DIGICIPHER™ HDTV SYSTEM DESCRIPTION

Submitted By:

GENERAL INSTRUMENT CORPORATION
VIDEOCIPHER DIVISION
6262 Lusk Boulevard
San Diego, California 92121

on behalf of

THE AMERICAN TELEVISION ALLIANCE

26 July 1991

Table of Contents

List of Figures

List of Tables

2. DIGICIPHER™ SYSTEM OVERVIEW

The DigiCipher™ HDTV system is an integrated system that can provide high definition digital video, CD-quality digital audio, data and text services over a single VHF or UHF channel. Bandwidth for an addressing signal that allows for conditional access of video, audio, and data services is also provided.

Figure 2-1 shows the overall system block diagram. At the HDTV station, the encoder accepts one high definition video and four audio signals and transmits one 32-QAM modulated data stream. The control computer can supply program related information such as program name, remaining times, and program rating. At the consumer's home, the DigiCipher™ HDTV receiver receives the 32-QAM data stream and provides video, audio, data, and text to the subscriber. On screen display can be used to display the program related information.

Figure 2-2 shows the block diagram of the encoder. The digital video encoder accepts YUV inputs with 16:9 aspect ratio and 1050-line interlace (1050/2:1) at a 59.94 field rate. The YUV signals are obtained from analog RGB inputs by low pass filtering, A/D conversion, and RGB-to-YUV matrix. The sampling frequency is 53.65 MHz for R,G, and B. The digital video encoder implements the compression algorithm and generates a video data stream. The digital audio encoder accepts four audio inputs and generates an audio data stream. The data/text processor accepts four data channels at 9600 baud and generates a data stream. The control channel processor interfaces with the control computer and generates control data stream.

The multiplexer combines the various data streams into one data stream at 18.24 Mbps. The FEC encoder adds error correction overhead bits and provide 24.39 Mbps of data to the 32-QAM modulator. The symbol rate of the 32-QAM signal is 4.88 MHz.

Figure 2-3 shows the block diagram of the decoder. The 32-QAM demodulator receives an IF signal from the VHF/UHF tuner and provides the demodulated data at 24.39 Mbps. The demodulator has an adaptive equalizer to effectively combat multipath distortions common in VHF or UHF terrestrial transmission. The FEC decoder corrects virtually all random or burst errors and provides the error-free data to the Sync/Data selector. The Sync/Data Selector maintains overall synchronization and provides video, audio, data/text, and control data streams to appropriate processing blocks.

The DigiCipher™ HDTV system can also support 16-QAM transmission that provides lower system threshold with a slight penalty in picture quality. The lower system threshold can be used to improve the ATV coverage area and/or to reduce the station spacing. Table 2-1 summarizes the key parameters of the DigiCipher™ HDTV system.

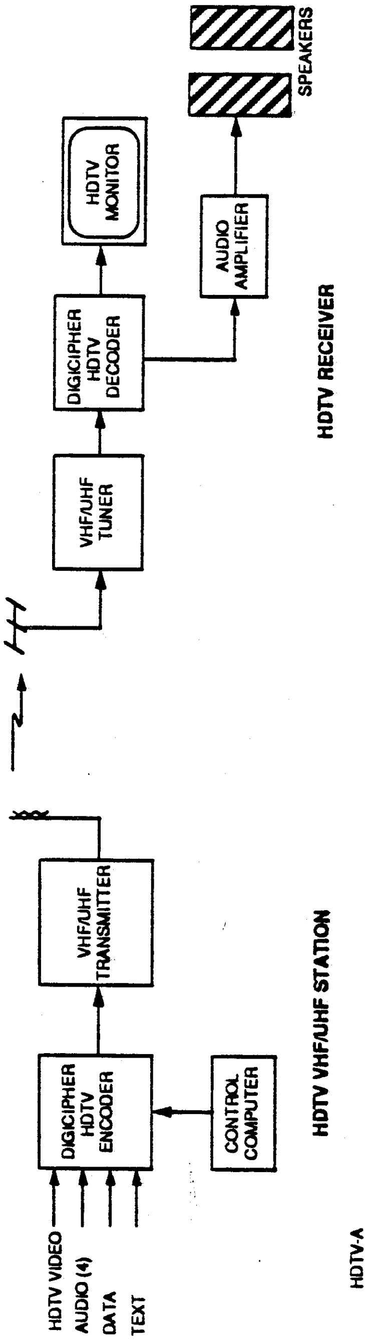

Figure 2-1. System Block Diagram

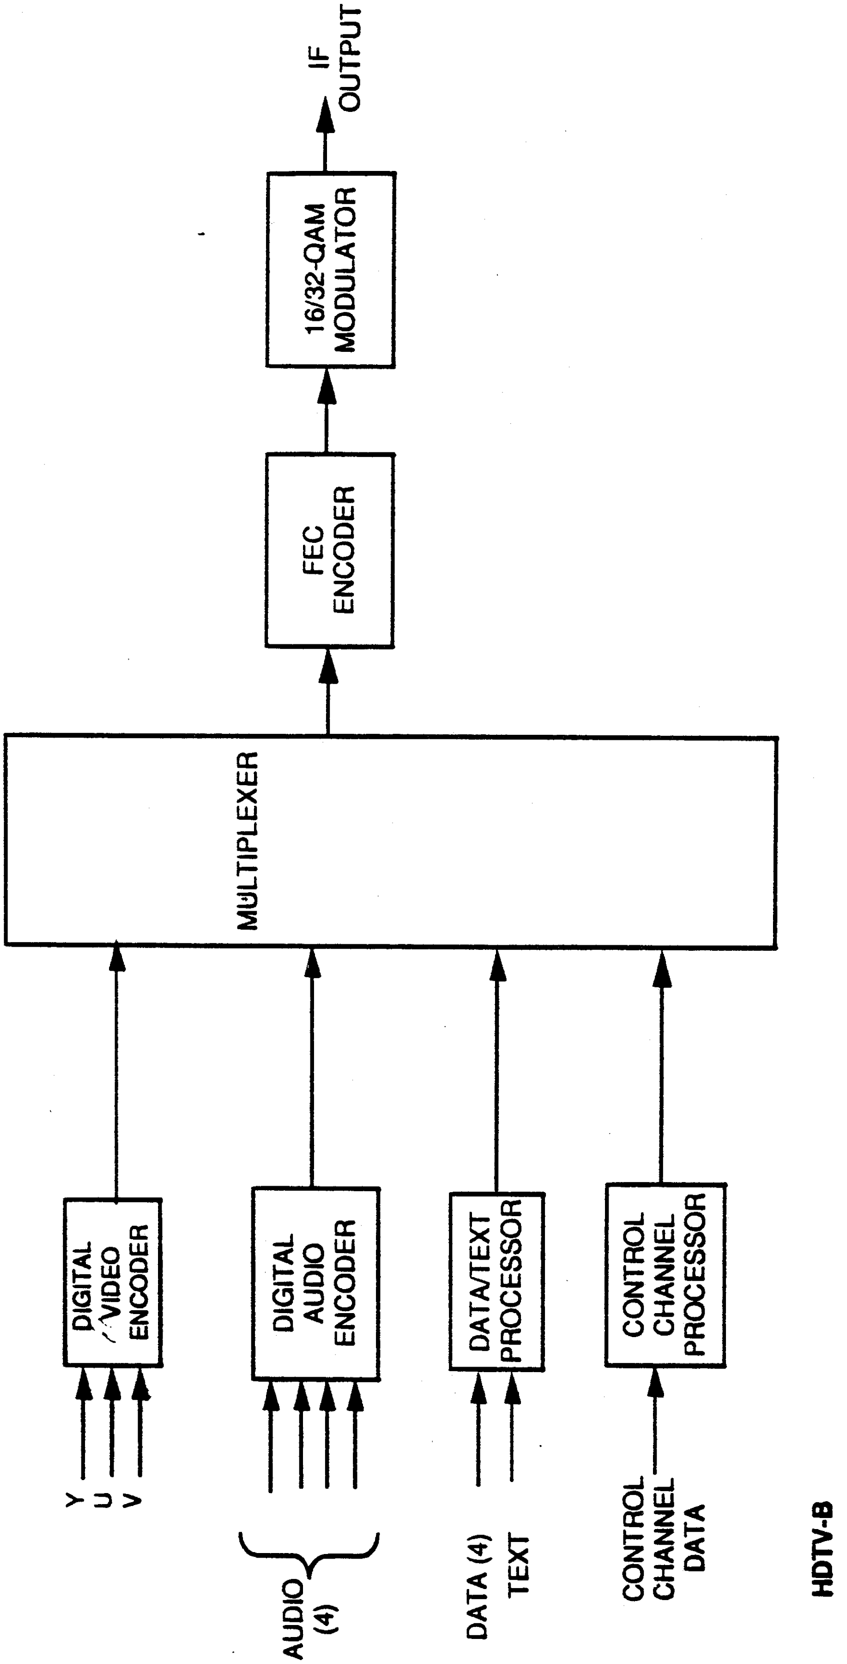

Figure 2-2. Encoder Block Diagram

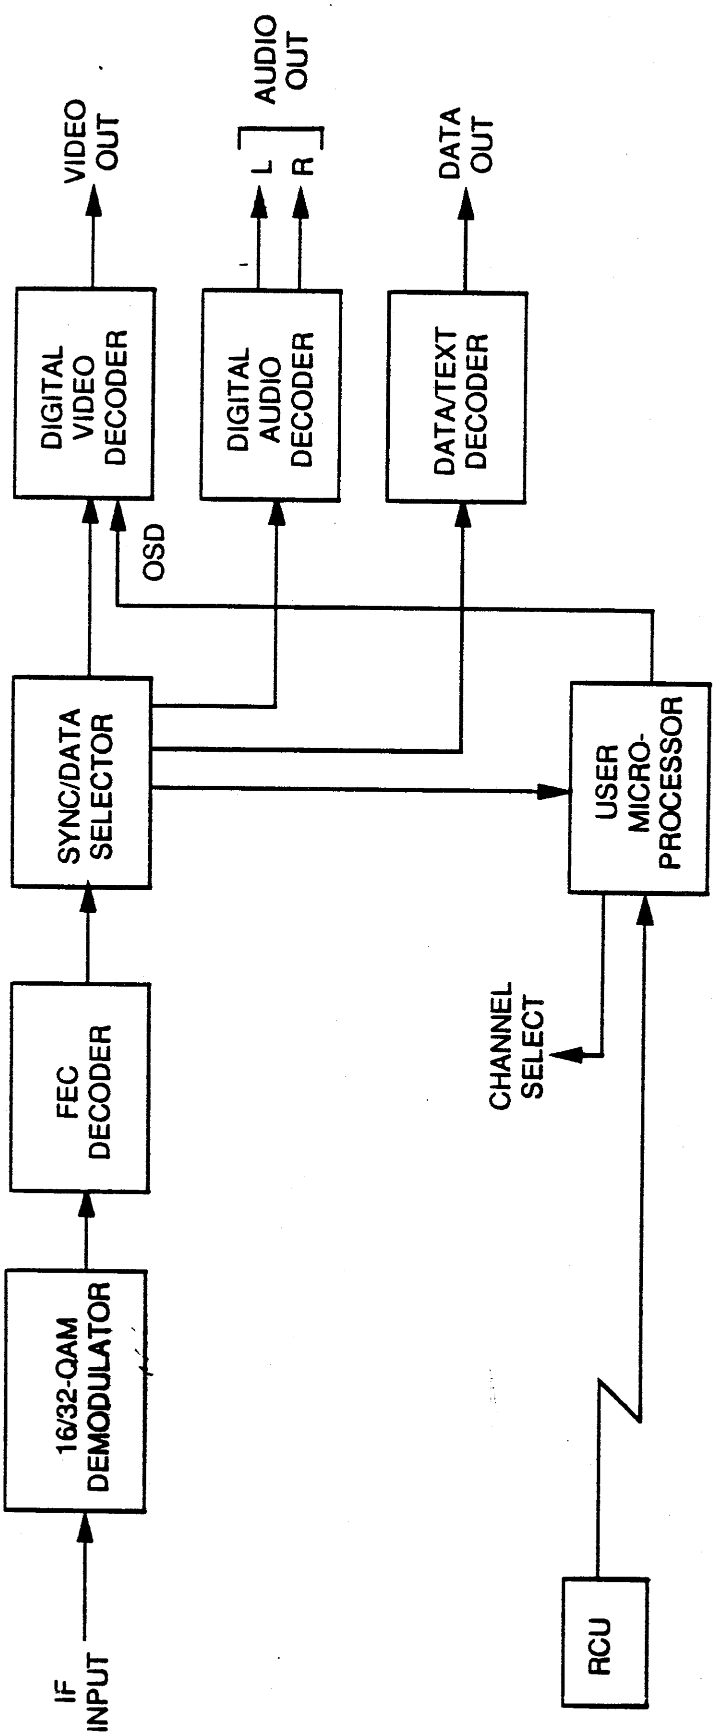

HDTV-C

Figure 2-3. Decoder Block Diagram

Table 2-1. DigiCipher™ System Parameters

Operating Mode	16-QAM	32-QAM
VIDEO		
Raster Format	1050/2:1 Interlaced	1050/2:1 Interlaced
Aspect Ratio	16:9	16.9
Frame Rate	29.97 Hz	29.97 Hz
Bandwidth		
Luminance	21.5 MHz	21.5 MHz
Chrominance	5.4 MHz	5.4 MHz
Active Pixels		
Luminance	960(V) x 1408(H)	960(V) x 1408(H)
Chrominance	480(V) x 352(H)	480(V) x 352(H)
Horizontal Resolution		
Static	660 Lines per Picture Height	660 Lines per Picture Height
Dynamic	660 Lines per Picture Height	660 Lines per Picture Height
Sampling Frequency	53.65 MHz	53.65 MHz
Colorimetry	SMPTE 240M	SMPTE 240M
Horizontal Line Time		
Active	26.24 µsec	26.24 µsec
Blanking	5.54 µsec	5.54 µsec
AUDIO		
Number of Channels	4	4
Bandwidth	20 kHz	20 kHz
Sampling Frequency	47.2 kHz	47.2 kHz
Dynamic Range	90 dB	90 dB
DATA		
Video Data	**12.59 Mbps**	**17.49 Mbps**
Audio Data	503 kbps	503 kbps
Async Data and Text	126 kbps	126 kbps
Control Channel Data	126 kbps	126 kbps
Total Data Rate	**13.34 Mbps**	**18.24 Mbps**
TRANSMISSION		
FEC Data	**6.17 Mbps**	**6.15 Mbps**
Data Transmission Rate	**19.51 Mbps**	**24.39 Mbps**
QAM Symbol Rate	4.88 MHz	4.88 MHz
Adaptive Equalizer Range	-2 to 24 µsec	-2 to 24 µsec
SYSTEM THRESHOLD		
Noise (C/iN)	**12.5 dB**	**16.5 dB**
ATV Interference (C/I)	**12.0 dB**	**16.0 dB**
NTSC Interference (C/I)	**2.0 dB**	**7.0 dB**

NOTE: The differences in the two operating modes are emphasized in **bold** letters.

3. DIGITAL VIDEO PROCESSING

The compression process can be broken down into the following different subprocesses:

1. A/D Conversion and RGB-to-YUV Matrix
2. Chrominance Preprocessor
3. Discrete Cosine Transform (DCT)
4. Coefficient Quantization
5. Huffman (Variable Length) Coding
6. Motion Estimation and Compensation
7. Integration of Motion Compensation with Intraframe Coding
8. Adaptive Field/Frame Processing
9. Motion Picture Processing
10. Rate Buffer Control

Basic block diagrams for the encoder and the decoder video processing are shown in Figures 3-1 and 3-2 respectively.

The subsequent discussions refer to certain basic picture processing elements:

- **Pixel:** An 8 bit active video sample (luminance or chrominance). Unless otherwise indicated, the term "pixel" refers to luminance pixels. Representing an image by digitized samples is generally referred to as PCM coding.

- **Block:** An image area 8 pixels horizontally by 8 pixels vertically.

- **Superblock:** An image area 4 luminance blocks horizontally by 2 luminance blocks vertically; associated with 1 chrominance block each for U and V derived from that image area.

- **Macroblock:** An image area 11 superblocks horizontally.

These elements are described further in the appropriate sections.

3.1 A/D Conversion and RGB-to-YUV Matrix

The analog R, G, B inputs are lowpass filtered and clamped before they are digitized. Figure 3-3 shows the characteristics of the lowpass filters employed. The lowpass filters are designed to provide adequate rejection of aliasing components and other spurious signals. The clamping restores proper DC-levels during the horizontal blanking interval.

The RGB-to-YUV matrix digitally converts the RGB signal into the YUV color space. The matrix conforms to the SMPTE 240M colorimetry.

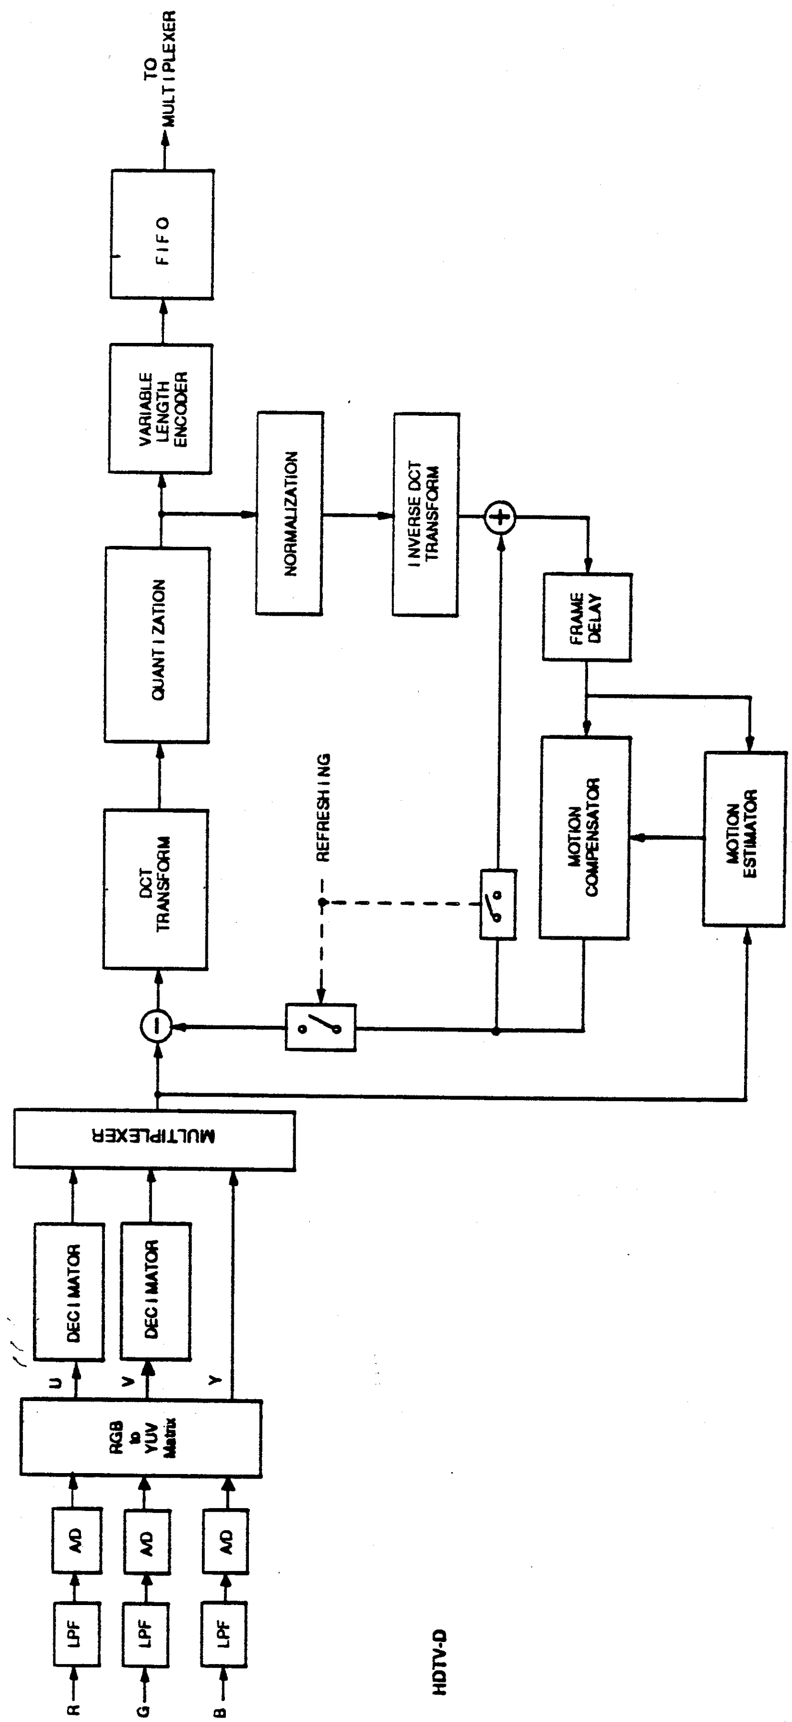

Figure 3-1. Digital Video Encoder Block Diagram

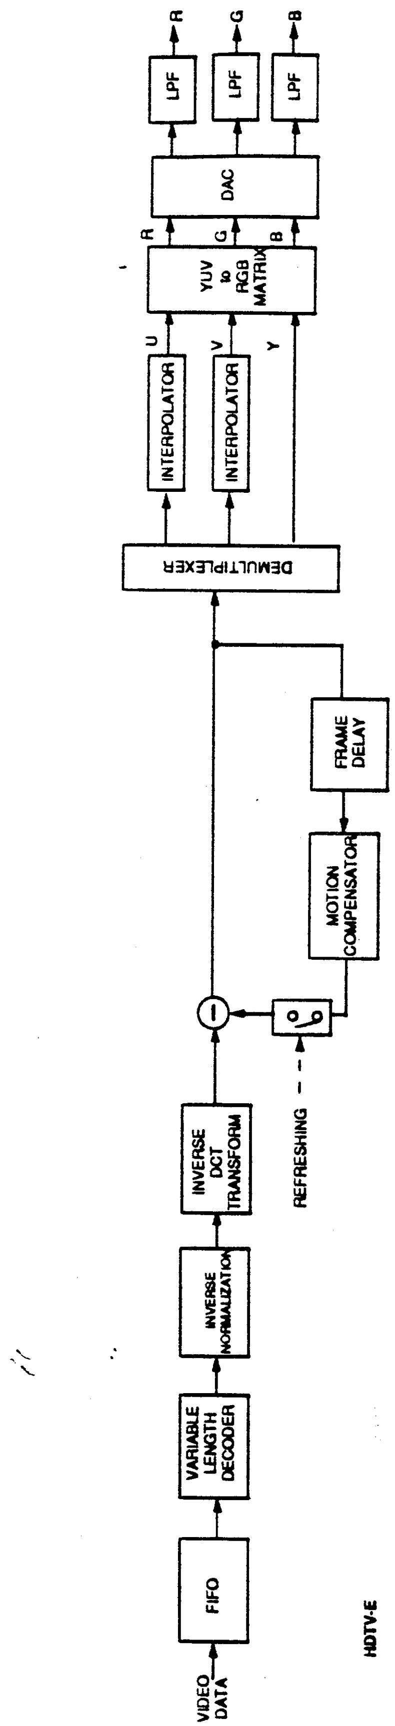

Figure 3-2. Digital Video Decoder Block Diagram

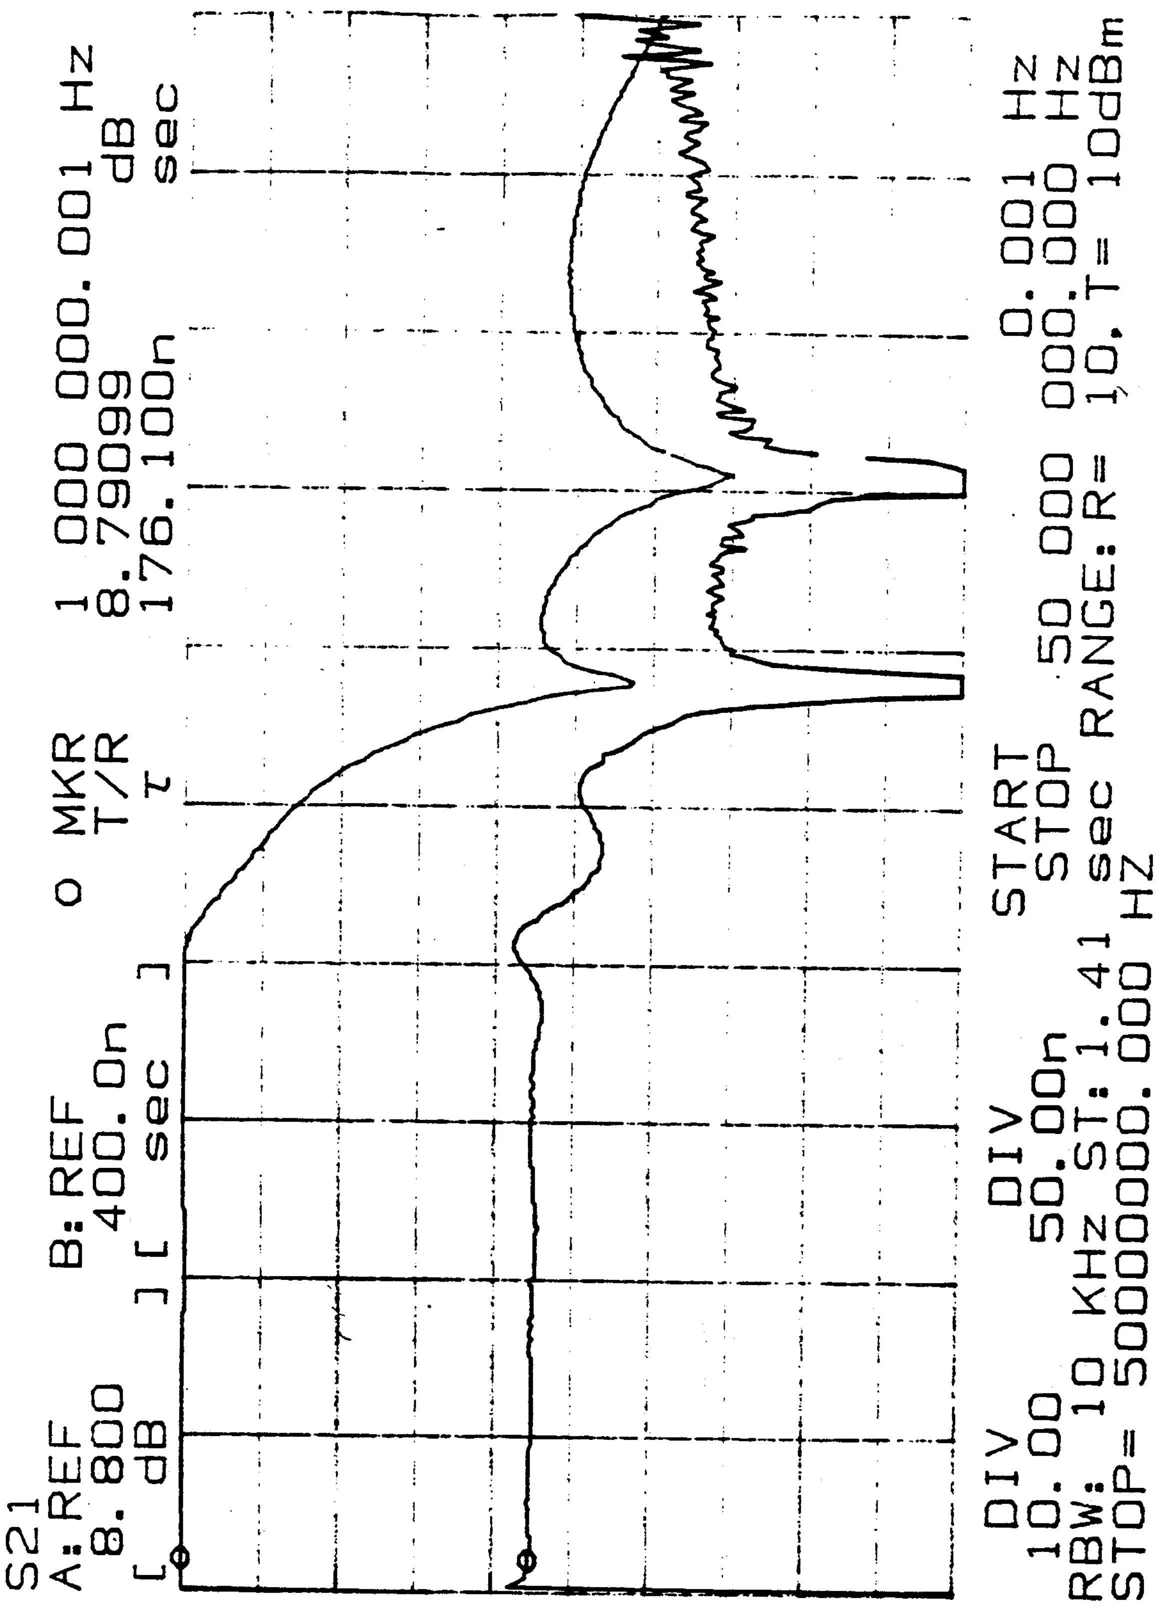

Figure 3-3. Video Lowpass Filter Characteristics

3.2 Chrominance Preprocessor

The resolution of chrominance information can be reduced relative to luminance resolution with only a slight effect on the perceived image quality. The U and V chrominance components are decimated horizontally by a factor of 4 and vertically by a factor of 2.

