

1 Compression Issues

The decision to use compression has a significant impact on the overall cost/performance balance within television production and post production operations as it will affect quality, storage/transmission efficiency, latency, editing/switching of the compressed stream as well as error resiliency.

Compression is the process of reducing the number of bits required to represent information by removing redundancy. In the case of information content such as video and audio it is usually necessary to extend this process by removing in addition information that is not redundant but is considered less important. Reconstruction from the compressed bit-stream thus leads to the addition of distortions or "artifacts." Compression for video and audio is therefore not normally lossless.

Thus it is important to make decisions about compression in the studio that take into account the additional production processes and additional compression generations that will follow. These decisions are quite likely to be different from choices that would be made if the compression were simply done only for presentation to a human observer.

This chapter considers a wide range of compression characteristics. Compression of video and audio allows functionality not viable with uncompressed processing. Through reduction of the number of bits required to represent given program content, it makes economical the support of applications such as storage of material, transmission of a multiplicity of program elements simultaneously through a common data network, and simultaneous access to the same content by a number of users for editing and other processes.

Choices made with regard to compression techniques and parameters have significant impacts on the performance that can be achieved in specific applications. Consequently, it is most important that those choices be made with the attributes of the associated application clearly understood. The application includes not only the processing that immediately follows the compression process but also any subsequent downstream processing operations. This chapter provides users with information about compression characteristics to assist in making judgments about appropriate solutions. It further recommends approaches to be taken to facilitate interoperation to the extent possible between systems within a single family of compression techniques and between families of compression methods.

This chapter includes examples of 525- and 625-line implementations of interlaced television systems that exchange program content as bit streams, but it is clearly expected that the techniques will be extensible to systems having higher numbers of lines, progressive scanning, and other advanced features.

The chapter is subdivided into six parts and one associated Annex:

Section 1.2 introduces the topic of compression through an examination of the relationship between its various aspects and image quality. This topic is very large and is developing as new techniques are devised.

Section 1.3 considers the need for several quality levels within a single facility.

Section 1.4 considers some of the operational aspects of compression.

Section 1.5 discusses the key attributes of storage media when used in conjunction with compression.

Section 1.6 deals briefly with the complex topic of interoperability.

Section 1.7 looks at the techniques of compliance testing and the benefits they produce.

Section 1.8 lists the recommendations for the application of compression within the television production and post production chain.

Annex C treats specific applications and their uses of compression techniques.

1.1 Image Quality

Selection of compression system parameters has a significant impact on overall image quality. These compression parameter choices must be optimized to preserve image quality while at the same time fitting the image data into the available bandwidth or storage space. Different combinations of compression parameters may be best for different specific applications.

Compression system parameters which should be considered include: the underlying coding methods, the coding sampling structure, pre-processing, data rates and the group of pictures (GOP) structure used. In choosing compression system parameters, interaction between parameter choices must also be considered. Finally, special operational issues such as editing the bit-stream or splicing new content into an incoming bit-stream should be considered.

1.1.1 Coding Method

The coding method is the most fundamental of compression choices. There are four main compression methods used in the television production and distribution chain: MPEG-2 Main Profile at Main Level (MP@ML), MPEG-2 4:2:2 Profile at Main Level (4:2:2@ML), Motion JPEG, and DV. All of these coding methods are based on the Discrete Cosine Transform (DCT). All of these coding methods use normalization and quantization of the transform coefficients, followed by variable length coding.

MPEG includes motion estimation and compensation in its tool kit of techniques. This allows improved coding efficiency, with some cost penalty in memory and processing latency. Motion JPEG and DV are both frame-bound, thereby minimizing coding cost, but these frame-bound coding methods do not take advantage of the coding efficiency of inter-frame motion estimation and compensation. MPEG and DV both allow motion adaptive processing in conjunction with intra-frame processing.

1.1.2 Sampling Structure

MPEG, Motion JPEG, and DV can all be used with the 4:2:2 pixel matrix of ITU-R BT.601. MPEG and Motion JPEG can both be used with other pixel matrices, multiple frame rates, and either interlace or progressive scan. Note that the 4:2:2 matrix is subsampled from the original full-bandwidth (4:4:4) signal. The pixel matrix can be further subsampled to reduce the signal data, with 4:2:2 sampling normally used for interchange between systems.

