bit rate. It is more important to appreciate that MPEG-1 introduced the great majority of the coding tools which would continue to be used in MPEG-2 and MPEG-4. These included an elementary stream syntax, bidirectional motion-compensated coding, buffering and rate control. Many of the spatial coding principles of MPEG-1 were taken from JPEG. MPEG-1 also specified audio compression of up to two channels.

## 1.9 MPEG-2: Profiles and Levels

MPEG-2 builds upon MPEG-1 by adding interlace capability as well as a greatly expanded range of picture sizes and bit rates. The use of scaleable systems is also addressed, along with definitions of how multiple MPEG bitstreams can be multiplexed. As MPEG-2 is an extension of MPEG-1, it is easy for MPEG-2 decoders to handle MPEG-1 data. In a sense an MPEG-1 bitstream is an MPEG-2 bitstream which has a restricted vocabulary and so can be readily understood by an MPEG-2 decoder.

MPEG-2 has too many applications to solve with a single standard and so it is subdivided into Profiles and Levels. Put simply a Profile describes a degree of complexity whereas a Level describes the picture size or resolution which goes with that Profile. Not all Levels are supported at all Profiles. Figure 1.11 shows the available combinations. In principle there

		Profiles					
		Simple	Main	4:2:2	SNR	Spatial	High
Levels	High		4:2:0 1920 × 1152 90 Mb/S				4:2:0 or 4:2:2 1920 × 1152 100 Mb/S
	High 1440		4:2:0 1440 × 1152 60 Mb/S			4:2:0 1440 × 1152 60 Mb/S	4:2:0 or 4:2:2 1440 × 1152 80 Mb/S
	Main	4:2:0 720 × 576 15 Mb/S NO B	4:2:0 720 × 576 15 Mb/S	4:2:2 720 × 608 50 Mb/S	4:2:0 720 × 576 15 Mb/S		4:2:0 or 4:2:2 720 × 576 20 Mb/S
	Low		4:2:0 352 × 288 4 Mb/S		4:2:0 352 × 288 4 Mb/S		

Figure 1.11 Profiles and Levels in MPEG-2. See text for details.

tion, down-  
although the

was due to  
that MPEG-1  
uld continue  
ntary stream  
ng and rate  
were taken  
f up to two

as well as a  
e of scaleable  
multiple MPEG  
f MPEG-1, it  
a sense an  
a restricted  
2 decoder.

standard and  
file describes  
ature size or  
ported at all  
inciples there

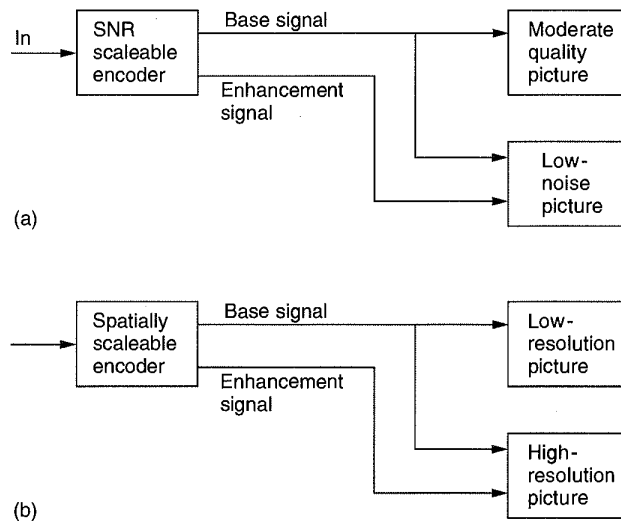
are twenty-four of these, but not all have been defined. An MPEG-2 decoder having a given Profile and Level must also be able to decode lower Profiles and Levels.

The simple Profile does not support bidirectional coding and so only *I* and *P* pictures will be output. This reduces the coding and decoding delay and allows simpler hardware. The simple Profile has only been defined at Main Level (SPML).

The Main Profile is designed for a large proportion of uses. The Low Level uses a low resolution input having only 352 pixels per line. The majority of broadcast applications will require the MPML (Main Profile at Main Level) subset of MPEG which supports SDTV (standard definition television). The High-1440 Level is a high-definition scheme which doubles the definition compared to Main Level. The High Level not only doubles the resolution but maintains that resolution with 16:9 format by increasing the number of horizontal samples from 1440 to 1920.

In compression systems using spatial transforms and requantizing it is possible to produce scaleable signals. A scaleable process is one in which the input results in a main signal and a 'helper' signal. The main signal can be decoded alone to give a picture of a certain quality, but if the information from the helper signal is added some aspect of the quality can be improved.

Figure 1.12(a) shows that in a conventional MPEG coder, by heavily requantizing coefficients a picture with moderate signal-to-noise ratio



**Figure 1.12** (a) An SNR scaleable encoder produces a 'noisy' signal and a noise cancelling signal. (b) A spatially scaleable encoder produces a low-resolution picture and a resolution-enhancing picture.

High
4:2:0 or 4:2:2
320 × 1152 100 Mb/S
4:2:0 or 4:2:2
440 × 1152 80 Mb/S
4:2:0 or 4:2:2
720 × 576 20 Mb/S

results. If, however, that picture is locally decoded and subtracted pixel by pixel from the original, a 'quantizing noise' picture would result. This can be compressed and transmitted as the helper signal. A simple decoder only decodes the main 'noisy' bitstream, but a more complex decoder can decode both bitstreams and combine them to produce a low-noise picture. This is the principle of SNR scalability.

As an alternative, Figure 1.12(b) shows that by coding only the lower spatial frequencies in a HDTV picture a base bitstream can be made which an SDTV receiver can decode. If the lower definition picture is locally decoded and subtracted from the original picture, a 'definition-enhancing' picture would result. This can be coded into a helper signal. A suitable decoder could combine the main and helper signals to recreate the HDTV picture. This is the principle of spatial scalability.

The High Profile supports both SNR and spatial scalability as well as allowing the option of 4:2:2 sampling (see Section 2.11).

The 4:2:2 Profile has been developed for improved compatibility with existing digital television production equipment. This allows 4:2:2 working without requiring the additional complexity of using the High Profile. For example a HP ML decoder must support SNR scalability which is not a requirement for production.

MPEG-2 increased the number of audio channels possible to five whilst remaining compatible with MPEG-1 audio. MPEG-2 subsequently introduced a more efficient audio coding scheme known as MPEG-2 AAC (advanced audio coding) which is not backwards compatible with the earlier audio coding schemes.

## 1.10 Introduction to MPEG-4

MPEG-4 introduces a number of new coding tools as shown in Figure 1.13. In MPEG-1 and MPEG-2 the motion compensation is based on regular fixed-size areas of image known as *macroblocks*. Whilst this works well at the designed bit rates, there will always be some inefficiency due to real moving objects failing to align with macroblock boundaries. This will increase the residual bit rate. In MPEG-4, moving objects can be coded as arbitrary shapes. Figure 1.14 shows that a background can be coded quite independently from objects in front of it. Object motion can then be described with vectors and much-reduced residual data. According to the Profile, objects may be two-dimensional, three-dimensional and opaque or translucent. The decoder must contain effectively a layering vision mixer which is capable of prioritizing image data as a function of how close it is to the viewer. The picture coding of MPEG-4 is known as texture coding and is more advanced than the

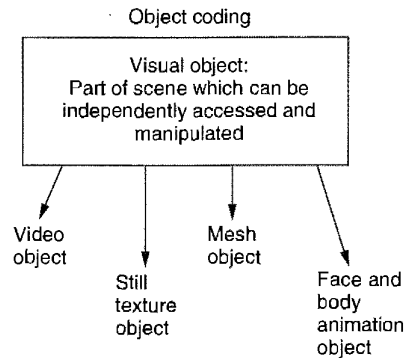
racted pixel  
result. This  
. A simple  
re complex  
duce a low-

y the lower  
n be made  
n picture is  
'definition-  
lper signal.  
gnals to re-  
eability.  
y as well as

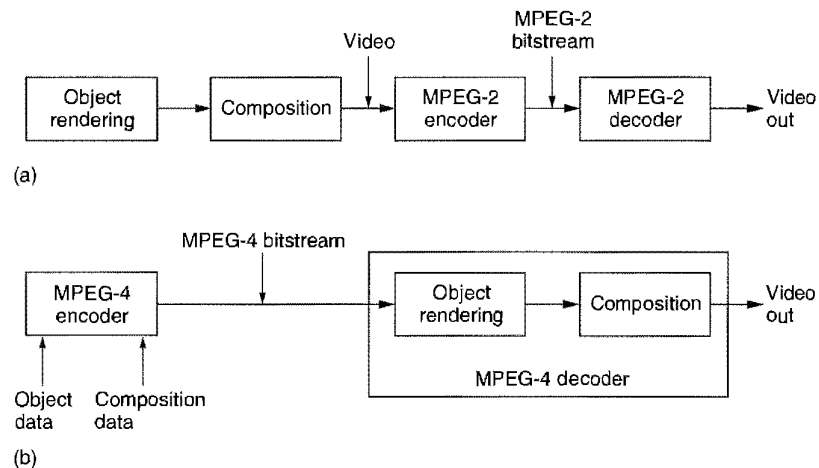
tibility with  
llows 4:2:2  
ng the High  
scaleability

o five whilst  
iently intro-  
PEG-2 AAC  
le with the

n in Figure  
s based on  
t this works  
iciency due  
daries. This  
ects can be  
und can be  
motion can  
idual data.  
mal, three-  
ust contain  
izing image  
e coding of  
d than the



**Figure 1.13** MPEG-4 introduces a number of new coding tools over those of earlier MPEG standards. These include object coding, mesh coding, still picture coding and face and body animation.



**Figure 1.14** (a) In MPEG-1 and MPEG-2, computer graphic images must be rendered to video before coding. (b) In contrast, MPEG-4 may move the rendering process to the decoder, reducing the bit rate needed with the penalty of increased decoder complexity.

MPEG-2 equivalent, using more lossless predictive coding for pixel values, coefficients and vectors.

In addition to motion compensation, MPEG-4 can describe how an object changes its perspective as it moves using a technique called mesh coding. By warping another image, the prediction of the present image is improved. MPEG-4 also introduces coding for still images using DCT or wavelets.

Although MPEG-2 supported some scaleability, MPEG-4 also takes this further. In addition to spatial and noise scaleability, MPEG-4 also allows

temporal scalability where a base level bitstream having a certain frame rate may be augmented by an additional enhancement bitstream to produce a decoder output at a higher frame rate. This is important as it allows a way forward from the marginal frame rates of today's film and television formats whilst remaining backwards compatible with traditional equipment. The comprehensive scalability of MPEG-4 is equally important in networks where it allows the user the best picture possible for the available bit rate.

MPEG-4 also introduces standards for face and body animation. Specialized vectors allow a still picture of a face and optionally a body to be animated to allow expressions and gestures to accompany speech at very low bit rates. In some senses MPEG-4 has gone upstream of the video signal which forms the input to MPEG-1 and MPEG-2 coders to analyse ways in which the video signal was rendered. Figure 1.14(a) shows that in a system using MPEG-1 and MPEG-2, all rendering and production steps take place before the encoder. Figure 1.14(b) shows that in MPEG-4, some of these steps can take place in the decoder. The advantage is that fewer data need to be transmitted. Some of these data will be rendering instructions which can be very efficient and result in a high compression factor. As a significant part of the rendering takes place in the decoder, computer graphics generators can be designed directly to output an MPEG-4 bitstream. In interactive systems such as simulators and video games, inputs from the user can move objects around the screen. The disadvantage is increased decoder complexity, but as the economics of digital processing continues to advance this is hardly a serious concern.

As might be expected, the huge range of coding tools in MPEG-4 is excessive for many applications. As with MPEG-2 this has been dealt with using Profiles and Levels. Figure 1.15 shows the range of Visual

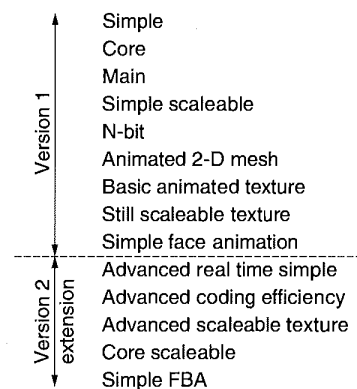


Figure 1.15 The visual object types supported by MPEG-4 in versions 1 and 2.

certain frame  
bitstream to  
portant as it  
y's film and  
with tradi-  
-4 is equally  
ure possible

animation.  
lly a body to  
ny speech at  
ream of the  
-2 coders to  
gure 1.14(a)  
ndering and  
) shows that  
ecoder. The  
f these data  
d result in a  
g takes place  
d directly to  
s simulators  
around the  
, but as the  
is hardly a

MPEG-4 is  
s been dealt  
ge of Visual

nd 2.

Version 1 visual profiles

	Object types	Simple	Core	Main	Simple scaleable	N-bit	Animated 2D mesh	Basic animated texture	Scaleable texture	Simple face
	Profiles									
1.	Simple	X								
2.	Simple scaleable	X			X					
3.	Core	X	X							
4.	Main	X	X	X					X	
5.	N-bit	X	X			X				
6.	Hybrid	X	X				X	X	X	X
7.	Basic animated texture							X	X	X
8.	Scaleable texture								X	
9.	Simple FA									X

Version 2 visual profiles

	Object types	Simple	Simple scaleable	Core	Core scaleable	Advanced real time simple	Advanced coding efficiency	Advanced scalable texture	Simple FBA
	Profiles								
V2-1.	Advanced real time simple	X				X			
V2-2.	Core scaleable	X	X	X	X				
V2-3.	Advanced coding efficiency	X		X			X		
V2-4.	Advanced core	X		X				X	
V2-5.	Advanced scalable texture							X	
V2-6.	Simple FBA								X

Figure 1.16 The visual object types supported by each visual profile of MPEG-4.

Object types in version 1 of MPEG-4 and as expanded in version 2. For each visual object type the coding tools needed is shown. Figure 1.16 shows the relationship between the Visual Profiles and the Visual Object types supported by each Profile. The crossover between computer-generated and natural images is evident in the Profile structure where Profiles 1–5 cover natural images, Profiles 8 and 9 cover rendered images and Profiles 6 and 7 cover hybrid natural/rendered images. It is only possible to give an introduction here and more detail is provided in Chapter 5. MPEG-4 also extends the boundaries of audio coding. The MPEG-2 AAC technique is extended in MPEG-4 by some additional tools. New tools are added which allow operation at very low bit rates for speech applications. Also introduced is the concept of *structured audio* in which the audio waveform is synthesized at the decoder from a bitstream which is essentially a digital musical score.

## 1.11 Audio compression

Perceptive coding in audio relies on the principle of auditory masking, which is treated in detail in section 4.1. Masking causes the ear/brain combination to be less sensitive to sound at one frequency in the presence of another at a nearby frequency. If a first tone is present in the input, then it will mask signals of lower level at nearby frequencies. The quantizing of the first tone and of further tones at those frequencies can be made coarser. Fewer bits are needed and a coding gain results. The increased quantizing error is allowable if it is masked by the presence of the first tone.

### 1.11.1 *Sub-band coding*

Sub-band coding mimics the frequency analysis mechanism of the ear and splits the audio spectrum into a large number of different bands. Signals in these bands can then be quantized independently. The quantizing error which results is confined to the frequency limits of the band and so it can be arranged to be masked by the program material. The techniques used in Layers I and II of MPEG audio are based on sub-band coding as are those used in DCC (Digital Compact Cassette).

### 1.11.2 *Transform coding*

In transform coding the time-domain audio waveform is converted into a frequency domain representation such as a Fourier, discrete cosine or wavelet transform (see Chapter 3). Transform coding takes advantage of the fact that the amplitude or envelope of an audio signal changes relatively slowly and so the coefficients of the transform can be transmitted relatively infrequently. Clearly such an approach breaks down in the presence of transients and adaptive systems are required in practice. Transients cause the coefficients to be updated frequently whereas in stationary parts of the signal such as sustained notes the update rate can be reduced. Discrete cosine transform (DCT) coding is used in Layer III of MPEG audio and in the compression system of the Sony MiniDisc.

### 1.11.3 *Predictive coding*

In a predictive coder there are two identical predictors, one in the coder and one in the decoder. Their job is to examine a run of previous data



values and to extrapolate forward to estimate or predict what the next value will be. This is subtracted from the *actual* next code value at the encoder to produce a prediction error which is transmitted. The decoder then adds the prediction error to its own prediction to obtain the output code value again.

Prediction can be used in the time domain, where sample values are predicted, or in the frequency domain where coefficient values are predicted. Time-domain predictive coders work with a short encode and decode delay and are useful in telephony where a long loop delay causes problems. Frequency prediction is used in AC-3 and MPEG AAC.

## 1.12 MPEG bitstreams

MPEG supports a variety of bitstream types for various purposes and these are shown in Figure 1.17. The output of a single compressor (video or audio) is known as an elementary stream. In transmission, many elementary streams will be combined to make a transport stream. Multiplexing requires blocks or packets of constant size. It is advantageous if these are short so that each elementary stream in the multiplex can receive regular data. A transport stream has a complex structure because it needs to incorporate metadata indicating which audio

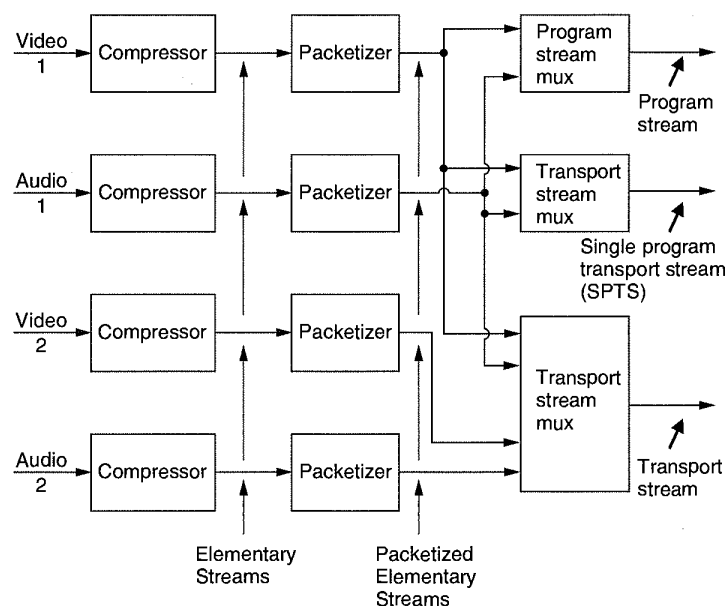


Figure 1.17 The bitstream types of MPEG-2. See text for details.

elementary streams and ancillary data are associated with which video elementary stream. It is possible to have a single program transport stream (SPTS) which carries only the elementary streams of one TV program.

For certain purposes, such as recording a single elementary stream, the transport stream is not appropriate. The small packets of the transport stream each require a header and this wastes storage space. In this case a program stream can be used. A program stream is a simplified bitstream which multiplexes audio and video for a single program together, provided they have been encoded from a common locked clock. Unlike a transport stream, the blocks are larger and are not necessarily of fixed size.

### 1.13 Drawbacks of compression

By definition, compression removes redundancy from signals. Redundancy is, however, essential to making data resistant to errors. As a result, compressed data are more sensitive to errors than uncompressed data. Thus transmission systems using compressed data must incorporate more powerful error-correction strategies and avoid compression techniques which are notoriously sensitive. As an example, the Digital Betacam format uses relatively mild compression and yet requires 20 per cent redundancy whereas the D-5 format does not use compression and only requires 17 per cent redundancy even though it has a recording density 30 per cent higher. Techniques using tables such as the Lempel-Ziv-Welch codes are very sensitive to bit errors as an error in the transmission of a table value results in bit errors every time that table location is accessed. This is known as error propagation. Variable-length techniques such as the Huffman code are also sensitive to bit errors. As there is no fixed symbol size, the only way the decoder can parse a serial bitstream into symbols is to increase the assumed wordlength a bit at a time until a code value is recognized. The next bit must then be the first bit in the next symbol. A single bit in error could cause the length of a code to be wrongly assessed and then all subsequent codes would also be wrongly decoded until synchronization could be re-established. Later variable-length codes sacrifice some compression efficiency in order to offer better resynchronization properties.

In non-real-time systems such as computers an uncorrectable error results in reference to the back-up media. In real-time systems such as audio and video this is impossible and concealment must be used. However, concealment relies on redundancy and compression reduces the degree of redundancy. Media such as hard disks can be verified so that uncorrectable errors are virtually eliminated, but tape is prone to

in which video  
ram transport  
ns of one TV

try stream, the  
the transport  
. In this case a  
ified bitstream  
ram together,  
clock. Unlike a  
sarily of fixed

gnals. Redun-  
rs. As a result,  
pressed data.  
it incorporate  
pression tech-  
e, the Digital  
equires 20 per  
pression and  
s a recording  
the Lempel-  
error in the  
me that table  
variable-length  
bit errors. As  
parse a serial  
length a bit at a  
en be the first  
ie length of a  
would also be  
lished. Later  
y in order to

rectable error  
tems such as  
ust be used.  
ssion reduces  
e verified so  
e is prone to

dropouts which will exceed the burst-correcting power of the replay system from time to time. For this reason the compression factors used on audio or video tape should be moderate.

As perceptive coders introduce noise, it will be clear that in a concatenated system the second codec could be confused by the noise due to the first. If the codecs are identical then each may well make, or better still be designed to make, the same decisions when they are in tandem. If the codecs are not identical the results could be disappointing. Signal manipulation between codecs can also result in artifacts which were previously undetectable becoming visible because the signal which was masking them is no longer present.

In general, compression should not be used for its own sake, but only where a genuine bandwidth or cost bottleneck exists. Even then the mildest compression possible should be used. Whilst high compression factors are permissible for final delivery of material to the consumer, they are not advisable prior to any post-production stages. For contribution material, lower compression factors are essential and this is sometimes referred to as *mezzanine* level compression.

One practical drawback of compression systems is that they are largely generic in structure and the same hardware can be operated at a variety of compression factors. Clearly the higher the compression factor, the cheaper the system will be to operate so there will be economic pressure to use high compression factors. Naturally the risk of artifacts is increased and so there is (or should be) counterpressure from those with engineering skills to moderate the compression. The way of the world at the time of writing is that the accountants have the upper hand. This was not a problem when there were fixed standards such as PAL and NTSC, as there was no alternative but to adhere to them. Today there is plenty of evidence that the variable compression factor control is being turned too far in the direction of economy.

It has been seen above that concatenation of compression systems should be avoided as this causes generation loss. Generation loss is worse if the codecs are different. Interlace is a legacy compression technique and if concatenated with MPEG, generation loss will be exaggerated. In theory and in practice better results are obtained in MPEG for the same bit rate if the input is progressively scanned. Consequently the use of interlace with MPEG coders cannot be recommended for new systems. Chapter 5 explores this theme in greater detail.

## 1.14 Compression pre-processing

Compression relies completely on identifying redundancy in the source material. Consequently anything which reduces that redundancy will

have a damaging effect. Noise is particularly undesirable as it creates additional spatial frequencies in individual pictures as well as spurious differences between pictures. Where noisy source material is anticipated some form of noise reduction will be essential.

When high compression factors must be used to achieve a low bit rate, it is inevitable that the level of artifacts will rise. In order to contain the artifact level, it is necessary to restrict the source entropy prior to the coder. This may be done by spatial low-pass filtering to reduce the picture resolution, and may be combined with downsampling to reduce the number of pixels per picture. In some cases, such as teleconferencing, it will also be necessary to reduce the picture rate. At very low bit rates the use of interlace becomes acceptable as a pre-processing stage providing downsampling prior to the MPEG compression.

A compression pre-processor will combine various types of noise reduction (see Chapter 3) with spatial and temporal downsampling.

### 1.15 Some guidelines

Although compression techniques themselves are complex, there are some simple rules which can be used to avoid disappointment. Used wisely, MPEG compression has a number of advantages. Used in an inappropriate manner, disappointment is almost inevitable and the technology could get a bad name. The next few points are worth remembering.

- Compression technology may be exciting, but if it is not necessary it should not be used.
- If compression is to be used, the degree of compression should be as small as possible; i.e. use the highest practical bit rate.
- Cascaded compression systems cause loss of quality and the lower the bit rates, the worse this gets. Quality loss increases if any post-production steps are performed between compressions.
- Avoid using interlaced video with MPEG.
- Compression systems cause delay.
- Compression systems work best with clean source material. Noisy signals or poorly decoded composite video give poor results.
- Compressed data are generally more prone to transmission errors than non-compressed data. The choice of a compression scheme must consider the error characteristics of the channel.
- Audio codecs need to be level calibrated so that when sound pressure level-dependent decisions are made in the coder those levels actually exist at the microphone.

- Low bit rate coders should only be used for the final delivery of post-produced signals to the end-user.
- Don't believe statements comparing codec performance to 'VHS quality' or similar. Compression artifacts are quite different from the artifacts of consumer VCRs.
- Quality varies wildly with source material. Beware of 'convincing' demonstrations which may use selected material to achieve low bit rates. Use your own test material, selected for a balance of difficulty.
- Don't be browbeaten by the technology. You don't have to understand it to assess the results. Your eyes and ears are as good as anyone's so don't be afraid to criticize artifacts. In the case of video, use still frames to distinguish spatial artifacts from temporal ones.

### References

1. ISO/IEC JTC1/SC29/WG11 MPEG, International standard ISO 11172, Coding of moving pictures and associated audio for digital storage media up to 1.5 Mbits/s (1992)
2. LeGall, D., MPEG: a video compression standard for multimedia applications. *Communications of the ACM*, 34, No.4, 46–58 (1991)
3. MPEG-2 Video Standard: ISO/IEC 13818–2: Information technology – generic coding of moving pictures and associated audio information: Video (1996) (aka ITU-T Rec. H-262 (1996))
4. MPEG-4 Standard: ISO/IEC 14496–2: Information technology – coding of audio-visual objects: Amd.1 (2000)
5. Huffman, D.A. A method for the construction of minimum redundancy codes. *Proc. IRE*, 40 1098–1101 (1952)
6. ISO Joint Photographic Experts Group standard JPEG-8-R8
7. Wallace, G.K., Overview of the JPEG (ISO/CCITT) still image compression standard. ISO/JTC1/SC2/WG8 N932 (1989)

# 4

## Audio compression

---

### 4.1 Introduction

Physics can tell us the mechanism by which disturbances propagate through the air. If this is our definition of sound, we have the problem that in physics there are no limits to the frequencies and levels which must be considered. Biology can tell that the ear only responds to a certain range of frequencies provided a threshold level is exceeded. This is a better definition of sound; reproduction is easier because it is only necessary to reproduce that range of levels and frequencies which the ear can detect.


Psychoacoustics can describe how our hearing has finite resolution in time, frequency and spatial domains such that what we perceive is an inexact impression. Some aspects of the original disturbance are inaudible to us and are said to be masked. If our goal is the highest quality, we can design our imperfect equipment so that the shortcomings are masked. Conversely if our goal is economy we can use compression and hope that masking will disguise the inaccuracies it causes.

By definition, the sound quality of a perceptive coder can only be assessed by human hearing. Equally, a useful perceptive coder can only be designed with a good knowledge of the human hearing mechanism.<sup>1</sup> The acuity of the human ear is astonishing. The frequency range is extremely wide, covering some ten octaves (an octave is a doubling of pitch or frequency) without interruption. It can detect tiny amounts of distortion, and will accept an enormous dynamic range. If the ear detects a different degree of impairment between two codecs having the same bit rate in properly conducted tests, we can say that one of them is superior.

Quality is completely subjective and can only be checked by listening tests, although these are meaningless if the loudspeakers are not of sufficient quality. However, any characteristic of a signal which can be heard can also be measured by a suitable instrument. The subjective tests can tell us how sensitive the instrument should be. Then the objective readings from the instrument give an indication of how acceptable a signal is in respect of that characteristic. Instruments and loudspeakers suitable for assessing the performance of codecs are currently extremely rare and there remains much work to be done.

## 4.2 The deciBel

The first audio signals to be transmitted were on analog telephone lines. Where the wiring is long compared to the electrical wavelength (not to be confused with the acoustic wavelength) of the signal, a transmission line exists in which the distributed series inductance and the parallel capacitance interact to give the line a characteristic impedance. In telephones this turned out to be about  $600\ \Omega$ . In transmission lines the best power delivery occurs when the source and the load impedance are



Ohm's law  $V = IR$



Power  $P = IV = \frac{V^2}{R}$  ← Power goes as square of voltage

$\therefore V = \sqrt{RP}$

In the case of  $600\ \Omega/1\ \text{mW}$ :

$V = \sqrt{600 \times 0.001} = 0.7746\ \text{V}$

(a)



Average power = peak power/2

As  $P \propto V^2$   $(V_{\text{peak}})^2 = 2(V_{\text{rms}})^2$

$\therefore V_{\text{rms}} = V_{\text{peak}}/\sqrt{2}$

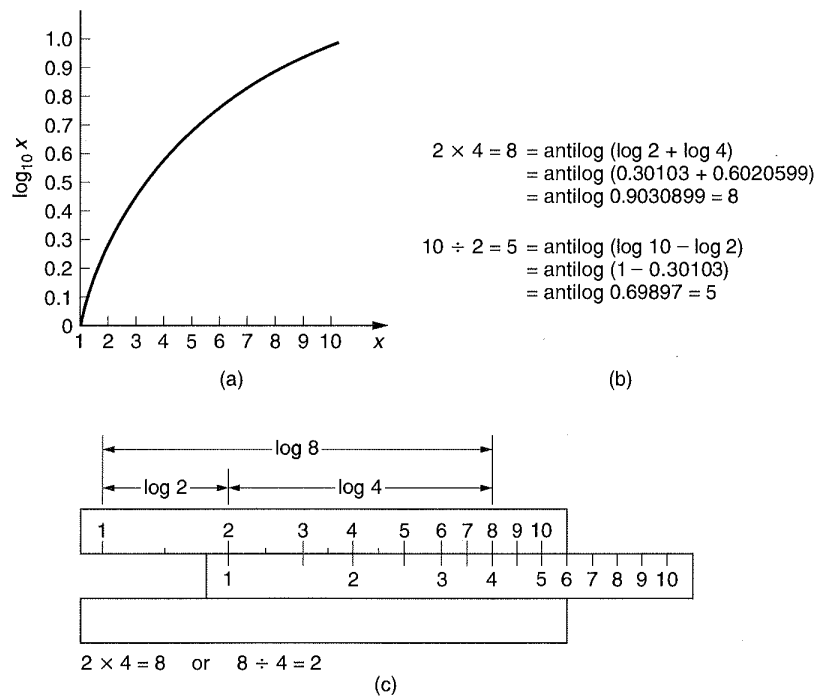
(b)

**Figure 4.1** (a) Ohm's law: the power developed in a resistor is proportional to the square of the voltage. Consequently, 1 mW in  $600\ \Omega$  requires 0.775 V. With a sinusoidal alternating input (b), the power is a sine square function which can be averaged over one cycle. A DC voltage which would deliver the same power has a value which is the square root of the mean of the square of the sinusoidal input.

the same; this is the process of matching. It was often required to measure the power in a telephone system, and 1 milliWatt was chosen as a suitable unit. Thus the reference against which signals could be compared was the dissipation of one milliWatt in  $600\ \Omega$ . Figure 4.1 shows that the dissipation of 1 mW in  $600\ \Omega$  will be due to an applied voltage of 0.775 V rms. This voltage is the reference against which all audio levels are compared.

The deciBel is a logarithmic measuring system and has its origins in telephony<sup>2</sup> where the loss in a cable is a logarithmic function of the length. Human hearing also has a logarithmic response with respect to sound pressure level (SPL). In order to relate to the subjective response audio signal level measurements have also to be logarithmic and so the deciBel was adopted for audio.

Figure 4.2 shows the principle of the logarithm. To give an example, if it is clear that  $10^2$  is 100 and  $10^3$  is 1000, then there must be a power between 2 and 3 to which 10 can be raised to give any value between 100 and 1000. That power is the logarithm to base 10 of the value. e.g.  $\log_{10} 300 = 2.5$  approx. Note that  $10^0$  is 1.



**Figure 4.2** (a) The logarithm of a number is the power to which the base (in this case 10) must be raised to obtain the number. (b) Multiplication is obtained by adding logs, division by subtracting. (c) The slide rule has two logarithmic scales whose length can easily be added or subtracted.



Logarithms were developed by mathematicians before the availability of calculators or computers to ease calculations such as multiplication, squaring, division and extracting roots. The advantage is that armed with a set of log tables, multiplication can be performed by adding, division by subtracting. Figure 4.2 shows some examples. It will be clear that squaring a number is performed by adding two identical logs and the same result will be obtained by multiplying the log by 2.

The slide rule is an early calculator which consists of two logarithmically engraved scales in which the length along the scale is proportional to the log of the engraved number. By sliding the moving scale two lengths can easily be added or subtracted and as a result multiplication and division are readily obtained.

The logarithmic unit of measurement in telephones was called the Bel after Alexander Graham Bell, the inventor. Figure 4.3(a) shows that the Bel was defined as the log of the *power* ratio between the power to be measured and some reference power. Clearly the reference power must have a level of 0 Bels since  $\log_{10} 1$  is 0.

The Bel was found to be an excessively large unit for many purposes and so it was divided into 10 deciBels, abbreviated to dB with a small d and a large B and pronounced 'deebee'. Consequently the number of dB is ten times the log of the power ratio. A device such as an amplifier can have a fixed power gain which is independent of signal level and this can

$$1 \text{ Bel} = \log_{10} \frac{P_1}{P_2} \quad 1 \text{ deciBel} = 1/10 \text{ Bel}$$

$$\text{Power ratio (dB)} = 10 \times \log_{10} \frac{P_1}{P_2}$$

(a)

As power  $\propto V^2$ , when using voltages:

$$\begin{aligned} \text{Power ratio (dB)} &= 10 \log \frac{V_1^2}{V_2^2} \\ &= 10 \times \log \frac{V_1}{V_2} \times 2 \\ &= 20 \log \frac{V_1}{V_2} \end{aligned}$$

(b)

**Figure 4.3** (a) The Bel is the log of the ratio between two powers, that to be measured and the reference. The Bel is too large so the deciBel is used in practice. (b) As the dB is defined as a power ratio, voltage ratios have to be squared. This is conveniently done by doubling the logs so the ratio is now multiplied by 20.

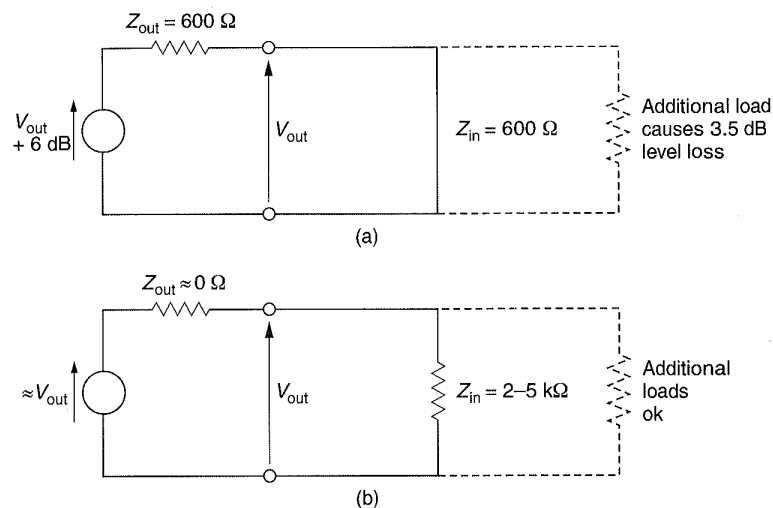
be measured in dB. However, when measuring the power of a signal, it must be appreciated that the dB is a ratio and to quote the number of dBs without stating the reference is about as senseless as describing the height of a mountain as 2000 without specifying whether this is feet or metres. To show that the reference is one milliWatt into  $600\ \Omega$ , the units will be dB(m). In radio engineering, the dB(W) will be found which is power relative to one Watt.

Although the dB(m) is defined as a power ratio, level measurements in audio are often done by measuring the signal voltage using 0.775 V as a reference in a circuit whose impedance is not necessarily  $600\ \Omega$ . Figure 4.3(b) shows that as the power is proportional to the square of the voltage, the power ratio will be obtained by squaring the voltage ratio. As squaring in logs is performed by doubling, the squared term of the voltages can be replaced by multiplying the log by a factor of two. To give a result in decibels, the log of the voltage ratio now has to be multiplied by 20.

Whilst  $600\ \Omega$  matched-impedance working is essential for the long distances encountered with telephones, it is quite inappropriate for analog audio wiring in the studio or the home. The wavelength of audio in wires at 20 kHz is 15 km. Most studios are built on a smaller scale than this and clearly analog audio cables are *not* transmission lines and their characteristic impedance is swamped by the devices connected at each end. Consequently the reader is cautioned that anyone who attempts to sell exotic analog audio cables by stressing their transmission line characteristics is more of a salesman than a physicist.


In professional analog audio systems impedance matching is not only unnecessary it is also undesirable. Figure 4.4(a) shows that when impedance matching is required the output impedance of a signal source must be artificially raised so that a potential divider is formed with the load. The actual drive voltage must be twice that needed on the cable as the potential divider effect wastes 6 dB of signal level and requires unnecessarily high power supply rail voltages in equipment. A further problem is that cable capacitance can cause an undesirable HF roll-off in conjunction with the high source impedance. In modern professional analog audio equipment, shown in Figure 4.4(b) the source has the lowest output impedance practicable. This means that any ambient interference is attempting to drive what amounts to a short circuit and can only develop very small voltages. Furthermore, shunt capacitance in the cable has very little effect. The destination has a somewhat higher impedance (generally a few  $k\Omega$ ) to avoid excessive currents flowing and to allow several loads to be placed across one driver.

In the absence of a fixed impedance it is now meaningless to consider power. Consequently only signal voltages are measured. The reference remains at 0.775 V, but power and impedance are irrelevant. Voltages




**Figure 4.4** (a) Traditional impedance matched source wastes half the signal voltage in the potential divider due to the source impedance and the cable. (b) Modern practice is to use low-output impedance sources with high-impedance loads.

measured in this way are expressed in dB(u); the commonest unit of level in modern systems. Most installations boost the signals on interface cables by 4 dB. As the gain of receiving devices is reduced by 4 dB, the result is a useful noise advantage without risking distortion due to the drivers having to produce high voltages. In order to make the difference between dB(m) and dB(u) clear, consider the lossless matching transformer shown in Figure 4.5. The turns ratio is 2:1 therefore the impedance matching ratio is 4:1. As there is no loss in the transformer, the power in is the same as the power out so that the transformer shows a gain of 0 dB(m). However, the turns ratio of 2:1 provides a voltage gain of 6 dB(u). The doubled output voltage will develop the same power into the quadrupled load impedance.



**Figure 4.5** A lossless transformer has no power gain so the level in dB(m) on input and output is the same. However, there is a voltage gain when measurements are made in dB(u).



**Figure 4.6** In complex systems each stage may have voltage gain measured in dB. By adding all of these gains together and adding to the input level in dB(u), the output level in dB(u) can be obtained.

In a complex system signals may pass through a large number of processes, each of which may have a different gain. Figure 4.6 shows that if one stays in the linear domain and measures the input level in volts rms, the output level will be obtained by multiplying by the gains of all the stages involved. This is a complex calculation.

The difference between the signal level with and without the presence of a device in a chain is called the *insertion loss* measured in dB. However, if the input is measured in dB(u), the output level of the first stage can be obtained by adding the insertion loss in dB. The output level of the second stage can be obtained by further adding the loss of the second stage in dB and so on. The final result is obtained by adding together all the insertion losses in dB and adding them to the input level in dB(u) to give the output level in dB(u). As the dB is a pure ratio it can multiply anything (by addition of logs) without changing the units. Thus dB(u) of level added to dB of gain are still dB(u).

In acoustic measurements, the sound pressure level (SPL) is measured in decibels relative to a reference pressure of  $2 \times 10^{-5}$  Pascals (Pa) rms. In order to make the reference clear the units are dB(SPL). In measurements which are intended to convey an impression of subjective loudness, a weighting filter is used prior to the level measurement which reproduces the frequency response of human hearing which is most sensitive in the midrange. The most common standard frequency response is the so-called A-weighting filter, hence the term dB(A) used when a weighted level is being measured. At high or low frequencies, a lower reading will be obtained in dB(A) than in dB(SPL).

### 4.3 Audio level metering

There are two main reasons for having level meters in audio equipment: to line up or adjust the gain of equipment, and to assess the amplitude of the program material. Gain line-up is especially important in digital systems where an incorrect analog level can result in ADC clipping.

Line-up is often done using a 1 kHz sine wave generated at an agreed level such as 0 dB(u). If a receiving device does not display the same level, then its input sensitivity must be adjusted. Tape recorders and other devices which pass signals through are usually lined up so that their input and output levels are identical, i.e. their insertion loss is 0 dB. Line-up is important in large systems because it ensures that inadvertent level changes do not occur. In measuring the level of a sine wave for the purposes of line-up, the dynamics of the meter are of no consequence, whereas on program material the dynamics matter a great deal. The simplest (and cheapest) level meter is essentially an AC voltmeter with a logarithmic response. As the ear is logarithmic, the deflection of the meter is roughly proportional to the perceived volume, hence the term Volume Unit (VU) meter.

In audio recording and broadcasting, the worst sin is to overmodulate the tape, the ADC or the transmitter by allowing a signal of excessive amplitude to pass. Real audio signals are rich in short transients which pass before the sluggish VU meter responds. Consequently the VU meter is also called the virtually useless meter in professional circles.

Broadcasters developed the Peak Program Meter (PPM) which is also logarithmic, but which is designed to respond to peaks as quickly as the ear responds to distortion. Consequently the attack time of the PPM is carefully specified. If a peak is so short that the PPM fails to indicate its true level, the resulting overload will also be so brief that the HAS will not hear it. A further feature of the PPM is that the decay time of the meter is very slow, so that any peaks are visible for much longer and the meter is easier to read because the meter movement is less violent.

The original PPM as developed by the BBC was sparsely calibrated, but other users have adopted the same dynamics and added dB scales, Figure 4.7 shows some of the scales in use.

In broadcasting, the use of level metering and line-up procedures ensures that the level experienced by the viewer/listener does not change significantly from program to program. Consequently in a transmission suite, the goal would be to broadcast tapes at a level identical to that which was obtained during production. However, when making a recording prior to any production process, the goal would be to modulate the tape as fully as possible without clipping as this would then give the best signal-to-noise ratio. The level would then be reduced if necessary in the production process.

Unlike analog recorders, digital systems do not have headroom, as there is no progressive onset of distortion until convertor clipping, the equivalent of saturation, occurs at 0 dBFs. Accordingly many digital recorders have level meters which read in dBFs. The scales are marked with 0 at the clipping level and all operating levels are below that. This causes no difficulty provided the user is aware of the consequences.




Figure 4.7 Some of the scales used in conjunction with the PPM dynamics. (After Francis Rumsey, with permission.)

However, in the situation where a digital copy of an analog tape is to be made, it is very easy to set the input gain of the digital recorder so that line-up tone from the analog tape reads 0 dB. This lines up digital clipping with the analog operating level. When the tape is dubbed, all signals in the headroom suffer convertor clipping.

In order to prevent such problems, manufacturers and broadcasters have introduced artificial headroom on digital level meters, simply by calibrating the scale and changing the analog input sensitivity so that 0 dB analog is some way below clipping. Unfortunately there has been little agreement on how much artificial headroom should be provided, and machines which have it are seldom labelled with the amount. There is an argument which suggests that the amount of headroom should be a function of the sample wordlength, but this causes difficulties when transferring from one wordlength to another. In sixteen-bit working, 12 dB of headroom is a useful figure, but now that eighteen- and twenty-bit convertors are available, 18 dB may be more appropriate.

## 4.4 The ear

The human auditory system, the sense called hearing, is based on two obvious transducers at the side of the head, and a number of less obvious mental processes which give us an impression of the world around us based on disturbances to the equilibrium of the air which we call sound. It is only possible briefly to introduce the subject here. The interested reader is referred to Moore<sup>3</sup> for an excellent treatment.

The HAS can tell us, without aid from any other senses, where a sound source is, how big it is, whether we are in an enclosed space and how big that is. If the sound source is musical, we can further establish information such as pitch and timbre, attack, sustain and decay. In order to do this, the auditory system must work in the time, frequency and space domains. A sound reproduction system which is inadequate in one of these domains will be unrealistic however well the other two are satisfied. Chapter 3 introduced the concept of uncertainty between the time and frequency domains and the ear cannot analyse both at once. The HAS circumvents this by changing its characteristics dynamically so that it can concentrate on one domain or the other.

The acuity of the HAS is astonishing. It can detect tiny amounts of distortion, and will accept an enormous dynamic range over a wide number of octaves. If the ear detects a different degree of impairment between two audio systems and an original or 'live' sound in properly conducted tests, we can say that one of them is superior. Thus quality is completely subjective and can only be checked by listening tests. However, any characteristic of a signal which can be heard can in principle also be measured by a suitable instrument, although in general the availability of such instruments lags the requirement and the use of such instruments lags the availability. The subjective tests will tell us how sensitive the instrument should be. Then the objective readings from the instrument give an indication of how acceptable a signal is in respect of that characteristic.

Figure 4.8 shows that the structure of the ear is traditionally divided into the outer, middle and inner ears. The outer ear works at low impedance, the inner ear works at high impedance, and the middle ear is an impedance-matching device. The visible part of the outer ear is called the pinna which plays a subtle role in determining the direction of arrival of sound at high frequencies. It is too small to have any effect at low frequencies. Incident sound enters the auditory canal or meatus. The pipe-like meatus causes a small resonance at around 4kHz. Sound vibrates the eardrum or tympanic membrane which seals the outer ear from the middle ear. The inner ear or cochlea works by sound travelling through a fluid. Sound enters the cochlea via a membrane called the oval window. If airborne sound were to be incident on the oval window




Figure 4.8 The structure of the human ear. See text for details.

directly, the serious impedance mismatch would cause most of the sound to be reflected. The middle ear remedies that mismatch by providing a mechanical advantage.

The tympanic membrane is linked to the oval window by three bones known as ossicles which act as a lever system such that a large displacement of the tympanic membrane results in a smaller displacement of the oval window but with greater force. Figure 4.9 shows that the




Figure 4.9 The malleus tensions the tympanic membrane into a conical shape. The ossicles provide an impedance-transforming lever system between the tympanic membrane and the oval.

malleus applies a tension to the tympanic membrane rendering it conical in shape. The malleus and the incus are firmly joined together to form a lever. The incus acts upon the stapes through a spherical joint. As the area of the tympanic membrane is greater than that of the oval window, there is a further multiplication of the available force. Consequently small pressures over the large area of the tympanic membrane are converted to high pressures over the small area of the oval window. The middle ear evolved to operate at natural sound levels and causes distortion at the high levels which can be generated with artificial amplification.


The middle ear is normally sealed, but ambient pressure changes will cause static pressure on the tympanic membrane which is painful. The



pressure is relieved by the Eustachian tube which opens involuntarily whilst swallowing. Some of the discomfort of the common cold is due to these tubes becoming blocked. The Eustachian tubes open into the cavities of the head and must normally be closed to avoid one's own speech appearing deafeningly loud. The ossicles are located by minute muscles which are normally relaxed. However, the middle ear reflex is an involuntary tightening of the *tensor tympani* and *stapedius* muscles which heavily damp the ability of the tympanic membrane and the stapes to transmit sound by about 12 dB at frequencies below 1 kHz. The main function of this reflex is to reduce the audibility of one's own speech. However, loud sounds will also trigger this reflex which takes some 60–120 milliseconds to operate; too late to protect against transients such as gunfire.

#### 4.5 The cochlea

The cochlea is the transducer proper, converting pressure variations in the fluid into nerve impulses. However, unlike a microphone, the nerve impulses are not an analog of the incoming waveform. Instead the cochlea has some analysis capability which is combined with a number of mental processes to make a complete analysis. As shown in Figure 4.10(a), the cochlea is a fluid-filled tapering spiral cavity within bony walls. The widest part, near the oval window, is called the *base* and the distant end is the *apex*. Figure 4.10(b) shows that the cochlea is divided lengthwise into three volumes by Reissner's membrane and the basilar membrane. The *scala vestibuli* and the *scala tympani* are connected by a small aperture at the apex of the cochlea known as the *helicotrema*. Vibrations from the stapes are transferred to the oval window and become fluid pressure variations which are relieved by the flexing of the round window. Effectively the basilar membrane is in series with the fluid motion and is



**Figure 4.10** (a) The cochlea is a tapering spiral cavity. (b) The cross-section of the cavity is divided by Reissner's membrane and the basilar membrane.




Figure 4.11 The basilar membrane tapers so its resonant frequency changes along its length.

driven by it except at very low frequencies where the fluid flows through the helicotrema, decoupling the basilar membrane.

To assist in its frequency domain operation, the basilar membrane is not uniform. Figure 4.11 shows that it tapers in width and varies in thickness in the opposite sense to the taper of the cochlea. The part of the basilar membrane which resonates as a result of an applied sound is a function of the frequency. High frequencies cause resonance near to the oval window, whereas low frequencies cause resonances further away. More precisely the distance from the apex where the maximum resonance occurs is a logarithmic function of the frequency. Consequently tones spaced apart in octave steps will excite evenly spaced resonances in the basilar membrane. The prediction of resonance at a particular location on the membrane is called *place theory*. Among other things, the basilar membrane is a mechanical frequency analyser. A knowledge of the way it operates is essential to an understanding of musical phenomena such as pitch discrimination, timbre, consonance and dissonance and to auditory phenomena such as critical bands, masking and the precedence effect.

The vibration of the basilar membrane is sensed by the organ of Corti which runs along the centre of the cochlea. The organ of Corti is active in that it contains elements which can generate vibration as well as sense it. These are connected in a regenerative fashion so that the  $Q$  factor, or frequency selectivity, of the ear is higher than it would otherwise be. The deflection of hair cells in the organ of Corti triggers nerve firings and these signals are conducted to the brain by the auditory nerve.

Nerve firings are not a perfect analog of the basilar membrane motion. A nerve firing appears to occur at a constant phase relationship to the basilar vibration, a phenomenon called phase locking, but firings do not necessarily occur on every cycle. At higher frequencies firings are intermittent, yet each is in the same phase relationship.

The resonant behaviour of the basilar membrane is not observed at the lowest audible frequencies below 50 Hz. The pattern of vibration does not appear to change with frequency and it is possible that the frequency is low enough to be measured directly from the rate of nerve firings.

## 4.6 Level and loudness

At its best, the HAS can detect a sound pressure variation of only  $2 \times 10^{-5}$  Pascals rms and so this figure is used as the reference against which sound pressure level (SPL) is measured. The sensation of loudness is a logarithmic function of SPL hence the use of the deciBel explained in section 4.2. The dynamic range of the HAS exceeds 130 dB, but at the extremes of this range, the ear is either straining to hear or is in pain.

The frequency response of the HAS is not at all uniform and it also changes with SPL. The subjective response to level is called loudness and is measured in *phons*. The phon scale and the SPL scale coincide at 1 kHz., but at other frequencies the phon scale deviates because it displays the actual SPLs judged by a human subject to be equally loud as a given level at 1 kHz. Figure 4.12 shows the so-called equal loudness contours which were originally measured by Fletcher and Munson and subsequently by Robinson and Dadson. Note the irregularities caused by resonances in the meatus at about 4 kHz and 13 kHz.

Usually, people's ears are at their most sensitive between about 2 kHz and 5 kHz, and although some people can detect 20 kHz at high level, there is much evidence to suggest that most listeners cannot tell if the upper frequency limit of sound is 20 kHz or 16 kHz.<sup>4,5</sup> For a long time it was thought that frequencies below about 40 Hz were unimportant, but it

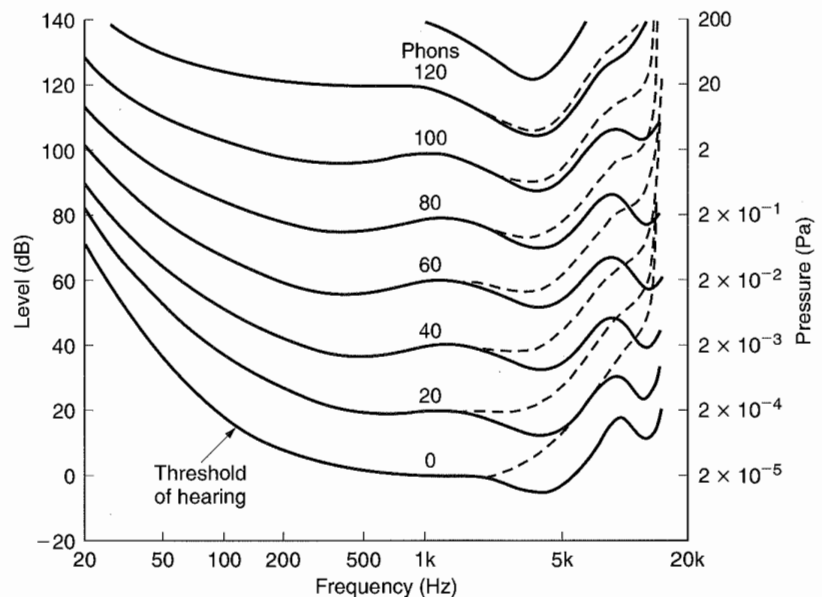


Figure 4.12 Contours of equal loudness showing that the frequency response of the ear is highly level dependent (solid line, age 20; dashed line, age 60).

is now clear that reproduction of frequencies down to 20 Hz improves reality and ambience.<sup>6</sup> The generally accepted frequency range for high-quality audio is 20 Hz to 20 000 Hz, although for broadcasting an upper limit of 15 000 Hz is often applied.

