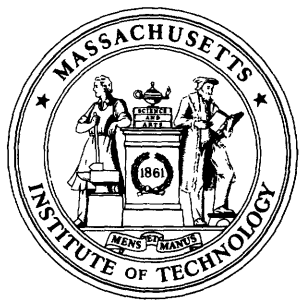

A Receiver-Compatible Noise Reduction System

RLE Technical Report No. 553

April 1990

Matthew M. Bace

Research Laboratory of Electronics
Massachusetts Institute of Technology
Cambridge, MA 02139 USA

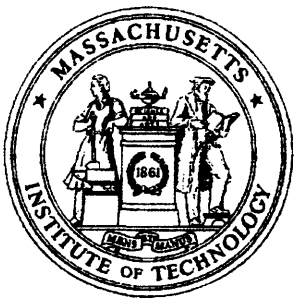

A Receiver-Compatible Noise Reduction System

RLE Technical Report No. 553

April 1990

Matthew M. Bace

Research Laboratory of Electronics
Massachusetts Institute of Technology
Cambridge, MA 02139 USA

This work was supported in part by the National Science Foundation under Contract No. MIP 87-14969 and in part by the Maryland Procurement Office under Contract No. MDA 904-89-C-3009.

A Receiver-Compatible Noise Reduction System

by

Matthew M. Bace

Submitted to the
Department of Electrical Engineering and Computer Science
on January 31, 1990 in partial fulfillment of the requirements
for the Degree of Master of Science

Abstract

This thesis presents an assessment of the feasibility of applying the techniques of two-channel processing and adaptive modulation to an NTSC-standard television signal in a receiver-compatible manner to reduce channel noise. As a result of this study, a receiver-compatible noise reduction system has been developed. The primary idea behind the system is to process the television signal on a frame-by-frame basis, dividing each frame into low and high spatial frequency components. The low-frequency component of each frame is subsampled, quantized, and transmitted digitally on a side channel. The high frequency component is adaptively modulated and added back to a compressed version of the low-frequency component to produce a reciever-compatible signal to be transmitted on the normal analog channel. An advanced receiver can use the information in the digital side channel, along with the received analog signal to produce a picture devoid of most channel degradations. A standard receiver can still receive an acceptable, but slightly contrasty, picture.

The major emphasis of this work is the optimization of the parameters of the system such as block size, interpolation scheme, and quantization scheme to achieve the highest levels of both compatibility and noise reduction. The results of this effort indicate that there is a very strict trade-off between compatibility and noise reduction in the selection of most of the system parameters. Achieving a significant level of noise reduction on an advanced receiver while preserving an acceptable level of performance on a standard NTSC receiver is extremely difficult, even in favorable environments. As a consequence as this trade-off, a transitional scheme is proposed in which the parameters of the system are slowly varied from those producing a high level of compatibility to those producing a high level of noise reduction.

Thesis Supervisor: Jae S. Lim

Title: Professor of Electrical Engineering

Acknowledgments

I would like to thank all of the people who have contributed, directly or indirectly, to the completion of this thesis.

In particular I would like to thank my thesis advisor, Professor Jae Lim, for the comments and suggestions he offered throughout the course of this research. Without his guidance, work on this thesis would not have even begun. I would also like to thank all of the members of Digital Signal Processing Group and of the Advanced Television Research Program at M.I.T. They were always willing to listen to me and offer advice or assistance whenever I needed it. These people have made my time at M.I.T. an enjoyable and rewarding experience.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation

In recent times, the effort devoted to improving the quality of broadcast television has increased dramatically. There have been numerous proposals for High-Definition (HDTV) and Extended-Defintion (EDTV) Television Systems. These proposed systems have take a wide varety of forms, ranging form those that are compatible with the existing (NTSC) standard [4,7,20,21,24], to those that augment or exist in cooperation with the existing standard [6,10], to those that are entirely incompatible [1,20,21]. Many interesting ideas have been introduced, including techniques to improve the spatial and temporal resolution, motion rendition, color, and sound. Few proposals, however, have included any provisions for combating the effects of channel noise.

For terrestrial broadcast, the generally low quality of the available transmission channels is the greatest obstacle to improved picture quality. In fact, many viewers would consider studio quality NTSC a significant improvement over what they now receive on their television sets at home. There are many degradation that can be imposed on a TV signal as it is transmitted through the atmosphere. In large urban areas, multipath is usually the greatest and perceptually most noticeable

degradation. In areas where the television bands are densely packed or where the transmitters are close together, intersymbol interference can significantly degrade the signal. Various atmospheric conditions and man-made electrical interferences can cause additive random noise and frequency distortion, which further degrade the signal quality.

The effect of these signal distortions on the received picture is so great that most of the above proposed improvements in resolution, color, and motion rendition will be almost completely lost. An effective means of noise cancellation should therefore be the top priority in any television system designed for terrestrial broadcast.

