process have inverse elements, and most combinations are lossless. For example, variable-length encoding and decoding is lossless, so it need not be considered as part of the system. The combination of DCT and inverse DCT (IDTC) is lossless with sufficient arithmetic precision. This pair of elements contributes round-off errors, but little else.

The lossy process is, of course, quantization. There is a block called dequantization, but it serves only to return the coefficients to the correct reconstruction values; the original information is irretrievably lost.

However, once this step has been performed, we can again look at a lossless system. Assuming there is no signal processing between the output of the decoder and the input to the next encoder, we can follow the values. Quantized coefficients are passed via variable-length encode/decode (lossless), then to the inverse DCT, thus producing a set of sample levels. If DCT/IDCT is nearly lossless, so is IDTC followed by DCT. After the DCT process of the next encoder, we will have the same coefficient set that was produced by the dequantizer. Each of these levels should be right in the center of the decision range for the new quantizer, so that it should produce the same result. Figure 19-1 illustrates.

Figure 19-1
Lossless groupings in concatenated compression systems.

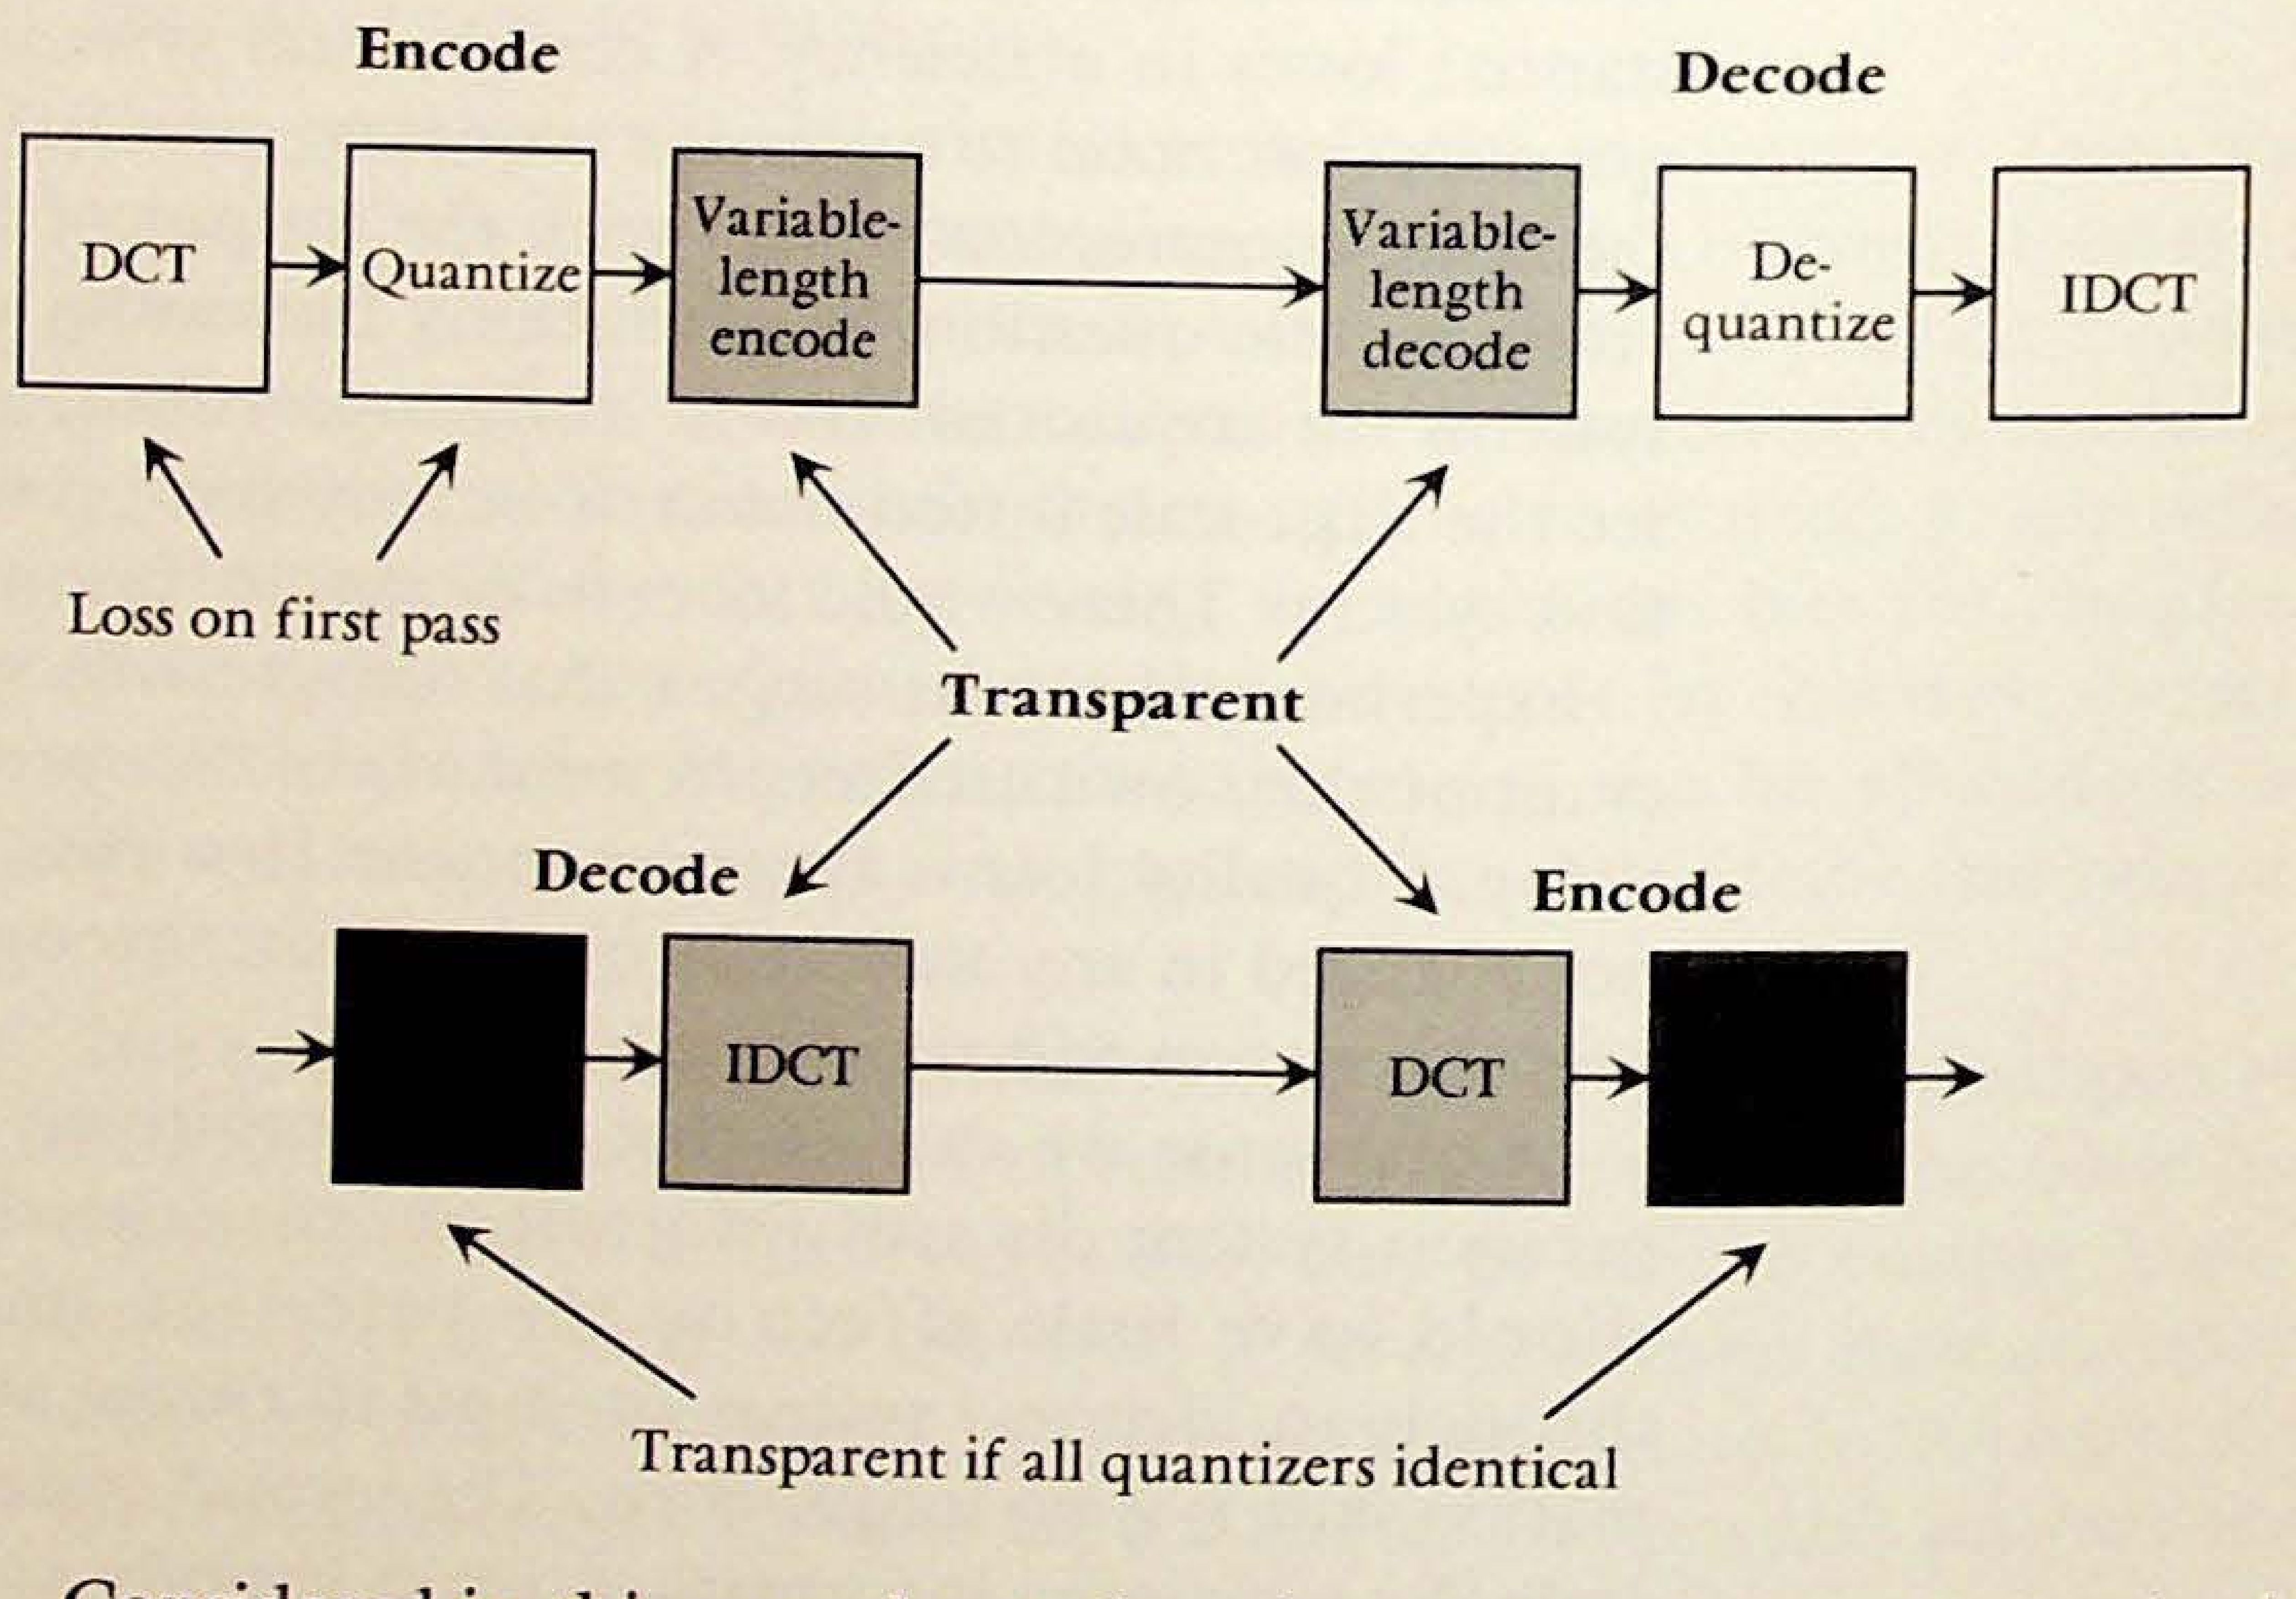

Considered in this way, the explanation of the observed behavior is simple. Quantization causes the loss observed on the first pass. On subsequent passes the same quantization results should occur every time.

The only issue is the DCT/IDCT match, which is never perfect. Arithmetic precision errors and round-off errors contribute the very gradual deterioration in performance on subsequent generations.

It will be seen that the quantizer is an essential element of this process. The lossless situation exists only if the quantizers are identical. Successive coding by two encoders that use different quantizers will likely result in quite rapid deterioration.

Researchers at Sarnoff Corporation have shown that compression systems can be very sensitive to the artifacts produced by other compression systems. In one experiment, two "light" compression systems intended for studio operations with high-definition television were compared. One was based on MPEG, the other on wavelets. The signals, both of which were free from visible artifacts, were fed to an ATSC transmission encoder that uses MPEG compression.

The transmission encoder performed normally when fed with the output of the MPEG system, but displayed a much higher level of artifacts when fed from the wavelet system.

This should be cause for concern. Already, television plants are full of compressed nonlinear edit systems and tape formats that use compression. We are about to see widespread deployment of MPEG-based delivery systems, yet we see the rapid growth of DV compression, which uses a quantization strategy quite different from MPEG. None of this will cause an overnight disaster, but we may be eroding our quality headroom without even realizing it.

Two parts of a European project are of great interest in this area. In one, developers are achieving essentially lossless recoding, even with a change of bit rate. All the coding decisions (motion vectors, etc.) of the original compression are extracted from the decoder and sent along with the video. (There is a proposal to code this data by using the least significant bit of one of the color difference signals.) The same coding decisions can then be used by the second encoder—a big step toward retaining quality.

The other work specifically concerns changing bit rates—perhaps from a distribution signal to a broadcast delivery signal. The encoder uses a special form of quantizer that recognizes the previous quantization decisions and attempts to minimize the combined effect of the two quantizers. Demonstrated results are spectacular.

Audio compression systems are not immune to the problems of concatenation. As with video, a compression system designed for final delivery to the point of use should be designed so that the effects of quantization are undetectable. If two identical systems are concatenated then, just as with video, the decisions in the second system should match those of the first and result in little or no additional impairment, *if the signal has not changed.* If the signal does change, the likely result is that a coarsely quantized signal is requantized differently, and it is very likely that the impairment will be unacceptable.