4:2:2 systems, such as the MPEG-2 4:2:2 Profile, 4:2:2 Motion JPEG, and the DV 4:2:2 system (which operates at 50 Mbits/s), all use half the number of color difference samples per line compared to the number used in the luminance channel. 4:2:2 provides half the horizontal bandwidth in the color difference channels compared to the luminance bandwidth and full vertical bandwidth.

4:1:1 systems such as DV 525 use one quarter the number of color difference samples per line compared to the number used in the luminance channel. 4:1:1 reduces the color difference horizontal bandwidth to one quarter that of the luminance channel while maintaining full vertical bandwidth. The filters used to achieve the 4:1:1 sub-sampled horizontal bandwidths, like other horizontal filters, generally have a flat frequency response within their passbands thereby enabling translation to and from 4:2:2 with no further degradation beyond that of 4:1:1 subsampling.

4:2:0 systems such as DV 625¹ and MPEG-2 Main Profile, use half the number of color difference samples horizontally and half the number of color difference samples vertically compared to the number used in the luminance channel. 4:2:0 therefore retains the same color difference horizontal bandwidth as 4:2:2 (i.e., half that of the luminance channel) but reduces the color difference vertical bandwidth to half that of the luminance channel. 4:2:0 coding, however, generally does not provide flat frequency response within its vertical passband, thereby precluding transparent translation to the other coding forms. Consequently, systems that use 4:2:0 sampling with intermediate processing generally will not retain the full 4:2:0 bandwidth of the prior coding.

Care must be exercised in selecting compression sampling structures where different compression coding techniques will be concatenated. In general, intermixing different sub-

sampled structures impacts picture quality, so cascading of these structures should be minimized. For example, while 4:1:1 or 4:2:0 signals will have their original quality maintained through subsequent 4:2:2 processing (analogous to "bumping up" of tape formats), cascading 4:1:1 and 4:2:0 will generally yield less than 4:1:0 performance.

1.1.3 Compression Pre-processing

Video compression systems have inherent limitations in their ability to compress images into finite bandwidth or storage space. The compression systems rely on removal of redundancy in the images, so when the images are very complex (having very little redundancy), the ability to fit into the available data space may be exceeded, leading to compression artifacts in the picture. In these cases, it may be preferable to reduce the complexity of the image through other methods before compression processing. These methods are called pre-processing, and include filtering and noise reduction.

When noise is present in the input signal, the compression system must expend some bits encoding the noise, leaving fewer bits for encoding the desired image. When either motion detection or motion estimation and compensation is used, noise can reduce the accuracy of the motion processing, which in turn reduces coding efficiency. Even in compression systems which do not use motion estimation and compensation, noise adds substantial energy in high frequency components of the DCT which might otherwise be zero. This not only wastes bits on extraneous DCT components, but degrades run length coding efficiency as well.

Compression system specifications generally define only the compression functions within equipment, but do not specify the pre-processing before the compression function. An exception is the shuffling which is an inherent part of the DV system, and is not to be confused with the shuffling used for error management in digital recorders.

Since most pre-processing, such as filtering or noise reduction, is not always required, the pre-processing parameters may be selected depending on the nature of the images and the capabilities of the compression system. These choices can be pre-set or can be adaptive.

1.1.4 Data Rate

The MPEG-2 4:2:2 Profile at Main Level defines data rates up to 50 Mbits/s. Motion JPEG 4:2:2 equipment typically operates at data rates up to 50 Mbits/s. DV 4:1:1 and DV 4:2:0 operate at 25 Mbits/s. DV 4:2:2 operating at 50 Mbits/s is undergoing standardization. MPEG-2 Main Profile at Main Level is defined at data rates up to 15 Mbits/s.

Selection of the data rate for MPEG-2 4:2:2@ML is inter-related with the Group Of Pictures (GOP) structure used. Lower bit rates will typically be used with longer, more efficient, GOP structures, while higher bit rates will be used with simpler, shorter GOP structures. Intra-coded images

¹Note that the proposed SMPTE D-7 format, although based on DV coding, will use 4:1:1 sampling for both the 525- and 625-line systems.

(MPEG-2 4:2:2 Profile I-pictures only, M-JPEG, and DV), at data rates of 50 Mbits/s, can yield comparable image quality. MPEG-2 4:2:2@ML with longer GOP structures and lower data rates can provide comparable quality to shorter GOP structures at higher data rates — albeit at the expense of latency. (See Group of Pictures below.)