Several of the currently proposed advanced television systems do include measures for combating the effects of channel degradations. The Zenith system [10] uses nonlinear companding and time dispersion of "peaky" signals to achieve some level of noise reduction. The MIT Channel Compatible System (MIT-CC) [20,21] introduces several more sophisticated noise reduction techniques, including adaptive modulation, adaptive channel equalization, and scrambling. Neither of these systems however, produces a signal that is compatible with the existing NTSC receivers (although the Zenith signal is intended to be simulcast with an NTSC-encoded signal).

On the other hand, none of the receiver-compatible schemes make any provision at all for noise cancellation. In fact, many receiver-compatible schemes, such as the Advanced Compatible Television (ACTV) System [7] and the systems of Fukinuki [4] and Yasumoto [24] from which the ACTV System is derived, may actually exacerbate the noise problem by using additional subcarriers and adding extra information in the vertical and horizontal blanking intervals without any special protection. These additional subcarriers are extremely vulnerable to degradation due to multipath.

Thus, it is the goal of this research to determine whether it is possible to achieve any level of noise reduction in a receiver compatible manner. The major idea around

which this work revolves is that it is possible to use signal level control to achieve noise reduction in such a way so that it is compatible with existing NTSC receivers. If such a receiver-compatible noise reduction system were to function ideally, then one would be able to receive a virtually degradation-free picture on an advanced receiver, while it would still be possible to receive a subjectively acceptable picture with a standard NTSC receiver.

1.2 Design Limitations

The desire to maintain compatibility with existing receivers and the power restrictions of the broadcast television industry place some fairly rigid constraints on what can be done to improve the quality of the current television picture. These constraints must be observed by any new television system that is to be considered seriously.

For instance, at the transmitter end of the system, the major limitations are bandwidth and power. No system will have much merit if it is not at least channel compatible. In the context of this work, channel compatible means that the amount of analog bandwidth required for transmission does exceed the 6 MHz currently allotted by the FCC. Furthermore, for any new television system to have any appeal to broadcasters, it must not use any more power at the transmitter than is currently being used. Broadcasters are most concerned with peak-to-peak power, or dynamic range, rather than average (RMS) power. Another consequence of this power restriction is that any proposed system must be able to perform well in the generally low SNR environments common in terrestrial transmission of television signals. In urban areas, it is not uncommon for SNRs to be as low as 24 dB.

The restrictions imposed at the receiver end of the system are primarily economic. Since manufacturers and consumers both desire to keep the prices of advanced television sets relatively low, there is a limit on the amount of hardware that

can be put in an "improved" or "advanced" receiver. This restriction means that processing at the receiver must remain fairly simple and should not require much storage (*i.e.* frame buffers). In the case of any receiver-compatible system, however, the most important receiver-end restriction is that the quality of the picture on conventional NTSC receivers not be perceptibly degraded by any changes made to the signal format. Because the NTSC signal occupies almost all of the available spectrum (although not very efficiently), this constraint implies that only a very narrow range of alterations may be made to the NTSC format without causing objectionable degradations to the picture received on an NTSC receiver.

1.3 Overview

The majority of this thesis will be spent describing the design of the receiver-compatible noise reduction system; however, some information must be presented first to provide the proper context. In Chapter 2, some of the pertinent properties of the human visual system will be discussed. Special attention will be given to the phenomena of spatial and temporal masking, upon which adaptive modulation algorithms rely heavily. Next, in Chapter 3, two-channel processing and adaptive modulation are introduced. A receiver-compatible version of adaptive modulation is described and extended into two dimensions.

Chapter 4 begins the discussion of how the receiver-compatible noise reduction system was designed. The NTSC format and its limitations are discussed briefly. An estimate is made of the size of the digital side channel that will be available to the receiver-compatible system. Then, the selection of the various system parameters is discussed. First the luminance processing is considered. Block sizes, bit allocation, and dynamic range restrictions are discussed and related to each other and to overall performance. Techniques for processing the chrominance components are then discussed. Finally, problems that arise when adaptive modulation is applied

to video, rather than still frames, are addressed. The performance of this system is analyzed in Chapter 5. Finally, Chapter 6 provides some concluding remarks and several ideas about how the system might be further improved.

Chapter 2

The Human Visual System

To understand how one may improve upon the current television system, it is necessary to examine the human visual system and how its characteristics are already exploited by the current television systems. The term "human visual system" is meant to include not only the low-level components of the eye such as the pupil, lens, and rods and cones, but also the higher-level components such as the optic nerve, the optic chiasma, and the visual cortex of the brain [8].

For the purpose of television system design, the most important characteristics of the human visual system are the spatial and temporal frequency responses and the spatial and temporal masking effects. As one can see from Figure 2.1, the response of human visual system has a bandpass nature in both the temporal and spatial domains. The peak sensitivity of the human visual system appears to be at about 2 Hz temporally and 2 cycles per degree (cpd) spatially [18], although it exhibits a fairly large response for temporal frequencies over the range of 1–30 Hz and spatial frequencies over the range 0.3–30 cpd. Below these frequency ranges, the human visual system acts as a differentiator, both temporally and spatially. Thus, the human visual system automatically sharpens overly blurry images. Above these frequency ranges, the human visual system functions as an integrator, so that overly busy images are automatically blurred. One can also see that in the spatial domain,