If nonidentical compression systems are concatenated, rapid deterioration is possible, again just like video. The artifacts of one system are likely to disrupt the efficient operation of the other. An interesting and very important example is the interaction of Dolby Surround Pro Logic (DSPL) with Dolby Digital (DD), if wrongly used, as described in the previous chapter.

Switching MPEG

MPEG was conceived for a relatively simple scenario in which information is encoded, sent to the decoder, and decoded, without any intervening steps other than transmission and/or storage. It became evident that there were situations where it would be desirable to switch from one program stream to another, or to a different part of the same stream. An obvious example in the context of MPEG's early objectives is the interactive movie—a story that switches to alternative routes and endings according to actions of the viewer.

Because of these considerations, a number of provisions were made within MPEG for switching or *splicing* of the bit stream. However, these techniques were really designed for the type of situation just described, where there is prior determination of the switch-out points and of the potential switch-in points. There is no mechanism in MPEG for a generalized switch from any point in one bit stream to any point in another.

Some closer thought makes it clear why this cannot be. At the lowest level, a variable-length—encoded bit stream is meaningful only if decoded from the beginning. Arbitrary switching provides no way to know what the data is, or what it means. If we tried to switch bit streams, like those shown in Chapter 7, we could (and likely would) end up interpreting extra bits as Huffman codes; what should have been motion vectors would be wrongly decoded and, perhaps, used as if the values were DCT coefficients.

We noted earlier that one of the important characteristics of the MPEG slice is that it is encoded without any outside references. All predictive encodings are reset at the beginning of a slice, so that we can begin decoding unambiguously at the start of any slice. Actually this demonstrates that we have not yet defined the problem. We could start decoding the new stream at any splice, but it is not productive to begin decoding a picture in the middle. Resynchronization at a slice boundary is useful for recovery from data errors, but not for switching between programs.

Perhaps we can use the picture boundaries for switching, just as we use vertical interval switching of baseband video today. For I-frame—only encoding we can do this, subject to certain conditions. I-frame—only streams, like motion JPEG streams, can be switched with reasonable ease. Even at this point there are complications. If we permit the same number of bits for each frame (the simplest case), a frame is 33.4 ms long in the nominal 60-Hz regions (such as North America) and 40 ms long in 50-Hz areas (such as Europe and Australasia). Neither lines up with the 32 ms used by a frame of AC-3 audio. Typically this means leaving a small gap in the audio whenever a switch is made, and this complication exists however or wherever we switch.

With other GOP structures the situation is more complex. Suppose we switch to the new bit stream, and the first picture is a P-frame. Some macroblocks will likely be intracoded, and these will decode successfully. Others, however, will be predictively encoded—the bit stream will carry a motion vector identifying a similar macroblock in the previous I-frame and a set of DCT-coded information representing the changes that should be made to that macroblock. The assumption made by the encoder is, of course, that the I-frame is present in storage at the decoder. Well, there is an I-frame in the decoder store, but it belongs to the previous video sequence—the one before the switch (Figure 19-2)! If other conditions are correct the decoder will use it, but it will produce nonsense.

If the frame following the switch is a B-frame, the situation is only slightly worse. The B-frame may use both forward and backward prediction, on the assumption that both of the surrounding pair of anchor frames (*I, I* or *I, P* or *P, P*) are in the decoder memory. Probably there are two frames in memory, and they may even be a pair of adjacent anchor frames, but both will belong to the previous video stream (Figure 19-3).

Figure 19-2
Incorrect P-frame reference after a splice.