Horizontal decimation is performed by applying a digital FIR filter prior to subsampling. The impulse response of the FIR filter is shown in Table 3-1. Horizontal interpolation is performed at the decoder by zero-padding and apply the same filter with the gain increased by a factor of four. Vertical decimation by a factor of two is performed by discarding one of every two fields. The decoder reconstructs the interlaced signal by repeating each chrominance field twice.

Since the vertical decimation is performed across two different fields, some degradation in motion rendition occurs. In practice, however, this degradation is very difficult to detect. We are not only less sensitive to reductions in chrominance spatial resolution, but in temporal resolution as well.

The luminance signal (Y) bypasses the chrominance preprocessor, and therefore full resolution is maintained. The chrominance components are then multiplexed with the luminance component, one block at a time, and all components are then subjected to the same processing. At the decoder, the components are again separated and the chrominance signals are interpolated back to full resolution.

3.3 Discrete Cosine Transform

The Discrete Cosine Transform (DCT) transforms a block of pixels into a new block of transform coefficients. A block size of 8 x 8 has been chosen because the efficiency of the transform coding doesn't improve much while the complexity grows substantially beyond the 8 x 8 block size. The transform is applied in turn to each such block until the entire image has been transformed. At the decoder, the inverse transformation is applied to recover the original image.

If $f(i,j)$ represents pixel intensity as a function of horizontal position j and vertical position i, and $F(u,v)$ represents the value of each coefficient after transformation, then the equations for the forward and inverse transformations are

$$F(u,v) = \frac{ii4C(u)C(v)}{} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} f(i,j) \cos\frac{(2i+1)u\pi}{} \cos\frac{(2j+1)v\pi}{}$$

$$F(i,j) = \sum_{u=0}^{N-1} \sum_{v=0}^{N-1} C(u)C(v)\,F(u,v) \cos\frac{(2i+1)u\pi}{} \cos\frac{(2j+1)v\pi}{}$$

where $C(w) = \begin{cases} 1/\sqrt{2} & \text{for } w=0 \\ 1 & \text{for } w=1,2,\cdots,N-1 \end{cases}$

where N is the horizontal and vertical dimension of the block.

There are instances when the DCT is not effective in compacting the energy into a small number of coefficients. For example, if the input signal was white noise, then the image energy would be no less randomly distributed after transformation than it was the in the pixel domain.

Table 3-1. Impulse Response of Horizontal Decimation/Interpolation Filter for Chrominance

COEFFICIENT INDEX	COEFFICIENT VALUE
55,56	0.2505
53,54,57,58	0.1590
51,52,59,60	0.0000
49,50,61,62	-0.0518
47,48,63,64	0.0000
45,46,65,66	0.0297
43,44,67,68	0.0000
41,42,69,70	-0.0197
39,40,71,72	0.0000
37,38,73,74	0.0140
35,36,75,76	0.0000
33,34,77,78	-0.0101
31,32,79,80	0.0000
29,30,81,82	0.0074
27,28,83,84	0.0000
25,26,85,86	-0.0054
23,24,87,88	0.0000
21,22,89,90	0.0038
19,20,91,92	0.0000
17,18,93,94	-0.0027
15,16,95,96	0.0000
13,14,97,98	0.0018
11,12,99,100	0.0000
9,10,101,102	-0.0011
7,8,103,104	0.0000
5,6,105,106	0.0007
3,4,107,108	0.0000
1,2,109,110	-0.0004

Under such conditions, the image becomes much more difficult to compress, and in fact, cannot be compressed without introducing artifacts of some form or another. Fortunately, under such conditions, artifacts tend to be much less conspicuous than they would be under more quiet conditions. Also, such conditions are not typical of television video. Generally a high degree of horizontal and vertical correlation exists among adjacent pixels. In the next six sections, the procedure for reducing the number of bits required to represent the DCT coefficients and the effect on the appearance of the image is described.

3.4 Coefficient Quantization

Coefficient quantization is a process that introduces small changes into the image in order to improve coding efficiency. This is done by first weighting each of the DCT coefficients and then selecting 8 bits for transmission to the decoder. Once assigned, the weights for each coefficient are fixed and are never changed. The current implementation uses the weighting matrix shown in Table 3-2. The matrix weighting coefficients represent individual scaling factors for the 8x8 DCT coefficients. As before, horizontal frequency increases from left to right and vertical frequency increases from top to bottom.

Table 3-2. Weighting Table for DCT Coefficients

16	16	19	22	26	27	29	34
16	16	22	24	27	29	34	37
19	22	26	27	29	34	34	38
22	22	26	27	29	34	37	40
22	26	27	29	32	35	40	48
26	27	29	32	35	40	48	58
26	27	29	34	38	46	56	69
27	29	35	38	46	56	69	83

Each coefficient is initially represented as a 12 bit number which is then divided by the respective weighting factor. However, additional scaling may still be necessary to achieve the desired data rate. Therefore, the weighted coefficients are next divided by a quantization factor. The quantization factor is determined by the quantization level that is periodically adjusted based on scene complexity and perceptual characteristics. The quantization level ranges from 0 to 31. Maximum precision occurs at quantization level 0 and minimum precision occurs at level 30. Level 31 is reserved and indicates to the decoder that no data will be transmitted.

After a 12 bit DCT coefficient is scaled by both the weighting factor and the quantization factor, the 8 least significant bits are selected. In almost all cases, the 4 MSB's will be 0 and therefore no information is lost. However, in some cases where both the weighting and quantization factors are small, it may be necessary to clip the resulting coefficient in order to prevent an overflow or underflow from occurring.

The quantization method described above does not apply to the DC coefficient. The 8 most significant bits of the DC coefficient are always selected, independent of the quantization level.

The subjective effect of excessively quantizing the DCT coefficients is evident, not so much within the block, but when the image is viewed as a whole. The most objectionable artifact almost always tends to be the blocking effect that arises due to the individual processing of each block. The system has been designed to prevent such artifacts from being visible at normal viewing distances (3 x picture height) and make them only occasionally visible at very close viewing distances.

3.5 Huffman Coding

Quantization improves the compressibility of an image by reducing the amplitude of the transform coefficients. In order to take advantage of the result, an algorithm for assigning a variable number of bits to these coefficients is required. At this stage, a statistical coding technique is used, which unlike the quantization process, is information preserving, and therefore, does not degrade the image.

Huffman coding is an optimum statistical coding procedure capable of approaching the theoretical entropy limit, given a priori knowledge of the probability of all possible events. The encoder can generate such probability distributions and send them to the decoder prior to the transmission of a given frame. This table is then used to derive Huffman code words where relatively short code words are assigned to events with the highest probability of occurrence. The decoder maintains an identical code book and is able to match each code word with the actual event.

In order to apply Huffman coding for this application, the 8 x 8 DCT coefficients are serialized into a sequence of 64, and "amplitude/runlength" coded. Scanning the sequence of 64, an event is defined to occur each time a coefficient is encountered with an amplitude not equal to zero. A code word is then assigned indicating the amplitude of the coefficient and the number of zeros preceding it (runlength). Table 3-3 shows the length of each code word in bits. It does not include the sign bit which must be also included with each code word.

When the coefficient amplitude is greater than 16 or the number of preceding zeros is more than 15, a special code word is used to tell the decoder not to use the code book to interpret the bits that follow. Instead, the runlength is sent uncoded. The coefficient amplitude is also sent uncoded with the number of bits determined by the quantization process described previously. In addition, it is sometimes more efficient to directly code the amplitude and runlength even if it can be coded through the use of the two-dimensional table. The encoder detects these occasions and will switch to direct coding if necessary to shorten the length of the code word. A special code word is also reserved to indicate the end of a block. It is always inserted after the last non-zero coefficient. In addition, the DC coefficient is Huffman coded after it is differentially coded within a superblock. This makes use of the high correlation of DC coefficients within a macroblock and further improves the compression efficiency.

The efficiency of this coding process is heavily dependent on the order in which the coefficients are scanned. By scanning from high amplitude to low amplitude, it is possible to reduce the number of runs of zero coefficients typically to a single long run at the end of the block. As defined above, any long run at the end of the block would be represented efficiently by the "end of block" code word.

3.6 Motion Estimation and Compensation

There is a limit to the amount of compression possible by spatial processing alone. An interframe coder, however, can benefit from temporal correlation as well as spatial correlation. A very high degree of temporal correlation exists whenever there is little movement from one

Table 3-3. Number of Bits Used for Two-Dimensional Huffman Code Book

	AMPLITUDE															
RUNLENGTH	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
0	2	3	5	5	6	7	7	8	8	9	9	9	10	10	10	10
1	3	5	6	8	8	9	10	11	11	12	12	13	13	14	14	14
2	4	7	8	9	11	11	12	13	14	14	15	15	16	16	17	17
3	5	8	10	11	12	13	14	15	16	16	17	18	20	19	18	19
4	6	9	11	13	14	15	16	17	18	19	20	21	21	21	24	20
5	6	10	12	13	15	16	18	18	20	20	22	22	24	24	21	22
6	7	10	13	14	16	18	19	20	22	21	24	24	22	24	22	21
7	7	11	13	15	16	17	20	20	21	24	22	21	21	18	20	22
8	8	12	14	16	18	20	20	21	24	24	24	24	24	24	24	24
9	8	13	16	18	20	20	24	24	24	24	24	24	24	24	24	24
10	9	13	17	19	20	19	18	20	22	24	24	24	24	24	24	24
11	9	13	16	16	20	21	22	21	24	24	24	24	24	24	24	24
12	9	15	18	21	22	24	24	24	24	24	24	24	24	24	24	24
13	10	16	20	24	24	24	24	24	24	24	24	24	24	24	24	24
14	10	17	20	24	24	24	24	24	24	24	24	24	24	24	24	24
15	11	18	21	24	24	24	24	24	24	24	24	24	24	24	24	24

frame to the next. Even if there is movement, high temporal correlation may still exist depending on the spatial characteristics of the image. If there is little spatial detail, then frame-to-frame correlation remains high even at high velocities. If the image is highly detailed, however, and contains high spatial frequencies, then even slight displacements of one pixel or less can significantly reduce the amount of correlation that exists.

In the DigiCipher™ system, we compress the signal by first predicting how the next frame will appear and then sending the difference between the prediction and the actual image. A reasonable predictor is simply the previous frame. This sort of temporal differential encoding (DPCM) will perform very well if little movement occurs or if there is little spatial detail. At other times, it will be less effective and occasionally worse than if the next frame had simply been encoded without prediction (PCM).

Motion compensation is a means of improving the performance of any temporal compression scheme when movement occurs. In order to apply motion compensation, it is first necessary to determine what has moved since the previous frame and where it has moved to. If this information is known at the decoder site, then the previous frame can be shifted or displaced in order to obtain a more accurate prediction of the next frame that has yet to be transmitted. The encoder would reproduce the same prediction as the decoder and then determine the difference between the prediction and the actual image. If the movements match the model used to estimate motion and if the motion estimates are accurate and the signal is free of noise, then this error signal would, in fact, be zero.

Displacement of the previous frame can be performed on a frame, partial frame, or pixel basis. That is, a unique displacement (motion vector) could be generated for every frame, part of a frame, or every pixel respectively. The usefulness of generating a single motion vector per frame, however, is limited since it can only model simple panning of the entire image. Ideally, a unique motion vector would be generated for each pixel. However, since motion estimation is a complex process and requires knowledge of the next frame, it can only be performed at the encoder, and the overhead involved in making this per-pixel motion information available to the decoder would be excessive. Therefore, the motion estimation is performed on a partial frame basis with the area of the portion chosen to equal a superblock. The superblock has a horizontal dimension equal to 4 DCT blocks and a vertical dimension equal to 2 DCT blocks. This sizing is compatible with the 4 times horizontal subsampling and 2 times vertical subsampling of the chrominance components, thus allowing the same motion vector to be used to displace a single chrominance DCT block.

The process of displacing portions of the previous frame in order to better predict the next frame is illustrated in Figure 3-4.

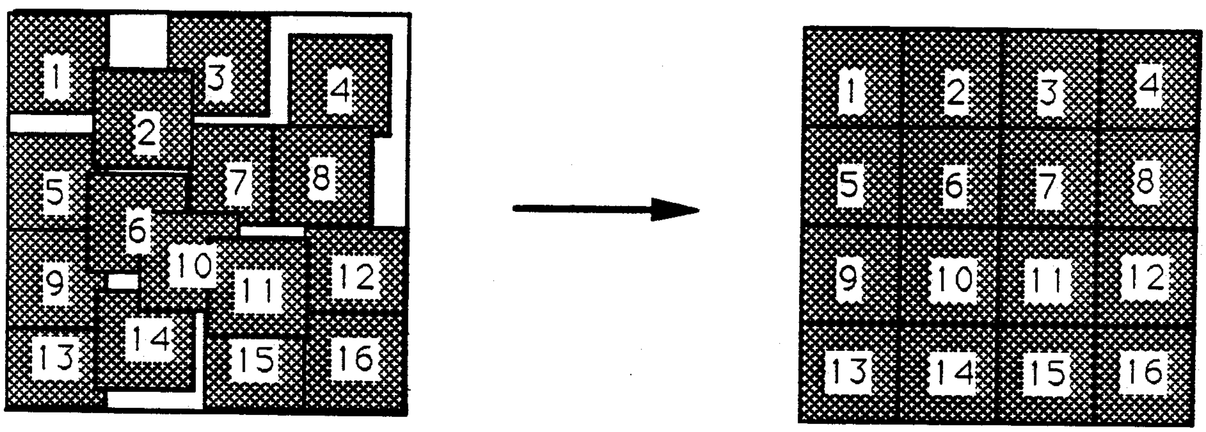

BLOCKS OF PREVIOUS FRAME USED TO PREDICT NEXT FRAME.

PREVIOUS FRAME AFTER USING MOTION VECTORS TO ADJUST BLOCK POSITIONS.

HDTV-G

Figure 3-4. Using Motion Compensation to Predict Next Frame

The search area covered by the full current frame/previous frame search algorithm is +31/-32 pixels horizontally and +7/-8 pixels vertically. The greater horizontal tracking range is due to the increased likelihood of rapid horizontal motion. These limits allow the tracking of objects moving at 0.68 picture widths per second and nearly 0.25 picture heights per second. The overhead required to send a single motion vector to the decoder is 10 bits per superblock (approximately 0.0195 bits/pixel).

3.7 Integration of Motion Compensation with Intraframe Coding

As shown in the encoder and decoder block diagrams in Figures 3-1 and 3-2 respectively, motion compensation is easily integrated into the overall system design. Instead of transform coding the image directly, an estimate of the image is first generated using motion compensation. The difference between this estimate and the actual image is then transform coded and the transform coefficients are then normalized and statistically coded as before. The second of the two frames from which the motion estimates are derived is always the previous frame *as it appears after reconstruction by the decoder.* The encoder thus must include a model of the decoder processing.

As stated previously, a lower bit rate is occasionally possible by direct PCM coding of a block instead of using motion compensation and coding the differences. Therefore, to obtain the lowest possible bit rate, the encoder determines the number of bits required for each of the two methods and then selects the method requiring the fewest bits, on a per-block basis. The overhead required to inform the decoder of the selection is thus one bit per block. For most scenes, the motion compensation rate averages between 85% to 100%. During scene changes, however, the compensation rate can drop to less than 10%.

Differential processing in general causes a basic problem for the decoder. When a decoder is tuned to a new channel, it has no "previous frame" information. Acquisition would be delayed until at least 1 PCM version of every block was received, which would result in an unbounded acquisition time. There are two basic schemes for achieving an acceptable acquisition time:

1. Every one second, all blocks of a frame can be sent in PCM form. This technique results in a DPCM-based acquisition time component of from 0 to 1 second, evenly distributed. However, the resulting large number of channel bits due to the less efficient PCM coding is difficult for the encoder buffer to handle, and may cause visible artifacts in the reconstructed image.

2. During each 0.37 second interval, process all blocks once in PCM form on a distributed basis. This technique results in a 0.37 second DPCM-based acquisition time component, but spreads the resulting increase in channel bits uniformly over time.

Note that 0.37 second parameter would imply a forced PCM block once every 11 frames, and there is a necessary but non-trivial reduction (about 9%) in the overall compression efficiency. The 0.37 second parameter can be varied to trade off acquisition time versus efficiency.

3.8 Adaptive Field/Frame Encoding

There are two options when processing interlaced signals. The first option is to separate each frame into its two fields and then process the two fields independently (Figure 3-5). The second option is to process the two fields as a single frame by interleaving the lines of corresponding even and odd fields (Figure 3-6).

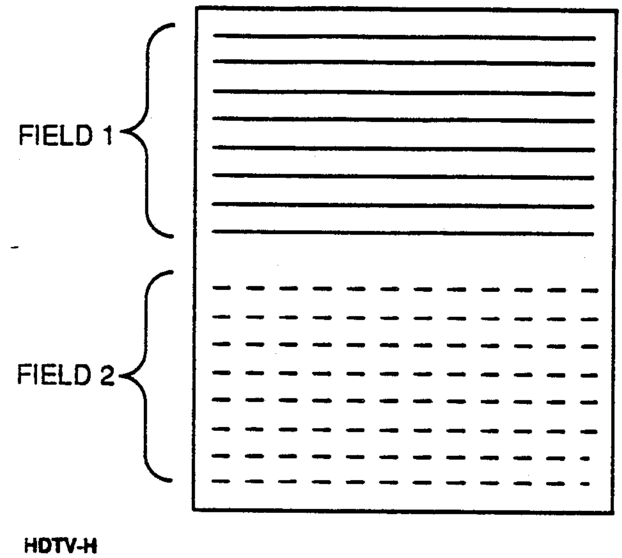

Figure 3-5. Field Processing

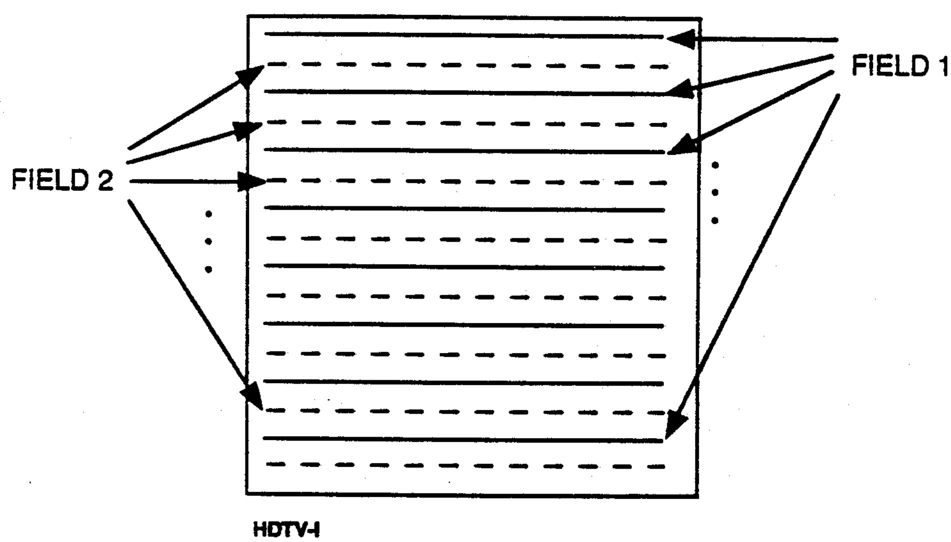

Figure 3-6. Frame Processing

Frame processing works better than field processing when there is little or no motion. Since each frame has twice as many lines or samples for a given picture height, there will be more correlation between samples and hence compressibility will be increased. Therefore, to achieve the same accuracy, field processing will require a higher bit rate, or alternatively, for equal bit rates, frame processing will achieve greater accuracy.

Similar advantages over field processing will be realized if horizontally moving features have little horizontal detail or if vertically moving features have little vertical detail. In other regions, where there is little detail of any sort, frame processing may still work better than field processing, no matter how rapidly changes occur.

Field processing generally works better than frame processing in detailed moving areas. In such cases, the interleaving of the even and odd fields would introduce spurious high vertical frequencies into the frame processing system. This would reduce the correlation between lines and therefore the effectiveness of the compression algorithm.

The DigiCipher™ HDTV System uses a novel method that has been developed to combine the advantages of both frame processing and field processing. It permits video signals to be compressed and then reconstructed with minimal degradation in motion rendition.

A selection between frame processing and field processing based on achieving minimum error has been found to be very effective. Still or slowly moving regions are rendered much more accurately than would be possible in a field-only processing system, while motion rendition is much better than would be possible in a frame-only processing system. Since the selection is made on a local basis, the system can adjust to scenes containing both moving and non-moving features.

Other simulated and less effective selection methods for the field/frame decision include simple motion detection schemes and criteria based on minimizing the number of bits used to represent the image.

Since one bit per block pair of side information is included in the encoded signal that is transmitted, decoder complexity can be significantly reduced. Field and frame processing decisions are not made at the decoder, and are instead extracted from the encoded video signal.

3.9 Motion Picture Processing

Almost all movies developed for the cinema and a significant amount of program material developed for television are initially acquired on film. Except for a few special cases, the display rate used for film is 24 frames/second, and therefore the motion rendition is significantly degraded in comparison to normal television video. Eventually, when this program material is converted to the NTSC television standard, a process called three-two pulldown is used. As shown in Figure 3-7, it involves alternating between three repetitions and two repetitions of each frame of the film.

Since the three-two pulldown process increases the number of video frames from 24 to 30 without increasing the amount of information in the signal, the first step in the source coding process is to restore the signal to its original state. Since one of every five fields is redundant, it can either be discarded, or averaged with the other identical field as shown in Figure 3-7. In the latter case, a 3 dB noise reduction results; however, its significance is questionable, since it will benefit only one of every four fields.

After the 24 frame/second signal is reconstructed at the decoder, it must be converted back to 60 fields/second before it can be displayed. This is easily accomplished by applying the three-two pulldown process once again.