The most dramatic effect of the curves of Figure 4.12 is that the bass content of reproduced sound is disproportionately reduced as the level is turned down. This would suggest that if a powerful yet high-quality reproduction system is available the correct tonal balance when playing a good recording can be obtained simply by setting the volume control to the correct level. This is indeed the case. A further consideration is that many musical instruments and the human voice change timbre with level and there is only one level which sounds correct for the timbre.

Oddly, there is as yet no standard linking the signal level in a transmission or recording system with the SPL at the microphone, although with the advent of digital microphones this useful information could easily be sent as metadata. Loudness is a subjective reaction and is almost impossible to measure. In addition to the level-dependent frequency response problem, the listener uses the sound not for its own sake but to draw some conclusion about the source. For example, most people hearing a distant motorcycle will describe it as being loud. Clearly at the source, it *is* loud, but the listener has compensated for the distance. Paradoxically the same listener may then use a motor mower without hearing protection.


The best that can be done is to make some compensation for the level-dependent response using *weighting curves*. Ideally there should be many, but in practice the A, B and C weightings were chosen where the A curve is based on the 40-phon response. The measured level after such a filter is in units of dBA. The A curve is almost always used because it most nearly relates to the annoyance factor of distant noise sources. The use of A-weighting at higher levels is highly questionable.

## 4.7 Frequency discrimination

Figure 4.13 shows an uncoiled basilar membrane with the apex on the left so that the usual logarithmic frequency scale can be applied. The envelope of displacement of the basilar membrane is shown for a single frequency at (a). The vibration of the membrane in sympathy with a single frequency cannot be localized to an infinitely small area, and nearby areas are forced to vibrate at the same frequency with an amplitude that decreases with distance. Note that the envelope is asymmetrical because the membrane is tapering and because of frequency-dependent losses in the propagation of vibrational energy down the cochlea. If the frequency is changed, as in (b), the position of

Pressure (Pa)  
 $10^{-1}$   
 $10^{-2}$   
 $10^{-3}$   
 $10^{-4}$   
 $10^{-5}$

f the ear



**Figure 4.13** The basilar membrane symbolically uncolled. (a) Single frequency causes the vibration envelope shown. (b) Changing the frequency moves the peak of the envelope.

maximum displacement will also change. As the basilar membrane is continuous, the position of maximum displacement is infinitely variable allowing extremely good pitch discrimination of about one twelfth of a semitone which is determined by the spacing of hair cells.

In the presence of a complex spectrum, the finite width of the vibration envelope means that the ear fails to register energy in some bands when there is more energy in a nearby band. Within those areas, other frequencies are mechanically excluded because their amplitude is insufficient to dominate the local vibration of the membrane. Thus the  $Q$  factor of the membrane is responsible for the degree of auditory masking, defined as the decreased audibility of one sound in the presence of another.

## 4.8 Critical bands

The term used in psychoacoustics to describe the finite width of the vibration envelope is *critical bandwidth*. Critical bands were first described by Fletcher.<sup>7</sup> The envelope of basilar vibration is a complicated function. It is clear from the mechanism that the area of the membrane involved will increase as the sound level rises. Figure 4.14 shows the bandwidth as a function of level.

As was shown in Chapter 3, the Heisenberg inequality teaches that the higher the frequency resolution of a transform, the worse the time accuracy. As the basilar membrane has finite frequency resolution measured in the width of a critical band, it follows that it must have finite time resolution. This also follows from the fact that the membrane is resonant, taking time to start and stop vibrating in response to a stimulus. There are many examples of this. Figure 4.15 shows the impulse response.




Figure 4.14 The critical bandwidth changes with SPL.




Figure 4.15 Impulse response of the ear showing slow attack and decay due to resonant behaviour.

Figure 4.16 shows the perceived loudness of a tone burst increases with duration up to about 200 ms due to the finite response time.

The HAS has evolved to offer intelligibility in reverberant environments which it does by averaging all received energy over a period of about 30 milliseconds. Reflected sound which arrives within this time is integrated to produce a louder sensation, whereas reflected sound which arrives after that time can be temporally discriminated and is perceived as an echo. Microphones have no such ability which is why we often need to have acoustic treatment in areas where microphones are used.

A further example of the finite time discrimination of the HAS is the fact that short interruptions to a continuous tone are difficult to detect. Finite time resolution means that masking can take place even when the masking tone begins after and ceases before the masked sound. This is referred to as forward and backward masking.<sup>8</sup>

As the vibration envelope is such a complicated shape, Moore and Glasberg have proposed the concept of equivalent rectangular bandwidth to simplify matters. The ERB is the bandwidth of a rectangular filter which passes the same power as a critical band. Figure 4.17(a) shows the




Figure 4.16 Perceived level of tone burst rises with duration as resonance builds up.





Figure 4.17 Effective rectangular bandwidth of critical band is much wider than the resolution of the pitch discrimination mechanism.

expression they have derived linking the ERB with frequency. This is plotted in (b) where it will be seen that one third of an octave is a good approximation. This is about thirty times broader than the pitch discrimination also shown in (b).

## 4.9 Beats

Figure 4.18 shows an electrical signal (a) in which two equal sine waves of nearly the same frequency have been linearly added together. Note



**Figure 4.18** (a) Result of adding two sine waves of similar frequency. (b) Spectrum of (a) to infinite accuracy. (c) With finite accuracy only a single frequency is distinguished whose amplitude changes with the envelope of (a) giving rise to beats.

that the envelope of the signal varies as the two waves move in and out of phase. Clearly the frequency transform calculated to infinite accuracy is that shown at (b). The two amplitudes are constant and there is no evidence of the envelope modulation. However, such a measurement requires an infinite time. When a shorter time is available, the frequency discrimination of the transform falls and the bands in which energy is detected become broader.

When the frequency discrimination is too wide to distinguish the two tones as in (c), the result is that they are registered as a single tone. The amplitude of the single tone will change from one measurement to the next because the envelope is being measured. The rate at which the envelope amplitude changes is called a *beat* frequency which is not actually present in the input signal. Beats are an artifact of finite frequency resolution transforms. The fact that the HAS produces beats from pairs of tones proves that it has finite resolution.

Measurement of when beats occur allows measurement of critical bandwidth. Figure 4.19 shows the results of human perception of a two-tone signal as the frequency  $\Delta F$  difference changes. When  $\Delta F$  is zero, described musically as *unison*, only a single note is heard. As  $\Delta F$  increases, beats are heard, yet only a single note is perceived. The limited frequency



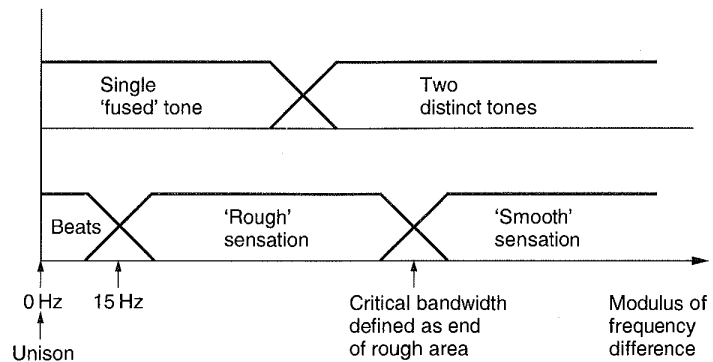


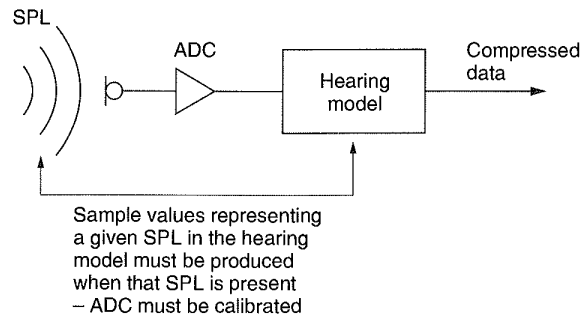
Figure 4.19 Perception of two-tone signal as frequency difference changes.

resolution of the basilar membrane has *fused* the two tones together. As  $dF$  increases further, the sensation of beats ceases at 12–15 Hz and is replaced by a sensation of roughness or *dissonance*. The roughness is due to parts of the basilar membrane being unable to decide the frequency at which to vibrate. The regenerative effect may well become confused under such conditions. The roughness which persists until  $dF$  has reached the critical bandwidth beyond which two separate tones will be heard because there are now two discrete basilar resonances. In fact this is the definition of critical bandwidth.

#### 4.10 Codec level calibration

Perceptive coding relies on the principle of auditory masking. Masking causes the ear/brain combination to be less sensitive to sound at one frequency in the presence of another at a nearby frequency. If a first tone is present in the input, then it will mask signals of lower level at nearby frequencies. The quantizing of the first tone and of further tones at those frequencies can be made coarser. Fewer bits are needed and a coding gain results. The increased quantizing distortion is allowable if it is masked by the presence of the first tone.

The functioning of the ear is noticeably level dependent and perceptive coders take this into account. However, all signal processing takes place in the electrical or digital domain with respect to electrical or numerical levels whereas the hearing mechanism operates with respect to true sound pressure level. Figure 4.20 shows that in an ideal system the overall gain of the microphones and ADCs is such that the PCM codes have a relationship with sound pressure which is the same as that assumed by the model in the codec. Equally the overall gain of the DAC and loudspeaker system should be such that the sound pressure levels which the codec assumes are those



**Figure 4.20** Audio coders must be level calibrated so that the psychoacoustic decisions in the coder are based on correct sound pressure levels.

actually heard. Clearly the gain control of the microphone and the volume control of the reproduction system must be calibrated if the hearing model is to function properly. If, for example, the microphone gain was too low and this was compensated by advancing the loudspeaker gain, the overall gain would be the same but the codec would be fooled into thinking that the sound pressure level was less than it really was and the masking model would not then be appropriate.

The above should come as no surprise as analog audio codecs such as the various Dolby systems have required and implemented line-up procedures and suitable tones. However obvious the need to calibrate coders may be, the degree to which this is recognized in the industry is almost negligible to date and this can only result in sub-optimal performance.

## 4.11 Quality measurement

As has been seen, one way in which coding gain is obtained is to requantize sample values to reduce the wordlength. Since the resultant requantizing error is a distortion mechanism it results in energy moving from one frequency to another. The masking model is essential to estimate how audible the effect of this will be. The greater the degree of compression required, the more precise the model must be. If the masking model is inaccurate, then equipment based upon it may produce audible artifacts under some circumstances. Artifacts may also result if the model is not properly implemented. As a result, development of audio compression units requires careful listening tests with a wide range of source material<sup>9,10</sup> and precision loudspeakers. The presence of artifacts at a given compression factor indicates only that performance is below expectations; it does not distinguish between the implementation and the model. If the implementation is verified, then a more detailed model must be sought. Naturally comparative listening tests are only valid if all the codecs have

been level calibrated and if the loudspeakers cause less loss of information than any of the codecs, a requirement which is frequently overlooked.

Properly conducted listening tests are expensive and time consuming, and alternative methods have been developed which can be used objectively to evaluate the performance of different techniques. The noise-to-masking ratio (NMR) is one such measurement.<sup>11</sup> Figure 4.21 shows how NMR is measured. Input audio signals are fed simultaneously to a data reduction coder and decoder in tandem and to a compensating delay whose length must be adjusted to match the codec

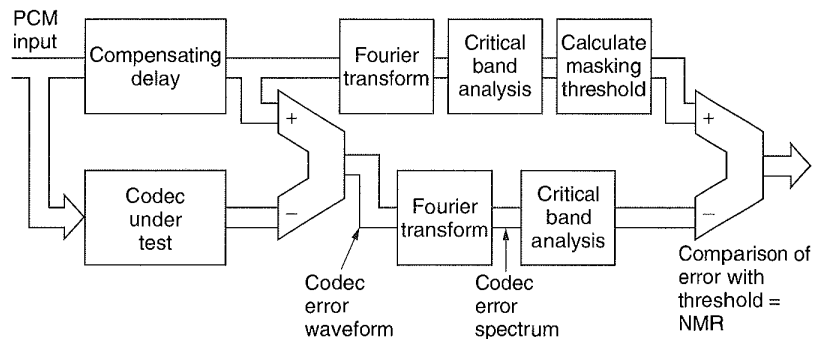


Figure 4.21 The noise-to-masking ratio is derived as shown here.

delay. At the output of the delay, the coding error is obtained by subtracting the codec output from the original. The original signal is spectrum-analysed into critical bands in order to derive the masking threshold of the input audio, and this is compared with the critical band spectrum of the error. The NMR in each critical band is the ratio between the masking threshold and the quantizing error due to the codec. An average NMR for all bands can be computed. A positive NMR in any band indicates that artifacts are potentially audible. Plotting the average NMR against time is a powerful technique, as with an ideal codec the NMR should be stable with different types of program material. If this is not the case the codec could perform quite differently as a function of the source material. NMR excursions can be correlated with the waveform of the audio input to analyse how the extra noise was caused and to redesign the codec to eliminate it.

Practical systems should have a finite NMR in order to give a degree of protection against difficult signals which have not been anticipated and against the use of post-codec equalization or several tandem codecs which could change the masking threshold. There is a strong argument that devices used for audio production should have a greater NMR than consumer or program delivery devices.

## 4.12 The limits


There are of course, limits to all technologies. Eventually artifacts will be heard as the amount of compression is increased which no amount of detailed modelling will remove. The ear is only able to perceive a certain proportion of the information in a given sound. This could be called the perceptual entropy,<sup>12</sup> and all additional sound is redundant or irrelevant. Data reduction works by removing the redundancy, and clearly an ideal system would remove all of it, leaving only the entropy. Once this has been done, the masking capacity of the ear has been reached and the NMR has reached zero over the whole band. Reducing the data rate further must reduce the entropy, because raising noise further at any frequency will render it audible. In practice the audio bandwidth will have to be reduced in order to keep the noise level acceptable. In MPEG-1 pre-filtering allows data from higher sub-bands to be neglected. MPEG-2 has introduced some low sampling rate options for this purpose. Thus there is a limit to the degree of data reduction which can be achieved even with an ideal coder. Systems which go beyond that limit are not appropriate for high-quality music, but are relevant in news gathering and communications where intelligibility of speech is the criterion. Interestingly, the data rate out of a coder is virtually independent of the input sampling rate unless the sampling rate is very low. This is because the entropy of the sound is in the waveform, not in the number of samples carrying it.

The compression factor of a coder is only part of the story. All codecs cause delay, and in general the greater the compression the longer the delay. In some applications, such as telephony, a short delay is required.<sup>13</sup> In many applications, the compressed channel will have a constant bit rate, and so a constant compression factor is required. In real program material, the entropy varies and so the NMR will fluctuate. If greater delay can be accepted, as in a recording application, memory buffering can be used to allow the coder to operate at constant NMR and instantaneously variable data rate. The memory absorbs the instantaneous data rate differences of the coder and allows a constant rate in the channel. A higher effective compression factor will then be obtained. Near-constant quality can also be achieved using statistical multiplexing.

## 4.13 Compression applications

One of the fundamental concepts of PCM audio is that the signal-to-noise ratio of the channel can be determined by selecting a suitable wordlength. In conjunction with the sampling rate, the resulting data rate is then determined. In many cases the full data rate can be transmitted or

information  
looked.  
consuming,  
can be used  
iques. The  
Figure 4.21  
simultane-  
and to a  
the codec



comparison of  
error with  
threshold =  
MR

obtained by  
al signal is  
e masking  
the critical  
is the ratio  
due to the  
A positive  
y audible.  
ue, as with  
t types of  
form quite  
ons can be  
e how the  
ite it.

re a degree  
anticipated  
lem codecs  
; argument  
ater NMR

recorded. A high-quality digital audio channel requires around one megabit per second, whereas a standard definition component digital video channel needs two hundred times as much. With mild video compression the video bit rate is still so much in excess of the audio rate that audio compression is not worth while. Only when high video compression factors are used does it become necessary to apply compression to the audio.

Professional equipment traditionally has operated with higher sound quality than consumer equipment in order to allow for some inevitable quality loss in production, and so there will be less pressure to use data reduction. In fact there is even pressure against the use of audio compression in critical applications because of the loss of quality, particularly in stereo and surround sound.

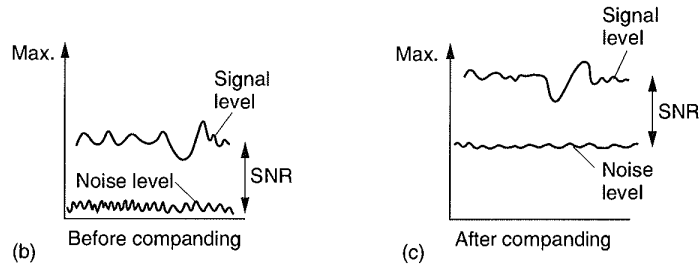
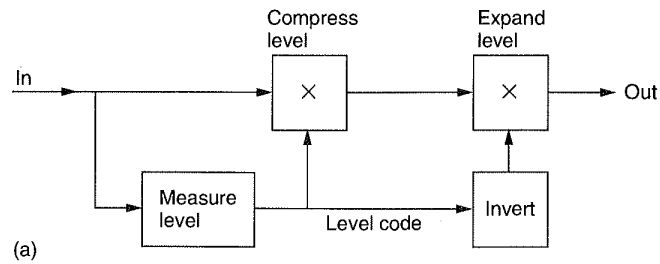
With the rising density and falling cost of digital recording media, the use of compression in professional audio storage applications is hard to justify. However, where there is a practical or economic restriction on channel bandwidth compression becomes essential. Now that the technology exists to broadcast television pictures using compressed digital coding, it is obvious that the associated audio should be conveyed in the same way.

#### 4.14 Audio compression tools

There are many different techniques available for audio compression, each having advantages and disadvantages. Real compressors will combine several techniques or tools in various ways to achieve different combinations of cost and complexity. Here it is intended to examine the tools separately before seeing how they are used in actual compression systems.

The simplest coding tool is companding which is a digital parallel of the noise reducers used in analog tape recording. Figure 4.22(a) shows that in companding the input signal level is monitored. Whenever the input level falls below maximum, it is amplified at the coder. The gain which was applied at the coder is added to the data stream so that the decoder can apply an equal attenuation. The advantage of companding is that the signal is kept as far away from the noise floor as possible. In analog noise reduction this is used to maximize the SNR of a tape recorder, whereas in digital compression it is used to keep the signal level as far as possible above the distortion introduced by various coding steps.

One common way of obtaining coding gain is to shorten the wordlength of samples so that fewer bits need to be transmitted. Figure 4.22(b) shows that when this is done, the distortion will rise by 6 dB for




**Figure 4.22** Digital companding. In (a) the encoder amplifies the input to maximum level and the decoder attenuates by the same amount. (b) In a companded system, the signal is kept as far as possible above the noise caused by shortening the sample wordlength (c).


every bit removed. This is because removing a bit halves the number of quantizing intervals which then must be twice as large, doubling the error amplitude.

Clearly if this step follows the compander of (a), the audibility of the distortion will be minimized. As an alternative to shortening the wordlength, the uniform quantized PCM signal can be converted to a non-uniform format. In non-uniform coding, shown at (c), the size of the quantizing step rises with the magnitude of the sample so that the distortion level is greater when higher levels exist.

Companding is a relative of floating-point coding shown in Figure 4.23 where the sample value is expressed as a mantissa and a binary exponent



**Figure 4.23** In this example of floating-point notation, the radix point can have eight positions determined by the exponent E. The point is placed to the left of the first '1', and the next four bits to the right form the mantissa M. As the MSB of the mantissa is always 1, it need not always be stored.




**Figure 4.24** In this example of an eight-bit mantissa, three-bit exponent system, the maximum SNR is  $6 \text{ dB} \times 8 = 48 \text{ dB}$  with maximum input of  $0 \text{ dB}$ . As input level falls by  $6 \text{ dB}$ , the convertor noise remains the same, so SNR falls to  $42 \text{ dB}$ . Further reduction in signal level causes the convertor to shift range (point A in the diagram) by increasing the input analog gain by  $6 \text{ dB}$ . The SNR is restored, and the exponent changes from 7 to 6 in order to cause the same gain change at the receiver. The noise modulation would be audible in this simple system. A longer mantissa word is needed in practice.

which determines how the mantissa needs to be shifted to have its correct absolute value on a PCM scale. The exponent is the equivalent of the gain setting or scale factor of a compandor.

Clearly in floating-point the signal-to-noise ratio is defined by the number of bits in the mantissa, and as shown in Figure 4.24, this will vary as a sawtooth function of signal level, as the best value, obtained when the mantissa is near overflow, is replaced by the worst value when the mantissa overflows and the exponent is incremented. Floating-point notation is used within DSP chips as it eases the computational problems involved in handling long wordlengths. For example, when multiplying floating-point numbers, only the mantissae need to be multiplied. The exponents are simply added.

A floating-point system requires one exponent to be carried with each mantissa and this is wasteful because in real audio material the level does not change so rapidly and there is redundancy in the exponents. A better alternative is floating-point block coding, also known as near-instantaneous companding, where the magnitude of the largest sample in a block is used to determine the value of an exponent which is valid for the whole block. Sending one exponent per block requires a lower data rate than in true floating-point.<sup>14</sup>

In block coding the requantizing in the coder raises the quantizing error, but it does so over the entire duration of the block. Figure 4.25 shows that if a transient occurs towards the end of a block, the decoder



**Figure 4.25** If a transient occurs towards the end of a transform block, the quantizing noise will still be present at the beginning of the block and may result in a pre-echo where the noise is audible before the transient.

will reproduce the waveform correctly, but the quantizing noise will start at the beginning of the block and may result in a burst of distortion products (also called pre-noise or pre-echo) which is audible before the transient. Temporal masking may be used to make this inaudible. With a 1 ms block, the artifacts are too brief to be heard.

Another solution is to use a variable time window according to the transient content of the audio waveform. When musical transients occur, short blocks are necessary and the coding gain will be low.<sup>15</sup> At other times the blocks become longer allowing a greater coding gain.

Whilst the above systems used alone do allow coding gain, the compression factor has to be limited because little benefit is obtained from masking. This is because the techniques above produce distortion which may be found anywhere over the entire audio band. If the audio input spectrum is narrow, this noise will not be masked.

Sub-band coding<sup>16</sup> splits the audio spectrum up into many different frequency bands. Once this has been done, each band can be individually processed. In real audio signals many bands will contain lower level signals than the loudest one. Individual companding of each band will be more effective than broadband companding. Sub-band coding also allows the level of distortion products to be raised selectively so that distortion is only created at frequencies where spectral masking will be effective.

It should be noted that the result of reducing the wordlength of samples in a sub-band coder is often referred to as noise. Strictly, noise is an unwanted signal which is decorrelated from the wanted signal. This is not generally what happens in audio compression. Although the original audio conversion would have been correctly dithered, the linearizing random element in the low-order bits will be some way below the end of the shortened word. If the word is simply rounded to the nearest integer



the linearizing effect of the original dither will be lost and the result will be quantizing distortion. As the distortion takes place in a bandlimited system the harmonics generated will alias back within the band. Where the requantizing process takes place in a sub-band, the distortion products will be confined to that sub-band as shown in Figure 3.71. Such distortion is anharmonic.


Following any perceptive coding steps, the resulting data may be further subject to lossless binary compression tools such as prediction, Huffman coding or a combination of both.

Audio is usually considered to be a time-domain waveform as this is what emerges from a microphone. As has been seen in Chapter 3, spectral analysis allows any periodic waveform to be represented by a set of harmonically related components of suitable amplitude and phase. In theory it is perfectly possible to decompose a periodic input waveform into its constituent frequencies and phases, and to record or transmit the transform. The transform can then be inverted and the original waveform will be precisely re-created.

Although one can think of exceptions, the transform of a typical audio waveform changes relatively slowly much of the time. The slow speech of an organ pipe or a violin string, or the slow decay of most musical sounds allow the rate at which the transform is sampled to be reduced, and a coding gain results. At some frequencies the level will be below maximum and a shorter wordlength can be used to describe the coefficient. Further coding gain will be achieved if the coefficients describing frequencies which will experience masking are quantized more coarsely.

In practice there are some difficulties, real sounds are not periodic, but contain transients which transformation cannot accurately locate in time. The solution to this difficulty is to cut the waveform into short segments and then to transform each individually. The delay is reduced, as is the computational task, but there is a possibility of artifacts arising because of the truncation of the waveform into rectangular time windows. A solution is to use window functions, and to overlap the segments as shown in Figure 4.26. Thus every input sample appears in just two transforms, but with variable weighting depending upon its position along the time axis.

The DFT (discrete frequency transform) does not produce a continuous spectrum, but instead produces coefficients at discrete frequencies. The frequency resolution (i.e. the number of different frequency coefficients)



**Figure 4.26** Transform coding can only be practically performed on short blocks. These are overlapped using window functions in order to handle continuous waveforms.

the result will be bandlimited and. Where the distortion is 3.71. Such


data may be predicted,

from as this is the 3, spectral by a set of the phase. In the waveform transmit the original waveform

typical audio below speech of musical sounds reduced, and a low maximum level. Further, frequencies

periodic, but located in time. Short segments coded, as is the case because of overlapping windows. A few segments as in just two of its position

a continuous frequencies. The coefficients)



the blocks. These waveforms.

is equal to the number of samples in the window. If overlapped windows are used, twice as many coefficients are produced as are theoretically necessary. In addition the DFT requires intensive computation, owing to the requirement to use complex arithmetic to render the phase of the components as well as the amplitude. An alternative is to use discrete cosine transforms (DCT) or the modified discrete cosine transform (MDCT) which has the ability to eliminate the overhead of coefficients due to overlapping the windows and return to the critically sampled domain.<sup>17</sup> Critical sampling is a term which means that the number of coefficients does not exceed the number which would be obtained with non-overlapping windows.

#### 4.15 Sub-band coding

Sub-band coding takes advantage of the fact that real sounds do not have uniform spectral energy. The wordlength of PCM audio is based on the dynamic range required and this is generally constant with frequency although any pre-emphasis will affect the situation. When a signal with an uneven spectrum is conveyed by PCM, the whole dynamic range is occupied only by the loudest spectral component, and all of the other components are coded with excessive headroom. In its simplest form, sub-band coding works by splitting the audio signal into a number of frequency bands and companding each band according to its own level. Bands in which there is little energy result in small amplitudes which can be transmitted with short wordlength. Thus each band results in variable-length samples, but the sum of all the sample wordlengths is less than that of PCM and so a coding gain can be obtained. Sub-band coding is not restricted to the digital domain; the analog Dolby noise-reduction systems use it extensively.

The number of sub-bands to be used depends upon what other compression tools are to be combined with the sub-band coding. If it is intended to optimize compression based on auditory masking, the sub-bands should preferably be narrower than the critical bands of the ear, and therefore a large number will be required. This requirement is frequently not met: ISO/MPEG Layers I and II use only 32 sub-bands. Figure 4.27 shows the critical condition where the masking tone is at the top edge of the sub-band. It will be seen that the narrower the sub-band, the higher the requantizing 'noise' that can be masked. The use of an excessive number of sub-bands will, however, raise complexity and the coding delay, as well as risking pre-ringing on transients which may exceed the temporal masking.

The bandsplitting process is complex and requires a lot of computation. One bandsplitting method which is useful is quadrature mirror filtering described in Chapter 3.




Figure 4.27 In sub-band coding the worst case occurs when the masking tone is at the top edge of the sub-band. The narrower the band, the higher the noise level which can be masked.

#### 4.16 Audio compression formats

There are numerous formats intended for audio compression and these can be divided into international standards and proprietary designs.

The ISO (International Standards Organization) and the IEC (International Electrotechnical Commission) recognized that compression would have an important part to play and in 1988 established the ISO/IEC/MPEG (Moving Picture Experts Group) to compare and assess various coding schemes in order to arrive at an international standard for compressing video. The terms of reference were extended the same year to include audio and the MPEG/Audio group was formed. MPEG audio coding is used for DAB (digital audio broadcasting) and for the audio content of digital television broadcasts to the DVB standard.

In the USA, it has been proposed to use an alternative compression technique for the audio content of ATSC (advanced television systems committee) digital television broadcasts. This is the AC-3<sup>18</sup> system developed by Dolby Laboratories. The MPEG transport stream structure has also been standardized to allow it to carry AC-3 coded audio. The Digital Video Disk can also carry AC-3 or MPEG audio coding.

Other popular proprietary codes include apt-X which is a mild compression factor/short delay codec and ATRAC which is the codec used in MiniDisc.

#### 4.17 MPEG audio compression

The subject of audio compression was well advanced when the MPEG/Audio group was formed. As a result it was not necessary for the group to produce *ab initio* codecs because existing work was considered suitable.

As part of the Eureka 147 project, a system known as MUSICAM<sup>19</sup> (Masking pattern adapted Universal Sub-band Integrated Coding And Multiplexing) was developed jointly by CCETT in France, IRT in Germany and Philips in the Netherlands. MUSICAM was designed to be suitable for DAB (digital audio broadcasting). As a parallel development, the ASPEC<sup>20</sup> (Adaptive Spectral Perceptual Entropy Coding) system was developed from a number of earlier systems as a joint proposal by AT&T Bell Labs, Thomson, the Fraunhofer Society and CNET. ASPEC was designed for use at high compression factors to allow audio transmission on ISDN.

These two systems were both fully implemented by July 1990 when comprehensive subjective testing took place at the Swedish Broadcasting Corporation.<sup>21,22</sup> As a result of these tests, the MPEG/Audio group combined the attributes of both ASPEC and MUSICAM into a standard<sup>23,24</sup> having three levels of complexity and performance.

These three different levels, which are known as *layers*, are needed because of the number of possible applications. Audio coders can be operated at various compression factors with different quality expectations. Stereophonic classical music requires different quality criteria from monophonic speech. The complexity of the coder will be reduced with a smaller compression factor. For moderate compression, a simple codec will be more cost effective. On the other hand, as the compression factor is increased, it will be necessary to employ a more complex coder to maintain quality.

MPEG Layer I is a simplified version of MUSICAM which is appropriate for the mild compression applications at low cost. Layer II is identical to MUSICAM and is used for DAB and for the audio content of DVB digital television broadcasts. Layer III is a combination of the best features of ASPEC and MUSICAM and is mainly applicable to telecommunications where high compression factors are required.

At each layer, MPEG Audio coding allows input sampling rates of 32, 44.1 and 48 kHz and supports output bit rates of 32, 48, 56, 64, 96, 112, 128, 192, 256 and 384 kbits/s. The transmission can be mono, dual channel (e.g. bilingual), or stereo. Another possibility is the use of joint stereo mode in which the audio becomes mono above a certain frequency. This allows a lower bit rate with the obvious penalty of reduced stereo fidelity.

The layers of MPEG Audio coding (I, II and III) should not be confused with the MPEG-1 and MPEG-2 television coding standards. MPEG-1 and MPEG-2 flexibly define a range of systems for video and audio coding, whereas the layers define types of audio coding.

The earlier MPEG-1 standard compresses audio and video into about 1.5 Mbits/s. The audio coding of MPEG-1 may be used on its own to encode one or two channels at bit rates up to 448 kbits/s. MPEG-2




Figure 4.28 To allow compatibility with two-channel systems, a stereo signal pair is derived from the five surround signals in this manner.

allows the number of channels to increase to five: Left, Right, Centre, Left surround and Right surround. In order to retain reverse compatibility with MPEG-1, the MPEG-2 coding converts the five-channel input to a compatible two-channel signal,  $L_0$ ,  $R_0$ , by matrixing<sup>25</sup> as shown in Figure 4.28. The data from these two channels are encoded in a standard MPEG-1 audio frame, and this is followed in MPEG-2 by an ancillary data frame which an MPEG-1 decoder will ignore. The ancillary frame contains data for another three audio channels. Figure 4.29 shows that there are eight modes in which these three channels can be obtained. The encoder will select the mode which gives the least data rate for the prevailing distribution of energy in the input channels. An MPEG-2 decoder will extract those three channels in addition to the MPEG-1 frame and then recover all five original channels by an inverse matrix which is steered by mode select bits in the bitstream.

L	R	C
L	$R_S$	C
$L_S$	R	C
$L_S$	$R_S$	C
L	R	$L_S$
L	R	$R_S$
$L_S$	R	$R_S$
L	$R_S$	$L_S$

Figure 4.29 In addition to sending the stereo compatible pair, one of the above combinations of signals can be sent. In all cases a suitable inverse matrix can recover the original five channels.

The requirement for MPEG-2 audio to be backward compatible with MPEG-1 audio coding was essential for some markets, but did compromise the performance because certain useful coding tools could not be used. Consequently the MPEG Audio group evolved a multi-channel standard which was not backward compatible because it incorporated additional coding tools in order to achieve higher performance. This came to be known as MPEG-2 AAC (advanced audio coding).

#### 4.18 MPEG Layer I audio coding

Figure 4.30 shows a block diagram of a Layer I coder which is a simplified version of that used in the MUSICAM system. A polyphase filter divides the audio spectrum into 32 equal sub-bands. The output of the filter bank is critically sampled. In other words the output data rate is no higher than the input rate because each band has been heterodyned to a frequency range from zero upwards. Sub-band compression takes advantage of the

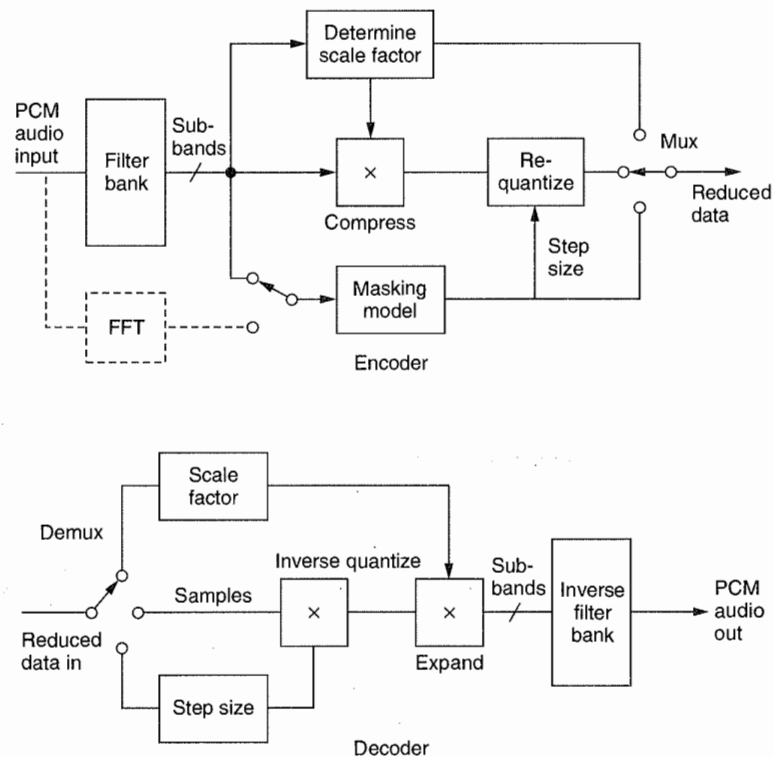


Figure 4.30 A simple sub-band coder. The bit allocation may come from analysis of the sub-band energy, or, for greater reduction, from a spectral analysis in a side chain.


fact that real sounds do not have uniform spectral energy. The wordlength of PCM audio is based on the dynamic range required and this is generally constant with frequency although any pre-emphasis will affect the situation. When a signal with an uneven spectrum is conveyed by PCM, the whole dynamic range is occupied only by the loudest spectral component, and all the other components are coded with excessive headroom. In its simplest form, sub-band coding works by splitting the audio signal into a number of frequency bands and companding each band according to its own level. Bands in which there is little energy result in small amplitudes which can be transmitted with short wordlength. Thus each band results in variable-length samples, but the sum of all the sample wordlengths is less than that of the PCM input and so a degree of coding gain can be obtained.

A Layer I-compliant encoder, i.e. one whose output can be understood by a standard decoder, can be made which does no more than this. Provided the syntax of the bitstream is correct, the decoder is not concerned with how the coding decisions were made. However, higher compression factors require the distortion level to be increased and this should only be done if it is known that the distortion products will be masked. Ideally the sub-bands should be narrower than the critical bands of the ear. Figure 4.27 showed the critical condition where the masking tone is at the top edge of the sub-band. The use of an excessive number of sub-bands will, however raise complexity and the coding delay. The use of 32 equal sub-bands in MPEG Layers I and II is a compromise.

Efficient polyphase bandsplitting filters can only operate with equal-width sub-bands and the result, in an octave-based hearing model, is that sub-bands are too wide at low frequencies and too narrow at high frequencies.

To offset the lack of accuracy in the sub-band filter a parallel fast Fourier transform is used to drive the masking model. The standard suggests masking models, but compliant bitstreams can result from other models. In Layer I a 512-point FFT is used. The output of the FFT is employed to determine the masking threshold which is the sum of all masking sources. Masking sources include at least the threshold of hearing which may locally be raised by the frequency content of the input audio. The degree to which the threshold is raised depends on whether the input audio is sinusoidal or atonal (broadband, or noise-like). In the case of a sine wave, the magnitude and phase of the FFT at each frequency will be similar from one window to the next, whereas if the sound is atonal the magnitude and phase information will be chaotic.

The masking threshold is effectively a graph of just noticeable noise as a function of frequency. Figure 4.31(a) shows an example. The masking threshold is calculated by convolving the FFT spectrum with the cochlea spreading function with corrections for tonality. The level of



**Figure 4.31** A continuous curve (a) of the just-noticeable noise level is calculated by the masking model. The levels of noise in each sub-band (b) must be set so as not to exceed the level of the curve.

the masking threshold cannot fall below the absolute masking threshold which is the threshold of hearing. The masking threshold is then superimposed on the actual frequencies of each sub-band so that the allowable level of distortion in each can be established. This is shown in Figure 4.31(b).

Constant size input blocks are used, containing 384 samples. At 48 kHz, 384 samples corresponds to a period of 8 ms. After the sub-band filter each band contains 12 samples per block. The block size is too long to avoid the pre-masking phenomenon of Figure 4.27. Consequently the masking model must ensure that heavy requantizing is not used in a block which contains a large transient following a period of quiet. This can be done by comparing parameters of the current block with those of the previous block as a significant difference will indicate transient activity.

The samples in each sub-band block or *bin* are companded according to the peak value in the bin. A six-bit scale factor is used for each sub-band which applies to all 12 samples. The gain step is 2 dB and so with a six-bit code over 120 dB of dynamic range is available.

A fixed output bit rate is employed, and as there is no buffering the size of the coded output block will be fixed. The wordlengths in each bin will have to be such that the sum of the bits from all the sub-bands equals the size of the coded block. Thus some sub-bands can have long wordlength coding if others have short wordlength coding. The process of determining the requantization step size, and hence the wordlength in each sub-band is known as bit allocation. In Layer I all sub-bands are treated in the same way and fourteen different requantization classes are used. Each one has an odd number of quantizing intervals so that all codes are referenced to a precise zero level. Where masking takes place, the signal is quantized more coarsely until the distortion level is




raised to just below the masking level. The coarse quantization requires shorter wordlengths and allows a coding gain. The bit allocation may be iterative as adjustments are made to obtain an equal NMR across all sub-bands. If the allowable data rate is adequate, a positive NMR will result and the decoded quality will be optimal. However, at lower bit rates and in the absence of buffering a temporary increase in bit rate is not possible. The coding distortion cannot be masked and the best the encoder can do is to make the (negative) NMR equal across the spectrum so that artifacts are not emphasized unduly in any one sub-band. It is possible that in some sub-bands there will be no data at all, either because such frequencies were absent in the program material or because the encoder has discarded them to meet a low bit rate.

The samples of differing wordlengths in each bin are then assembled into the output coded block. Unlike a PCM block, which contains samples of fixed wordlength, a coded block contains many different wordlengths and which may vary from one sub-band to the next. In order to deserialize the block into samples of various wordlength and demultiplex the samples into the appropriate frequency bins, the decoder has to be told what bit allocations were used when it was packed, and some synchronizing means is needed to allow the beginning of the block to be identified.


The compression factor is determined by the bit-allocation system. It is trivial to change the output block size parameter to obtain a different compression factor. If a larger block is specified, the bit allocator simply iterates until the new block size is filled. Similarly, the decoder need only deserialize the larger block correctly into coded samples and then the expansion process is identical except for the fact that expanded words contain less noise. Thus codecs with varying degrees of compression are available which can perform different bandwidth/performance tasks with the same hardware.

Figure 4.32(a) shows the format of the Layer I elementary stream. The frame begins with a sync pattern to reset the phase of deserialization, and a header which describes the sampling rate and any use of pre-emphasis. Following this is a block of 32 four-bit allocation codes. These specify the wordlength used in each sub-band and allow the decoder to deserialize the sub-band sample block. This is followed by a block of 32 six-bit scale factor indices, which specify the gain given to each band during companding. The last block contains 32 sets of 12 samples. These samples vary in wordlength from one block to the next, and can be from 0 to 15 bits long. The deserializer has to use the 32 allocation information codes to work out how to deserialize the sample block into individual samples of variable length.

The Layer I MPEG decoder is shown in Figure 4.33. The elementary stream is deserialized using the sync pattern and the variable-length



(a) MPEG Layer I frame



(b) MPEG Layer II frame

Figure 4.32 (a) The MPEG Layer I data frame has a simple structure. (b) In the Layer II frame, the compression of the scale factors requires the additional SCFSI code described in the text.

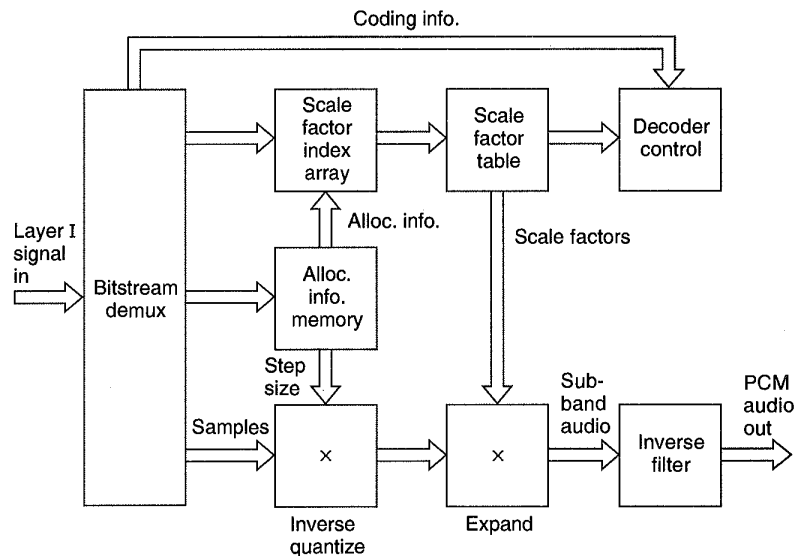


Figure 4.33 The Layer I decoder. See text for details.

samples are assembled using the allocation codes. The variable-length samples are returned to fifteen-bit wordlength by adding zeros. The scale factor indices are then used to determine multiplication factors used to return the waveform in each sub-band to its original amplitude. The 32 sub-band signals are then merged into one spectrum by the synthesis filter. This is a set of bandpass filters which heterodynes every sub-band to the correct place in the audio spectrum and then adds them to produce the audio output.

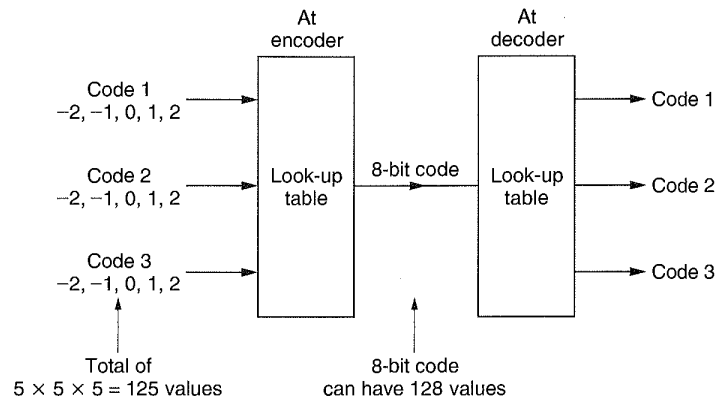
## 4.19 MPEG Layer II audio coding

MPEG Layer II audio coding is identical to MUSICAM. The same 32-band filterbank and the same block companding scheme as Layer I is used. In order to give the masking model better spectral resolution, the side-chain FFT has 1024 points. The FFT drives the masking model which may be the same as is suggested for Layer I. The block length is increased to 1152 samples. This is three times the block length of Layer I, corresponding to 24 ms at 48 kHz.

Figure 4.32(b) shows the Layer II elementary stream structure. Following the sync pattern the bit-allocation data are sent. The requantizing process of Layer II is more complex than in Layer I. The sub-bands are categorized into three frequency ranges, low, medium and high, and the requantizing in each range is different. Low-frequency samples can be quantized into 15 different wordlengths, mid-frequencies into seven different wordlengths and high frequencies into only three different wordlengths. Accordingly the bit-allocation data uses words of four, three and two bits depending on the sub-band concerned. This reduces the amount of allocation data to be sent. In each case one extra combination exists in the allocation code. This is used to indicate that no data are being sent for that sub-band.

The 1152 sample block of Layer II is divided into three blocks of 384 samples so that the same companding structure as Layer I can be used. The 2 dB step size in the scale factors is retained. However, not all the scale factors are transmitted, because they contain a degree of redundancy. In real program material, the difference between scale factors in successive blocks in the same band exceeds 2 dB less than 10 per cent of the time. Layer II coders analyse the set of three successive scale factors in each sub-band. On a stationary program, these will be the same and only one scale factor out of three is sent. As the transient content increases in a given sub-band, two or three scale factors will be sent. A two-bit code known as SCFSI (scale factor select information) must be sent to allow the decoder to determine which of the three possible scale factors have been sent for each sub-band. This technique effectively halves the scale factor bit rate.

As for Layer I, the requantizing process always uses an odd number of steps to allow a true centre zero step. In long-wordlength codes this is not a problem, but when three, five or nine quantizing intervals are used, binary is inefficient because some combinations are not used. For example, five intervals need a three-bit code having eight combinations leaving three unused. The solution is that when three-, five- or nine-level coding is used in a sub-band, sets of three samples are encoded into a *granule*. Figure 4.34 shows how granules work. Continuing the example of five quantizing intervals, each sample could have five different values,

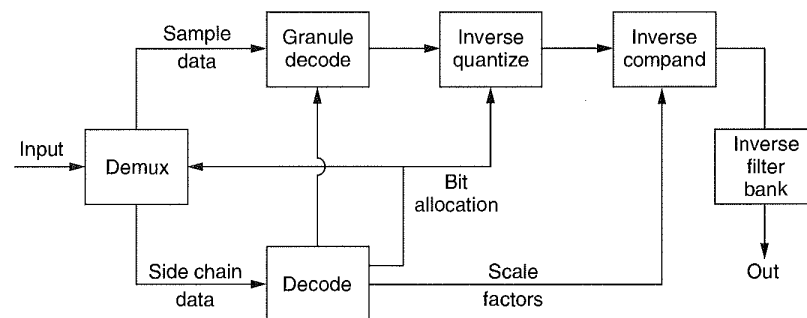


**Figure 4.34** Codes having ranges smaller than a power of two are inefficient. Here three codes with a range of five values which would ordinarily need  $3 \times 3$  bits can be carried in a single eight-bit word.

therefore all combinations of three samples could have 125 different values. As 128 values can be sent with a seven-bit code, it will be seen that this is more efficient than coding the samples separately as three five-level codes would need nine bits. The three requantized samples are used to address a look-up table which outputs the granule code. The decoder can establish that granule coding has been used by examining the bit-allocation data.

The requantized samples/granules in each sub-band, bit allocation data, scale factors and scale factor select codes are multiplexed into the output bit stream.

The Layer II decoder is shown in Figure 4.35. This is not much more complex than the Layer I decoder. The demultiplexing will separate the sample data from the side information. The bit-allocation data will specify the wordlength or granule size used so that the sample block can



**Figure 4.35** A Layer II decoder is slightly more complex than the Layer I decoder because of the need to decode granules and scale factors.

be deserialized and the granules decoded. The scale factor select information will be used to decode the compressed scale factors to produce one scale factor per block of 384 samples. Inverse quantizing and inverse sub-band filtering takes place as for Layer I.

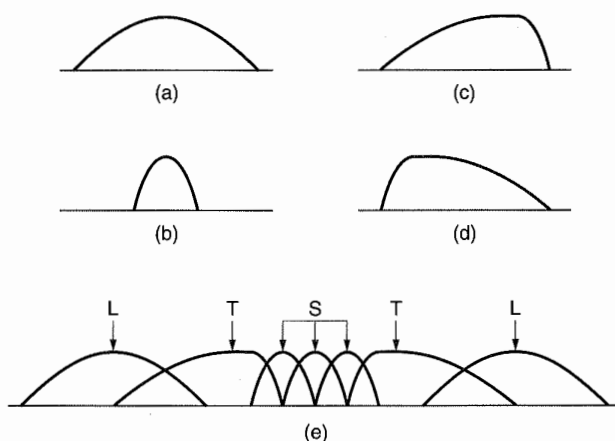
## 4.20 MPEG Layer III audio coding

Layer III is the most complex layer, and is only really necessary when the most severe data rate constraints must be met. It is also known as MP3 in its application of music delivery over the Internet. It is a transform code based on the ASPEC system with certain modifications to give a degree of commonality with Layer II. The original ASPEC coder used a direct MDCT on the input samples. In Layer III this was modified to use a hybrid transform incorporating the existing polyphase 32 band QMF of Layers I and II and retaining the block size of 1152 samples. In Layer III, the 32 sub-bands from the QMF are further processed by a critically sampled MDCT.

The windows overlap by two to one. Two window sizes are used to reduce pre-echo on transients. The long window works with 36 sub-band samples corresponding to 24 milliseconds only at 48 kHz and resolves 18 different frequencies, making 576 frequencies altogether. Coding products are spread over this period which is acceptable in stationary material but not in the vicinity of transients. In this case the window length is reduced to 8 ms. Twelve sub-band samples are resolved into six different frequencies making a total of 192 frequencies. This is the Heisenberg inequality: by increasing the time resolution by a factor of three, the frequency resolution has fallen by the same factor.

Figure 4.36 shows the available window types. In addition to the long and short symmetrical windows there is a pair of transition windows, known as start and stop windows which allow a smooth transition between the two window sizes. In order to use critical sampling, MDCTs must resolve into a set of frequencies which is a multiple of four. Switching between 576 and 192 frequencies allows this criterion to be met. Note that an 8 ms window is still too long to eliminate pre-echo. Pre-echo is eliminated using buffering. The use of a short window minimizes the size of the buffer needed.

Layer III provides a suggested (but not compulsory) psychoacoustic model which is more complex than that suggested for Layers I and II, primarily because of the need for window switching. Pre-echo is associated with the entropy in the audio rising above the average value and this can be used to switch the window size. The perceptive model is used to take advantage of the high-frequency resolution available from the DCT which allows the noise floor to be shaped much more accurately



**Figure 4.36** The window functions of Layer III coding. At (a) is the normal long window, whereas (b) shows the short window used to handle transients. Switching between window sizes requires transition windows (c) and (d). An example of switching using transition windows is shown in (e).

than with the 32 sub-bands of Layers I and II. Although the MDCT has high-frequency resolution, it does not carry the phase of the waveform in an identifiable form and so is not useful for discriminating between tonal and atonal inputs. As a result a side FFT which gives conventional amplitude and phase data is still required to drive the masking model.

Non-uniform quantizing is used, in which the quantizing step size becomes larger as the magnitude of the coefficient increases. The quantized coefficients are then subject to Huffman coding. This is a technique where the most common code values are allocated the shortest wordlength. Layer III also has a certain amount of buffer memory so that pre-echo can be avoided during entropy peaks despite a constant output bit rate.