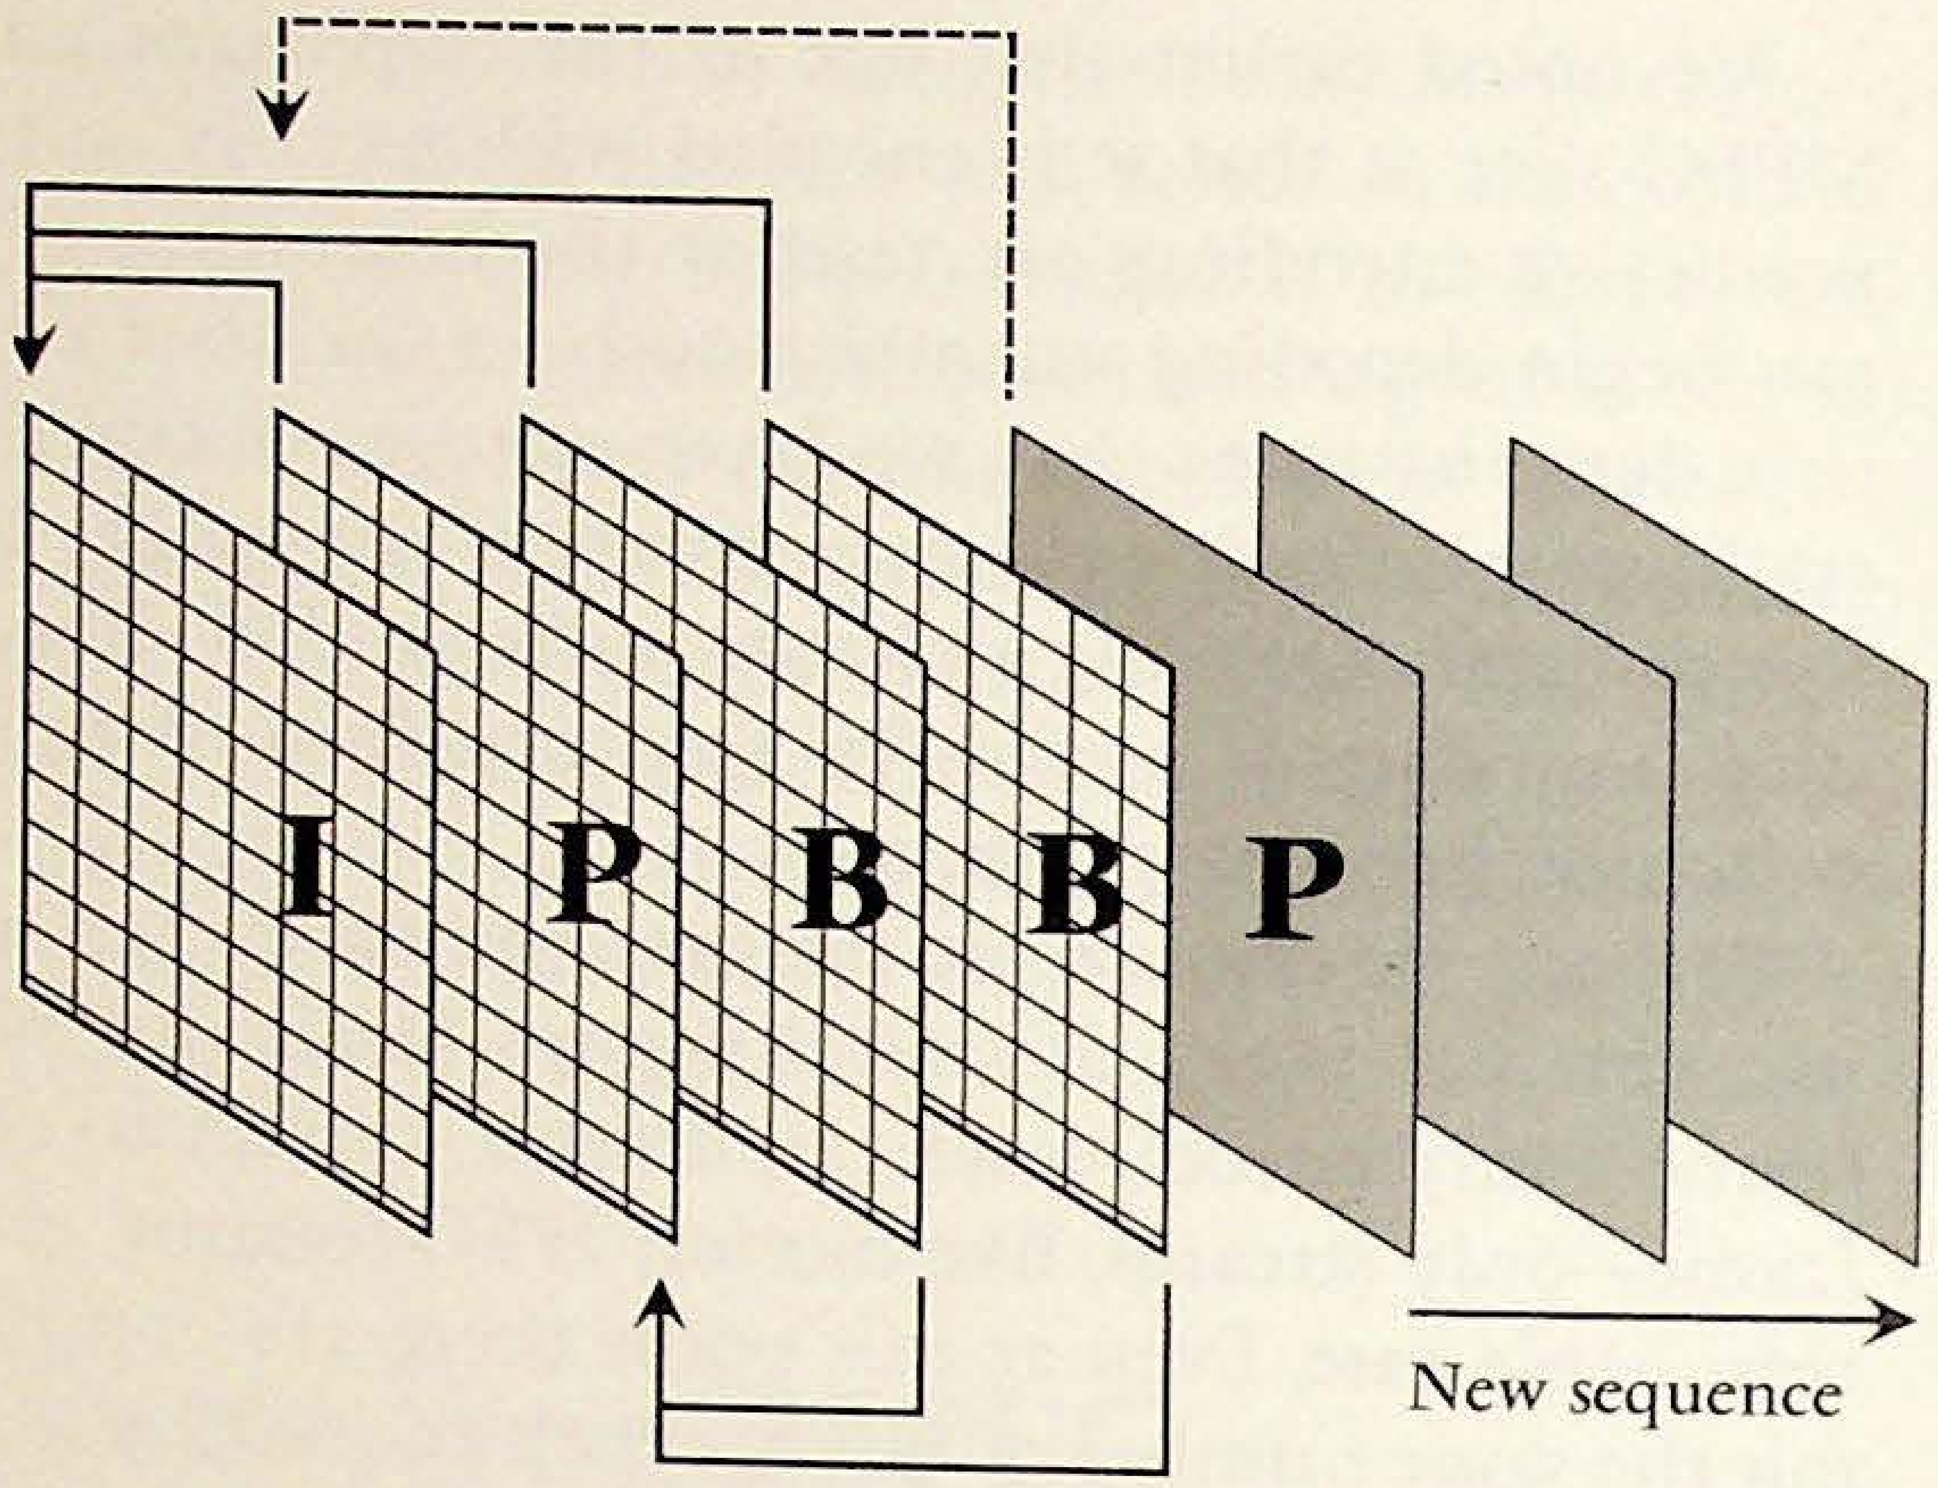

Figure 19-3
Incorrect references from a B-frame.