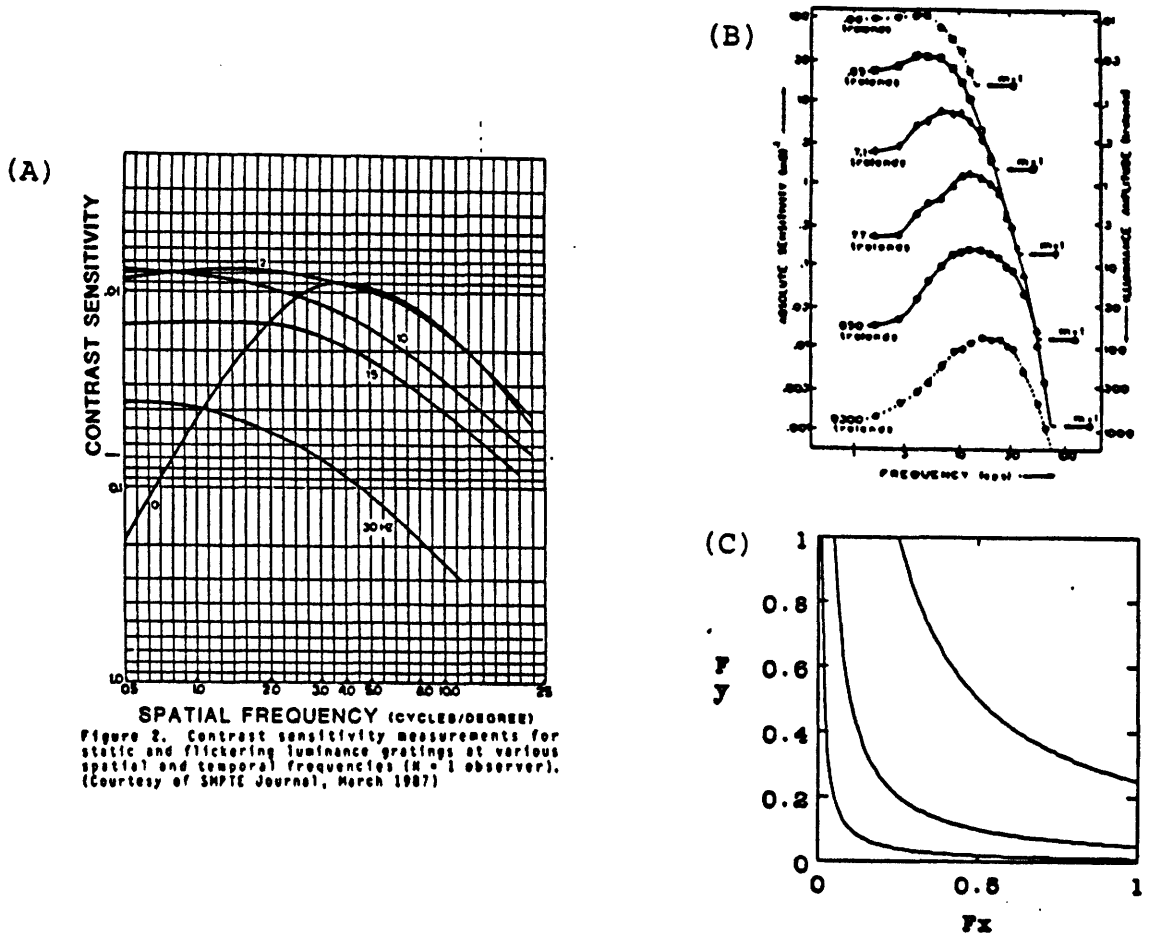

Figure 2.1: Spatial and temporal frequency response of the human visual system: (a) Spatial frequency response in horizontal (or vertical) direction. (b) Temporal frequency response. (c) Equisensitivity contours of the two-dimensional spatial frequency response.

the response of the human visual system is not isotropic. In particular, the human visual system is much more sensitive to high spatial frequencies in the horizontal and vertical directions than in the diagonal directions.

Another important feature of the human visual system is that its sensitivity to foreground objects depends heavily on what is in the background. One example of this phenomenon is Weber's Law, which holds that the differential luminous intensity required to cause a noticeable difference in the foreground is directly proportional to the intensity of the background. Weber's Law holds over a fairly large

dynamic range of intensities, so that one could model the contrast sensitivity of the human visual system as roughly logarithmic.

Two more examples of this phenomenon are spatial and temporal masking.. The term "masking" refers to the fact that fine spatial details (low-amplitude high spatial frequency components) are "masked" by large-scale spatial or temporal transitions in the luminance. Spatial masking is the phenomenon that the human visual system is less sensitive to fine spatial detail in regions where there is an abrupt and large-scale change in luminance than in regions of constant or slowly-varying luminance. Thus, fine spatial detail is less visible near edges and in highly detailed regions than in uniform background regions. Temporal masking is a similar effect. Fine spatial detail is less visible during the time intervals in which sudden transitions in the luminance occur (as is the case when there is rapidly moving region of widely varying luminous intensity) than during periods of relatively constant luminance (as is the case with a still picture).