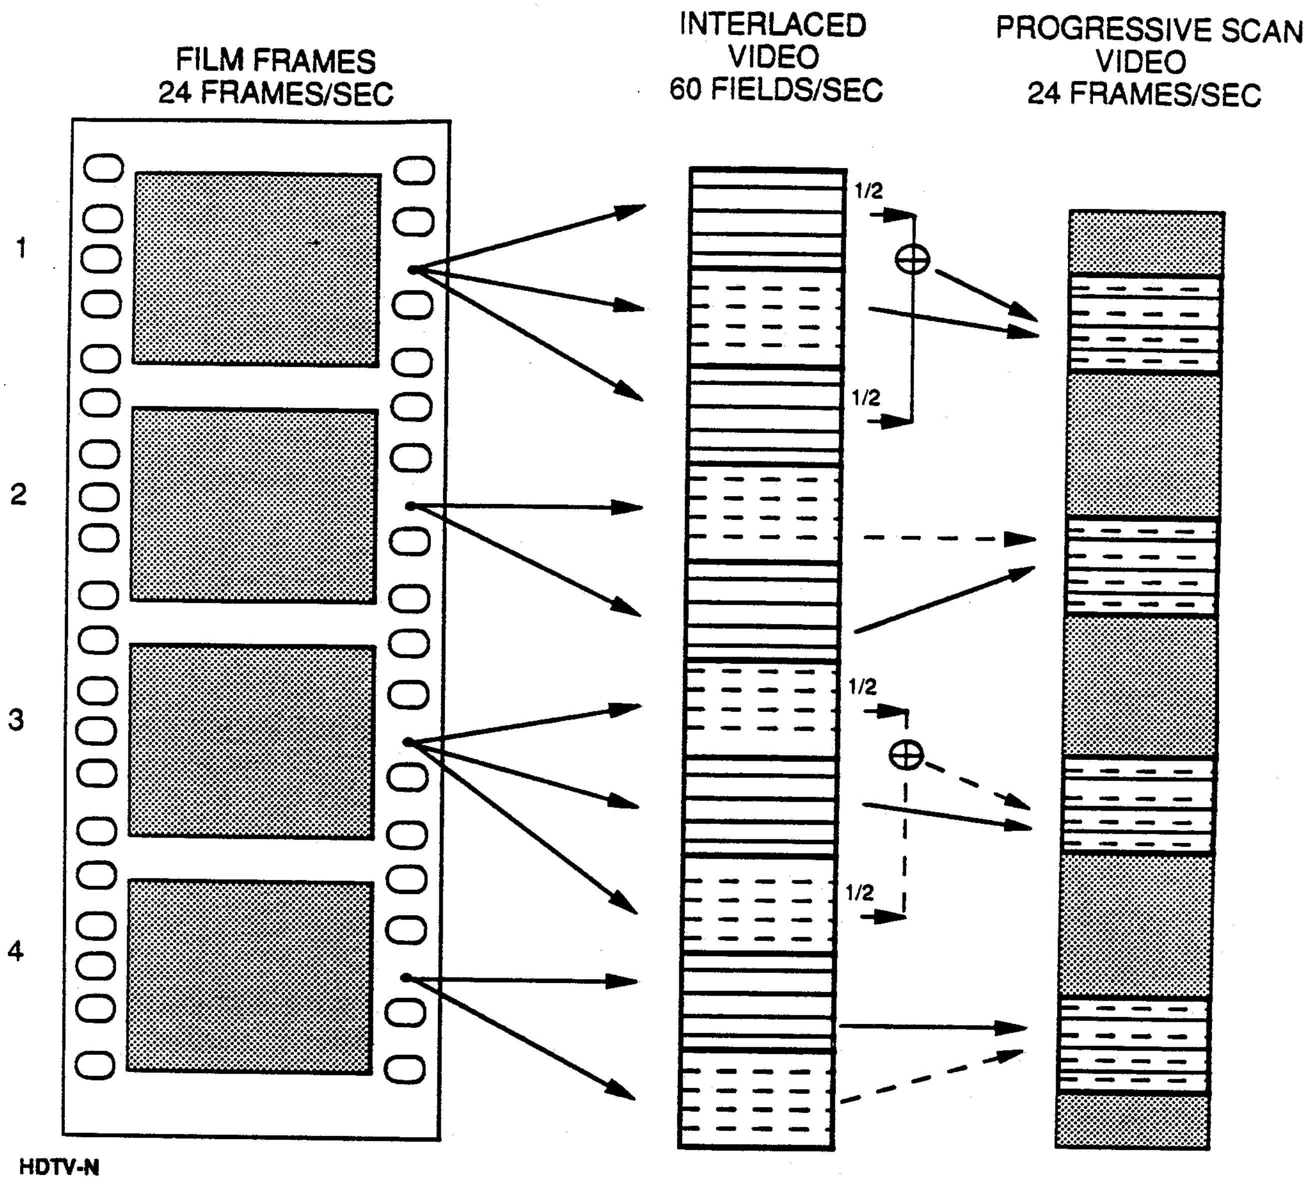

Figure 3-7. Conversion of Film to Interlaced Videoand Restoration of Progressive Scan

The DigiCipher™ System processes material shot on film in this matter to further improve performance. Other video source material is, of course, handled without going through this process.

3.10 Rate Buffer Control

Each single channel video processing section in the encoder requires a rate buffer in order to match the variable rate of the Huffman-coded data to the fixed output rate necessary for channel transmission. This rate buffer is implemented as a one frame FIFO located after the Huffman encoder. The total storage size is large enough to handle variations of plus and minus one field.

In order to prevent the video output buffer FIFO from overflowing or underflowing, the FIFO input block rate must be continuously adjusted. This is the purpose of the multi-quantization level coding structure. As the quantization level is incremented, quantization becomes coarser, blocks are shortened, and an increase in the FIFO input block rate results. As the quantization level is decremented to a minimum level of 0, finer quantization results in longer blocks, and a reduced FIFO input block rate. This adjustment has the required effect of keeping the bit rate into the FIFO relatively constant. The status of the buffer is continuously monitored, and as long as the number of stored blocks remains within a predetermined window, the quantization level will remain unchanged. If the buffer level drops below the lower threshold or rises above the higher threshold, then the quantization level will decrement or increment respectively. Fill bits may need to be inserted into the channel in order to prevent underflows during the transmission of very simple images.

4. DIGITAL AUDIO PROCESSING

The DigiCipher™ HDTV system uses Dolby Laboratories' AC-2 digital audio system that combines highly efficient data compression with professional-quality audio transparency. It uses frequency-domain signal processing in a multiplicity of narrow bands to take full advantage of psychoacoustics and noise making. The functions of the encoder and the decoder are described in the following. Refer to Appendix B for additional detail.

4.1 Digital Audio Encoder

Figure 4-1 shows the block diagram of the encoder. The encoder accepts two channels of analog input signals. They are lowpass filtered and synchronously quantized to 16-bit precision using dual high-quality A/D converters. The 16-bit PCM output data representing both audio channels is processed using the Dolby AC-2 algorithm at 24-bit precision and formatted along with the 1200 bps data into a single serial output data stream at 252 kbits/second.

The dynamic range of the Dolby AC-2 process exceeds that of the current A/D and D/A converter technology. For protection against A/D converter overload, independent safety limiters provide up to 10 dB of gain reduction. The limiter control signals can be linked to preserve image stability with stereo program signals.

The Dolby AC-2 process is designed to be tolerant of random data errors which may occur between encode and decode. In addition, it incorporates Reed-Solomon error protection for certain bits in the AC-2 format. This protection scheme provides sufficient protection for random bit errors to 1×10^{-5}.

4.2 Digital Audio Decoder

Figure 4-2 shows the block diagram of the decoder. The decoder accepts the 252 kbits/second AC-2 serial data stream and a clock signal. Complementary processing is performed at 24-bit precision to recover the two main audio signals which are passed to high-quality 16-bit D/A converters. The two audio signals and 1200 bps data are available as outputs on the rear panel.

Both encoder and decoder provide front panel accessible input and output level controls, system status indicators and LED calibration displays. An internally generated 1125 Hz calibration signal is provided for system alignment.

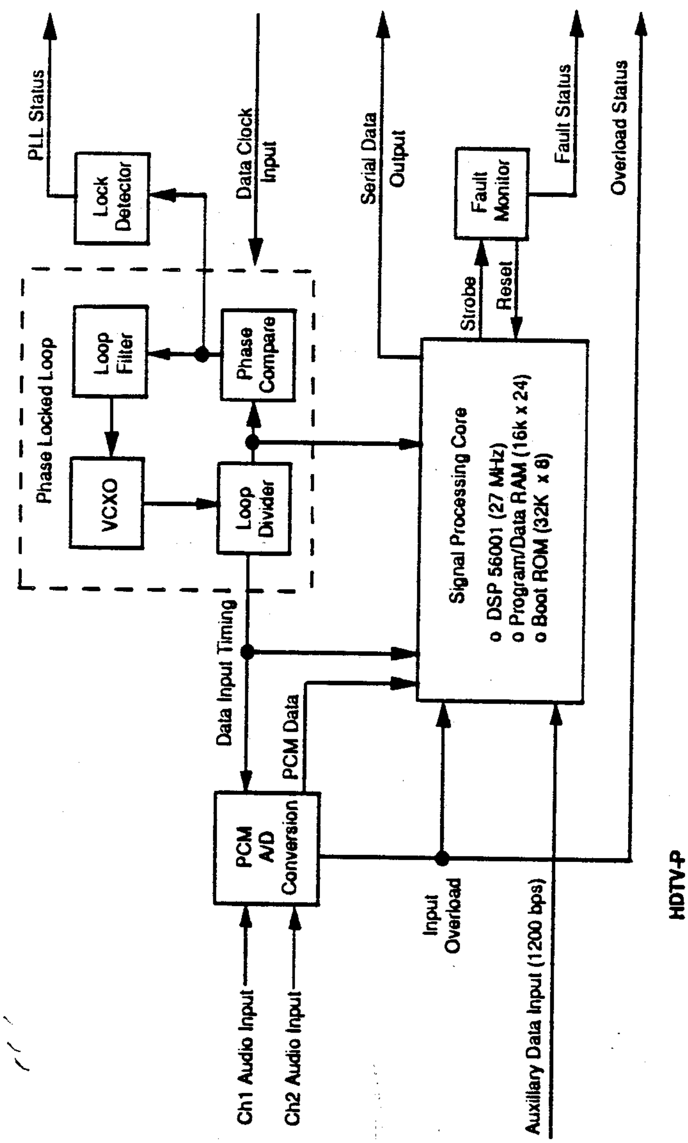

Figure 4-1. Digital Audio Encoder Block Diagram

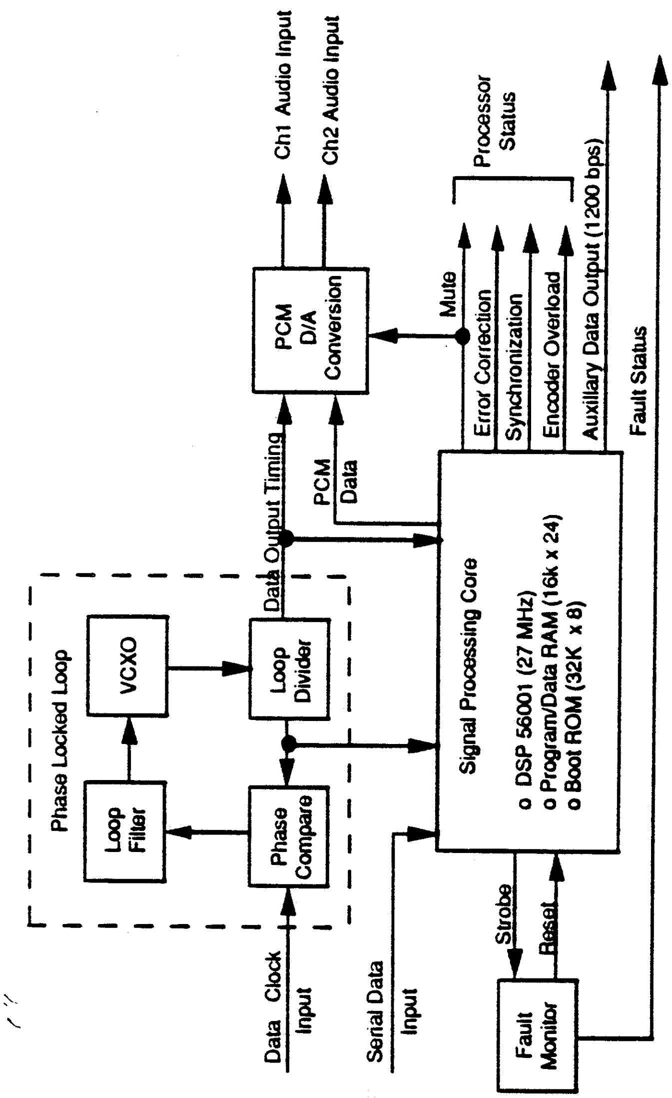

HDTV-Q

Figure 4-2. Digital Audio Decoder Block Diagram

5. DATA PROCESSING AND DATA MULTIPLEX FORMAT

The DigiCipher™ HDTV system can allocate data capacity to support various services such as additional audio channels, 5.1 channel surround sound system (Dolby AC-3 system, for example), descriptive video, special audio for the hearing impaired, expander control data, program guide, closed captioning, program mode control, conditional access, and teletext. The encoder can adaptively allocate unused data capacity back to the digital video processor to optimize the video performance. The prototype DigiCipher™ HDTV system supports the data processing described as follows.

5.1 Data Channel Processing

The DigiCipher™ transmission format has 4 bits per line, (8 bits per 2 line time), assigned to data channel capacity. This allocation of 125.87 kbps is sufficient capacity for 13 9600 baud data streams. The initial DigiCipher™ design will support 4 such data streams, with the remaining capacity reserved.

5.2 Control Channel Processing

The DigiCipher™ transmission format has 4 bits per line time, (8 bits per 2 line time) which amounts to 125.87 kbps to for the control channel (subscriber addressing). There is great flexibility allowed in message mixing within the control channel on a bit stream.

5.3 Data Multiplex Format

This section defines the data multiplex for video, audio, data, text, and control channel.

Prior to forward error coding at the encoder, each pair video line time includes 848 information bits for 16-QAM and 1,160 information bits for 32-QAM. Figure 5-1 shows the data transmission format. During lines 1 & 2, a 24-bit sequence is transmitted for maintaining overall synchronization. The system control contains a 24-bit control word, a 16-bit frame count, and a 16-bit NMP (next macroblock position) word. The NMP indicates the beginning position of the next macroblock and it is used during the initial acquisition or when the overall synchronization is lost.

6. DIGITAL TRANSMISSION

Modulation and channel coding are key elements of the DigiCipher™ system. The modulation technique must be efficient in order to send the required number of information bits reliably through a single 6 MHz VHF or UHF channel. The channel coding technique must be powerful in order to maintain a very low error rate; the more compression (source coding), the more serious the effect of a single error.

Figure 6-1 shows the basic communication system blocks, including coding, modulation, pulse shaping (transmit filtering), receive filtering, demodulation, tracking, and decoding.

Adaptive equalization is employed to handle the reflections (multipath) found in typical VHF or UHF reception. Details of the DigiCipher™ digital transmission system are described in the following sections.

6.1 Forward Error Correction

The DigiCipher™ system uses concatenated trellis coding and block coding to protect against the effect of channel errors. Concatenated coding is one of the most powerful error correction techniques and it offers error free operation at low C/N and/or low C/I condition.

Figure 6-2 and Figure 6-3 show the block diagrams of the FEC encoder and the FEC decoder, respectively. A trellis decoder (rate 3/4 for 16-QAM and rate 4/5 for 32-QAM)is used for the inner code, as it supports the use of soft decisions easily. A Reed-Solomon decoder (rate 106/116, t=5 for 16-QAM and rate 145/155, t=5 for 32-QAM) is used for the outer code, as its built-in burst error correcting capability can handle burst errors produced by the trellis decoder. Interleaver #1 is used to improve the performance of the Reed-Solomon decoder by dispersing the burst errors generated by the trellis decoder. Interleaver #2 is used to combat burst or impulsive noises such as car ignition noise. It can effectively handle 3 μsec long impulse noises.

Figure 6-4 shows measured performance of the DigiCipher™ transmission system. Notice that the system threshold is approximately 12.5 dB C/N for 16-QAM and 16.5 dB C/N for 32-QAM where the noise power is measured over a 5 MHz bandwidth. At this threshold there will be one uncorrected error per minute and the video degradation is not quite perceptible because the digital video decoder almost always detect uncorrected errors and provide error concealment.

6.2 Modulation

The modulation selected for digital transmission over the VHF or UHF channel is 16/32-QAM at 4.88 Msps. The 16-QAM has extremely low threshold thus providing wider coverage while the 32-QAM provides higher video quality with 4 dB higher threshold. Consumer HDTV receivers can be easily designed to accept either signals and automatically detect the transmitting mode. Therefore, the choice between the two operating modes can be left to local broadcasters as a tradeoff between the coverage versus video quality.

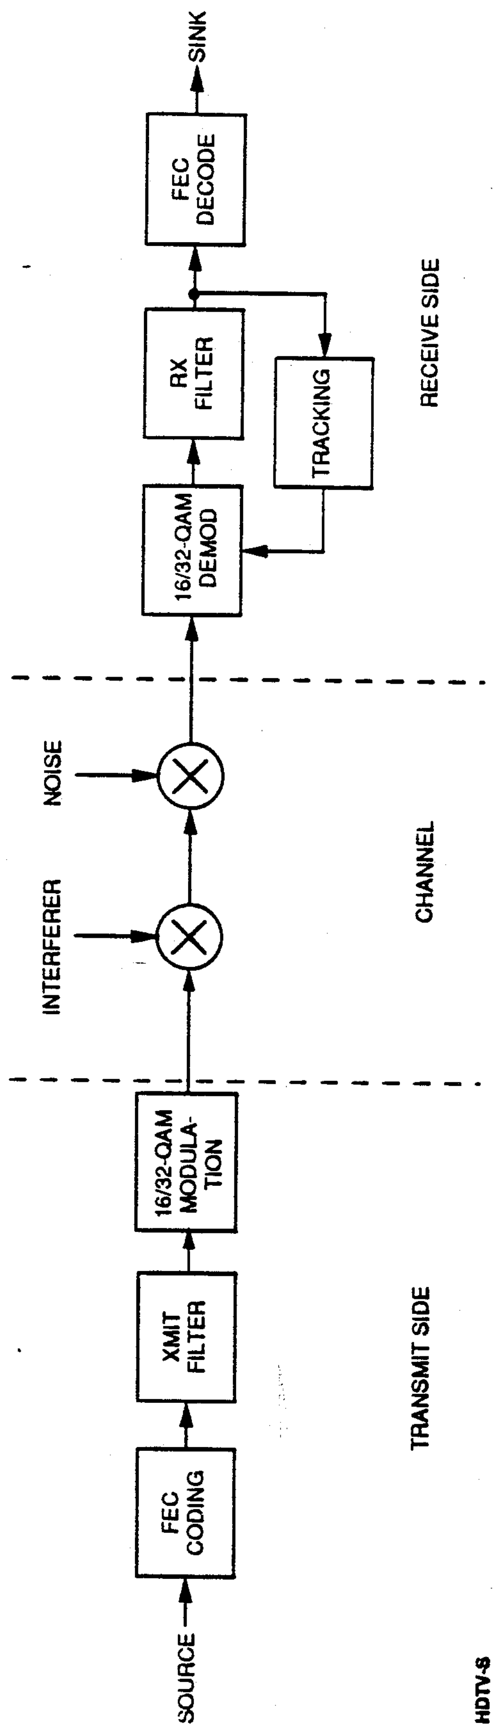

Figure 6-1. Communication System Blocks

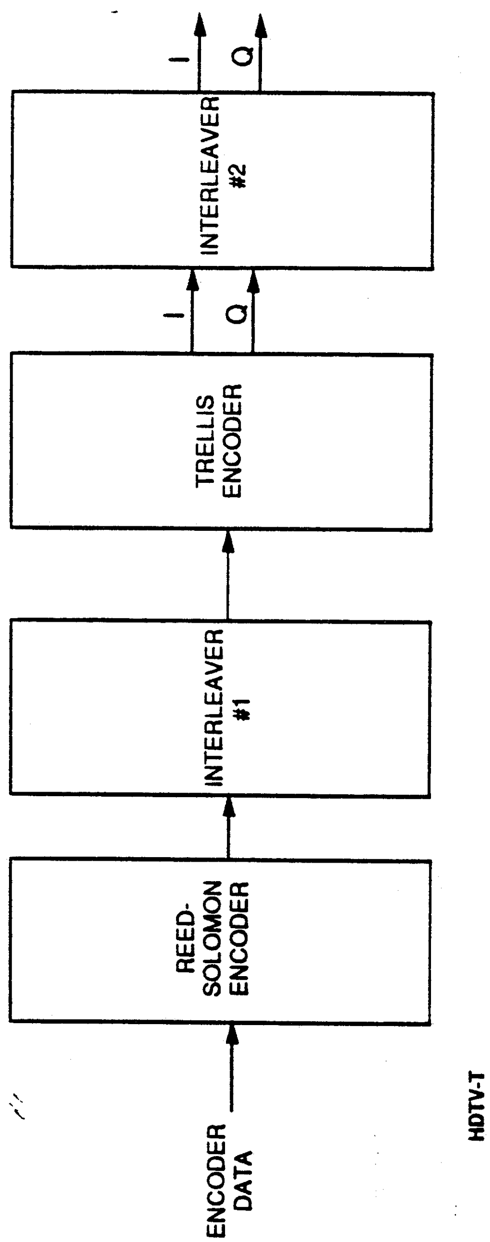

Figure 6-2. FEC Encoder

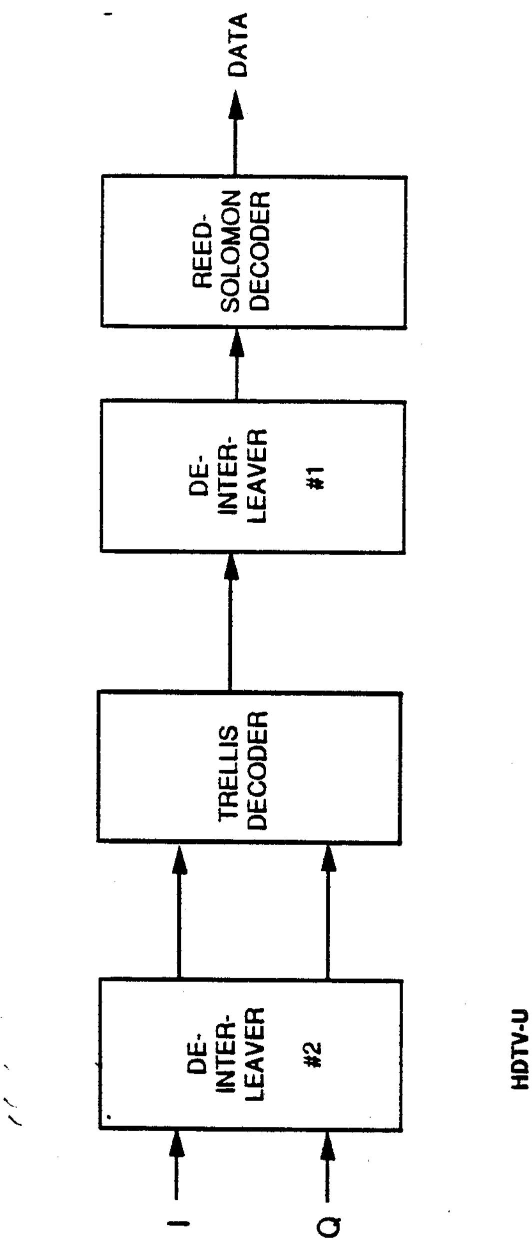

Figure 6-3. FEC Decoder

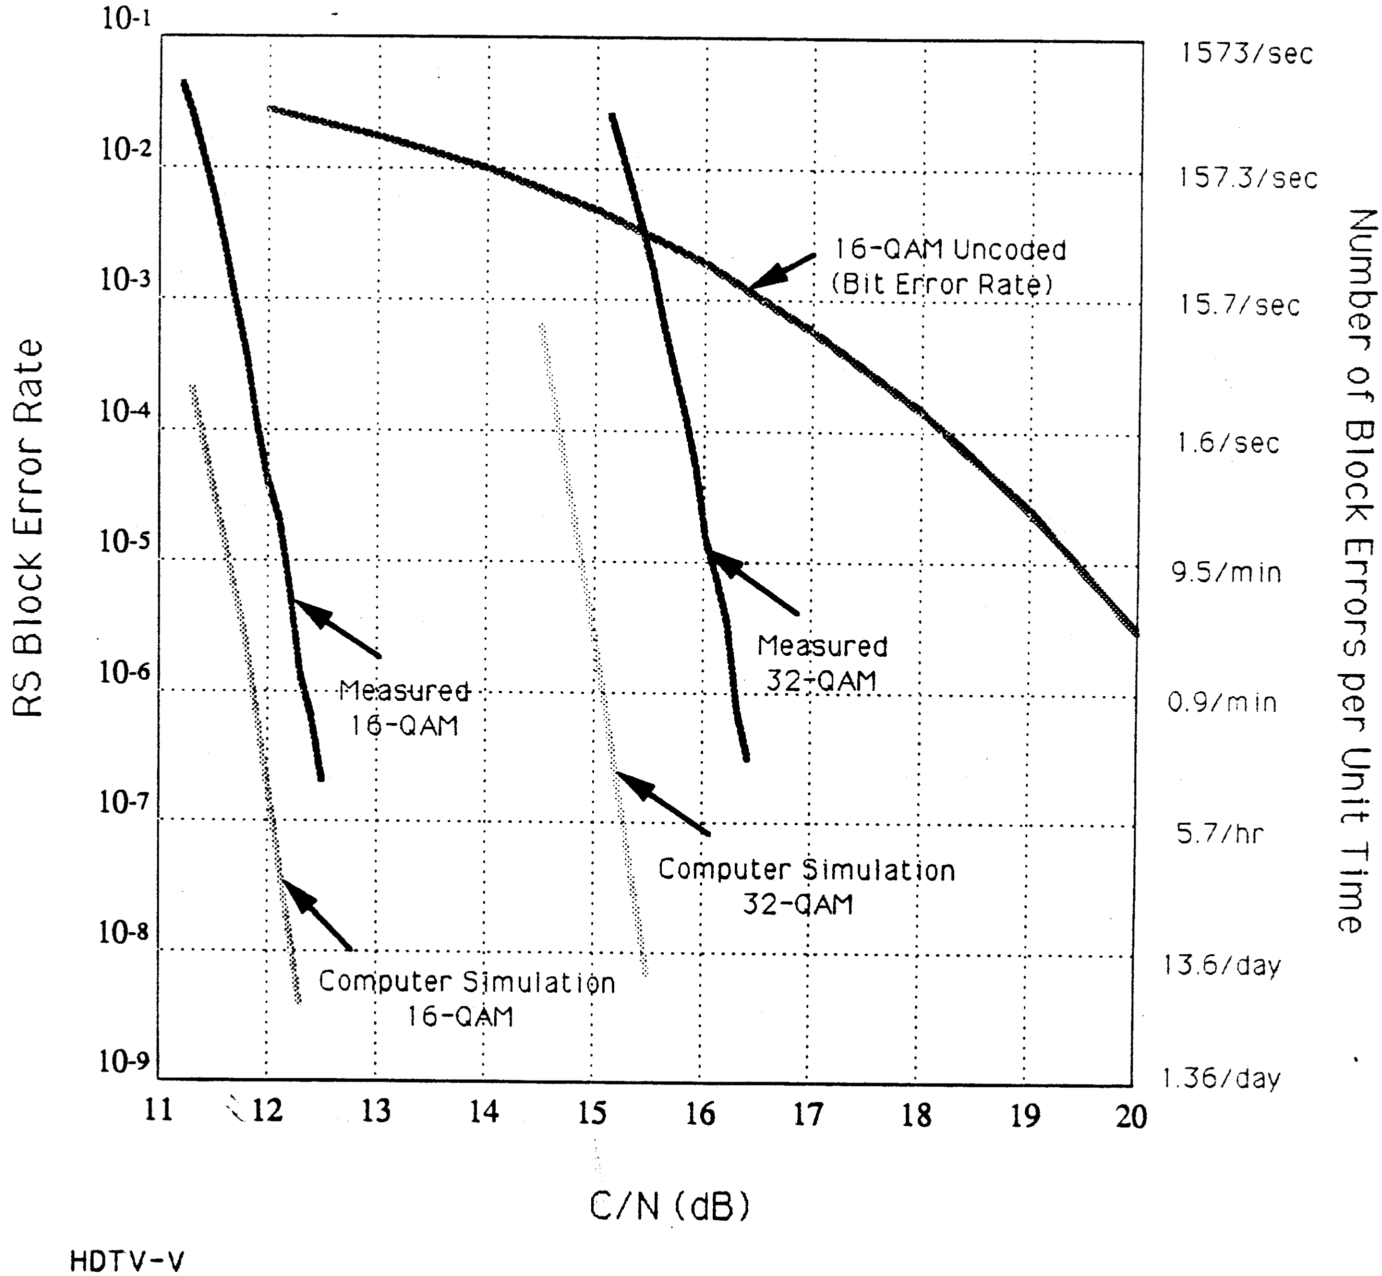

Figure 6-4. Performance of DigiCipher™ Transmission System

Figure 6-5 shows the signal constellations of the QAM signal. Digital filtering with a 21% roll-off raised cosine is used to prevent adjacent channel interference. SAW filters are used to further remove spurious signals. Figure 6-6 shows the output spectrum of the encoder. Notice that the output spectrum is well contained within a 6 MHz bandwidth.