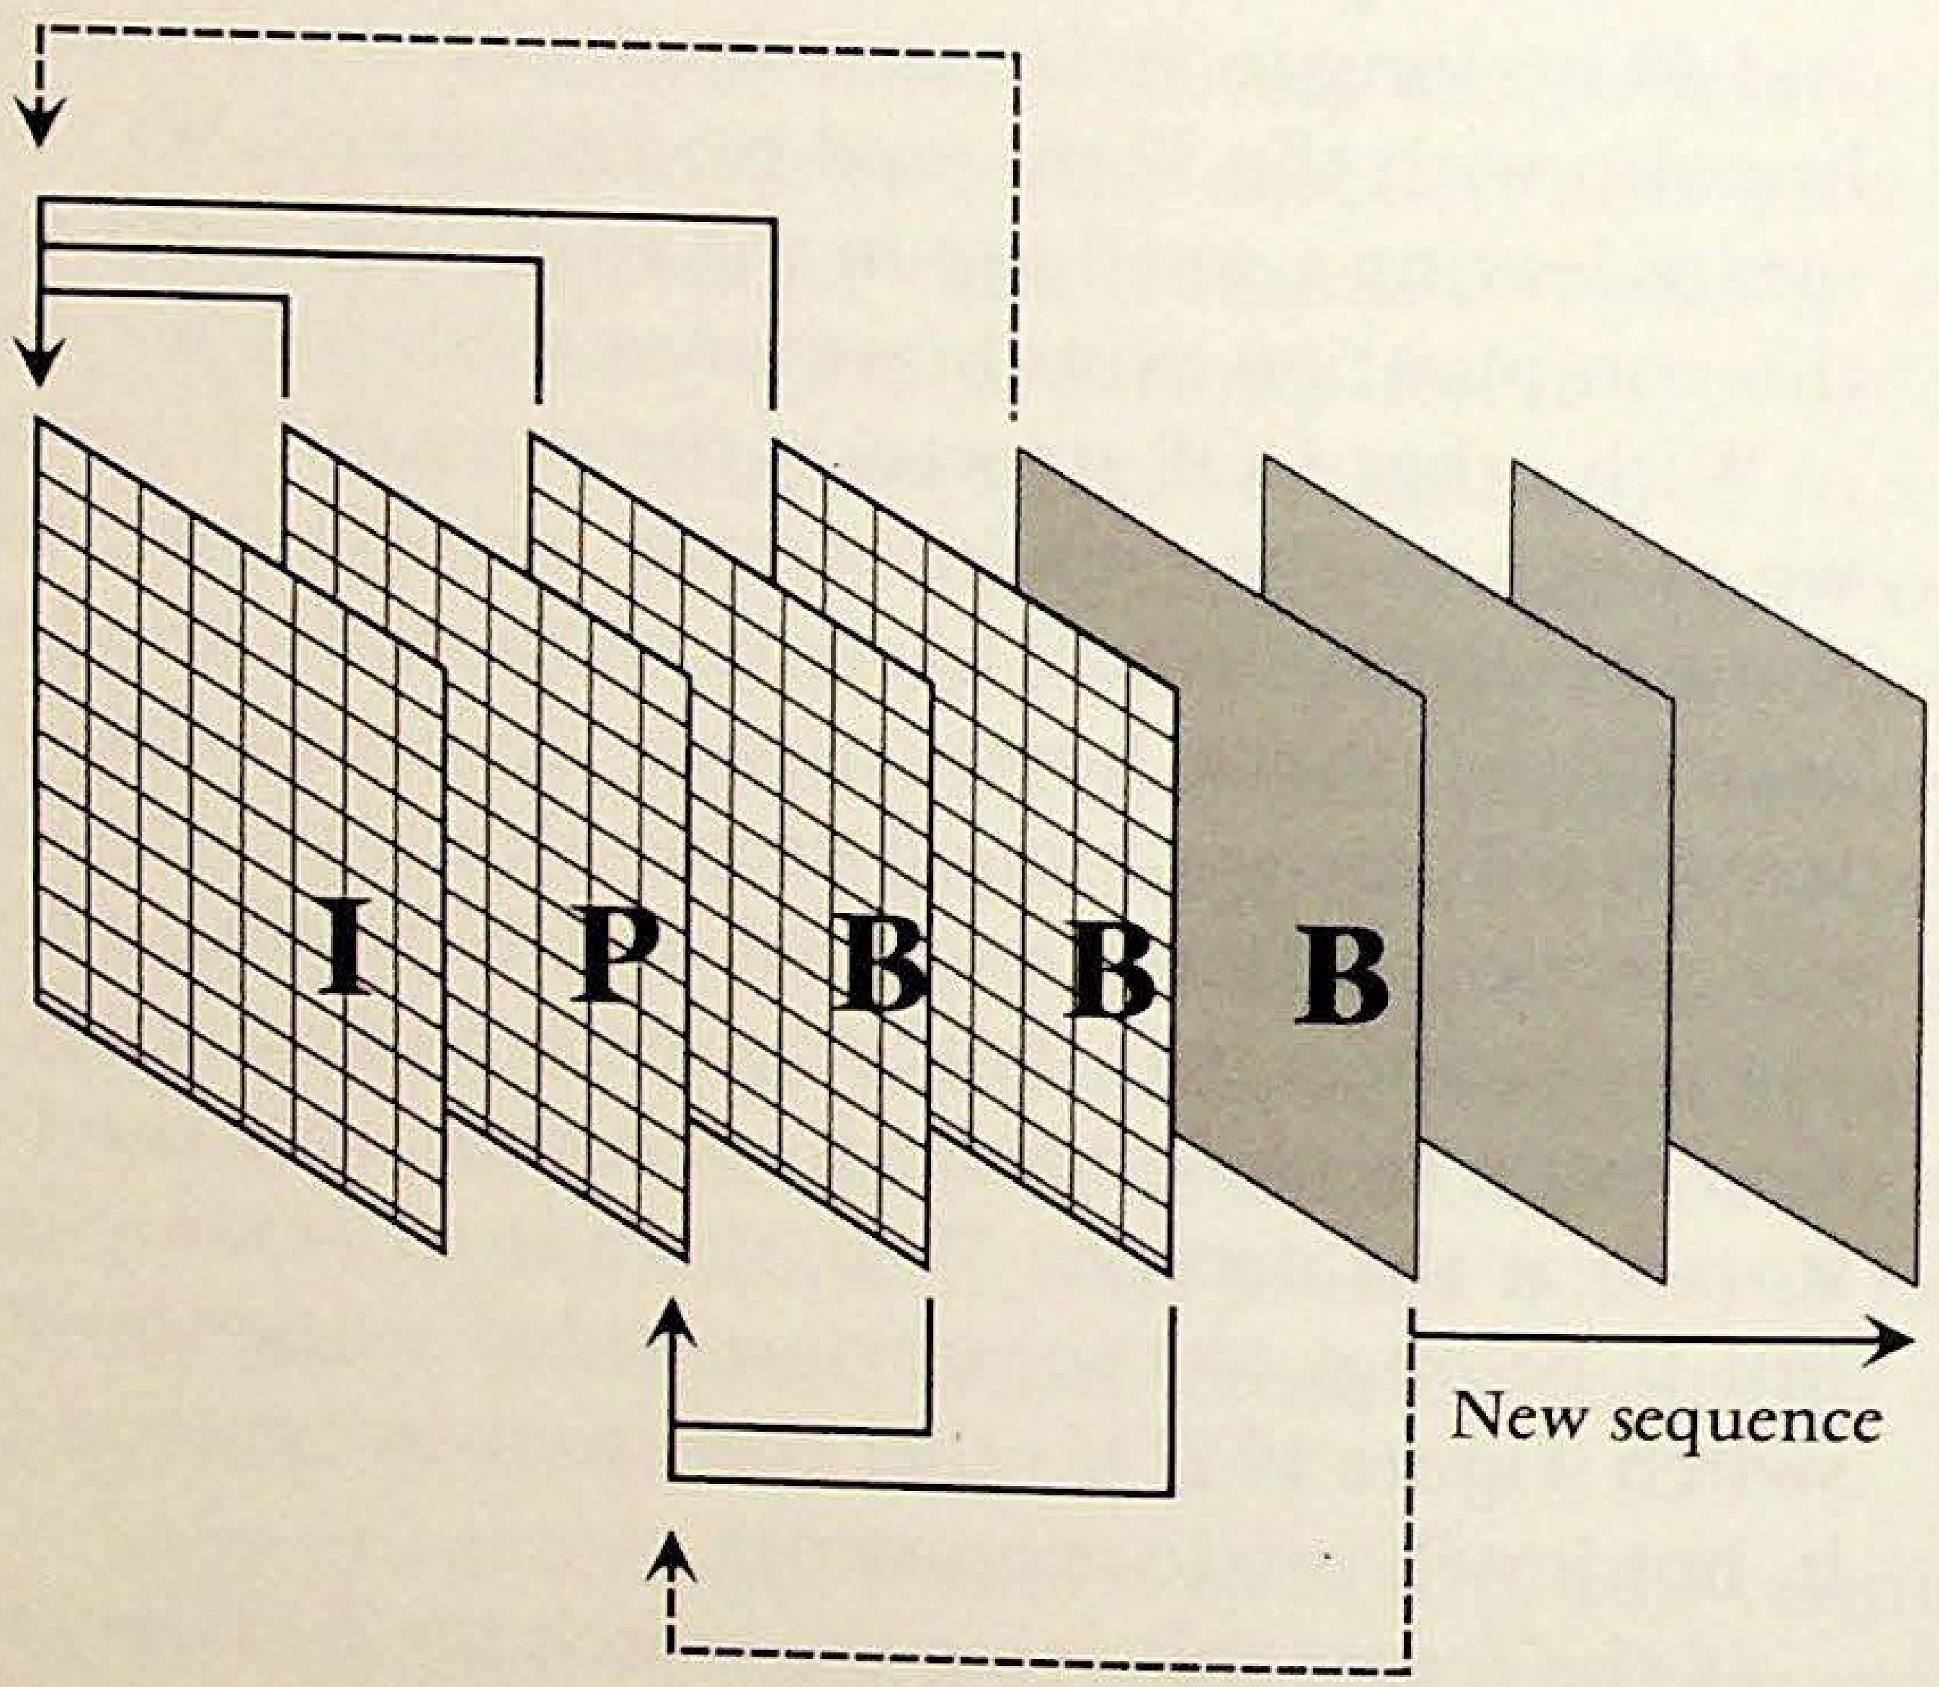

It appears that we may be able to switch MPEG bit streams only at GOP boundaries, even though this may not provide sufficient resolution for some applications. We will look at some applications in a moment, but there is one more major issue awaiting us, even at the GOP boundary.

In Chapter 9 we looked at MPEG rate control and discussed the concept of the VBV buffer model. Although not based on real buffer designs, VBV provides a means of specifying decoder buffer performance and a means of checking that a bit stream is decodable by a standard decoder. The basic conditions are that VBV should never overflow or underflow.

VBV is filled at a nominally constant rate by the transmission channel. It is emptied at regular (frame) intervals, but by varying amounts depending on whether the frame being extracted is an I-, P-, or B-frame, and on the complexity of the particular frame. VBV represents an important degree of freedom in the battle for constant quality, and a good decoder will use the full gamut of the buffer to help achieve this. At times the buffer will be close to overflow, at other times close to underflow.

It all works because the encoder tracks VBV fullness at all times; except, of course, when we switch bit streams! The encoder for the second stream is managing VBV for hypothetical decoders that have been receiving the bit stream *since the beginning of the sequence.* If any (or all) of the decoders have been receiving a different sequence, there is no way for the encoder to "know" that the switch has occurred, or that the decoder VBV is in an unknown and unknowable state. We might switch to an almost full VBV at a time when the second stream encoder has calculated that VBV should be almost empty; overflow is inevitable. Similarly, switching to a near-empty VBV when the encoder is tracking VBV as close to full will result in underflow.

What happens when VBV overflows or underflows is undefined by MPEG, and depends on the design of a particular decoder. Most likely (and this is probably the least disturbing solution) the picture will freeze until the decoder finds that it again has valid data waiting to be decoded. Unfortunately, the complexity of the process makes it uncertain how quickly this point will be reached. It is also quite possible that the decoder will start decoding again, but with a VBV state that still differs from that assumed by the encoder. In this case another overflow or underflow is likely, but perhaps some considerable time after the switch.

As a final twist, statistical multiplexing might be the issue that makes everything else look easy. It is difficult to imagine how to switch into one program stream—that is, to replace all packets associated with that one program in the multiplexed bit stream—when the number of bits allocated to that program changes dynamically under the control of some device that may be on the other side of the country. A working scenario would have to include rules about bit rate floors, at least at commercial insertion times, and bit rate ceilings for any material to be inserted at the switch.

Clearly, switching MPEG bit streams is a complex and uncertain process. In fact, the preceding discussion presents only a gross simplification of the problem. Almost certainly the general problem for two unconstrained bit streams is insoluble. An MPEG expert would suggest that MPEG is not designed for such treatment, so "Don't do that!" Unfortunately, in the world of television broadcasting there are a number of areas where switching of MPEG streams is the only practical solution to an operational requirement. Let's look at some applications and see how they might be handled.

MPEG Applications

MPEG is used by recording devices and it would be useful to edit segments together without decoding. On disk recorders particularly, the ability to perform edits merely by using the random access capabilities of the disk is a major convenience. At the time of writing, almost all disk systems use motion JPEG, or some other intracoding scheme where splicing is not a problem. However, there are strong incentives to use MPEG for increased recording time and/or higher image quality. MPEG's increase in efficiency comes mainly from temporal compression, so intracoding is not an option that satisfies these requirements.

The ATSC *Digital Television Standard* uses MPEG-2 encoding for video and provides for (nominally) one high-definition program or several multiplexed program streams of standard definition. Normally about 18 Mbits/s is available for video in the transmission signal. Initially it was assumed that this signal would be distributed by networks, and that television stations would switch to local commercials by switching the MPEG bit stream. The lack of production values offered to the station, and the difficulties of the switching, have contributed to the unpopularity of this scheme, and most networks are now proposing a different approach (see below).