Figure 4.37 shows a Layer III encoder. The output from the sub-band filter is 32 continuous band-limited sample streams. These are subject to 32 parallel MDCTs. The window size can be switched individually in each sub-band as required by the characteristics of the input audio. The parallel FFT drives the masking model which decides on window sizes as well as producing the masking threshold for the coefficient quantizer. The distortion control loop iterates until the available output data capacity is reached with the most uniform NMR. The available output capacity can vary owing to the presence of the buffer.

Figure 4.38 shows that the buffer occupancy is fed back to the quantizer. During stationary program material, the buffer contents are deliberately run down by slight coarsening of the quantizing. The buffer empties because the output rate is fixed but the input rate has been

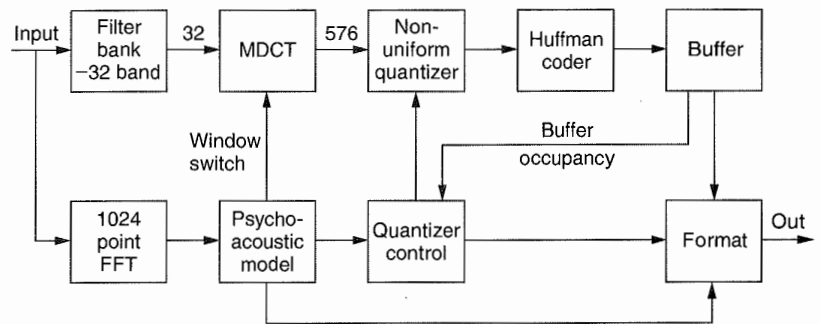


Figure 4.37 The Layer III coder. Note the connection between the buffer and the quantizer which allows different frames to contain different amounts of data.





Figure 4.38 The variable rate coding of Layer III. An approaching transient via the perceptual entropy signal causes the coder to quantize more heavily in order to empty the buffer. When the transient arrives, the quantizing can be made more accurate and the increased data can be accepted by the buffer.

reduced. When a transient arrives, the large coefficients which result can be handled by filling the buffer, avoiding raising the output bit rate whilst also avoiding the pre-echo which would result if the coefficients were heavily quantized.

In order to maintain synchronism between encoder and decoder in the presence of buffering, headers and side information are sent synchronously at frame rate. However, the position of boundaries between the main data blocks which carry the coefficients can vary with respect to the position of the headers in order to allow a variable frame size. Figure 4.39 shows that the frame begins with a unique sync pattern which is followed by the side information. The side information contains a parameter called *main data begin* which specifies where the main data for the present frame began in the transmission. This parameter allows the decoder to find the coefficient block in the decoder buffer. As the frame headers are at fixed locations, the main data blocks may be interrupted by the headers.



**Figure 4.39** In Layer III, the logical frame rate is constant and is transmitted by equally spaced sync patterns. The data blocks do not need to coincide with sync. A pointer after each sync pattern specifies where the data block starts. In this example block 2 is smaller whereas 1 and 3 have enlarged.

## 4.21 MPEG-2 AAC – advanced audio coding

The MPEG standards system subsequently developed an enhanced system known as advanced audio coding (AAC).<sup>26,27</sup> This was intended to be a standard which delivered the highest possible performance using newly developed tools that could not be used in any backward-compatible standard. AAC also forms the core of the audio coding of MPEG-4.

AAC supports up to 48 audio channels with default support of monophonic, stereo and 5.1 channel (3/2) audio. The AAC concept is based on a number of coding tools known as *modules* which can be combined in different ways to produce bitstreams at three different profiles.

The main profile requires the most complex encoder which makes use of all the coding tools. The low-complexity (LC) profile omits certain tools and restricts the power of others to reduce processing and memory requirements. The remaining tools in LC profile coding are identical to those in main profile such that a main profile decoder can decode LC profile bitstreams. The scalable sampling rate (SSR) profile splits the input audio into four equal frequency bands each of which results in a self-contained bitstream. A simple decoder can decode only one, two or three of these bitstreams to produce a reduced bandwidth output. Not all the AAC tools are available to SSR profile.

The increased complexity of AAC allows the introduction of lossless coding tools. These allow a lower bit rate for the same or improved quality at a given bit rate where the reliance on lossy coding is reduced. There is greater attention given to the interplay between time-domain and frequency-domain precision in the human hearing system.



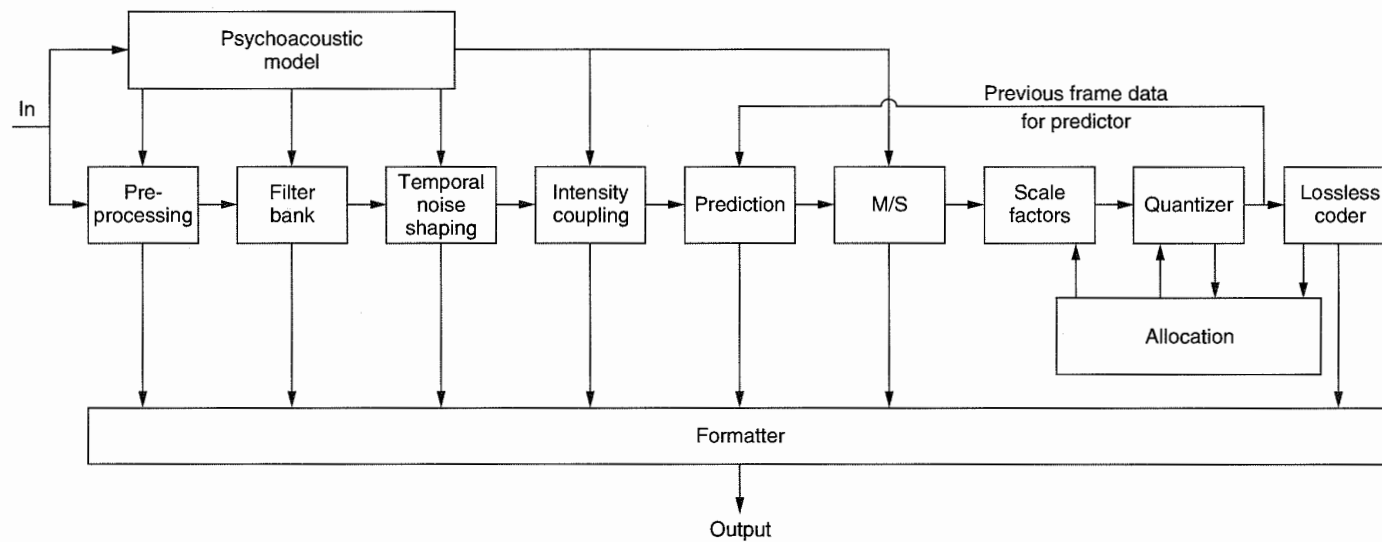


Figure 4.40 The AAC encoder. Signal flow is from left to right whereas side-chain data flow is vertical.

Figure 4.40 shows a block diagram of an AAC main profile encoder. The audio signal path is straight through the centre. The formatter assembles any side chain data along with the coded audio data to produce a compliant bitstream. The input signal passes to the filter bank and the perceptual model in parallel. The filter bank consists of a 50 per cent overlapped critically sampled MDCT which can be switched between block lengths of 2048 and 256 samples. At 48 kHz the filter allows resolutions of 23 Hz and 21 ms or 187 Hz and 2.6 ms. As AAC is a multichannel coding system, block length switching cannot be done indiscriminately as this would result in loss of block phase between channels. Consequently if short blocks are selected, the coder will remain in short block mode for integer multiples of eight blocks. This is shown in Figure 4.41 which also shows the use of transition windows between the block sizes as was done in Layer III.





Figure 4.41 In AAC short blocks must be used in multiples of 8 so that the long block phase is undisturbed. This keeps block synchronism in multichannel systems.

The shape of the window function interferes with the frequency selectivity of the MDCT. In AAC it is possible to select either a sine window or a Kaiser-Bessel-derived (KBD) window as a function of the input audio spectrum. As was seen in Chapter 3, filter windows allow different compromises between bandwidth and rate of roll-off. The KBD window rolls off later but is steeper and thus gives better rejection of frequencies more than about 200 Hz apart whereas the sine window rolls off earlier but less steeply and so gives better rejection of frequencies less than 70 Hz.

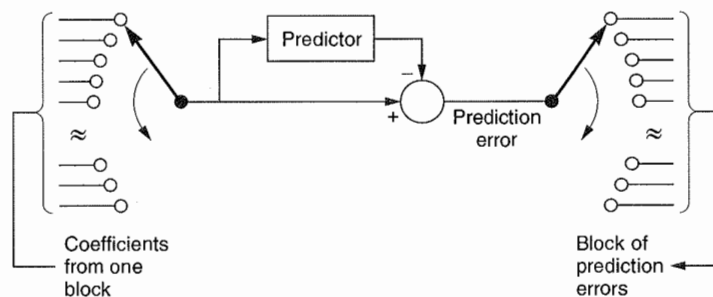
Following the filter bank is the intra-block predictive coding module. When enabled this module finds redundancy between the coefficients within one transform block. In Chapter 3 the concept of transform duality was introduced, in which a certain characteristic in the frequency domain would be accompanied by a dual characteristic in the time domain and vice versa. Figure 4.42 shows that in the time domain, predictive coding works well on stationary signals but fails on transients. The dual of this characteristic is that in the frequency domain, predictive coding works well on transients but fails on stationary signals.



**Figure 4.42** Transform duality suggests that predictability will also have a dual characteristic. A time predictor will not anticipate the transient in (a), whereas the broad spectrum of signal (a), shown in (b), will be easy for a predictor advancing down the frequency axis. In contrast, the stationary signal (c) is easy for a time predictor, whereas in the spectrum of (c) shown at (d) the spectral spike will not be predicted.

Equally, a predictive coder working in the time domain produces an error spectrum which is related to the input spectrum. The dual of this characteristic is that a predictive coder working in the frequency domain produces a prediction error which is related to the input time domain signal. This explains the use of the term *temporal noise shaping* (TNS) used in the AAC documents.<sup>28</sup> When used during transients, the TNS module produces distortion which is time-aligned with the input such that pre-echo is avoided. The use of TNS also allows the coder to use longer blocks more of the time. This module is responsible for a significant amount of the increased performance of AAC.

Figure 4.43 shows that the coefficients in the transform block are serialized by a commutator. This can run from the lowest frequency to the



**Figure 4.43** Predicting along the frequency axis is performed by running along the coefficients in a block and attempting to predict the value of the current coefficient from the values of some earlier ones. The prediction error is transmitted.

highest or in reverse. The prediction method is a conventional forward predictor structure in which the result of filtering a number of earlier coefficients (20 in main profile) is used to predict the current one. The prediction is subtracted from the actual value to produce a prediction error or residual which is transmitted. At the decoder, an identical predictor produces the same prediction from earlier coefficient values and the error in this is cancelled by adding the residual.

Following the intra-block prediction, an optional module known as the intensity/coupling stage is found. This is used for very low bit rates where spatial information in stereo and surround formats is discarded to keep down the level of distortion. Effectively over at least part of the spectrum a mono signal is transmitted along with amplitude codes which allow the signal to be panned in the spatial domain at the decoder.

The next stage is the inter-block prediction module. Whereas the intra-block predictor is most useful on transients, the inter-block predictor module explores the redundancy between successive blocks on stationary signals.<sup>29</sup> This prediction only operates on coefficients below 16 kHz. For each DCT coefficient in a given block, the predictor uses the quantized coefficients from the same locations in two previous blocks to estimate the present value. As before the prediction is subtracted to produce a residual which is transmitted. Note that the use of quantized coefficients to drive the predictor is necessary because this is what the decoder will have to do. The predictor is adaptive and calculates its own coefficients from the signal history. The decoder uses the same algorithm so that the two predictors always track. The predictors run all the time whether prediction is enabled or not in order to keep the prediction coefficients adapted to the signal.

Audio coefficients are associated into sets known as *scale factor bands* for later companding. Within each scale factor band inter-block prediction can be turned on or off depending on whether a coding gain results.

Protracted use of prediction makes the decoder prone to bit errors and drift and removes decoding entry points from the bitstream. Consequently the prediction process is reset cyclically. The predictors are assembled into groups of 30 and after a certain number of a frames a different group is reset until all have been reset. Predictor reset codes are transmitted in the side data. Reset will also occur if short frames are selected.

In stereo and 3/2 surround formats there is less redundancy because the signals also carry spatial information. The effecting of masking may be up to 20 dB less when distortion products are at a different location in the stereo image from the masking sounds. As a result, stereo signals require much higher bit rate than two mono channels, particularly on transient material which is rich in spatial clues.

In some cases a better result can be obtained by converting the signal to a mid-side (M/S) or sum/difference format before quantizing. In



al  
the broad  
vn the  
whereas

luces an  
al of this  
domain  
domain  
(S) used  
module  
that pre-  
er blocks  
nount of

lock are  
cy to the




g the  
ient from

surround sound the M/S coding can be applied to the front L/R pair and the rear L/R pair of signals.

The M/S format can be selected on a block-by-block basis for each scale factor band.

Next comes the lossy stage of the coder where distortion is selectively introduced as a function of frequency as determined by the masking threshold. This is done by a combination of amplification and requantizing. As mentioned, coefficients (or residuals) are grouped into scale factor bands. As Figure 4.44 shows, the number of coefficients varies in order to divide the coefficients into approximate critical bands. Within each scale factor band, all coefficients will be multiplied by the same scale factor prior to requantizing.

Coefficients which have been multiplied by a large scale factor will suffer less distortion by the requantizer whereas those which have been



**Figure 4.44** In AAC the fine-resolution coefficients are grouped together to form scale factor bands. The size of these varies to loosely mimic the width of critical bands.

multiplied by a small scale factor will have more distortion. Using scale factors, the psychoacoustic model can shape the distortion as a function of frequency so that it remains masked. The scale factors allow gain control in 1.5 dB steps over a dynamic range equivalent to twenty-four-bit PCM and are transmitted as part of the side data so that the decoder can re-create the correct magnitudes. The scale factors are differentially coded with respect to the first one in the block and the differences are then Huffman coded.

The requantizer uses non-uniform steps which give better coding gain and has a range of  $\pm 8191$ . The global step size (which applies to all scale factor bands) can be adjusted in 1.5 dB steps. Following requantizing the coefficients are Huffman coded.

There are many ways in which the coder can be controlled and any which results in a compliant bitstream is acceptable although the highest performance may not be reached. The requantizing and scale factor stages will need to be controlled in order to make best use of the available bit rate and the buffering. This is non-trivial because the use of Huffman coding after the requantizer makes it impossible to predict the exact amount of data which will result from a given step size. This means that the process must iterate.

Whatever bit rate is selected, a good encoder will produce consistent quality by selecting window sizes, intra- or inter-frame prediction and using the buffer to handle entropy peaks. This suggests a connection between buffer occupancy and the control system. The psychoacoustic model will analyse the incoming audio entropy and during periods of average entropy it will empty the buffer by slightly raising the quantizer step size so that the bit rate entering the buffer falls. By running the buffer down, the coder can temporarily support a higher bit rate to handle transients or difficult material.

Simply stated, the scale factor process is controlled so that the distortion spectrum has the same shape as the masking threshold and the quantizing step size is controlled to make the level of the distortion spectrum as low as possible within the allowed bit rate. If the bit rate allowed is high enough, the distortion products will be masked.

## 4.22 Dolby AC-3

Dolby AC-3<sup>18</sup> is in fact a family of transform coders based on time-domain aliasing cancellation (TDAC) which allow various compromises between coding delay and bit rate to be used. In the modified discrete cosine transform (MDCT), windows with 50 per cent overlap are used. Thus twice as many coefficients as necessary are produced. These are sub-sampled by a factor of two to give a critically sampled transform, which results in potential aliasing in the frequency domain. However, by making a slight change to the transform, the alias products in the second half of a given window are equal in size but of opposite polarity to the alias products in the first half of the next window, and so will be cancelled on reconstruction. This is the principle of TDAC.

Figure 4.45 shows the generic block diagram of the AC-3 coder. Input audio is divided into 50 per cent overlapped blocks of 512 samples. These are subject to a TDAC transform which uses alternate modified sine and

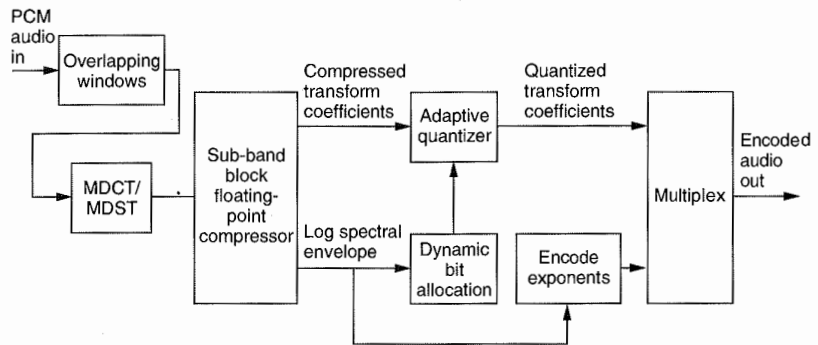


Figure 4.45 Block diagram of the Dolby AC-3 coder. See text for details.

cosine transforms. The transforms produce 512 coefficients per block, but these are redundant and after the redundancy has been removed there are 256 coefficients per block. The input waveform is constantly analysed for the presence of transients and if these are present the block length will be halved to prevent pre-noise. This halves the frequency resolution but doubles the temporal resolution.

The coefficients have high-frequency resolution and are selectively combined in sub-bands which approximate the critical bands. Coefficients in each sub-band are normalized and expressed in floating-point block notation with common exponent. The exponents in fact represent the logarithmic spectral envelope of the signal and can be used to drive the perceptive model which operates the bit allocation. The mantissae of the transform coefficients are then requantized according to the bit allocation.

The output bitstream consists of the requantized coefficients and the log spectral envelope in the shape of the exponents. There is a great deal of redundancy in the exponents. In any block, only the first exponent, corresponding to the lowest frequency, is transmitted absolutely. Remaining coefficients are transmitted differentially. Where the input has a smooth spectrum the exponents in several bands will be the same and the differences will then be zero. In this case exponents can be grouped using flags.

Further use is made of temporal redundancy. An AC-3 sync frame contains six blocks. The first block of the frame contains absolute exponent data, but where stationary audio is encountered, successive blocks in the frame can use the same exponents.

The receiver uses the log spectral envelope to deserialize the mantissae of the coefficients into the correct wordlengths. The highly redundant exponents are decoded starting with the lowest frequency coefficient in the first block of the frame and adding differences to create the remainder.

The exponents are then used to return the coefficients to fixed-point notation. Inverse transforms are then computed, followed by a weighted overlapping of the windows to obtain PCM data.

#### 4.23 MPEG-4 Audio

The audio coding options of MPEG-4 parallel the video coding in complexity. In the same way that the video coding in MPEG-4 has moved in the direction of graphics with rendering taking place in the decoder, MPEG-4 audio introduces *structured audio* in which audio synthesis takes place in the decoder, taking MPEG-4 audio into the realm of interactive systems and virtual reality. It is now necessary to describe the audio of earlier formats as *natural sound*, i.e. that which could have come from a microphone. Natural sound is well supported in MPEG-4 with a development of AAC which is described in the next section.

Like video coding, MPEG-4 audio coding may be object based. For example, instead of coding the waveforms of a stereo mix, each sound source in the mix may become a sound object which is individually coded. At the decoder each sound object is then supplied to the composition stage where it will be panned and mixed with other objects. At the moment techniques do not exist to create audio objects from a natural stereo signal pair, but where the audio source is a synthetic, or a mixdown of multiple natural tracks, object coding can be used with a saving in data rate.


It is also possible to define virtual instruments at the decoder and then to make each of these play by transmitting a score.

Speech coding is also well supported. Natural speech may be coded at very low bit rates where the goal is intelligibility of the message rather than fidelity. This can be done with one of two tools: HVXC (Harmonic Vector eXcitation Coding) or CELP (Code Excited Linear Prediction). Synthetic speech capability allows applications such as text-to-speech (TTS). MPEG-4 has standardized transmission of speech information in IPA (International Phonetic Alphabet) or as text.

#### 4.24 MPEG-4 AAC

The MPEG-2 AAC coding tools described in section 4.21 are extended in MPEG-4. The additions are perceptual noise substitution (PNS) and vector quantization. All coding schemes find noise difficult because it contains no redundancy. Real audio program material may contain a certain proportion of noise from time to time, traditionally requiring a high bit rate to transmit without artifacts.

However, experiments have shown that under certain conditions, the listener is unable to distinguish between the original noise-like waveform



lock, but  
there are  
lysed for  
th will be  
ation but

electively  
s. Coeffi-  
ing-point  
represent  
l to drive  
ntissae of  
o the bit

s and the  
great deal  
xponent,  
Remain-  
ut has a  
e and the  
ed using

nc frame  
absolute  
uccessive

nantissae  
edundant  
ficient in  
mainder.



and noise locally generated in the decoder. This is the basis for PNS. Instead of attempting to code a difficult noise sequence, the PNS process will transmit the amplitude of the noise which the decoder will then create.

At the encoder, PNS will be selected if, over a certain range of frequencies, there is no dominant tone and the time-domain waveform remains stable, i.e. there are no transients. On a scale factor band and group basis the Huffman-coded symbols describing the frequency coefficients will be replaced by the PNS flag. At the decoder the missing coefficients will be derived from random vectors. The amplitude of the noise is encoded in 1.5 dB steps with noise energy parameters for each group and scale factor band.

In stereo applications, where PNS is used at the same time and frequency in both channels, the random processes in each channel will be different in order to avoid creating a mid-stage noise object. PNS cannot be used with M/S audio coding.

In MPEG-2 AAC, the frequency coefficients (or their residuals) are quantized according to the adaptive bit-allocation system and then Huffman coded. At low bit rates the heavy requantizing necessary will result in some coefficients being in error. At bit rates below 16 kbit/s/channel an alternative coding scheme known as TwinVQ (Transform Domain Weighted Interleaved Vector Quantization) may be used. Vector quantization, also known as block quantization,<sup>30</sup> works on blocks of coefficients rather than the individual coefficients used by the Huffman code. In vector quantizing, one transmitted symbol represents the state of a number of coefficients. In a lossless system, this symbol would need as many bits as the sum of the coefficients to be coded. In practice the symbol has fewer bits because the coefficients are quantized (and thereby contain errors). The encoder will select a symbol which minimizes the errors over the whole block.

Error minimization is assisted by interleaving at the encoder. After interleaving, adjacent coefficients in frequency space are in different blocks. After the symbol look-up process in the decoder, de-interleaving is necessary to return the coefficients to their correct frequencies. In Twin VQ the transmitted symbols have constant wordlength because the vector table has a fixed size for a given bit rate. Constant size symbols have an advantage in the presence of bit errors because it is easier to maintain synchronization.

#### **4.25 Compression in stereo and surround sound**

Once hardware became economic, the move to digital audio was extremely rapid. One of these reasons was the sheer sound quality available. When well engineered, the PCM digital domain does so little

damage to the sound quality that the problems of the remaining analog parts usually dominate. The one serious exception to this is lossy compression which does not preserve the original waveform and must therefore be carefully assessed before being used in high-quality applications.

In a monophonic system, all the sound is emitted from a single point and psychoacoustic masking operates to its fullest extent. Audio compression techniques of the kind described above work well in mono. However, in stereophonic (which in this context also includes surround sound) applications different criteria apply. In addition to the timbral information describing the nature of the sound sources, stereophonic systems also contain spatial information describing their location.

The greatest cause for concern is that in stereophonic systems, masking is not as effective. When two sound sources are in physically different locations, the degree of masking is not as great as when they are co-sited. Unfortunately all the psychoacoustic masking models used in today's compressors assume co-siting. When used in stereo or surround systems, the artifacts of the compression process can be revealed. This was first pointed out by the late Michael Gerzon who introduced the term *unmasking* to describe the phenomenon.

The hearing mechanism has an ability to concentrate on one of many simultaneous sound sources based on direction. The brain appears to be able to insert a controllable time delay in the nerve signals from one ear with respect to the other so that when sound arrives from a given direction the nerve signals from both ears are coherent causing the binaural threshold of hearing to be 3–6 dB better than monaural at around 4 kHz. Sounds arriving from other directions are incoherent and are heard less well. This is known as *attentional selectivity*, or more colloquially as the *cocktail party effect*.<sup>31</sup>

Human hearing can locate a number of different sound sources simultaneously by constantly comparing excitation patterns from the two ears with different delays. Strong correlation will be found where the delay corresponds to the interaural delay for a given source. This delay-varying mechanism will take time and the ear is slow to react to changes in source direction. Oscillating sources can only be tracked up to 2–3 Hz and the ability to locate bursts of noise improves with burst duration up to about 700 milliseconds.

Monophonic systems prevent the use of either of these effects completely because the first version of all sounds reaching the listener come from the same loudspeaker. Stereophonic systems allow attentional selectivity to function in that the listener can concentrate on specific sound sources in a reproduced stereophonic image with the same facility as in the original sound. When two sound sources are spatially separated, if the listener uses attentional selectivity to concentrate on one of them,

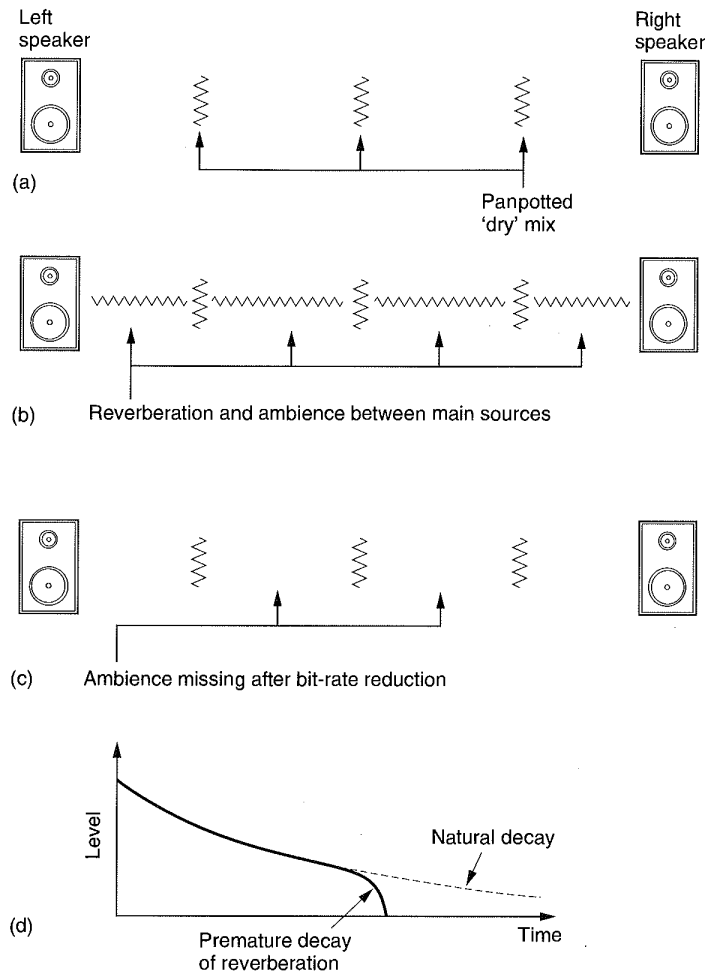
the contributions from both ears will correlate. This means that the contributions from the other sound will be decorrelated, reducing its masking ability significantly. Experiments showed long ago that even technically poor stereo was always preferred to pristine mono. This is because we are accustomed to sounds and reverberation coming from all different directions in real life and having them all superimposed in a mono speaker convinces no-one, however accurate the waveform.

We live in a reverberant world which is filled with sound reflections. If we could separately distinguish every different reflection in a reverberant room we would hear a confusing cacaphony. In practice we hear very well in reverberant surroundings, far better than microphones can, because of the transform nature of the ear and the way in which the brain processes nerve signals. Because the ear has finite frequency discrimination ability in the form of critical bands, it must also have finite temporal discrimination. When two or more versions of a sound arrive at the ear, provided they fall within a time span of about 30 ms, they will not be treated as separate sounds, but will be fused into one sound. Only when the time separation reaches 50–60 ms do the delayed sounds appear as echoes from different directions. As we have evolved to function in reverberant surroundings, most reflections do not impair our ability to locate the source of a sound.

Clearly the first version of a transient sound to reach the ears must be the one which has travelled by the shortest path and this must be the direct sound rather than a reflection. Consequently the ear has evolved to attribute source direction from the time of arrival difference at the two ears of the first version of a transient. This phenomenon is known as the precedence effect. Intensity stereo, the type of signal format obtained with coincident microphones or panpots, works purely by amplitude differences at the two loudspeakers. The two signals should be exactly in phase. As both ears hear both speakers the result is that the space between the speakers and the ears turns the intensity differences into time of arrival differences. These give the illusion of virtual sound sources.

A virtual sound source from a panpot has zero width and on ideal loudspeakers would appear as a virtual point source. Figure 4.46(a) shows how a panpotted dry mix should appear spatially on ideal speakers whereas (b) shows what happens when artificial stereo reverb is added. Figure 4.46(b) is also what is obtained with real sources using a coincident pair of high-quality mikes. In this case the sources are the real sources and the sound between is reverb/ambience.

When listening on high-quality loudspeakers, audio compressors change the characteristic of Figure 4.46(b) to that shown in (c). Even at high bit rates, corresponding to the smallest amount of compression, there is an audible difference between the original and the compressed



**Figure 4.46** Compression is less effective in stereo. In (a) is shown the spatial result of a 'dry' panpotted mix. (b) shows the result after artificial reverberation which can also be obtained in an acoustic recording with coincident mikes. After compression (c) the ambience and reverberation may be reduced or absent. (d) Reverberation may also decay prematurely.

result. The dominant sound sources are reproduced fairly accurately, but what is most striking is that the ambience and reverb between is dramatically reduced or absent, making the decoded sound much drier than the original. It will also be found that the rate of decay of reverberation is accelerated as shown in (d).

These effects are heard because the reverberation exists at a relatively low level. The coder will assume that it is inaudible due to masking and remove or attenuate it. The effect is apparent to the same extent with both

MPEG Layer II and Dolby AC-3 coders even though their internal workings are quite different. This is not surprising because both are probably based on the same psychoacoustic masking model.

MPEG Layer III fares quite badly in stereo because the bit rate is lower. Transient material has a peculiar effect whereby the ambience would come and go according to the entropy of the dominant source. A percussive note would narrow the sound stage and appear dry but afterwards the reverb level would come back up.

The effects are not subtle and do not require 'golden ears' or special source material. The author has successfully demonstrated these effects under conference conditions.

All these effects disappear when the signals to the speakers are added to make mono as this prevents attentional selectivity and unmasking cannot occur. The above compression artifacts are less obvious or even absent when poor-quality loudspeakers are used. Loudspeakers are part of the communication channel and have an information capacity, both timbral and spatial. If the quality of the loudspeaker is poor, it may remove more information from the signal than the preceding compressor and coding artifacts will not be revealed. The precedence effect allows the listener to locate the source of a sound by concentrating on the first version of the sound and rejecting later versions. Versions which may arrive from elsewhere simply add to the perceived loudness but do not change the perceived location of the source. The precedence effect can only reject reverberant sounds which arrive after the inter-aural delay. When reflections arrive within the inter-aural delay of about  $700\mu\text{s}$  the precedence effect breaks down and the perceived direction can be pulled away from that of the first arriving source by an increase in level. Figure 4.47 shows that this area is known as the time-intensity trading region.





Figure 4.47 Time-intensity trading occurs within the inter-aural delay period.

Once the maximum inter-aural delay is exceeded, the hearing mechanism knows that the time difference must be due to reverberation and the trading ceases to change with level.

Unfortunately reflections with delays of the order of  $700\ \mu\text{s}$  are exactly what are provided by the traditional rectangular loudspeaker with flat sides and sharp corners. The discontinuities between the panels cause impedance changes which act as acoustic reflectors. The loudspeaker becomes a multiple source producing an array of signals within the time-intensity trading region and instead of acting as a point source the loudspeaker becomes a distributed source.

Figure 4.48 shows that when a loudspeaker is a distributed source, it cannot create a point image. Instead the image is also distributed, an effect known as smear. Note that the point sources have spread so that there are almost no gaps between them, effectively masking the ambience. If a compressor removes it the effect cannot be heard. It may erroneously be assumed that the compressor is transparent when in fact it is not.



**Figure 4.48** A loudspeaker which is a distributed source cannot produce a point stereo image; only a spatially spread or smeared image.

For the above reasons, the conclusions of many of the listening tests carried out on audio compressors are simply invalid. The author has read descriptions of many such tests, and in almost all cases little is said about the loudspeakers used or about what steps were taken to measure their spatial accuracy. This is simply unscientific and has led to the use of compression being more widespread than is appropriate.

A type of flat panel loudspeaker which uses chaotic bending modes of the diaphragm has recently been introduced. It should be noted that these speakers are distributed sources and so are quite unsuitable for assessing compression quality.

Whilst compression may be adequate to deliver post-produced audio to a consumer with mediocre loudspeakers, these results underline that it has no place in a high-quality production environment. When assessing codecs, loudspeakers having poor diffraction design will conceal arti-

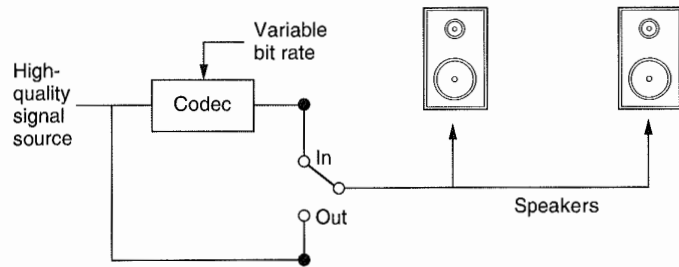


Figure 4.49 Using compression to test loudspeakers. The better the loudspeaker, the less compression can be used before it becomes audible.

facts. When mixing for a compressed delivery system, it will be necessary to include the codec in the monitor feeds so that the results can be compensated. Where high-quality stereo is required, either full bit rate PCM or lossless (packing) techniques must be used.

An interesting corollary of the above is that compressors can be used to test stereo loudspeakers. The arrangement of Figure 4.49 is used. Starting with the highest bit rate possible, the speakers are switched between the source material and the output of the codec. The compression factor is increased until the effects noted above are heard. When the effects are just audible, the bit rate of the compressor is at the information capacity of the loudspeakers. Domestic bookshelf loudspeakers measure approximately 100 kbits/s. Traditional square box hi-fi speakers measure around 200 kbits/s but a pair of state-of-the-art speakers will go above 500 kbits/s.

## References

1. Johnston, J.D., Transform coding of audio signals using perceptual noise criteria. *IEEE J. Selected Areas in Comms.*, JSAC-6, 314-323 (1988)
2. Martin, W.H., Decibel - the new name for the transmission unit. *Bell System Tech. J.* (Jan.1929)
3. Moore, B.C., *An Introduction to the Psychology of Hearing*, section 6. 12, London: Academic Press (1989)
4. Muraoka, T., Iwahara, M. and Yamada, Y., Examination of audio bandwidth requirements for optimum sound signal transmission. *J. Audio Eng. Soc.*, 29, 2-9 (1982)
5. Muraoka, T., Yamada, Y. and Yamazaki, M., Sampling frequency considerations in digital audio. *J. Audio Eng. Soc.*, 26, 252-256 (1978)
6. Fincham, L.R., The subjective importance of uniform group delay at low frequencies. Presented at the 74th Audio Engineering Society Convention (New York, 1983), Preprint 2056(H-1)
7. Fletcher, H., Auditory patterns. *Rev. Modern Physics*, 12, 47-65 (1940)
8. Zwicker, E., Subdivision of the audible frequency range into critical bands. *J. Acoust. Soc. Amer.*, 33, 248 (1961)
9. Grewin, C. and Ryden, T., Subjective assessments on low bit-rate audio codecs. *Proc. 10th. Int. Audio Eng. Soc. Conf.*, 91-102, New York: Audio Eng. Soc. (1991)

10. Gilchrist, N.H.C., Digital sound: the selection of critical programme material and preparation of the recordings for CCIR tests on low bit rate codecs. *BBC Res. Dept. Rep. RD 1993/1*
11. Colomes, C. and Faucon, G., A perceptual objective measurement system (POM) for the quality assessment of perceptual codecs. Presented at 96th Audio Eng. Soc. Conv. Amsterdam (1994), Preprint No. 3801 (P4.2)
12. Johnston, J., Estimation of perceptual entropy using noise masking criteria. *ICASSP*, 2524–2527 (1988)
13. Gilchrist, N.H.C., Delay in broadcasting operations. Presented at 90th Audio Eng. Soc. Conv. (1991), Preprint 3033
14. Caine, C.R., English, A.R. and O'Clarey, J.W.H. NICAM-3: near-instantaneous compressed digital transmission for high-quality sound programmes. *J. IERE*, 50, 519–530 (1980)
15. Davidson, G.A. and Bosi, M., AC-2: High quality audio coding for broadcast and storage. *Proc. 46th Ann. Broadcast Eng. Conf.*, Las Vegas, 98–105 (1992)
16. Crochiere, R.E., Sub-band coding. *Bell System Tech. J.*, 60, 1633–1653 (1981)
17. Princen, J.P., Johnson, A. and Bradley, A.B., Sub-band/transform coding using filter bank designs based on time domain aliasing cancellation. *Proc. ICASSP*, 2161–2164 (1987)
18. Davis, M.F., The AC-3 multichannel coder. Presented at 95th AES Conv., Preprint 2774.
19. Wiese, D., MUSICAM: flexible bitrate reduction standard for high quality audio. Presented at Digital Audio Broadcasting Conference (London, March 1992)
20. Brandenburg, K., ASPEC coding. *Proc. 10th. Audio Eng. Soc. Int. Conf.*, 81–90, New York: Audio Eng. Soc. (1991)
21. ISO/IEC JTC1/SC2/WG11 N0030: MPEG/AUDIO test report. Stockholm (1990)
22. ISO/IEC JTC1/SC2/WG11 MPEG 91/010 The SR report on: The MPEG/AUDIO subjective listening test. Stockholm (1991)
23. ISO/IEC JTC1/SC29/WG11 MPEG, International standard ISO 11172–3 Coding of moving pictures and associated audio for digital storage media up to 1.5 Mbits/s, Part 3: Audio (1992)
24. Brandenburg, K. and Stoll, G., ISO-MPEG-1 Audio: A generic standard for coding of high quality audio. *JAES*, 42, 780–792 (1994)
25. Bonicel, P. *et al.*, A real time ISO/MPEG2 Multichannel decoder. Presented at 96th Audio Eng. Soc. Conv. (1994), Preprint No. 3798 (P3.7)4.30
26. ISO/IEC 13818–7, Information Technology – Generic coding of moving pictures and associated audio, Part 7: Advanced audio coding (1997)
27. Bosi, M. *et al.*, ISO/IEC MPEG-2 Advanced Audio Coding. *JAES*, 45, 789–814 (1997)
28. Herre, J. and Johnston, J.D., Enhancing the performance of perceptual audio coders by using temporal noise shaping (TNS). Presented at 101st AES Conv., Preprint 4384 (1996)
29. Fuchs, H., Improving MPEG audio coding by backward adaptive linear stereo prediction. Presented at 99th AES Conv., Preprint 4086 (1995)
30. Gersho, A., Asymptotically optimal block quantization. *IEEE Trans. Info. Theory*, IT-25, No.4, 373–380 (1979)
31. Cao, Y., Sridharan, S. and Moody, M., Co-talker separation using the cocktail party effect. *J. Audio Eng. Soc.*, 44, No. 12, 1084–1096 (1996)



# 5


## MPEG video compression

---

In this chapter the principles of video compression are explored, leading to descriptions of MPEG-1, MPEG-2 and MPEG-4. MPEG-1 supports only progressively scanned images, whereas MPEG-2 and MPEG-4 support both progressive and interlaced scan. MPEG uses the term 'picture' to mean a full-screen image of any kind at one point on the time axis. This could be a field or a frame in interlaced systems but only a frame in non-interlaced systems. The terms field and frame will be used only when the distinction is important. MPEG-4 introduces object coding which can handle entities that may not fill the screen. In MPEG-4 the picture becomes a plane in which one or more video objects can be displayed, hence the term video object plane (VOP).

### 5.1 The eye

All imaging signals ultimately excite some response in the eye and the viewer can only describe the result subjectively. Familiarity with the functioning and limitations of the eye is essential to an understanding of image compression. The simple representation of Figure 5.1 shows that the eyeball is nearly spherical and is swivelled by muscles. The space between the cornea and the lens is filled with transparent fluid known as *aqueous humour*. The remainder of the eyeball is filled with a transparent jelly known as *vitreous humour*. Light enters the cornea, and the amount of light admitted is controlled by the pupil in the iris. Light entering is involuntarily focused on the retina by the lens in a process called *visual accommodation*. The lens is the only part of the eye which is not nourished by the bloodstream and its centre is technically dead. In a young person the lens is flexible and muscles distort it to perform the focusing action.




**Figure 5.1** The eyeball is in effect a living camera except for the lens which receives no blood supply and is technically dead.

In old age the lens loses some flexibility and causes *presbyopia* or limited accommodation. In some people the length of the eyeball is incorrect resulting in *myopia* (short-sightedness) or *hypermetropia* (long-sightedness). The cornea should have the same curvature in all meridia, and if this is not the case, *astigmatism* results.


The retina is responsible for light sensing and contains a number of layers. The surface of the retina is covered with arteries, veins and nerve fibres and light has to penetrate these in order to reach the sensitive layer. This contains two types of discrete receptors known as *rods* and *cones* from their shape. The distribution and characteristics of these two receptors are quite different. Rods dominate the periphery of the retina whereas cones dominate a central area known as the *fovea* outside which their density drops off. Vision using the rods is monochromatic and has poor resolution but remains effective at very low light levels, whereas the cones provide high resolution and colour vision but require more light. Figure 5.2 shows how the sensitivity of the retina slowly increases in response to entering darkness. The first part of the curve is the adaptation of cone or *photopic* vision. This is followed by the greater adaptation of the rods in *scotopic* vision. At such low light levels the fovea is essentially blind and small objects which can be seen in the peripheral rod vision disappear when stared at.

The cones in the fovea are densely packed and directly connected to the nervous system allowing the highest resolution. Resolution then falls off away from the fovea. As a result the eye must move to scan large areas of



**Figure 5.2** Vision adapts to darkness in two stages known as photopic and scotopic vision.

detail. The image perceived is not just a function of the retinal response, but is also affected by processing of the nerve signals. The overall acuity of the eye can be displayed as a graph of the response plotted against the degree of detail being viewed. Detail is generally measured in lines per millimetre or cycles per picture height, but this takes no account of the distance from the eye. A better unit for eye resolution is one based upon the subtended angle of detail as this will be independent of distance. Units of cycles per degree are then appropriate. Figure 5.3 shows the



**Figure 5.3** The response of the eye to static detail falls off at both low and high spatial frequencies.

response of the eye to static detail. Note that the response to very low frequencies is also attenuated. An extension of this characteristic allows the vision system to ignore the fixed pattern of shadow on the retina due to the nerves and arteries.

The retina does not respond instantly to light, but requires between 0.15 and 0.3 second before the brain perceives an image. The resolution of the eye is primarily a spatio-temporal compromise. The eye is a spatial sampling device; the spacing of the rods and cones on the retina represents a spatial sampling frequency.

The measured acuity of the eye exceeds the value calculated from the sample site spacing because a form of oversampling is used. The eye is in a continuous state of subconscious vibration called saccadic motion. This causes the sampling sites to exist in more than one location, effectively increasing the spatial sampling rate provided there is a temporal filter which is able to integrate the information from the various different positions of the retina. This temporal filtering is responsible for 'persistence of vision'. Flashing lights are perceived to flicker only until the critical flicker frequency (CFF) is reached; the light appears continuous for higher frequencies. The CFF is not constant but changes with brightness (see Figure 5.4). Note that the field rate of European television at 50 fields per second is marginal with bright images.

Figure 5.5 shows the two-dimensional or spatio-temporal response of the eye. If the eye were static, a detailed object moving past it would give rise to temporal frequencies, as Figure 5.6(a) shows. The temporal





Figure 5.4 The critical flicker frequency is not constant but rises as brightness increases.

scotopic

l response,  
erall acuity  
against the  
n lines per  
ount of the  
ased upon  
f distance.  
shows the

50

high spatial



**Figure 5.5** The response of the eye shown with respect to temporal and spatial frequencies. Note that even slow relative movement causes a serious loss of resolution. The eye tracks moving objects to prevent this loss.



frequency is given by the detail in the object, in lines per millimetre, multiplied by the speed. Clearly a highly detailed object can reach high temporal frequencies even at slow speeds, yet Figure 5.5 shows that the eye cannot respond to high temporal frequencies.

However, the human viewer has an interactive visual system which causes the eyes to track the movement of any object of interest. Figure 5.6(b) shows that when eye tracking is considered, a moving object is rendered stationary with respect to the retina so that temporal frequencies fall to zero and much the same acuity to detail is available despite motion. This is known as *dynamic resolution* and it describes how humans judge the detail in real moving pictures.

## 5.2 Dynamic resolution

As the eye uses involuntary tracking at all times, the criterion for measuring the definition of moving-image portrayal systems has to be *dynamic resolution*, defined as the apparent resolution perceived by the viewer in an object moving within the limits of accurate eye tracking. The traditional metric of static resolution in film and television has to be abandoned as unrepresentative of the subjective results.


Figure 5.7(a) shows that when the moving eye tracks an object on the screen, the viewer is watching with respect to the optic flow axis, not the time axis, and these are not parallel when there is motion. The optic flow axis is defined as an imaginary axis in the spatio-temporal volume which joins the same points on objects in successive frames. Clearly when many objects move independently there will be one optic flow axis for each.



**Figure 5.6** In (a) a detailed object moves past a fixed eye, causing temporal frequencies beyond the response of the eye. This is the cause of motion blur. In (b) the eye tracks the motion and the temporal frequency becomes zero. Motion blur cannot then occur.

The optic flow axis is identified by motion-compensated standards converters to eliminate judder and also by MPEG compressors because the greatest similarity from one picture to the next is along that axis. The success of these devices is testimony to the importance of the theory.

According to sampling theory, a sampling system cannot properly convey frequencies beyond half the sampling rate. If the sampling rate is considered to be the picture rate, then no temporal frequency of more than 25 or 30 Hz can be handled (12 Hz for film). With a stationary camera and scene, temporal frequencies can only result from the brightness of lighting changing, but this will not approach the limit. However, when there is relative movement between camera and scene, detailed areas develop high temporal frequencies, just as was shown in Figure 5.6(a) for




**Figure 5.7** The optic flow axis (a) joins points on a moving object in successive pictures. (b) When a tracking eye follows a moving object on a screen, that screen will be seen in a different place at each picture. This is the origin of background strobing.

the eye. This is because relative motion results in a given point on the camera sensor effectively scanning across the scene. The temporal frequencies generated are beyond the limit set by sampling theory, and aliasing should take place.

However, when the resultant pictures are viewed by a human eye, this aliasing is not perceived because, once more, the eye tracks the motion of the scene.<sup>1</sup>

Figure 5.8 shows what happens when the eye follows correctly. Although the camera sensor and display are both moving through the field of view, the original scene and the retina are now stationary with respect to one another. As a result the temporal frequency at the eye due to the object being followed is brought to zero and no aliasing is perceived by the viewer due to sampling in the image-portrayal system.

Whilst this result is highly desirable, it does not circumvent sampling theory because the effect only works if several assumptions are made, including the requirement for the motion to be smooth.



**Figure 5.8** An object moves past a camera, and is tracked on a monitor by the eye. The high temporal frequencies cause aliasing in the TV signal, but these are not perceived by the tracking eye as this reduces the temporal frequency to zero.

Figure 5.7(b) shows that when the eye is tracking, successive pictures appear in different places with respect to the retina. In other words if an object is moving down the screen and followed by the eye, the raster is actually moving up with respect to the retina. Although the tracked object is stationary with respect to the retina and temporal frequencies are zero, the object is moving with respect to the sensor and the display and in those units high temporal frequencies will exist. If the motion of the object on the sensor is not correctly portrayed, dynamic resolution will suffer. Dynamic resolution analysis confirms that both interlaced television and conventionally projected cinema film are both seriously sub-optimal. In contrast, progressively scanned television systems have no such defects.

In real-life eye tracking, the motion of the background will be smooth, but in an image-portrayal system based on periodic presentation of frames, the background will be presented to the retina in a different position in each frame. The retina separately perceives each impression of the background leading to an effect called *background strobing*.

The criterion for the selection of a display frame rate in an imaging system is sufficient reduction of background strobing. It is a complete myth that the display rate simply needs to exceed the critical flicker frequency.

Manufacturers of graphics displays which use frame rates well in excess of those used in film and television are doing so for a valid reason: it gives better results!



The traditional reason against high picture rates is that they require excessive bandwidth. In PCM form this may be true, but there are major exceptions. First, in the MPEG domain, it is the information which has to be sent, not the individual picture. Raising the picture rate does not raise the MPEG bit rate in proportion. This is because pictures which are closer together in time have more redundancy and smaller motion distances. Second, the display rate and the transmission rate need not be the same in an advanced system. Motion-compensated up-conversion at the display can correctly render the background at intermediate positions to those transmitted.

### 5.3 Contrast

The contrast sensitivity of the eye is defined as the smallest brightness difference which is visible. In fact the contrast sensitivity is not constant, but increases proportionally to brightness. Thus whatever the brightness of an object, if that brightness changes by about 1 per cent it will be equally detectable.

The true brightness of a television picture can be affected by electrical noise on the video signal. As contrast sensitivity is proportional to brightness, noise is more visible in dark picture areas than in bright areas. In practice the gamma characteristic of the CRT is put to good use in making video noise less visible. Instead of having linear video signals which are subjected to an inverse gamma function immediately prior to driving the CRT, the inverse gamma correction is performed at the camera. In this way the video signal is non-linear for most of its journey.

Figure 5.9 shows a reverse gamma function. As a true power function requires infinite gain near black, a linear segment is substituted. It will be seen that contrast variations near black result in larger signal amplitude than variations near white. The result is that noise picked up by the video signal has less effect on dark areas than on bright areas. After the gamma of the CRT has acted, noise near black is compressed with respect to noise near white. Thus a video transmission system using gamma correction at source has a better perceived noise level than if the gamma correction is performed near the display.

In practice the system is not rendered perfectly linear by gamma correction and a slight overall exponential effect is usually retained in order further to reduce the effect of noise in darker parts of the picture. A gamma correction factor of 0.45 may be used to achieve this effect. If another type of display is to be used with signals designed for CRTs, the gamma characteristic of that display will probably be different and some gamma conversion will be required.




Figure 5.9 CCIR Rec.709 reverse gamma function used at camera has a straight line approximation at the lower part of the curve to avoid boosting camera noise. Note that the output amplitude is greater for modulation near black.





Figure 5.10 The sensitivity of the eye to noise is greatest at low frequencies and drops rapidly with increasing frequency. This can be used to mask quantizing noise caused by the compression process.


Sensitivity to noise is also a function of spatial frequency. Figure 5.10 shows that the sensitivity of the eye to noise falls with frequency from a maximum at zero. Thus it is vital that the average brightness should be correctly conveyed by a compression system, whereas higher spatial frequencies can be subject to more requantizing noise. Transforming the image into the frequency domain allows this characteristic to be explored.

## 5.4 Colour vision

Colour vision is due to the cones on the retina which occur in three different types, responding to different colours. Figure 5.11 shows that human vision is restricted to range of light wavelengths from 400



(a)



(b)

**Figure 5.11** (a) The response of the eye to various wavelengths. There is a pronounced peak in the green region. (b) The different types of cone in the eye have the approximate responses shown here which allow colour vision.

nanometres to 700 nanometres. Shorter wavelengths are called ultra-violet and longer wavelengths are called infra-red. Note that the response is not uniform, but peaks in the area of green. The response to blue is very poor and makes a nonsense of the use of blue lights on emergency vehicles which owe much to tradition and little to psycho-optics.

Figure 5.11 shows an approximate response for each of the three types of cone. If light of a single wavelength is observed, the relative responses of the three sensors allows us to discern what we call the colour of the light. Note that at both ends of the visible spectrum there are areas in which only one receptor responds; all colours in those areas look the

same. There is a great deal of variation in receptor response from one individual to the next and the curves used in television are the average of a great many tests. In a surprising number of people the single receptor zones are extended and discrimination between, for example, red and orange is difficult.

The triple receptor characteristic of the eye is extremely fortunate as it means that we can generate a range of colours by adding together light sources having just three different wavelengths in various proportions. This process is known as *additive colour matching* which should be clearly distinguished from the subtractive colour matching that occurs with paints and inks. Subtractive matching begins with white light and selectively removes parts of the spectrum by filtering. Additive matching uses coloured light sources which are combined.