1.1.5 Group of Pictures

There are three fundamental ways in which to code or compress an image. The most basic is coding a field or frame with reference only to elements contained within that field or frame. This is called intra-coding (I-only coding for short). The second way in which to code an image uses motion compensated prediction of a picture (called a P-picture) from a preceding I-coded picture. Coding of the prediction error information allows the decoder to reconstruct the proper output image. The third method also uses motion compensated prediction, but allows the prediction reference (called an anchor frame) to precede and/or follow the image being coded (bi-directional or B-picture coding). The selection of the reference for each picture or portion of a picture is made to minimize the number of bits required to code the image.

Sequences of images using combinations of the three coding types, as defined by MPEG, are called groups of pictures (GOPs). Both Motion JPEG and DV use only intra-frame coding and therefore are not described in terms of GOPs.

MPEG-2 allows many choices of GOP structures, some more commonly used than others. In general, a GOP is described in terms of its total length and the repetition sequence of the picture coding types (e.g. 15 frames of IBBP). The optimal choice of GOP structure is dependent on the specific application, the data rate used, and latency considerations.

Since I-only pictures are least efficient and B-pictures are most efficient, longer GOPs with more B- and P-pictures will provide higher image quality for a given data rate. This effect is pronounced at lower data rates and diminished at higher data rates. At 20 Mbits/s, the use of long GOPs (e.g. IBBP) may prove useful, while at 50 Mbits/s shorter GOPs can provide the required quality.

Besides affecting image quality, the choice of GOP structure also impacts latency. Since a B-picture cannot be coded until the subsequent anchor picture is available, delay is introduced in the coding process. Note, however, that this delay is dependent on the distance between anchor frames, not the total length of the GOP structure. This means that a blend of the coding efficiency of long GOP structures together with the lower latency of short GOP structures can be obtained by judicious use of P-picture anchors.

1.1.6 Constant Quality vs. Constant Data Rate

Compression systems are sometimes referred to as variable

bit rate (VBR) or constant bit rate (CBR). MPEG-2 and Motion JPEG can operate in either variable bit rate or constant bit rate modes; DV operates only with constant bit rate. In practice, even those systems commonly believed to be constant data rate have data rate variation, but over shorter periods of time. Another way to characterize the compression systems is to compare constant quality with constant data rate systems.

Constant quality systems attempt to maintain a uniform picture quality by adjusting coded data rate, typically within the constraint of a maximum data rate. Since simpler images are easier to code, they are coded at lower data rates. This results in more efficient compression of simpler images and can be a significant advantage in storage systems and in non-real-time transfer of images. Constant quality operation is useful for disk recording and some tape recording systems such as tape streamers.

Constant data rate systems attempt to maintain a constant average data rate at the output of the compression encoder. This will result in higher quality with simpler images and lower quality with more complex images. In addition to maintaining a constant average data rate, some constant data rate systems also maintain the data rate constant over a GOP. Constant data rate compression is useful for videotape recording and for fixed data rate transmission paths, such as common carrier services.

Constant data rate processing will, of course, be characterized by a target data rate. Variable data rate processing can be constrained to have a maximum data rate. By ensuring that this maximum data rate is less than the target rate of the constant data rate device, constant quality coding can operate into a constant data rate environment.

1.1.7 Editing

Consideration of compression parameters relating to editing falls into two general applications categories: complex editing and simple cuts-only editing (seamless splicing). In the case of complex editing involving effects or sophisticated image processing and analysis, many of the processes will require decoding back to the ITU-R BT. 601 domain. In these cases, the coding efficiency advantage of complex GOP structures may merit consideration. In the case of cuts-only editing, however, it may be desirable to perform the edits entirely in the compressed domain using bit-stream splicing. Bit-stream splicing can be done between two bit-streams which both use the same compression method. Data rates and other parameters of the compression scheme may need to be bounded in order to facilitate splicing. Some existing compressed streams can be seamlessly spliced (to provide cuts-only edits) in the compressed domain with signals of different data rates.

Techniques for operating directly in the compressed domain are still being developed. Issues relating to editing in the compressed domain are being addressed. It has even been suggested that carrying out more complex operations in the compressed domain may be possible. It should be noted,

however, that much of the image degradation encountered in decompressing and re-compressing for special effects will similarly be encountered if those effects operations are performed directly in the compressed domain, since the relationships of the DCT coefficients will still be altered by the effects.

If all the compression coding methods used in an editing environment are well defined in open standards, systems could include multi-format decoding. Multi-format decoding would allow receiving devices to process compressed streams based on a limited number of separate compression standards, thereby mitigating the existence of more than one compression standard.