As we move down the transmission food chain, however, there is a continued need for switching, and although neither is attractive, eventually there is no option but to decode and re-encode, or to switch the transmission signal. The most obvious example is the cable head-end. As cable systems extend the digital signal to the home, economics demand that the head-end operate in pass-through mode, without decoding the video and audio. Yet insertion of local commercials at cable head-ends is now a thriving business, and this requirement will certainly not disappear.

Most of these problems are as yet unsolved, and it is clear that there is no universal solution, but potential solutions for various applications are on the horizon.

Some Solutions

There are significant developments in the field of re-encoding, based on reusing the original coding decisions. These developments will make a decode/re-encode approach viable in many areas by drastically reducing the losses normally associated with process.

Another approach requires partial decoding of the signal to change frame types to preserve the integrity of motion compensation. In particular, B-frames may be modified so that vectors are used in only one direction. At the International Broadcasting Convention (IBC) in 1977, the European Atlantic project demonstrated switching of this type.

Conventional wisdom might suggest that the first frame of the new sequence (after a switch) should be an I-frame, requiring no references from previous frames. However, B-frames may be used if they employ only backward prediction—that is, prediction based on an anchor frame later in time. (But remember, the I-frame still must be transmitted first, even though it is for later presentation.) The advantage is that these B-frames can, at the expense of image quality, be squeezed to a very low number of bits if required, and it may be necessary to reduce the number of bits at the beginning of the second sequence to correct the VBV fullness.

Taking bits from the B-frames has two advantages. First, the B-frames are a part of the sequence that can be deprived of bits without affecting the quality of the I-frame or introducing errors that will propagate. Second, a scene change provides a substantial degree of temporal masking in the human visual system. If the low-quality B-frames occur immediately after the switch, it is most unlikely that the quality loss will be noticed.

Certainly the IBC demonstration supported this view. An example switch between two sequences used two unidirectional B-frames immediately after the cut. Examining the result frame by frame revealed that these B-frames were of very poor quality, but at normal play speed the same sequence was quite acceptable, with no apparent quality loss.

It is not certain what impact such developments will have on decisions on signal switching at cable head-ends and similar locations. If a switch can be made by decoding and re-encoding (switching at video) with no significant loss of quality, this mechanism may be attractive for many applications. It may be that, by combining the techniques of reusing coding decisions and modifying frame types, the motion detector could be eliminated from the re-encoder. This would eliminate the most significant cost element and make this approach financially attractive.

In the meantime, most workers assume that some method of direct switching of MPEG bit streams will be required for some applications. A working group of SMPTE, under the chairmanship of the author, investigated possible techniques for some two years. The real work was performed by an ad hoc group chaired by Katie Cornog of Avid, with members from many different industries. Eventually the committee produced a standard for splicing MPEG bitstreams, SMPTE 312M, published in 1999.

The standard specifies how MPEG transport streams may be constructed to permit splicing. Splice points, in-points, and out-points are inserted at the desired points in the bit stream, and messages may be sent in the bit stream to advise splicing devices of upcoming splice points.

The key to necessary VBV buffer management is to relate buffer fullness to a delay. As the "input" to VBV is a constant rate bit stream, each value of buffer fullness corresponds to a certain delay: the time it would take for the buffer to reach fullness from empty. This is known as the *splice decoding delay*. A *seamless* splice is defined as one where there is no discernible artifact when the spliced bit stream is played by a standard decoder. To achieve a seamless splice, both bit streams are constrained to have a splice decode delay equal to a defined nominal value at the splice point. (There are many other constraints, but this is the essential factor for VBV management.)

To provide for cases where a seamless splice cannot be achieved, the proposed standard also defines a *nonseamless* splice where VBV may undergo a controlled underflow. A well-designed decoder will play a seamless splice with minimal disturbance, usually a freeze of the video for a few frames.

Since the publication of the SMPTE Standard, it has become increasingly apparent that for most applications, seamless splicing is not practical. The bitstreams require considerable preparation to include splice points and to manage the buffers in the vicinity. It appears that other techniques, making use of mezzanine compression, transparent decode/encode, and/or GOP modification will provide practical solutions for studio applications.

For lower-budget applications, such a local commercial insertion in cable systems, the cable industry has developed economical approaches that will likely be incorporated into a future version of the SMPTE Standard.

Mezzanine Compression Systems

It is clear from the preceding sections that switching compressed signals is not simple. Furthermore, production effects other than cuts are needed frequently at intermediate points. For example, a local television station must switch local commercials to a network feed, but also must key station identification and alert messages over the video and insert voiceovers in the audio. Some stations perform more sophisticated production effects such as squeezing the network image as the credits roll to allow a "promo" for the next program. At present none of these effects can be performed on the compressed signal—to produce spatial effects we need the signal in the spatial domain, not the frequency domain. It is possible that solutions may be found for the simple cases, but in general it is safe to assume that production effects require a baseband (uncompressed) signal, whether in video or audio.

This presents a problem. Although we have seen that identical compression systems concatenate almost losslessly, this is not true if the signal is changed between decoding and re-encoding. In the case of MPEG delivery systems, it is fair to assume that insufficient headroom exists to be able to decode, perform effects, and re-encode without substantial degradation.

The issue arises at many points in a television broadcast chain, in contribution and distribution channels, and in the TV station when recording devices are encountered. Most standard-definition digital video recorders use some form of compression, and it is highly probably that compressed recording will dominate the world of high definition. How do we design a viable system that includes compression at so many points, and where there is a need for production effects at many different points in the chain?

The concept of *mezzanine*, or in-between, compression is evolving to meet these needs for broadcast systems based on the ATSC Digital Television Standard. (To quote a colleague, a mezzanine is the floor in a hotel you can never find, and where your meeting is!)

Requirements for a mezzanine system cover quite a large range. At one extreme, the mezzanine signal must be carried over a single satellite transponder (around 45 Mbits/s maximum) as part of the network distribution, and must have sufficient quality overhead to permit decoding, production effects, and re-encoding to the ATSC emission standard (approximately 18 Mbits/s for video).

At the other extreme, compression used for recording, and possibly in-plant distribution, must be sufficiently robust to permit a significant number of decode/re-encode operations (perhaps 6 to 10) without noticeable degradation. It is also required that this signal be capable of simple switching and editing, so intra-only coding should be used, preferably with a fixed number of bits per frame. Color encoding should be at least 4:2:2 for studio effects.

From the discussion of concatenation, we can draw some conclusions about the required characteristics of the mezzanine systems. Working backward from the emission system defined, we must see the same algorithm (DCT), the same DCT block placement, and the same quantization strategy.

Obviously, there is a scheme that meets these requirements—MPEG-2 itself—but as published there is no MPEG profile/level combination that meets all these needs. Because of the urgent need for a working document, SMPTE is in the process of standardizing a new operating point, the 4:2:2 profile at high level (4:2:2@HL) permitting bit rates up to 300 Mbits/s. Table 19-1 shows some examples of how this might be used. (SDTI is *serial data transmission interface*—a 1998 SMPTE standard for transmission of packetized data over serial digital links in the studio.)

TABLE 19-1

Possible scenarios for the use of MPEG for pre-transmission encoding of high-definition TV signals.

Application	Possible bit rates	Possible GOP structures	Possible circuits
Contribution (low delay)	155 Mbits/s	I only, or IP or IB	OC-3 fiber
Contribution (other)	60-80 Mbits/s	Long GOP if delay unimportant, otherwise balance between quality and delay requirements	8-PSK satellite
In-plant	200-270 Mbits/s	I only	Recorders and SDTI
Distribution	40-45 Mbits/s	Same as emission standard	QPSK satellite or DS-3

Similar issues arise with audio distribution, but there are some different twists. Like any compression system, AC-3 (Dolby Digital) can be decoded and re-encoded, but this is not desirable when compression has been performed on the final delivery rate. However, production requirements are similar—a station needs to perform voiceovers, and these cannot be done with audio in a compressed form. It is possible to use the same compression scheme at a higher data rate for distribution, and this is certainly one possibility. However, once the need to deviate from the transmission standard is seen, it may be profitable to explore further. In this case requirements in the studio suggest an alternative approach that could also be used for distribution.

Two problems arise in the studio. One is the number of channels required. The six channels could be carried by three AES/EBU circuits, but most recording devices will not handle this many channels. Also, it is difficult if not impossible to carry three AES/EBU circuits in perfect synchronism through all the operations and equipment of a television studio. It is probable that the signals will slip by at least 1 bit. Unfortunately 1 bit represents a phase change of some 35° at 10 kHz, and this is a significant error in a true surround-sound system. The other problem is that there is no convenient mechanism to carry the metadata (discussed in Chapter 17) that is required for the AC-3 encoder.

Dolby Laboratories has designed a new compression scheme to satisfy these requirements. The new system, Dolby-E, uses an algorithm based on that used for AC-3 and specifically designed to be benign when used ahead of AC-3. It will compress 5.1 channels into the payload of a single AES/EBU circuit, together with the metadata needed for the AC-3 encoder. Some versions of Dolby-E provide additional channels within the same bitstream. This can be particularly useful when a program needs 5.1 audio for digital transmission and stereo or DSPL for analog transmission.

Dolby-E also addresses an issue common to all other digital audio systems—the fact that the length of the frame or access unit does not relate easily to the length of the video frame. Dolby-E has versions corresponding to all the major video formats, and in each case the length of the access unit matches that of the video frame. This removes many of the problems traditionally associated with switching and editing digital audio and video.

This audio mezzanine approach has many advantages. The bit rate is acceptable for distribution, yet high enough to permit a useful number of decode/reencode operations; stations need not triple the layers of audio routing, videotape recorders will be able to record the signal, and the metadata will be delivered as part of the package.

A Glimpse into the Future

Even though MPEG-2 is just going into widespread service, great things are being planned for the future. The MPEG-4 standard is approved and applications are appearing.

MPEG-7 is well under way and, in conjunction with new technology, promises to greatly increase the real value of multimedia archives, MPEG-21 is intended to provide a complete structure to manage digital assets.

We saw that MPEG-4 can transmit a model of a face, and then animate that face to match dialog. Other researchers are working on the implementation of *agents*—models that will not only accurately mimic the delivery of a speech, but will do so in whatever language is requested.

That's one of the interesting things about compression and its allied technologies. The progress over the last few years has been dramatic, yet the more we know, the more we realize that we have only scratched the surface. I find the subject fascinating—I just hope that I have managed to convey some of my fascination with this book.