## 5.5 Colour difference signals

An effective colour television system can be made in which only three pure or single wavelength colours or *primaries* can be generated. The primaries need to be similar in wavelength to the peaks of the three receptor responses, but need not be identical. Figure 5.12 shows a rudimentary colour television system. Note that the colour camera is in fact three cameras in one, where each is fitted with a different coloured filter. Three signals, *R*, *G* and *B* must be transmitted to the display which produces three images that must be superimposed to obtain a colour picture.

A monochrome camera produces a single luminance signal *Y* whereas a colour camera produces three signals, or *components*, *R*, *G* and *B* which are essentially monochrome video signals representing an image in each primary colour. *RGB* and *Y* signals are incompatible, yet when colour television was introduced it was a practical necessity that it should be possible to display colour signals on a monochrome display and vice versa. Creating or *transcoding* a luminance signal from *R*, *G* and *B* is relatively easy. Figure 5.11 showed the spectral response of the eye which has a peak in the green region. Green objects will produce a larger stimulus than red objects of the same brightness, with blue objects producing the least stimulus. A luminance signal can be obtained by adding *R*, *G* and *B* together, not in equal amounts, but in a sum which is *weighted* by the relative response of the eye. Thus:

$$Y = 0.3R + 0.59G + 0.11B$$

If *Y* is derived in this way, a monochrome display will show nearly the same result as if a monochrome camera had been used in the first place.



$$\begin{bmatrix} Y \\ \text{Cr-128} \\ \text{Cr-128} \end{bmatrix} = \begin{bmatrix} \frac{77}{256} & \frac{150}{256} & \frac{29}{256} \\ \frac{131}{256} & -\frac{110}{256} & -\frac{21}{256} \\ -\frac{44}{256} & -\frac{67}{256} & \frac{131}{256} \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

(b)

**Figure 5.12** (a) A simple colour system uses three primaries and transmits a complete picture for each. This is incompatible with monochrome and uses too much bandwidth. Practical systems use colour difference and luminance signals which are obtained by a weighted calculation as shown in (b).

The results are not identical because of the non-linearities introduced by gamma correction.

As colour pictures require three signals, it should be possible to send  $Y$  and two other signals which a colour display could arithmetically convert back to  $R$ ,  $G$  and  $B$ . There are two important factors which restrict the form which the other two signals may take. One is to achieve reverse compatibility. If the source is a monochrome camera, it can only produce  $Y$  and the other two signals will be completely absent. A colour display should be able to operate on the  $Y$  signal only and show a monochrome picture. The other is the requirement to conserve bandwidth for economic reasons.




Figure 5.13 Colour components are converted to colour difference signals by the transcoding shown here.

These requirements are met by sending two *colour difference signals* along with  $Y$ . There are three possible colour difference signals,  $R-Y$ ,  $B-Y$  and  $G-Y$ . As the green signal makes the greatest contribution to  $Y$ , then the amplitude of  $G-Y$  would be the smallest and would be most susceptible to noise. Thus  $R-Y$  and  $B-Y$  are used in practice as Figure 5.13 shows.

$R$  and  $B$  are readily obtained by adding  $Y$  to the two colour difference signals.  $G$  is obtained by rearranging the expression for  $Y$  above such that:

$$G = \frac{Y - 0.3R - 0.11B}{0.59}$$

If a colour CRT is being driven, it is possible to apply inverted luminance to the cathodes and the  $R-Y$  and  $B-Y$  signals directly to two of the grids so that the tube performs some of the matrixing. It is then only necessary to obtain  $G-Y$  for the third grid, using the expression:

$$G-Y = 0.51(R-Y) - 0.186(B-Y)$$

If a monochrome source having only a  $Y$  output is supplied to a colour display,  $R-Y$  and  $B-Y$  will be zero. It is reasonably obvious that if there are no colour difference signals the colour signals cannot be different from one another and  $R = G = B$ . As a result the colour display can produce only a neutral picture.

The use of colour difference signals is essential for compatibility in both directions between colour and monochrome, but it has a further advantage that follows from the way in which the eye works. In order to produce the highest resolution in the fovea, the eye will use signals from

all types of cone, regardless of colour. In order to determine colour the stimuli from three cones must be compared. There is evidence that the nervous system uses some form of colour difference processing to make this possible. As a result the full acuity of the human eye is only available in monochrome. Differences in colour cannot be resolved so well. A further factor is that the lens in the human eye is not achromatic and this means that the ends of the spectrum are not well focused. This is particularly noticeable on blue.

If the eye cannot resolve colour very well there is no point in expending valuable bandwidth sending high-resolution colour signals. Colour difference working allows the luminance to be sent separately at a bit rate which determines the subjective sharpness of the picture. The colour difference signals can be sent with considerably reduced bit rate, as little as one quarter that of luminance, and the human eye is unable to tell.

The overwhelming advantages obtained by using downsampled colour difference signals mean that in broadcast and production facilities their use has become almost universal. The technique is equally attractive for compression applications and is retained in MPEG. The outputs from the RGB sensors in most cameras are converted directly to  $Y$ ,  $R-Y$  and  $B-Y$  in the camera control unit and output in that form. Whilst signals such as  $Y$ ,  $R$ ,  $G$  and  $B$  are *unipolar* or positive only, it should be stressed that colour difference signals are *bipolar* and may meaningfully take on levels below zero volts.

The downsampled colour formats used in MPEG were described in section 2.9.

## 5.6 Progressive or interlaced scan?

Analog video samples in the time domain and vertically down the screen so a two-dimensional vertical/temporal sampling spectrum will result. In a progressively scanned system there is a rectangular matrix of sampling sites vertically and temporally. The rectangular sampling structure of progressive scan is *separable* which means that, for example, a vertical manipulation can be performed on frame data without affecting the time axis. The sampling spectrum will be obtained according to section 2.6 and consists of the baseband spectrum repeated as sidebands above and below harmonics of the two-dimensional sampling frequencies. The corresponding spectrum is shown in Figure 5.14. The baseband spectrum is in the centre of the diagram, and the repeating sampling sideband spectrum extends vertically and horizontally. The vertical aspects of the star-shaped spectrum result from vertical spatial frequencies in the image. The horizontal aspect is due to image movement. Note that the star shape is rather hypothetical; the actual shape depends heavily on the




Figure 5.14 The vertical/temporal spectrum of a non-interlaced video signal.

source material. On a still picture the horizontal dimensions collapse to a line structure. In order to return a progressive scan video signal to a continuous moving picture, a two-dimensional low-pass filter having a rectangular response is required. This is quite feasible as persistence of vision acts as a temporal filter and by sitting far enough away from the screen the finite acuity of the eye acts as a spatial reconstruction filter.

Interlace is actually a primitive form of compression in which the system bandwidth is typically halved by sending only half the frame lines in the first field, with the remaining lines being sent in the second field. The use of interlace has a profound effect on the vertical/temporal spectrum. Figure 5.15 shows that the lowest sampling frequency on the time axis is the frame rate, and the lowest sampling frequency on the vertical axis is the number of lines in a field. The arrangement is called a quincunx pattern because of the similarity to the five of dice. The triangular passband has exactly half the area of the rectangular passband of Figure 5.14 illustrating that half the information rate is available.

As a consequence of the triangular passband of an interlaced signal, if the best vertical frequency response is to be obtained, no motion is allowed. It should be clear from Figure 5.16 that a high vertical spatial frequency resulting from a sharp horizontal edge in the picture is only repeated at frame rate, resulting in an artifact known as *interlace twitter*. Conversely, to obtain the best temporal response, the vertical resolution must be impaired. Thus interlaced systems have poor resolution on moving images, in other words their *dynamic resolution* is poor.

In order to return to a continuous signal, a quincuncial spectrum requires a triangular spatio-temporal low-pass filter. In practice no such filter can be realized. Consequently the sampling sidebands are not






Figure 5.15 The vertical temporal spectrum of monochrome video due to interlace.


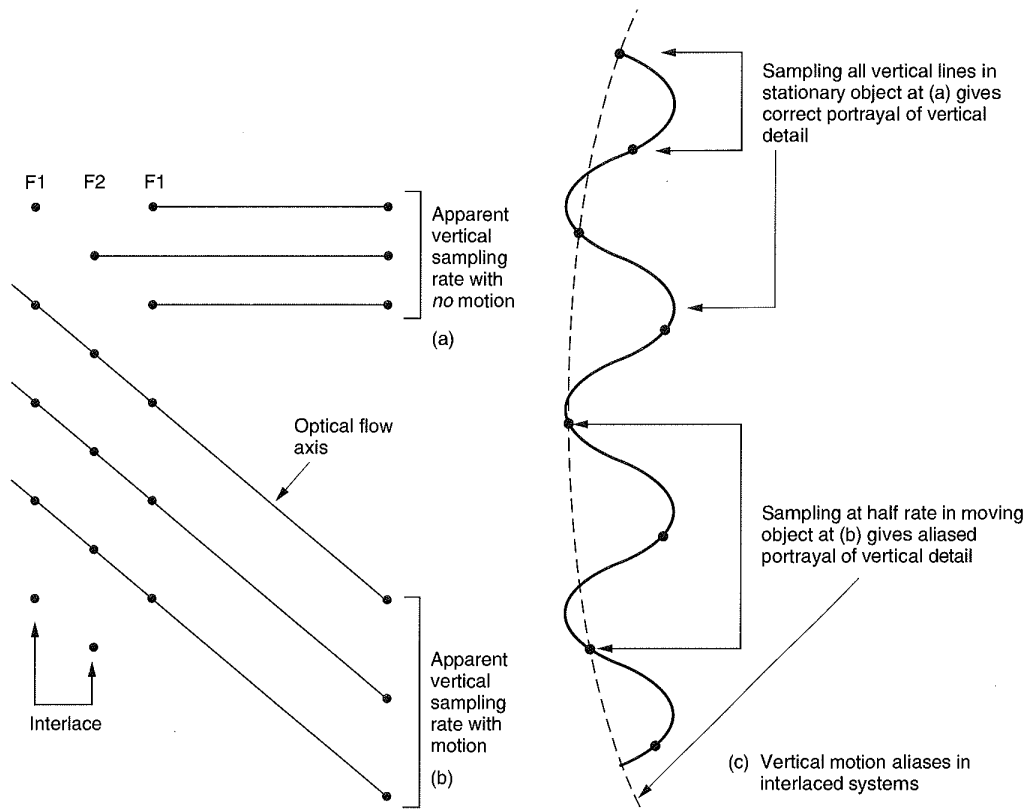


Figure 5.16 In an interlaced system, a horizontal edge is present in *one* field only and so is presented to the viewer at frame rate not field rate. This is below the critical flicker frequency and is visible as *twitter*.



**Figure 5.17** When an interlaced picture is stationary, viewing takes place along the time axis as shown in (a). When a vertical component of motion exists, viewing takes place along the optic flow axis. (b) The vertical sampling rate falls to one half its stationary value. (c) The halving in sampling rate causes high spatial frequencies to alias.

filtered out and are visible, particularly the frame rate component. This artifact is visible on an interlaced display even at distances so great that resolution cannot be assessed.

Figure 5.17(a) shows a dynamic resolution analysis of interlaced scanning. When there is no motion, the optic flow axis and the time axis are parallel and the apparent vertical sampling rate is the number of lines in a frame. However, when there is vertical motion, (b), the optic flow axis turns. In the case shown, the sampling structure due to interlace results in the vertical sampling rate falling to one half of its stationary value.

Consequently interlace does exactly what would be expected from a half-bandwidth filter. It halves the vertical resolution when any motion with a vertical component occurs. In a practical television system, there is no anti-aliasing filter in the vertical axis and so when the vertical sampling rate of an interlaced system is halved by motion, high spatial

frequencies will alias or heterodyne causing annoying artifacts in the picture. This is easily demonstrated.

Figure 5.17(c) shows how a vertical spatial frequency well within the static resolution of the system aliases when motion occurs. In a progressive scan system this effect is absent and the dynamic resolution due to scanning can be the same as the static case.

This analysis also illustrates why interlaced television systems must have horizontal raster lines. This is because in real life, horizontal motion is more common than vertical. It is easy to calculate the vertical image motion velocity needed to obtain the half-bandwidth speed of interlace, because it amounts to one raster line per field. In 525/60 (NTSC) there are about 500 active lines, so motion as slow as one picture height in 8 seconds will halve the dynamic resolution. In 625/50 (PAL) there are about 600 lines, so the half-bandwidth speed falls to one picture height in 12 seconds. This is why NTSC, with fewer lines and lower bandwidth, doesn't look as soft as it should compared to PAL, because it actually has better dynamic resolution.

The situation deteriorates rapidly if an attempt is made to use interlaced scanning in systems with a lot of lines. In 1250/50, the resolution is halved at a vertical speed of just one picture height in 24 seconds. In other words on real moving video a 1250/50 interlaced system has the same dynamic resolution as a 625/50 progressive system. By the same argument a 1080 I system has the same performance as a 480 P system.

Interlaced signals are not separable and so processes which are straightforward in progressively scanned systems become more complex in interlaced systems. Compression systems should not be cascaded indiscriminately, especially if they are different. As digital compression techniques based on transforms are now available, it makes no sense to use an interlaced, i.e. compressed, video signal as an input.

Interlaced signals are harder for MPEG to compress.<sup>2</sup> The confusion of temporal and spatial information makes accurate motion estimation more difficult and this reflects in a higher bit rate being required for a given quality. In short, how can a motion estimator accurately measure motion from one field to another when differences between the fields can equally be due to motion, vertical detail or vertical aliasing?

Computer-generated images and film are not interlaced, but consist of discrete frames spaced on a time axis. As digital technology is bringing computers and television closer the use of interlaced transmission is an embarrassing source of incompatibility. The future will bring image-delivery systems based on computer technology and oversampling cameras and displays which can operate at resolutions much closer to the theoretical limits.

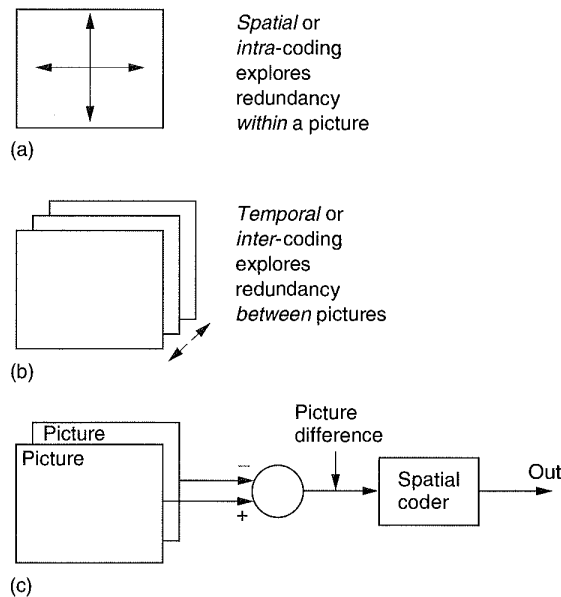
Interlace was the best that could be managed with thermionic valve technology sixty years ago, and we should respect the achievement of its

developers at a time when things were so much harder. However, we must also recognize that the context in which interlace made sense has disappeared.

## 5.7 Spatial and temporal redundancy in MPEG

Chapter 1 introduced these concepts in a general sense and now they will be treated with specific reference to MPEG. Figure 5.18(a) shows that spatial redundancy is redundancy within a single picture or object, for example repeated pixel values in a large area of blue sky. Temporal redundancy (b) exists between successive pictures or objects.

In MPEG, where temporal compression is used, the current picture/object is not sent in its entirety; instead the difference between the current picture/object and the previous one is sent. The decoder already has the previous picture/object, and so it can add the difference, or residual image, to make the current picture/object. A residual image is created by subtracting every pixel in one picture/object from the corresponding pixel in another. This is trivially easy when pictures are restricted to progressive scan, as in MPEG-1, but MPEG-2 had to develop greater complexity (continued in MPEG-4) so that this can also be done with interlaced pictures. The handling of interlace in MPEG will be detailed later.



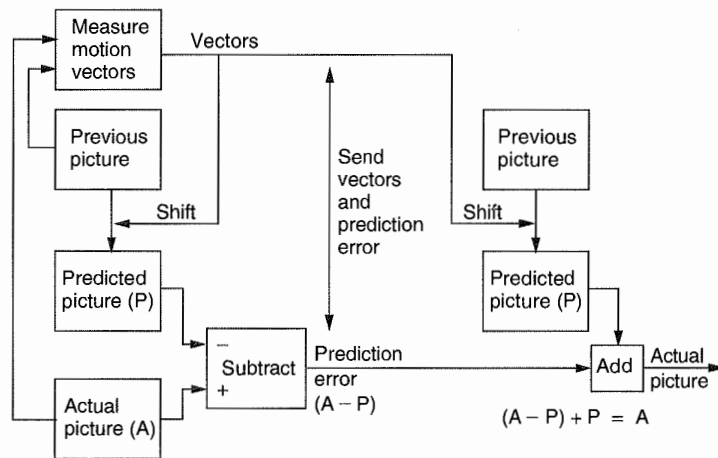
**Figure 5.18** (a) Spatial or intra-coding works on individual images. (b) Temporal or inter-coding works on successive images. (c) In MPEG inter-coding is used to create difference images. These are then compressed spatially.

A residual is an image of a kind, although not a viewable one, and so should contain some kind of spatial redundancy. Figure 5.18(c) shows that MPEG takes advantage of both forms of redundancy. Residual images are spatially compressed prior to transmission. At the decoder the spatial compression is decoded to re-create the residual, then this is added to the previous picture/object to complete the decoding process.

Whenever objects move they will be in a different place in successive pictures. This will result in large amounts of difference data. MPEG overcomes the problem using motion compensation. The encoder contains a motion estimator which measures the direction and distance of motion between pictures and outputs these as vectors which are sent to the decoder. When the decoder receives the vectors it uses them to shift data in a previous picture to more closely resemble the current picture. Effectively the vectors are describing the optic flow axis of some moving screen area, along which axis the image is highly redundant. Vectors are bipolar codes which determine the amount of horizontal and vertical shift required.

In real images, moving objects do not necessarily maintain their appearance as they move. For example, objects may move into shade or light. Consequently motion compensation can never be ideal and it is still necessary to send a residual to make up for any shortcomings in the motion compensation.

Figure 5.19 shows how this works in MPEG-1 and MPEG-2. In addition to the motion-encoding system, the coder also contains a motion decoder. When the encoder outputs motion vectors, it also uses them locally in the



**Figure 5.19** A motion-compensated compression system. The coder calculates motion vectors which are transmitted as well as being used locally to create a predicted picture. The difference between the predicted picture and the actual picture is transmitted as a prediction error.

ne, and so  
(c) shows

Residual  
ecoder the  
en this is  
g process.  
successive  
ta. MPEG  
: encoder  
istance of  
re sent to  
m to shift  
it picture.  
ie moving  
ectors are  
rtical shift

tain their  
shade or  
d it is still  
gs in the

n addition  
n decoder.  
ally in the

ctual  
cture

s motion  
ed picture.  
tted as a



same way that a real decoder will, and is able to produce a *predicted picture* based solely on the previous picture shifted by motion vectors. This is then subtracted from the *actual* current picture to produce a *prediction error* or *residual* which is an image of a kind that can be spatially compressed. The decoder takes the previous picture, shifts it with the vectors to re-create the predicted picture and then decodes and adds the prediction error to produce the actual picture. Picture data sent as vectors plus prediction error are said to be *P* coded.

The concept of sending a prediction error is a useful approach because it allows both the motion estimation and compensation to be imperfect. An ideal motion compensation system will send just the right amount of vector data. With insufficient vector data, the prediction error will be large, but transmission of excess vector data will also cause the the bit rate to rise. There will be an optimum balance which minimizes the sum of the prediction error data and the vector data.


In MPEG-1 and MPEG-2 the balance is obtained by dividing the screen into areas called *macroblocks* which are 16 luminance pixels square. Each macroblock is associated with a vector. The vector has horizontal and vertical components so that, used in combination, these can shift pixel data in any direction. The location of the boundaries of a macroblock are fixed with respect to the display screen and so clearly the vector cannot move the macroblock. Instead the vector tells the decoder where to look in another picture to find pixel data which will be fetched to the macroblock. Figure 5.20(a) shows this concept. The shifting process is generally done by modifying the read address of a RAM using the vector. This can shift by one-pixel steps. MPEG-2 vectors have half-pixel resolution so it is necessary to interpolate between pixels from RAM to obtain half-pixel shifted values. Chapter 3 showed how interpolations of this kind can be performed.

MPEG-1 and MPEG-2 also use a simple compression scheme on the vectors. Figure 5.21 shows that in the case of a large moving object, all the vectors within the object will be identical or nearly so. A horizontal run of macroblocks can be associated to form a structure called a *slice* (see section 5.14). The first vector of a slice is sent as an absolute value, but subsequent vectors will be sent as differences. No advantage is taken of vertical vector redundancy.

Real moving objects will not coincide with macroblocks and so the motion compensation will not be ideal but the prediction error makes up for any shortcomings. Figure 5.20(b) shows the case where the boundary of a moving object bisects a macroblock. If the system measures the moving part of the macroblock and sends a vector, the decoder will shift the entire block making the stationary part wrong. If no vector is sent, the moving part will be wrong. Both approaches are legal in MPEG-1 and MPEG-2 because the prediction error compensates for the incorrect



**Figure 5.20** (a) In motion compensation, pixel data are brought to a fixed macroblock in the target picture from a variety of places in another picture. (b) Where only part of a macroblock is moving, motion compensation is non-ideal. The motion can be coded (c), causing a prediction error in the background, or the background can be coded (d) causing a prediction error in the moving object.



**Figure 5.21** In a slice, the first vector is sent as an absolute value, whereas subsequent vectors are sent as differences. When a large moving object is to be coded, several macroblocks may have the same velocity and this is efficiently coded in a slice.

values. An intelligent coder might try both approaches to see which required the least prediction error data.

The prediction error concept also allows the use of simple but inaccurate motion estimators in low-cost systems. The greater prediction error data are handled using a higher bit rate. On the other hand if a precision motion estimator is available, a very high compression factor may be achieved because the prediction error data are minimized. MPEG does not specify how motion is to be measured; it simply defines how a decoder will interpret the vectors. Encoder designers are free to use any motion-estimation system provided that the right vector protocol is created. Chapter 3 contrasted a number of motion estimation techniques.


Figure 5.22(a) shows that a macroblock contains both luminance and colour difference data at different resolutions. Most of the MPEG-2 Profiles use a 4:2:0 structure which means that the colour is downsampled by a factor of two in both axes. Thus in a  $16 \times 16$  pixel block, there are only  $8 \times 8$  colour difference sampling sites. MPEG-2 is based upon the  $8 \times 8$  DCT (see section 3.7) and so the  $16 \times 16$  block is the screen area which contains an  $8 \times 8$  colour difference sampling block. Thus in 4:2:0 in each macroblock there are four luminance DCT blocks, one  $R-Y$  DCT block and one  $B-Y$  DCT block, all steered by the same vector.

In the 4:2:2 Profile of MPEG-2, shown in Figure 5.22(b), the chroma is not downsampled vertically, and so there is twice as much chroma data in each macroblock which is otherwise substantially the same.

macroblock in  
part of a  
coded (c),  
l (d)

→  
↓  
Sync  
↓  
k

subsequent  
verbal  
e.



**Figure 5.22** The structure of a macroblock. (A macroblock is the screen area steered by one vector.) (a) In 4:2:0, there are two chroma DCT blocks per macroblock whereas in 4:2:2 (b) there are four; 4:2:2 needs 33 per cent more data than 4:2:0.



In MPEG-4 a number of refinements are made to motion compensation in addition to the MPEG-2 tools which are still available. One of these tools is a more complex form of vector compression which will be described in section 5.21. This cuts down the bit rate needed for a given set of vector data and has a number of benefits. By retaining the existing macroblock size, a higher compression factor results from the reduced vector data. However, the opportunity then arises to increase the vector density. Optionally, MPEG-4 allows one vector to be associated with an  $8 \times 8$  DCT block rather than with a  $16 \times 16$  macroblock. This allows much more precise prediction of complex motion, especially near the edge of objects, and will result in a reduced residual bit rate. On a macroblock-by-macroblock basis an MPEG-4 coder can decide whether to code four vectors per macroblock or only one. When the compressed vector data and residual are combined, one of these approaches will result in the least data.

## 5.8 I and P coding

Predictive (*P*) coding cannot be used indefinitely, as it is prone to error propagation. A further problem is that it becomes impossible to decode the transmission if reception begins part-way through. In real video signals, cuts or edits can be present across which there is little redundancy and which make motion estimators throw up their hands.

In the absence of redundancy over a cut, there is nothing to be done but to send the new picture information in absolute form. This is called *I* coding where *I* is an abbreviation of *intra* coding. As *I* coding needs no previous picture for decoding, then decoding can begin at *I* coded information.

MPEG is effectively a toolkit. The bigger the MPEG-number, the more tools are available. However, there is no compulsion to use all the tools available. Thus an encoder may choose whether to use *I* or *P* coding, either once and for all or dynamically on a macroblock-by-macroblock basis. For practical reasons, an entire frame may be encoded as *I* macroblocks periodically. This creates a place where the bitstream might be edited or where decoding could begin.

Figure 5.23 shows a typical application of the Simple Profile of MPEG-2. Periodically an *I* picture is created. Between *I* pictures are *P* pictures which are based on the picture before. These *P* pictures predominantly contain macroblocks having vectors and prediction errors. However, it is perfectly legal for *P* pictures to contain *I* macroblocks. This might be useful where, for example, a camera pan introduces new material at the edge of the screen which cannot be created from an earlier picture.




Figure 5.23 A Simple Profile MPEG-2 signal may contain periodic  $I$  pictures with a number of  $P$  pictures between.

Note that although what is sent is called a  $P$  picture, it is not a picture at all. It is a set of instructions to convert the previous picture into the current picture. If the previous picture is lost, decoding is impossible.

An  $I$  picture together with all of the pictures before the next  $I$  picture form a *group of pictures* (GOP).

## 5.9 Bidirectional coding

Motion-compensated predictive coding is a useful compression technique, but it does have the drawback that it can only take data from a previous picture.

Where moving objects reveal a background this is completely unknown in previous pictures and forward prediction fails. However, more of the background is visible in later pictures. Figure 5.24 shows the concept. In


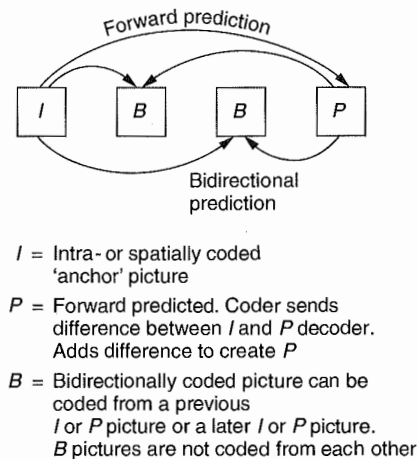


Figure 5.24 In bidirectional coding the revealed background can be efficiently coded by bringing data back from a future picture.



**Figure 5.25** In bidirectional coding, a number of *B* pictures can be inserted between periodic forward predicted pictures. See text.


the centre of the diagram, a moving object has revealed some background. The previous picture can contribute nothing, whereas the next picture contains all that is required.

Bidirectional coding is shown in Figure 5.25. A bidirectional or *B* macroblock can be created using a combination of motion compensation and the addition of a prediction error. This can be done by forward prediction from a previous picture or backward prediction from a subsequent picture. It is also possible to use an average of both forward and backward prediction. On noisy material this may result in some reduction in bit rate. The technique is also a useful way of portraying a dissolve.

The averaging process in MPEG is a simple linear interpolation which works well when only one *B* picture exists between the reference pictures before and after. A larger number of *B* pictures would require weighted interpolation but MPEG does not support this.

Typically two *B* pictures are inserted between *P* pictures or between *I* and *P* pictures. As can be seen, *B* pictures are never predicted from one another, only from *I* or *P* pictures. A typical GOP for broadcasting purposes might have the structure *IBBPBBPBBPBB*. Note that the last *B* pictures in the GOP require the *I* picture in the next GOP for decoding and so the GOPs are not truly independent. Independence can be obtained by creating a *closed GOP* which may contain *B* pictures but which ends with a *P* picture. It is also legal to have a *B* picture in which every macroblock is forward predicted, needing no future picture for decoding.

Bidirectional coding is very powerful. Figure 5.26 is a constant quality curve showing how the bit rate changes with the type of coding. On the



**Figure 5.26** Bidirectional coding is very powerful as it allows the same quality with only 40 per cent of the bit rate of intra-coding. However, the encoding and decoding delays must increase. Coding over a longer time span is more efficient but editing is more difficult.

left, only *I* or spatial coding is used, whereas on the right a *IBBP* structure is used. This means that there are two bidirectionally coded pictures in between a spatially coded picture (*I*) and a forward predicted picture (*P*). Note how for the same quality the system which only uses spatial coding needs two and a half times the bit rate that the bidirectionally coded system needs.

Clearly information in the future has yet to be transmitted and so is not normally available to the decoder. MPEG gets around the problem by sending pictures in the wrong order. Picture reordering requires delay in the encoder and a delay in the decoder to put the order right again. Thus the overall codec delay must rise when bidirectional coding is used. This is quite consistent with Figure 1.5 which showed that as the compression factor rises the latency must also rise.

Figure 5.27 shows that although the original picture sequence is *IBBPBBPBBIBB ...*, this is transmitted as *IPBBPBBIBB ...* so that the future picture is already in the decoder before bidirectional decoding begins. Note that the *I* picture of the next GOP is actually sent before the last *B* pictures of the current GOP.

Figure 5.27 also shows that the amount of data required by each picture is dramatically different. *I* pictures have only spatial redundancy and so need a lot of data to describe them. *P* pictures need fewer data because they are created by shifting the *I* picture with vectors and then adding a prediction error picture. *B* pictures need the least data of all because they can be created from *I* or *P*.

With pictures requiring a variable length of time to transmit, arriving in the wrong order, the decoder needs some help. This takes the form of picture-type flags and time stamps which will be described in section 6.2.




Figure 5.27 Comparison of pictures before and after compression showing sequence change and varying amount of data needed by each picture type. *I*, *P*, *B* pictures use unequal amounts of data.

## 5.10 Coding applications

Figure 5.28 shows a variety of possible GOP (group of pictures in MPEG-1 and MPEG-2) or GOV (group of video object planes in MPEG-4) structures. The simplest is the *III*... sequence in which every picture (or object in MPEG-4) is intra-coded. These can be fully decoded without reference to any other picture or object and so editing is straightforward. However, this approach requires about two-and-one-half times the bit rate of a full bidirectional system.

Bidirectional coding is most useful for final delivery of post-produced material either by broadcast or on prerecorded media as there is then no editing requirement. As a compromise the *IBIB*... structure can be used which has some of the bit rate advantage of bidirectional coding but without too much latency. It is possible to edit an *IBIB* stream by performing some processing. If it is required to remove the video following a *B* picture, that *B* picture could not be decoded because it needs *I* pictures either side of it for bidirectional decoding. The solution

<i>IIIIII</i> ...	<i>I</i> only freely editable, needs high bit rate
<i>IPPPPI</i> ...	Forward predicted only, needs less decoder memory, used in Simple Profile
<i>IBBPBBPB</i> ...	Forward and bidirectional, best compression factor, needs large decoder memory, hard to edit
<i>IBIBIB</i> ...	Lower bit rate than <i>I</i> only, editable with moderate processing

Figure 5.28 Various possible GOP structures used with MPEG. See text for details.

is to decode the *B* picture first, and then re-encode it with forward prediction only from the previous *I* picture. The subsequent *I* picture can then be replaced by an edit process. Some quality loss is inevitable in this process but this is acceptable in applications such as ENG and industrial video.

## 5.11 Intra-coding

Intra-coding or spatial compression in MPEG is used in *I* pictures on actual picture data and in *P* and *B* pictures on prediction error data. MPEG-1 and MPEG-2 use the discrete cosine transform described in section 3.13. In still pictures, MPEG-4 may also use the wavelet transform described in section 3.14.

Entering the spatial frequency domain has two main advantages. It allows dominant spatial frequencies which occur in real images to be efficiently coded, and it allows noise shaping to be used. As the eye is not uniformly sensitive to noise at all spatial frequencies, dividing the information up into frequency bands allows a different noise level to be produced in each.

The DCT works on blocks and in MPEG these are  $8 \times 8$  pixels. Section 5.7 showed how the macroblocks of the motion compensation structure are designed so they can be broken down into  $8 \times 8$  DCT blocks. In a 4:2:0 macroblock there will be six DCT blocks whereas in a 4:2:2 macroblock there will be eight.

Figure 5.29 shows the table of basis functions or *wave table* for an  $8 \times 8$  DCT. Adding these two-dimensional waveforms together in different proportions will give any original  $8 \times 8$  pixel block. The coefficients of the DCT simply control the proportion of each wave which is added in the inverse transform. The top-left wave has no modulation at all because it conveys the DC component of the block. This coefficient will be a unipolar (positive only) value in the case of luminance and will typically be the largest value in the block as the spectrum of typical video signals is dominated by the DC component.

Increasing the DC coefficient adds a constant amount to every pixel. Moving to the right the coefficients represent increasing horizontal spatial frequencies and moving downwards the coefficients represent increasing vertical spatial frequencies. The bottom-right coefficient represents the highest diagonal frequencies in the block. All these coefficients are bipolar, where the polarity indicates whether the original spatial waveform at that frequency was inverted.

Figure 5.30 shows a one-dimensional example of an inverse transform. The DC coefficient produces a constant level throughout the pixel block. The remaining waves in the table are AC coefficients. A zero coefficient




mitted

g sequence  
ctures use

es in MPEG-  
in MPEG-4)  
y picture (or  
led without  
ghtforward.  
imes the bit

st-produced  
e is then no  
can be used  
coding but  
stream by  
e the video  
d because it  
The solution


r details.



**Figure 5.29** The discrete cosine transform breaks up an image area into discrete frequencies in two dimensions. The lowest frequency can be seen here at the top-left corner. Horizontal frequency increases to the right and vertical frequency increases downwards.

would result in no modulation, leaving the DC level unchanged. The wave next to the DC component represents the lowest frequency in the transform which is half a cycle per block. A positive coefficient would make the left side of the block brighter and the right side darker whereas a negative coefficient would do the opposite. The magnitude of the coefficient determines the amplitude of the wave which is added. Figure 5.30 also shows that the next wave has a frequency of one cycle per block. i.e. the block is made brighter at both sides and darker in the middle. Consequently an inverse DCT is no more than a process of mixing various pixel patterns from the wave table where the relative amplitudes and polarity of these patterns are controlled by the coefficients. The original transform is simply a mechanism which finds the coefficient amplitudes from the original pixel block.

The DCT itself achieves no compression at all. Sixty-four pixels are converted to sixty-four coefficients. However, in typical pictures, not all coefficients will have significant values; there will often be a few dominant coefficients. The coefficients representing the higher two-



rete  
top-left  
reases

nged. The  
ncy in the  
ent would  
r whereas  
de of the  
ed. Figure  
per block.  
ie middle.  
of mixing  
mplitudes  
ients. The  
coefficient

pixels are  
es, not all  
be a few  
gher two-

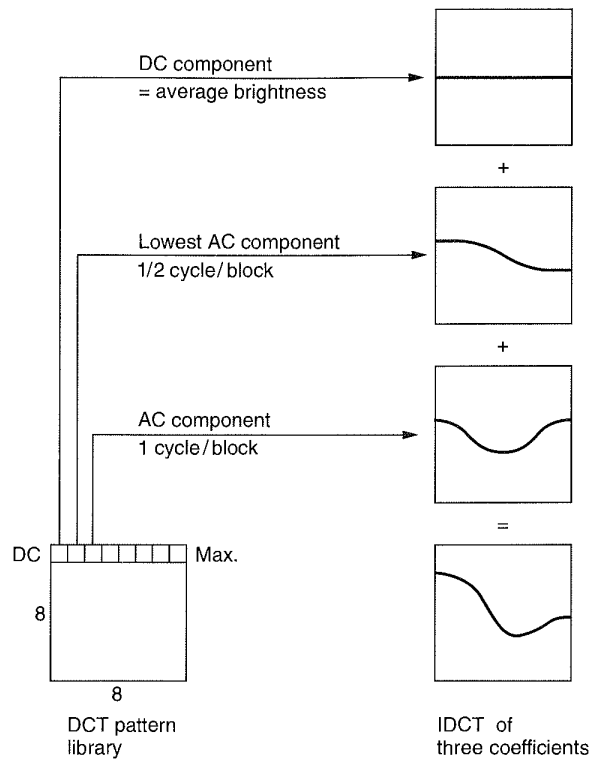


Figure 5.30 A one-dimensional inverse transform. See text for details.

dimensional spatial frequencies will often be zero or of small value in large areas, due to blurring or simply plain undetailed areas before the camera.

Statistically, the further from the top-left corner of the wave table the coefficient is, the smaller will be its magnitude. Coding gain (the technical term for reduction in the number of bits needed) is achieved by transmitting the low-valued coefficients with shorter wordlengths. The zero-valued coefficients need not be transmitted at all. Thus it is not the DCT which compresses the data, it is the subsequent processing. The DCT simply expresses the data in a form which makes the subsequent processing easier.

Higher compression factors require the coefficient wordlength to be further reduced using requantizing. Coefficients are divided by some factor which increases the size of the quantizing step. The smaller number of steps which results permits coding with fewer bits, but of course with an increased quantizing error. The coefficients will be multiplied by a reciprocal factor in the decoder to return to the correct magnitude.

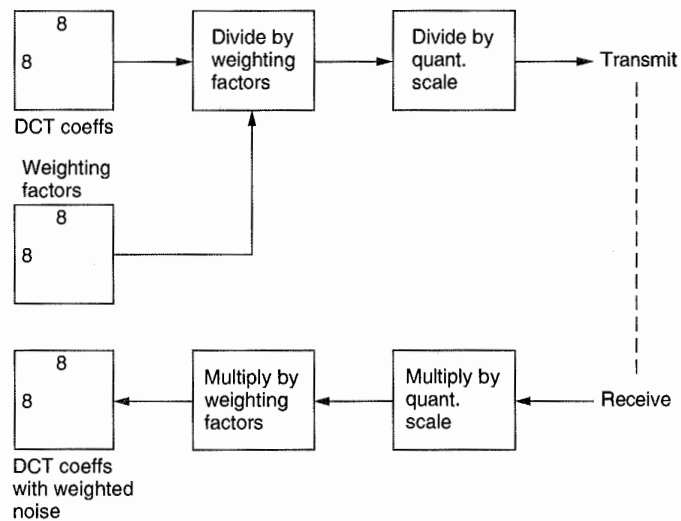


Inverse transforming a requantized coefficient means that the frequency it represents is reproduced in the output with the wrong amplitude. The difference between original and reconstructed amplitude is regarded as a noise added to the wanted data. Figure 5.10 showed that the visibility of such noise is far from uniform. The maximum sensitivity is found at DC and falls thereafter. As a result the top-left coefficient is often treated as a special case and left unchanged. It may warrant more error protection than other coefficients.

Transform coding takes advantage of the falling sensitivity to noise. Prior to requantizing, each coefficient is divided by a different weighting constant as a function of its frequency. Figure 5.31 shows a typical weighting process. Naturally the decoder must have a corresponding inverse weighting. This weighting process has the effect of reducing the magnitude of high-frequency coefficients disproportionately. Clearly different weighting will be needed for colour difference data as colour is perceived differently.

*P* and *B* pictures are decoded by adding a prediction error image to a reference image. That reference image will contain weighted noise. One purpose of the prediction error is to cancel that noise to prevent tolerance build-up. If the prediction error were also to contain weighted noise this result would not be obtained. Consequently prediction error coefficients are flat weighted.

When forward prediction fails, such as in the case of new material introduced in a *P* picture by a pan, *P* coding would set the vectors to zero



**Figure 5.31** Weighting is used to make the noise caused by requantizing different at each frequency.

and encode the new data entirely as an unweighted prediction error. In this case it is better to encode that material as an *I* macroblock because then weighting can be used and this will require fewer bits.

Requantizing increases the step size of the coefficients, but the inverse weighting in the decoder results in step sizes which increase with frequency. The larger step size increases the quantizing noise at high frequencies where it is less visible. Effectively the noise floor is shaped to match the sensitivity of the eye. The quantizing table in use at the encoder can be transmitted to the decoder periodically in the bitstream.

## 5.12 Intra-coding in MPEG-1 and MPEG-2

Study of the signal statistics gained from extensive analysis of real material is used to measure the probability of a given coefficient having a given value. This probability turns out to be highly non-uniform suggesting the possibility of a variable-length encoding for the coefficient values. On average, the higher the spatial frequency, the lower the value of a coefficient will be. This means that the value of a coefficient tends to fall as a function of its radius from the DC coefficient. DC coefficients also have certain characteristics. As they represent the average brightness of a block, in many picture areas, adjacent blocks will have similar DC coefficient values. This can be exploited using differential coding which will be dealt with in section 5.14. However, errors in the magnitude of DC coefficients cause wide area flicker in the picture and so the coding must be accurate.

Typical material often has many coefficients which are zero valued, especially after requantizing. The distribution of these also follows a pattern. The non-zero values tend to be found in the top-left-hand corner of the DCT block, but as the radius increases, not only do the coefficient values fall, but it becomes increasingly likely that these small coefficients will be interspersed with zero-valued coefficients. As the radius increases further it is probable that a region where all coefficients are zero will be entered.

MPEG uses all these attributes of DCT coefficients when encoding a coefficient block. By sending the coefficients in an optimum order, by describing their values with Huffman coding and by using run-length encoding for the zero-valued coefficients it is possible to achieve a significant reduction in coefficient data which remains entirely lossless. Despite the complexity of this process, it does contribute to improved picture quality because for a given bit rate lossless coding of the coefficients must be better than requantizing, which is lossy. Of course, for lower bit rates both will be required.




Figure 5.32 The zig-zag scan for a progressively scanned image.

It is an advantage to scan in a sequence where the largest coefficient values are scanned first. Then the next coefficient is more likely to be zero than the previous one. With progressively scanned material, a regular zig-zag scan begins in the top-left corner and ends in the bottom-right corner as shown in Figure 5.32. Zig-zag scanning means that significant values are more likely to be transmitted first, followed by the zero values. Instead of coding these zeros, an unique 'end of block' (EOB) symbol is transmitted instead.

As the zig-zag scan approaches the last finite coefficient it is increasingly likely that some zero value coefficients will be scanned. Instead of transmitting the coefficients as zeros, the *zero-run-length*, i.e. the number of zero-valued coefficients in the scan sequence, is encoded into the next non-zero coefficient which is itself variable-length coded. This combination of run-length and variable-length coding is known as RLC/VLC in MPEG.

The DC coefficient is handled separately because it is differentially coded and this discussion relates to the AC coefficients. Three items need to be handled for each coefficient: the zero-run-length prior to this coefficient, the wordlength and the coefficient value itself. The wordlength needs to be known by the decoder so that it can correctly parse the bitstream. The wordlength of the coefficient is expressed directly as an integer called the *size*.

Figure 5.33(a) shows that a two-dimensional run/size table is created. One dimension expresses the zero-run-length; the other the size. A run length of zero is obtained when adjacent coefficients are non-zero, but a code of 0/0 has no meaningful run/size interpretation and so this bit pattern is used for the end-of-block (EOB) symbol.

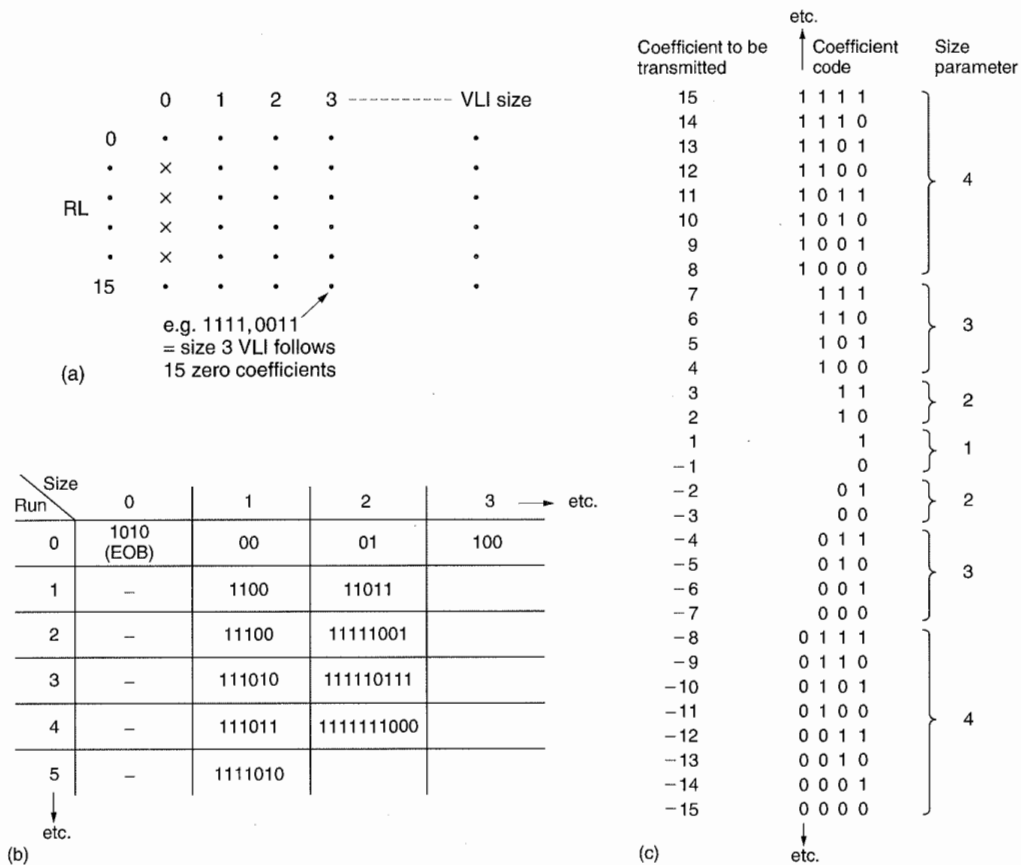


Figure 5.33 Run-length and variable-length coding simultaneously compresses runs of zero-valued coefficients and describes the wordlength of a non-zero coefficient.

In the case where the zero-run-length exceeds 14, a code of 15/0 is used signifying that there are fifteen zero-valued coefficients. This is then followed by another run/size parameter whose run-length value is added to the previous fifteen.

The run/size parameters contain redundancy because some combinations are more common than others. Figure 5.33(b) shows that each run/size value is converted to a variable-length Huffman codeword for transmission. As was shown in section 1.5, the Huffman codes are designed so that short codes are never a prefix of long codes so that the decoder can deduce the parsing by testing an increasing number of bits until a match with the look-up table is found. Having parsed and decoded the Huffman run/size code, the decoder then knows what the coefficient wordlength will be and can correctly parse that.

The variable-length coefficient code has to describe a bipolar coefficient, i.e. one which can be positive or negative. Figure 5.33(c) shows that for a particular size, the coding scale has a certain gap in it. For example, all values from  $-7$  to  $+7$  can be sent by a size 3 code, so a size 4 code only has to send the values of  $-15$  to  $-8$  and  $+8$  to  $+15$ . The coefficient code is sent as a pure binary number whose value ranges from all zeros to all ones where the maximum value is a function of the size. The number range is divided into two, the lower half of the codes specifying negative values and the upper half specifying positive.

In the case of positive numbers, the transmitted binary value is the actual coefficient value, whereas in the case of negative numbers a constant must be subtracted which is a function of the size. In the case of a size 4 code, the constant is  $15_{10}$ . Thus a size 4 parameter of  $0111_2 (7_{10})$  would be interpreted as  $7 - 15 = -8$ . A size of 5 has a constant of 31 so a transmitted coded of  $01010_2 (10_2)$  would be interpreted as  $10 - 31 = -21$ .

This technique saves a bit because, for example, 63 values from  $-31$  to  $+31$  are coded with only five bits having only 32 combinations. This is possible because that extra bit is effectively encoded into the run/size parameter.

Figure 5.34 shows the entire MPEG-1 and MPEG-2 spatial coding subsystem. Macroblocks are subdivided into DCT blocks and the DCT is calculated. The resulting coefficients are multiplied by the weighting matrix and then requantized. The coefficients are then reordered by the zig-zag scan so that full advantage can be taken of run-length and variable-length coding. The last non-zero coefficient in the scan is followed by the EOB symbol.

In predictive coding, sometimes the motion-compensated prediction is nearly exact and so the prediction error will be almost zero. This can also happen on still parts of the scene. MPEG takes advantage of this by sending a code to tell the decoder there is no prediction error data for the macroblock concerned.

The success of temporal coding depends on the accuracy of the vectors. Trying to reduce the bit rate by reducing the accuracy of the vectors is false economy as this simply increases the prediction error. Consequently for a given GOP structure it is only in the the spatial coding that the




Figure 5.34 A complete spatial coding system which can compress an *I* picture or the prediction error in *P* and *B* pictures. See text for details.

olar coeffi-  
shows that  
or example,  
4 code only  
ficient code  
zeros to all  
he number  
ng negative


is the actual  
stant must  
4 code, the  
interpreted  
d coded of

rom - 31 to  
ons. This is  
ie run/size

tial coding  
the DCT is  
weighting  
red by the  
length and  
he scan is

rediction is  
is can also  
of this by  
data for the

he vectors.  
e vectors is  
nsequently  
ig that the



ture or the

overall bit rate is determined. The RLC/VLC coding is lossless and so its contribution to the compression cannot be varied. If the bit rate is too high, the only option is to increase the size of the coefficient-requantizing steps. This has the effect of shortening the wordlength of large coefficients, and rounding small coefficients to zero, so that the bit rate goes down. Clearly if taken too far the picture quality will also suffer because at some point the noise floor will become visible as some form of artifact.

### 5.13 A bidirectional coder

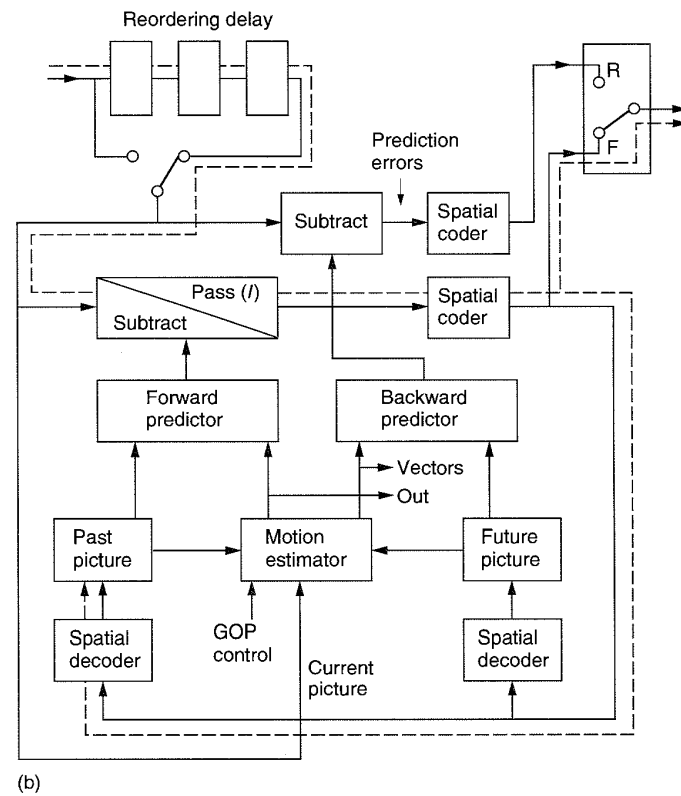
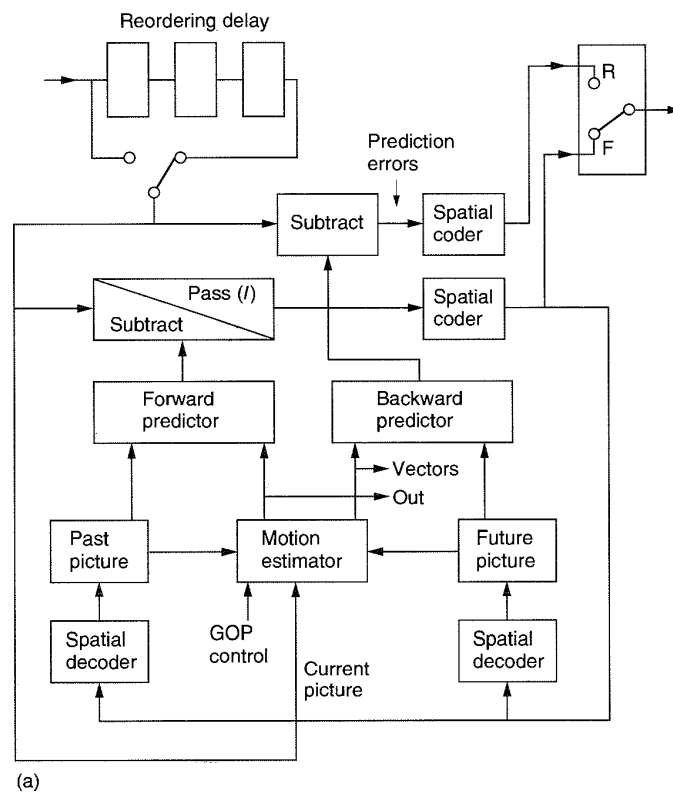
MPEG does not specify how an encoder is to be built or what coding decisions it should make. Instead it specifies the protocol of the bitstream at the output. As a result the coder shown in Figure 5.35 is only an example. Conceptually MPEG-2 Main Profile coding with a progressive scan input is almost identical to that of MPEG-1 as the differences are to be found primarily in the picture size and the resulting bit rate. The coder shown here will create MPEG-1 and MPEG-2 bidirectionally coded bitstreams from progressively scanned inputs. It also forms the basis of interlaced MPEG-2 coding and for MPEG-4 texture coding, both of which will be considered later.