These masking effects have been well-documented, most notably by Glenn [5,6], and have even been measured quantitatively. Glenn also points out that in addition to fine detail, color detail is also masked. According to his results, a time interval of roughly 200 msec is required to discern both chrominance and high-frequency details following a sudden change in the viewed scene. This lag time can be observed in Figure 2.2, which shows how sensitivity to high spatial frequencies varies in the vicinity of a sharp spatial or temporal luminance transition. Since noise is typically a high-frequency type of degradation, one can conjecture that it is best to hide noise in regions where the luminance is varying rapidly in either the spatial or temporal directions.

The NTSC television system already depends heavily on the properties of the human visual system. For instance, the purpose of the interlaced display format is to achieve a field rate of approximately 60 fields/sec. This field rate is set deliberately much higher than the 30 frames/sec or so which is adequate for motion rendition

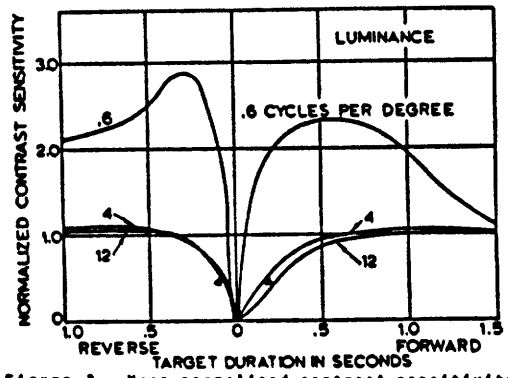

Figure 3. Mean normalized contrast sensitivity measurements for perception of luminance gratings of varying spatial frequency and exposure duration. Gratings appeared immediately before (reverse) or after (forward) a 0.1-sec masking stimulus (N = 7 observers). ▲ indicates reduction in sensitivity by a factor of 2, relative to unmarked conditions. (Courtesy of SMPTE Journal, March 1987)

Figure 2.2: The temporal masking effect. Contrast sensitivity is shown as a function of the time before and after a sudden change in the background luminance.

because the human visual system is very sensitive to temporal frequencies of up to roughly 30 Hz, so that a frame rate of only 30 frames/sec would result in an unpleasant effect known as flicker. At 60 fields per second, the human visual system integrates fields together, whereas at 30 fields per second, a flicker is noticeable, as the image sequence is essentially undersampled in the temporal dimension.

The NTSC system also makes use of the fact that the human visual system is not capable of resolving color details as well as variations in the luminance to achieve some bandwidth compression in the transmission of the chrominance signals. The I component of the NTSC signal is bandlimited to about 1.3 MHz and the Q component is bandlimited to about 0.6 MHz. Few receivers, however, are able to reconstruct more than the lowest 0.5 MHz of either of these components. Thus, each of the reconstructed chrominance components uses only about one-eighth the 4.2 Mhz of bandwidth that the luminance signal uses.

NTSC does not, however, take advantage of all of the characteristics of the

human visual system. For instance, NTSC does nothing to exploit the anisotropic nature of the human visual system. In fact, the NTSC encoding scheme actually allows for more spatial frequency bandwidth in the diagonal directions than in the horizontal or vertical directions, which is exactly opposite the requirements of the human visual system.

More importantly for this research, however, NTSC does nothing to exploit the spatial and temporal masking properties of the human visual system. Schreiber [17, 19] has explained how it is possible to design a television system to take advantage of masking effects to improve performance when the transmission channel is noisy. Also, Glenn [5,6] has shown that these making effects may be used to reduce the amount of visual information that must be transmitted while retaining the same picture quality. The ideas of two-channel processing and adaptive modulation, which are presented in the next chapter, are based entirely on the spatial masking effect.

Chapter 3

Signal Processing Building Blocks

3.1 Adaptive Modulation

Adaptive modulation is a noise reduction technique that takes advantage of under-utilized dynamic range in the original signal. It is particularly useful in applications where the only constraint is peak-to-peak power output (*i.e.* dynamic range) so that an increase in average power output incurs little or no additional cost. If the proper preprocessing is performed on the original signal, adaptive modulation can reduce the effects of a wide variety of interferences, including additive random noise, multipath, and intersymbol interference. As long as compatibility is not an issue, the only drawback to adaptive modulation is that it requires the presence of a second channel (usually digital) to transmit the adaptation factors.

Adaptive modulation achieves the greatest amount of noise reduction when the input signal is approximately stationary with a mean near the center of the allowed dynamic range and a standard deviation which is small compared to the dynamic range. Thus, signals with large low-frequency components generally must be preprocessed in order to achieve maximum performance with an adaptive modulation system, since only the high-frequency portion is suitable for the application of adaptive modulation.