GLOSSARY

This glossary is based upon that included in the ATSC Digital Television Standard, A53, and permission to reproduce it here is gratefully acknowledged. The ATSC glossary, and some of my additions, include terms that are not strictly applicable to video compression and will not be found elsewhere in this book. Others refer to advanced topics beyond the scope of this book. All of these terms are, however, relevant to the field of digital television. Some of the definitions are those from the MPEG standards; in these cases I have retained the formal syntax (e.g., group_start_code).

8-PSK—A variant of QPSK used for satellite links to provide greater data capacity under low-noise conditions.

8 VSB—Vestigial sideband modulation with 8 discrete amplitude levels.

16 VSB—Vestigial sideband modulation with 16 discrete amplitude levels.

ACATS—Advisory Committee on Advanced Television Service.

AC-3—The audio compression scheme invented by Dolby Laboratories and specified for the ATSC *Digital Television Standard*. In the world of consumer equipment it is called *Dolby Digital*.

access unit—A coded representation of a presentation unit. In the case of audio, an access unit is the coded representation of an audio frame. In the case of video, an access unit includes all the coded data for a picture and any stuffing that follows it, up to, but not including, the start of the next access unit. If a picture is not preceded by a group_start_code or a sequence_header_code, the access unit begins with a picture_start_code. If a picture is preceded by a group_start_code and/or a sequence_header_code, the access unit begins with the first byte of the first of these start codes. If it is the last picture preceding a sequence_end_code in the bit stream, all bytes between the last byte of the coded picture and the sequence_end_code (including the sequence_end_code) belong to the access unit.

A/D—Analog-to-digital converter.

AES—Audio Engineering Society.

anchor frame—A video frame used for prediction. I-frames and P-frames are generally used as anchor frames, but B-frames are never anchor frames.

ANSI—American National Standards Institute.

asynchronous transfer mode (ATM)—A digital signal protocol for efficient transport of both constant-rate and bursty information in broadband digital networks. The ATM digital stream consists of fixed-length packets called *cells,* each containing 53 8-bit bytes—a 5-byte header and a 48-byte information payload.

ATEL—Advanced Television Evaluation Laboratory.

ATM—*See* asynchronous transfer mode.

ATSC—Advanced Television Systems Committee, the organization charged with standards documentation for the digital television system for the United States. ATSC is now an international organization with members from many countries and is working to facilitate the smooth introduction of digital television and to promote the adoption of the ATSC system in other countries.

ATTC—Advanced Television Test Center.

ATV—The U. S. advanced television system.

bidirectional pictures, B-pictures, B-frames—Pictures that use both future and past pictures as a reference. This technique is termed *bidirectional prediction.* B-pictures provide the most compression. B-pictures do not propagate coding errors, because they are never used as a reference.

bit rate—The rate at which the compressed bit stream is delivered from the channel to the input of a decoder.

bits/s—Bits per second.

block—A block is an 8×8 array of pixel values or DCT coefficients representing luminance or chrominance information.

bps—Bits per second.

byte-aligned—A bit in a coded bit stream is byte-aligned if its position is a multiple of 8 bits from the first bit in the stream.

CDTV—*See* conventional-definition television.

channel—A digital medium that stores or transports a digital television stream.

coded orthogonal frequency-division multiplex—A modified form of OFDM. A modulation scheme used for digital transmission that is employed by the European DVB system. It uses a very large number of carriers (hundreds or thousands), each carrying data at a very low rate. The system is relatively insensitive to Doppler frequency shifts and can use multipath signals constructively. It is, therefore, particularly suited for mobile reception and for single-frequency networks.

coded representation—A data element as represented in its encoded form.

COFDM—*See* coded orthogonal frequency division multiplex.

compression—Reduction in the number of bits used to represent an item of data.

conventional-definition television (CDTV)—This term is used to signify the *analog* NTSC television system as defined in ITU-R Recommendation 470. *See also* standard-definition television and ITU-R Recommendation 1125.

CRC—Cyclic redundancy check to verify the correctness of the data

D-frame—Frame coded according to an MPEG-1 mode, which uses DC coefficients only.

data element—An item of data as represented before encoding and after decoding.

DCT—*See* discrete cosine transform.

decoded stream—The decoded reconstruction of a compressed bit stream.

decoder—An embodiment of a decoding process.

decoding (process)—The process defined in the Digital Television Standard that reads an input coded bit stream and outputs decoded pictures or audio samples.

decoding time stamp (DTS)—A field that may be present in a PES packet header that indicates the time that an access unit is decoded in the system target decoder.

Digital Video Broadcasting (DVB)—A system developed in Europe for digital television transmission, originally for standard definition only, although high-definition modes have been added to the specification. DVB defines a complete system for terrestrial, satellite, and cable transmission. Like the ATSC system, DVB uses MPEG-2 compression for video, but it uses MPEG audio compression and COFDM modulation for terrestrial transmission.

discrete cosine transform (DCT)—A mathematical transform that can be perfectly undone and which is useful in image compression.

Dolby Digital—*See* AC-3.

DTS—*See* decoding time stamp.

DVB—*See* Digital Video Broadcasting.

DVCR—Digital video cassette recorder.

ECM—*See* entitlement control message.

editing—A process by which one or more compressed bit streams are manipulated to produce a new compressed bit stream. Conforming edited bit streams are understood to meet the requirements defined in the Digital Television Standard.

EIA—Electronic Industries Association.

elementary stream (ES)—A generic term for one of the coded video, coded audio, or other coded bit streams. One elementary stream is carried in a sequence of PES packets with one and only one stream_id.

elementary stream clock reference (ESCR)—A time stamp in the PES from which decoders of PES may derive timing.

EMM—*See* entitlement management message.

encoder—An embodiment of an encoding process.

encoding (process)—A process that reads a stream of input pictures or audio samples and produces a valid coded bit stream as defined in the Digital Television Standard.

entitlement control message (ECM)—Private conditional access information that specifies control words and possibly other stream-specific, scrambling, and/or control parameters.

entitlement management message (EMM)—Private conditional access information that specifies the authorization level or the services of specific decoders. They may be addressed to single decoders or groups of decoders.

entropy coding—Variable-length lossless coding of the digital representation of a signal to reduce redundancy.

entry point—A point in a coded bit stream after which a decoder can become properly initialized and commence syntactically correct decoding. The first transmitted picture after an entry point is either an I-picture or a P-picture. If the first transmitted picture is not an I-picture, the decoder may produce one or more pictures during acquisition.

ES—*See* elementary stream.

ESCR—*See* elementary stream clock reference.

event—A collection of elementary streams with a common time base, an associated start time, and an associated end time.

Federal Communications Commission (FCC)—The government agency responsible for (among other things) the regulation of spectrum utilization in the United States, and the body that licenses radio and television broadcast stations.

field—For an interlaced video signal, a field is the assembly of alternate lines of a frame. Therefore, an interlaced frame is composed of two fields, a top field and a bottom field.

forbidden—This term, when used in clauses defining the coded bit stream, indicates that the value shall never be used. This is usually to avoid emulation of start codes.

FPLL—Frequency- and phase-locked loop.

frame—A frame contains lines of spatial information of a video signal. For progressive video, these lines contain samples starting from one time instant and continuing through successive lines to the bottom of the frame. For interlaced video a frame consists of two fields, a top field and a bottom field. One of these fields will commence one field later than the other.

group of pictures (GOP)—One or more pictures in sequence.

high-definition television (HDTV)—High-definition television has a resolution of approximately twice that of standard definition television in both the horizontal (H) and vertical (V) dimensions and a picture aspect ratio (HxV) of 16:9. ITU-R Recommendation 1125 further defines "HDTV quality" as the delivery of a television picture that is subjectively identical with the interlaced HDTV studio standard.

high level—A range of allowed picture parameters defined by the MPEG-2 video coding specification, which corresponds to high-definition television.

Huffman coding—A type of source coding that uses codes of different lengths to represent symbols that have unequal likelihood of occurrence.

IEC—International Electrotechnical Commission.

intracoded pictures, I-pictures, I-frames—Pictures that are coded by using information present only in the picture itself and without depending on information from other pictures. I-pictures provide a mechanism for random access into the compressed video data. I-pictures employ transform coding of the pixel blocks and provide only moderate compression.

ISO—International Organization for Standardization.

ITU—International Telecommunication Union.

JEC—Joint Engineering Committee of EIA and NCTA.

JPEG—Joint Photographic Experts' Group.

layer—One of the levels in the data hierarchy of the video and system specification.

level—A range of allowed picture parameters and combinations of picture parameters.

macroblock—The fundamental unit within MPEG for motion prediction. A macroblock includes all data relating to a picture area defined by 16×16 luminance pixels. In the advanced television system a macroblock consists of four blocks of luminance and one each C_B and C_R block.

main level—A range of allowed picture parameters defined by the MPEG-2 video coding specification with maximum resolution equivalent to ITU-R Recommendation 601.

main profile—A subset of the syntax of the MPEG-2 video coding specification that is expected to be supported over a large range of applications.

motion vector—A pair of numbers that represent the vertical and horizontal displacement of a region of a reference picture for prediction.

MP@HL—Main Profile at High Level.

MP@ML—Main Profile at Main Level.

MP3—An abbreviation for MPEG audio, Level III. A popular format for compressing music files for distribution over the Internet.

MPEG—Refers to standards developed by the ISO/IEC JTC1/SC29 WG11, Moving Picture Experts Group. MPEG may also refer to the group.

MPEG-1—Refers to ISO/IEC standards 11172-1 (Systems), 11172-2 (Video), 11172-3 (Audio), 11172-4 (Compliance Testing), and 11172-5 (Technical Report).

MPEG-2—Refers to ISO/IEC standards 13818-1 (Systems), 13818-2 (Video), 13818-3 (Audio), and 13818-4 (Compliance).

NCTA—National Cable Television Association.

NTSC—National Television Systems Committee. The committee that defined 525-line monochrome television in the 1940s and the U.S. standard-definition color television system in the 1950s.

OFDM—Orthogonal frequency division multiplex. *See also* COFDM.

pack—A pack header followed by zero or more packets; it is a layer in the system coding syntax.

packet—A header followed by a number of contiguous bytes from an elementary data stream; it is a layer in the system coding syntax.

packet data—Contiguous bytes of data from an elementary data stream present in the packet.

packet identifier (PID)—A unique integer value used to associate elementary streams of a program in a single- or multi-program transport stream.

padding—A method to adjust the average length of an audio frame in time to the duration of the corresponding PCM samples by continuously adding a slot to the audio frame.

payload—The bytes that follow the header byte in a packet. For example, the payload of a transport stream packet includes the PES_packet_header and its PES_packet_data_bytes or pointer_field and PSI sections, or private data. A PES_packet_payload, however, consists only of PES_packet_data_bytes. The transport stream packet header and adaptation fields are not payload.