Figure 5.35(a) shows the component parts of the coder. At the input is a chain of picture stores which can be bypassed for reordering purposes. This allows a picture to be encoded ahead of its normal timing when bidirectional coding is employed.

At the centre is a dual-motion estimator which can simultaneously measure motion between the input picture, an earlier picture and a later picture. These reference pictures are held in frame stores. The vectors from the motion estimator are used locally to shift a picture in a frame store to form a predicted picture. This is subtracted from the input picture to produce a prediction error picture which is then spatially coded.

The bidirectional encoding process will now be described. A GOP begins with an *I* picture which is intra-coded. In Figure 5.35(b) the *I* picture emerges from the reordering delay. No prediction is possible on an *I* picture so the motion estimator is inactive. There is no predicted picture and so the prediction error subtractor is set simply to pass the input. The only processing which is active is the forward spatial coder which describes the picture with DCT coefficients. The output of the forward spatial coder is locally decoded and stored in the past picture frame store.

The reason for the spatial encode/decode is that the past picture frame store now contains exactly what the decoder frame store will contain, including the effects of any requantizing errors. When the same picture is



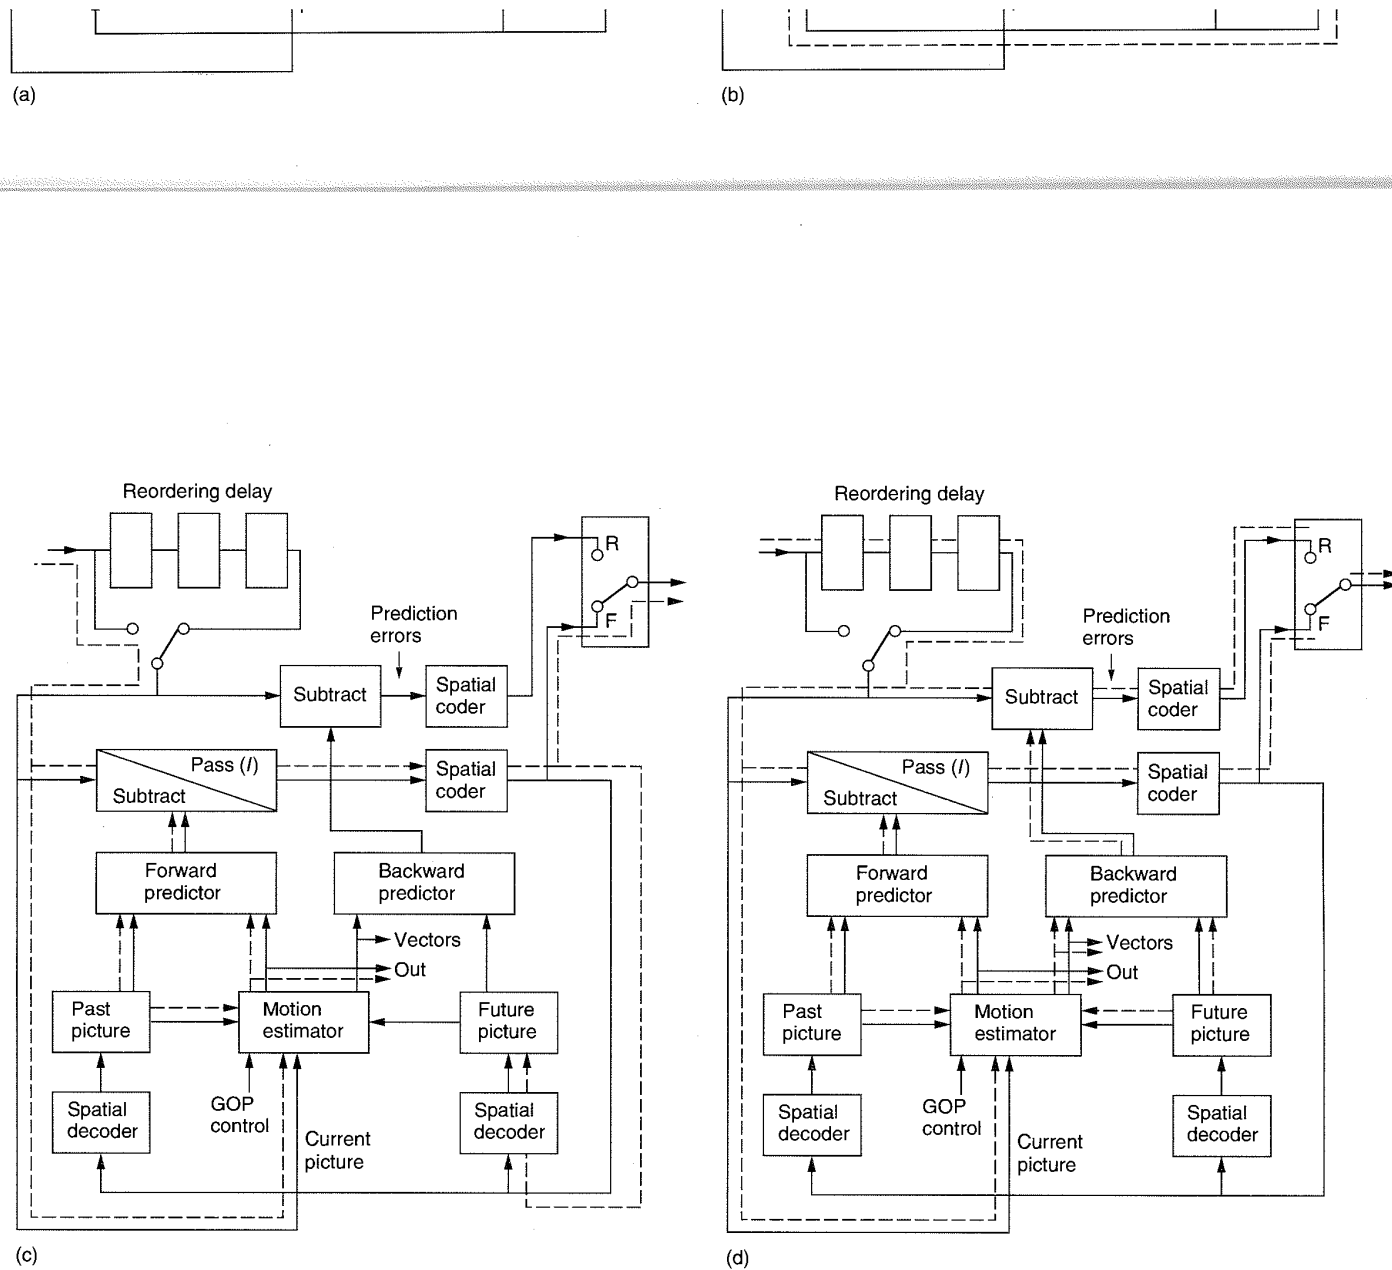


Figure 5.35 A bidirectional coder. (a) The essential components. (b) Signal flow when coding an I picture. (c) Signal flow when coding a P picture. (d) Signal flow when bidirectional coding.



used as a reference at both ends of a differential coding system, the errors will cancel out.

Having encoded the *I* picture, attention turns to the *P* picture. The input sequence is *IBBP*, but the transmitted sequence must be *IPBB*. Figure 5.35(c) shows that the reordering delay is bypassed to select the *P* picture. This passes to the motion estimator which compares it with the *I* picture and outputs a vector for each macroblock. The forward predictor uses these vectors to shift the *I* picture so that it more closely resembles the *P* picture. The predicted picture is then subtracted from the actual picture to produce a forward prediction error. This is then spatially coded. Thus the *P* picture is transmitted as a set of vectors and a prediction error image.

The *P* picture is locally decoded in the right-hand decoder. This takes the forward-predicted picture and adds the decoded prediction error to obtain exactly what the decoder will obtain.

Figure 5.35(d) shows that the encoder now contains an *I* picture in the left store and a *P* picture in the right store. The reordering delay is reselected so that the first *B* picture can be input. This passes to the motion estimator where it is compared with both the *I* and *P* pictures to produce forward and backward vectors. The forward vectors go to the forward predictor to make a *B* prediction from the *I* picture. The backward vectors go to the backward predictor to make a *B* prediction from the *P* picture. These predictions are simultaneously subtracted from the actual *B* picture to produce a forward prediction error and a backward prediction error. These are then spatially encoded. The encoder can then decide which direction of coding resulted in the best prediction; i.e. the smallest prediction error.

In the encoder shown, motion estimation sometimes takes place between an input picture and a decoded picture. With increased complexity, the original input pictures may be stored for motion-estimation purposes. As these contain no compression artifacts the motion vectors may be more accurate.

Not shown in the interests of clarity is a third signal path which creates a predicted *B* picture from the average of forward and backward predictions. This is subtracted from the input picture to produce a third prediction error. In some circumstances this prediction error may use fewer data than either forward or backward prediction alone.

As *B* pictures are never used to create other pictures, the decoder does not locally decode the *B* picture. After decoding and displaying the *B* picture the decoder will discard it. At the encoder the *I* and *P* pictures remain in their frame stores and the second *B* picture is input from the reordering delay.

Following the encoding of the second *B* picture, the encoder must reorder again to encode the second *P* picture in the GOP. This will be

locally decoded and will replace the *I* picture in the left store. The stores and predictors switch designation because the left store is now a future *P* picture and the right store is now a past *P* picture. *B* pictures between them are encoded as before.

## 5.14 Slices

In MPEG-1 and MPEG-2 *I* pictures and in MPEG-4 *I-VOPs*, the DC coefficient describes the average brightness of an entire DCT block. In real video the DC component of adjacent blocks will be similar much of the time. A saving in bit rate can be obtained by differentially coding the DC coefficient.

In *P* and *B* pictures this is not done because these are prediction errors, not actual images, and the statistics are different. However, *P* and *B* pictures send vectors and instead the redundancy in these is explored. In a large moving object, many macroblocks will be moving at the same velocity and their vectors will be the same. Thus differential vector coding will be advantageous.

As has been seen above, differential coding cannot be used indiscriminately as it is prone to error propagation. Periodically absolute DC coefficients and vectors must be sent and the *slice* is the logical structure which supports this mechanism. In *I* coding, the first DC coefficient in a slice is sent in absolute form, whereas the subsequent coefficients are sent differentially. In *P* or *B* coding, the first vector in a slice is sent in absolute form, but the subsequent vectors are differential.

Slices are horizontal picture strips which are one macroblock (16 pixels) high and which proceed from left to right across the screen and may run on to the next horizontal row. The encoder is free to decide how big slices should be and where they begin and end.

In the case of a central dark building silhouetted against the bright sky, there would be two large changes in the DC coefficients, one at each edge of the building. It may be advantageous to the encoder to break the width of the picture into three slices, one each for the left and right areas of sky and one for the building. In the case of a large moving object, different slices may be used for the object and the background.

Each slice contains its own synchronizing pattern, so following a transmission error, correct decoding can resume at the next slice. Slice size can also be matched to the characteristics of the transmission channel. For example, in an error-free transmission system the use of a large number of slices in a packet simply wastes data capacity on surplus synchronizing patterns. However, in a non-ideal system it might be advantageous to have frequent resynchronizing. In DVB, for example, additional constraints are placed on the slice size.


### 5.15 Handling interlaced pictures

MPEG-1 does not support interlace, whereas MPEG-2 and MPEG-4 do. Spatial coding, predictive coding and motion compensation can still be performed using interlaced source material at the cost of considerable complexity. Despite that complexity, MPEG cannot be expected to perform as well with interlaced material.

Figure 5.36 shows that in an incoming interlaced frame there are two fields each of which contain half of the lines in the frame. In MPEG-2 and MPEG-4 these are known as the *top field* and the *bottom field*. In video from a camera, these fields represent the state of the image at two different times. Where there is little image motion, this is unimportant and the fields can be combined, obtaining more effective compression. However, in the presence of motion the fields become increasingly decorrelated because of the displacement of moving objects from one field to the next.

This characteristic determines that MPEG must be able to handle fields independently or together. This dual approach permeates all aspects of MPEG and affects the definition of pictures, macroblocks, DCT blocks and zig-zag scanning.

Figure 5.36 also shows how MPEG-2 designates interlaced fields. In picture types *I*, *P* and *B*, the two fields can be superimposed to make a



**Figure 5.36** An interlaced frame consists of top and bottom fields. MPEG-2 can code a frame in the ways shown here.

*frame-picture* or the two fields can be coded independently as two *field-pictures*. As a third possibility, in *I* pictures only, the bottom field-picture can be predictively coded from the top field-picture to make an *IP* frame-picture.

A frame-picture is one in which the macroblocks contain lines from both field types over a picture area 16 scan lines high. Each luminance macroblock contains the usual four DCT blocks but there are two ways in which these can be assembled. Figure 5.37(a) shows how a frame is divided into *frame DCT* blocks. This is identical to the progressive scan approach in that each DCT block contains eight contiguous picture lines. In 4:2:0, the colour difference signals have been downsampled by a factor of two and shifted as was shown in Figure 2.17. Figure 5.37(a) also shows how one 4:2:0 DCT block contains the chroma data from 16 lines in two fields.


Even small amounts of motion in any direction can destroy the correlation between odd and even lines and a frame DCT will result in an excessive number of coefficients. Figure 5.37(b) shows that instead the luminance component of a frame can also be divided into *field DCT* blocks. In this case one DCT block contains odd lines and the other contains even lines. In this mode the chroma still produces one DCT block from both fields as in Figure 5.37(a).

When an input frame is designated as two *field-pictures*, the macroblocks come from a screen area which is 32 lines high. Figure 5.37(c) shows that the DCT blocks contain the same data as if the input frame had been designated a *frame-picture* but with *field DCT*. Consequently it is only frame-pictures which have the option of field or frame DCT. These may be selected by the encoder on a macroblock-by-macroblock basis and, of course, the resultant bitstream must specify what has been done.


In a frame which contains a small moving area, it may be advantageous to encode as a frame-picture with frame DCT except in the moving area where field DCT is used. This approach may result in fewer bits than coding as two field-pictures.

In a field-picture and in a frame-picture using field DCT, a DCT block contains lines from one field type only and this must have come from a screen area 16 scan lines high, whereas in progressive scan and frame DCT the area is only 8 scan lines high. A given vertical spatial frequency in the image is sampled at points twice as far apart which is interpreted by the field DCT as a doubled spatial frequency, whereas there is no change in the horizontal spectrum.

Following the DCT calculation, the coefficient distribution will be different in field-pictures and field DCT frame-pictures. In these cases, the probability of coefficients is not a constant function of radius from the DC coefficient as it is in progressive scan, but is elliptical where the ellipse is twice as high as it is wide.



**Figure 5.37** (a) In Frame-DCT, a picture is effectively de-interlaced. (b) In Field-DCT, each DCT block only contains lines from one field, but over twice the screen area. (c) The same DCT content results when field-pictures are assembled into blocks.



**Figure 5.38** The zig-zag scan for an interlaced image has to favour vertical frequencies twice as much as horizontal.

Using the standard  $45^\circ$  zig-zag scan with this different coefficient distribution would not have the required effect of putting all the significant coefficients at the beginning of the scan. To achieve this requires a different zig-zag scan, which is shown in Figure 5.38. This scan, sometimes known as the Yeltsin walk, attempts to match the elliptical probability of interlaced coefficients with a scan slanted at  $67.5^\circ$  to the vertical.

Motion estimation is more difficult in an interlaced system. Vertical detail can result in differences between fields and this reduces the quality of the match. Fields are vertically subsampled without filtering and so contain alias products. This aliasing will mean that the vertical waveform representing a moving object will not be the same in successive pictures and this will also reduce the quality of the match.

Even when the correct vector has been found, the match may be poor so the estimator fails to recognize it. If it is recognized, a poor match means that the quality of the prediction in *P* and *B* pictures will be poor and so a large prediction error or residual has to be transmitted. In an attempt to reduce the residual, MPEG-2 allows field-pictures to use motion-compensated prediction from either the adjacent field or from the same field type in another frame. In this case the encoder will use the better match. This technique can also be used in areas of frame-pictures which use field DCT.


The motion compensation of MPEG-2 has half-pixel resolution and this is inherently compatible with an interlace because an interpolator must be present to handle the half-pixel shifts. Figure 5.39(a) shows that in an interlaced system, each field contains half of the frame lines and so interpolating half-way between lines of one field type will actually create values lying on the sampling structure of the other field type. Thus it is

8  
screen  
lines

16  
screen  
lines

16  
screen  
lines

field-DCT,  
1 area. (c)  
i.



**Figure 5.39** (a) Each field contains half of the frame lines and so interpolation is needed to create values lying on the sampling structure of the other field type. (b) Prediction can use data from the previous field or the one before that.

equally possible for a predictive system to decode a given field type based on pixel data from the other field type or of the same type.

If when using predictive coding from the other field type the vertical motion vector contains a half-pixel component, then no interpolation is needed because the act of transferring pixels from one field to another results in such a shift.

Figure 5.39(b) shows that a macroblock in a given *P* field-picture can be encoded using a vector which shifts data from the previous field or from the field before that, irrespective of which frames these fields occupy. As noted above, field-picture macroblocks come from an area of screen 32 lines high and this means that the vector density is halved, resulting in larger prediction errors at the boundaries of moving objects.

As an option, field-pictures can restore the vector density by using  $16 \times 8$  motion compensation where separate vectors are used for the top and bottom halves of the macroblock. Frame-pictures can also use  $16 \times 8$  motion compensation in conjunction with field DCT. Whilst the  $2 \times 2$  DCT block luminance structure of a macroblock can easily be divided vertically in two, in 4:2:0 the same screen area is represented by only one

chroma macroblock of each component type. As it cannot be divided in half, this chroma is deemed to belong to the luminance DCT blocks of the upper field. In 4:2:2 no such difficulty arises.

MPEG supports interlace simply because interlaced video exists in legacy systems and there is a requirement to compress it. However, where the opportunity arises to define a new system, interlace should be avoided. Legacy interlaced source material should be handled using a motion-compensated de-interlacer prior to compression in the progressive domain.

## 5.16 MPEG-1 and MPEG-2 coders

Figure 5.40 shows a complete coder. The bidirectional coder outputs coefficients and vectors, and the quantizing table in use. The vectors of *P* and *B* pictures and the DC coefficients of *I* pictures are differentially encoded in slices and the remaining coefficients are RLC/VLC coded. The multiplexer assembles all these data into a single bitstream called an elementary stream. The output of the encoder is a buffer which absorbs the variations in bit rate between different picture types. The buffer output has a constant bit rate determined by the demand clock. This comes from the transmission channel or storage device. If the bit rate is low, the buffer will tend to fill up, whereas if it is high the buffer will tend to empty. The buffer content is used to control the severity of the requantizing in the spatial coders. The more the buffer fills, the bigger the requantizing steps get.

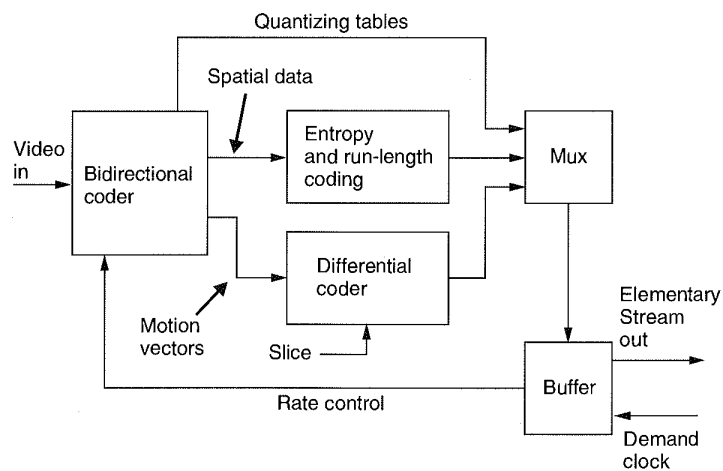


Figure 5.40 An MPEG-2 coder. See text for details.



The buffer in the decoder has a finite capacity and the encoder must model the decoder's buffer occupancy so that it neither overflows nor underflows. An overflow might occur if an *I* picture is transmitted when the buffer content is already high. The buffer occupancy of the decoder depends somewhat on the memory access strategy of the decoder. Instead of defining a specific buffer size, MPEG-2 defines the size of a particular mathematical model of a hypothetical buffer. The decoder designer can use any strategy which implements the model, and the encoder can use any strategy which doesn't overflow or underflow the model. The elementary stream has a parameter called the video buffer verifier (VBV) which defines the minimum buffering assumptions of the encoder.

As was seen in Chapter 1, buffering is one way of ensuring constant quality when picture entropy varies. An intelligent coder may run down the buffer contents in anticipation of a difficult picture sequence so that a large amount of data can be sent.

MPEG-2 does not define what a decoder should do if a buffer underflow or overflow occurs, but since both irrecoverably lose data it is obvious that there will be more or less of an interruption to the decoding. Even a small loss of data may cause loss of synchronization and in the case of long GOP the lost data may make the rest of the GOP undecodable. A decoder may choose to repeat the last properly decoded picture until it can begin to operate correctly again.

Buffer problems occur if the VBV model is violated. If this happens then more than one underflow or overflow can result from a single violation. Switching an MPEG bitstream can cause a violation because the two encoders concerned may have radically different buffer occupancy at the switch.

## 5.17 The elementary stream

Figure 5.41 shows the structure of the elementary stream from an MPEG-2 encoder. The structure begins with a set of coefficients representing a DCT block. Six or eight DCT blocks form the luminance and chroma content of one macroblock. In *P* and *B* pictures a macroblock will be associated with a vector for motion compensation. Macroblocks are associated into slices in which DC coefficients of *I* pictures and vectors in *P* and *B* pictures are differentially coded. An arbitrary number of slices forms a picture and this needs *I/P/B* flags describing the type of picture it is. The picture may also have a global vector which efficiently deals with pans.

Several pictures form a group of pictures (GOP). The GOP begins with an *I* picture and may or may not include *P* and *B* pictures in a structure which may vary dynamically.

ncoder must  
verflows nor  
mitted when  
the decoder  
oder. Instead  
f a particular  
designer can  
oder can use  
model. The  
erifier (VBV)  
ncoder.  
ring constant  
ay run down  
nce so that a

if a buffer  
lose data it is  
he decoding.  
n and in the  
of the GOP  
erly decoded

this happens  
om a single  
n because the  
occupancy at

m an MPEG-  
representing a  
and chroma  
block will be  
robblocks are  
nd vectors in  
ber of slices  
e of picture it  
ly deals with

' begins with  
n a structure

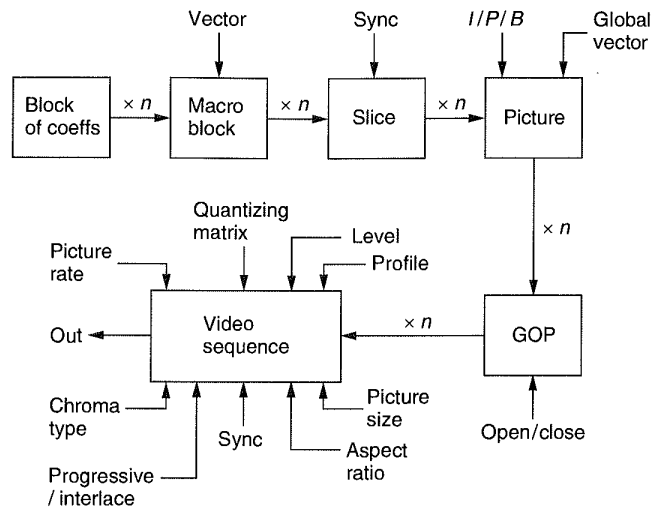


Figure 5.41 The structure of an elementary stream. MPEG defines the syntax precisely.

Several GOPs form a sequence which begins with a sequence header containing important data to help the decoder. It is possible to repeat the header within a sequence, and this helps lock-up in random access applications. The sequence header describes the MPEG-2 profile and level, whether the video is progressive or interlaced, whether the chroma is 4:2:0 or 4:2:2, the size of the picture and the aspect ratio of the pixels. The quantizing matrix used in the spatial coder can also be sent. The sequence begins with a standardized bit pattern which is detected by a decoder to synchronize the deserialization.

## 5.18 An MPEG-2 decoder

The decoder is only defined by implication from the definitions of syntax and any decoder which can correctly interpret all combinations of syntax at a particular profile will be deemed compliant however it works. Compliant MPEG-2 decoders must be able to decode an MPEG-1 elementary stream. This is not particularly difficult as MPEG-2 uses all the tools of MPEG-1 and simply adds more of its own. Consequently an MPEG-1 bitstream can be thought of as a subset of an MPEG-2 bitstream.

The first problem a decoder has is that the input is an endless bitstream which contains a huge range of parameters many of which have variable length. Unique synchronizing patterns must be placed periodically throughout the bitstream so that the decoder can identify known starting points. The pictures which can be sent under MPEG are so flexible that

the decoder must first find a sequence header so that it can establish the size of the picture, the frame rate, the colour coding used, etc.

The decoder must also be supplied with a 27 MHz system clock. In a DVD player, this would come from a crystal, but in a transmission system this would be provided by a numerically locked loop running from a clock reference parameter in the bitstream (see Chapter 6). Until this loop has achieved lock the decoder cannot function properly.

Figure 5.42 shows a bidirectional decoder. The decoder can only begin decoding with an *I* picture and as this only uses intra-coding there will be no vectors. An *I* picture is transmitted as a series of slices which begin with subsidiary synchronizing patterns. The first macroblock in the slice contains an absolute DC coefficient, but the remaining macroblocks code the DC coefficient differentially so the decoder must subtract the differential values from the previous value to obtain the absolute value.

The AC coefficients are sent as Huffman coded run/size parameters followed by coefficient value codes. The variable-length Huffman codes are decoded by using a look-up table and extending the number of bits considered until a match is obtained. This allows the zero-run-length and the coefficient size to be established. The right number of bits is taken from the bitstream corresponding to the coefficient code and this is decoded to the actual coefficient using the size parameter.

If the correct number of bits has been taken from the stream, the next bit must be the beginning of the next run/size code and so on until the EOB symbol is reached. The decoder uses the coefficient values and the zero-run-lengths to populate a DCT coefficient block following the appropriate zig-zag scanning sequence. Following EOB, the bitstream then continues with the next DCT block. Clearly this Huffman decoding will work perfectly or not at all. A single bit slippage in synchronism or a single corrupted data bit can cause a spectacular failure.

Once a complete DCT coefficient block has been received, the coefficients need to be inverse quantized and inverse weighted. Then an inverse DCT can be performed and this will result in an  $8 \times 8$  pixel block. A series of DCT blocks will allow the luminance and colour information for an entire macroblock to be decoded and this can be placed in a frame store. Decoding continues in this way until the end of the slice when an absolute DC coefficient will once again be sent. Once all the slices have been decoded, an entire picture will be resident in the frame store.

The amount of data needed to decode the picture is variable and the decoder just keeps going until the last macroblock is found. It will obtain data from the input buffer. In a constant bit rate transmission system, the decoder will remove more data to decode an *I* picture than has been received in one picture period, leaving the buffer emptier than it began. Subsequent *P* and *B* pictures need much fewer data and allow the buffer to fill again.

establish the  
clock. In a  
ion system  
ng from a  
il this loop

only begin  
ere will be  
hich begin  
in the slice  
locks code  
btract the  
lute value.  
parameters  
man codes  
ber of bits  
length and  
ts is taken  
nd this is

n, the next  
n until the  
es and the  
owing the  
bitstream  
decoding  
ironism or

eived, the  
d. Then an  
ixel block.  
formation  
in a frame  
e when an  
lices have  
store.  
le and the  
will obtain  
ystem, the  
has been  
n it began.  
the buffer




Figure 5.42 A bidirectional MPEG-2 decoder. See text for details.

The picture will be output when the time stamp (see Chapter 6) sent with the picture matches the state of the decoder's time count.

Following the *I* picture may be another *I* picture or a *P* picture. Assuming a *P* picture, this will be predictively coded from the *I* picture. The *P* picture will be divided into slices as before. The first vector in a slice is absolute, but subsequent vectors are sent differentially. However, the DC coefficients are not differential.

Each macroblock may contain a forward vector. The decoder uses this to shift pixels from the *I* picture into the correct position for the predicted *P* picture. The vectors have half-pixel resolution and where a half-pixel shift is required, an interpolator will be used.

The DCT data is sent much as for an *I* picture. It will require inverse quantizing, but not inverse weighting because *P* and *B* coefficients are flat-weighted. When decoded this represents an error-cancelling picture which is added pixel-by-pixel to the motion-predicted picture. This results in the output picture.

If bidirectional coding is being used, the *P* picture may be stored until one or more *B* pictures have been decoded. The *B* pictures are sent essentially as a *P* picture might be, except that the vectors can be forward, backward or bidirectional. The decoder must take pixels from the *I* picture, the *P* picture, or both, and shift them according to the vectors to make a predicted picture. The DCT data decodes to produce an error-cancelling image as before.

In an interlaced system, the prediction mechanism may alternatively obtain pixel data from the previous field or the field before that. Vectors may relate to macroblocks or to  $16 \times 8$  pixel areas. DCT blocks after decoding may represent frame lines or field lines. This adds up to a lot of different possibilities for a decoder handling an interlaced input.

## 5.19 MPEG-4

As was seen in Chapter 1, MPEG-4 advances the coding art in a number of ways. Whereas MPEG-1 and MPEG-2 were directed only to coding the video pictures which resulted after shooting natural scenes or from computer synthesis, MPEG-4 also moves further back in the process of how those scenes were created. For example, the rotation of a detailed three-dimensional object before a video camera produces huge changes in the video from picture to picture which MPEG-2 would find difficult to code. Instead, if the three-dimensional object is re-created at the decoder, rotation can be portrayed by transmitting a trivially small amount of vector data.

If the above object is synthetic, effectively the synthesis or rendering process is completed in the decoder. However, a suitable if complex

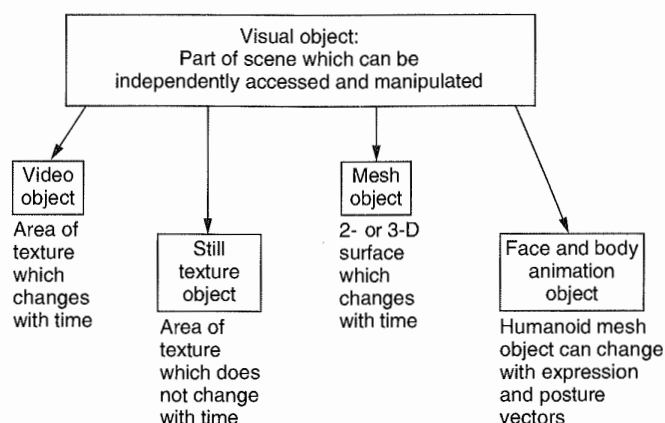


Figure 5.43 In MPEG-4 four types of *objects* are coded.

image processor at the encoder could identify such objects in natural scenes. MPEG-4 objects are defined as a part of a scene which can independently be accessed or manipulated. An object is an entity that exists over a certain time span. The pictures of conventional imaging become *object planes* in MPEG-4. Where an object intersects an object plane, it can be described by the coding system using intra-coding, forward prediction or bidirectional prediction.

Figure 5.43 shows that MPEG-4 has four object types. A video object is an arbitrarily shaped planar pixel array describing the appearance or *texture* of part of a scene. A still texture object or *sprite* is a planar video object in which there is no change with respect to time. A mesh object describes a two- or three-dimensional shape as a set of points. The shape and its position can change with respect to time. Using computer graphics techniques, texture can be mapped onto meshes, a process known as warping, to produce rendered images.

Using two-dimensional warping, a still texture object can be made to move. In three-dimensional graphic rendering, mesh coding allows an arbitrary solid shape to be created which is then covered with texture. Perspective computation then allows this three-dimensional object to be viewed in correct perspective from any viewpoint. MPEG-4 provides tools to allow two- or three-dimensional meshes to be created in the decoder and then oriented by vectors. Changing the vectors then allows realistic moving images to be created with an extremely low bit rate.


Face and body animation is a specialized subset of three-dimensional mesh coding in which the mesh represents a human face and/or body. As the subject moves, carefully defined vectors carry changes of expression which allow rendering of an apparently moving face and/or body which has been almost entirely synthesized from a single still picture.

In addition to object coding, MPEG-4 refines the existing MPEG tools by increasing the efficiency of a number of processes using lossless prediction. This improves the performance of the motion compensation and coefficient coding allowing either a lower bit rate or improved quality. MPEG-4 also extends the idea of scalability introduced in MPEG-2. Multiple scalability is supported, where a low-bit-rate base-level picture may optionally be enhanced by adding information from one or more additional bitstreams. This approach is useful in network applications where the content creator cannot know the bandwidth which a particular user will have available. Scalability allows the best quality in the available bandwidth.

Although most of the spatial compression of MPEG-4 is based on the DCT as in earlier MPEG standards, MPEG-4 also introduces wavelet coding of still objects. Wavelets are advantageous in scaleable systems because they naturally decompose the original image into various resolutions.

## 5.20 Video objects

Figure 5.44 shows an example of a video object intersecting *video object planes*, or VOPs. At each plane, the shape and the texture of the object must be portrayed. Figure 5.45 shows this can be done using appropriate



**Figure 5.44** Effectively objects exist through continuous time, but are sampled for display purposes. A video object plane (VOP) is a temporal sample of an object. VOPs correspond to pictures of MPEG-1 and MPEG-2.

3 MPEG tools  
using lossless  
compensation  
or improved  
ntroduced in  
-bit-rate base-  
rmation from  
al in network  
dwidth which  
best quality in

based on the  
lucuses wavelet  
eable systems  
into various

ng video object  
: of the object  
g appropriate

as  
e  
t

ampled for  
a object. VOPs




Figure 5.45 VOPs can be encoded using the same intra- and inter-coding approach of MPEG-1 and MPEG-2.





Figure 5.46 An MPEG-4 object-based decoder is essentially a multi-layered compositing device, using keying to give the illusion that objects are in front of a background.



combinations of intra- and inter-coding as described for the earlier standards. This gives rise to I-VOPs, P-VOPs and B-VOPs. A group of VOPs is known as a GOV.

Figure 5.46 shows how video objects are handled at the decoder which effectively contains a multi-layer compositing stage. The shape of each object is transmitted using *alpha data* which is decoded to produce a key signal. This has the effect of keying the texture data into the overall picture. Several objects can be keyed into the same picture. The background may be yet another video object or a still texture object which can be shifted by vectors in the case of pans etc.

Figure 5.47 shows that MPEG-4 codes each arbitrarily shaped video object by creating a *bounding rectangle* within which the object resides. The shape and the texture of each video object are re-created at the decoder by two distinct but interrelated processes. The bounding rectangle exists on macroblock boundaries and can change from VOP to VOP. At each VOP,



**Figure 5.47** The alpha plane is used in MPEG-4 to describe the shape of an object. Within a bounding rectangle are three different types of macroblock: those within the object, those outside and those which are intersected by the boundary.

the bitstream contains horizontal and vertical references specifying the position of the top-left corner of the bounding rectangle and its width and height.

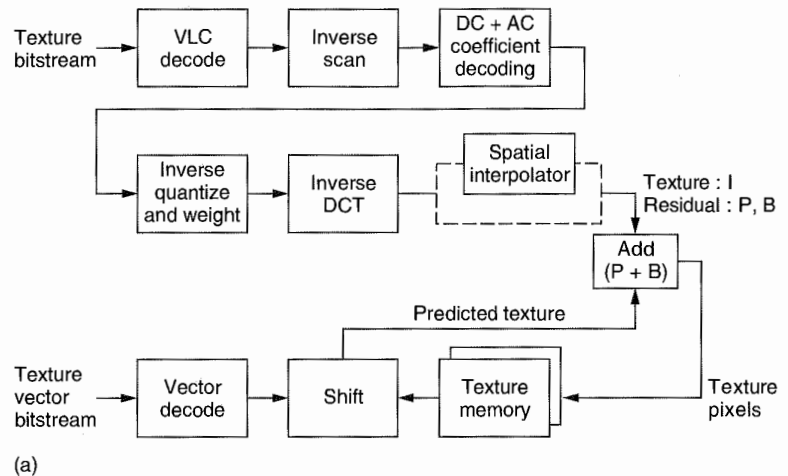
Within the bounding rectangle are three kinds of macroblock. *Transparent macroblocks* are entirely outside the object and contain no texture data. The shape data are identical over the entire block and result in a key signal which deselects this block in the compositing process. Such shape data are trivial to compress. *Opaque macroblocks* are entirely within the object and are full of texture data. The shape data are again identical over the entire block and result in the block being keyed at the compositor. *Boundary macroblocks* are blocks through which the edge of the object passes. They contain essentially all the shape information and somewhat fewer texture data than an opaque macroblock.

## 5.21 Texture coding

In MPEG-1 and MPEG-2 the only way of representing an image is with pixels and this requires no name. In MPEG-4 there are various types of image description tools and it becomes necessary to give the pixel representation of the earlier standards a name. This is *texture coding* which is that part of MPEG-4 that operates on pixel-based areas of image. Coming later than MPEG-1 and MPEG-2, the MPEG-4 texture-coding system can afford additional complexity in the search for higher performance. Figure 5.48(a) shows the texture decoding system of MPEG-4 and in (b) is a comparison with MPEG-2 showing that it is a refinement of the earlier technique. These refinements are lossless in that the reduction in bit rate they allow does not result in a loss of quality.

When inter-coding, MPEG-4 may use motion-compensated bidirectional coding as for MPEG-2, but in doing so it also uses a more advanced form of lossless vector compression which may then be used to increase the vector density to one vector per DCT block. In inter-coded pictures the prediction of the picture is improved so that the residual to be coded is smaller.

When intra-coding, MPEG-4 looks for further redundancy between coefficients using prediction. When a given DCT block is to be intra-coded, certain of its coefficients will be predicted from adjacent blocks. The choice of the most appropriate block is made by measuring the picture gradient, defined as the rate of change of the DC coefficient. Figure 5.49(a) shows that the three adjacent blocks, A, B and C, are analysed to decide whether to predict from the DCT block above (vertical prediction) or to the left (horizontal prediction). Figure 5.49(b) shows that in vertical prediction the top row of coefficients is predicted from the block above so that only the differences between them need to be coded.



MPEG-2 Picture coding	MPEG-4 Texture coding
DCT	DCT
DC coefficient prediction in slice	DC coefficient prediction
Single coefficient scan	AC coefficient prediction
Fixed resolution weighting	$3 \times$ coefficient scans
VLC/RLC	Reduced resolution weighting
	VLC/RLC
One vector/macroblock	Result : more efficient I coding
Differential vectors in slice	$4 \times$ vectors/macroblock option
	Vector prediction
	Result: more efficient P, B coding

(b)

Figure 5.48 (a) The texture decoding system of MPEG-4. (b) A comparison of MPEG-2 and MPEG-4 shows that the latter is more refined.

Figure 5.49(c) shows that in horizontal prediction the left column of coefficients is predicted from the block on the left so that again only the differences need be coded.

Choosing the blocks above and to the left is important because these blocks will already be available in both the encoder and decoder. By making the same picture gradient measurement, the decoder can establish whether vertical or horizontal prediction has been used and so no flag is needed in the bitstream.

Some extra steps are needed to handle the top row and the left column of a picture or object where true prediction is impossible. In these cases both encoder and decoder assume standardized constant values for the missing prediction coefficients.

The picture gradient measurement determines the direction in which there is the least change from block to block. There will generally be fewer DCT coefficients present in this direction. There will be more coefficients in the other axis where there is more change. Consequently it is

ture : I  
idual : P, B

Texture  
pixels

of MPEG-2

column of  
in only the

ause these  
ecoder. By  
coder can  
sed and so

eft column  
these cases  
ues for the

n in which  
ly be fewer  
coefficients  
ently it is





Figure 5.49 (a) DC coefficients are used to measure the picture gradient. (b) In vertical prediction the top row of coefficients are predicted using those above as a basis. (c) In horizontal prediction the left column of coefficients are predicted from those to the left.

advantageous to alter the scanning sequence so that the coefficients which are likely to exist are transmitted earlier in the sequence.

Figure 5.50 shows the two alternate scans for MPEG-4. The alternate horizontal scan concentrates on horizontal coefficients early in the scan and will be used in conjunction with vertical prediction. Conversely the alternate vertical scan concentrates on vertical coefficients early in the scan and will be used in conjunction with horizontal prediction. The decoder can establish which scan has been used in the encoder from the picture gradient.

Coefficient prediction is not employed when inter-coding because the statistics of residual images are different. Instead of attempting to predict residual coefficients, in inter-coded texture, pixel-based prediction may

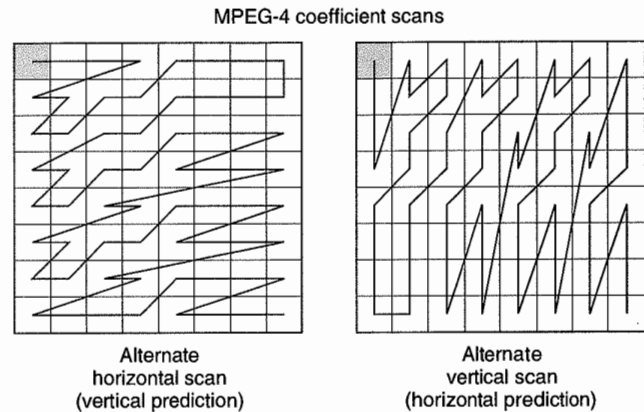


Figure 5.50 The alternate zig-zag scans employed with vertical or horizontal prediction.

be used to reduce the magnitude of texture residuals. This technique is known as *overlapped block motion compensation* (OBMC) which is only used in P-VOPs. With only one vector per DCT block, clearly in many cases the vector cannot apply to every pixel in the block. If the vector is considered to describe the motion of the centre of the block, the vector accuracy falls towards the edge of the block. A pixel in the corner of a block is almost equidistant from a vector in the centre of an adjacent block.

OBMC uses vectors from adjacent blocks, known as *remote vectors* in addition to the vector of the current block for prediction. Figure 5.51 shows that the motion-compensation process of MPEG-1 and MPEG-2 which uses a single vector is modified by the addition of the pixel prediction system which considers three vectors. A given pixel in the block to be coded is predicted from the weighted sum of three motion-compensated pixels taken from the previous I- or P-VOP. One of these pixels is obtained in the normal way by accessing the previous VOP with a shift given by the vector of this block. The other two are obtained by accessing the same VOP pixels using the remote vectors of two adjacent blocks.

The remote vectors which are used and the weighting factors are both a function of the pixel position in the block. Figure 5.51 shows that the block to be coded is divided into quadrants. The remote vectors are selected from the blocks closest to the quadrant in which the pixel resides. For example, a pixel in the bottom-right quadrant would be predicted using remote vectors from the DCT block immediately below and the block immediately to the right.

Not all blocks can be coded in this way. In P-VOPs it is permissible to have blocks which are not coded or intra-blocks which contain no vector. Remote vectors will not all be available at the boundaries of a VOP. In the




horizontal prediction.

technique is only used in many cases the accuracy falls back is almost


note vectors in Figure 5.51 and MPEG-2 of the pixel pixel in the three motion- One of these us VOP with obtained by two adjacent

vectors are both rows that the e vectors are pixel resides. be predicted below and the

permissible to ain no vector. a VOP. In the



(a)



(b)

**Figure 5.51** (a) MPEG-4 inter-coded pixel values may be predicted from three remote vectors as well as the true vector. (b) The three remote vectors are selected according to the macroblock quadrant. Macroblocks may have one or four vectors and the prediction mechanism allows prediction between blocks of either type.

normal sequence of macroblock transmission, vectors from macroblocks below the current block are not yet available. Some additional steps are needed to handle these conditions. Adjacent to boundaries where a remote vector is not available it is replaced by a copy of the actual vector. This is also done where an adjacent block is intra-coded and for blocks at the bottom of a macroblock where the vectors for the macroblocks below will not be available yet. In the case of non-coded blocks the remote vector is set to zero.

Figure 5.52(a) shows that the weighting factors for pixels near the centre of a block favour that block. In the case of pixels at the corner of the block, the weighting is even between the value obtained from the true vector and the sum of the two pixel values obtained from remote vectors.

The weighted sum produces a predicted pixel which is subtracted from the actual pixel in the current VOP to be coded to produce a residual pixel. Blocks of residual pixels are DCT coded as usual. OBMC reduces the magnitude of residual pixels and gives a corresponding reduction in the number or magnitude of DCT coefficients to be coded.

4	5	5	5	5	5	5	4
5	5	5	5	5	5	5	5
5	5	6	6	6	6	5	5
5	5	6	6	6	6	5	5
5	5	6	6	6	6	5	5
5	5	6	6	6	6	5	5
5	5	5	5	5	5	5	5
4	5	5	5	5	5	5	4

(a) Weighting values  
for use with motion  
vector of current  
Y block (true vector)

2	2	2	2	2	2	2	2
1	1	2	2	2	2	1	1
1	1	1	1	1	1	1	1
1	1	1	1	1	1	1	1
1	1	1	1	1	1	1	1
1	1	1	1	1	1	1	1
1	1	2	2	2	2	1	1
2	2	2	2	2	2	2	2

(b) Weighting values for  
use with motion vector  
of Y blocks above or  
below current Y block  
(remote vector)

2	1	1	1	1	1	1	2
2	2	1	1	1	1	2	2
2	2	1	1	1	1	2	2
2	2	1	1	1	1	2	2
2	2	1	1	1	1	2	2
2	2	1	1	1	1	2	2
2	2	1	1	1	1	2	2
2	1	1	1	1	1	1	2

(c) Weighting values  
for use with motion  
vectors of Y blocks  
left or right of current  
Y block (remote vector)

Figure 5.52 Vector prediction.

OBMC is lossless because the decoder already has access to all the vectors and knows the weighting tables. Consequently the only overhead is the transmission of a flag which enables or disables the mechanism.

MPEG-4 also has the ability to downsample prediction error or residual macroblocks which contain little detail. A  $16 \times 16$  macroblock block is downsampled to  $8 \times 8$  and flagged. The decoder will identify the flag and interpolate back to  $16 \times 16$ .

In vector prediction, each macroblock may have only one or four vectors as the coder decides. Consequently the prediction of a current vector may have to be done from either macroblock or DCT block vectors. In the case of predicting one vector for an entire macroblock, or the top-left DCT block vector, the process shown in Figure 5.51(b) is used. Three earlier vectors, which may be macroblock or DCT block vectors, as available, are used as the input to the prediction process. In the diagram the large squares show the macroblock vectors to be selected and the small squares show the DCT block vectors to be selected. The three vectors are passed to a median filter which outputs the vector in the centre of the range unchanged.

A median filter is used because the same process can be performed in the decoder with no additional data transmission. The median vector is used as a prediction, and comparison with the actual vector enables a residual to be computed and coded for transmission. At the decoder the same prediction can be made and the received residual is added to re-create the original vector.

The remaining parts of Figure 5.51(b) show how the remaining three DCT block vectors are predicted from adjacent DCT block vectors. If the relevant block is only macroblock coded, that vector will be substituted.

## 5.22 Shape coding

1	1	1	2
1	1	2	2
1	1	2	2
1	1	2	2
1	1	2	2
1	1	2	2
1	1	1	2

g values  
ith motion  
f Y blocks  
nt of current  
emote vector)

s to all the  
only over-  
isables the

or residual  
ck block is  
he flag and

ne or four  
of a current  
ck vectors.

or the top-  
used. Three  
vectors, as  
he diagram  
ed and the  
. The three  
ector in the

erformed in  
an vector is  
r enables a  
decoder the  
dded to re-

ining three  
ctors. If the  
r will be

Shape coding is the process of compressing alpha or keying data. Most objects are opaque and so a binary alpha signal is adequate for the base-level shape system. For each texture pixel, an alpha bit exists forming a *binary alpha map* which is effectively a two-dimensional mask through which the object can be seen. At the decoder binary alpha data are converted to 000 or 255 levels in an eight-bit keying system.

Optionally the object can be faded in the compositing process by sending a *constant alpha value* at each VOP. As a further option, variable transparency can be supported by sending *alpha texture data*. This is coded using the usual MPEG spatial coding tools such as DCT and scanning etc.

Binary data such as alpha data do not respond to DCT coding, and another compression technique has been developed for binary alpha blocks (bab). This is known as *context-based coding* and it effectively codes the location of the boundary between alpha one bits and alpha zero bits. Clearly once the boundary is located the values of all remaining alpha bits are obvious.

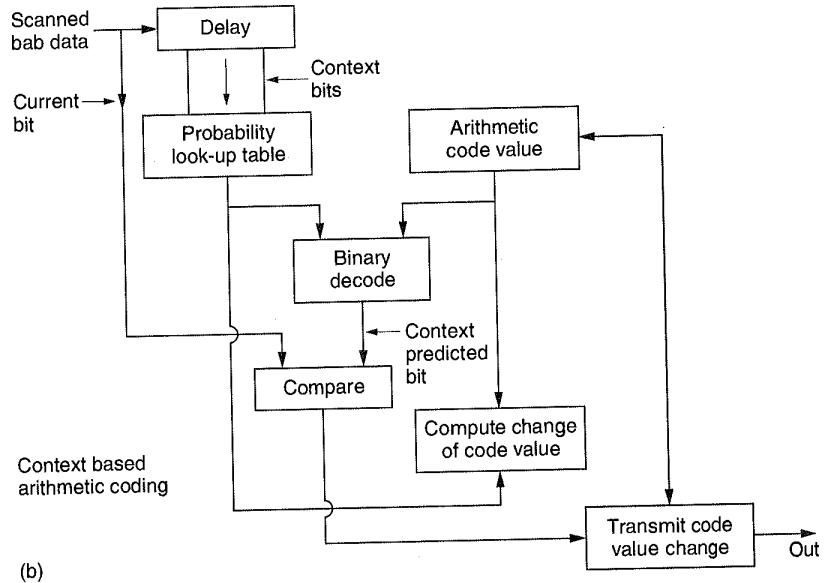

The babs are raster scanned into a serial bitstream. Context coding works by attempting to predict the state of the current bit from a set of bits which have already been decoded. Figure 5.53(a) shows the set of bits, known as a context, used in an intra-coded VOP. There are ten bits in the context and so there can be 1024 different contexts. Extensive analysis of real shapes shows that for each context there is a certain probability that the current bit will be zero. This probability exists as a standardized look-up table in the encoder and decoder.

Figure 5.53(b) shows that the encoder compares the probability with a parameter called the arithmetic code value to predict the value of the alpha bit. If the prediction is correct, the arithmetic code value needs no change and nothing needs to be transmitted. However, if the predicted bit is incorrect, a change to the arithmetic code value must be transmitted, along with the position of the bit in the scan to which it applies. This change is calculated to be just big enough to make the prediction correct. In this way only prediction errors, which represent the boundary location where alpha changes, are transmitted and the remaining bits self-predict.


Figure 5.54 shows that the spatial extent of the context requires alpha bits from adjacent babs to code the present bab. This results in the concept of the *bordered bab* which incorporates a two-pixel deep border of alpha bits above and to each side. In the case of babs adjacent to top or sides of the bounding rectangle, some of the bounding bits are outside that rectangle and their value is obviously always zero.