6.3 Adaptive Equalizer

The DigiCipher™ HDTV receiver uses a 256 tap adaptive equalizer to handle multipath distortions. Figure 6-7 shows the block diagram of the adaptive equalizer. The effective range of the equalizer is -2 to +24 µsec and it can handle single or multiple echoes within the range. The level of the multipath can be as high as -6 dB for close-in echoes (-2 to +4 µsec) and -12 dB for long echoes (+4 to +24 µsec). The adaptive equalizer also compensates for non-ideal frequency response or group delay distortions caused by the transmitting amplifier, antennas, and the tuner.

6.4 Tuner

The 16/32-QAM demodulator does not place any special requirements on the tuner but the DigiCipher™ HDTV prototype has a double-conversion tuner. It has improved characteristics compared to a conventional single-conversion tuner in the area of spurious response rejection capabilities.

Figure 6-8 shows the block diagram of the tuner. The first IF frequency is 1200 MHz and the second IF frequency is 43.5 MHz. The first IF frequency has been chosen to be high enough for spurious free reception of all VHF and UHF frequencies, yet low enough for low cost implementation of the tuner. The selection of the 43.5 MHz IF as compared to 44 MHz is due to the availability of off-the-shelf SAW filters. The operating range of the VHF/UHF input to the tuner is -70 to -5 dBm.

The channel selection can be made by entering a two-digit number or by the depressing channel up/down key on the front panel.

6.5 Transmitting System Requirements

The DigiCipher™ system has been designed to be able to tolerate distortions that can be introduced by ATV transmitting systems. Table 6-1 shows the specification that ATV transmitting system needs to meet for proper operation of the system.

The system will operate at the max tolerable limit but with additional degradation in C/N of several decibel. The rank indicates the importance of the parameters where rank of 1 indicates the most critical parameter and the rank of 6 indicates the least critical parameter.

16-QAM Constellation

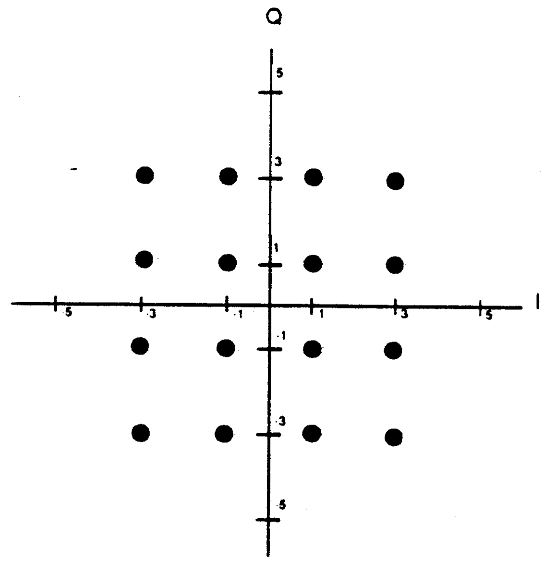

32-QAM Constellation

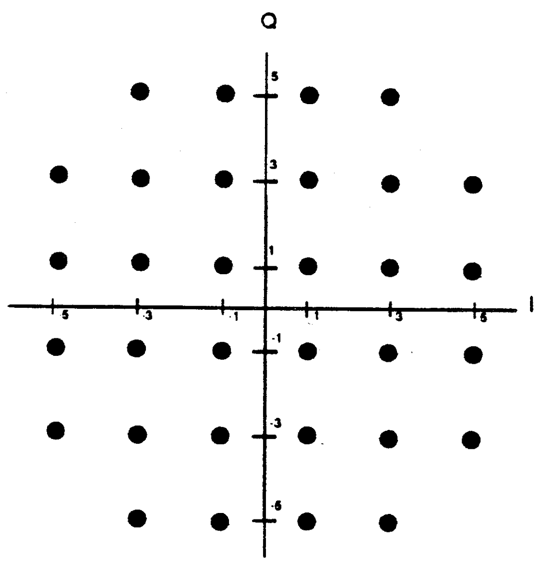

Figure 6-5. Signal Constellations of 16-QAM & 32-QAM Signals

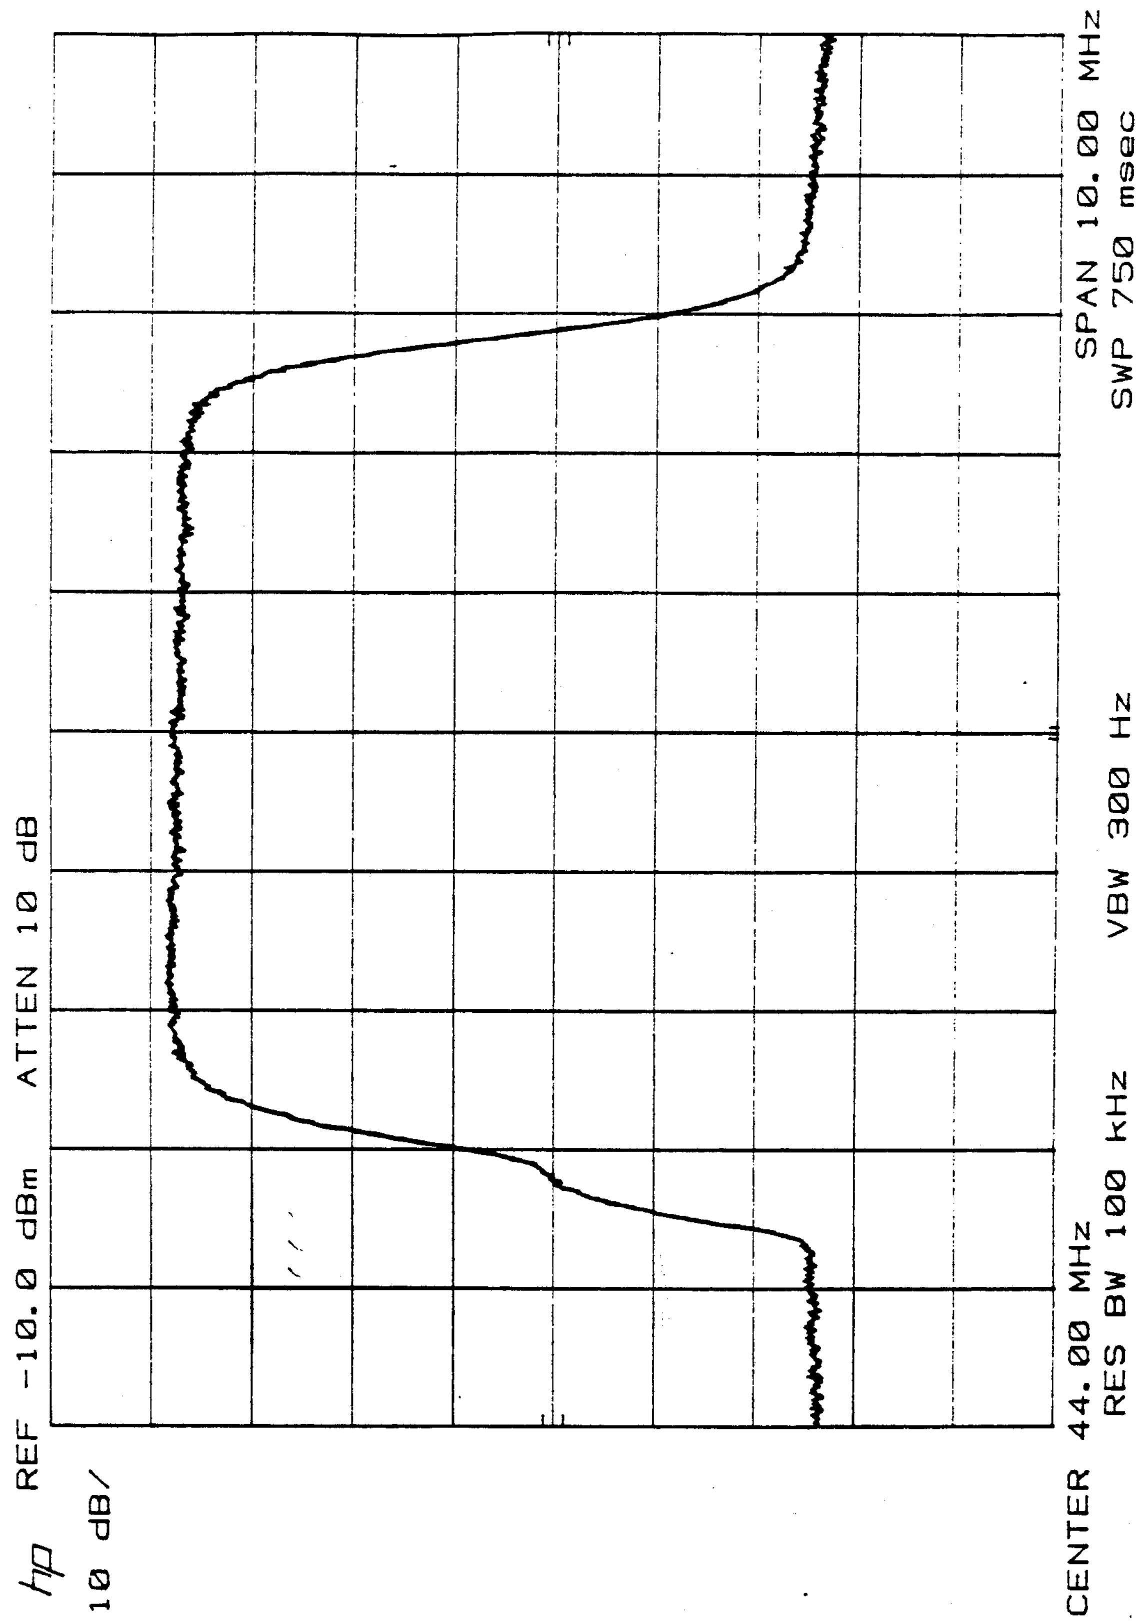

Figure 6-6. IF Output Spectrum

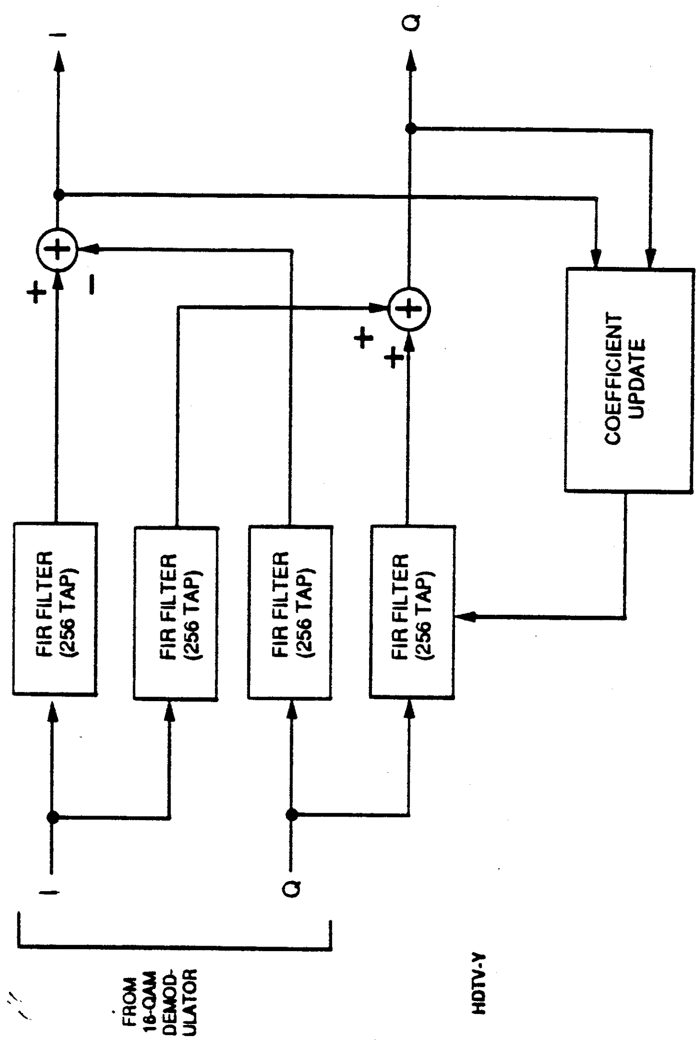

Figure 6-7. Adaptive Equalizer Block Diagram

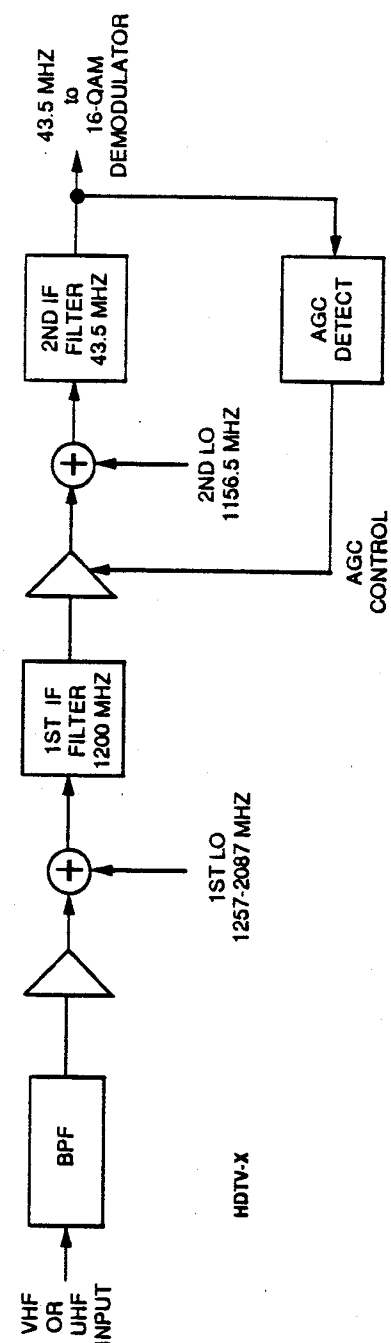

Figure 6-8. HDTV Tuner Block Diagram

Table 6-1. Transmitting System Specifications

Parameter	Specification	Max Tolerable Limit	Rank
Frequency Response (over 6 MHz):	± 1 dB	± 2 dB	4
Group Delay (over 6 MHz):	± 50 nsec	± 100 nsec	3
Amplitude Non-Linearity:	3%	5%	1
Incidental Carrier Phase Mod:	± 2 degrees	± 3 degrees	2
Output System Return Loss (over 6 MHz):	-20 dB	-10 dB	6
AC Hum:	-40 dB	-30 dB	5

7. SYNCHRONIZATION

The DigiCipher™ HDTV System has been designed to provide fast, reliable acquisition in the presence of noise, multipath, and interference. Through the use of advanced digital processing, the acquisition and synchronization time has been minimized.

7.1 Clock Synchronization

Figure 7-1 shows the block diagram of the clock/sync generator used in the encoder and the decoder. A VCXO generates a master clock at 214.6154 MHz. All the clocks used for the digital video processing and the 16/32-QAM modem are derived from the master clock. The encoder synchronizes the master clock into the 31.47 kHz horizontal sync provided externally. The decoder synchronizes the master clock into the QAM symbol rate. The phase-locked loops have been designed to minimize the clock jitter (a few nsec).

The encoder also makes use of vertical sync externally provided to synchronize the frame vertical sync. The decoder synchronizes the internal vertical sync by using the 24-bit sync generated by the encoder.

7.2 Acquisition

When the HDTV receiver is tuned into a new channel, the decoder goes through a number of synchronization processes. The DigiCipher™ HDTV system requires approximately 0.40 seconds for QAM demodulator and overall synchronization to occur, and 0.37 seconds for complete refreshing of the decoded video. Total acquisition time of the iDigicipher™ HDTV system, therefore, is 0.77 seconds. However, subscribers will be able to start to see portions of the new channel video and hear the audio in 0.4 seconds. Apparent acquisition time is thus significantly faster than 0.77 seconds. Table 7-1 lists the breakdown of the synchronization process.

Table 7-1. Signal Acquisition time

QAM Demodulator and Overall Synchronization	
AGC	0.05 sec
Bit Sync	0.10 sec
Adaptive Equalizer	0.10 sec
Carrier Sync	0.10 sec
Interleaver	0.02 sec
24-bit Sync	0.03 sec
Total	0.40 sec
Digital Video Decoder	
DPCM Refreshing	0.37 sec
Total Acquisition Time	**0.77 sec**

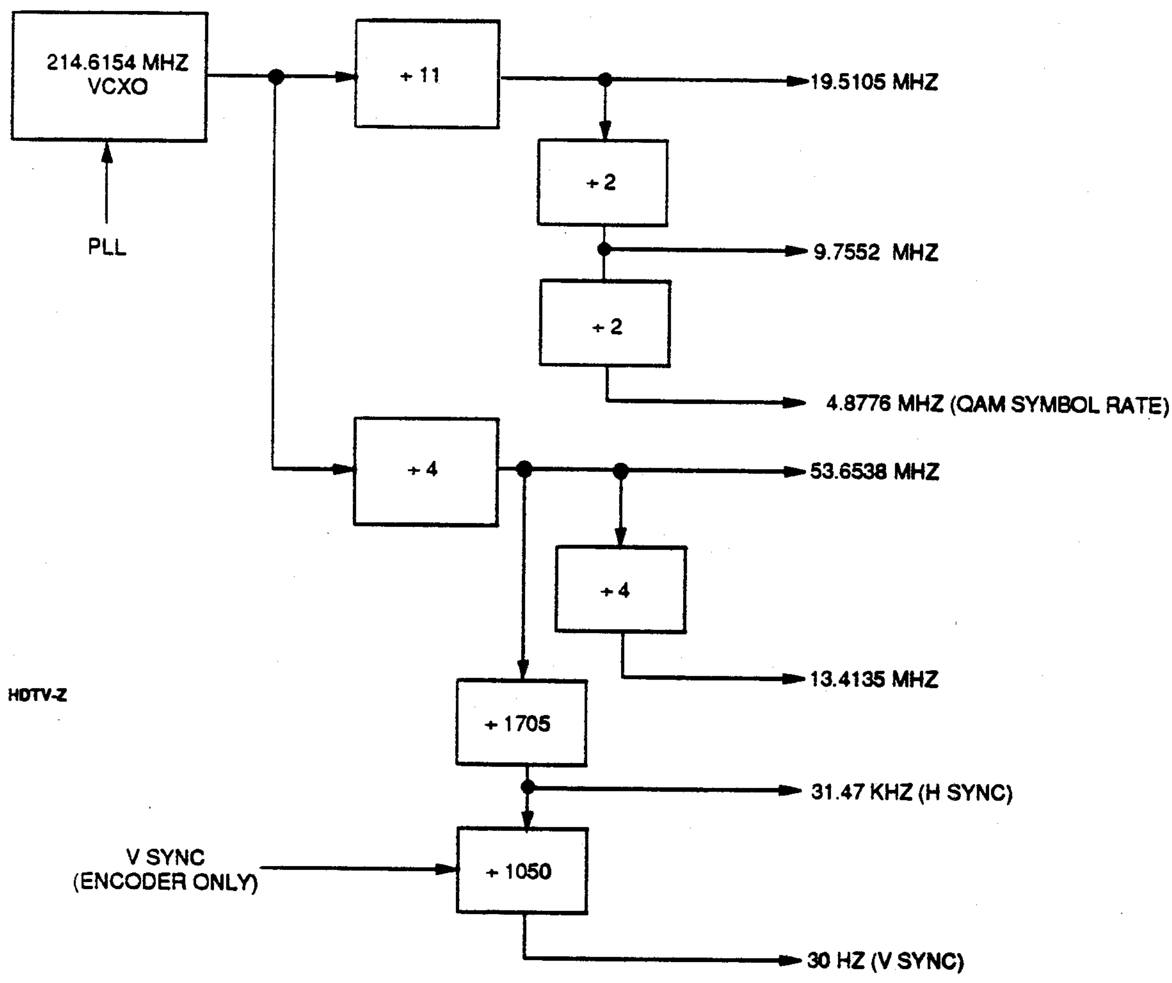

Figure 7-1. Clock/Sync Generator

8. HARDWARE DESCRIPTION

The DigiCipher™ HDTV prototype hardware has been designed to fully comply with the ATTC interface specifications. The system consists of two 6', EMI shielded racks; one for the encoder and one for the decoder. The encoder and the decoder draw approximately 15 amps each at 120 volts AC.

8.1 Encoder

The encoder consists of 16 multi-layer printed circuit boards in a rack mountable cage. The encoder also has a power supply, a fan, 2 audio synchronizers, 2 Dolby® AC-2 digital audio encoders, video distribution amplifiers, a control panel, and an interface panel mounted in a rack as shown in Figure 8-1.

8.2 Decoder

The decoder consists of 11 multi-layer printed circuit boards in a rack mountable cage. The decoder also has a power supply, a fan, a VHF/UHF tuner, video distribution amplifiers, a control panel, and an interface panel mounted in a rack as shown in Figure 8-2.

8.3 Consumer HDTV Receiver

The Digicipher™ system has been designed to provide optimum performance while the hardware complexity is kept low. This allows the production of low cost consumer HDTV sets as well as low cost broadcasting equipment. Consumer HDTV sets can be designed using a total of 12 custom VLSIs as shown in Table 8-1. In addition, 2 Mbytes of memory are required for the implementation of the digital video decoder.

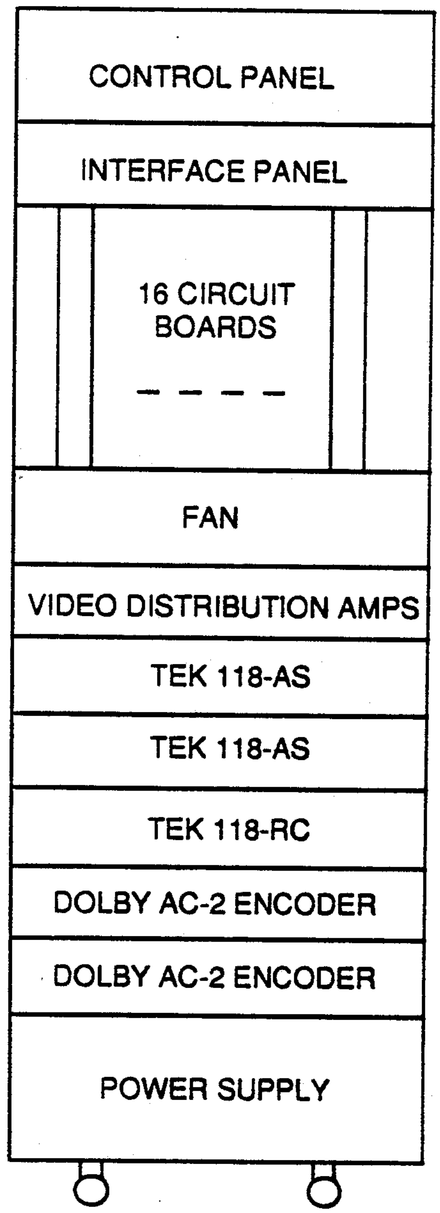

HDTV-BB

Figure 8-1. DigiCipher™ HDTV Encoder

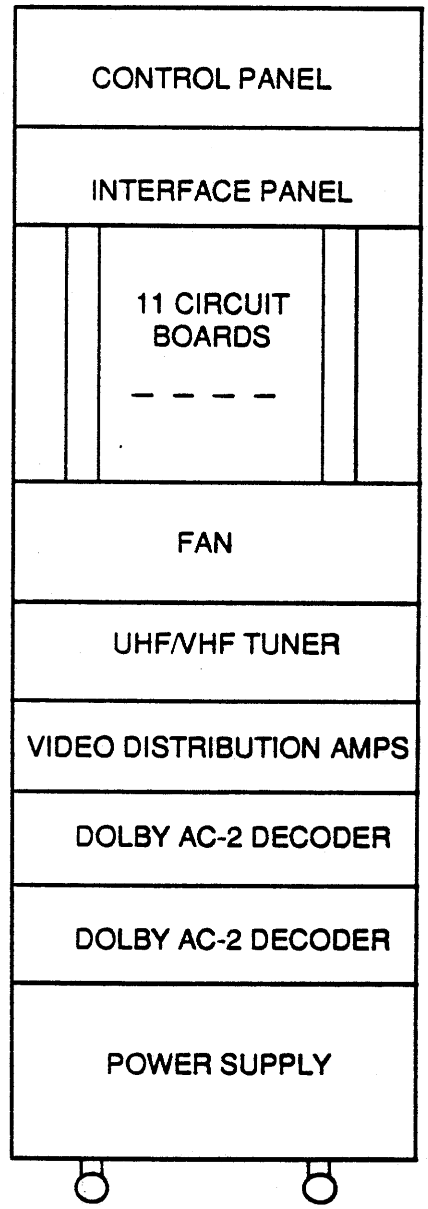

HDTV-CC

Figure 8-2. DigiCipher™ HDTV Decoder

Table 8-1. HDTV Receiver VLSI Chip Set

VLSI Chip	Qty Per HDTV Receiver
16/32-QAM Demod	1
Adaptive Equalizer	4
FEC Decoder/Sync	1
Decompression	4
Video Mux/Filter/OSD	1
Digital Audio Decoder	1
Total	**12**

9. COVERAGE AREA ANALYSIS

In this section, an analysis of the coverage area of the DigiCipher™ HDTV System is provided. The analysis is based on a set of assumptions as well as measured performance of the DigiCipher™ HDTV transmission system.

9.1 Measured Performance

NTSC Interference-Into-ATV-Service

The DigiCipher™ HDTV transmission system provides essentially error free performance at 2 dB C/I in the 16-QAM mode and 7 dB C/I in the 32-QAM mode where C is the average carrier level of the ATV signal and I is the peak sync level of NTSC interference signal. An HP3780A Noise & Interference Test Set was used for this measurement.[1] This superior NTSC interference rejection capability is mainly attributable to the powerful forward error correction and two levels of interleaving.

ATV Interference-Into-ATV Service

ATV-to-ATV C/I performance has been measured to be 0.5 dB better than the C/N performance of the DigiCipher ™ HDTV system. The difference can be attributed to the non-Gaussian nature of the QAM signal.

9.2 ATV Coverage Area Calculations

Key transmission performance parameters are provided in sections 6.1 and 9.1. PS/WP-3 has the responsibility to project coverage areas, using ATTC measurements of the parameters. Realizing the high level of interest in those projections, and that they will not be available for quite some time, we provide our own here.