PCM—Pulse-code modulation.

PCR—*See* program clock reference.

pel—*See* pixel.

PES—Packetized elementary stream.

PES packet—The data structure used to carry elementary stream data. It consists of a packet header followed by PES packet payload.

PES packet header—The leading fields in a PES packet up to but not including the PES_packet_data_byte fields where the stream is not a padding stream. In the case of a padding stream, the PES packet header is defined as the leading fields in a PES packet up to but not including the padding_byte fields.

PES stream—PES packets, all of whose payloads consist of data from a single elementary stream, and all of which have the same stream_id.

picture—Source, coded, or reconstructed image data. A source or reconstructed picture consists of three rectangular matrices representing the luminance and two chrominance signals.

PID—*See* packet identifier.

pixel—Picture element, or pel. A pixel is a digital sample of the color intensity values of a picture at a single point.

predicted pictures, P-pictures, P-frames—Pictures that are coded with respect to the nearest *previous* I- or P-picture. This technique is termed *forward prediction.* P-pictures provide more compression than I-pictures and serve as a reference for future P-pictures or B-pictures. P-pictures can propagate coding errors when P-pictures (or B-pictures) are predicted from prior P-pictures where the prediction is flawed.

presentation time stamp (PTS)—A field that may be present in a PES packet header that indicates the time that a presentation unit is presented in the system target decoder.

presentation unit (PU)—A decoded audio access unit or a decoded picture.

profile—A defined subset of the syntax specified in the MPEG-2 video coding specification.

program—A collection of program elements. Program elements may be elementary streams. Program elements need not have any defined time base; those that do have a common time base and are intended for synchronized presentation.

program clock reference (PCR)—A time stamp in the transport stream from which decoder timing is derived.

program element—A generic term for one of the elementary streams or other data streams that may be included in the program.

program-specific information (PSI)—Normative data that is necessary for the demultiplexing of transport streams and the successful regeneration of programs.

PSI—*See* program specific information.

PTS—*See* presentation time stamp.

PU—*See* presentation unit.

QAM: *See* quadrature amplitude modulation.

QPSK—Quadrature phase-shift keying, a modulation scheme for digital transmission particularly suitable for satellite links.

quadrature amplitude modulation (QAM)—A modulation scheme for digital transmission that uses both amplitude and phase information. Favored by the cable industry for digital television distribution.

quantizer—A processing step that intentionally reduces the precision of DCT coefficients.

random access—The process of beginning to read and decode the coded bit stream at an arbitrary point.

SCR—*See* system clock reference.

scrambling—The alteration of the characteristics of a video, audio, or coded data stream to prevent unauthorized reception of the information in a clear form. This alteration is a specified process under the control of a conditional access system.

SDTV—*See* standard-definition television.

slice—A series of consecutive macroblocks.

SMPTE—Society of Motion Picture and Television Engineers.

source stream—A single, nonmultiplexed stream of samples before compression coding.

splicing—Concatenation of, or switching between, two different streams of compressed data.

sprite—A static video object, usually a background, that is larger than the presentation area. Once the sprite has been transmitted, a degree of camera movement may be represented without transmission of a new background.

standard-definition television (SDTV)—Used to signify a digital television system in which the quality is approximately equivalent to that of NTSC. Also called standard digital television. *See also* conventional-definition television.

start codes—32-bit codes embedded in the coded bit stream that are unique. They are used for several purposes including identifying some of the layers in the coding syntax. Start codes consist of a 24-bit prefix (0×000001) and an 8-bit stream_id.

STD—*See* system target decoder.

STD input buffer—A first-in, first-out buffer at the input of a system target decoder for storage of compressed data from elementary streams before decoding.

still picture—A coded still picture consists of a video sequence containing exactly one coded picture, which is intracoded. This picture has an associated PTS, and the presentation time of succeeding pictures, if any, is later than that of the still picture by at least two picture periods.

streaming, streaming media—A technique for transmitting digital content over networks and links so that the content may be presented as it is received without waiting for the end of the transmission.

system clock reference (SCR)—A time stamp in the program stream from which decoder timing is derived.

system header—A data structure that carries information summarizing the system characteristics of the Digital Television Standard multiplexed bit stream.

system target decoder (STD)—A hypothetical reference model of a decoding process used to describe the semantics of the Digital Television Standard multiplexed bit stream.

time stamp—The time of a specific action, such as the arrival of a byte or the presentation of a presentation unit.

TOV—Threshold of visibility.

transport stream packet header—The leading fields in a transport stream packet up to and including the continuity_counter field.

variable bit rate—Operation where the bit rate varies with time during the decoding of a compressed bit stream.

video buffering verifier (VBV)—A hypothetical decoder that is conceptually connected to the output of an encoder. Its purpose is to provide a constraint on the variability of the data rate that an encoder can produce.

video sequence—A video sequence is represented by a sequence header, one or more groups of pictures, and an end_of_sequence code in the data stream.

BIBLIOGRAPHY

Abramson, Norman. *Information Theory and Coding*. McGraw-Hill: New York, 1963.

Bartlett, Eugene R. *Cable Television Handbook*. McGraw-Hill: New York, 2000.

Benson, K. Blair, and Fink, Donald G. HDTV: *Advanced Television for the 1990s*. McGraw-Hill: New York, 1991.

Benson, K. Blair, revised by Whitaker, Jerry. *Television Engineering Handbook*. McGraw-Hill: New York, 1992.

Benson, K. Blair, and Whitaker, Jerry. *Television and Audio Handbook for Technicians and Engineers*. McGraw-Hill: New York, 1990.

Clarke, Roger C. *Digital Compression of Still Images and Video*. Academic Press: San Diego, 1995.

De Bruin, Ronald and Smits, Jan. *Digital Video Broadcasting: Technology, Standards and Regulations*. Artech House: Norwood, Mass., 1999.

Habibi, A. "Comparison of N-th Order DPCM Encoder With Linear Transformations and Block Quantization Techniques." *IEEE Transactions on Communications*, COM-12(6): 948—956, 1971.

Hendee, William R. and Wells, Peter N.T. *The Perception of Visual Information*, 2d ed. Springer-Verlag: New York, 1997.

Hubbard, Barbara Burke. *The World According to Wavelets*. 2d ed. A.K. Peters, Ltd.: Wellesley, Mass., 1996.

Inglis, Andrew F. *Video Engineering*, 2d ed. McGraw-Hill: New York, 1996.

Jack, Keith. *Video Demystified*, 2d ed. HighText Publications: San Diego, 1996.

Lim, Jae S. *Two Dimensional Signal and Image Processing*. Prentice Hall: Englewood Cliffs, N.J., 1990.

Lohscheller, H. "Subjectively Adapted Image Communication System," *IEEE Transactions on Communications*, COM-32(12): 1316—22, December 1984.

Microsoft Corporation. *Inside Windows Media.* Macmillan USA: Indianapolis, 1999.

Mitchell, Joan L., Pennebaker, William B., Fogg, Chad E., and LeGall, Didier J. *MPEG Video Compression Standard.* Chapman & Hall: New York, 1996.

Netravali, Arun N., and Haskell, Barry G. *Digital Pictures.* Plenum Press: New York, 1988.

Pennebaker, William B., and Mitchell, Joan L. *JPEG Digital Image Compression Standard.* Van Nostrand Reinhold: New York, 1993.

Poynton, Charles A. *A Technical Introduction to Digital Video.* John Wiley & Sons: New York, 1996 (out of print; see superceding edition below).

Poynton, Charles. *Digital Video and HDTV: Pixels, Pictures, and Perception.* John Wiley & Sons: New York, 2001.

Pratt, William K. *Digital Image Processing.* John Wiley & Sons: New York, 1978.

Rabbani, Majid, and Jones, Paul W. *Digital Image Compression Techniques.* SPIE Optical Engineering Press: Bellingham, Wash., 1991.

Reza, F. M. *An Introduction to Information Theory.* McGraw-Hill: New York, 1961.

Robin, Michael and Poulin, Michel. *Digital Television Fundamentals: Design and Installation of Video and Audio Systems,* 2d ed. McGraw-Hill: New York, 2000.

Rogowitz, Bernice E., and Allebach, Jan P., eds. *Human Vision and Electronic Imaging,* SPIE Proceedings, Bellingham, Wash., 1996.

Schreiber, William F. *Fundamentals of Electronic Imaging Systems,* 2d ed. Springer-Verlag: New York, 1991.

Shannon, C. E. "Prediction and Entropy of Printed English," *Bell System Technical Journal,* 30(1): 50—64, January 1951.

Taylor, Jim. *DVD Demystified.* McGraw-Hill: New York, 1998.

Turin, William. *Digital Transmission Systems: Performance Analysis and Modeling.* McGraw-Hill: New York, 1999.

Watkinson, John. *The Engineer's Guide to Motion Compensation.* Snell & Wilcox Ltd.: Petersfield, U.K., 1994.

Watkinson, John. *Compression in Video and Audio.* Focal Press: Oxford, U.K., 1995.

Watkinson, John. *The Engineer's Guide to Compression.* Snell & Wilcox Ltd.: Petersfield, U.K., 1996.

Watson, Andrew B., ed. *Digital Images and Human Vision.* MIT Press: Cambridge, Mass., 1993.

Weiss, S. Merrill. *Issues in Advanced Television Technology.* Focal Press: Boston, 1996.

Welch, Terry A. "A Technique for High Performance Data Compression," *IEEE Computer,* 17(6): 8—19, June 1984.

Whitaker, Jerry C., *DTV: The Revolution in Digital Video,* McGraw-Hill: New York, 1999.

Whitaker, Jerry C. and Benson, K. Blair, eds. *Standard Handbook of Video and Television Engineering,* revised ed. McGraw-Hill: New York, 2000.

Whitaker, Jerry, ed., *Electronics Handbook.* CRC Press: Boca Raton, Fla., 1996.