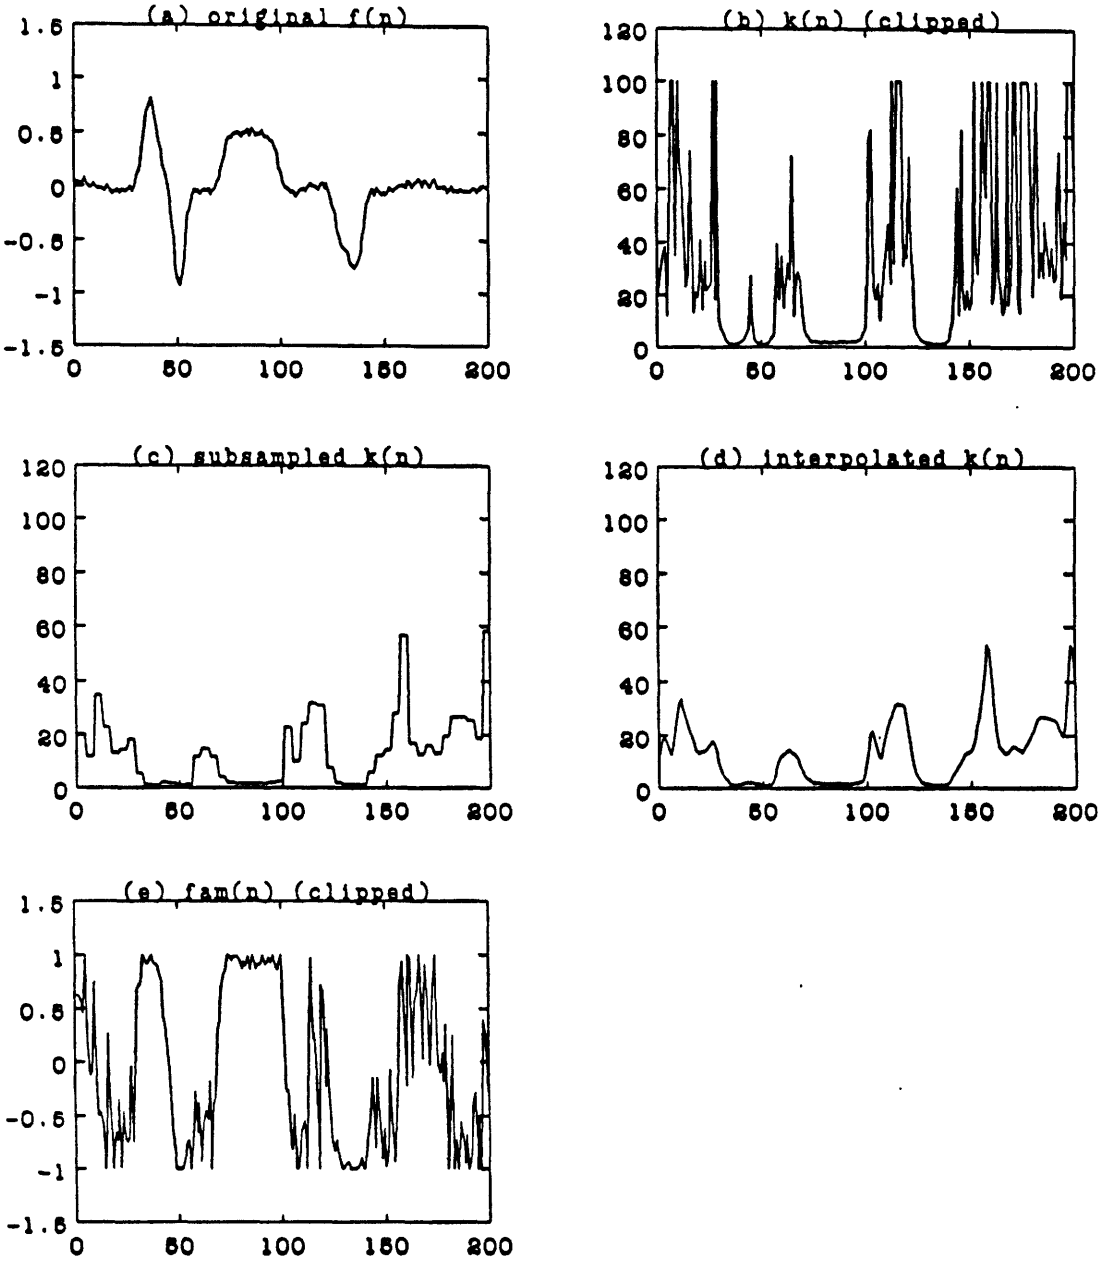

Figure 3.1: Adaptive modulation of a one-dimensional signal. (a) The original zero-mean signal $f_{\mathrm{H}}(n)$. (b) The maximum possible adaptation factors $k(n)$ for each point. (c) The subsampled adaptation factors $k_{\mathrm{L}}^{\mathrm{sub}}(m)$. (d) The interpolated adaptation factors $k_{\mathrm{L}}^{\mathrm{int}}(n)$. (e) The adaptively modulated signal $f_{\mathrm{H}}^{\mathrm{am}}(n)$.