Making projections is complicated by the fact that planning factors for the ATV service are not yet resolved. We have the choice of (1) using the same assumptions others have used, with the result that our projections can be more readily compared, or (2) using a more conservative set of assumptions which we believe to be more realistic, but which do not support direct comparisons. We have chosen to provide both. In doing so, we are attempting to influence the discussion to be more responsible. We are certainly open to further refining our planning factor assumptions.

In providing an ATV service, the key factors are spectrum availability, ATV service area, and the impact on current NTSC service. Consistent with our understanding of work to date in the area, we have made the following central assumptions:

[1]For proper calibration of the NTSC interference level, a CW carrier at the peak sync level of the NTSC signal has to be supplied to the I input when the HP3780A enters into the calibration mode.

1. The ATV service area should be comparable to the current NTSC service area.

2. Interference to existing NTSC service by an ATV signal should be no worse than the interference into NTSC by another NTSC signal now.

3. ATV service will be co-channel interference limited, and, in order to meet simulcast channel allocation goals, co-channel transmitters between ATV and NTSC, and ATV to ATV, may have to be as close as 100 miles.

In what follows, we first present our planning factor assumptions in two categories: 9.2.1, System Independent, presents factors which do no depend on the specifics of the ATV system. 9.2.2, System Dependent, presents factors which do depend on the performance of the specific system being analyzed.

Section 9.2.3 presents our coverage calculations. We are very pleased with the performance of our DigiCipher™ transmission system. Depending on the specific assumptions, results range from showing that DigiCipher™ system will more than meet broadcasters' requirements, to DigiCipher™ system coming close to, but not quite meeting those requirements. As expected, 16-QAM performance exceeds that of 32-QAM, but the difference is not great.

9.2.1 System Independent Planning Factors

9.2.1.1 Transmitting Antenna Height and NTSC Effective Radiated Power

We conform to PS/WP-3's choice of 1200 feet as the UHF transmitter Height Above Average Terrain (HAAT). Assumed NTSC Effective Radiated Power (ERP) is 37 dBk.

9.2.1.2 NTSC Service Assumptions

We have used the FCC planning factors assumed for NTSC receivers for UHF grade B service:

Antenna gain: 13 dB

Antenna front-to-back (F/B) ratio: 6 dB

Downlead loss: 5 dB

Receiver noise figure: 15 dB

9.2.1.3 NTSC Carrier-to-Noise Ratio Requirements

The NTSC grade B service boundary is the locus of points for which the carrier-to-noise ratio (C/N) is 28.5 dB (peak sync to RMS noise in a 6 MHz band). This level of service is based on an F(50, 90) field, i.e., a field strength available at least 50% of locations at at least 90% of the time, for a receiving antenna height of 30 feet above ground level. F(50, 90) field strength data are derived from F(50, 50) data provided in FCC Rules (Part 73.699). For an NTSC UHF transmitter with 37 dBk peak ERP at 1200 feet HAAT, the grade B service contour is at a nominal radius of 56 miles. See Figure 9-1.

9.2.1.4 NTSC Interference Requirements

Co-channel NTSC interference is measured in terms of a carrier-to-interference (C/I) ratio. For the case of NTSC carriers employing nominal frequency offset, a 28 dB Desired/Undesired (D/U) ratio defines the acceptable limit of interference. Given the assumptions used in this study, the point at which a 28 dB D/U penetrates furthest into the NTSC service area occures 41.5 miles from the desired NTSC transmitter, along a straight line between the desired and undesired transmitters. Outside this point, i.e., points at which the D/U is less than 28 dB, service is considered to be interference limited. See Figure 9-1.

Some analyses of ATV service areas have used nominal frequency offset to define the interference limited service area which may be duplicated by an undesired ATV signal intruding into an NTSC service area. For comparability of results, we have also done so.

However, we believe a more realistic, though harsher, requirement is to utilizes precise frequency offset to define the NTSC interference limited service area. In such case, a 22 dB D/U yields comparable subjective results to a 28 dB D/U in the nominal frequency offset case. As Figure 9-1 also shows, the 22 dB point occurs 46.5 miles from the desired transmitter. The result is a smaller interference limited service area. We have also provided analyses assuming precise frequency offset.

9.2.1.5 ATV Service Assumptions

Planning factors for ATV are similar to those for NTSC receivers, but with some exceptions:

Antenna gain: 13 dB (same as NTSC)

Antenna F/B ratio: 16 dB (NTSC = 6 dB)

Downlead loss: 5 dB (same as NTSC)

Receiver noise figure: 10 dB (NTSC = 15 dB)

A 10 dB noise figure is generally accepted as characteristic of current RF front-end performance.

We have assumed a 16 dB front-to-back ratio for an ATV receiving antenna, as have some others. While that is significantly greater than 6 dB used in NTSC planning, we believe that such performance is feasible, and reasonable in cost. Essentially all UHF antennas which meet the 13 dB gain value exhibit F/B ratios that far exceed the 6 dB value used in NTSC planning. Consumer antennas with values exceeding 20 dB have been available from several manufacturers for many years. (See reference 8).

Practical space limitations and objects near the antenna can degrade the effective F/B ratio in a specific installation. Much of the degradation observed in the past is attributable to unshielded twin-lead antenna wire used. Modern usage of shielded coaxial cable addresses this problem. In assuming a 16 dB F/B ratio, we are assuming an antenna F/B ratio of 22-24, which then degraded by the installation to 16.

9.2.2.4 ATV Into ATV Co-Channel Interference Requirements

Based on measurements, the ATV-ATV co-channel C/I threshold is 12 dB for the 16-QAM mode, and 16 dB for the 32-QAM mode. This measure is based on average desired ATV carrier power to average undesired ATV interference carrier power.

9.2.3 Coverage Calculations

The UHF planning factors used are summarized in the Table 9-1:

Table 9-1. Link Budget Calculation Results for NTSC and ATV UHF Service [3]

Factor	NTSC	16-QAM ATV	32-QAM ATV
Noise Figure (dB)	15	10	10
C/N (dB)	28.5	12.5	16.5
Antenna Gain (dB)	13	13	13
Antenna F/B (dB)	6	16	16
Line Loss (dB)	5	5	5
Req'd Local Field (dBuv/m)	60	38.5	43
Time Availability Adjust (dB)2	4	7.2	7.2
Req'd F(50, 50) Field (dBuv/m)	64	45.7	50.2

Given the required fields specified in the table, Figure 9-2 shows the 16-QAM ATV-NTSC 100 mile separation interference-limited case, whereupon the ATV transmitter power is allowed to reach a level of 19.5 kBk average ERP, so as to match the interference penetration depth into the NTSC service area to that of an NTSC-NTSC nominal frequency offset, i.e., 41.5 miles from the desired transmitter. The 16-QAM ATV coverage area under 99% time availability criteria also matches the NTSC grade B service area, while ATV with 90% time availability reaches 61 miles.

In Figure 9-2, the NTSC interference into ATV reaches its boundary level at 42 miles, in close match with the 41.5 mile point of ATV into NTSC. From this figure, given the assumptions, it may be concluded that 16-QAM service exceeds the performance of NTSC in every respect at a separation of 100 miles.

ATV-ATV co-channel service with 16-QAM ATV at 100 miles sepapration is shown in Figure 9-3. As can be seen, in this case the mutual interference zones are smaller than those encountered in the case of NTSC-NTSC using nominal frequency offset.

[3]NTSC: Under the FCC practice with 90% time availability. ATV: Digital HDTV transmission using 16-and 32-QAM with 5 MHz noise bandwidth with 99% time availability