INTERNET RESOURCES

There is a wealth of information available from the Internet on the subjects of video and audio compression. I have listed just a few of the key sites; those that I expect to remain stable. The best way to research the Internet is to use the links from these sites; generally these are kept up to date. To save typing, these links may be found in the top directory of the CD-ROM.

http://www.symes.tv This is where I will post updates to the book, errata, and links to new or newly discovered interesting Web sites. It also provides a mechanism for feedback on the book and related topics. If you find errors, or have suggestions for improvements, please use the Web site and let me know.

http://www.mpeg.org/index.html/ This is not the "official" MPEG site, but it is without doubt the richest source of information on almost anything connected with compression. There is even a site search engine! The site is published by MpegTV LLC, a San Francisco-based privately owned company founded in 1997, and maintained by Tristan Savatier, a long-time contributor to MPEG. Most, if not all, of the following sites may be accessed from the approximately 2,000 links indexed in mpeg.org.

http://drogo.cselt.stet.it/mpeg/ This is the official MPEG site. A visit to this site is particularly useful for those interested in the ongoing work of MPEG. There are some excellent tutorials on completed and current work.

http://www.m4if.org/ The industry forum for MPEG-4. News, information, and links to companies with MPEG-4 products.

http://www.jpeg.org/ A great deal of information on JPEG, including the continuing work of the committee.

http://www.dvpa.com/links/ Links to many sites covering DV compression.

http://www.real.com/ The home of Real Video and Real Audio; free players and trial versions of authoring programs, many sources of streaming media.

http://www.microsoft.com/windows/windowsmedia/en/default.asp Everything you ever wanted to know about Windows Media, and all the tools you need to do it. Free (at the time of writing) players, authoring software, and servers.

http://www.vide.gatech.edu/cookbook2.0/ Excellent source of information on video conferencing.

http://www.internz.com/compression-pointers.html Many links to sites addressing every aspect of compression technology.

http://www.atsc.org/ The home of the Advanced Television Systems Committee and the *ATSC Digital Television Standard* (A/53). Other documents available from this site include *Digital Audio Compression (AC-3) Standard* (A/52), and *Guide to the Use of the ATSC Digital Television Standard* (A/54) and an excellent tutorial on compression and digital transmission.

http://www.dvb.org/ The European Digital Video Broadcasting organization maintains extensive information on the DVB standards.

http://www.unik.no/˜robert/hifi/dvd/ Extensive information on *DVD*; what was *digital video disk* or digital versatile disk is now just *DVD*.

http://www.davic.org/ The Digital Audio Visual Council is an association of some 200 companies seeking interoperability in the multimedia communications world.

http://www.smpte.org/ The Society of Motion Picture and Television Engineers is an international standards organization. SMPTE publishes the studio standards that form the input to standardized compression systems and many standards relating to compression.

http://www.iec.ch http://www.iso.ch http://www.itu.ch Official sites of the standards organizations, including information on the purchase of standards documents.

http://www.dolby.com/ Many general interest and technical papers on audio systems and audio compression.

http://www.inforamp.net/˜poynton/ Charles Poynton's home page with a fascinating collection of articles on digital signal processing, color science, and other aspects of vision and television.

http://www.math.auth.gr/~axonis/studies/audio.htm A site created by Panos Stokas of the Department of Mathematics at Aristotle University in Thessaloniki, Greece. It describes a long-term exercise to find the best algorithm for audio compression, one that combines high quality at low bit rates with processing speed and portability.

http://www.sv.philips.com/newtech/ Good information on many audio and video related technologies.

ABOUT THE CD-ROM

The CD-ROM included with this book includes a number of software applications for the compression of still images, video, and audio, together with some sample source material. All applications are intended for the Microsoft Windows® operating system. All of the content is provided for educational and experimental purposes, and no license is granted to use any material or application for any commercial purpose.

The contents of the CD are listed in a readme file in the top-level directory. On Windows 2000 and Windows ME systems, this file should be opened automatically in your browser when you insert the CD. If not, open this file manually and proceed from there. All of the content is described in linked html files, so you can investigate the CD from your browser. If you do not have Version 5.5 of Internet Explorer installed, you might want to start by installing this program from the CD.

Because the content has been supplied by multiple organizations, you will find considerable variety in the organization and presentation. Generally the content from each source will be in separate directories, together with any particular applicable instructions, copyright notices, and the like.

Some of the applications are freeware (but still subject to licensing agreements); others are demonstration versions that may have limited functionality, or that may be used only for a specified time after installation. In most cases the application or the accompanying files will contain explanations, links to the company that owns the application, and instructions on how to purchase unrestricted versions (and some of the applications are very inexpensive).

I believe it is helpful to gather all this content together in a way that allows easy exploration, but I strongly suggest that if you find an application useful you should contact the manufacturer to obtain the latest version. Please be aware of the licensing conditions for each product, and respect the rights of those who own the content. In particular, please do not redistribute any of the content of the CD unless the license for that content explicitly permits redistribution.

INDEX

E

F

G

H

I

J

L

M

W

ABOUT THE AUTHOR

Peter Symes is Manager of Advanced Technology at Grass Valley Group. He is an SMPTE Fellow, a Senior Member of the IEEE, and the author of *Video Compression* (McGraw-Hill). At the time of this writing, he served as the Engineering Director, Television for SMPTE.

SOFTWARE AND INFORMATION LICENSE

The software and information on this diskette (collectively referred to as the "Product") are the property of The McGraw-Hill Companies, Inc. ("McGraw-Hill") and are protected by both United States copyright law and international copyright treaty provision. You must treat this Product just like a book, except that you may copy it into a computer to be used and you may make archival copies of the Products for the sole purpose of backing up our software and protecting your investment from loss.

By saying "just like a book," McGraw-Hill means, for example, that the Product may be used by any number of people and may be freely moved from one computer location to another, so long as there is no possibility of the Product (or any part of the Product) being used at one location or on one computer while it is being used at another. Just as a book cannot be read by two different people in two different places at the same time, neither can the Product be used by two different people in two different places at the same time (unless, of course, McGraw-Hill's rights are being violated).

McGraw-Hill reserves the right to alter or modify the contents of the Product at any time.

This agreement is effective until terminated. The Agreement will terminate automatically without notice if you fail to comply with any provisions of this Agreement. In the event of termination by reason of your breach, you will destroy or erase all copies of the Product installed on any computer system or made for backup purposes and shall expunge the Product from your data storage facilities.

LIMITED WARRANTY

McGraw-Hill warrants the physical diskette(s) enclosed herein to be free of defects in materials and workmanship for a period of sixty days from the purchase date. If McGraw-Hill receives written notification within the warranty period of defects in material or workmanship, and such notification is determined by McGraw-Hill to be correct, McGraw-Hill will replace the defective diskette(s). Send request to:

Customer Service
McGraw-Hill
Gahanna Industrial Park
860 Taylor Station Road
Blacklick, OH 43004-9615

The entire and exclusive liability and remedy for breach of this Limited Warranty shall be limited to replacement of defective diskette(s) and shall not include or extend to any claim for or right to cover any other damages, including but not limited to, loss of profit, data, or use of the software, or special, incidental, or consequential damages or other similar claims, even if McGraw-Hill has been specifically advised as to the possibility of such damages. In no event will McGraw-Hill's liability for any damages to you or any other person ever exceed the lower of suggested list price or actual price paid for the license to use the Product, regardless of any form of the claim.

THE McGRAW-HILL COMPANIES, INC. SPECIFICALLY DISCLAIMS ALL OTHER WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, ANY IMPLIED WARRANT OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Specifically, McGraw-Hill makes no representation or warranty that the Product is fit for any particular purpose and any implied warranty of merchantability is limited to the sixty day duration of the Limited Warranty covering the physical diskette(s) only (and not the software or information) and is otherwise expressly and specifically disclaimed.

This Limited Warranty gives you specific legal rights, you may have others which may vary from state to state. Some states do not allow the exclusion of incidental or consequential damages, or the limitation on how long an implied warranty lasts, so some of the above may not apply to you.

This Agreement constitutes the entire agreement between the parties relating to use of the Product. The terms of any purchase order shall have no effect on the terms of this Agreement. Failure of McGraw-Hill to insist at any time on strict compliance with this Agreement shall not constitute a waiver of any rights under this Agreement. This Agreement shall be construed and governed in accordance with the laws of New York. If any provision of this Agreement is held to be contrary to law, that provision will be enforced to the maximum extent permissible and the remaining provisions will remain in force and effect.


DEMYSTIFIED
PETER
SYMES
Mc Graw Hill
P/N 136383-1 Copyright © 2001 The McGraw-Hill Companies, Inc. All rights
621.388
S986
MRDF


CD-ROM TO ACCOMPANY
VIDEO COMPRESSIO
DEMYSTIFIED
PETER SYMES
621.388
S986
MRDF
Mc Graw Hill
P/N 136383-1 Copyright © 2001 The McGraw-Hill Companies, Inc. All rights reserved. Made in U.S.A.

waiver of any rights
and governed in accordance with the laws of
of this Agreement is held to be contrary to law, that provision will be enforced
to the maximum extent permissible and the remaining provisions will remain in force and effect.

Video/Audio $49.95 U.S.A.

"Symes's well-written explanations and heavily illustrated examples of video compression make this important topic accessible to readers at all levels. His book is recommended reading for all new members of my group."

Dr. Michael Isnardi
Head, Compression Systems, Sarnoff Corporation

TAKE THE MYSTERY OUT OF—

The core technology of digital TV ...
Digital HDTV ...
Movies on DVD ...
Videoconferencing, streaming video, video on the web ...
Delivering video on demand...
Scrunching vast image files through narrow bandwidths quickly...

THE COMPLETE GUIDE TO VIDEO COMPRESSION FOR INTERNET, HOBBY, AND PROFESSIONAL USE

Video compression is how enormous image data files get shrunk for fast transmission without significant loss of visual quality. It's the magic behind digital TV, videoconferencing, DVDs, and Internet video.

Now is the time to learn about video compression—how to do it and how to use it. Peter Symes, a video compression specialist, paints the full video compression picture in this friendly, not-too-big, not-too-technical introduction. This is not a detailed implementation book, but it does explain what happens in compression, why compression works, and what will stop it from working.

You MUST read this book if

- you want to understand the fundamental concepts of compression
- you read the MPEG Standard but didn't understand it
- you read books on MPEG and didn't understand them
- you're a Web site or multimedia developer
- you're in the audio/video business
- you're a video technology professional or IT consultant
- you produce or sell video products, or are thinking of doing so
- you cover the entertainment, Internet, or technology beat as a reporter or financial analyst
- you can't afford expensive compression courses
- you're a technovideophile
- you simply want to be up to speed on this fascinating new technology of key importance to the entertainment and communications industries

ON THE CD:
ENCODERS AND DECODERS FOR DCT, WAVELET, AND FRACTAL ALGORITHMS; VIDEO SAMPLES GALORE.

Cover design: Todd Radom

McGraw-Hill
A Division of The McGraw-Hill Companies

Visit us on the World Wide Web at
www.books.mcgraw-hill.com

P/N 136382-3 PART OF
ISBN 0-07-136324-6
90000
9 780071 363242

6