For the purpose of illustration, assume that we have a signal $f_{\mathrm{H}}(n)$ which is approximately stationary and whose mean is approximately zero, as in Figure 3.1(a). Assume also that the allowed dynamic range of the transmitted signal is limited to the interval $[-A, A]$. The first step at the transmitter is to compute for each point n the maximum possible adaptation factor, *i.e.* the largest number $k(n) \geq 1$ such that $k(n) f_{\mathrm{H}}(n) \in [-A, A]$. These values $k(n)$ are shown in Figure 3.1(b). In order to keep the amount of transmitted side information down to a reasonable level, the adaptation factors $k(n)$ are subsampled by a factor of K. This subampling usually does not impede performance since the adaptation factors tend to be low-pass in nature. The subampling is achieved by dividing the signal into blocks and then choosing an adaptation factor $k_{\mathrm{L}}^{\mathrm{sub}}(m)$ for each block which is the minimum of all the $k(n)$ within the block. The result of the subsampling is a piecewise constant function, such as the one in Figure 3.1(c). To prevent blocking effects, these subsampled adaptation factors $k_{\mathrm{L}}^{\mathrm{sub}}(m)$are interpolated in a continuous manner (typically by linear interpolation) to get the adaptation factors $k_{\mathrm{L}}^{\mathrm{int}}(n)$ which are actually used, as shown in Figure 3.1(d). Finally, the original input signal is multiplied by the interpolated adaptation factors $k_{\mathrm{L}}^{\mathrm{int}}(n)$ to produce the adaptively modulated signal which is transmitted over the normal channel. The modulated signal $f_{\mathrm{H}}^{\mathrm{am}}(n)$ is shown in Figure 3.1(e). The subsampled adaptation factors $k_{\mathrm{L}}^{\mathrm{sub}}(m)$ are transmitted on a separate channel (usually digital). Before the adaptive modulation process is applied, $f_{\mathrm{H}}(n)$ may optionally be compressed or expanded in a nonlinear manner, but the improvement in performance due to any such companding will be slight compared to the gain in SNR achieved through adaptive modulation.

It should be pointed out that in practice there are several more subtle points to consider when choosing the subsampled adaptation factors. At the boundary between a block with a small adaptation factor and a block with a large adaptation factor, multiplication of the signal by the interpolated adaptation factor will often cause the resulting signal to stray outside of the allotted dynamic range. This

problem can be treated either by clipping the modulated signal to the limits of the dynamic range or by decreasing the adaptation factors of the blocks neighboring blocks. If clipping is used, there is no way to exactly reconstruct the original signal at the receiver, even if no noise is present in the channel. However, as long as the number of clipped data points is small, the effects of the clipping will probably not be noticed, especially in an application such as television. Reducing the adaptation factors will, of course, make perfect reconstruction possible, but only at the cost of reduced performance in the presence of noise.

It is assumed that the adaptation factors are transmitted without error, so the only distortions are due to interference in the original analog channel. Thus, at the receiver, the adaptive modulation process may be inverted to recover the original signal. The received adaptation factor samples $k_{\mathrm{L}}^{\mathrm{sub}}(m)$ are interpolated in the same manner as they were at the transmitter to produce $k_{\mathrm{L}}^{\mathrm{int}}(n)$. The received analog signal is $f_{\mathrm{H}}^{\mathrm{am}}(n) + w(n)$, where $w(n)$accounts for all channel degradations. This received signal is divided by the interpolated adaptation factors to yield the noise-reduced recovered signal $\hat{f}_{\mathrm{H}}(n)$. Note that the amount of additive noise is reduced by a factor of $k_{\mathrm{L}}^{\mathrm{int}}(n)$ at each point, since

$$\begin{aligned} \hat{f}_{\mathrm{H}}(n) &= \frac{1}{k_{\mathrm{L}}^{\mathrm{int}}(n)} \left(f_{\mathrm{H}}^{\mathrm{am}}(n) + w(n)\right) \\ &= \frac{1}{k_{\mathrm{L}}^{\mathrm{int}}(n)} \left(k_{\mathrm{L}}^{\mathrm{int}}(n) f(n) + w(n)\right) \\ &= f(n) + \frac{w(n)}{k_{\mathrm{L}}^{\mathrm{int}}(n)}. \end{aligned} \tag{3.1}$$

A one-dimensional example demonstrating how the adaptive modulation process reduces the effects of random additive noise and multipath is given in Figures 3.2 and 3.3. In each figure, the original signal $f(n)$ is shown in (a) and the degraded signal $f(n) + w(n)$ is shown in (b). The degraded adaptively modulated signal $f_{\mathrm{H}}^{\mathrm{am}}(n) + w(n)$ is shown in (c), while the signal $\hat{f}_{\mathrm{H}}(n)$ recovered from it is shown in (d).