Shape data can also be inter-coded. This is handled using a context which spreads across two VOPs. As Figure 5.55 shows, four context bits

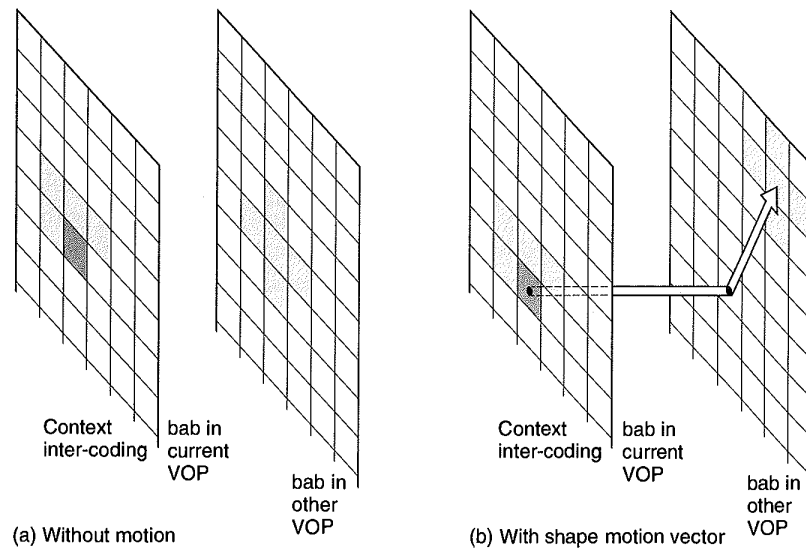




**Figure 5.53** (a) The context is a set of bits used for prediction of the current bab bit. (b) The encoder compares the actual bit with the predicted value and modifies the arithmetic code in order to make the prediction correct at the decoder.



**Figure 5.54** A bordered bab contains bits from adjacent babs so that context bits are available.



**Figure 5.55** (a) In bab inter-coding, the context spreads over the current and the previous VOP. (b) Shape motion vectors shift the context with respect to the reference VOP to improve the accuracy of prediction.

are in the current VOP, whereas five context bits are in a previous VOP. This nine-bit context requires a different probability table, but otherwise coding proceeds as before. If the shape changes from one VOP to another, motion compensation can be used. A shape vector is transmitted. It will be seen from Figure 5.55(b) that this vector shifts the context relative to the previous VOP, thus altering (hopefully improving) the chances of correctly predicting the current bit.

Within a boundary block, shape vectors are highly redundant with the motion vectors of texture and texture vectors may be used to predict shape vectors.

It might be thought that binary shape data would result in objects having ratcheted edges because the key signal can only exist in pixel steps. However, texture coding can overcome this problem. If a binary keyed image is used as a predictor, the texture prediction error will deliver the necessary accuracy in the final image.

## 5.23 Padding

The peculiar conditions in boundary macroblocks where some of the block is inside the object and some outside requires special handling and this is the function of padding. There are two types of padding which

code  
age  
Out  
t bab bit. (b)  
he

t bits are

perform distinct functions; texture padding and motion compensation (MC) padding.

Texture padding makes the spatial coding of boundary blocks more efficient. As the frequency coefficients of the DCT relate to the entire block area, detail outside the object can result in coefficients which are not due to the object itself. Texture padding consists of replacing those pixels in a boundary block which are outside the object with pixel values that are selected to minimize the amount of coefficient data. The exact values are unimportant as they will be discarded in the decoder.

When objects move they may also change their shape. Motion compensation will be used whereby texture and shape from another VOP is used to predict texture and shape in the current VOP. If there is a small error in the boundary data, background pixels could leak into the object causing a large prediction error. MC padding is designed to prevent this. In the case of VOPs used as anchors for inter-coding, in boundary blocks, after decoding, texture padding is replaced with MC padding. In MC padding, pixel values at the very edge of the object are copied to pixels outside the object so that small shape errors can be corrected with minor residual data.

## 5.24 Video object coding

Figure 5.56 shows the video object coding process. At the encoder, image analysis dissects the images into objects, in this case one small object (1) and a background object (2). The picture and shape data for object 1 are used to create a bounding rectangle whose coordinates are coded. Pixel values outside the boundary are replaced with texture padding. The texture macroblocks are then DCT coded and the shape data are context coded. This can be on an intra- or motion-compensated inter-basis according to the type of VOP being coded. Background data are coded by another process and multiplexed into the bitstream as another object.

At the decoder the texture and shape data are decoded, and in anchor VOPs the shape data are used to strip off the texture padding and replace it with MC padding. MC padded object data are used to predict intra-coded objects. Finally the shape data key the object into the background.

Figure 5.57 shows an object coder. The object analysis of the incoming image provides texture and texture motion as well as object shape and shape motion. Texture coding takes place in much the same way as MPEG picture coding. Intra-VOPs are spatially coded, whereas P- and B-VOPs use motion-compensated prediction and a spatially coded residual is transmitted. The vectors in MPEG-4 are subject to a prediction process not found in MPEG-2. Note the use of texture padding before the



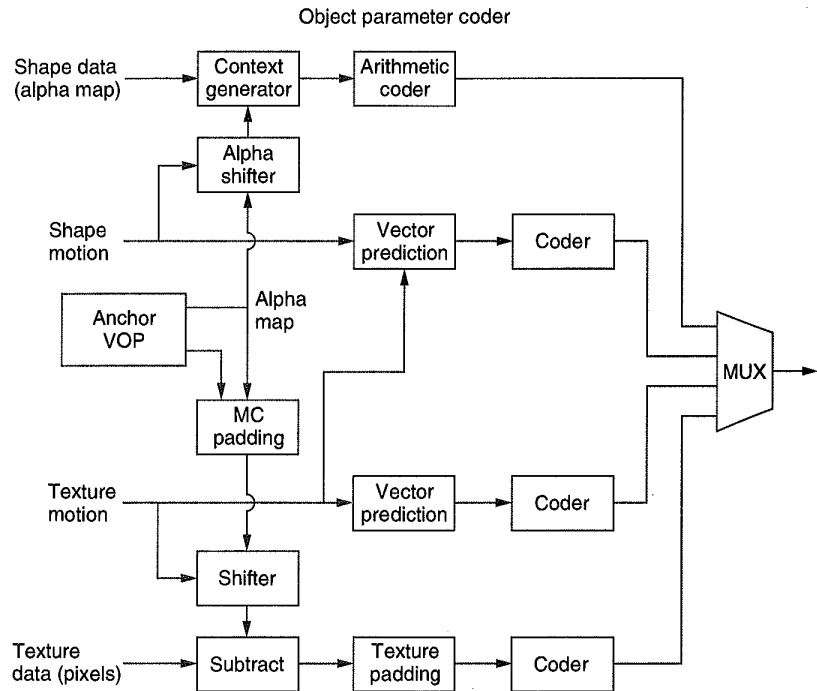


Figure 5.57 An MPEG-4 object coder. See text for details.

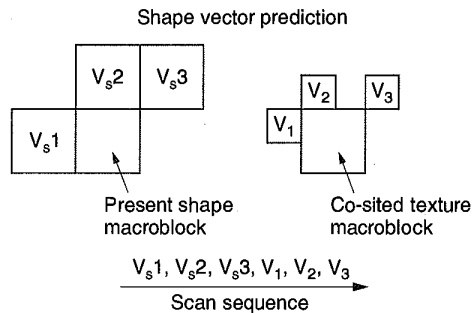


Figure 5.58 Shape vectors may be coded or taken from texture vectors. The decoder scans vectors in the sequence shown and takes the first ones available.

Shape babs can exist in a variety of forms as Figure 5.47 showed. Opaque and transparent blocks are easily and efficiently coded. Intra babs have no shape vector but need a lot of data. Inter babs may or may not contain a shape vector according to object motion and only a residual is sent. On the other hand, under certain object conditions, sufficient shape accuracy may result by sending only the shape vector.

## 5.25 Two-dimensional mesh coding

Mesh coding was developed in computer graphics to aid the rendering of synthetic images with perspective. As Figure 5.59 shows, as a three-dimensional body turns, its two-dimensional appearance changes. If the body has a flat side square-on to the optical axis whose texture is described by uniformly spaced pixels, after a rotation the pixels will no longer be uniformly spaced. Nearly every pixel will have moved, and clearly an enormous amount of data would be needed to describe the motion of each one. Fortunately this is not necessary; if the geometry of the object is correctly sampled before and after the motion, then the shift of every pixel can be computed.




Figure 5.59 As a body turns, its shape on a two-dimensional screen may change.




Figure 5.60 All the nodes of a two-dimensional mesh lie in the same plane.

The geometric sampling results in a structure known as a mesh. Effectively a mesh is a sampled surface. The smoother the surface, the fewer samples are necessary because the remainder of the surface can be interpolated. In MPEG-4, meshes can be two- or three-dimensional. Figure 5.60 shows a two-dimensional mesh which is a set of points known as *nodes* which must remain in the video object plane. Figure 5.61 shows a three-dimensional mesh which is a set of nodes describing a non-flat surface.

Like other MPEG-4 parameters, meshes can be scaleable. A base mesh has relatively few nodes; adding enhancement data will increase the number of nodes so that the shape is better described.

ecoder

showed.  
ed. Intra  
y or may  
residual  
sufficient




Figure 5.61 A three-dimensional mesh samples the surface of a solid body.




Figure 5.62 In mesh coding, an object is coded as texture, shape and mesh data. The I-VOP mesh may be regular and need few data to code it.

Figure 5.62 shows an example of how two-dimensional meshes are used to predict an object from a previous VOP. The first object may be sent as an I-VOP, including texture and alpha outline data, both of which will be intra-coded. The texture pixels in the I-VOP are uniformly spaced and so a uniform mesh can be sent as a further part of the same I-VOP. Uniform meshes require very few data. Figure 5.63(a) shows that little





Figure 5.63 (a) The initial mesh may be regular which requires few data. (b) A mesh may be warped by transmitting vectors which displace the nodes.

more than the coordinates of the top-left corner of the mesh and the mesh pitch need be sent. There are also some choices of how the mesh boxes are divided into triangles.

As an alternative to a regular mesh, an irregular or *Delaunay* mesh may be created by specifying the location of each vertex. This will require more data initially, but may prove to be more efficient in the case of an irregular object which exists for a long time.

Figure 5.63(b) shows how the shape of the mesh can be changed by transmitting mesh vectors in P-VOPs. These specify the new position of





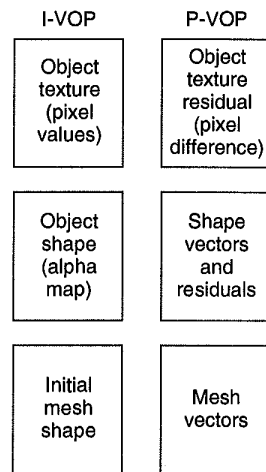
**Figure 5.64** Mesh vectors are predicted from the previous two vertices of the triangle concerned.

each vertex with respect to the original position. The mesh motion vectors are predictively coded as shown in Figure 5.64. In two adjacent triangles there are two common nodes. If the first triangle has been fully coded, only one node of the second triangle remains to be coded. The vector for the remaining node is predicted by taking the average of the two common nodes and rounding to half-pixel accuracy. This prediction is subtracted from the actual vector to produce a residual which is then variable-length coded for transmission. As each node is coded, another triangle is completed and a new pair of common nodes become available for further prediction.

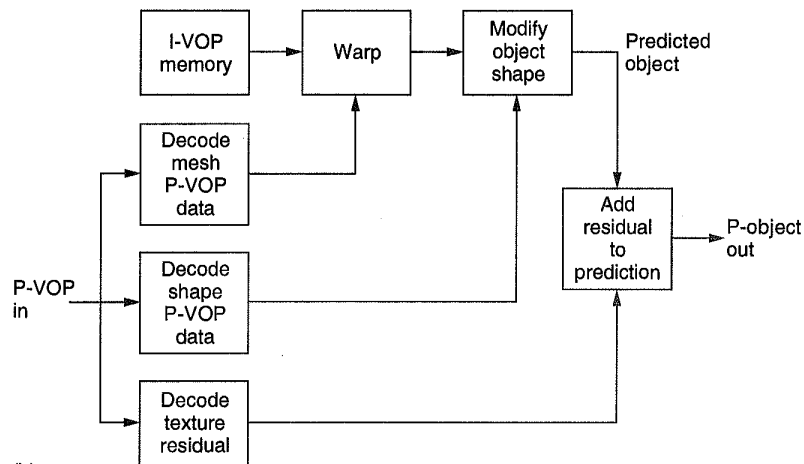
Figure 5.65 shows an object coding system using meshes. At (a), following the transmission of the I-VOP containing shape, outline and a regular mesh, the I-VOP and the next VOP, which will be a P-VOP, are supplied to a motion estimator. The motion estimator will determine the motion of the object from one VOP to the next. As the object moves and turns, its shape and perspective will change. The new shape may be encoded using differential alpha data which the decoder uses to change the previous shape into the current shape. The motion may also result in a perspective change. The motion estimator may recognize features in the object which are at a new location and this information is used to create mesh vectors.

At the encoder, mesh vectors are used to warp the pixels in the I-VOP to predict the new perspective. The differential shape codes are used to predict the new shape of the object. The warped, reshaped object is then compared with the actual object, and the result is a texture residual which will be DCT coded.

At the decoder, (b) the I-VOP is in memory and the mesh motion, the shape difference and the texture residual arrive as a P-VOP. The mesh motion is decoded and used to warp the object pixels taken from the



(a)



(b)

**Figure 5.65** (a) Mesh coding requires texture, shape and mesh data to be multiplexed into the bitstream. (b) An object coding system using meshes. See text.

I-VOP memory. The shape difference is decoded to produce a new shape around the warped pixels, resulting in a predicted object. The texture residual is decoded and added to the predicted object to produce the final decoded object which is then composited into the output picture along with other objects.

It should be clear that by transmitting warping vectors the pixels from the I-VOP can be used to create a very accurate prediction of the texture of the P-VOP. The result is that the texture residual can be encoded with a very small number of bits.

## 5.26 Sprites

Sprites are images or parts of images which do not change with time. Such information only needs to be sent once and is obviously attractive for compression. However, it is important to appreciate that it is only the transmitted sprite which doesn't change. The decoder may manipulate the sprite in a different way for each picture, giving the illusion of change.

Figure 5.66 shows the example of a large sprite used as a background. The pixel array size exceeds the size of the target picture. This can be sent as an I-VOP (*intra-coded video object plane*). Once the texture of the sprite is available at the decoder, vectors can be sent to move or warp the picture. These instructions are sent as S-VOPs (*static video object planes*). Using S-VOP vectors, a different part of the sprite can be delivered to the display at each picture, giving the illusion of a moving background. Other video objects can then be keyed over the background.





Figure 5.66 A large sprite used as a background may be bigger than the screen. If shifted by vectors, the illusion of a moving background is given.

Sending a large background sprite in one I-VOP may cause difficulty at high compression factors because the one-off image distorts the bit rate. The solution here is *low-latency* sprite coding. Although the sprite is a single still texture object, it need not all be transmitted at once. Instead the sprite is transmitted in *pieces*. The first of these is an *object piece*, whereas subsequently *update pieces* may be sent.




**Figure 5.67** An object piece may be used in several VOPs. (a) To allow a zoom-in, part of the sprite object may have its resolution enhanced by an additional object piece. (b) An additional object piece may extend an object to allow a pan.

Update pieces may coincide spatially with the object piece and result in a refinement of image quality, or may be appended spatially to increase the size of a sprite. Figure 5.67 shows the example of an object piece which is transmitted once but used in a series of pictures. In Figure 5.67(a) an update piece is transmitted which improves the quality of the sprite so that a zoom-in can be performed.

Figure 5.67(b) shows a large sprite sent to create a moving background. Just before the edge of the sprite becomes visible, the sprite is extended by sending an update piece. Update pieces can be appended to pixel accuracy because they are coded with offset parameters which show where the update fits with respect to the object piece.

Figure 5.68 shows a sprite decoder. The basic sprite is decoded from an I-VOP and placed in the sprite buffer. It does not become visible until an S-VOP arrives carrying shifting and warping instructions. Low-latency decoding sends sprite pieces as S-VOPs which are assembled in the buffer as each sprite piece arrives.



**Figure 5.68** An MPEG-4 sprite decoder. See text for details.

## 5.27 Wavelet-based compression

MPEG-4 introduces the use of wavelets in still picture object coding. Wavelets were introduced in Chapter 3 where it was shown that the transform itself does not achieve any compression. Instead the picture information is converted to a form in which redundancy is easy to find. Figure 5.69 shows three of the stages in a wavelet transform. Each of these




Figure 5.69 Three of the stages in a wavelet transform in which different levels of detail exist.

represents more detail in the original image, assuming, of course, that there was any detail there to represent. In real images, the highest spatial frequencies may not be present except in localized areas such as edges or in detailed objects. Consequently many of the values in the highest resolution difference data will be zero or small. An appropriate coding scheme can achieve compression on data of this kind. As in DCT coding, wavelet coefficients can be quantized to shorten their wordlength, advantageous scanning sequences can be employed and arithmetic coding based on coefficient statistics can be used.

MPEG-4 provides standards for scaleable coding. Using scaleable coding means that images can be reproduced both at various resolutions and with various noise levels. In one application a low-cost simple decoder will decode only the low-resolution noisy content of the bitstream whereas a more complex decoder may use more or all of the data to enhance resolution and to lower the noise floor. In another application where the bit rate is severely limited, the early part of a transmission may be decoded to produce a low-resolution image whose resolution increases as further detail arrives.

The wavelet transform is naturally a multi-resolution process and it is easy to see how a resolution scaleable system could be built from Figure 5.69 by encoding the successive difference images in separate bitstreams. As with DCT-based image compression, the actual bit rate reduction in wavelet compression is achieved by a combination of lossy and lossless coding. The lossless coding consists of scanning the parameters in a manner which reflects their statistics, followed by variable-length coding.

The lossy coding consists of quantizing the parameters to reduce their wordlength. Coarse quantisation may be employed on a base level bitstream, whereas an enhancement bitstream effectively transmits the quantizing error of that coarse quantization.

The wavelet transform produces a decomposition of the image which is ideal for compression. Each sub-band represents a different band of spatial frequencies, but in the form of a spatial array. In real images, the outline of every object consists of an edge which will contain a wide range of frequencies. There will be energy in all sub-bands in the case of a sharp edge. As each sub-band is a spatial array, edge information will be found in the same place in each sub-band. For example, an edge placed one quarter of the way across the picture will result in energy one quarter of the way across each sub-band.

This characteristic can be exploited for compression if the sequence in which the coefficients are coded starts at a given coefficient in the lowest sub-band and passes through each sub-band in turn. Figure 5.70 shows that for each coefficient in a given sub-band there will be four *children* in the next, sixteen *descendants* in the next and so on, in a tree structure. In the case of a sharp edge, all the relevant coefficients will generally lie on a single route through the tree. In plain areas containing only low frequencies, the corollary of this characteristic is that generally if a given


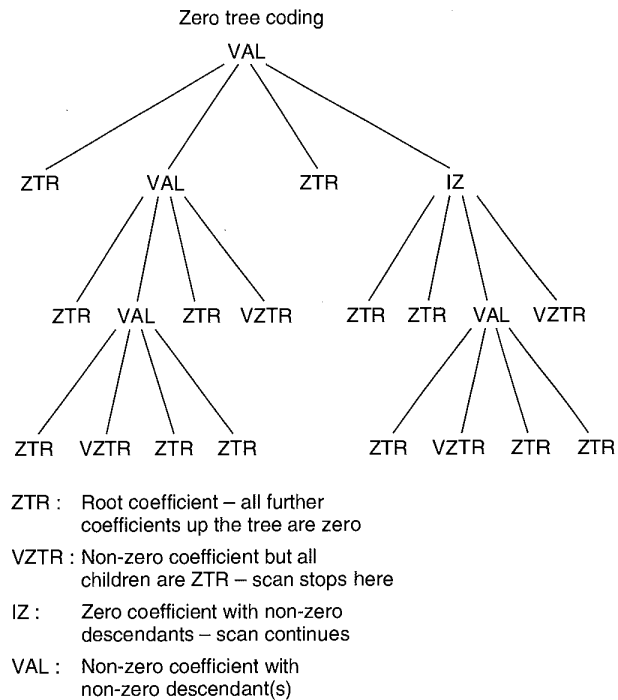


Figure 5.70 In wavelet coding, every coefficient at a given resolution corresponds to four coefficients at the next highest level.




**Figure 5.71** Zerotree coding relies on the redundancy in wavelet coefficients from the lowest resolution upwards in a tree structure.

coefficient has value zero, then all of the coefficients in higher sub-bands will frequently also have value zero, producing a *zerotree*.<sup>3</sup> Zerotrees are easy to code because they simply need a bit pattern which is interpreted as meaning there are no further non-zero coefficients up this tree.

Figure 5.71 shows how zerotree coding works. Zero frequency coefficients are coded separately using predictive coding, so zerotree coding starts with coefficients in the lowest AC sub-band. Each coefficient represents the root of a tree and leads to four more nodes in the next sub-band and so on. At nodes, one of four codes may be applied. ZTR means that this node is at the bottom of a zerotree and that this all further coefficients up the tree are zero and need not be coded. VZTR means that this coefficient is non-zero but all its descendants are. In this case only the valid coefficient at the foot of the tree is sent. IZ means that this coefficient is zero but it has non-zero descendants and the scan must continue. Finally VAL means a non-zero coefficient with non-zero descendants. This coefficient and subsequent ones must be coded.

Zerotree coding is efficient because in real images the energy is concentrated at low frequency and most trees do not need to be coded as far as the leaf node.



**Figure 5.72** (a) In tree-depth scanning the coefficients are scanned up each tree in turn. (b) In band-by-band scanning the coefficients are transmitted in order of increasing resolution.

It is important to appreciate that the scanning order in which zerotree coding takes place must go from the lowest AC sub-band upwards. However, when the coefficients have been designated by the zerotree coder into the four categories above, they can be transmitted in any order which is convenient. MPEG-4 supports either tree depth scanning or band-by-band scanning. Figure 5.72(a) shows that in tree-depth scanning the coefficients are transmitted one tree at a time from the lowest spatial frequency to the highest. The coding is efficient but the picture will not be available until all the data have been received. Figure 5.72(b) shows that in band-by-band scanning the coefficients are transmitted in order of ascending spatial frequency. At the decoder a soft picture becomes available first, and the resolution gradually increases as more sub-bands are decoded. This coding scheme is less efficient but the decoding latency is reduced which may be important in image database browsing applications.

Figure 5.73(a) shows a simple wavelet coder operating in single\_quant mode. After the DWT, the DC or root coefficients are sent to a predictive coder where the values of a given coefficient are predicted from adjacent coefficients. The AC coefficients are zerotree coded. Coefficients are then quantized and arithmetic coded. The decoder is shown in (b). After arithmetic decoding, the DC coefficients are obtained by inverse prediction and the AC coefficients are re-created by zerotree decoding. After inverse quantizing the coefficients control an inverse DWT to re-create the texture.



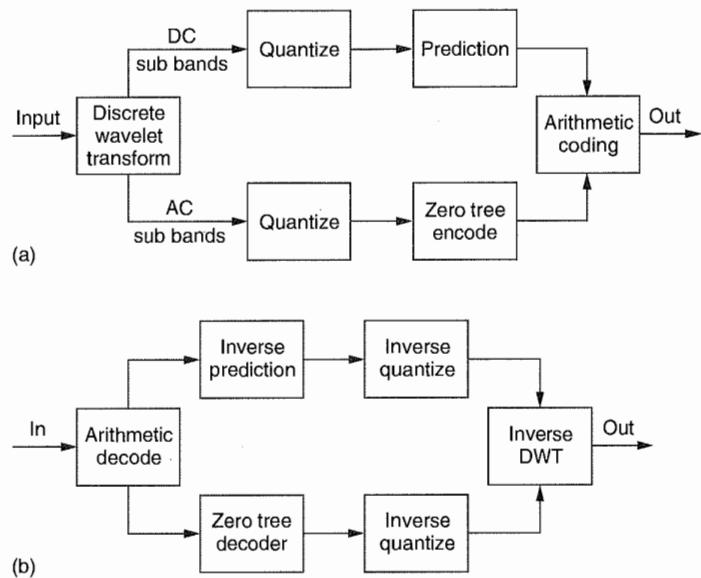



Figure 5.73 (a) A simple wavelet encoder. (b) The corresponding decoder.

MPEG-4 supports scaleable wavelet coding in `multi_quant` and `bi-level_quant` modes. Multi-quant mode allows resolution and noise scaleability. In the first layer, a low-resolution and/or noisy texture is encoded. The decoder will buffer this texture and optionally may display it. Higher layers may add resolution by delivering higher sub-band coefficients or reduce noise by delivering the low-order bits removed by the quantizing of the first layer. These refinements are added to the contents of the buffer.

In `bi-level_quant` mode the coefficients are not sent as binary numbers, but are sent bitplane-by-bitplane. In other words the MSBs of all coefficients are sent first followed by all the second bits and so on. The decoder would reveal a very noisy picture at first in which the noise floor falls as each enhancement layer arrives. With enough enhancement layers the compression could be lossless.

## 5.28 Three-dimensional mesh coding

In computer-generated images used, for example, in simulators and virtual reality, the goal is to synthesize the image which would have been seen by a camera or a single eye at a given location with respect to the virtual objects. Figure 5.74 shows that this uses a process known as *ray tracing*. From a fixed point in the virtual camera, a ray is projected




**Figure 5.74** In ray tracing, imaginary rays are projected from the viewer's eye through each screen pixel in turn to points on the scene to be rendered. The characteristics of the point in the scene are transferred to the pixel concerned.

outwards through every pixel in the image to be created. These rays will strike either an object or the background. Rays which strike an object must result in pixels which represent that object. Objects which are three-dimensional will reflect the ray according to their geometry and reflectivity, and the reflected ray has to be followed in case it falls on a source of light which must appear as a reflection in the surface of the object. The reflection may be sharp or diffuse according to the surface texture of the object. The colour of the reflection will be a function of the spectral reflectivity of the object and the spectrum of the incident light.

If the geometry, surface texture and colour of all objects and light sources are known, the image can be computed in a process called *rendering*. If any of the objects are moving, or if the viewpoint of the virtual camera changes, then a new image will need to be rendered for each video frame. If such synthetic video were to be compressed by, say MPEG-2, the resultant bit rate would reflect the large number of pixels changing from frame to frame. However, the motion of virtual objects can be described by very few bits. A solid object moves in three dimensions and can rotate in three axes and this is trivial to transmit. A flexible object such as a jellyfish requires somewhat more information to describe how its geometry changes with time, but still much less than describing the changes in the rendered video.

Figure 5.75 shows that in MPEG-4, the rendering process is transferred to the decoder. The bitstream describes objects to the decoder which then renders the virtual image. The bitstream can then describe the motion of the objects with a very low bit rate, and the decoder will render a new image at each frame. In interactive applications such as simulators, CAD and video games, the user may change the viewpoint of the virtual camera and/or the position of virtual objects and result in an unique




**Figure 5.75** Three-dimensional mesh coding allows the rendering process to be completed in the decoder.

rendering. This does not require any communication between the user and the encoder because the user is only interacting with the rendering process.

Transmitting raw geometry requires a large amount of data. Space is quantized in three axes and for every location in, say, the  $x$ -axis, the  $y$  and  $z$  coordinates of the surface must be transmitted. Three-dimensional mesh coding is an efficient way of describing the geometry and properties of bodies in a compressed bitstream. Audio and video compression relies upon the characteristics of sounds, images and the human senses, whereas three-dimensional mesh coding relies on the characteristics of topology.<sup>4</sup>

Figure 5.76(a) shows a general body in which the curvature varies. Sampling this body at a constant spacing is obviously inefficient because if the samples are close enough together to describe the sharply curved areas, the gently curved areas will effectively be heavily oversampled. Three-dimensional mesh coding uses irregular sampling locations called *vertices* which can be seen in Figure 5.76(b). A vertex is a point in the surface of an object which has an unique position in  $x$ ,  $y$  and  $z$ . Shape coding consists of efficiently delivering the coordinates of a number of




**Figure 5.76** (a) An arbitrary body with variable curvature. (b) The surface can be sampled by vertices which are not necessarily evenly spaced. (c) Filtering the vertices re-creates the shape of the body.

vertices to the decoder. Figure 5.76(c) shows that the continuous shape can then be reconstructed by filtering as all sampled processes require.

Filtering of irregularly spaced samples is more complex than that of conventional reconstruction filters and uses algorithms called splines.<sup>5</sup> Failure correctly to filter the shape results in synthetic objects which appear multi-faceted instead of smoothly curved.

Figure 5.77(a) shows that a set of three vertices may be connected by straight lines called *edges* to form a triangle. In fact this is how the term 'vertex' originates. Note that these edges are truly straight and are thus not necessarily on the surface of the object except at the vertices. There are a number of advantages to creating these triangles. Wherever the vertices are, a triangle is always flat, which is not necessarily the case for shapes with four or more sides. A flat surface needs no further description because it can be interpolated from the vertices. Figure 5.77(b) shows that by cutting along certain edges and hingeing along others, a triangulated body can be transformed into a plane. (A knowledge of origami is useful here.) Figure 5.77(c) shows that once a triangle has been transmitted to the decoder, it can be used as the basis for further information about the surface geometry. In a process known as a *forest split operation*, new vertices are located with respect to the original triangle.<sup>6</sup>



**Figure 5.77** (a) Vertices are joined by edges to form triangles. (b) Hingeing along some edges and opening others allows the triangles to be unfolded into a plane. (c) Greater resolution can be obtained by locating additional vertices with respect to each triangle in a forest split operation.

Forest split operations are efficient because the displacements of the new vertices from the original triangle plane are generally small. Additionally a forest split is only needed in triangles in the vicinity of sharp changes. Forest splits can be cascaded allowing the geometry to be reconstructed to arbitrary accuracy. By encoding the forest split data in different layers, a scaleable system can be created where a rough shape results from the base layer which can be refined by subsequent layers.

Figure 5.78 shows an MPEG-4 three-dimensional mesh coding system at high level. The geometry of the object is coded by the vertices which are points, whereas the properties are coded with respect to the triangles between the vertices. The dual relationship between the triangles and the vertices which they share is known as the *connectivity*. Geometry is compressed using topological principles and properties are compressed by quantizing. These parameters are then subject to entropy (variable-length) coding prior to multiplexing into a bitstream. At the decoder the three information streams are demultiplexed and decoded to drive the rendering and composition process. The connectivity coding is logical and must be lossless whereas the vertex positions and triangle properties may be lossy or lossless according to the available bit rate.

Generic geometry coding is based on meshes using polygons. Each polygon is defined by the position of each vertex and its surface properties. MPEG-4 breaks such polygons down into triangles for coding the vertices and associates the properties with the correct triangles at the decoder using the connectivity data so that the polygons can be re-created. The geometry coding is based on triangles because they have useful properties.

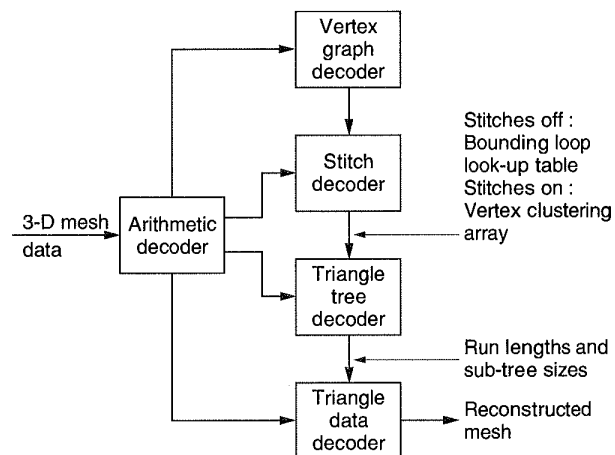


Figure 5.78 An MPEG-4 three-dimensional mesh coder. See text for details.




Figure 5.79 Triangles are binary: entering by one side leaves a choice of two sides by which to leave.


Figure 5.79 shows that triangles are fundamentally binary. Digital transmission is serial, and compression works by finding advantageous sequences in which to send information, for example the scanning of DCT coefficients. When scanning triangles, if a scan enters via one edge, there are then only two sides by which to exit. This decision can be represented by a single bit for each triangle, known as a *marching bit*.

Figure 5.80(a) shows a representative object, in this case a spheroid. Its surface has been roughly sampled to create twelve vertices, and edges are drawn between the vertices to create triangles. The flat triangles create a dodecahedron, and this may later be refined more closely to the spheroid with a forest split operation. For the moment the goal is to transmit the vertices so that the triangles can be located. As adjacent triangles share vertices, there are fewer vertices than triangles. In this example there are twenty triangles but only twelve vertices.

When a vertex is transmitted, the information may be required by more than one triangle. Transferring this information to the next triangle in the scan sequence is obvious, but where the scan branches it is non-obvious and a structure known as *vertex loop look-up table* is needed to connect vertex data. Only one vertex needs to have its absolute coordinates transmitted. This is known as the *root vertex*. The remaining vertices can be coded by prediction from earlier vertices so that only the residual needs to be sent.

In general it is impossible for a continuous scan sequence to visit all of the triangles. Figure 5.80(b) shows that if cuts are made along certain edges, the triangles can be hinged at their edges. This is continued until they all lie in the same plane to form a polygon (c). This process has created a binary tree along which the triangles can be scanned, but it has also separated some of the vertices. Figure 5.80(c) apparently has 22 vertices, but ten of these are redundant because they simply repeat the locations of the real vertices. Figure 5.80(d) shows how the vertex loop look-up process can connect a single coded vertex to all its associated triangles.

In graphics terminology, the edges along which the polyhedron was cut form a rooted vertex tree. Each run of the vertex tree terminates at a point where a cut ended. The extremities of the vertex tree correspond to the



**Figure 5.80** (a) A spheroid which has been sampled by vertices and triangulated. (b) Partially unfolding the triangles. (c) Completely unfolded triangles form a planar polygon. (d) The vertex loop which drives the vertex look-up process in order to link the vertices which were disconnected during unfolding.

concavities in the polygon. When the vertex loop is collapsed to zero area in the process of folding the polygon into a solid body it becomes a vertex tree.

A solid body such as a spheroid is described as *manifold* because its surface consists of a continuous layer of faces with a regular vertex at each corner. However, some bodies which need to be coded are not manifold. Figure 5.81 shows a non-manifold body consisting of a pair of intersecting planes. The two planes share an edge joining shared vertices. The two planes can be coded separately and joined using *stitches* which indicate common vertices.




Figure 5.81 A non-manifold body in which planes are joined by stitches.

Stitches can also be used to transmit meshes in discrete stages known as *partitions*. Partitions can be connected together by stitches such that the result is indistinguishable from a mesh coded all at once. There are many uses of mesh partitions, but an example is in the coding of a rotating solid body. Only the part of the body which is initially visible needs to be transmitted before decoding can begin. Extra partitions can be transmitted which describe that part of the surface of the body that will come into view as the rotation continues. This approach is the three-dimensional equivalent of transmitting static sprites as pieces; both allow the bit rate to remain more constant.

Figure 5.82 shows that when scanning, it is possible to enter four types of triangle. Having entered by one edge, the other two edges will either be internal edges, leading to another triangle, or a boundary edge. This can be described with a two-bit code; one bit for the left edge and one bit for the right edge. The scanning process is most efficient if it can create long *runs* which are linear sequences of triangles between nodes. Starting at a boundary edge, the run will continue as long as type 01 or 10 triangles are found. If a type 11 triangle is found, there are two runs from this node.





Figure 5.82 The four types of triangle which can be encountered when scanning.



The encoder must remember the existence of nodes in a stack, and pick one run until it ends at a type 00 triangle. It must then read the stack and encode the other run. When the current run has reached a 00 triangle and the stack is empty, all the triangles must have been coded.

There is redundancy in the two-bit code describing the triangles. Except for the end triangles, the triangles in a run will have the value 01 or 10. By transmitting the *run length*, the decoder knows where the run ends. As only two values are possible in the run, 01 can be coded as 0 and 10 can be coded as 1 and only this needs be coded for each triangle. This bit is, of course, the marching bit which determines whether the run leaves the triangle by the left edge or the right. Clearly in type 00 or 11 triangles there is no marching bit.

Figure 5.83 shows how the triangles are coded. There are seven runs in the example shown, and so seven run lengths are transmitted. Marching



**Figure 5.83** An example of triangle coding using the body of Figure 5.80. Runs of triangles are created. See text for details.

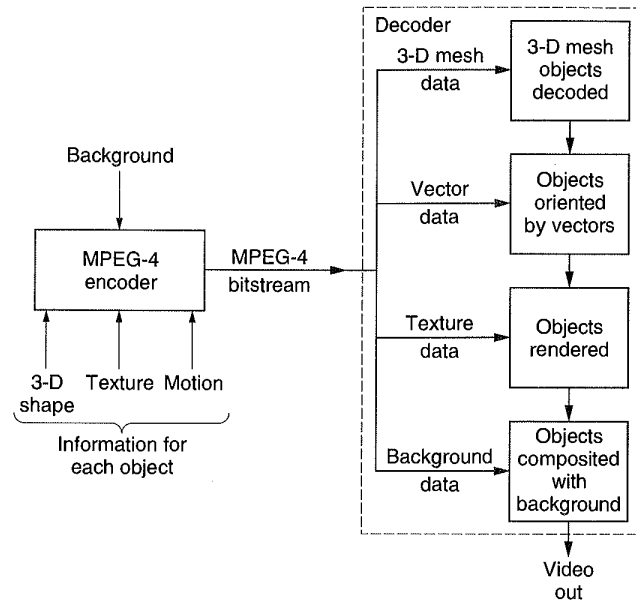


Figure 5.84 A three-dimensional mesh codec. See text for details.

bits are only transmitted for type 01 or 10 triangles and the decoder can determine how the sequence of marching bits is to be decoded using the run lengths.

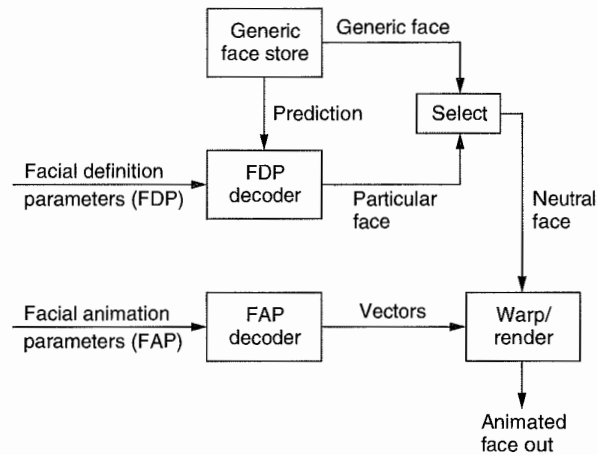
Figure 5.84 shows a three-dimensional mesh codec. This decoder demultiplexes the mesh object bitstream into data types and reverses the arithmetic coding. The first item to be decoded is the vertex graph which creates the vertex loop look-up table. If stitches are being used these will also be decoded so that in conjunction with the vertex graph look-up table a complete list of how triangles are to be joined to create the body will result.

The triangle topology is output by the triangle tree decoder. This decodes the run lengths and the marching bits of the runs which make up the polygon. Finally the triangle data are decoded to produce the vertex coordinates. The coordinates are linked by the connectivity data to produce the finished mesh (or partition).

The faces of the mesh then need to be filled in with properties so that rendering can take place.

## 5.29 Animation

In applications such as videophones and teleconferencing the images to be coded frequently comprise the human face in close-up or the entire



**Figure 5.85** Face animation decoder used facial definition parameters (FDP) to create a neutral face. Facial animation parameters (FAPs) are specialized vectors which move the face vertices to reproduce expressions.

body. Video games may also contain virtual human actors. Face and body animation (FBA) is a specialized subset of mesh coding which may be used in MPEG-4 for coding images of, or containing, humans at very low bit rate.

In face coding, the human face is constructed in the decoder as a three-dimensional mesh onto which the properties of the face to be coded are mapped. Figure 5.85 shows a facial animation decoder. As all faces have the same basic layout, the decoder contains a generic face which can be rendered immediately. Alternatively the generic face can be modified by *facial definition parameters* (FDP) into a particular face. The data needed to do this can be moderate.

The FDP decoder creates what is known as a neutral face; one which carries no expression. This face is then animated to give it dynamically changing expressions by moving certain vertices and re-rendering. It is not necessary to transmit data for each vertex. Instead face-specific vectors known as *face animation parameters* (FAPs) are transmitted. As certain combinations of vectors are common in expressions such as a smile, these can be coded as *visemes* which are used either alone or as predictions for more accurate FAPs. The resulting vectors control algorithms which are used to move the vertices to the required places. The data rate needed to do this is minute; 2–3 kilobits per second is adequate.


In body animation, the three-dimensional mesh created in the decoder is a neutral human body; i.e. one with a standardized posture. All the degrees of freedom of the body can be replicated by transmitting codes corresponding to the motion of each joint.

In both cases if the source material is video from a camera, the encoder will need sophisticated algorithms which recognize the human features and output changing animation parameters as expression and posture change.

### 5.30 Scalability

MPEG-4 offers extensive support for scalability. Given the wide range of objects which MPEG-4 can encode, offering scalability for virtually all of them must result in considerable complexity. In general, scalability requires the information transmitted to be decomposed into a base layer along with one or more enhancement layers. The base layer alone can be decoded to obtain a picture of a certain quality. The quality can be improved in some way by adding enhancement information.

Figure 5.86 shows a generic scalable decoder. There is a decoder for the base layer information and a decoder for the enhancement layer. For efficient compression, the enhancement layer decoder will generally require a significant number of parameters from the base layer to assist with predictively coded data. These are provided by the *mid-processor*. Once the base and enhancement decoding is complete, the two layers are combined in a post-processor.



**Figure 5.86** A generic scalable decoder contains decoding for base and enhancement layers as well as a mid-processor which adapts the base layer data to the needs of the enhancement process.

Figure 5.87 shows a spatial scalability system in which the enhancement improves resolution. Spatial scalability may be applied to rectangular VOPs or to arbitrarily shaped objects using alpha data. In the case of irregular objects the alpha data will also need to be refined. The mid-processor will upsample the base layer images to provide base pixel data at the new, higher resolution, sampling grid. With respect to a sampling array upsampled by a factor of two, motion vectors appear half as accurate and there are one quarter as many as are needed. Consequently to improve prediction in the enhancement layer, the motion compensation must be made more accurate.

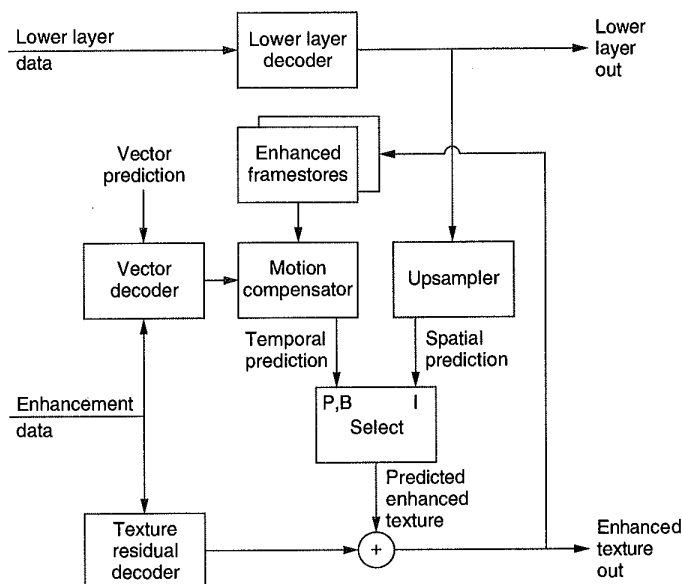


Figure 5.87 Spatial scalability system requires the enhancement layer to improve the vector density in order to minimize the residual bit rate.

The output of the base level decoder is considered to be a spatial prediction by the enhancement process. This may be shifted by vectors from the enhancement bitstream. Figure 5.88 shows that a base level macroblock of  $16 \times 16$  pixels predicted with one vector is upsampled to four macroblocks, each of which can be shifted with its own enhancement

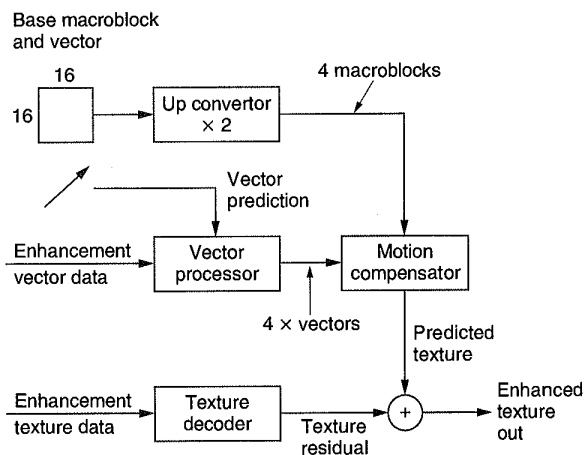
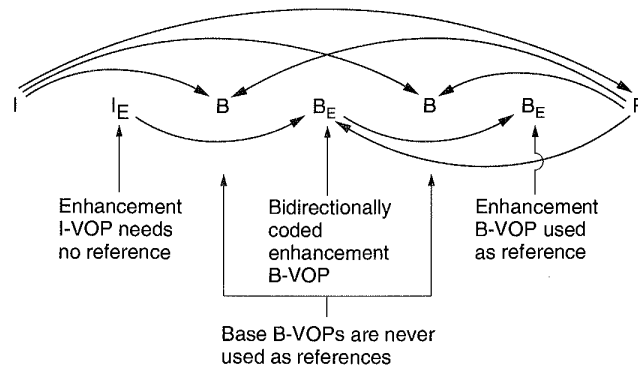


Figure 5.88 Each base level macroblock corresponds to four enhancement macroblocks. Base vectors are used to predict enhancement vectors.




**Figure 5.89** Temporal scalability allows new VOPs to be created between base layer VOPs thereby increasing the picture rate.

vector. Upsampled base layer pixels which have been shifted by enhancement vectors form a new high-resolution prediction. The encoder transmits the spatial residual from this prediction.

Temporally scalable decoding allows new VOPs to be created in between the VOPs of the base layer to increase the frame rate. Figure 5.89 shows that enhancement I-VOPs are intra-coded and so are decoded without reference to any other VOP just like a base I-VOP. P- and B-VOPs in the enhancement layer are motion predicted from other VOPs. These may be in the base layer or the enhancement layer according to prediction reference bits in the enhancement data. Base layer B-VOPs are never used as enhancement references, but enhancement B-VOPs can be references for other enhancement VOPs.

Figure 5.90 shows the example of an enhancement B-VOP placed between two VOPs in the base layer. Enhancement vectors are used to



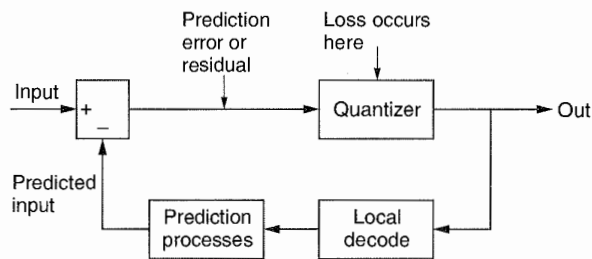
**Figure 5.90** This enhancement B-VOP contains vectors to shift base level anchor VOPs. Residuals are then added to correct the prediction error.

shift base layer pixels to new locations in the enhancement VOP. The enhancement residual is then added to these predictions to create the output VOP.

### 5.31 Coding artifacts

This section describes the visible results of imperfect coding. This may be where the coding algorithm is sub-optimal, where the coder latency is too short or where the compression factor in use is simply too great for the material.

Figure 5.91 shows that all MPEG coders can be simplified to the same basic model in which the predictive power of the decoder is modelled in the encoder. The current picture/object is predicted as well as possible, and the prediction error or residual is transmitted. As the decoder can make the same prediction, if the residual is correctly transmitted, the decoded output will be lossless.



**Figure 5.91** All MPEG coders can be reduced to a lossless prediction process and a lossy coding of the prediction error or residual.

The MPEG coding family contains a large number of compression tools and many of these are lossless. Lossless coding tools include the motion compensation, zig-zag scanning and run-length coding of MPEG-1 and MPEG-2 and the context coding and various types of predictive coding used in MPEG-4. It is a simple fact that the more sophisticated the prediction, the fewer data will be needed to transmit the residual.

However, in many, perhaps most, applications of MPEG, the bit rate is not decided by the encoder, but by some other decision process which may be economically biased. All MPEG coders contain an output buffer, and if the allocated bit rate is not adequate, this will start to fill up. To prevent data loss and decoder crashes, the encoder must reduce its bit rate.

It cannot make any economy in the lossless coding. Reducing the losslessly coded information would damage the prediction mechanism

and *increase* the residual. Reducing the accuracy of vectors, for example, is counterproductive. Consequently all that can be done is to reduce the bit rate of the residual. This must be done where the DCT coefficients are quantized because this is the only variable mechanism at the encoder's disposal. When in an unsympathetic environment, the coder has to reduce the bit rate the best it can.

The result is that the output pictures are visually different from the input pictures. The difference between the two is classified as *coding noise*. This author has great difficulty with this term, because true noise should be decorrelated from the message whereas the errors of MPEG coders are signal dependent and the term *coding distortion* is more appropriate. This is supported by the fact that the subjective visual impact of coding errors is far greater than would be caused by the same level of true noise. Put another way, quantizing distortion is irritating and pretending that it is noise gives no assistance in determining how bad it will look.

In motion-compensated systems such as MPEG, the use of periodic intra-fields means that the coding error may vary from picture to picture and this may be visible as *noise pumping*. The designer of the coder is in a difficult position as the user reduces the bit rate. If the *I* data are excessive, the *P* and *B* data will have to be heavily quantized resulting in errors. However, if the data representing the *I* pictures/objects is reduced too much, *P* and *B* pictures may look better than the *I* pictures. Consequently it is necessary to balance the requantizing in the *I*, *P* and *B* pictures/objects so that the level of visible artifacts in each remains roughly the same in each.

Noise pumping may also be visible where the amount of motion changes. If a pan is observed, as the pan speed increases the motion vectors may become less accurate and reduce the quality of the prediction processes. The prediction errors will get larger and will have to be more coarsely quantized. Thus the picture gets noisier as the pan accelerates and the noise reduces as the pan slows down. The same result may be apparent at the edges of a picture during zooming. The problem is worse if the picture contains fine detail. Panning on grass or trees waving in the wind taxes most coders severely. Camera shake from a hand-held camera also increases the motion vector data and results in more noise, as does film weave.

Input video noise or film grain degrades inter-coding as there is less redundancy between pictures and the difference data become larger, requiring coarse quantizing and adding to the existing noise.

Where a codec is really fighting the quantizing may become very coarse and as a result the video level at the edge of one DCT block may not match that of its neighbour. As a result the DCT block structure becomes visible as a mosaicing or tiling effect. MPEG-4 introduces some decoding




techniques to filter out blocking effects and in principle these could be applied to MPEG-1 and MPEG-2 decoders.

Coarse quantizing also causes some coefficients to be rounded up and appear larger than they should be. High-frequency coefficients may be eliminated by heavy quantizing and this forces the DCT to act as a steep-cut low-pass filter. This causes fringeing or ringing around sharp edges and extra shadowy edges which were not in the original and is most noticeable on text.

Excess compression may also result in colour bleed where fringeing has taken place in the chroma or where high-frequency chroma coefficients have been discarded. Graduated colour areas may reveal banding or posterizing as the colour range is restricted by requantizing. These artifacts are almost impossible to measure with conventional test gear.

Neither noise pumping nor blocking are visible on analog video recorders and so it is nonsense to liken the performance of a codec to the quality of a VCR. In fact noise pumping is extremely objectionable because, unlike steady noise, it attracts attention in peripheral vision and may result in viewing fatigue.

In addition to highly detailed pictures with complex motion, certain types of video signal are difficult for MPEG to handle and will usually result in a higher level of artifacts than usual. Noise has already been mentioned as a source of problems. Timebase error from, for example, VCRs is undesirable because this puts successive lines in different horizontal positions. A straight vertical line becomes jagged and this results in high spatial frequencies in the DCT process. Spurious coefficients are created which need to be coded.



**Figure 5.92** In composite video the subcarrier frequency is arranged so that inversions occur between adjacent lines and pictures to help reduce the visibility of the chroma.

Much archive video is in composite form and MPEG can only handle this after it has been decoded to components. Unfortunately many general-purpose composite decoders have a high level of residual subcarrier in the outputs. This is normally not a problem because the subcarrier is designed to be invisible to the naked eye. Figure 5.92 shows that in PAL and NTSC the subcarrier frequency is selected so that a phase reversal is achieved between successive lines and frames.

Whilst this makes the subcarrier invisible to the eye, it is not invisible to an MPEG decoder. The subcarrier waveform is interpreted as a horizontal frequency, the vertical phase reversals are interpreted as a vertical spatial frequency and the picture-to-picture reversals increase the magnitude of the prediction errors. The subcarrier level may be low but it can be present over the whole screen and require an excess of coefficients to describe it.

Composite video should not in general be used as a source for MPEG encoding, but where this is inevitable the standard of the decoder must be much higher than average, especially in the residual subcarrier specification. Some MPEG preprocessors support high-grade composite decoding options.