1.1.8 Concatenated Compression

To the extent possible, television systems using video compression should maintain the video in compressed form rather than employing islands of compression which must be interconnected in uncompressed form. Since several compression and decompression steps are likely, the ability to withstand concatenated compression and decompression is a key consideration in the choice of a compression system. The results of concatenated compression systems will be influenced by whether the systems are identical or involve differing compression techniques and parameters.

There are a number of factors which influence the quality of concatenated compression systems. All the systems considered here rely on the DCT technique. Anything which changes the input to the respective DCT operations between concatenated compression systems can result in the transformed data being quantized differently, which in turn means that additional image information is lost. Further, any changes which result in different buffer management will result in different quantization.

In the case of MPEG coding, any change in the alignment of GOP structure between cascaded compression steps will result in different quantization, since the P and B picture transforms operate on motion compensated image predictions, while the I picture transforms operate on the full image.

For MPEG, M-JPEG, and DV, any change in the spatial alignment of the image between cascaded compression steps will result in different quantization, since the input to any particular DCT block will have changed. Any effects or other processing between cascaded compression steps will similarly change the quantization

Concatenated compression processes interconnected through ITU-R BT.601 will have minimal image degradation through successive generations if the compression coding method and compression parameters, including spatial alignment and temporal alignment, are identical in each compression stage.

It is not always possible to avoid mixing compression meth-

ods and/or parameters. In some applications the total image degradation due to cascaded compression and decompression will be minimized by attempting to maintain the highest quality compression level throughout and only utilizing lower quality compression levels where occasionally necessary, such as in acquisition or in using common carrier services. For other applications, however, which must make greater use of lower quality compression levels, the best overall image quality may be maintained by returning to the higher compression quality level only where dictated by image processing requirements.

Beyond the quality issues just discussed, there are operational advantages realized by staying in the compressed domain. Faster than real-time transfers, as well as slower than real-time transfers, can be facilitated in the compressed domain. Further, some users would welcome image processing in the compressed domain as a potential means to achieving faster-than-real-time image processing.

1.2 Quality Levels

While different compression performance levels will be used in different application categories, users will attempt to minimize the total number of performance levels within their operation. Performance differences will be accompanied by differences in cost of equipment and of operation that are appropriate to the application category. For example, a typical broadcast operation might have three levels of compression quality.

The highest compression quality level, generally requiring the highest data rate, would be used in applications which require the highest picture quality and in applications which involve extensive post production manipulation. A key attribute of this quality level is the ability to support multiple generation processing with little image degradation. The highest compression quality level might therefore be used in some higher-quality production applications, but production applications which require the very highest quality will continue to use uncompressed storage and processing. The highest compression quality would also be used for critical imagery and to archive program content which is likely to be re-used in conjunction with subsequent further production processing.

A middle compression quality level would be used in applications which require good picture quality and in applications which involve some limited post production manipulation. This quality level would support a limited number of processing generations, and might be used for news acquisition, news editing, network program distribution, and local program production. The quality level would also be used to archive program content which may be re-used but is not likely to involve significant additional production processing.

A lower compression quality level would be used in applications which are more sensitive to cost than to quality. This quality level would not normally support subsequent processing but might be used for program presentation or mass storage for rapid-access browsing. The lower compression

quality would not generally be used to archive program content which might be re-used.

These examples of highest, middle, and lower compression quality levels do not necessarily correspond to any particular absolute performance categories, but rather should be taken as relative quality levels to be interpreted according to the specific requirements of a particular user's criteria. Further details on particular applications and their use of compression can be found in Annex C: Applications.

1.3 Operational Considerations

Systems of all compression performance levels must be fully functional in their intended applications. Equipment employing compression should function and operate in the same or better manner as similar analog and non-compressed digital equipment. The use of compression in any system should not impede the operation of that system.

If it is possible to select and alter compression characteristics as part of the regular operation of a compressed system, such selection and alteration should be made easy by deliberate design of the manufacturer. Variable compression characteristic systems should possess user interfaces that are easy to learn and intuitive to operate. In addition, selections and alterations made to a compressed system must not promote confusion or compromise the function and performance of systems connected to it.

More than a single compression method or parameter set might be employed in a television production facility. Where this is the case, these should be made interoperable. Compression characteristics used in the post production process must concatenate and inter-operate with MPEG-2 MP@ML for emission.