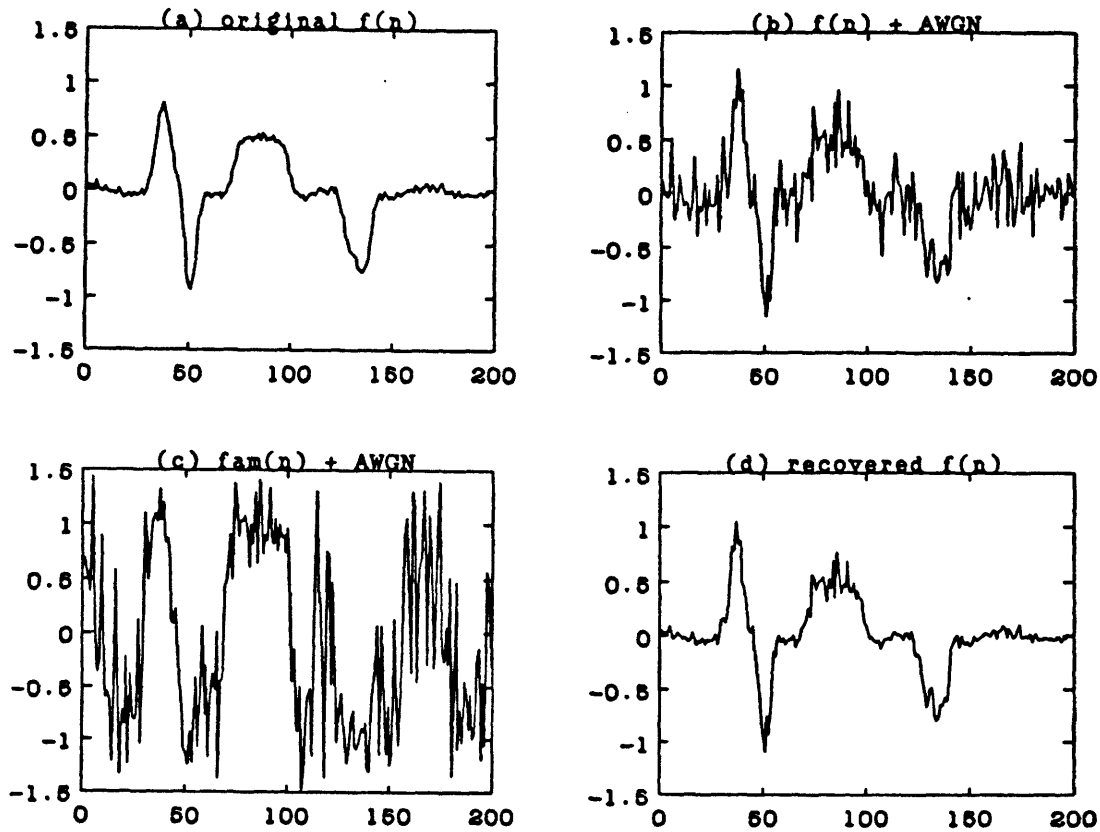

Figure 3.2: Performance of adaptive modulation in the presence of additive white Gaussian noise. (a) Original signal $f(n)$. (b) Original signal $f(n)$ degraded by AWGN $w(n)$. (c) Adaptively modulated signal $f_{\mathrm{H}}^{\mathrm{am}}(n)$ degraded by AWGN $w(n)$. (d) Recovered signal $\hat{f}_{\mathrm{H}}(n)$ obtained by performing inverse adaptive modulation on the signal in (c).