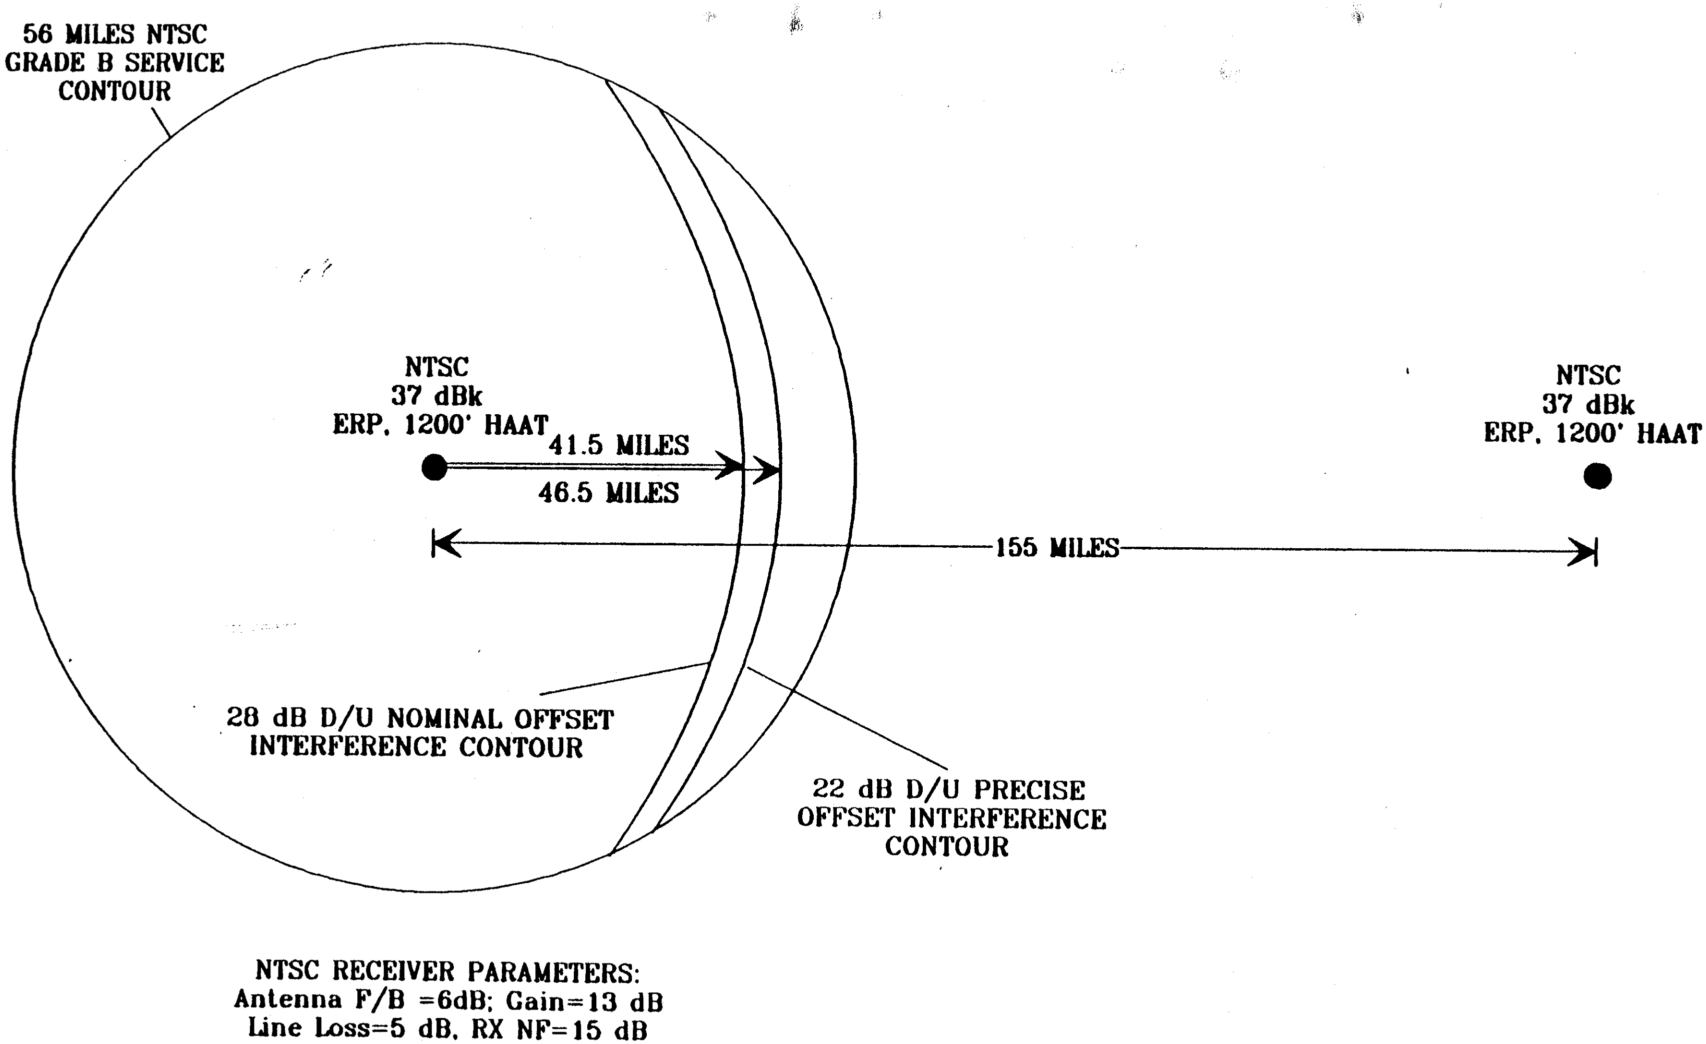

Figure 9-1. Minimum NTSC Co-Channel Separation and the Interference Boundary

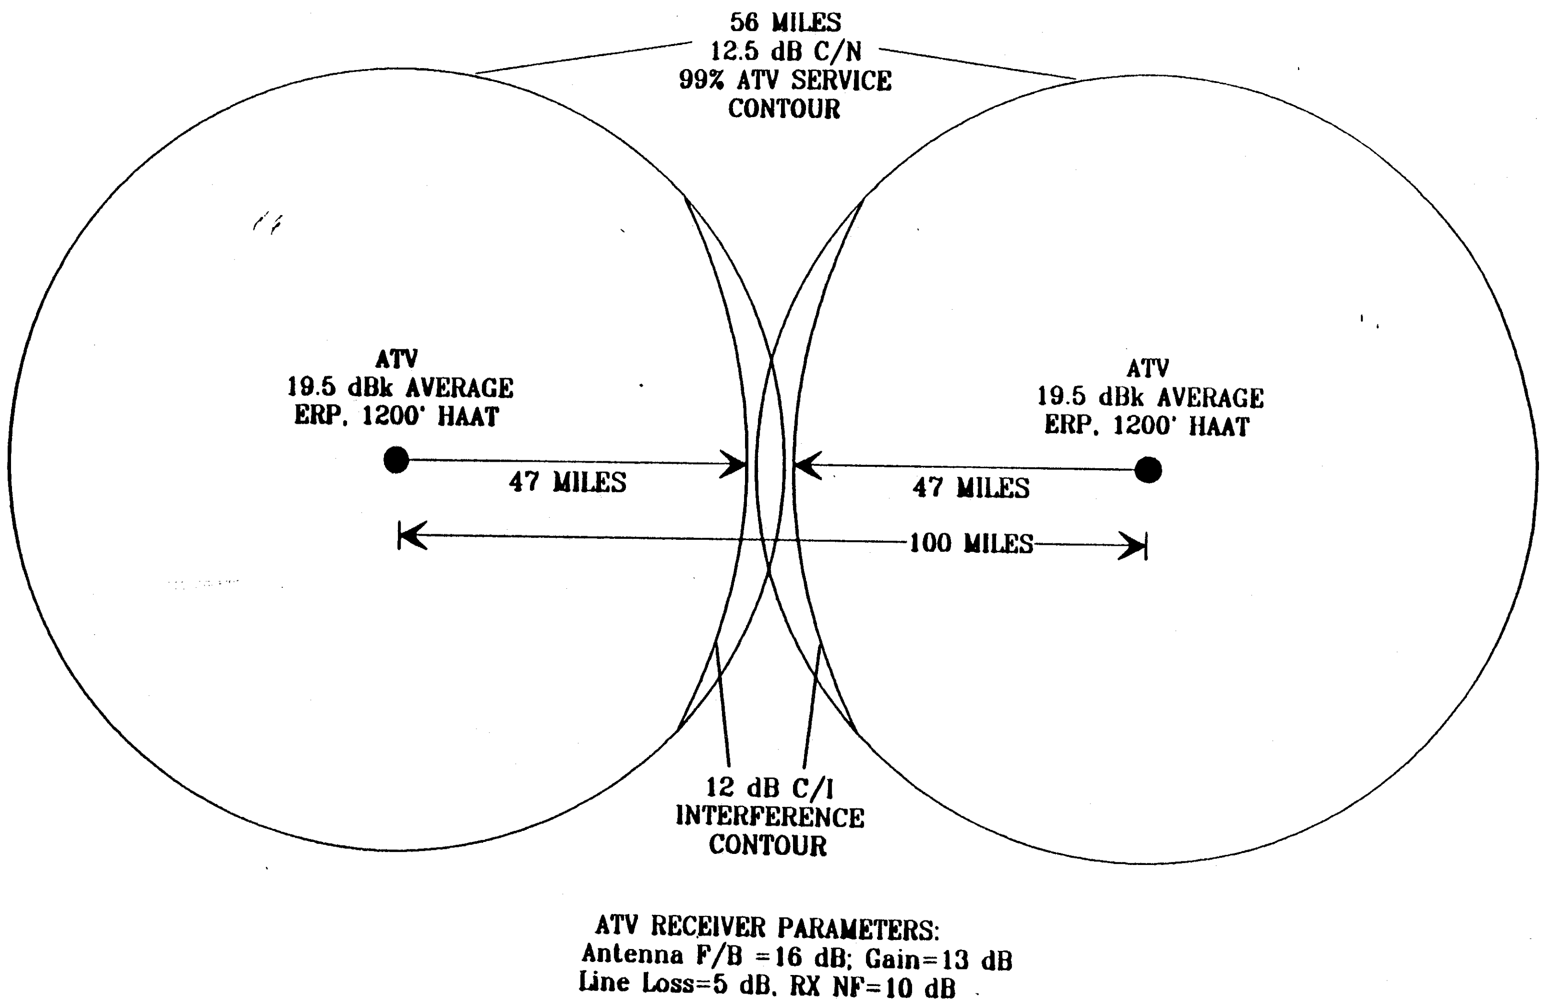

Figure 9-3. Coverage for 16-QAM ATV-ATV Interference Limited Service

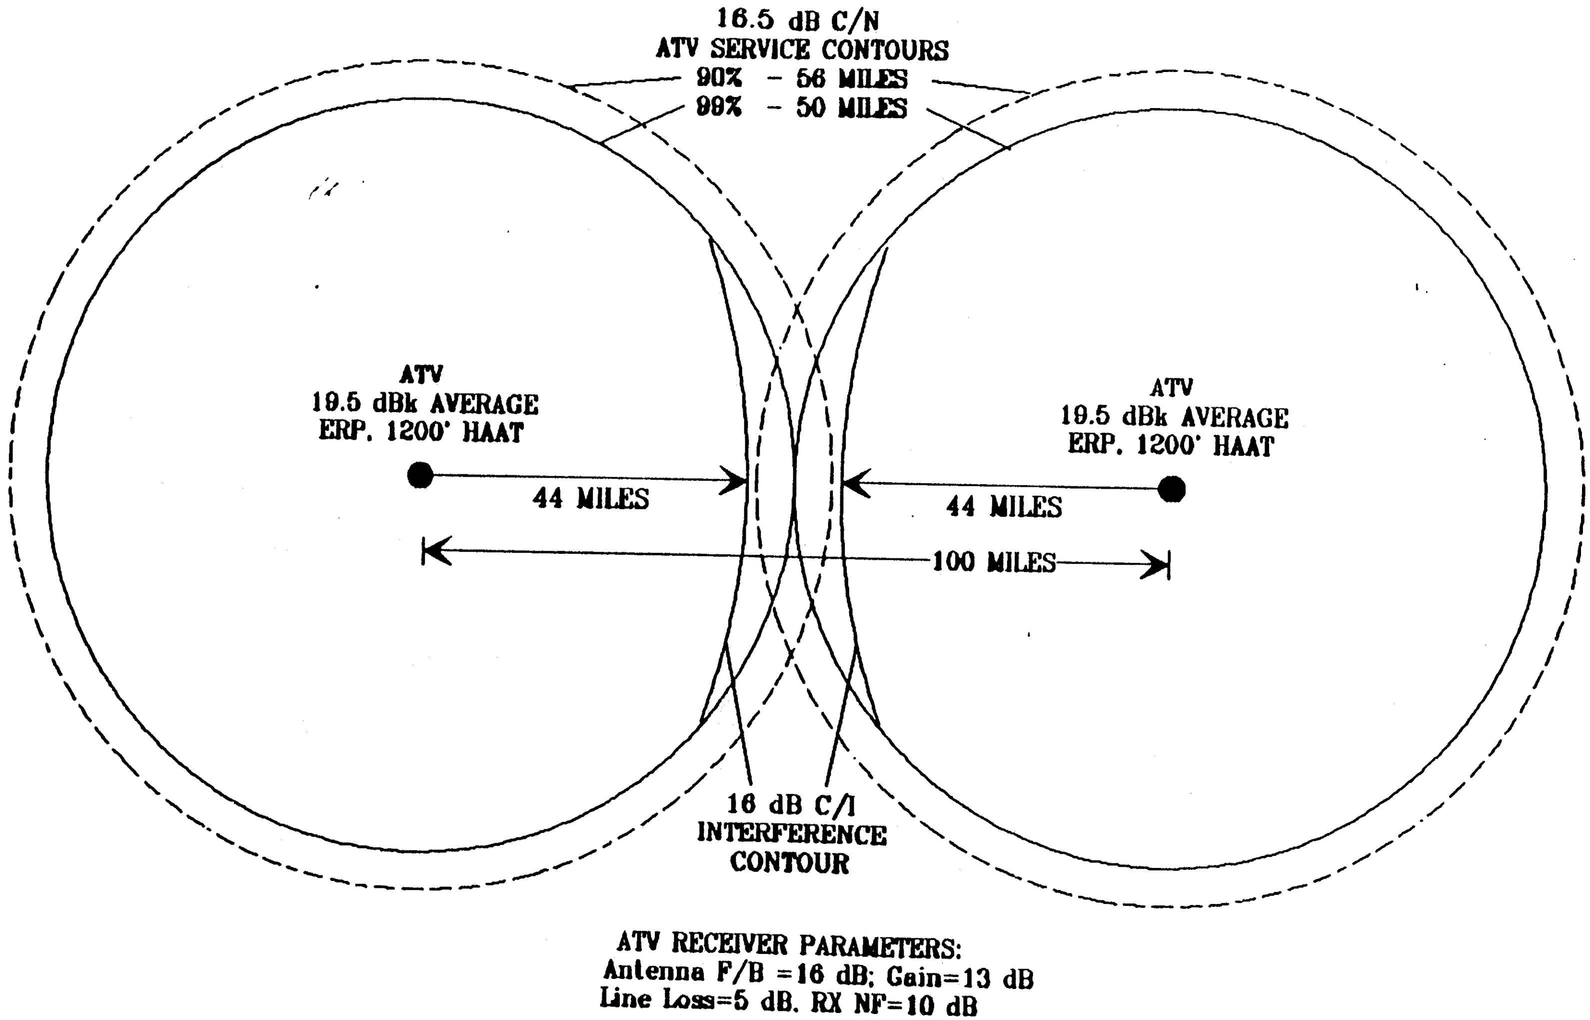

Figure 9-5. Coverage for 32-QAM ATV-ATV Interference Limited Service

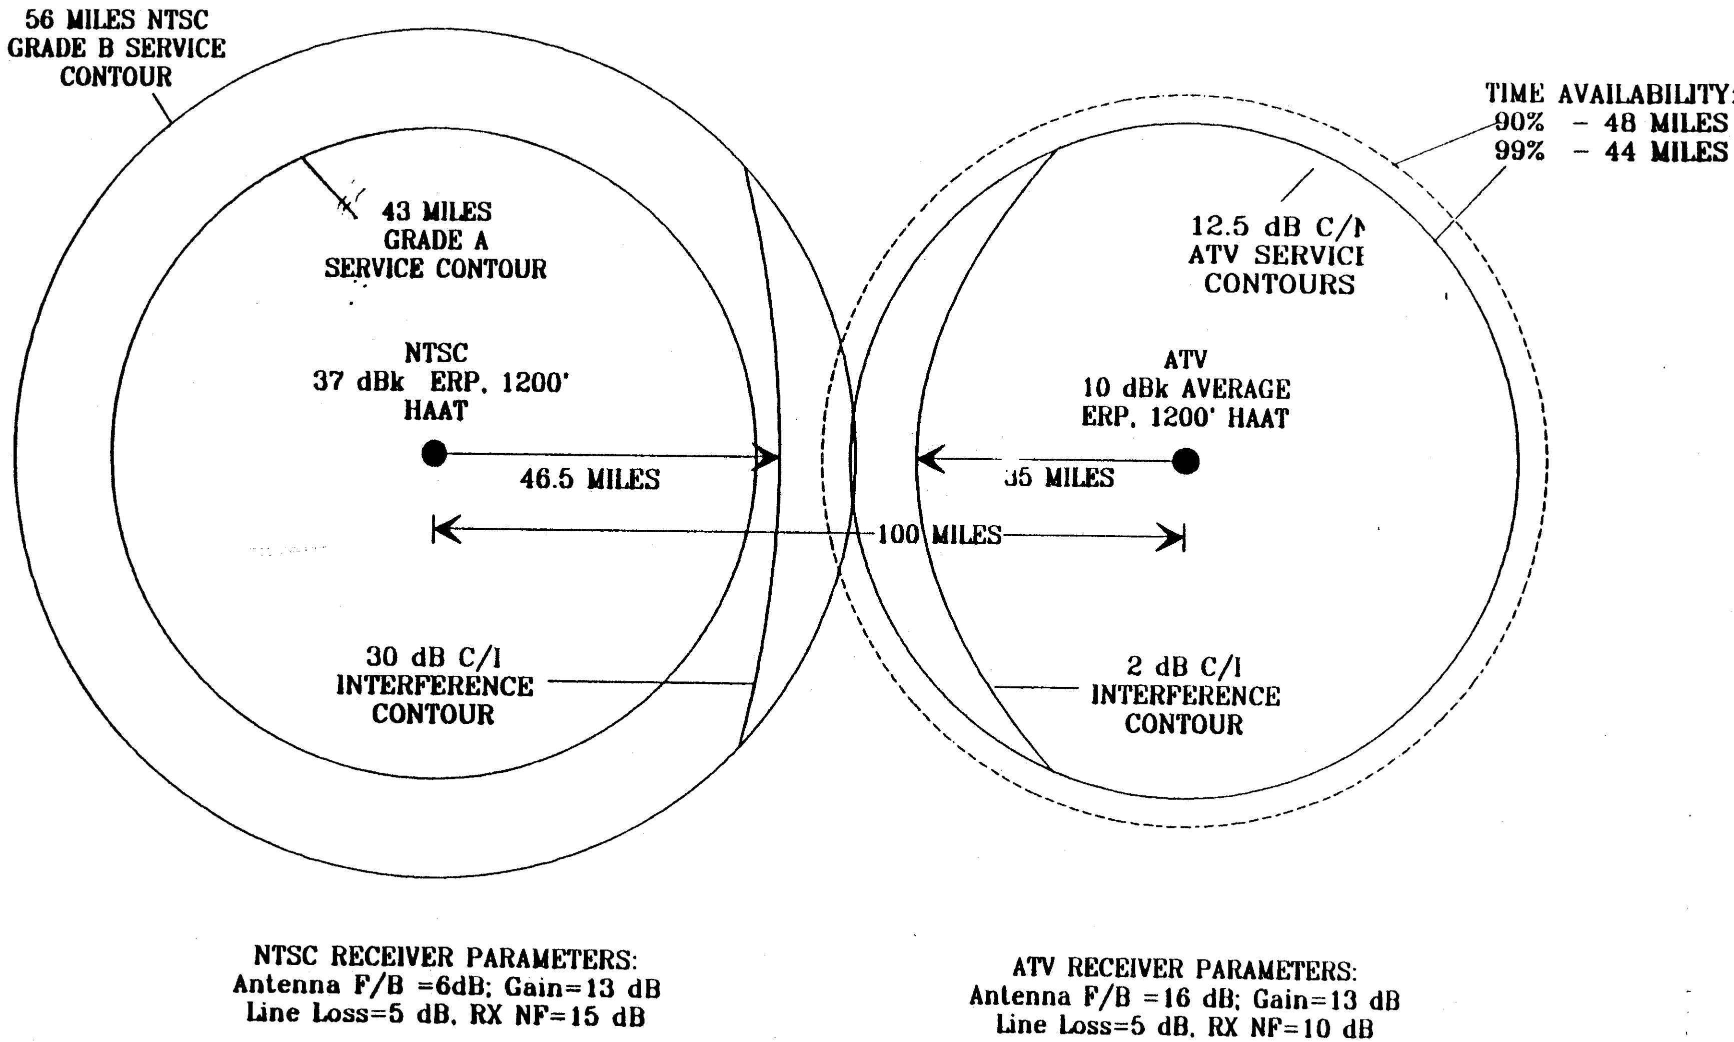

Figure 9-7. Co-Channel 16-QAM ATV and NTSC with 100 Mile Separation

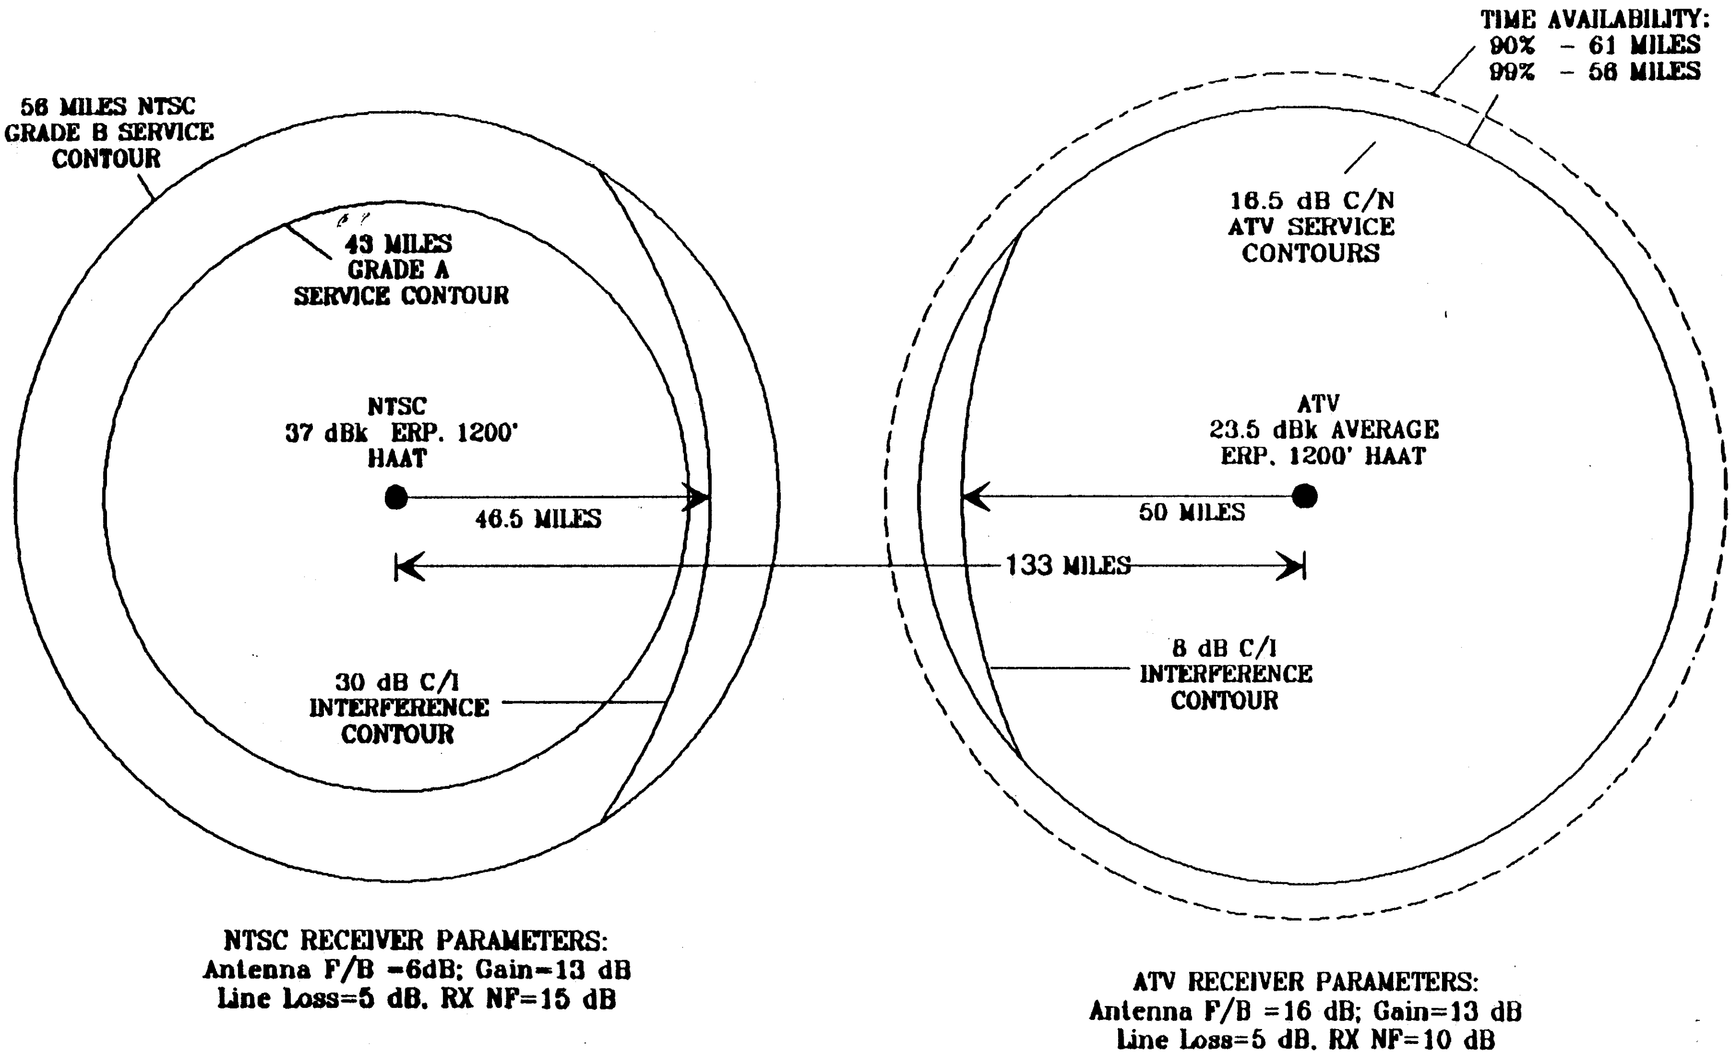

Figure 9-9. Co-Channel 32-QAM ATV and NTSC with 133 Mile Separation

10. ALTERNATE MEDIA DISTRIBUTION

10.1 Cable Transmission

The DigiCipher™ HDTV System is completely cable transmission compatible. The DigiCipher™ HDTV signal can be placed in a 6 MHz cable channel adjacent to other DigiCipher™ HDTV signals or NTSC VSB-AM signals. Features of the DigiCipher™ HDTV System for cable applications include:

- Pass through of satellite or broadcaster delivered signals to the cable subscriber without signal decompression and recompression at the cable headend.

- Lower power requirements than VSB-AM NTSC will result in an unimpaired HDTV signal delivered to the subscriber with all of the advantages of lower system power loading

- Channel transmission compatibility with the DigiCipher™ multi-channel NTSC system allows reception/access control of both signals with the same cable converter.

10.2 Satellite Transmission

The DigiCipher™ HDTV System can be transmitted over C-band or Ku-band satellite channels using QPSK modulation.

The system can support both FSS and BSS satellite transponders. The threshold C/N is 7.5 dB measured over a 24 MHz bandwidth, therefore the DigiCipher™ HDTV System allows the use of smaller dish size compared to other analog or hybrid HDTV systems.

10.3 Other Terrestrial Distribution

Since the DigiCipher™ HDTV System is an all-digital system, it can be readily applied to other transmission media such as microwave distribution service (MDS), multi-channel MDS (MMDS) and fiberoptic cables (FO).

An inherent characteristic of the all-digital system is that the HDTV service is free from transmission artifacts caused by various transmission media. Also, the complexity of the interface equipment between various transmission media is substantially lower.

11. REFERENCES

1. W. K. Pratt, Digital Image Processing, John Wiley and Sons, Inc., New York, NY, 1978.

2. S. Ericson, "Fixed and Adaptive Predictors for Hybrid Predictive/Transform Coding," IEEE Trans., on Comm., Vol. CQM-33, No. 12, December 1985.

3. W. Chen and W. K. Pratt, "Scene Adaptive Coder," IEEE Trans. on Comm., Vol. Com-32, No. 3, March 1984.

4. A. Fernandez, R. Ansari, D. J. Gall and C. T. Chen, "HDTV Subband/DCT Coding: Analysis of System Complexity," IEEE Globecom Proceedings, 343.1.1 - 343.1.4, 1990.

5. K. Feher, Advanced Digital Communications, Prentice-Hall, Inc. Englewood Cliff, NY, 1987.

6. V. Phargave, D. Haccoun, R. Matyas, P. Nuspl, Digital Communications by Satellite, John Wiley and Sons, Inc. New York, NY 1981.

7. W. Paik, "DigiCipher™ - All Digital, Channel Compatible, HDTV Broadcast System," IEEE Trans. on Broadcasting, Vol. 36, No. 4, December 1990.

8. Jerrold GT Catalog, pp. 2-6, Jerrold Electronics Corporation, Horsham, PA, 1976.

Appendix A. List of Clock and Carrier Frequencies Used in the Prototype HDTV System.

DIGITAL VIDEO ENCODER/DECODER	
Master Oscillator	214.6154 MHz
A/D & D/A Sampling Clock	53.6538 MHz
Digital Video Processing Clock	13.4135 MHz
H Sync	31.47 KHz
V Sync	29.97
DIGITAL AUDIO ENCODER/DECODER	
DSP CLOCK	27 MHz
Master Oscillator	12.0839 MHz
A/D & D/A Sampling Clock	47.2 KHz
Digital Audio Data Clock	251.75 KHz
DATA INTERFACE	
Master Oscillator	18.525 MHz
Digital Data Board Rate	9600 Baud
DIGITAL MODULATOR/DEMODULATOR	
QAM Symbol Rate	4.8776 MHz
2 x Symbol Rate	9.7553 MHz
4 x Symbol Rate	19.5105 MHz
Local Oscillator	44 MHz (mod)/43.5 MHz (demod)
TUNER	
1st IF	1200 MHz $\pm$3 MHz
2nd IF	43.5 MHz $\pm$3 MHz
1st LO	1257 - 2087 MHz
2nd LO	1156.5 MHz

Appendix B. Technical Description of Dolby AC-2 System

SEE ATTACHED DESCRIPTION

AC-2: A FAMILY OF LOW COMPLEXITY TRANSFORM BASED MUSIC CODERS

Louis D. Fielder
Grant A. Davidson
Dolby Laboratories, Inc.
San Francisco, California 94103

ABSTRACT

Two high-quality data rate reduction music coders from a family of TDAC transform based coders are discussed. An overview of the psychoacoustic principals used in their design is given and their limitations discussed. The use of psychoacoustics and DSP technology are combined to yield a low complexity approach to music coding. Issues of complexity, word length requirements, and memory usage are examined for both general-purpose DSP and custom IC implementations.

0. INTRODUCTION

The use of data rate reduction coders for digital audio applications shows great promise for a large variety of storage and transmission applications. Since Compact Disc digital audio employs a data rate greater than 1.4 Mbits/sec., this type of digital audio has been limited only to areas that can maintain a high data rate. Fortunately, the development of high quality data rate reduction technology for music applications has changed this situation. Now lower data rates may be used for audio in radio and television broadcast, computer hard disk storage, and telephone line connections. This paper will describe two coders from the Dolby AC-2 family, developed for different applications, that have the desired characteristics of data rate reduction, excellent sound quality, and computational simplicity.

The need to reduce the data rate for the practical application of digital audio into many areas has resulted in much work in the field of data rate reduction for music, as typified by Brandenburg et al. [1990], Johnston [1988], Schroeder et al. [1987], Stoll and Dehery [1990], Davidson et al. [1990], and Fielder [1989]. The fundamental approach of these techniques is to divide the audible frequency range into sub-bands which approximate auditory critical bands. Crucial elements in the design of these coders are the bit allocation and quantization schemes in which perceptually relevant sub-bands are identified, and the appropriate fraction of the available bit rate assigned to their representation. Many of these algorithms require a great deal of processing power to perform the frequency division and quantization operations (e.g., multiple DSP chip implementations for a single audio channel). Furthermore, they all extrapolate published models of human hearing and masking to a broader class of signals than those upon which the models were based.

This paper builds on the work described by Davidson et al. [1990] and Fielder [1989] which described 15 kHz bandwidth coders with resultant data rates between 128 and 192 kbits/sec. per channel. The two coders described here have 20 kHz bandwidth, require less than one programmable DSP chip to implement one stereo pair, and possess excellent sound quality. In particular, one coder, which will be called the low-delay coder, achieves excellent subjective and objective quality at 4:1 compression, exhibits robust tandem coding performance (i.e., where a number of encode/decode processes occur in series) and has a coding/decoding delay less than 9 msec. This low-delay feature is essential for applications requiring that announcers monitor their own coded voice signals. The other coder trades coding delay for a lower bit-rate (6:1 compression) and will be called the moderate delay coder. The coding systems described here can be applied for either 44.1 k or 48 ksamples/sec., however the remaining discussions will center on 48 ksample/sec. results.

A general overview of the psychoacoustics of masking as it effects the design of data rate reduction music coder technology will be given. Next, the details of the two coding systems resulting from this psychoacoustic examination will also be presented. Issues of implementation will also be discussed and the use of 24-bit and 16-bit DSP chips will be examined and processor speed/memory requirements determined. The use of custom DSP chips will also be considered. It will be shown that the two systems described are quite low in complexity while at the same time providing excellent sound quality.

1. APPLICATION OF PSYCHOACOUSTIC MODELS TO CODER DESIGN

The basis of all good rate reduction music coders is the application of the psychophysics of the human auditory system. As a result, a discussion of the present state of knowledge in this area is essential for the understanding of coders of this type. Masking effects for simple signals will be extended to the development of the filter bank design and quantization technology used in music coders. It will be seen that the targeted application will greatly influence the way the psychoacoustic principals are utilized. Next, these principles will be extended to more complex signals and discussed for AC-2 coding. An indication of the effectiveness of the AC-2 coding system in controlling the amount and frequency characteristics of the errors due to the reduction of word-lengths for data rate reduction will be given by a spectral comparison between both coder's performance and frequency characteristics of auditory masking.

1.1 Critical-Band Model of Hearing

Central to the development of a workable model of the auditory system is the critical-band concept and its relationship to the masking characteristics of the ear. The critical-band model of the human auditory system was first developed by Fletcher [1940] to explain why masking experiments showed that signals covering a frequency range less than a certain threshold bandwidth produced the same masking and detection properties as other signals with smaller bandwidths. The fundamental approximation is that the ear acts as a multi-channel real-time analyzer with varying sensitivities and bandwidths throughout the audio range. Despite the intrinsic simplicity of the model, it has been shown to be very enduring. Effective data rate reduction coders for music rely heavily on this model.

The critical-bandwidth represents the minimum frequency bandwidth resolvable for masked signals. For example, the masking of a low level error signal caused by a larger level tone nearby in frequency is maximal and continues at a constant level until the frequency separation between them exceeds this bandwidth. Detection of a signal component takes place based on the entire energy within a critical-bandwidth, whether it is tonal in nature, noise-like, or a combination of the two. Later workers have further refined this concept; Zwicker et al. [1957] examined this resolution bandwidth via various detection and masking experiments. Later Zwicker [1961] established 24 fixed critical-bands over the 20 Hz - 15 kHz frequency range.

1.2 The Use of Single Tone Masking Curves

Information on the masking effect of signal components is available primarily for single tones or bands of noise. As a result, coder design depends greatly on principles derived from these simple masking experiments. These typically generate masking curves of single high level component masking the presence of another smaller component and are quite useful because they can be used to derive an upper bound on the levels of permissible error signals due to the data rate reduction process. Since the masking effect varies significantly depending on whether the large level component or masker is tone-like or noise-like in character, the more demanding situation of sinewave masking curves are shown in Figures 1, 2, and 3. The figures present various 1/3 octave hearing thresholds when subjects are subjected to various levels of 100 Hz, 500 Hz, and 4 kHz sinewave maskers, as described by Fielder [1987]. For more information on the variation of the masking effect for tonal or noise signals, see Ehmer [1959].

The most appropriate way to examine masking phenomena is to perform a spectrum analysis based on critical-bandwidths. Since critical-band analyzers are not common, a good approximation can be made with the use of 1/3 octave bands; see Fielder [1987] for further details. These spectral analyses of masking are then used as a basis for the design of the coder filter bank structures and the methods to reduce the bit rate via word-length reduction.

The first observation from Figures 1-3 is that masking is generally greatest at the masker's frequency. This indicates that the coder design should concentrate error energy directly adjacent to the signal frequency. The next property the figures have in common is that the masking effect slowly decreases with increasing frequency separation, if the smaller signal is higher in frequency than the masker. The masking effect for signals at a 70 dB acoustic level may extend only a few octaves upward in frequency while higher level situations may produce six upward octaves of significant masking.

Looking at masking of signals lower in frequency than the masker shows a very different situation. For these signals, the masking effect falls off much more quickly. This is particularly evident for frequencies between 500 Hz - 2 kHz when evaluated in a dB per Hz fall-off from the masker frequency; in this frequency region the slope can be as steep as 100 dB per 350 Hz below 500 Hz (i.e. 90 dB/octave) and drop as deep as 40 dB within 1/2 octave. This rapid decrease in masking for components lower in frequency than the masker has significant consequences in coder design, and has been one of the primary reasons that data rate reductions of 4:1 or greater have awaited the practical availability of powerful DSP architectures which can practically implement the necessary complementary filter structures with sharp frequency characteristics that are suitable for music coders.

The differences in the masking characteristics versus frequency are also significant. In Figure 1 the masked threshold falls off only for frequencies above 100 Hz. The upward frequency fall-off in masking above 100 Hz is rapid on a dB per Hz basis, with a slope that is as much as 100 dB per 400 Hz. In the case of 100 Hz masking curves, it is important to note that a ratio of as much as 100 dB may be necessary between the 100 Hz masker and a resultant error component, if the error is to be inaudible. This means that any filter bank used by a rate reduction coder is most effective if its ultimate attenuation spans this 100 dB range. The masking curves of 100 Hz are typical for masking situations for maskers at or below 200 Hz.

The masking curves for 500 Hz, depicted in Figure 2, show a different situation. In this case there is a rapid reduction of downward frequency masking of up to 100 dB per 360 Hz, while having a much slower reduction at higher frequencies. In addition, high sound levels between 90-110 dB cause a very large masking effect at the second harmonic, causing the masking effect to be significantly extended upward in frequency. These 500 Hz curves are typical for the masking properties of midrange signals in the 500 Hz - 2 kHz region. Although not shown, at 2 kHz the slope of the masking curves have only 1/2 - 1/3 the slope of masking curves at 500 Hz, but the total fall-off has increased to 60 dB.

Figure 3 shows masking that is typical for high frequency signals. Masking for lower frequency error components falls off fast but not as fast a dB per rate as in the case of midrange signals. However, the total may exceed 70 dB for maskers at 8 kHz and above. As in the case of midrange signals, upward frequency masking reduces slowly with frequency but covers a more extended frequency range.

1.3 Temporal Masking and Time vs. Frequency Trade-Off

Sinewave masking experiments and the shape of masking curves derived from them indicate the requirements for the filter bank of a low bit rate coder under steady-state signals. Another requirement is the accommodation of human auditory characteristics during transient events. Although the frequency resolution for steady state sinewave signals is extremely sharp, the characteristics of auditory masking for transient events involves time resolutions on the order of a few milliseconds. The temporal characteristics of masking are important because the filter banks used for data rate reduction coders can disperse error signals in time. This spreading occurs because of the fundamental trade-off between temporal and frequency resolution of filters. For this reason, filter bank design typically involves a trade-off between these conflicting goals.

Just as in the case of the frequency characteristics of auditory masking under steady state signal conditions, there is a basic asymmetry in the characteristics of temporal masking. The masking of small signal components occurring during in time before a masker (i.e. backward masking) is substantially less than the forward masking effect in which the same small signals occur after the masker. Backward masking remains strong for about 4 milliseconds and disappears for time separations larger than 10's of milliseconds, while forward masking lasts approximately ten times as long. For further information on the temporal masking characteristics of the ear, see Carterette and Friedman [1978]. The temporal resolution characteristics of a filter bank used for data rate reduction of music signals should maximize the masking effect so that the largest data rate reduction induced errors are tolerated by the ear. Since a transient event can occur anywhere within the effective time window of a particular filter, this argues strongly for filter banks with time resolutions less than 4 milliseconds.

1.4 Filter Bank Design and Auditory Masking

The filter bank of a coder is the primary element that allows rate reduction to occur with minimal audible consequences. It does that by confining the error temporally and spectrally in such a way as to allow the greatest errors to occur. This spectral and temporal confinement must satisfy the following conditions. First, the ideal filter bank should have a frequency selectivity less than one critical band in any part of the audio band, have a fall-off rate of 100 dB per 360 Hz, with an ultimate rejection of 100 dB, and finally, have a temporal spreading effect of less than 4 msec. A filter bank which is easy and efficient to implement is also desirable. Unfortunately, the attainment of all the previously mentioned goals is extremely difficult and a compromise is necessary. As a result, further discussion will concentrate on the compromises and results for the low and moderate time delay AC-2 coders.

The design of the AC-2 coding technology is strongly influenced by the desire to keep the implementation as low in complexity as possible, while preserving coder effectiveness. For this reason, the AC-2 coders use Time Domain Aliasing Cancellation (TDAC), as developed by Princen and Bradley [1986]. This transform has the computational complexity advantages of an FFT and has excellent frequency selectivity characteristics. Unfortunately, the resultant filter bank is constant bandwidth, rather than having the varying bandwidths of the auditory system. This disadvantage of the TDAC can be overcome by approximating the nonuniform bandwidths of the human auditory system by grouping transform coefficients together to form sub-bands with bandwidths approximately that of the auditory system.

Consider first the TDAC filter bank for the moderate time delay coder, useful in applications where a low data rate is more important than low time delay. In this case the transform length is chosen to be 512 samples, which is found to be the best compromise between frequency and temporal selectivity. The resultant filter bank has a frequency selectivity that is sufficient for most of the audio band, while at the same time having a time resolution on the order of 10 msec. This compromise is acceptable since limitations in the temporal or spectral resolution are minor and can be greatly improved by a quantization process that allocates additional data to mitigate the increased audibility of errors during transient circumstances.

The other AC-2 coder is targeted for applications were low time delay is important, such as disk based storage applications requiring fine time resolution editing or for broadcast applications where an announcer may listen to the transmitted signal as a verification of proper system operation. Monitoring of the transmitted voice signal is problematic for the announcer if the time delay is too long, because it interferes with the cognitive process of speaking. The time delay at which speech difficulties begin to occur is not well defined, but 10 msec. appears to be a reasonable compromise, see Gilchrist [1990] for more details. The transform block length for this coder is set at 128 samples by this requirement and the resultant encode/decode delay is 8 msec.

This restriction in the block length has important consequences in the coder design because it moves the filter bank temporal- frequency resolution trade-off away from the optimal compromise. As a result, the frequency resolution is inadequate for masking the error signals for frequencies below 3 kHz. Insufficient frequency selectivity translates to either reduced audio quality or increased data rate. For this reason, this coder uses a higher data rate of 192 kbits/sec per channel. The time resolution of the system is 2.7 msec. and the resultant coder has excellent performance under transient conditions.

The loss of frequency selectivity to satisfy time resolution or computational complexity issues is very important in coder design. Figure 4 demonstrates this point by comparing the filter bank selectivity of three filter banks used in music coders to that of a masking curve for a 100 dB S.P.L. 1 kHz sinewave. This masking curve for 1 kHz was chosen since it is nearly a worst case for the selectivity requirements of a single tone situation. Both filter banks used in the two AC-2 coder are shown, and in addition, a typical uniform bandwidth sub-band filter band, having 750 Hz bandwidth, is included.

Examination of Figure 4 shows that none of the filter banks presented have ideal frequency selectivity when compared to this most demanding requirement of the human auditory system. The consequences of this fact is that all the coders implemented with these filter banks must either have a higher data rate than ideal or have lower sound quality. Inspection of this figure shows that the moderate delay version of AC-2 has the selectivity closest to that required, implying that little additional data rate is required to preserve sound quality. Next in selectivity is the uniform sub-band filter; the sharpness of the filter roll-off is excellent but it has the limitation that the filter's bandwidth is too wide for low frequency and midrange signals. This lack of frequency selectivity will result in quantization error that is spread over a wide frequency range (i.e. 550-1500 Hz) and must be accommodated by an increase in data rate. This increase in data rate results in an additional word-length requirement because the overall level of the error must be lowered until all of its spectrum lies below the masking curve. Finally, the short time delay AC-2 filter bank frequency selectivity is considered. In this case, additional data rate is seen to be required to mitigate the insufficient frequency selectivity of the low time delay filter bank. This, along with the desire for excellent multi-generation sound quality results in a data rate for this coder of 192 kbits/sec.