It is well recognized that integration of compressed video systems into complex systems must be via standardized interfaces. Even with standardized interfaces, however, signal input/output delays due to compression processing (encoding/decoding) occur. System designers are advised that this compression latency, as well as stream synchronization and synchronization of audio, video, and metadata must be considered. Efficient video coding comes at the cost of codec delay, so balance must be achieved between minimum codec delay and the required picture quality. This may be particularly important for live interview feeds, especially where the available bandwidth is low and the real time requirement is high. Compressed systems must be designed to prevent the loss of synchronization or disruption of time relationships.

Compressed signal bit streams should be designed so that they can be formatted and packaged to permit transport over as many communications circuits and networks as possible. Note that compressed bit streams are very sensitive to errors and therefore appropriate channel coding methods and error protection must be employed where necessary.

Provision should be made for selected analog vertical inter-

val information to be carried through the compression system, although not necessarily compressed with the video. Additionally, selected parts of the ancillary data space of digital signals may carry data (e.g., metadata), and provision should be made to carry selected parts of this data through a transparent path synchronously with the video and audio data.

1.4 Storage

Where a compressed video bit stream is stored and accessed on a storage medium, there may be storage and compression attributes required of the storage medium depending on the intended application.

1.4.1 Interfaces

ITU-R BT.601 is the default method of interfacing. However, as network interfaces become available with the required guaranteed bandwidth access and functionality, they will allow methods of digital copying between storage devices. Because storage devices can both accept and deliver data representing video at non-real time, the network should also allow the transfer of files at both faster and slower than true time for greater flexibility. The network interface should allow the options of Variable Bit Rate (VBR or constant quality) and Constant Bit Rate (CBR) at different transfer bit rates, and optionally transfer specialized bit streams optimized for stunt modes. This will allow a downstream device to copy a file directly from a primary device for stunt mode replay on the secondary device.

1.4.2 Data Rate Requirements

Where possible, users would prefer to record incoming data directly as files on a data storage device rather than decoding and re-encoding for storage. Since there will be different compressed video bit rates depending on the application, any network connection to the device must be capable of a wide variety of input and output data rates.

Both a tape streaming device and a disc based video server will need to be able to store variable bit rate compressed video streams. This will require an interface which can accommodate the requirements of a VBR data stream.

Furthermore, compressed video streams may be stored on a tape streamer or disc server with each stream recorded at a different average bit rate.

1.4.3 Resource Management

A tape streamer needs to be able to accept and present compressed video files over a range of values. An integrated system will need to know how to control the streaming device for an I/O channel which may have a programmable data rate rather than a constant data rate.

The storage devices should specify the range of data rates which can be recorded and played back. A disc based video server additionally has the capability of accepting multiple I/O channels. Further signaling may be necessary to ensure

that both channel bandwidth and the number of channels can be adequately signaled to the system.

1.4.4 Audio, Video and Metadata Synchronization

Many storage devices may record video data, audio data and metadata on different parts of the media or on separate media for various reasons. Synchronization information should be included to facilitate proper timing of the reconstructed data at normal playback speed.

1.4.5 VTR Emulation

Where a storage device using compressed video is intended to be or to mimic a VTR, it may implement VTR stunt modes. Such stunt modes may include viewing in shuttle mode for the purpose of identifying the content, pictures in jog and slow motion for the purpose of identifying editing points, and broadcast-quality slow-motion. However, the removal of redundancy from the video signal by compression will naturally reduce the possibilities for high quality stunt mode reproduction. Compression methods and parameters must allow stunt mode capability where required in the user's application. If the recording device is required to reconfigure the data onto the recording media to provide better stunt mode functionality, such conversion should be transparent and impose no conversion loss.

1.5 Interoperability

Interoperability can be a confusing term because it has different meanings in different fields of work. Compression systems further confuse the meaning of interoperability because of the issues of program transfers, concatenation, cascading, encoding and decoding quality, and compliance testing. Program exchange requires interoperability at three levels: the physical level, the protocols used, and the compression characteristics. This chapter considers only compression while other chapters address physical links and protocols.

Considering program transfers, the Task Force has identified that there are several types of interoperability. The first example identified is interoperation through ITU-R BT.601 by decoding the compressed signals to a raster and re-encoding them. This is the current default method and is well understood. Additional methods of interoperation are expected to be identified in the future. Further work for this Task Force is to categorize methods of interoperation, explore their characteristics and relate them to various applications, and develop possible constraints on device and system characteristics necessary to ensure predictable levels of performance sought by users for specific applications.