Judder from conventional linear standards convertors degrades the performance of MPEG. The optic flow axis is corrupted and linear filtering causes multiple images which confuse motion estimators and result in larger prediction errors. If standards conversion is necessary, the MPEG system must be used to encode the signal in its original format and the standards convertor should be installed after the decoder. If a standards convertor has to be used before the encoder, then it must be a type which has effective motion compensation.

Film weave causes movement of one picture with respect to the next and this results in more vector activity and larger prediction errors. Movement of the centre of the film frame along the optical axis causes magnification changes which also result in excess prediction error data. Film grain has the same effect as noise: it is random and so cannot be compressed.

Perhaps because it is relatively uncommon, MPEG-2 cannot handle image rotation well because the motion-compensation system is only designed for translational motion. Where a rotating object is highly detailed, such as in certain fairground rides, the motion compensation failure requires a significant amount of prediction error data and if a suitable bit rate is not available the level of artifacts will rise. The additional coding tools of MPEG-4 allow rotation be handled more effectively. Meshes can be rotated by repositioning their vertices with a few vectors and this improves the prediction of rotations dramatically.

Flash guns used by still photographers are a serious hazard to MPEG especially when long GOPs are used. At a press conference where a series of flashes may occur, the resultant video contains intermittent white

frames. These are easy enough to code, but a huge prediction error is required to return the white frame to the next picture. The output buffer fills and heavy requantizing is employed. After a few flashes the picture has generally gone to tiles.

### 5.32 MPEG and concatenation

Concatenation loss occurs when the losses introduced by one codec are compounded by a second codec. All practical compressors, MPEG included, are lossy because what comes out of the decoder is not bit-identical to what went into the encoder. The bit differences are controlled so that they have minimum visibility to a human viewer.

MPEG is a toolbox which allows a variety of manipulations to be performed both in the spatial and the temporal domain. MPEG-2 has more tools than MPEG-1 and MPEG-4 has more still. There is a limit to the compression which can be used on a single picture/object, and if higher compression factors are needed, temporal coding will have to be used. The longer the run of time considered, the lower the bit rate needed, but the harder it becomes to edit.

The most editable form of MPEG is to use *I* data only. As there is no temporal coding, pure cut edits can be made between pictures. The next best thing is to use a repeating *IB* structure which is locked to the odd/even field structure. Cut edits cannot be made as the *B* pictures are bidirectionally coded and need data from both adjacent *I* pictures for decoding. The *B* picture has to be decoded prior to the edit and re-encoded after the edit. This will cause a small concatenation loss.

Beyond the *IB* structure processing gets harder. If a long GOP is used for the best compression factor, an *IBBPBBP* . . . structure results. Editing this is very difficult because the pictures are sent out of order so that bidirectional decoding can be used. MPEG allows closed GOPs where the last *B* picture is coded wholly from the previous pictures and does not need the *I* picture in the next GOP. The bitstream can be switched at this point but only if the GOP structures in the two source video signals are synchronized (makes colour framing seem easy). Consequently in practice a long GOP bitstream will need to be decoded prior to any production step. Afterwards it will need to be re-encoded.

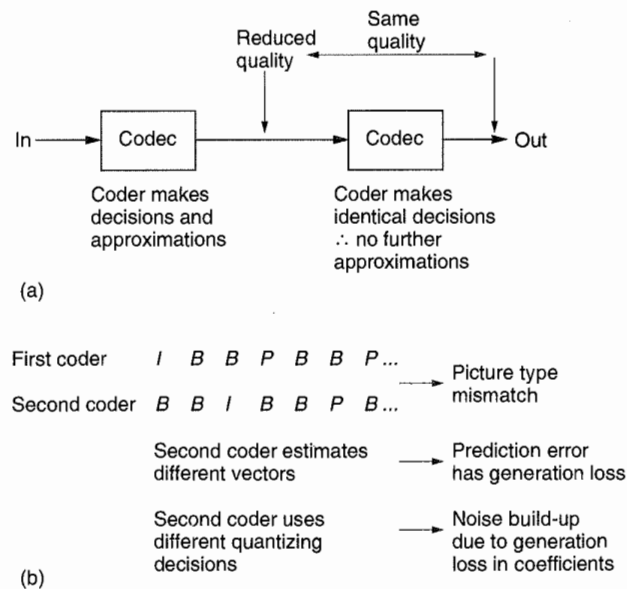
This is known as *naive* concatenation and an enormous pitfall awaits. Unless the GOP structure of the output is identical to and synchronized with the input the results will be disappointing. The worst case is where an *I* picture is encoded from a picture which was formerly a *B* picture. It is easy enough to lock the GOP structure of a coder to a single input, but if an edit is made between two inputs, the GOP timings could well be different.

As there are so many structures allowed in MPEG, there will be a need to convert between them. If this has to be done, it should only be in the direction which increases the GOP length and reduces the bit rate. Going the other way is inadvisable. The ideal way of converting from, say, the *IB* structure of a news system to the *IBBP* structure of an emission system is to use a recompressor. This is a kind of standards convertor which will give better results than a decode followed by an encode.

The DCT part of MPEG itself is lossless. If all the coefficients are preserved intact an inverse transform yields the same pixel data. Unfortunately this does not yield enough compression for many applications. In practice the coefficients are made less accurate by removing bits starting at the least significant end and working upwards. This process is weighted, or made progressively more aggressive as spatial frequency increases.

Small-value coefficients may be truncated to zero and large-value coefficients are most coarsely truncated at high spatial frequencies where the effect is least visible.

Figure 5.93(a) shows what happens in the ideal case where two *identical* coders are put in tandem and synchronized. The first coder quantizes the coefficients to finite accuracy and causes a loss on decoding. However, when the second coder performs the DCT calculation, the coefficients obtained will be identical to the quantized coefficients in the first coder



**Figure 5.93** (a) Two identical coders in tandem which are synchronized make similar coding decisions and cause little loss. (b) There are various ways in which concatenated coders can produce non-ideal performance.

and so if the second weighting and requantizing step is identical the same truncated coefficient data will result and there will be no further loss of quality.<sup>7</sup>

In practice this ideal situation is elusive. If the two DCTs become non-identical for any reason, the second requantizing step will introduce further error in the coefficients and the artifact level goes up. Figure 5.93(b) shows that non-identical concatenation can result from a large number of real-world effects.

An intermediate processing step such as a fade will change the pixel values and thereby the coefficients. A DVE resize or shift will move pixels from one DCT block to another. Even if there is no processing step, this effect will also occur if the two codecs disagree on where the MPEG picture boundaries are within the picture. If the boundaries are correct there will still be concatenation loss if the two codecs use different weighting.

One problem with MPEG is that the compressor design is unspecified. Whilst this has advantages, it does mean that the chances of finding identical coders is minute because each manufacturer will have his own views on the best compression algorithm. In a large system it may be worth obtaining the coders from a single supplier.

It is now increasingly accepted that concatenation of compression techniques is potentially damaging, and results are worse if the codecs are different. Clearly feeding a digital coder such as MPEG-2 with a signal which has been subject to analog compression comes into the category of worse. Using interlaced video as a source for MPEG coding is sub-optimal and using decoded composite video is even worse.

One way of avoiding concatenation is to stay in the compressed data domain. If the goal is just to move pictures from one place to another, decoding to traditional video so that an existing router can be used is not ideal, although it is substantially better than going through the analog domain.

Figure 5.94 shows some possibilities for picture transport. Clearly if the pictures exist as a compressed file on a server, a file transfer is the right way to do it as there is no possibility of loss because there has been no concatenation. File transfer is also quite indifferent to the picture format. It doesn't care about the frame rate, whether the pictures are interlaced or not or whether the colour is 4:2:0 or 4:2:2

Decoding to SDI (serial digital interface) standard is sometimes done so that existing serial digital routing can be used. This is concatenation and has to be done carefully. The compressed video can only use interlace with non-square pixels and the colour coding has to be 4:2:2 because SDI only allows that. If a compressed file has 4:2:0 the chroma has to be interpolated up to 4:2:2 for SDI transfer and then subsampled back to 4:2:0 at the second coder and this will cause generation loss. An SDI

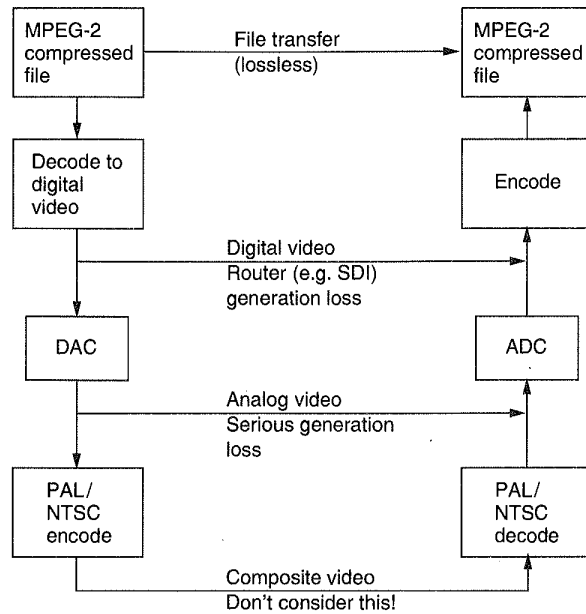



Figure 5.94 Compressed picture transport mechanisms contrasted.

transfer also can only be performed in real time, thus negating one of the advantages of compression. In short traditional SDI is not really at home with compression.

As 4:2:0 progressive scan gains popularity and video production moves steadily towards non-format-specific hardware using computers and data networks, use of the serial digital interface will eventually decline. In the short term, if an existing SDI router has to be used, one solution is to produce a bitstream which is sufficiently similar to SDI that a router will pass it. One example of this is known as SDTI. The signal level, frequency and impedance of SDTI is pure SDI, but the data protocol is different so that a bit accurate file transfer can be performed. This has two advantages over SDI. First, the compressed data format can be anything appropriate and non-interlaced and/or 4:2:0 can be handled in any picture size, aspect ratio or frame rate. Second, a faster than real-time transfer can be used depending on the compression factor of the file.

Equipment which allows this is becoming available and its use can mean that the full economic life of a SDI routing installation can be obtained.

An improved way of reducing concatenation loss has emerged from the ATLANTIC research project.<sup>8</sup> Figure 5.95 shows that the second encoder in a concatenated scheme does not make its own decisions from the incoming video, but is instead steered by information from the first



**Figure 5.95** In an ATLANTIC system, the second encoder is steered by information from the decoder.

bitstream. As the second encoder has less intelligence, it is known as a *dim* encoder.

The information bus carries all the structure of the original MPEG-2 bitstream which would be lost in a conventional decoder. The ATLANTIC decoder does more than decode the pictures. It also places on the information bus all parameters needed to make the dim encoder re-enact what the initial MPEG-2 encode did as closely as possible.

The GOP structure is passed on so that pictures are re-encoded as the same type. Positions of macroblock boundaries become identical so that DCT blocks contain the same pixels and motion vectors relate to the same screen data. The weighting and quantizing tables are passed so that coefficient truncation is identical. Motion vectors from the original bitstream are passed on so that the dim encoder does not need to perform motion estimation. In this way predicted pictures will be identical to the original prediction and the prediction error data will be the same.

One application of this approach is in recompression, where an MPEG-2 bitstream has to have its bit rate reduced. This has to be done by heavier requantizing of coefficients, but if as many other parameters as possible can be kept the same, such as motion vectors, the degradation will be minimized. In a simple recompressor just requantizing the coefficients means that the predictive coding will be impaired. In a proper encode, the quantizing error due to coding say an *I* picture is removed from the *P* picture by the prediction process. The prediction error of *P* is obtained by subtracting the decoded *I* picture rather than the original *I* picture.

In simple recompression this does not happen and there may be a tolerance buildup known as drift.<sup>9</sup> A more sophisticated recompressor will need to repeat the prediction process using the decoded output pictures as the prediction reference.

MPEG-2 bitstreams will often be decoded for the purpose of switching. Local insertion of commercial breaks into a centrally originated bitstream is one obvious requirement. If the decoded video signal is switched, the information bus must also be switched. At the switch point identical re-encoding becomes impossible because prior pictures required for predictive coding will have disappeared. At this point the dim encoder

has to become bright again because it has to create an MPEG-2 bitstream without assistance.

It is possible to encode the information bus into a form which allows it to be invisibly carried in the serial digital interface. Where a production process such as a vision mixer or DVE performs no manipulation, i.e. becomes bit transparent, the subsequent encoder can extract the information bus and operate in 'dim' mode. Where a manipulation is performed, the information bus signal will be corrupted and the encoder has to work in 'bright' mode. The encoded information signal is known as a 'mole'<sup>10</sup> because it burrows through the processing equipment!

There will be a generation loss at the switch point because the re-encode will be making different decisions in bright mode. This may be




Figure 5.96 (a) A bitstream switch at a different level of buffer occupancy can cause a decoder to overflow. (b) Recompression after a switch to return to correct buffer occupancy.



difficult to detect because the human visual system is slow to react to a vision cut and defects in the first few pictures after a cut are masked.

In addition to the video computation required to perform a cut, the process has to consider the buffer occupancy of the decoder. A downstream decoder has finite buffer memory, and individual encoders model the decoder buffer occupancy to ensure that it neither overflows nor underflows. At any instant the decoder buffer can be nearly full or nearly empty without a problem provided there is a subsequent correction. An encoder which is approaching a complex *I* picture may run down the buffer so it can send a lot of data to describe that picture. Figure 5.96(a) shows that if a decoder with a nearly full buffer is suddenly switched to an encoder which has been running down its buffer occupancy, the decoder buffer will overflow when the second encoder sends a lot of data.

An MPEG-2 switcher will need to monitor the buffer occupancy of its own output to avoid overflow of downstream decoders. Where this is a possibility the second encoder will have to recompress to reduce the output bit rate temporarily. In practice there will be a recovery period where the buffer occupancy of the newly selected signal is matched to that of the previous signal. This is shown in Figure 5.96(b).

### References

1. Kelly, D.H., Visual processing of moving stimuli. *J. Opt. Soc. America*, **2** 216–225 (1985)
2. Uyttendaele, A., Observations on scanning formats. Presented at HDTV '97 Montreux (June 1997)
3. Shapiro, J.M., Embedded image coding using zerotrees of wavelet coefficients. *IEEE Trans. SP*, **41** 3445–3462 (1993)
4. Taubin, G. and Rossignac, J., Geometric compression through topological surgery. *ACM Trans. on Graphics*, 88–115 (April 1998)
5. de Boor, C., *A Practical Guide to Splines*, Berlin: Springer (1978)
6. Taubin, G., Guezic, A., Horn, W.P. and Lazarus, F., *Progressive forest split compression*, *Proc. Siggraph '98*, 123–132 (1998)
7. Stone, J. and Wilkinson, J., Concatenation of video compression systems. Presented at 137th SMPTE Tech. Conf., New Orleans (1995)
8. Wells, N.D., The ATLANTIC project: Models for programme production and distribution. *Proc. Euro. Conf. Multimedia Applications Services and Techniques (ECMAST)*, 243–253 (1996)
9. Werner, O., Drift analysis and drift reduction for multiresolution hybrid video coding. *Image Communication*, **8**, 387–409 (1996)
10. Knee, M.J. and Wells, N.D., Seamless concatenation – a 21st century dream. Presented at Int. Television. Symp., Montreux (1997)

# 6

## Program and transport streams

---

After compression, audio and video signals are simply data and only differ from generic data in that they relate to real-time signals. MPEG bitstreams are basically means of transporting data whilst allowing the time axis of the original signal to be re-created.


### 6.1 Introduction

There are two basic applications of MPEG bitstreams – recording and transmission – and these have quite different requirements. In the multichannel recording application the encoders and decoders can all share the same clock. The transport speed can be altered so that the incoming bit rate can be slaved to a reference generated by the decoder itself.

In the case of multichannel transmission the program sources are not necessarily synchronous. The source of timing is the encoder and the decoder must synchronize or genlock to that. Thus the difference between a program stream and a transport stream is that the latter must contain additional synchronizing information to lock the encoder and decoder clocks together independently in each program.

### 6.2 Packets and time stamps


The video elementary stream is an endless bitstream representing pictures which are not necessarily in the correct order and which take a variable length of time to transmit. Storage and transmission systems prefer discrete blocks of data and so elementary streams are packetized to form a PES (packetized elementary stream). Audio elementary streams



**Figure 6.1** A PES packet structure is used to break up the continuous elementary stream.

are also packetized. A packet is shown in Figure 6.1. It begins with a header containing an unique packet start code and a code which identifies the type of data stream. Optionally the packet header also may contain one or more time stamps which are used for synchronizing the video decoder to real time and for obtaining lip-sync.

Figure 6.2 shows that a time stamp is a sample of the state of a counter which is driven by a 90 kHz clock. This is obtained by dividing down the master 27 MHz clock of MPEG-2. There are two types of time stamp: PTS and DTS. These are abbreviations for presentation time stamp and decode time stamp. A presentation time stamp determines when the associated picture should be displayed on the screen, whereas a decode time stamp determines when it should be decoded. In bidirectional coding these times can be quite different.



**Figure 6.2** Time stamps are the result of sampling a counter driven by the encoder clock.

Audio packets only have presentation time stamps. Clearly if lip-sync is to be obtained, the audio and the video streams of a given program must have been locked to the same master 27 MHz clock and the time stamps must have come from the same counter driven by that clock.

With reference to Figure 6.3, the GOP begins with an *I* picture, and then *P1* is sent out of sequence prior to the *B* pictures. *P1* has to be decoded before *B1* and *B2* can be decoded. As only one picture can be decoded at a time, the *I* picture is decoded at time *N*, but not displayed until time *N*




Figure 6.3 An example of using PTS/DTS to synchronize bidirectional decoding.

+ 1. As the *I* picture is being displayed, *P*<sub>1</sub> is being decoded at *N* + 1. *P*<sub>1</sub> will be stored in RAM. At time *N* + 2, *B*<sub>1</sub> is decoded and displayed immediately. For this reason *B* pictures need only PTS. At *N* + 3, *B*<sub>2</sub> is decoded and displayed. At *N* + 4, *P*<sub>1</sub> is displayed, hence the large difference between PTS and DTS in *P*<sub>1</sub>. Simultaneously *P*<sub>2</sub> is decoded and stored ready for the decoding of *B*<sub>3</sub> and so on.

In practice the time between input pictures is constant and so there is a certain amount of redundancy in the time stamps. Consequently PTS/DTS need not appear in every PES packet. Time stamps can be up to 700 ms apart in program streams and up to 100 ms apart in transport streams. As each picture type (*I*, *P* or *B*) is flagged in the bitstream, the decoder can infer the PTS/DTS for every picture from the ones actually transmitted.

Figure 6.4 shows that one or more PES packets can be assembled in a pack whose header contains a sync pattern and a system clock reference code which allows the decoder to re-create the encoder clock. Clock reference operation is described in section 6.4.


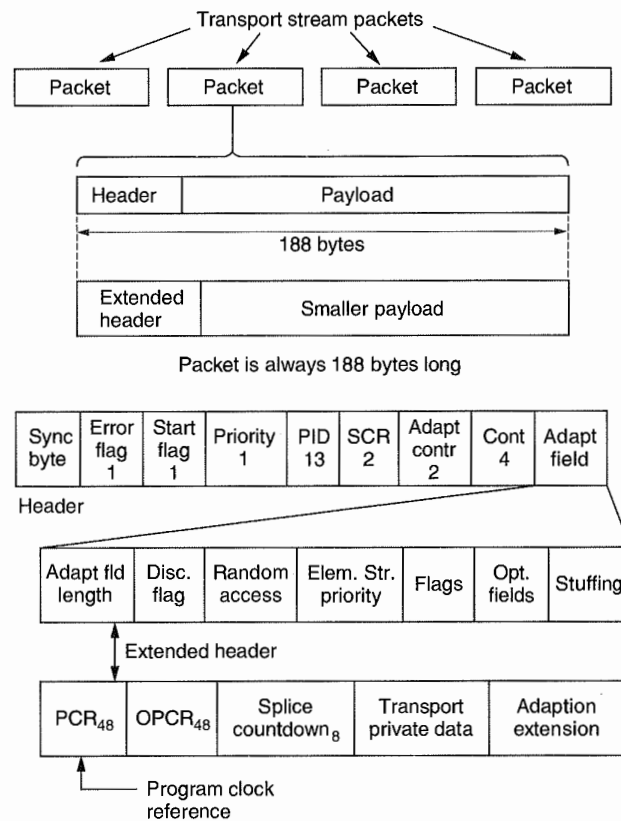


Figure 6.4 A pack is a set of PES packets. The pack header contains a clock reference code.

### 6.3 Transport streams

The MPEG-2 transport stream is intended to be a multiplex of many TV programs with their associated sound and data channels, although a single program transport stream (SPTS) is possible. The transport stream is based upon packets of constant size so that adding error-correction codes and interleaving<sup>1</sup> in a higher layer is eased. Figure 6.5 shows that these are always 188 bytes long. Transport stream packets should not be confused with PES packets which are larger and vary in size.



**Figure 6.5** Transport stream packets are always 188 bytes long to facilitate multiplexing and error correction.

Transport stream packets always begin with a header. The remainder of the packet carries data known as the payload. For efficiency, the normal header is relatively small, but for special purposes the header may be extended. In this case the payload gets smaller so that the overall size of the packet is unchanged.

The header begins with a sync byte which is an unique pattern detected by a demultiplexer. A transport stream may contain many different elementary streams and these are identified by giving each an unique 13-bit Packet Identification Code or PID which is included in the header. A multiplexer seeking a particular elementary stream simply checks the PID of every packet and accepts those which match, rejecting the rest.

In a multiplex there may be many packets from other programs in between packets of a given PID. To help the demultiplexer, the packet header contains a continuity count. This is a four-bit value which increments at each new packet having a given PID.

This approach allows statistical multiplexing as it does not matter how many or how few packets have a given PID; the demux will still find them. Statistical multiplexing has the problem that it is virtually impossible to make the sum of the input bit rates constant. Instead the multiplexer aims to make the average data bit rate slightly less than the maximum and the overall bit rate is kept constant by adding stuffing or null packets. These packets have no meaning, but simply keep the bit rate constant. Null packets always have a PID of 8191 (all ones) and the demultiplexer discards them.

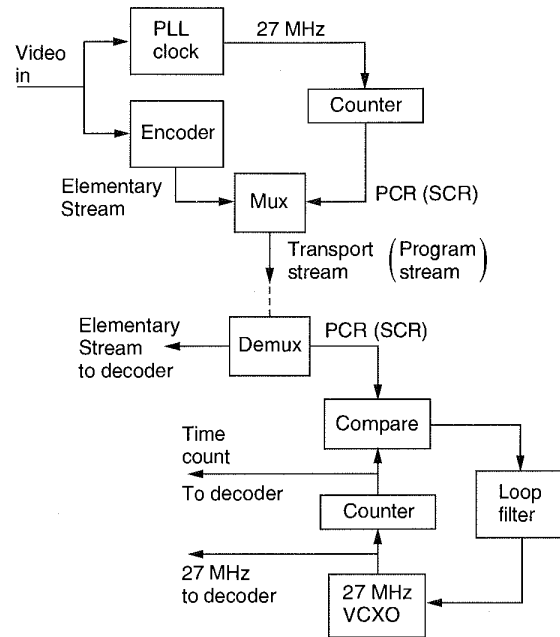
## 6.4 Clock references

A transport stream is a multiplex of several TV programs and these may have originated from widely different locations. It is impractical to expect all the programs in a transport stream to be synchronous and so the stream is designed from the outset to allow asynchronous programs. A decoder running from a transport stream has to genlock to the encoder and the transport stream has to have a mechanism to allow this to be done independently for each program. The synchronizing mechanism is called Program Clock Reference (PCR).

In program streams all the programs must be synchronous so that only one clock is required at the decoder. In this case the synchronizing mechanism is called System Clock Reference (SCR).

Figure 6.6 shows how the PCR/SCR system works. The goal is to recreate at the decoder a 27MHz clock which is synchronous with that at the encoder. The encoder clock drives a forty-eight-bit counter which continuously counts up to the maximum value before overflowing and beginning again.

A transport stream multiplexer will periodically sample the counter and place the state of the count in an extended packet header as a PCR (see Figure 6.5). The demultiplexer selects only the PIDs of the required program, and it will extract the PCRs from the packets in which they were inserted. In a program stream the count is placed in a pack header as an SCR which the decoder can identify.



**Figure 6.6** Program or System Clock Reference codes regenerate a clock at the decoder. See text for details.

The PCR/SCR codes are used to control a numerically locked loop (NLL). This is similar to a phase-locked loop, except that the two phases concerned are represented by the state of a binary number. The NLL contains a 27MHz VCXO (voltage controlled crystal oscillator), a variable-frequency oscillator based on a crystal which has a relatively small frequency range.

The VCXO drives a forty-eight-bit counter in the same way as in the encoder. The state of the counter is compared with the contents of the PCR/SCR and the difference is used to modify the VCXO frequency. When the loop reaches lock, the decoder counter would arrive at the same value as is contained in the PCR/SCR and no change in the VCXO would then occur. In practice the transport stream packets will suffer from transmission jitter and this will create phase noise in the loop. This is removed by the loop filter so that the VCXO effectively averages a large number of phase errors.

A heavily damped loop will reject jitter well, but will take a long time to lock. Lock-up time can be reduced when switching to a new program if the decoder counter is jammed to the value of the first PCR received in the new program. The loop filter may also have its time constants shortened during lock-up. Once a synchronous 27 MHz clock is available

at the decoder, this can be divided down to provide the 90 kHz clock which drives the time stamp mechanism. The entire timebase stability of the decoder is no better than the stability of the clock derived from PCR/SCR. MPEG-2 sets standards for the maximum amount of jitter which can be present in PCRs in a real transport stream.

## 6.5 Program Specific Information (PSI)

In a real transport stream, each elementary stream has a different PID, but the demultiplexer has to be told what these PIDs are and what audio belongs with what video before it can operate. This is the function of PSI. Figure 6.7 shows the structure of PSI. When a decoder powers up, it knows nothing about the incoming transport stream except that it must search for all packets with a PID of zero. PID zero is reserved for the Program Association Table (PAT). The PAT is transmitted at regular intervals and contains a list of all the programs in this transport stream. Each program is further described by its own Program Map Table (PMT) and the PIDs of the PMTs are contained in the PAT.

Figure 6.7 also shows that the PMTs fully describe each program. The PID of the video elementary stream is defined, along with the PID(s) of the associated audio and data streams. Consequently when the viewer selects a particular program, the demultiplexer looks up the program

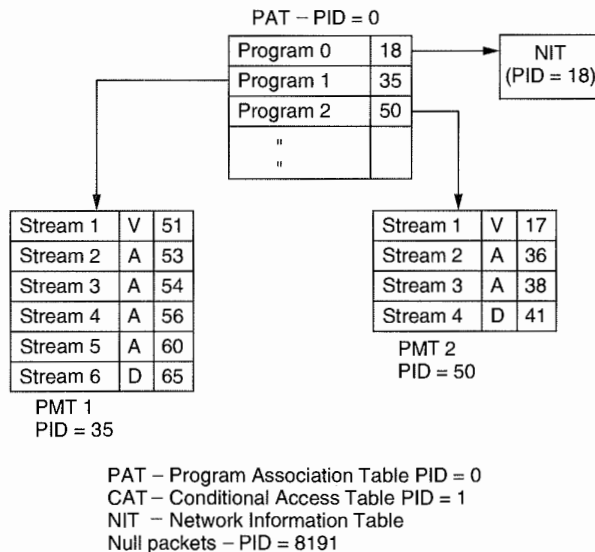


Figure 6.7 MPEG-2 Program Specific Information (PSI) is used to tell a demultiplexer what the transport stream contains.



number in the PAT, finds the right PMT and reads the audio, video and data PIDs. It then selects elementary streams having these PIDs from the transport stream and routes them to the decoders.


Program 0 of the PAT contains the PID of the Network Information Table (NIT). This contains information about what other transport streams are available. For example, in the case of a satellite broadcast, the NIT would detail the orbital position, the polarization, carrier frequency and modulation scheme. Using the NIT a set-top box could automatically switch between transport streams.

Apart from 0 and 8191, a PID of 1 is also reserved for the Conditional Access Table (CAT). This is part of the access control mechanism needed to support pay per view or subscription viewing.

## 6.6 Multiplexing

A transport stream multiplexer is a complex device because of the number of functions it must perform. A fixed multiplexer will be considered first. In a fixed multiplexer, the bit rate of each of the programs must be specified so that the sum does not exceed the payload bit rate of the transport stream. The payload bit rate is the overall bit rate less the packet headers and PSI rate.

In practice the programs will not be synchronous to one another, but the transport stream must produce a constant packet rate given by the bit rate divided by 188 bytes, the packet length. Figure 6.8 shows how this is handled. Each elementary stream entering the multiplexer passes through a buffer which is divided into payload-sized areas. Note that



**Figure 6.8** A transport stream multiplexer can handle several programs which are asynchronous to one another and to the transport stream clock. See text for details.

periodically the payload area is made smaller because of the requirement to insert PCR.

MPEG-2 decoders also have a quantity of buffer memory. The challenge to the multiplexer is to take packets from each program in such a way that neither its own buffers nor the buffers in any decoder either overflow or underflow. This requirement is met by sending packets from all programs as evenly as possible rather than bunching together a lot of packets from one program. When the bit rates of the programs are different, the only way this can be handled is to use the buffer contents indicators. The fuller a buffer is, the more likely it should be that a packet will be read from it. This a buffer content arbitrator can decide which program should have a packet allocated next.

If the sum of the input bit rates is correct, the buffers should all slowly empty because the overall input bit rate has to be less than the payload bit rate. This allows for the insertion of Program Specific Information. While PATs and PMTs are being transmitted, the program buffers will fill up again. The multiplexer can also fill the buffers by sending more PCRs as this reduces the payload of each packet. In the event that the multiplexer has sent enough of everything but still cannot fill a packet then it will send a null packet with a PID of 8191. Decoders will discard null packets and as they convey no useful data, the multiplexer buffers will all fill while null packets are being transmitted.

In a statistical multiplexer or statmux, the bit rate allocated to each program can vary dynamically. Figure 6.9 shows that there must be tight connection between the statmux and the associated compressors. Each compressor has a buffer memory which is emptied by a demand clock from the statmux. In a normal, fixed bit rate, coder the buffer contents feeds back and controls the requantizer. In statmuxing this process is less severe and only takes place if the buffer is very close to full, because the buffer content is also fed to the statmux.

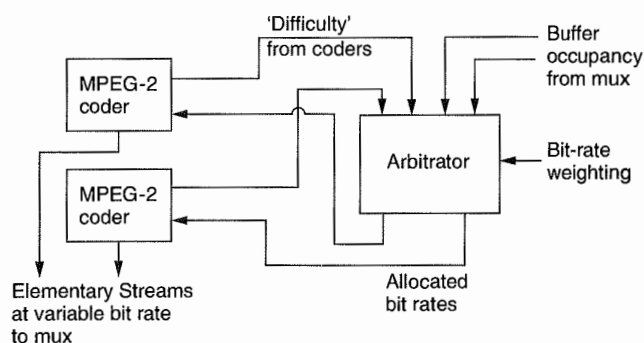


Figure 6.9 A statistical multiplexer contains an arbitrator which allocates bit rate to each program as a function of program difficulty.

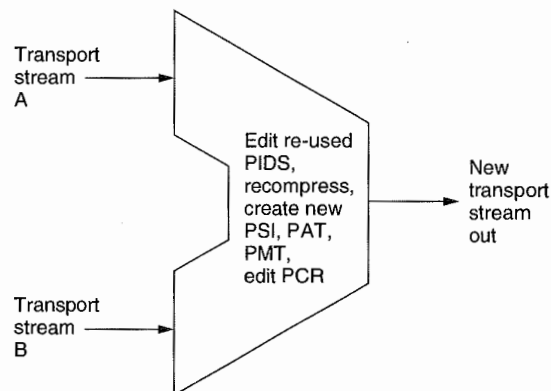
The statmux contains an arbitrator which allocates packets to the program with the fullest buffer. Thus if a particular program encounters difficult material it will produce large prediction errors and begin to fill its output buffer. This will cause the statmux to allocate more packets to that program. In order to fill more packets, the statmux clocks more data out of that buffer, causing the level to fall again. Of course, this is only possible if the other programs in the transport stream are handling typical video.

In the event that several programs encounter difficult material at once, clearly the buffer contents will rise and the requantizing mechanism will have to operate.

## 6.7 Remultiplexing

In real life a program creator may produce a transport stream which carries all its programs simultaneously. A service provider may take in several such streams and create its own transport stream by selecting different programs from different sources. In an MPEG-2 environment this requires a remultiplexer, also known as a transmultiplexer. Figure 6.10 shows what a remultiplexer does.

Remultiplexing is easier when all the incoming programs have the same bit rate. If a suitable combination of programs is selected it is obvious that the output transport stream will always have sufficient bit rate. Where statistical multiplexing has been used, there is a possibility that the sum of the bit rates of the selected programs will exceed the bit rate of the output transport stream. To avoid this, the remultiplexer will have to employ recompression.



**Figure 6.10** A remultiplexer creates a new transport stream from selected programs in other transport streams.

Recompression requires a partial decode of the bitstream to identify the DCT coefficients. These will then be requantized to reduce the bit rate until it is low enough to fit the output transport stream.

Remultiplexers have to edit the Program Specific Information (PSI) such that the Program Association Table (PAT) and the Program Map Tables (PMT) correctly reflect the new transport stream content. It may also be necessary to change the packet identification codes (PIDs) since the incoming transport streams could have inadvertently used the same values.

When Program Clock Reference (PCR) data are included in an extended packet header, they represent a real-time clock count and if the associated packet is moved in time the PCR value will be wrong. Remultiplexers have to recreate a new multiplex from a number of other multiplexes and it is inevitable that this process will result in packets being placed in different locations in the output transport stream than they had in the input. In this case the remultiplexer must edit the PCR values so that they reflect the value the clock counter would have had at the location the packet at which the packet now resides.

#### Reference

1. Watkinson, J.R., *The Art of Digital Video*, second edition, Ch. 6, Oxford: Focal Press (1994)

# 7

## MPEG applications

---

### 7.1 Introduction

It should be borne in mind that however exciting it may be, MPEG is only a technology. Technology *per se* is not useful and only becomes useful when it is incorporated into an appropriate and affordable product or service. If the use is not appropriate or affordable, there is no point in blaming the technology.

In this chapter a number of useful applications of MPEG are outlined. In each case designers have followed the rules of high technology to design systems which maximize the potential of the technology whilst protecting against its weaknesses.

MPEG is an information technology. It cannot deliver anything physical: you get the movie but no popcorn. The strength of MPEG is that it delivers audio-visual material with economy of bandwidth. Its weaknesses are sensitivity to bit errors and the requirement for near-real-time transmission. The variable-length coding of MPEG can lose synchronization in the presence of a single bit error. Decoders can take a considerable time to recover from a buffer underflow or overflow.

Wherever bandwidth or storage capacity is in short supply, MPEG is useful. Where bandwidth is plentiful, as in optical fibre networks, MPEG is pointless. As the capacity and economics of storage devices continue to increase, MPEG becomes less important.

MPEG will remain important in any area where there will be a permanent shortage of bandwidth; this includes any mobile equipment or service which must use radio communication.

Figure 7.1 shows that well-engineered applications of MPEG will always include an error-correction system appropriate for the characteristics of the channel and suitable mechanisms to ensure that decoder

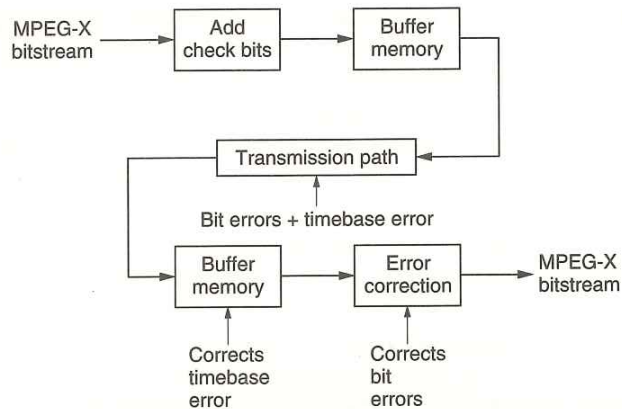


Figure 7.1 Any real-time MPEG delivery system must include means to maintain a low error rate and buffering to deliver data within the decoder's time window.

buffers neither overflow nor underflow. Constant-bit rate channels having constant delay are easy to implement because the buffering mechanisms of the standard MPEG encoders and decoders are able to operate with no further assistance. Digital TV transmitters provide essentially constant bit-rate because this is related to the standardized (and regulated) channel bandwidth allowed.

In packet networks the delivery delay is not constant because other types of messages are multiplexed with the audio-visual message of interest. Variable delay channels cause MPEG decoders difficulty. Not only is there a greater danger of buffer overflow or underflow, but the program clock reference information is no longer correct if a packet is




Figure 7.2 There are a large number of conflicting demands in the design of a compression system. Here the effect of a number of choices is shown to alter the balance between economy and bandwidth.



received with a time shift. In this case additional buffering is needed at both ends of the channel so that the variable delay characteristic of the channel is invisible to the MPEG layer.

There are now many members of the MPEG family and it can be difficult to decide which type of MPEG coding to use. Figure 7.2 shows some general guidelines. As a decoder becomes more complex, the bit rate which can be used for a given quality will fall. However, this will be accompanied by an increase in power consumption at the encoder and decoder as well as an increased delay. The designer will have to balance the cost of the bandwidth with the cost of the encoding and decoding equipment.

## 7.2 Video phones

The telephone is a classic example of limited bandwidth. If there is a market for video phones, then compression is the only way forward. The facial animation coding developed for MPEG-4 is an obvious contender as it allows a reasonable reproduction of the caller's face with a data rate of only a few kilobits per second. This is considerably less than the data rate needed for the speech. Video telephony is relatively easy between desk-based equipment, where the camera can be mounted on top of the display, but it is not clear how this can be implemented in a cellular telephone which is held against the user's ear. Another important issue for cellular telephones is the amount of signal processing needed to decode and display facial animation data. Essentially the decoder is a real-time rendering engine and this is only practicable in a portable product if the current consumption is within the range of available battery technology.

## 7.3 Digital television broadcasting

Digital television broadcasting relies on the combination of a number of fundamental technologies. These are: MPEG-2 compression to reduce the bit rate, multiplexing to combine picture and sound data into a common bitstream, digital modulation schemes to reduce the RF bandwidth needed by a given bit rate and error correction to reduce the error statistics of the channel down to a value acceptable to MPEG data.

MPEG is a compression and multiplexing standard and does not specify how error correction should be performed. Consequently a transmission standard must define a system which has to correct essentially all errors such that the delivery mechanism is transparent.

DVB (digital video broadcasting)<sup>1</sup> is a standard which incorporates MPEG-2 picture and sound coding but which also specifies all the additional steps needed to deliver an MPEG transport stream from one place to another. This transport stream will consist of a number of elementary streams of video and audio, where the audio may be coded according to MPEG audio standard or AC-3. In a system working within its capabilities, the picture and sound quality will be determined only by the performance of the compression system and not by the RF transmission channel. This is the fundamental difference between analog and digital broadcasting. In analog television broadcasting, the picture quality may be limited by composite video encoding artifacts as well as transmission artifacts such as noise and ghosting. In digital television broadcasting the picture quality is determined instead by the compression artifacts and interlace artifacts if interlace has been retained.

If the received error rate increases for any reason, once the correcting power is used up, the system will degrade rapidly as uncorrected errors enter the MPEG decoder. In practice decoders will be programmed to recognize the condition and blank or mute to avoid outputting garbage. As a result digital receivers tend either to work well or not at all.

It is important to realize that the signal strength in a digital system does not translate directly to picture quality. A poor signal will increase the number of bit errors. Provided that this is within the capability of the error-correction system, there is no visible loss of quality. In contrast, a very powerful signal may be unusable because of similarly powerful reflections due to multipath propagation.

Whilst in one sense an MPEG transport stream is only data, it differs from generic data in that it must be presented to the viewer at a particular rate. Generic data are usually asynchronous, whereas baseband video and audio are synchronous. However, after compression and multiplexing audio and video are no longer precisely synchronous and so the term *isochronous* is used. This means a signal which was at one time synchronous and will be displayed synchronously, but which uses buffering at transmitter and receiver to accommodate moderate timing errors in the transmission.

Clearly another mechanism is needed so that the time axis of the original signal can be re-created on reception. The time stamp and program clock reference system of MPEG does this.

Figure 7.3 shows that the concepts involved in digital television broadcasting exist at various levels which have an independence not found in analog technology. In a given configuration a transmitter can radiate a given payload data bit rate. This represents the useful bit rate and does not include the necessary overheads needed by error correction, multiplexing or synchronizing. It is fundamental that the transmission system does not care what this payload bit rate is used for. The entire



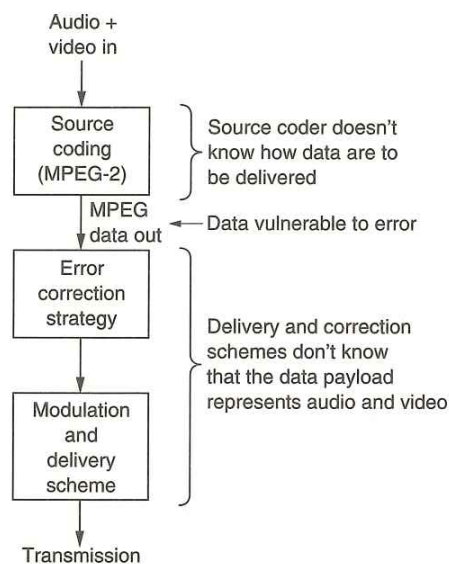


Figure 7.3 Source coder doesn't know delivery mechanism and delivery doesn't need to know what the data mean.

capacity may be used up by one high-definition channel, or a large number of heavily compressed channels may be carried. The details of this data usage are the domain of the *transport stream*. The multiplexing of transport streams is defined by the MPEG standards, but these do not define any error correction or transmission technique.

At the lowest level in Figure 7.3 is the source coding scheme, in this case MPEG-2 compression results in one or more elementary streams, each of which carries a video or audio channel. Figure 7.4 shows that elementary streams are multiplexed into a transport stream. The viewer then selects the desired elementary stream from the transport stream. Metadata in the transport stream ensure that when a video elementary stream is chosen, the appropriate audio elementary stream will automatically be selected.

A key difference between analog and digital transmission is that the transmitter output is switched between a number of discrete states rather than continuously varying. The process is called channel coding which is the digital equivalent of modulation.<sup>2</sup> A good code minimizes the channel bandwidth needed for a given bit rate. This quality of the code is measured in bits/s/Hz.

Where the SNR is poor, as in satellite broadcasting, the amplitude of the signal will be unstable, and phase modulation is used. Figure 7.5 shows that phase-shift keying (PSK) can use two or more phases. When four

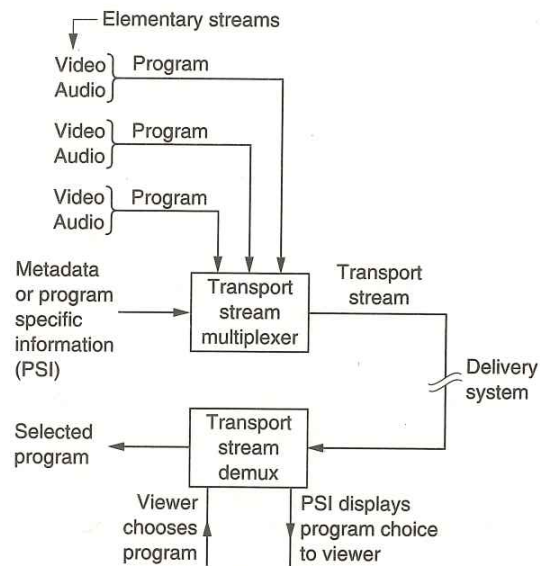


Figure 7.4 Program Specific Information helps the demultiplexer to select the required program.




Figure 7.5 Differential quadrature phase shift keying (DQPSK).

phases in quadrature are used, the result is quadrature phase shift keying or QPSK. Each period of the transmitted waveform can have one of four phases and therefore conveys the value of two data bits. 8-PSK uses eight phases and can carry three bits per symbol where the SNR is adequate. PSK is generally encoded in such a way that a knowledge of absolute phase is not needed at the receiver. Instead of encoding the signal phase directly, the data determine the magnitude of the phase shift between symbols.

In terrestrial transmission more power is available than, for example, from a satellite and so a stronger signal can be delivered to the receiver. Where a better SNR exists, an increase in data rate can be had using




Figure 7.6 In 8-VSB the transmitter operates in eight different states enabling three bits to be sent per symbol.

multi-level signalling instead of binary. Figure 7.6 shows that the ATSC system uses an eight-level signal (8-VSB) allowing three bits to be sent per symbol. Four of the levels exist with normal carrier phase and four exist with inverted phase so that a phase-sensitive rectifier is needed in the receiver. Clearly the data separator must have a three-bit ADC which can resolve the eight signal levels. The gain and offset of the signal must be precisely set so that the quantizing levels register precisely with the centres of the eyes. The transmitted signal contains sync pulses which are encoded using specified code levels so that the data separator can set its gain and offset.

Multi-level signalling systems have the characteristic that the bits in the symbol have different error probability. Figure 7.7 shows that a small noise level will corrupt the low-order bit, whereas twice as much noise will be needed to corrupt the middle bit and four times as much will be needed to corrupt the high-order bit. In ATSC the solution is that the lower two bits are encoded together in an inner error-correcting scheme so that they represent only one bit with similar reliability to the top bit. As a result the 8-VSB system actually delivers two data bits per symbol even though eight-level signalling is used.




Figure 7.7 In multi-level signalling the error probability is not the same for each bit.

Multi-level signalling can be combined with PSK to obtain multi-level quadrature amplitude modulation (QUAM). Figure 7.8 shows the example of 64-QUAM. Incoming six-bit data words are split into two three-bit words and each is used to amplitude modulate a pair of sinusoidal carriers which are generated in quadrature. The modulators

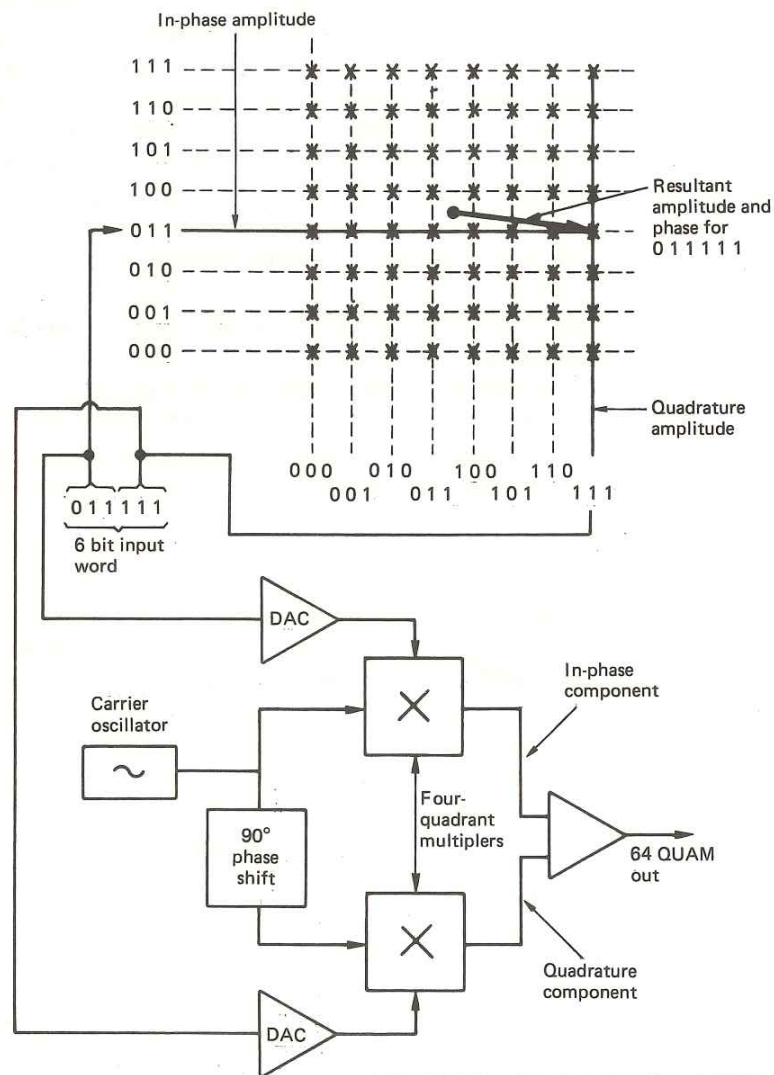


Figure 7.8 In 64-QUAM, two carriers are generated with a quadrature relationship. These are independently amplitude modulated to eight discrete levels in four quadrant multipliers. Adding the signals produces a QUAM signal having 64 unique combinations of amplitude and phase. Decoding requires the waveform to be sampled in quadrature like a colour TV subcarrier.



are four-quadrant devices such that  $2^3$  amplitudes are available, four which are in phase with the carrier and four which are antiphase. The two AM carriers are linearly added and the result is a signal which has  $2^6$  or 64 combinations of amplitude and phase. There is a great deal of similarity between QUAM and the colour subcarrier used in analog television in which the two colour difference signals are encoded into one amplitude- and phase-modulated waveform. On reception, the waveform is sampled twice per cycle in phase with the two original carriers and the result is a pair of eight-level signals. 16-QUAM is also possible, delivering only four bits per symbol but requiring a lower SNR.

The data bit patterns to be transmitted can have any combinations whatsoever, and if nothing were done, the transmitted spectrum would be non-uniform. This is undesirable because peaks cause interference with other services, whereas energy troughs allow external interference in. A randomizing technique known as energy dispersal is used to overcome the problem. The signal energy is spread uniformly throughout the allowable channel bandwidth so that it has less energy at a given frequency.

A pseudo-random sequence (PRS) generator is used to generate the randomizing sequence. Figure 7.9 shows the randomizer used in DVB. This sixteen-bit device has a maximum sequence length of 65 535 bits, and is preset to a standard value at the beginning of each set of eight transport

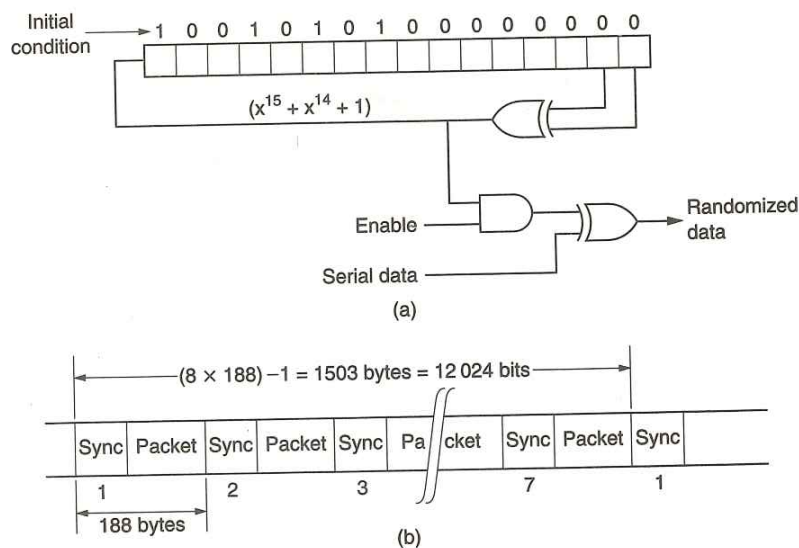


Figure 7.9 The randomizer of DVB is pre-set to the initial condition once every eight transport stream packets. The maximum length of the sequence is 65 535 bits, but only the first 12 024 bits are used before resetting again (b).

stream packets. The serialized data are XORed with the sequence, which randomizes the output that then goes to the modulator. The spectrum of the transmission is now determined by the spectrum of the PRS.

On reception, the de-randomizer must contain the identical ring counter which must also be set to the starting condition to bit accuracy. Its output is then added to the data stream from the demodulator. The randomizing will effectively then have been added twice to the data in modulo-2, and as a result is cancelled out leaving the original serial data.

The way that radio signals interact with obstacles is a function of the relative magnitude of the wavelength and the size of the object. AM sound radio transmissions with a wavelength of several hundred metres can easily diffract around large objects. The shorter the wavelength of a transmission, the larger objects in the environment appear to it and these objects can then become reflectors. Reflecting objects produce a delayed signal at the receiver in addition to the direct signal. In analog television transmissions this causes the familiar ghosting. In digital transmissions, the symbol rate may be so high that the reflected signal may be one or more symbols behind the direct signal, causing intersymbol interference. As the reflection may be continuous, the result may be that almost every symbol is corrupted. No error-correction system can handle this. Raising the transmitter power is no help at all as it simply raises the power of the reflection in proportion.