In conclusion, the examination of sinewave masking shows that the frequency selectivity of the moderate delay AC-2 is somewhat less than the worst case condition of 1 kHz masking. This indicates that its computationally efficient filter bank does not significantly limit the performance. The low time delay AC-2 coder selectivity is examined and shown to be too broad for use in the lowest possible data rate system. Fortunately, this increase in data rate is modest because the selectivity of the human auditory system is poorer than this filter bank over most of the audio band (i.e. 4 kHz - 20 kHz). One additional benefit of the short time delay AC-2 coder is that it possesses a temporal resolution substantially below that at which either forward or backward masking effects occur. The disadvantage of having too wide a filter bank bandwidth was demonstrated by the 750 Hz sub-band filter example.

1.5 Extension to Complex Signals

The use of simple stimuli masking models has determined the basic requirements of frequency and time resolution. This is done because there is not a widely accepted model of hearing for more complex signals. Unfortunately, real music signals are complex, so coder design must extend these simple masking models to the complex conditions of music signals. In the case of the AC-2 coding systems, simple stimuli masking principles are extended in a very conservative manner. Although many coding systems adaptively allocate most of the available data rate in a signal dependent manner to produce errors that are just below predicted masking, this was found to be an unnecessarily aggressive approach for applications with data rates at or above 128 kbits/sec.

The conservative approach of the AC-2 coder family is as follows: The appropriate TDAC transform filter bank is first combined with a trial quantization process that has a fixed number of bits assigned to each band, which are adjusted to simultaneously satisfy the masking requirements of simple and complex signals. Once this fixed allocation scheme is properly adjusted for optimal audible effect, a modest amount of the data responsible for this representation of the audio signal is removed and replaced by a smaller amount of adaptively allocated data, resulting in 20% or less data of this type. The advantage of the largely non-adaptive nature of most of the data is that problems in the extension of simple masking models are not nearly as serious as in the case of coders that have a more adaptive allocation strategy. This prevents serious audible mistakes from occurring: in fact the audible performances of the AC-2 coders without any adaptive bits are quite good.

This method of extension to more complex signals is evaluated and optimized by both objective and subjective means. This includes comparison of computed noise spectra with psychoacoustic masking threshold data, and conducting A:B listening tests. Subjects are asked to evaluate signals coded by hardware in real-time to facilitate exposing the coder to a wide variety of instrumental, vocal, and synthetic audio signals.

Although coder performance is more rigorously evaluated using complex music signals, many important features are revealed by the sinewave error spectrum. Figures 5 and 6 are a comparison of both coder's 1 kHz error spectra with a 100 dB S.P.L., 1 kHz masking curve. The moderate delay AC-2 coder results are shown in Figure 5 and those of the low delay AC-2 coder in Figure 6. Both figures give an indication of the worst case performance of the coder because the 1 kHz auditory selectivity is the most severe. These comparisons assume a consumer playback sound level at 108 dB peak acoustic level, being limited by the maximum loudness capabilities of typical home loudspeakers and amplifiers. In both figures, the error spectra are shown for coder operation with, and without, the adaptively allocated portion of the data. Both coding systems are interfaced to 16-bit ADC's and DAC's so the noise of the conversion process also is present.

Examination of Figure 5 shows that the error signal and converter noise under normal operation is just at or below audibility. The error spectra above 2.5 kHz for both situations are limited by the ADC/DAC noise floors and indicative of the 91 dB dynamic range of typical 16-bit conversion systems. The frequency region below 2.5 kHz is a result of coder operation and a significant deviation from an ADC/DAC noise floor results. In this region, the 1 kHz error spectrum under normal operation is substantially below the masked threshold curve, except for frequencies between 400 Hz - 700 Hz, where the error spectrum is comparable to the masked threshold. This indicates that a slight modulation noise may be audible, although in practice this has not been heard. The error spectrum shown without the adaptive portion of the data shows that modulation noise is now quite audible since the error spectrum is significantly above the masked threshold in the frequency range of 400 Hz - 1200 Hz. The generation of audible modulation noise indicates that the adaptive bit allocation process is necessary to preserve excellent sound quality. Notice that the error spectrum falls off less rapidly than that of the downward frequency portion of the 1 kHz masking curve.

This is exactly as predicted by the earlier discussion of the requirements of filter bank selectivity.

Figure 6 shows the same comparison for the low time delay AC-2 coder. In this case, the audio performance is essentially noise-free in normal operation, but limited by modulation noise without the adaptively allocated bits. This time the coder dependent part of the spectrum extends to 5 kHz and the extension of the range where the coder affects the noise spectrum is due to the more gradual frequency selectivity of the low time delay filter bank. Similar to the case of the moderate delay AC-2 coder, the normal operation spectrum slightly exceeds the masked threshold curve in the region of 400-600 Hz, indicating the presence of a small amount of masking noise. As before, actual listening tests determine that no modulation noise is audible. The use of adaptive bits is shown to be important since the situation with no adaptively allocated bits indicates the presence of substantial modulation noise. In this case, the error spectrum exceeds the masked threshold by 25 dB at 500 Hz.

2. AC-2 CODING ALGORITHM

In Section 1, some of the groundwork for audio coder design was established. In this section, we build upon this presentation by exploring the AC-2 coding algorithm in more detail. The description generally applies to all members of the AC-2 family; differences between the low and moderate delay versions are described where appropriate.

Figure 7 presents the block diagram of a generic AC-2 digital audio encoder. In the first stage of processing, PCM audio is buffered into frames of length N samples. Each new frame overlaps the previous one by 50%, i.e., the first N/2 samples in each frame are comprised of the last N/2 samples from the previous and present one. Consequently, each input sample is contained within exactly two consecutive frames. Next, the buffered samples are multiplied by a window function to reduce the effect of frame boundary discontinuities on the spectral estimate provided by the transform. The window also significantly improves the frequency analysis properties of the encoder.

The time-to-frequency domain transformation is based on evenly- stacked TDAC, consisting of alternating Modified Discrete Cosine (MDCT) and Modified Discrete Sine (MDST) transforms. A crucial advantage of this approach is that 50% frame overlap is achieved without increasing the required bit-rate. In a critically- sampled analysis technique such as TDAC, exactly N unique nonzero transform coefficients are generated on the average in an interval of time representing N input PCM samples. In TDAC, each MDCT or MDST transform of frame size N generates only N/2 unique nonzero transform coefficients, so critical sampling is achieved with 50% frame overlap. Any nonzero overlap used with conventional transforms (such as the DFT or standard DCT) precludes critical sampling, since each N-point transform generates N unique nonzero transform coefficients. Additionally, several memory and computation-efficient techniques are available for implementing the MDCT and MDST transforms.

TDAC is applied to model the auditory system by grouping adjacent transform coefficients into sub-bands for further decomposition and analysis. The number of coefficients per sub-band is computed a priori to approximate the nonuniform critical-bands. Transform coefficients within one sub-band are converted to a frequency block floating-point representation, with one or more mantissas per exponent, depending upon the sub-band center frequency. Each exponent represents the quantized peak log-amplitude for its associated sub-band. The exponents collectively provide an estimate of the log-spectral envelope for the current audio frame, computed on a critical-band frequency scale.

From a psychoacoustic perspective, the log-spectral envelope provides an ideal framework for estimating which sub-bands of a given audio frame are perceptually most relevant, and for ranking them in relative order of importance for dynamic bit allocation. Furthermore, the nonuniform frequency division scheme offers key advantages compared to one based on uniform-width filter banks. Accordingly, the AC-2 frequency division scheme reduces the need both for relying upon a complex masking model, and for using a second, higher-resolution filter bank in the encoder.

The dynamic bit allocation routine is completely feed-forward in nature and is constrained to produce a constant bit-rate as required for transmission applications. Bits are allocated in accordance with a set of deterministic rules derived from conservative use of single-tone masking curves. A portion of the routine employs a water-filling procedure in which sub-bands are ranked and allocated bits on a band-by-band basis.

The allocation routine provides step-size information for an adaptive quantizer. Each sub-band mantissa is quantized to a bit resolution defined by the sum of a fixed allocation and a dynamic allocation. The total fixed allocation for one frame outweighs the dynamic allocation in approximately a 4:1 ratio. For a given level of error protection overhead, this approach was found to provide more robust coding and error performance, since the number of most-significant mantissa bits is known a priori in the decoder. In the final stage of the encoder, exponents are multiplexed and interleaved with mantissa bits for transmission to the decoder. Optional error correction codes may be added at this step. The amount of overhead information reserved for error control coding can be adjusted to give greater or lesser protection depending upon channel error performance for a given application.

Serial bitstream formats can be optimized for the application. In the DP501/DP502 digital audio encoder/decoder products employing AC-2, two independent channels are interleaved in a regular pattern of alternating 16-bit segments. This format allows for straightforward demultiplexing of the encoder bitstream into separate channels, and for recombining monophonic bitstreams from different encoder units. Provision is also made for the insertion of a 1200 bit/s auxiliary data stream, algorithm identification bits, ADC overload status, and other information.

In the AC-2 decoder, shown in Figure 8, the input bitstream is demultiplexed and errors, if any, are corrected. The received log spectral envelope is processed in a stage identical to the encoder bit allocation routine, which generates step-size information for the adaptive inverse quantizer. The fixed and dynamically-allocated portions of each mantissa are concatenated to regenerate compressed transform coefficients. A sub-band block floating-point expander then linearizes the compressed transform coefficients and passes them to an inverse MDCT/MDST transform stage. After the inverse transformation, a window identical to that used in the encoder is used to post-multiply the reconstructed time-domain samples for each frame. Adjacent windowed frames are overlapped by 50% and then added together to reconstruct the PCM output.

Total coding/decoding time delay is determined by the frame size N, the manner in which frames are processed, and the processor speed. In the low-delay coder, input frames are processed one-by-one, resulting in a theoretical minimum total coding and transmission delay of 2.5N samples when employing infinitely-fast encoder and decoder processors. In actual practice, the delay increases to about 3N samples for fully-utilized (finite-speed) encoder and decoder processors. With a frame size of N = 128 samples and a sample rate of 48 kHz, a delay of 8 msec. is obtained. In the moderate-delay coder, two successive frames from one channel are buffered and processed jointly. In this case, the total delay when using fully-utilized processors is about 4N samples, which results in less than 45 msec. of delay at a sampling rate of 48 kHz and with N = 512.

3. HARDWARE IMPLEMENTATION

All coders within the AC-2 family have been optimized for very low hardware implementation cost. By today's standards, cost is ultimately measured by the die size and package cost of a custom VLSI implementation. Accordingly, the cost equation must not only include such traditional complexity measures as multiply-add count and RAM/ROM memory usage, but regularity of computation and minimum word-length requirements as well. Considerable attention has been given to structuring the computations in AC-2 to minimize VLSI implementation cost and simultaneously achieve the audio performance objectives. At a sampling rate of 48 kHz, the total number of multiplies and adds per second in a stereo AC-2 encoder is about 2.7 million. The decoder complexity is slightly lower. This compares to calculations by Reader [1991], estimating a total of about 35 million multiplies and adds per second for a straightforward implementation of a current generation sub-band encoder, and 16 million multiplies and adds per second in the decoder.

The low computational complexity can be attributed to several factors. First, the computational structures employed are highly regular in nature. Second, an efficient technique has been found for implementing the evenly-stacked TDAC transform by combining a core FFT routine with pre-twiddle and post-twiddle operations. Third, the nonuniform frequency division stage and log spectral energy representation enables the use of a low-complexity dynamic bit allocation routine. Finally, the use of functions which are inefficiently implemented on programmable DSPs or in custom-ICs, such as logarithms, square roots, and divides, have been found unnecessary. The only functions required are multiply, add, integer left/right shift, normalize, and compare.

3.1 General-Purpose Programmable DSPs

Programmable DSPs provide a flexible and expedient path to real-time algorithm development, and as such provide an attractive means for a first implementation. An early embodiment of AC-2 based on the Fourier transform was implemented using six Texas Instruments TMS32010s by Fielder [1989]. This work subsequently led to an implementation employing TDAC and based on the Motorola DSP56001, as detailed by Davidson et al. [1990]. In the latter case, a single 27 MHz chip could either encode or decode two independent channels. Recent improvements in software run-time efficiency have reduced this speed requirement to 20 MHz.

3.1.1 24-Bit Fixed-Point

Since it's inception in 1987, the Motorola DSP56001 has proven to be a capable platform for implementation of a wide variety of audio processing algorithms. This general trend has been supported by several audio compression implementations, including AC-2. The DSP56001's 24-bit data path, flexible addressing modes, and dual-accumulator arithmetic logic unit (ALU) are keys to its successful application in audio.

In particular for AC-2, we found that the 24-bit word-length was sufficient for all arithmetic tasks. Furthermore, no elaborate scaling or rounding procedures were required. The dynamic range of the implementation, as measured from PCM input to output, is 108 dB. This figure greatly exceeds the theoretically-achievable dynamic range of 16-bit ADC and DAC converters, and is commensurate with next-generation 18 and 20-bit converter technologies.

One of the more time-intensive processing blocks of those shown in Figures 7 and 8 is the inverse transform, which requires about 18% of the total DSP processing time. Surprisingly, however, the most time-intensive tasks are bit multiplexing and demultiplexing. This indicates that a custom IC could save significant ALU resources compared to a DSP if dedicated logic performed the multiplexing and demultiplexing. This topic is discussed further in Section 3.2.

3.1.2 16-Bit Fixed-Point

A study was made to determine the feasibility of implementing an AC-2 decoder on a 16-bit DSP chip. The motivation for this work was to identify a lower-cost platform for the implementation of an AC-2 decoder, while maintaining the flexibility of a programmable DSP. Our results indicate that current generation 16-bit DSPs, such as the Texas Instruments TMS320C5x, Analog Devices ADSP-2105, and Motorola DSP56116, are sufficiently powerful to implement a single-chip stereo encoder or decoder.

An analysis of finite word-length effects was conducted in part by modifying the real-time AC-2 DSP56001 software to emulate a reduced word-length processor. The data word-length was selected on-the-fly with switches. Coefficient word-lengths could also be varied. This approach allowed us to independently adjust, and jointly minimize, the data and coefficient word-lengths in each processing stage of the coder. The real-time variable word-length simulation served as a valuable tool for rapid objective and subjective evaluation of finite precision arithmetic effects.

Figure 9 presents a plot of the spectral error between an original and a coded 100 dB S.P.L., 100 Hz sinewave as processed by both 24-bit and a 16-bit ALUs in the decoder. Results from the moderate-delay decoder are shown since arithmetic round-off noise in the inverse transform is highest for long frame lengths; round-off noise in the low-delay coder is more than 6 dB lower. The idle channel noise produced by 16-bit ADC and DAC converters is included to show when the coder is limited by the conversion process. The low frequency sinewave represents a demanding test signal since minimal masking of the 4 to 6 kHz region occurs, where the ear's hearing threshold is low.

At frequencies below 500 Hz, noise introduced by transform coefficient quantization dominates arithmetic round-off noise. This region is perceptually insignificant because both noise curves are below the masking curve. Above 2 kHz, round-off noise for the 16-bit ALU significantly exceeds the masking curve, indicating that 16-bit single-precision (SP) arithmetic is inadequate.

Most of the noise shown in Figure 9 is generated during the inverse FFT computation of the inverse MDCT/MDST transform computation. Therefore, conventional techniques for reducing round-off error in fixed-point FFTs apply, such as those described by Meyer [1989]. We found that the combination of dynamic scaling between IFFT stages, optimal rounding, and optimal placement of quantizers in the butterfly produced a significant, but still insufficient, reduction in round-off noise. Furthermore, such techniques may impose a three-fold increase in IFFT butterfly computation time within a general-purpose DSP.

Based on these results, a preferred approach is to employ an extended-precision (EP) scheme based on 16 x 32-bit multiplies, which for many 16-bit DSPs results in a fixed two-fold increase in butterfly computation time, and provides a digital noise floor which is more than 40 dB lower than that obtainable with 24-bit SP multiplies. All other processing stages of the decoder can be implemented with 16-bit SP arithmetic. The minimum required DSP clock speed using 16-bit EP is only about 18% higher than the equivalent rating for a 24-bit fixed-point or 32-bit floating-point device.

3.2 Full-Custom VLSI

In order for an audio processor to be utilized in high volume applications, the device cost must usually be low. Since programmable DSP chips frequently contain more hardware logic than required for a given application, we have considered the design of a special-purpose VLSI architecture for implementing an AC-2 decoder. The architecture is capable of implementing any of the coders in the AC-2 family with one IC.

The architecture consists of three sections: a bit demultiplexer, a quantizer step-size control, and an inverse transform and reconstruction processor. The chip inputs are a serial bitstream and data clock, and the output is one or more 20-bit PCM digital audio channels. The bit demultiplexer performs such functions as data de-scrambling and bit de-interleaving. The demultiplexer directs the unpacked exponent data to the quantizer step size control, and the unpacked fixed and adaptive mantissa bits to a dedicated state machine/barrel shifter. The quantizer step size control, composed of a simple programmable microcontroller, processes incoming exponents and directs the state machine and barrel shifter to concatenate fixed and adaptive transform coefficient mantissa bits. The reconstruction processor performs either an IMDCT or IMDST, producing one frame of PCM samples. These samples are then windowed and overlap/added with the previous windowed block of PCM data to reconstruct audio samples. Since the multiply-add rate of the audio synthesis stage is quite low, a bit-serial multiplier has been employed. The serial multiplier requires significantly less chip area than a single-cycle array multiplier of the same word-length.

4. CONCLUSIONS

Adaptive transform coding of audio signals with AC-2 technology offers a high-quality, low complexity approach for data rate reduction of professional grade audio. Two 20 kHz bandwidth examples of the AC-2 coding family have been discussed, providing 4:1 and 6:1 bit-rate compression at low and moderate time delays, respectively. The excellent sound quality and computational ease of implementation of the AC-2 technology make it a natural candidate for broadcast, computer multimedia, and digital storage applications. The 128 kbits/sec. data rate of the moderate delay coder make it very appropriate for Digital Audio Broadcast and High Definition Television applications. The low delay coder is optimized for music material contribution applications (i.e. studio to transmitter and contribution quality links) requiring excellent multi-generational sound quality and a time delay acceptable for off-air monitoring during voice announcing.

The performance of these systems has been quantified by examination of simple stimuli masking models which have been the driving force shaping the design of the employed filter bank structures. Sinewave masking models have been used because a comprehensive and complete model for complex signals is not widely agreed upon. As a result, extension of the simple models is necessary for the design of practical coding systems. It was shown that the AC-2 family used a conservative extension process which resulted a relatively small amount of adaptively allocated data. As a consequence, these coder techniques were robust with respect to difficult program material. Other benefits created by this approach were a relative insensitivity to the effects of data-stream errors and low computational complexity.

Issues of computational complexity and practical implementation were discussed in some detail. It was shown that the AC-2 coder family is straightforward to implement at 128 and 192 kbits/sec. In particular, implementation of a stereo encoder or decoder was readily accomplished in one 20 MHz Motorola DSP56001. It was also shown that a practical modification of the frequency division algorithm permitted the realization of full fidelity realizations on 16-bit fixed-point DSP chips. A custom approach was also presented. It was shown that the AC-2 algorithms lend themselves well to dedicated chip hardware because of their reliance on simple shift operations and a low-complexity bit allocation strategy.

In conclusion, the AC-2 coder family represents one of the most cost effective solutions to very high-quality music coding applications at a 4:1 to 6:1 compression ratio. Although only two coders with data rates of 128 and 192 kbits/sec. were discussed, this technology can be applied to other sample rates, lower data rates (i.e. 64 kbits/sec.), and other signal bandwidths as well.

5. REFERENCES

K. Brandenburg, H. Gerhaeuser, D. Seitzer, and T. Sporer, "Transform Coding of High-Quality Digital Audio at Low Bit Rates: Algorithms and Implementation," Proc. of IEEE Intl. Conf. on Comm., Atlanta, pp. 932-936, April 1990.

E. C. Carterette and M. P. Friedman, Handbook of Perception, pp. 305-319, Academic Press Inc., New York, 1978.

G. Davidson, L. Fielder, and M. Antill, "Low-Complexity Transform Coder for Satellite Link Applications," presented at the 89th Convention of the Audio Engineering Society, preprint 2966, Sept. 1990.

R. H. Ehmer, "Masking of Tones vs. Noise Bands," J. of the Audio Eng. Soc., vol. 31, pp. 1253-1256, Sept. 1959.

L. D. Fielder, "Evaluation of the Audible Distortion and Noise Produced by Digital Audio Converters," J. Audio Eng. Soc., vol. 35, pp. 517-535, July/August 1987.

L. D. Fielder, "Low-complexity Transform Coder for Music Applications," presented at the IEEE ASSP Workshop on Appl. of Sig. Proc. to Audio and Acous., New Paltz, New York, session 2- 1, Oct. 1989.

H. Fletcher, "Auditory Patterns," Reviews of Modern Physics, vol. 12, pp. 47-65, Jan. 1940.
N. H. C. Gilchrist, "Delay in Broadcasting Operations,"
presented at the 90th Convention of the Audio Engineering Society, preprint 3033, Feb. 1991.

J. D. Johnston, "Transform Coding of Audio Signals Using Perceptual Noise Criteria," IEEE J. on Sel. Areas in Comm., vol. JSAC-6, no. 2, pp. 314-323, Feb. 1988.

R. Meyer, "Error Analysis and Comparison of FFT Implementation Structures," Proc of IEEE Conf. on Acous, Speech and Sig. Proc., Glasgow, pp. 888-891, 1989.

J. P. Princen and A. B. Bradley, "Analysis/Synthesis Filter Bank Design Based on Time-domain Aliasing Cancellation," IEEE Trans. on Acous., Speech, and Sig. Proc., vol. ASSP-34, no. 5, pp. 1153-1161, Oct. 1986.

C. Reader, "Comments on Audio CD and Analysis of Audio Complexity," ISO/MPEG Document ISO/IEC/JTC1/SC2/WG11/MPEG91/082, May 20, 1991.

E. F. Schroeder, H. J. Platte, and D. Krahe, "MSC: Stereo Audio Coding with CD-Quality and 256 kbits/sec," IEEE Trans. on Cons. Elec., vol. CE-33, no. 4, pp. 512-519, Nov. 1987.

G. Stoll and Y. F. Dehery, "High-Quality Audio Bit-Rate Reduction System Family for Different Applications," Proc. of IEEE Intl. Conf. on Comm., Atlanta, pp. 937-941, April 1990.

E. Zwicker, "Subdivision of the Audible Frequency Range into Critical Bands," J. of the Acous. Soc. of Am., vol. 33, no. 2, p. 248, Feb. 1961.

E. Zwicker, G. Flottorp, and S. S. Stevens, "Critical Band Width in Loudness Summation," J. Acous. Soc. of Am. vol. 29, pp. 548- 557, May 1957.