1.6 Compliance Testing

Interoperability between compressed video products is essential to successful implementation of systems using compression. Although inter-operation is possible via ITU-R BT.601, it is desirable to have inter-operation at the compressed level to minimize concatenation losses.

Compressed inter-operation can involve encoders and decoders using the same compression method and parameters, the same compression method with different parameters, or even different compression methods. Compliance testing is a fundamental step toward ensuring proper inter-operability.

Compliance testing can be employed by manufacturers and users of compression systems in a variety of ways. Encoders can be tested to verify that they produce valid bit-streams. Decoders can be tested to verify that a range of compliant bit-streams can be properly decoded. Applications can be tested to verify that the characteristics of a given bit-stream meet the application requirements, for example whether the amount of data used to code a picture is within specified limits. In practice, defining and generating compliance tests is more involved than applying those tests, so tests employed by manufacturers might be identical to tests employed by users.

In the case of MPEG-2, compliance testing focuses on the bit-stream attributes without physical compliance testing, since MPEG-2 does not assume a particular physical layer. A number of standardized tests are described in ISO/IEC 13818-4. The concepts for tests specified in the MPEG-2 documents may be extended to other compression methods, including Motion JPEG and DV. These compliance tests include transport stream tests, program stream tests, timing accuracy tests, video bit-stream tests, and audio bit-stream tests. The MPEG-2 video bit-stream tests include a number of tests specific to the 4:2:2 Profile at Main Level. Other tests will need to be developed.

1.7 Compression Issues – Recommendations

The deliberations that went into preparation of this chapter led to the following recommendations for the application of compression within television production, contribution and distribution:

1. Compression algorithms and transport schemes should be based on Open Standards. This implies availability of the intellectual property necessary to implement those standards to all interested parties on a fair and equitable basis. Availability in the marketplace of chip sets and/or algorithms for software encoding and decoding may give users confidence in the adoption of particular compression methods.
2. The number of compression methods and parameters should be minimized for each uniquely defined application in order to maximize compatibility and interoperability.
3. Compliance testing methods should be available for those building equipment to standards for algorithms and transport schemes and for users purchasing and installing equipment to those standards. Standards bodies should adopt standards for compliance testing meth-

- ods to support both manufacturer and user needs.
4. A single compression scheme used with different compression parameters throughout the chain should be decodable by a single decoder.
 5. To support use of more than one compression family, the development of a common (“agile”) decoder is desirable.
 6. Integration of video compression into more complex systems must be via standardized interfaces. Translating through ITU-R BT.601, i.e., decoding and re-encoding, is the default method of concatenating video signals compressed using different techniques and/or parameters, although other methods are possible.
 7. The compression scheme chosen should not preclude the use of infrastructures based on the serial digital interface (SDI) as embodied in SMPTE 259M and ITU-R BT.656.
 8. Issues related to interoperability must be further explored and standards developed to allow predictable levels of performance to be achieved in the implementation of specific applications.
 9. Bit streams carrying compressed signals should be designed so that they can be formatted and packaged for transport over as many types of communications circuits and networks as possible.
 10. Compressed bit streams are very sensitive to errors, and therefore it is recommended that appropriate channel coding methods and error protection be employed where necessary.
 11. Compression systems should be designed so that, in normal operation, signal timing relationships (e.g., audio/video lip sync) and synchronization presented at encoder inputs are reproduced at decoder outputs.
 12. Signal delays through compression processing (encoding/decoding) must be limited to durations that are practical for specific applications, e.g., live interview situations.
 13. Provision should be made for selected analog vertical interval information to be carried through the compression system, although not necessarily compressed with the video. Additionally, selected parts of the ancillary data space of digital signals may carry data (e.g., Metadata), and provision should be made to carry selected parts of this data through a transparent path synchronously with the video and audio data.
 14. The compression scheme chosen for devices that mimic VTRs should allow for the reproduction of pictures in shuttle mode for identifying content and of pictures in jog and slow motion modes for selecting edit points.
 15. Network interfaces and storage devices should provide for both Variable Bit Rate (VBR) and Constant Bit Rate (CBR) options and must be capable of supporting a wide variety of data rates, as required by particular applications.
 16. Storage devices should allow recording and playing back of streams and files as data rather than decoding to baseband for recording and re-encoding upon playback.
 17. The compression strategy chosen for standard television should be extensible to high definition applications to allow for commonality in the transitional phase.