The only solution is to change the characteristics of the RF channel in some way to either prevent the multipath reception or to prevent it being a problem. The RF channel includes the modulator, transmitter, antennae, receiver and demodulator.

As with analog UHF TV transmissions, a directional antenna is useful with digital transmission as it can reject reflections. However, directional antennae tend to be large and they require a skilled permanent installation. Mobile use on a vehicle or vessel is simply impractical.

Another possibility is to incorporate a ghost canceller in the receiver. The transmitter periodically sends a standardized known waveform known as a training sequence. The receiver knows what this waveform looks like and compares it with the received signal. In theory it is possible for the receiver to compute the delay and relative level of a reflection and so insert an opposing one. In practice if the reflection is strong it may prevent the receiver finding the training sequence.

The most elegant approach is to use a system in which multipath reception conditions cause only a small increase in error rate which the error-correction system can manage. This approach is used in DVB. Figure 7.10(a) shows that when using one carrier with a high bit rate, reflections can easily be delayed by one or more bit periods, causing interference between the bits. Figure 7.10(b) shows that instead, OFDM



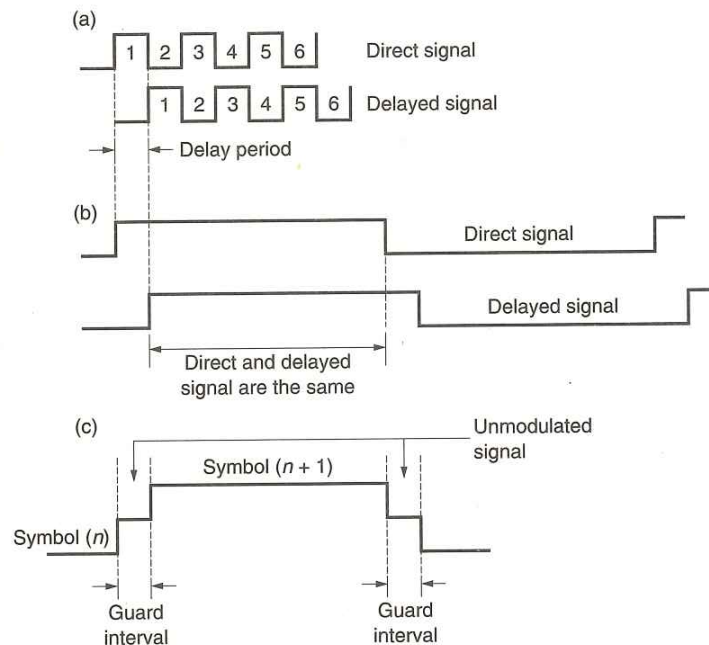


Figure 7.10 (a) High bit rate transmissions are prone to corruption due to reflections. (b) If the bit rate is reduced the effect of reflections is eliminated, in fact reflected energy can be used. (c) Guard intervals may be inserted between symbols.

sends many carriers each having a low bit rate. When a low bit rate is used, the energy in the reflection will arrive during the same bit period as the direct signal. Not only is the system immune to multipath reflections, but the energy in the reflections can actually be used. This characteristic can be enhanced by using guard intervals shown in Figure 7.10(c). These reduce multipath bit overlap even more.

Note that OFDM is not a modulation scheme, and each of the carriers used in a OFDM system still needs to be modulated using any of the digital coding schemes described above. What OFDM does is to provide an efficient way of packing many carriers close together without mutual interference.

A serial data waveform basically contains a train of rectangular pulses. The transform of a rectangle is the function  $\sin x/x$  and so the baseband pulse train has a  $\sin x/x$  spectrum. When this waveform is used to modulate a carrier the result is a symmetrical  $\sin x/x$  spectrum centred on the carrier frequency. Figure 7.11(a) shows that nulls in the spectrum appear spaced at multiples of the bit rate away from the carrier.

Further carriers can be placed at spacings such that each is centred at the nulls of the others as is shown in Figure 7.11(b). The distance between




Figure 7.11 In OFDM the carrier spacing is critical, but when correct the carriers become independent and most efficient use is made of the spectrum. (a) Spectrum of bitstream has regular nulls. (b) Peak of one carrier occurs at null of another.

the carriers is equal to  $90^\circ$  or one quadrant of  $\sin x$ . Owing to the quadrant spacing, these carriers are mutually orthogonal, hence the term orthogonal frequency division. A large number of such carriers (in practice several thousand) will be interleaved to produce an overall spectrum which is almost rectangular and which fills the available transmission channel.

When guard intervals are used, the carrier returns to an unmodulated state between bits for a period which is greater than the period of the reflections. Then the reflections from one transmitted bit decay during the guard interval before the next bit is transmitted. The use of guard intervals reduces the bit rate of the carrier because for some of the time it is radiating carrier, not data. A typical reduction is to around 80 per cent of the capacity without guard intervals.

This capacity reduction does, however, improve the error statistics dramatically, such that much less redundancy is required in the error-correction system. Thus the effective transmission rate is improved. The use of guard intervals also moves more energy from the sidebands back to the carrier. The frequency spectrum of a set of carriers is no longer perfectly flat but contains a small peak at the centre of each carrier.

The ability to work in the presence of multipath cancellation is one of the great strengths of OFDM. In DVB, more than 2000 carriers are used in single transmitter systems. Provided there is exact synchronism, several transmitters can radiate exactly the same signal so that a single-frequency network can be created throughout a whole country. SFNs require a variation on OFDM which uses over 8000 carriers.



With OFDM, directional antennae are not needed and, given sufficient field strength, mobile reception is perfectly feasible. Of course, directional antennae may still be used to boost the received signal outside of normal service areas or to enable the use of low-powered transmitters.

An OFDM receiver must perform fast Fourier transforms (FFTs) on the whole band at the symbol rate of one of the carriers. The amplitude and/or phase of the carrier at a given frequency effectively reflects the state of the transmitted symbol at that time slot and so the FFT partially demodulates as well.

In order to assist with tuning in, the OFDM spectrum contains pilot signals. These are individual carriers which are transmitted with slightly more power than the remainder. The pilot carriers are spaced apart through the whole channel at agreed frequencies which form part of the transmission standard. Practical reception conditions, including multipath reception, will cause a significant variation in the received spectrum and some equalization will be needed. Figure 7.12 shows what the




Figure 7.12 Multipath reception can place notches in the channel spectrum. This will require equalization at the receiver.

possible spectrum looks like in the presence of a powerful reflection. The signal has almost been cancelled at certain frequencies. However, the FFT performed in the receiver is effectively a spectral analysis of the signal and so the receiver computes for free the received spectrum. As in a flat spectrum the peak magnitude of all the coefficients would be the same (apart from the pilots), equalization is easily performed by multiplying the coefficients by suitable constants until this characteristic is obtained.

Although the use of transform-based receivers appears complex, when it is considered that such an approach simultaneously allows effective equalization the complexity is not significantly higher than that of a conventional receiver which needs a separate spectral analysis system just for equalization purposes. The only drawback of OFDM is that the transmitter must be highly linear to prevent intermodulation between the carriers. This is readily achieved in terrestrial transmitters by derating the transmitter so that it runs at a lower power than it would in analog service.

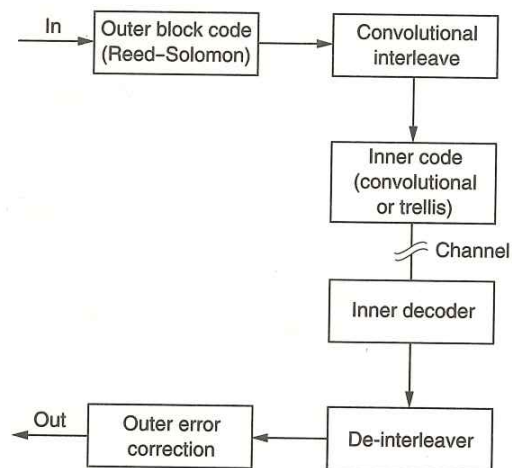


Figure 7.13 Error-correcting strategy of digital television broadcasting systems.

This is not practicable in satellite transmitters which are optimized for efficiency, so OFDM is not really suitable for satellite use.

Broadcast data suffer from both random and burst errors and the error-correction strategies of digital television broadcasting have to reflect that. Figure 7.13 shows a typical system in which inner and outer codes are employed. The Reed-Solomon codes are universally used for burst-correcting outer codes, along with an interleave which will be convolutional rather than the block-based interleave used in recording media. The inner codes will not be R-S, as more suitable codes exist for the statistical conditions prevalent in broadcasting. DVB uses a parity based variable-rate system in which the amount of redundancy can be adjusted according to reception conditions.

Figure 7.14 shows a block diagram of a DVB-T (terrestrial) transmitter. Incoming transport stream packets of 188 bytes each are first subject to

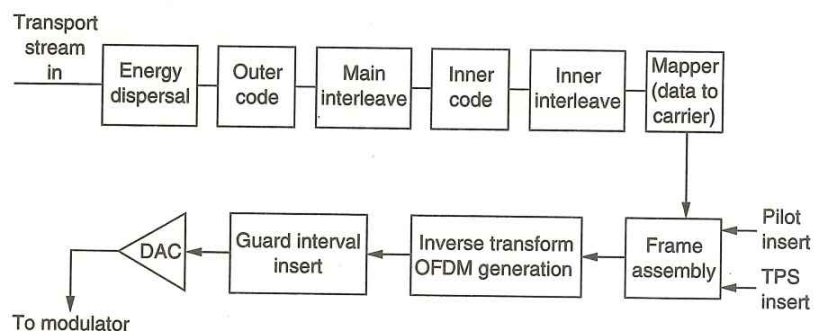
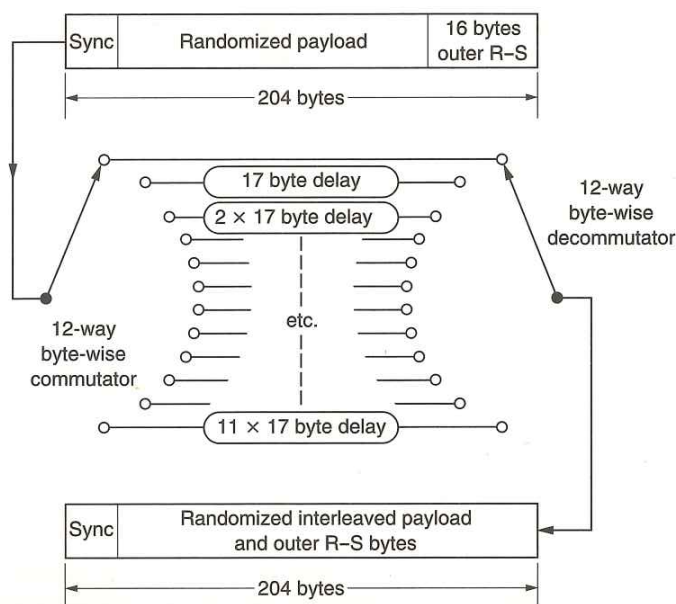


Figure 7.14 DVB-T transmitter block diagram. See text for details.



**Figure 7.15** The interleaver of DVB uses 12 incrementing delay channels to reorder the data. The sync byte passes through the undelayed channel and so is still at the head of the packet after interleave. However, the packet now contains non-adjacent bytes from 12 different packets.

R-S outer coding. This adds 16 bytes of redundancy to each packet, resulting in 204 bytes. Outer coding is followed by interleaving. The interleave mechanism is shown in Figure 7.15. Outer code blocks are commutated on a byte basis into twelve parallel channels. Each channel contains a different amount of delay, typically achieved by a ring-buffer RAM. The delays are integer multiples of 17 bytes, designed to skew the data by one outer block ( $12 \times 17 = 204$ ). Following the delays, a commutator reassembles interleaved outer blocks. These have 204 bytes as before, but the effect of the interleave is that adjacent bytes in the input are 17 bytes apart in the output. Each output block contains data from twelve input blocks making the data resistant to burst errors.

Following the interleave, the energy-dispersal process takes place. The pseudo-random sequence runs over eight outer blocks and is synchronized by inverting the transport stream packet sync symbol in every eighth block. The packet sync symbols are not randomized.

The inner coding process of DVB is convolutional. The output bit rate of the inner coder is twice the input bit rate. Under the worst reception conditions, this 100 per cent redundancy offers the most powerful correction with the penalty that a low data rate is delivered. However, a



variety of inner redundancy factors can be used from  $1/2$  down to  $1/8$  of the transmitted bit rate.

The DVB standard allows the use of QPSK, 16-QAM or 64-QAM coding in an OFDM system. There are five possible inner code rates, and four different guard intervals which can be used with each modulation scheme. Thus for each modulation scheme there are twenty possible transport stream bit rates in the standard DVB channel, each of which requires a different receiver SNR. The broadcaster can select any suitable balance between transport stream bit rate and coverage area. For a given transmitter location and power, reception over a larger area may require a channel code with a smaller number of bits/s/Hz and this reduces the bit rate which can be delivered in a standard channel.

Alternatively a higher amount of inner redundancy means that the proportion of the transmitted bit rate which is data goes down. Thus for wider coverage the broadcaster will have to send fewer programs in the multiplex or use higher compression factors.

## 7.4 The DVB receiver

Figure 7.16 shows a block diagram of a DVB receiver. The off-air RF signal is fed to a mixer driven by the local oscillator. The IF output of the mixer is bandpass filtered and supplied to the ADC which outputs a digital IF signal for the FFT stage. The FFT is analysed initially to find the higher-level pilot signals. If these are not in the correct channels the local

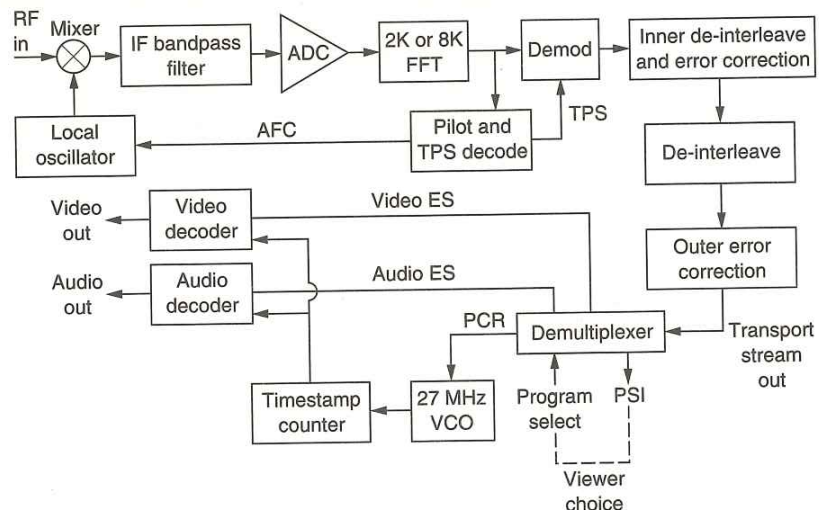


Figure 7.16 DVB receiver block diagram. See text for details.

oscillator frequency is incorrect and it will be changed until the pilots emerge from the FFT in the right channels. The data in the pilots will be decoded in order to tell the receiver how many carriers, what inner redundancy rate, guard band rate and modulation scheme are in use in the remaining carriers. The FFT magnitude information is also a measure of the equalization required.

The FFT outputs are demodulated into 2K or 8K bitstreams and these are multiplexed to produce a serial signal. This is subject to inner error correction which corrects random errors. The data are then de-interleaved to break up burst errors and then the outer R-S error correction operates. The output of the R-S correction will then be de-randomized to become an MPEG transport stream once more. The de-randomizing is synchronized by the transmission of inverted sync patterns.

The receiver must select a PID of 0 and wait until a Program Association Table (PAT) is transmitted. This will describe the available programs by listing the PIDs of the Program Map Tables (PMT). By looking for these packets the receiver can determine what PIDs to select to receive any video and audio elementary streams.

When an elementary stream is selected, some of the packets will have extended headers containing Program Clock Reference (PCR). These codes are used to synchronize the 27 MHz clock in the receiver to the one in the MPEG encoder of the desired program. The 27 MHz clock is divided down to drive the time stamp counter so that audio and video emerge from the decoder at the correct rate and with lip sync.

It should be appreciated that time stamps are relative, not absolute. The time stamp count advances by a fixed amount each picture, but the exact count is meaningless. Thus the decoder can only establish the frame rate of the video from time stamps, but not the precise timing. In practice the receiver has finite buffering memory between the demultiplexer and the MPEG decoder. If the displayed video timing is too late, the buffer will tend to overflow whereas if the displayed video timing is too early the decoding may not be completed. The receiver can advance or retard the time stamp counter during lock-up so that it places the output timing mid-way between these extremes.

## 7.5 CD-Video and DVD

CD-Video and DVD are both digital optical disks which store video and audio program material. CD-Video is based on the Compact Disc, with which it shares the same track dimensions and bit rate. MPEG-1 coding is used on video signals adhering to the common intermediate format (CIF – see Chapter 1). The input subsampling needed to obtain the CIF imaging and the high compression factor needed to comply with the low



bit rate of a digital audio disk means that the quality of CD-Video is marginal.

DVD is based on the same concepts as Compact Disc, but with a number of enhancements. The laser in the pickup has a shorter wavelength. In conjunction with a larger lens aperture this allows the readout spot significantly to be reduced in size. This in turn allows the symbols on the tracks to be shorter as well as allowing the track spacing to be reduced. The EFM code of CD is modified to EFMPlus which improves the raw error rate.

The greatly increased capacity of DVD means that moderate compression factors can be employed. MPEG-2 coding is used, so that progressively scanned or interlaced material can be handled. As most DVDs are mastered from movie film, the use of progressive scan is common.

The perceived quality of DVDs can be very high indeed. This is partly due to the fact that a variable bit rate is supported. When difficult frames are encountered, the bit rate can rise above the average. Another advantage of DVD is that it can be mastered in non-real time. This allows techniques such as aligning cuts with *I*-pictures to be used. Also the quality obtained can be assessed and if artifacts are observed the process can be modified.

Figure 7.17 shows the block diagram of a typical DVD player, and illustrates the essential components. A CD-Video player is similar. The most natural division within the block diagram is into the control/servo system and the data path. The control system provides the interface between the user and the servo mechanisms, and performs the logical interlocking required for safety and the correct sequence of operation.

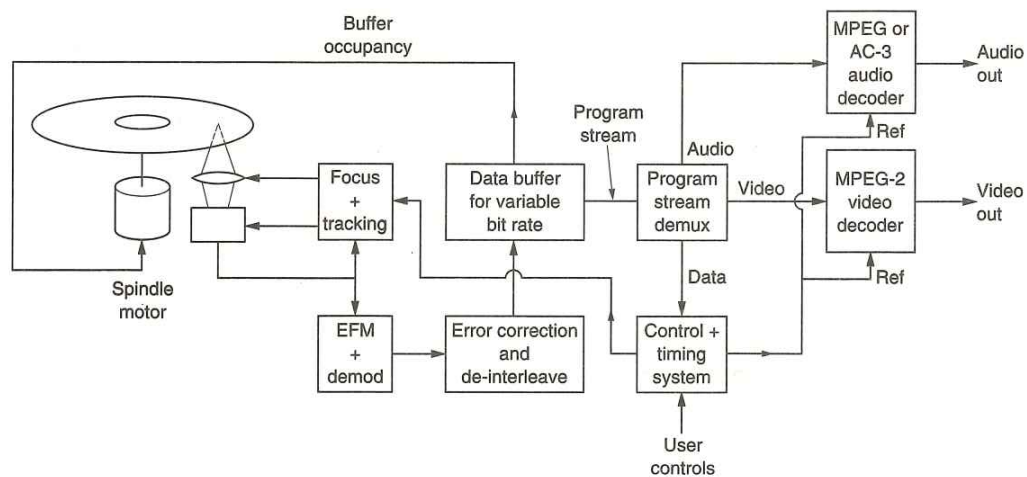


Figure 7.17 A DVD player's essential parts. See text for details.

The servo systems include any power-operated loading drawer and chucking mechanism, the spindle-drive servo, and the focus and tracking servos already described.

Power loading is usually implemented on players where the disk is placed in a drawer. Once the drawer has been pulled into the machine, the disk is lowered onto the drive spindle, and clamped at the centre, a process known as chucking. In the simpler top-loading machines, the disk is placed on the spindle by hand, and the clamp is attached to the lid so that it operates as the lid is closed. The lid or drawer mechanisms have a safety switch which prevents the laser operating if the machine is open. This is to ensure that there can be no conceivable hazard to the user. In actuality there is very little hazard in a DVD pickup. This is because the beam is focused a few millimetres away from the objective lens, and beyond the focal point the beam diverges and the intensity falls rapidly. It is almost impossible to position the eye at the focal point when the pickup is mounted in the player, but it would be foolhardy to attempt to disprove this.

The data path consists of the data separator, the de-interleaving and error-correction process followed by a RAM buffer which supplies the MPEG decoder.

The data separator converts the EFMplus readout waveform into data. Following data separation the error-correction and de-interleave processes take place. Because of the interleave system, there are two opportunities for correction, first, using the inner code prior to de-interleaving, and second, using the outer code after de-interleaving. Interleaving is designed to spread the effects of burst errors among many different codewords, so that the errors in each are reduced. However, the process can be impaired if a small random error, due perhaps to an imperfection in manufacture, occurs close to a burst error caused by surface contamination. The function of the inner redundancy is to correct single-symbol errors, so that the power of interleaving to handle bursts is undiminished, and to generate error flags for the outer system when a gross error is encountered.

The EFMplus coding is a group code which means that a small defect which changes one channel pattern into another could have corrupted up to eight data bits. In the worst case, if the small defect is on the boundary between two channel patterns, two successive bytes could be corrupted. However, the final odd/even interleave on encoding ensures that the two bytes damaged will be in different inner codewords; thus a random error can never corrupt two bytes in one inner codeword, and random errors are therefore always correctable.

The de-interleave process is achieved by writing sequentially into a memory and reading out using a sequencer. The outer decoder will then correct any burst errors in the data. As MPEG data are very sensitive to error the error-correction performance has to be extremely good.



Following the de-interleave and outer error-correction process an MPEG program stream (see Chapter 6) emerges. Some of the program stream data will be video, some will be audio and this will be routed to the appropriate decoder. It is a fundamental concept of DVD that the bit rate of this program stream is not fixed, but can vary with the difficulty of the program material in order to maintain consistent image quality. However, the disk uses constant linear velocity rather than constant angular velocity. It is not possible to obtain a particular bit rate with a fixed spindle speed.

The solution is to use a RAM buffer between the transport and the MPEG decoders. The RAM is addressed by counters which are arranged to overflow, giving the memory a ring structure. Writing into the memory is done using clocks derived from the disk whose frequency rises and falls with runout, whereas reading is done by the decoder which, for each picture, will take as much data as is required from the buffer.

The buffer will only function properly if the two addresses are kept apart. This implies that the amount of data read from the disk over the long term must equal the amount of data used by the MPEG decoders. This is done by analysing the address relationship of the buffer. If the bit rate from the disk is too high, the write address will move towards the read address; if it is too low, the write address moves away from the read address. Subtraction of the two addresses produces an error signal which can be fed to the drive.

In some DVD drives, the disk spins at a relatively high speed, resulting in an excessively high continuous data rate. However, this is reduced by jumping the pickup back. Repeating a previous track does not produce any new data and so the average rate falls. As all the disk blocks are labelled it is a simple matter to reassemble the discontinuous bitstream in memory.

The exact speed of the motor is unimportant. The important factor is that the data rate needed by the decoder is correct, and the system will skip the pickup as often as necessary so that the buffer neither underflows nor overflows.

The MPEG-2 decoder will convert the compressed elementary streams into PCM video and audio and place the pictures and audio blocks into RAM. These will be read out of RAM whenever the time stamps recorded with each picture or audio block match the state of a time stamp counter. If bidirectional coding is used, the RAM readout sequence will convert the recorded picture sequence back to the real-time sequence. The time stamp counter is derived from a crystal oscillator in the player which is divided down to provide the 90 kHz time stamp clock. As a result the frame rate at which the disk was mastered will be replicated as the pictures are read from RAM. Once a picture buffer is read out, this will trigger the decoder to decode another picture. It will read data from the



buffer until this has been completed and thus indirectly influence the disk bit rate.

Owing to the use of constant linear velocity, the disk speed will be wrong if the pickup is suddenly made to jump to a different radius using manual search controls. This may force the data separator out of lock, or cause a buffer overflow and the decoder may freeze briefly until this has been remedied. The control system of a DVD player is inevitably microprocessor-based, and as such does not differ greatly in hardware terms from any other microprocessor-controlled device. Operator controls will simply interface to processor input ports and the various servo systems will be enabled or overridden by output ports. Software, or more correctly firmware, connects the two. The necessary controls are Play and Eject, with the addition in most players of at least Pause and some buttons which allow rapid skipping through the program material.

Although machines vary in detail, the flowchart of Figure 7.18 shows the logic flow of a simple player, from start being pressed to pictures and sound emerging. At the beginning, the emphasis is on bringing the various servos into operation. Towards the end, the disk subcode is read in order to locate the beginning of the first section of the program material.

When track following, the tracking-error feedback loop is closed, but for track crossing, in order to locate a piece of music, the loop is opened,

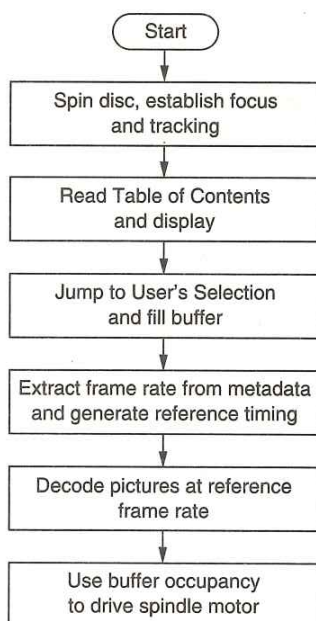


Figure 7.18 Simple processes required for a DVD player to operate.


and a microprocessor signal forces the laser head to move. The tracking error becomes an approximate sinusoid as tracks are crossed. The cycles of tracking error can be counted as feedback to determine when the correct number of tracks have been crossed. The 'mirror' signal obtained when the readout spot is half a track away from target is used to brake pickup motion and re-enable the track following feedback.

## 7.6 Personal video recorders

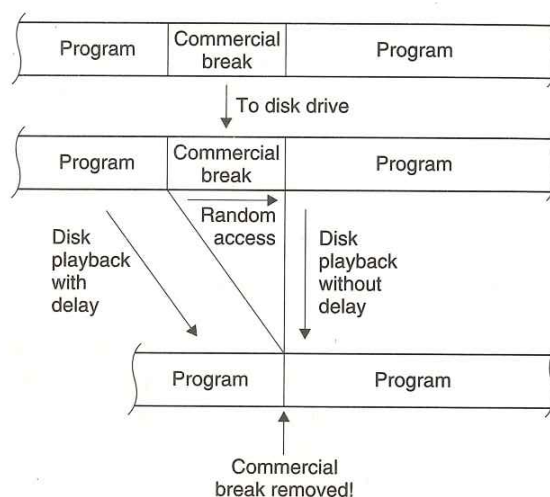
The development of the consumer VCR was a small step to end the linearity of commercial television. Increasing numbers of viewers use VCRs not just as time shifters but also as a means of fast-forwarding through the commercial breaks. The non-linear storage of video was until recently restricted by economics to professional applications. However, with the falling cost of hard disk drives and the availability of MPEG compression, non-linear video storage is now a consumer product.

The personal video recorder (PVR) is based on the random access storage of a hard disk drive. As disk drives can move their heads from track to track across the disk surface, they can access data anywhere very quickly. Unlike tape, which can only record or play back but not both at the same time, a PVR can do both simultaneously at arbitrary points on a time line.

Figure 7.19 shows that the disk drive can transfer data much faster than the required bit rate, and so it transfers data in bursts which are smoothed out by RAM buffers. It is straightforward to interleave read and write



**Figure 7.19** In a hard disk recorder, a large-capacity memory is used as a buffer or timebase corrector between the converters and the disk. The memory allows the converters to run constantly despite the interruptions in disk transfer caused by the head moving between tracks.



**Figure 7.20** The ability of the consumer to use random access to skip over commercial breaks may destroy commercial television as it exists today.

functions so that it gives the impression of reading and writing simultaneously. The read and write processes can take place from anywhere on the disk.

Although the PVR can be used as an ordinary video recorder, it can do some other tricks. Figure 7.20 shows the most far-reaching trick. The disk drive starts recording an off-air commercial TV station. A few minutes later the viewer starts playing back the recording. When the commercial break is transmitted, the disk drive may record it, but the viewer can skip over it using the random access of the hard drive. With suitable software the hard drive could skip over the commercial break automatically by simply not recording it.

When used with digital television systems, the PVR can simply record the transmitted transport stream data and replay it into an MPEG decoder. In this way the PVR has no quality loss whatsoever. The picture quality will be the same as off-air. Optionally PVRs may also have MPEG encoders so that they can record analog video inputs.

Real-time playback in a PVR is trivial as it is just a question of retrieving recorded data as the decoder buffer demands it. However, the consumer is accustomed to the ability of the VCR to produce a picture over a range of speeds as well as a freeze frame. PVRs incorporate additional processing which modifies the recorded MPEG bitstream on playback so that a standard MPEG decoder will reproduce it with a modified timebase. Figure 7.21 shows that a typical off-air MPEG elementary stream uses bidirectional coding with a moderately long GOP structure. An increase in apparent playback speed can be obtained by






Figure 7.21 By manipulating the elementary stream, a compliant bitstream can be produced which will play at a different speed from the original.

discarding some of the pictures in the group. This has to be done with caution as the pictures are heavily interdependent. However, if the group is terminated after a *P* picture, all pictures up to that point are decodable. The resultant motion portrayal will not be smooth because pictures are being omitted, but the goal is simply to aid the viewer in finding the right place in the recording so this is of no consequence. In practice the results are good, not least because the noise bar of the analog VCR is eliminated.

High forward speeds can be obtained by selecting only the *I* pictures from the recorded bitstream. These cannot be transmitted to the decoder directly, because the *I* pictures contain a lot of data and the decoder buffer might overflow. Consequently the data rate is diluted by sending null *P* pictures. These are *P* pictures in which all of the vectors are zero and the residual data are also zero. Null *P* pictures require very few data but have the effect of converting an *I* picture into another identical picture. A similar technique can be used to decode in reverse. With a normal MPEG bidirectionally coded bitstream, reverse decoding is impossible, but by taking *I* pictures in reverse order and padding them with null-*P* pictures a reverse replay is obtained.

## 7.7 Networks

A network is basically a communication resource which is shared for economic reasons. Like any shared resource, decisions have to be made somewhere and somehow about how the resource is to be used. In the

absence of such decisions the resultant chaos will be such that the resource might as well not exist.

In communications networks the resource is the ability to convey data from any node or port to any other. On a particular cable, clearly only one transaction of this kind can take place at any one instant even though in practice many nodes will simultaneously be wanting to transmit data. Arbitration is needed to determine which node is allowed to transmit.

Data networks originated to serve the requirements of computers and it is a simple fact that most computer processes don't need to be performed in real time or indeed at a particular time at all. Networks tend to reflect that background as many of them, particularly the older ones, are asynchronous. Asynchronous means that the time taken to deliver a given quantity of data is unknown. A TDM system may chop the data into several different transfers and each transfer may experience delay according to what other transfers the system is engaged in. Ethernet and most storage system buses are asynchronous. For broadcasting purposes an asynchronous delivery system is no use at all, but for copying an MPEG data file between two storage devices an asynchronous system is perfectly adequate.

The opposite extreme is the synchronous system in which the network can guarantee a constant delivery rate and a fixed and minor delay. An AES/EBU digital audio router or an SDI digital video router is a synchronous network. In between asynchronous and synchronous networks reside the isochronous approaches which cause a fixed moderate delay. These can be thought of as sloppy synchronous networks or more rigidly controlled asynchronous networks.

These three different approaches are needed for economic reasons. Asynchronous systems are very efficient because as soon as one transfer completes another one can begin. This can only be achieved by making every device wait with its data in a buffer so that transfer can start immediately. Asynchronous systems also make it possible for low bit rate devices to share a network with high bit rate devices. The low bit rate device will only need a small buffer and will send few cells, whereas the high bit rate device will send more cells.

Isochronous systems try to give the best of both worlds, generally by sacrificing some flexibility in block size. Modern networks are tending to be part isochronous and part asynchronous so that the advantages of both are available.

There are a number of different arbitration protocols and these have evolved to support the needs of different types of network. In small networks, such as LANs, a single point failure which halts the entire network may be acceptable, whereas in a public transport network owned by a telecommunications company, the network will be redundant so that if a particular link fails data may be sent via an alternative route.



A link which has reached its maximum capacity may also be supplanted by transmission over alternative routes.

In physically small networks, arbitration may be carried out in a single location. This is fast and efficient, but if the arbitrator fails it leaves the system completely crippled. The processor buses in computers work in this way.

In centrally arbitrated systems the arbitrator needs to know the structure of the system and the status of all the nodes. Following a configuration change, due perhaps to the installation of new equipment, the arbitrator needs to be told what the new configuration is, or have a mechanism which allows it to explore the network and learn the configuration. Central arbitration is only suitable for small networks which change their configuration infrequently.


In other networks the arbitration is distributed so that some decision-making ability exists in every node. This is less efficient but it does allow at least some of the network to continue operating after a component failure.

Distributed arbitration also means that each node is self-sufficient and so no changes need to be made if the network is reconfigured by adding or deleting a node. This is the only possible approach in wide area networks where the structure may be very complex and changes dynamically in the event of failures or overload.

Ethernet uses distributed arbitration. FireWire is capable of using both types of arbitration. A small amount of decision-making ability is built into every node so that distributed arbitration is possible. However, if one of the nodes happens to be a computer, it can run a centralized arbitration algorithm.

The physical structure of a network is subject to some variation as Figure 7.22 shows. In radial networks, (a), each port has an unique cable connection to a device called a *hub*. The hub must have one connection for every port and this limits the number of ports. However, a cable failure will only result in the loss of one port. In a ring system (b) the nodes are connected like a daisy chain with each node acting as a feedthrough. In this case the arbitration requirement must be distributed. With some protocols, a single cable break doesn't stop the network operating. Depending on the protocol, simultaneous transactions may be possible provided they don't require the same cable. For example, in a storage network a disk drive may be outputting data to an editor whilst another drive is backing up data to a tape streamer. For the lowest cost, all nodes are physically connected in parallel to the same cable. Figure 7.22(c) shows that a cable break would divide the network into two halves, but it is possible that the impedance mismatch at the break could stop both halves working.

One of the concepts involved in arbitration is priority which is fundamental to providing an appropriate quality of service. If two



**Figure 7.22** Network configurations. At (a) the radial system uses one cable to each node. (b) Ring system uses less cable than radial. (c) Linear system is simple but has no redundancy.

processes both want to use a network, the one with the highest priority would normally go first.

Attributing priority must be done carefully because some of the results are non-intuitive. For example, it may be beneficial to give a high priority to a humble device which has a low data rate for the simple reason that if it is given use of the network it won't need it for long. In a television environment transactions concerned with on-air processes would have priority over file transfers concerning production and editing.

When a device gains access to the network to perform a transaction, generally no other transaction can take place until it has finished. Consequently it is important to limit the amount of time that a given port can stay on the bus. In this way when the time limit expires, a further arbitration must take place. The result is that the network resource rotates between transactions rather than one transfer hogging the resource and shutting everyone else out. It follows from the presence of a time (or data quantity) limit that ports must have the means to break large files up into frames or cells and reassemble them on reception. This process is sometimes called *adaptation*. If the data to be sent originally exist at a fixed bit rate, some buffering will be needed so that the data can be time-compressed into the available frames. Each frame must be contiguously numbered and the system must transmit a file size or word count so that the receiving node knows when it has received every frame in the file.



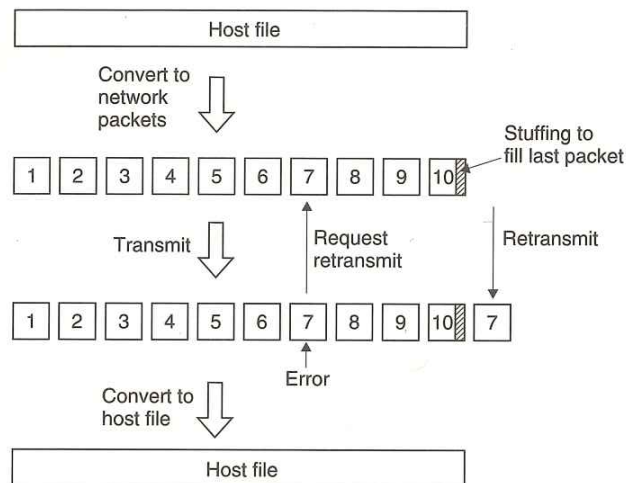


Figure 7.23 Receiving a file which has been divided into packets allows for the retransmission of just the packet in error.

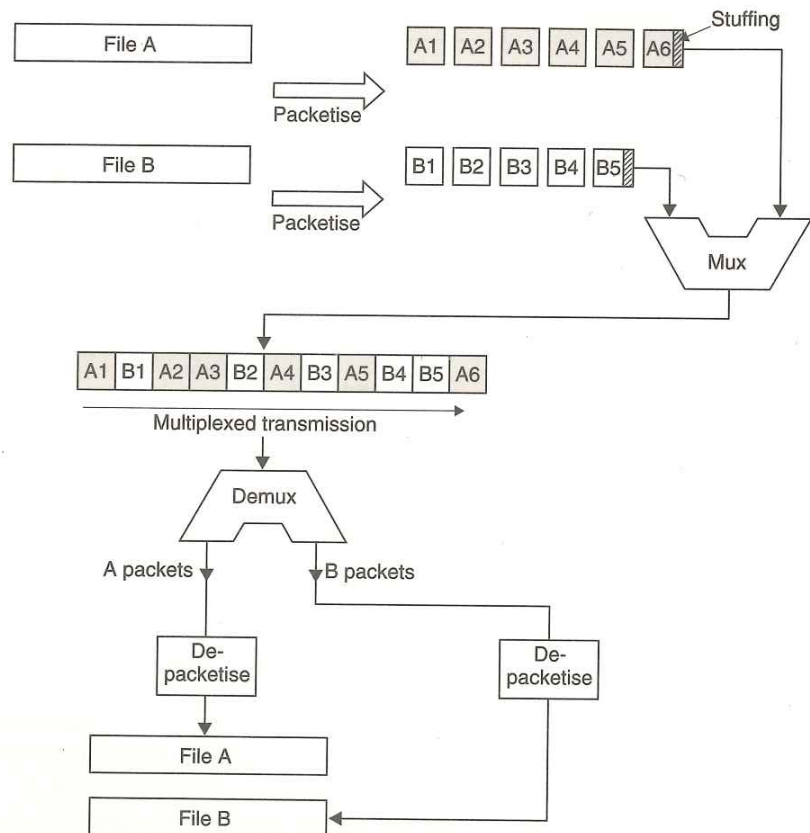
The error-detection system interacts with this process because if any frame is in error on reception, the receiving node can ask for a retransmission of the frame. This is more efficient than retransmitting the whole file. Figure 7.23 shows the mechanism of retransmission where one packet has been lost or is in error. The retransmitted packet is inserted at the correct place in the bitstream by the receiver.

Breaking files into frames helps to keep down the delay experienced by each process using the network. Figure 7.24 shows a frame-multiplexed transmission. Each frame may be stored ready for transmission in a silo memory. It is possible to make the priority a function of the number of frames in the silo, as this is a direct measure of how long a process has been kept waiting. Isochronous systems must do this in order to meet maximum delay specifications.

When delivering an MPEG bitstream over an isochronous network, the network transmitter and receiver will require buffering over and above that provided in the MPEG encoder and decoder. This will enable the receiver to re-create a bitstream at the decoder which has essentially the same temporal characteristics as that leaving the encoder, except for a fixed delay. The greater the maximum delay of the network, the greater that fixed delay will have to be. Providing buffering is expensive and delays cause difficulty in applications such as videoconferencing. In this case it is better to use a network protocol which can limit delay by guaranteeing bandwidth.

A central arbitrator is relatively simple to implement because when all decisions are taken centrally there can be no timing difficulty (assuming





**Figure 7.24** Files are broken into frames or packets for multiplexing with packets from other users. Short packets minimize the time between the arrival of successive packets. The priority of the multiplexing must favour isochronous data over asynchronous data.

a well-engineered system). In a distributed system, there is an extra difficulty due to the finite time taken for signals to travel down the data paths between nodes.

Figure 7.25 shows the structure of Ethernet which uses a protocol called CSMA/CD (carrier sense multiple access with collision detect) developed by DEC and Xerox. This is a distributed arbitration network where each node follows some simple rules. The first of these is not to transmit if an existing bus signal is detected. The second is not to transmit more than a certain quantity of data before releasing the bus. Devices wanting to use the bus will see bus signals and so will wait until the present bus transaction finishes. This must happen at some point because of the frame size limit. When the frame is completed, signalling on the bus should cease. The first device to sense the bus becoming free and to assert its own

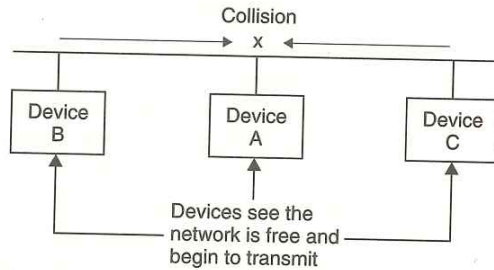



Figure 7.25 In Ethernet collisions can occur because of the finite speed of the signals. A 'back-off' algorithm handles collisions, but they do reduce the network throughput.

signal will prevent any other nodes transmitting according to the first rule. Where numerous devices are present it is possible to give them a priority structure by providing a delay between sensing the bus coming free and beginning a transaction. High-priority devices will have a short delay so they get in first. Lower-priority devices will only be able to start a transaction if the high-priority devices don't need to transfer.

It might be thought that these rules would be enough and everything would be fine. Unfortunately the finite signal speed means that there is a flaw in the system. Figure 7.25 shows why. Device A is transmitting and devices B and C both want to transmit and have equal priority. At the end of A's transaction, devices B and C see the bus become free at the same instant and start a transaction. With two devices driving the bus, the resultant waveform is meaningless. This is known as a collision and all nodes must have means to recover from it. First, each node will read the bus signal at all times. When a node drives the bus, it will also read back the bus signal and compare it with what was sent. Clearly if the two are the same all is well, but if there is a difference, this must be because a collision has occurred and two devices are trying to determine the bus voltage at once.

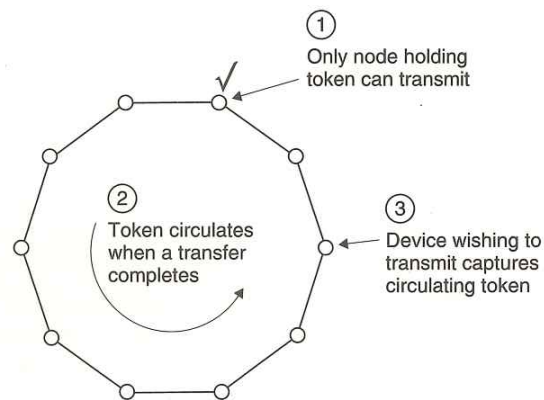
If a collision is detected, both colliding devices will sense the disparity between the transmitted and readback signals, and both will release the bus to terminate the collision. However, there is no point in adhering to the simple protocol to reconnect because this will simply result in another collision. Instead each device has a built-in delay which must expire before another attempt is made to transmit. This delay is not fixed, but is controlled by a random number generator and so changes from transaction to transaction.

The probability of two node devices arriving at the same delay is infinitesimally small. Consequently if a collision does occur, both devices will drop the bus, and they will start their back-off timers. When the first timer expires, that device will transmit and the other will see the

transmission and remain silent. In this way the collision is not only handled, but is prevented from happening again.

The performance of Ethernet is usually specified in terms of the bit rate at which the cabling runs. However, this rate is academic because it is not available all of the time. In a real network bit rate is lost by the need to send headers and error-correction codes and by the loss of time due to interframe spaces and collision handling. As the demand goes up, the number of collisions increases and throughput goes down. Collision-based arbitrators do not handle congestion well.

An alternative method of arbitration developed by IBM is shown in Figure 7.26. This is known as a *token ring* system. All the nodes have an



**Figure 7.26** In a token ring system only the node in possession of the token can transmit so collisions are impossible. In very large rings the token circulation time causes loss of throughput.

input and an output and are connected in a ring which must be complete for the system to work. Data circulate in one direction only. If data are not addressed to a node which receives it, the data will be passed on. When the data arrive at the addressed node, that node will capture the data as well as passing them on with an acknowledge added. Thus the data packet travels right around the ring back to the sending node. When the sending node receives the acknowledge, it will transmit a token packet. This token packet passes to the next node, which will pass it on if it does not wish to transmit. If no device wishes to transmit, the token will circulate endlessly. However, if a device has data to send, it simply waits until the token arrives again and captures it. This node can now transmit data in the knowledge that there cannot be a collision because no other node has the token.



In simple token ring systems, the transmitting node transmits idle characters after the data packet has been sent in order to maintain synchronization. The idle character transmission will continue until the acknowledge arrives. In the case of long packets the acknowledge will arrive before the packet has all been sent and no idle characters are necessary. However, with short packets idle characters will be generated. These idle characters use up ring bandwidth.

Later token ring systems use *early token release* (ETR). After the packet has been transmitted, the sending node sends a token straight away. Another node wishing to transmit can do so as soon as the current packet has passed. It might be thought that the nodes on the ring would transmit in their physical order, but this is not the case because a priority system exists. Each node can have a different priority if necessary. If a high-priority node wishes to transmit, as a packet from elsewhere passes through that node, the node will set *reservation bits* with its own priority level. When the sending node finishes and transmits a token, it will copy that priority level into the token.

In this way nodes with a lower priority level will pass the token on instead of capturing it. The token will ultimately arrive at the high-priority node.

The token ring system has the advantage that it does not waste throughput with collisions and so the full capacity is always available. However, if the ring is broken the entire network fails.


In Ethernet the performance is degraded by the number of transactions, not the number of nodes, whereas in token ring the performance is degraded by the number of nodes.

## 7.8 FireWire

FireWire<sup>3</sup> is actually an Apple Computers Inc. trade name for the interface which is formally known as IEEE 1394-1995. It was originally intended as a digital audio network, but grew out of recognition. FireWire is more than just an interface as it can be used to form networks and if used with a computer effectively extends the computer's data bus. Figure 7.27 shows that devices are simply connected together as any combination of daisy-chain or star network.

Any pair of devices can communicate in either direction, and arbitration ensures that only one device transmits at once. Intermediate devices simply pass on transmissions. This can continue even if the intermediate device is powered down as the FireWire carries power to keep repeated functions active.

Communications are divided into *cycles* which have a period of 125  $\mu$ s. During a cycle, there are 64 time slots. During each time slot, any one



**Figure 7.27** FireWire supports radial (star) or daisy-chain connection. Two-port devices pass on signals destined for a more distant device – they can do this even when powered down.

node can communicate with any other, but in the next slot, a different pair of nodes may communicate. Thus FireWire is best described as a time-division multiplexed (TDM) system. There will be a new arbitration between the nodes for each cycle.

FireWire is eminently suitable for video/computer convergent applications because it can simultaneously support asynchronous transfers of non-real-time computer data and isochronous transfers of real-time audio/video data. It can do this because the arbitration process allocates a fixed proportion of slots for isochronous data (about 80 per cent) and these have a higher priority in the arbitration than the asynchronous data. The higher the data rate a given node needs, the more time slots it will be allocated. Thus a given bit rate can be guaranteed throughout a transaction; a prerequisite of real-time A/V data transfer.

It is the sophistication of the arbitration system which makes FireWire remarkable. Some of the arbitration is in hardware at each node, but some is in software which only needs to be at one node. The full functionality requires a computer somewhere in the system which runs the isochronous bus management arbitration. Without this only asynchronous transfers are possible. It is possible to add or remove devices whilst the system is working. When a device is added the system will recognize it through a periodic learning process. Essentially every node on the system transmits in turn so that the structure becomes clear.

The electrical interface of FireWire is shown in Figure 7.28. It consists of two twisted pairs for signalling and a pair of power conductors. The twisted pairs carry differential signals of about 220 mV swinging around a common mode voltage of about 1.9 V with an impedance of 112  $\Omega$ . Figure 7.29 shows how the data are transmitted. The host data are simply serialized and used to modulate twisted pair A. The other twisted pair (B)






Figure 7.28 FireWire uses twin twisted pairs and a power pair.




Figure 7.29 The strobe signal is the X-OR of the data and the bit clock. The data and strobe signals together form a self-clocking system.

carries a signal called *strobe*, which is the exclusive-OR of the data and the clock. Thus whenever a run of identical bits results in no transitions in the data, the strobe signal will carry transitions. At the receiver another exclusive-OR gate adds data and strobe to re-create the clock.

This signalling technique is subject to skew between the two twisted pairs and this limits cable lengths to about 10 metres between nodes. Thus FireWire is not a long-distance interface technique, instead it is very useful for interconnecting a large number of devices in close proximity. Using a copper interconnect, FireWire can run at 100, 200 or 400 megabits/s, depending on the specific hardware. It is proposed to create an optical fibre version which would run at gigabit speeds.

## 7.9 Broadband networks and ATM

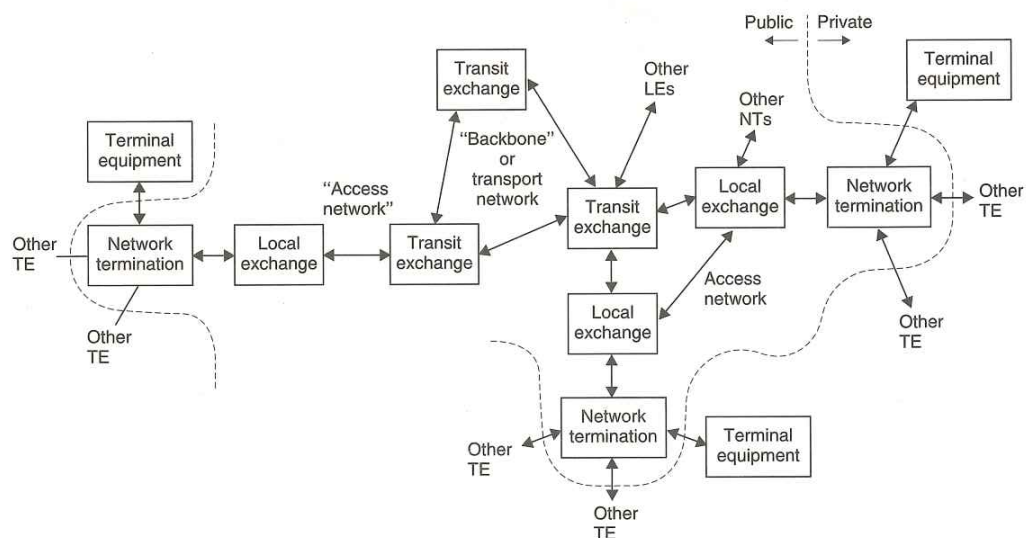
Broadband ISDN (B-ISDN) is the successor to N-ISDN and in addition to providing more bandwidth, offers practical solutions to the delivery of any conceivable type of data. The flexibility with which ATM operates means that intermittent or one-off data transactions which only require

asynchronous delivery can take place alongside isochronous MPEG video delivery. This is known as *application independence* whereby the sophistication of isochronous delivery does not raise the cost of asynchronous data. In this way, generic data, video, speech and combinations of the above can co-exist.

ATM is multiplexed, but it is not time-division multiplexed. TDM is inefficient because if a transaction does not fill its allotted bandwidth, the capacity is wasted. ATM does not offer fixed blocks of bandwidth, but allows infinitely variable bandwidth to each transaction. This is done by converting all host data into small fixed size cells at the adaptation layer. The greater the bandwidth needed by a transaction, the more cells per second are allocated to that transaction. This approach is superior to the fixed bandwidth approach, because if the bit rate of a particular transaction falls, the cells released can be used for other transactions so that the full bandwidth is always available.

As all cells are identical in size, a multiplexer can assemble cells from many transactions in an arbitrary order. The exact order is determined by the quality of service required, where the time positioning of isochronous data would be determined first, with asynchronous data filling the gaps.

Figure 7.30 shows how a broadband system might be implemented. The transport network would typically be optical fibre based, using SONET (Synchronous optical network) or SDH (synchronous digital hierarchy). These standards differ in minor respects. Figure 7.31 shows



**Figure 7.30** Structure and terminology of a broadband network. See text.