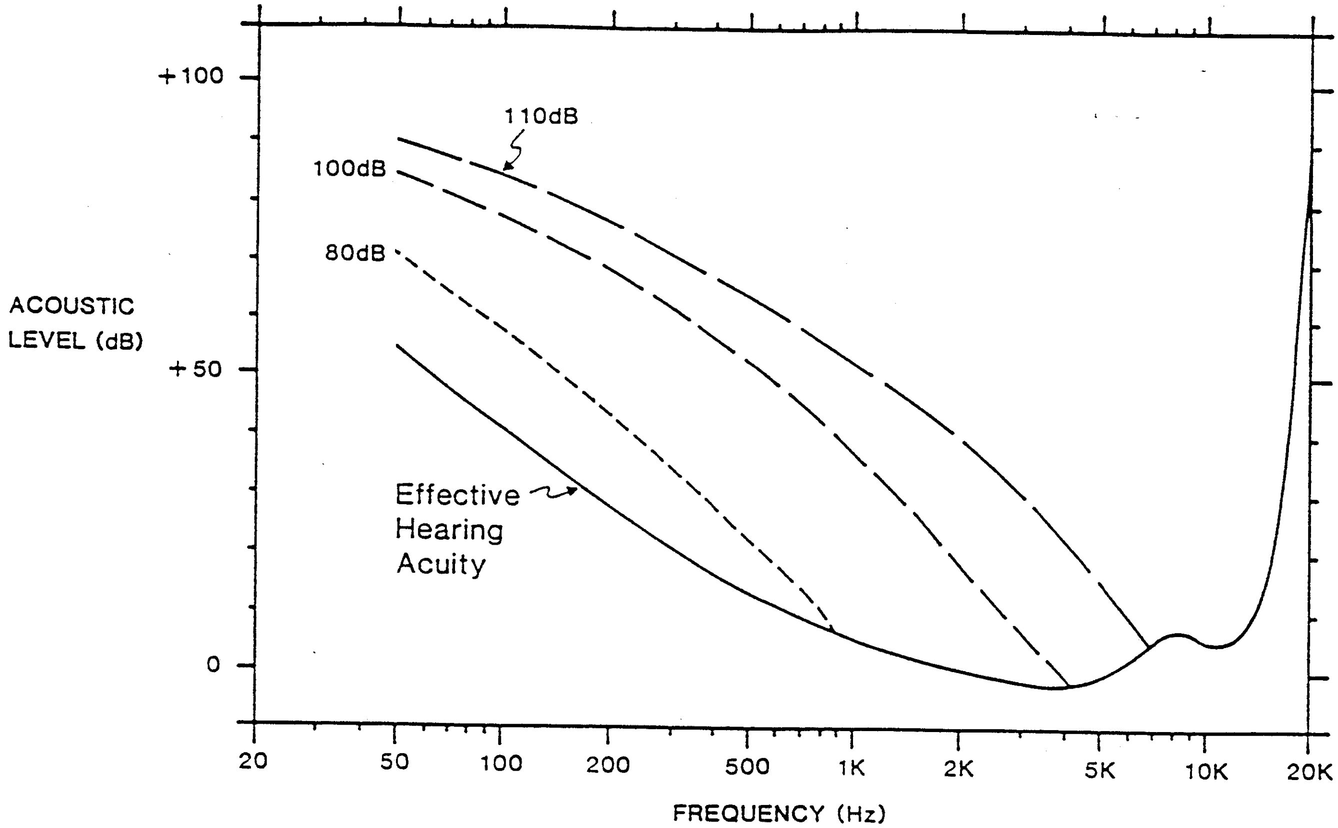

Fig. 1. 100 Hz Masked Threshold Curves (0 dB = 20 micropascals)

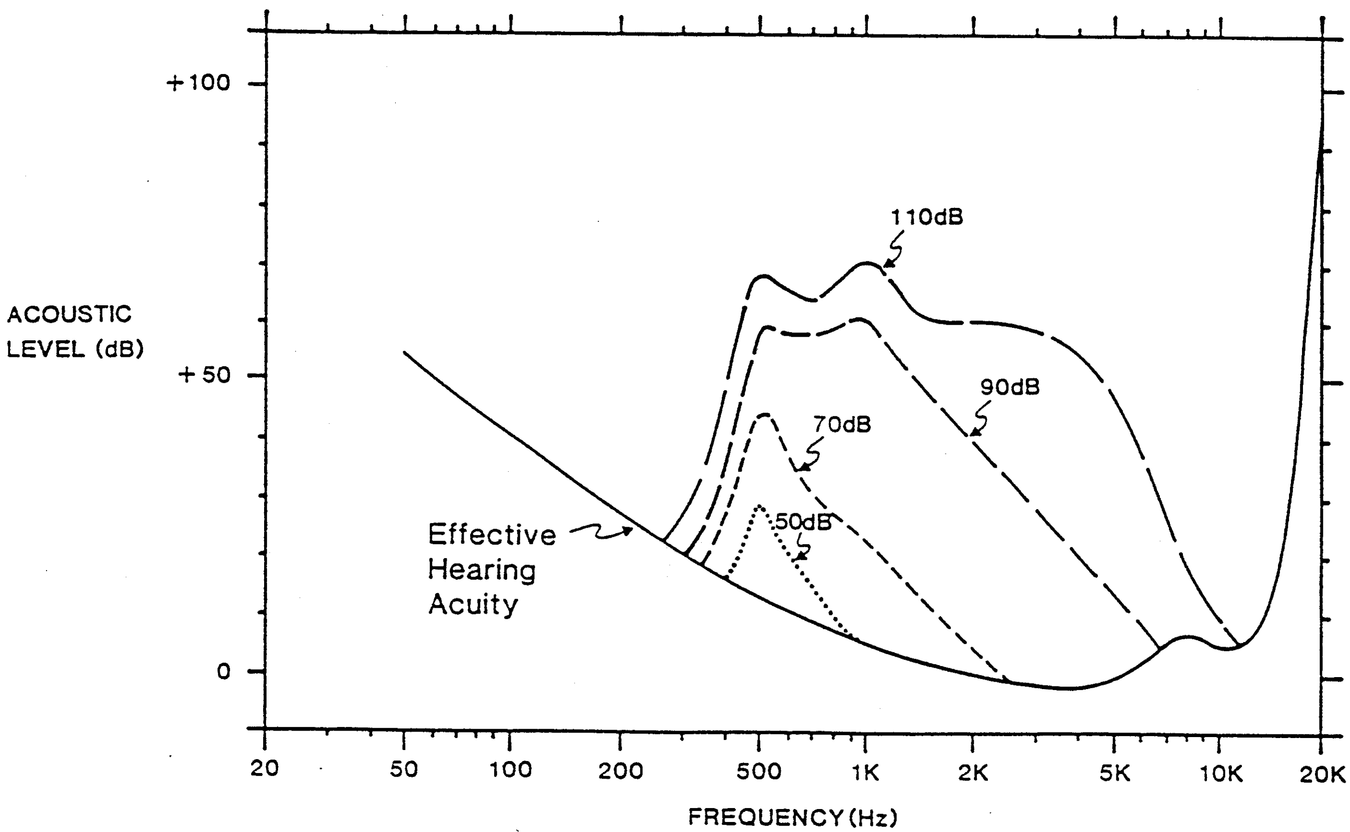

Fig. 2. 500 Hz Masked Threshold Curves (0 dB = 20 micropascals)

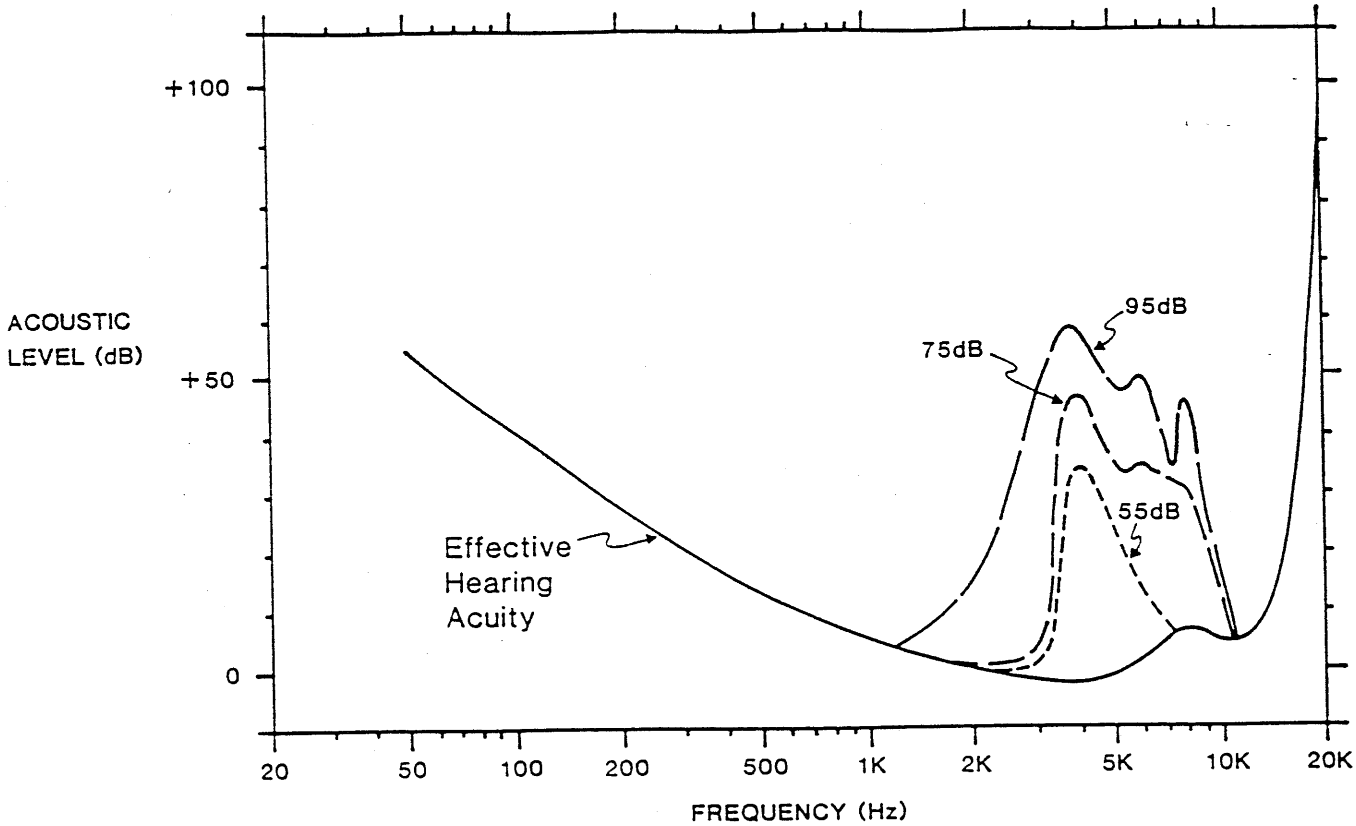

Fig. 3 4 kHz Masked Threshold Curves (0 dB = 20 micropascals)

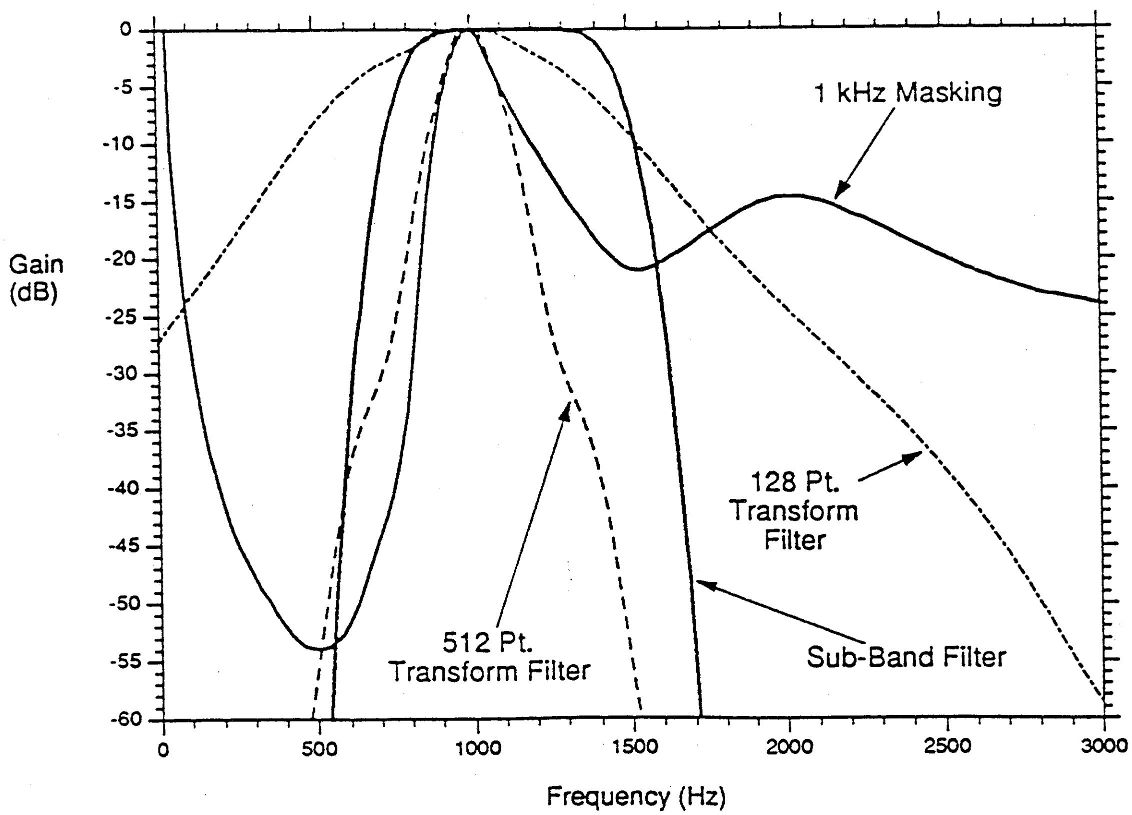

Fig. 4 Comparison Between Various Filter Banks and 1 kHz Human Auditory selectivities

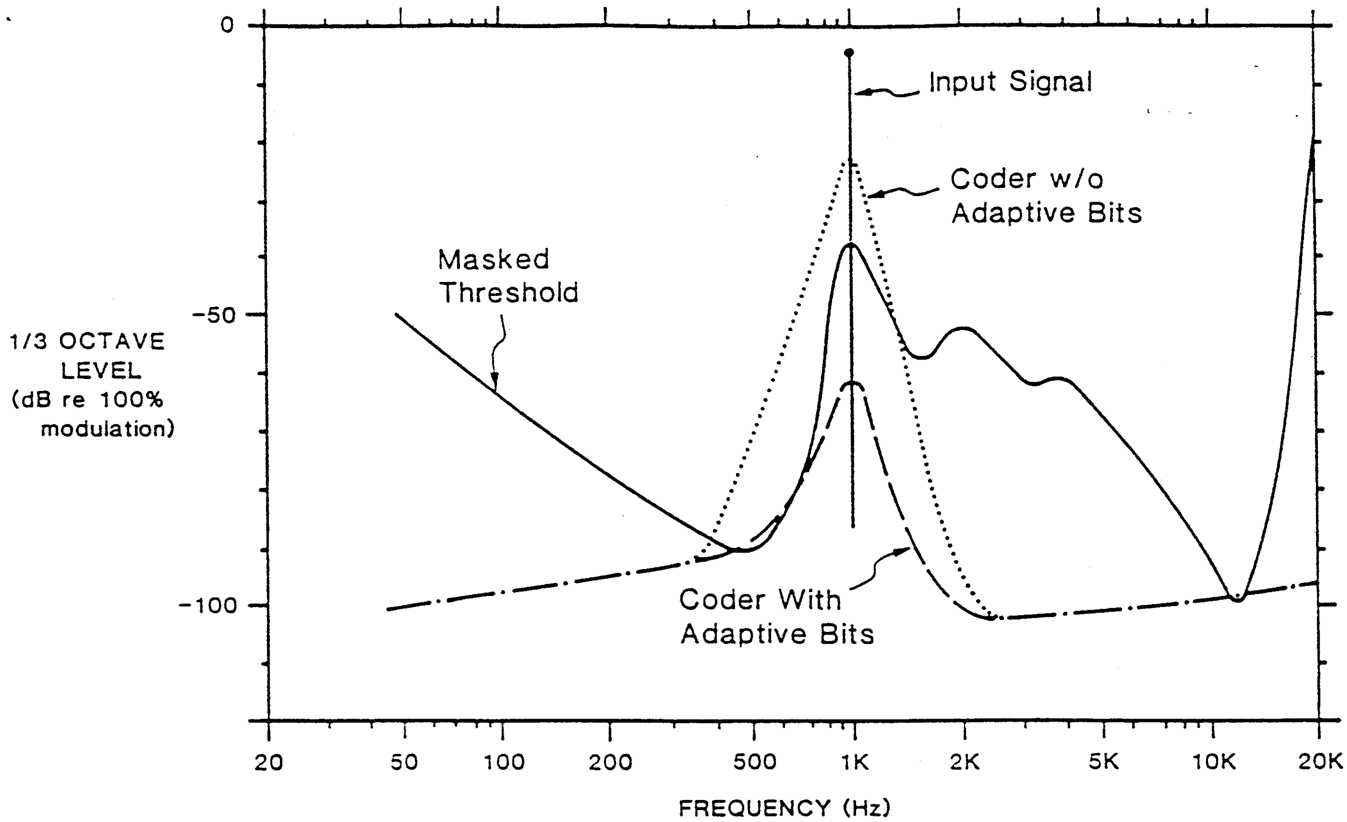

Fig. 5 Moderate Time Delay AC-2 Coder Performance with the Application of a 1 kHz Sinewave

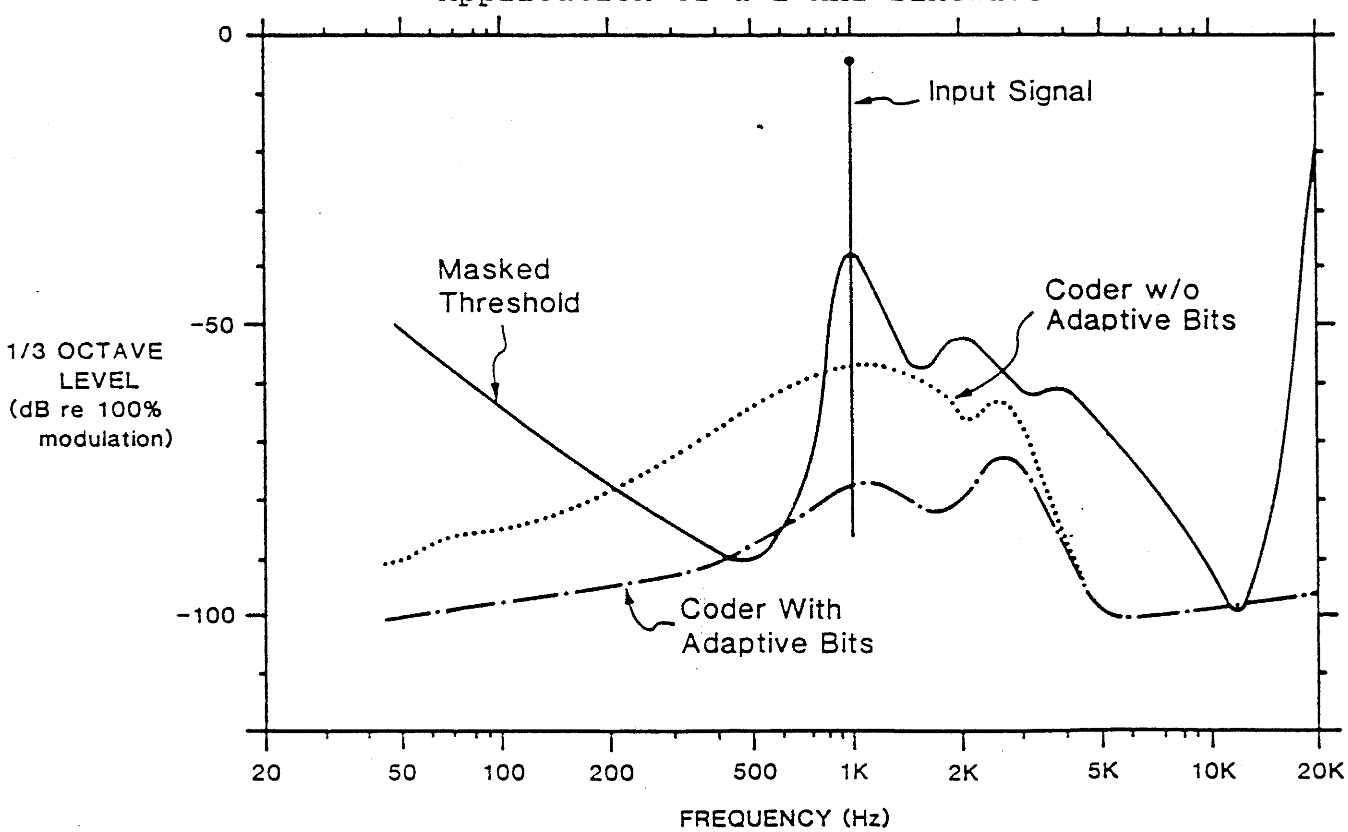

Fig. 6 Low Time Delay AC-2 Coder Performance with the Application of a 1 kHz Sinewave

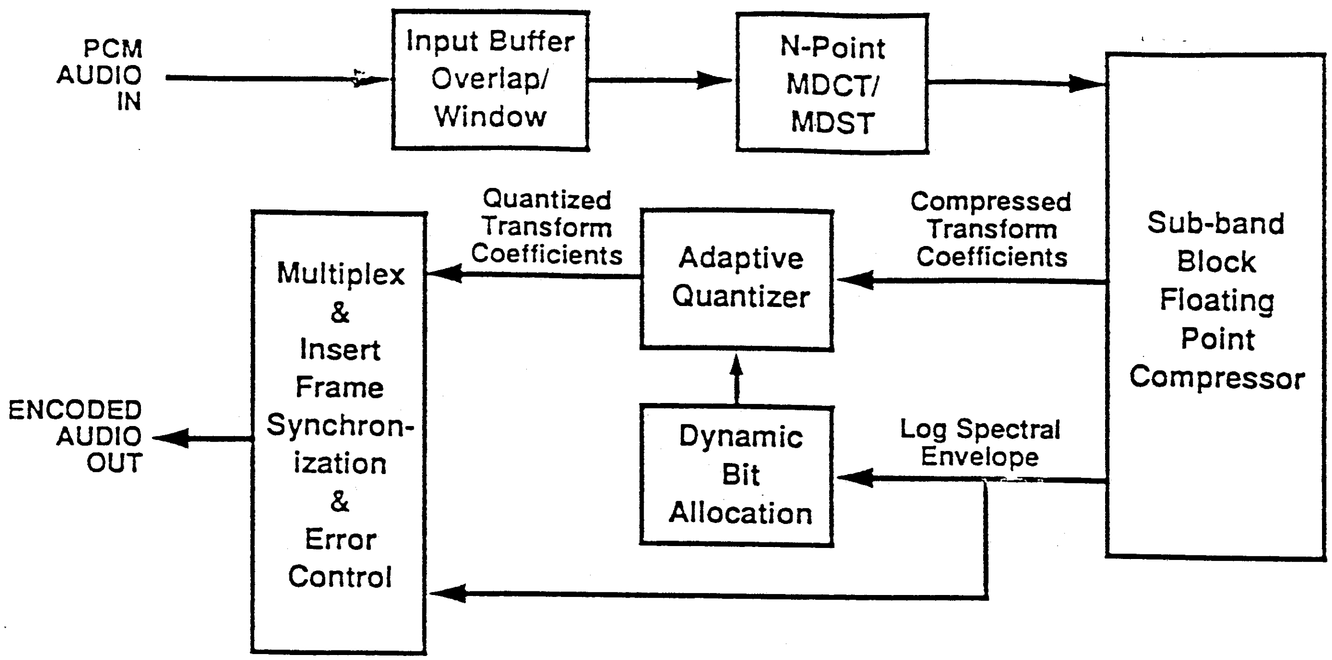

Fig. 7 AC-2 Digital Audio Encoder Family Block Diagram

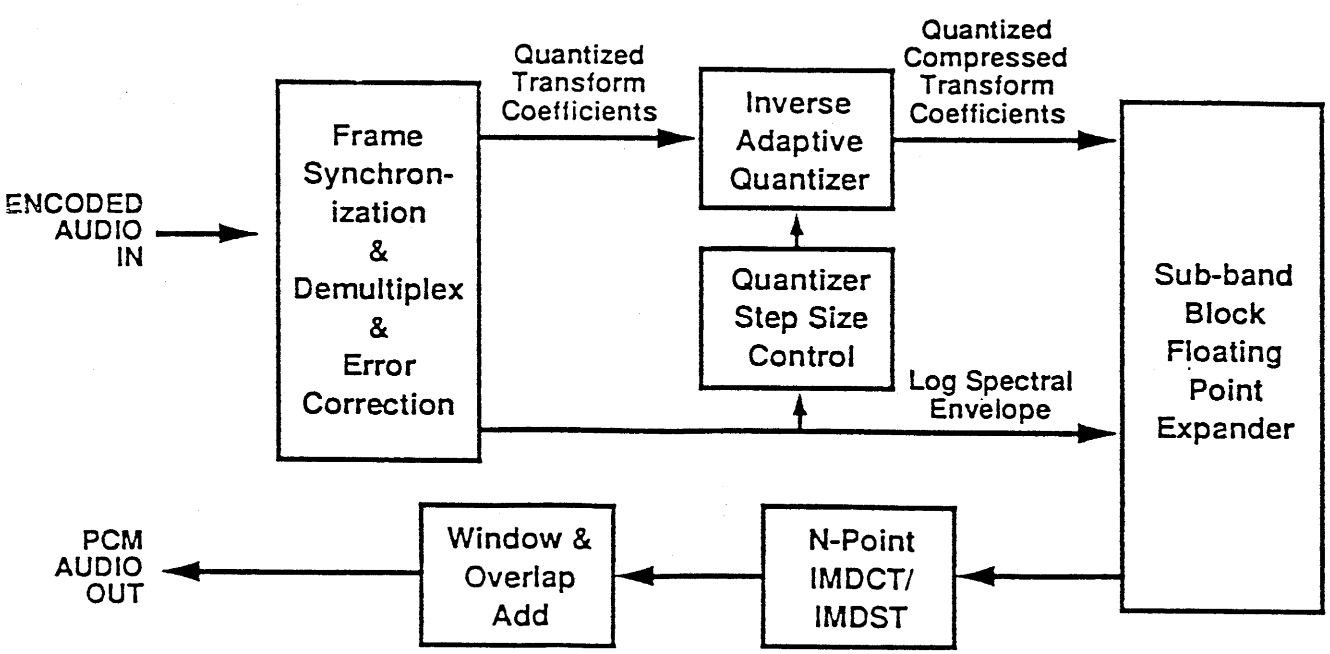

Fig. 8 AC-2 Digital Audio Decoder Family Block Diagram

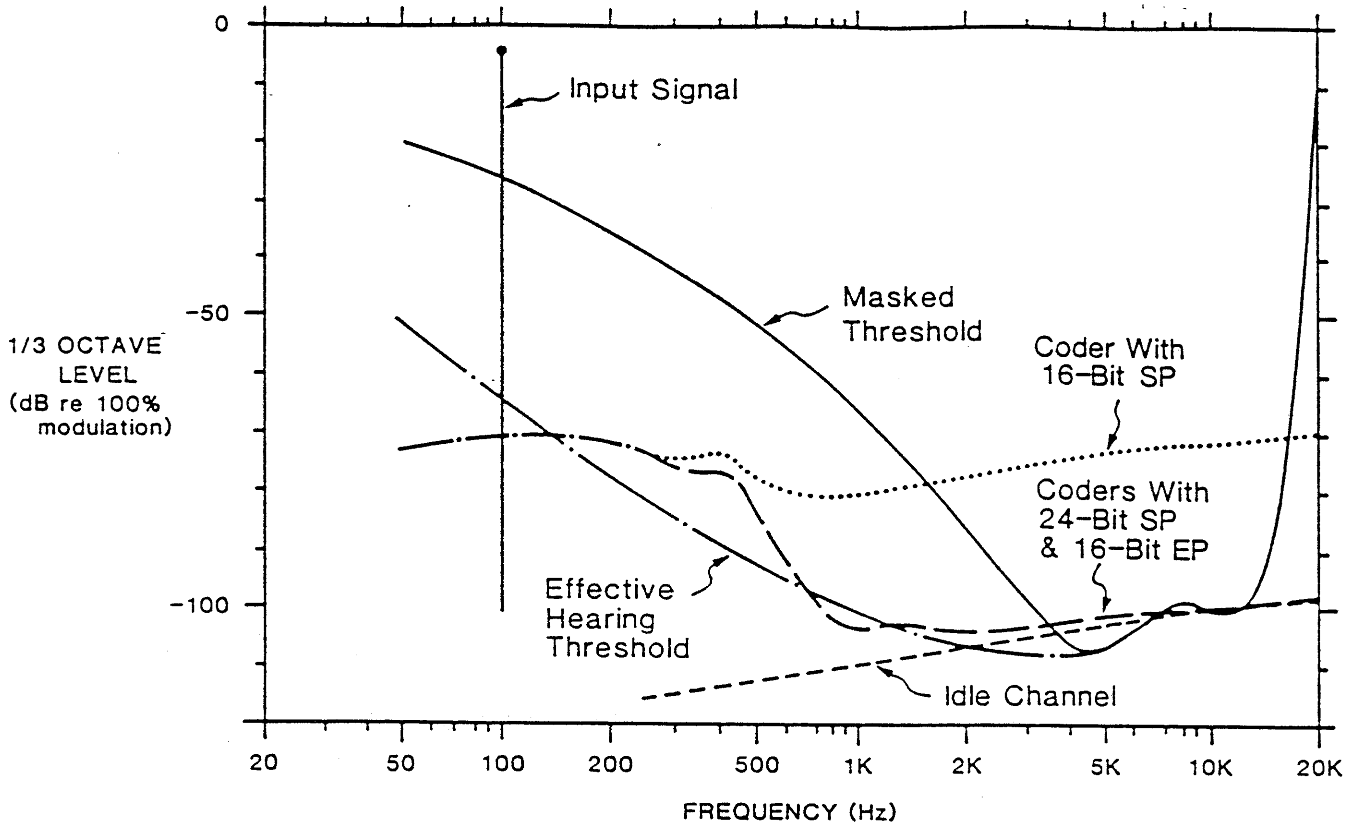

Fig. 9 Comparison of Filter Bank Arithmetic Noise of Single Precision (SP) 16-Bit Arithmetic Versus Extended Precision (EP) 16-Bit and Single Precision 24-Bit Arithmetic