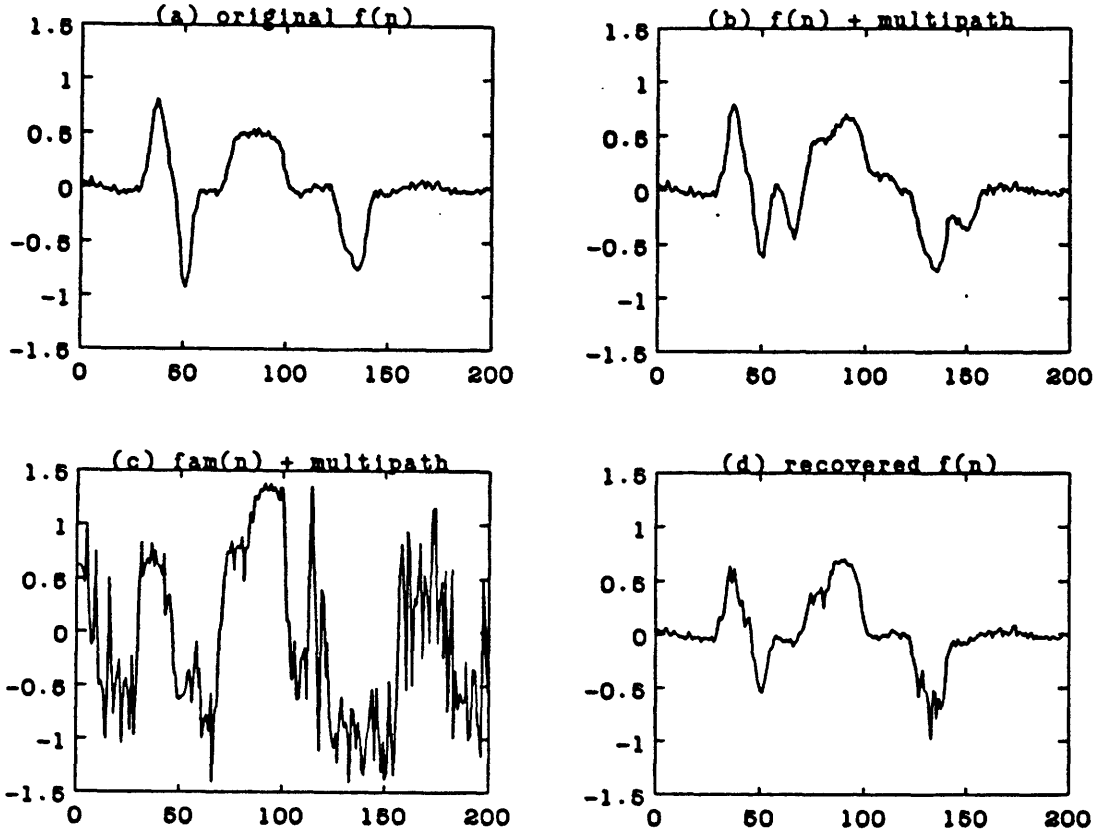

Figure 3.3: Performance of adaptive modulation in the presence of multipath. (a) Original signal $f(n)$. (b) Original signal $f(n)$ degraded by multipath. (c) Adaptively modulated signal $f_H^{am}(n)$degraded by multipath. (d) Recovered signal $\hat{f}_H(n)$ obtained by performing inverse adaptive modulation on the signal in (c).

3.2 Two-Channel Processing

Two-channel processing is a method which allows adaptive modulation to be applied to signals with arbitary statistics, not just those which are roughly zero-mean and stationary. The major idea behind two-channel processing is to apply adaptive modulation only to the high-frequency portion of the signal (which is the portion that is most suitable for adaptive modulation) prior to transmission. This encoding process makes it possible to remove a large percentage of the high-frequency noise at the receiving end by performing the inverse processing.

Two-channel processing can be combined with adaptive modulation in a fairly straightforward manner. In its simplest form, two-channel processing is not receiver-compatible. However, it will be shown later how adjustments may be made to achieve compatibility. The transmitter processing for the simple two-channel system may be divided into three major tasks: transmitter processing:

- The signal is separated into low-frequency and high-frequency portions.
- The low-frequency portion is subsampled and efficiently encoded.
- The high-frequency portion is adaptively modulated.

Then, the low-frequency information, along with some additional side information, is transmitted digitally, and the high-frequency portion of the signal is transmitted over the standard analog channel that would normally be used to transmit the entire signal. At the receiver, the inverse operations are carried out:

- The received (high-frequency) analog signal is adaptively demodulated.
- The low-frequency information is decoded and the low-frequency portion of the signal is recovered by interpolation.
- The demodulated high-frequency and interpolated low-frequency portions are combined to form the recovered signal.

Consider, for example, a one-dimensional signal $f(n)$, as shown in Figure 3.4(a). The first step in two-channel processing is to process $f(n)$ with a low-pass filter (often a simple boxcar filter) to get $f_{\mathrm{L}}(n)$. The low-frequency signal $f_{\mathrm{L}}(n)$ is then subsampled according to

$$f_{\mathrm{L}}^{\mathrm{sub}}(m) = f_{\mathrm{L}}(n)|_{n=Lm}, \tag{3.2}$$

where L is the sampling interval (it is assumed that the low-pass filtering limits the power of the signal to lie approximately within the frequency interval $\omega \in [-\frac{\pi}{L}, \frac{\pi}{L}]$. To ensure that perfect reconstruction is possible at the receiver in the case of no noise, an interpolated low-frequency signal $f_{\mathrm{L}}^{\mathrm{int}}(n)$ is constructed from the subsampled signal, as shown in Figure 3.4(b). Any reasonable interpolation scheme may be used here, although simple first-order linear interpolation appears to be the best. Unlike ideal bandlimited interpolation or QMF inverse-filtering, linear interpolation does not cause any perceptually undesirable ringing at sharp discontinuities.

The high-frequency portion of the original signal $f_{\mathrm{H}}(n)$ is obtained by subtracting $f_{\mathrm{L}}^{\mathrm{int}}(n)$ from $f(n)$, as shown in Figure 3.4(c). The high-frequency signal then undergoes an adaptive modulation process in the manner described in the previous section. The outputs of the adaptive modulation process are the subsampled adaptation factors $k_{\mathrm{L}}^{\mathrm{sub}}(m)$ and the modulated high-frequency portion of the signal $f_{\mathrm{H}}^{\mathrm{am}}(n)$. The interpolated adaptation factors $k_{\mathrm{L}}^{\mathrm{int}}(n)$ for the example in Figure 3.4 are shown in (d) and the modulated high-frequency portion $f_{\mathrm{H}}^{\mathrm{am}}(n)$of the signal is shown in (e).

At the receiving end, the inverse processing is performed. The low-frequency portion of the signal $f_{\mathrm{L}}^{\mathrm{int}}(n)$ is exactly reconstructed by interpolating the subsampled version $f_{\mathrm{L}}^{\mathrm{sub}}(m)$, which is assumed to be transmitted digitally, without error. The subampled adaptation factors are also transmitted digitally and interpolated at the receiver to produce $k_{\mathrm{L}}^{\mathrm{int}}(n)$exactly. They are then used to produce the recovered high-frequency signal $\hat{f}_{\mathrm{H}}(n)$ through inverse adaptive modulation on the received

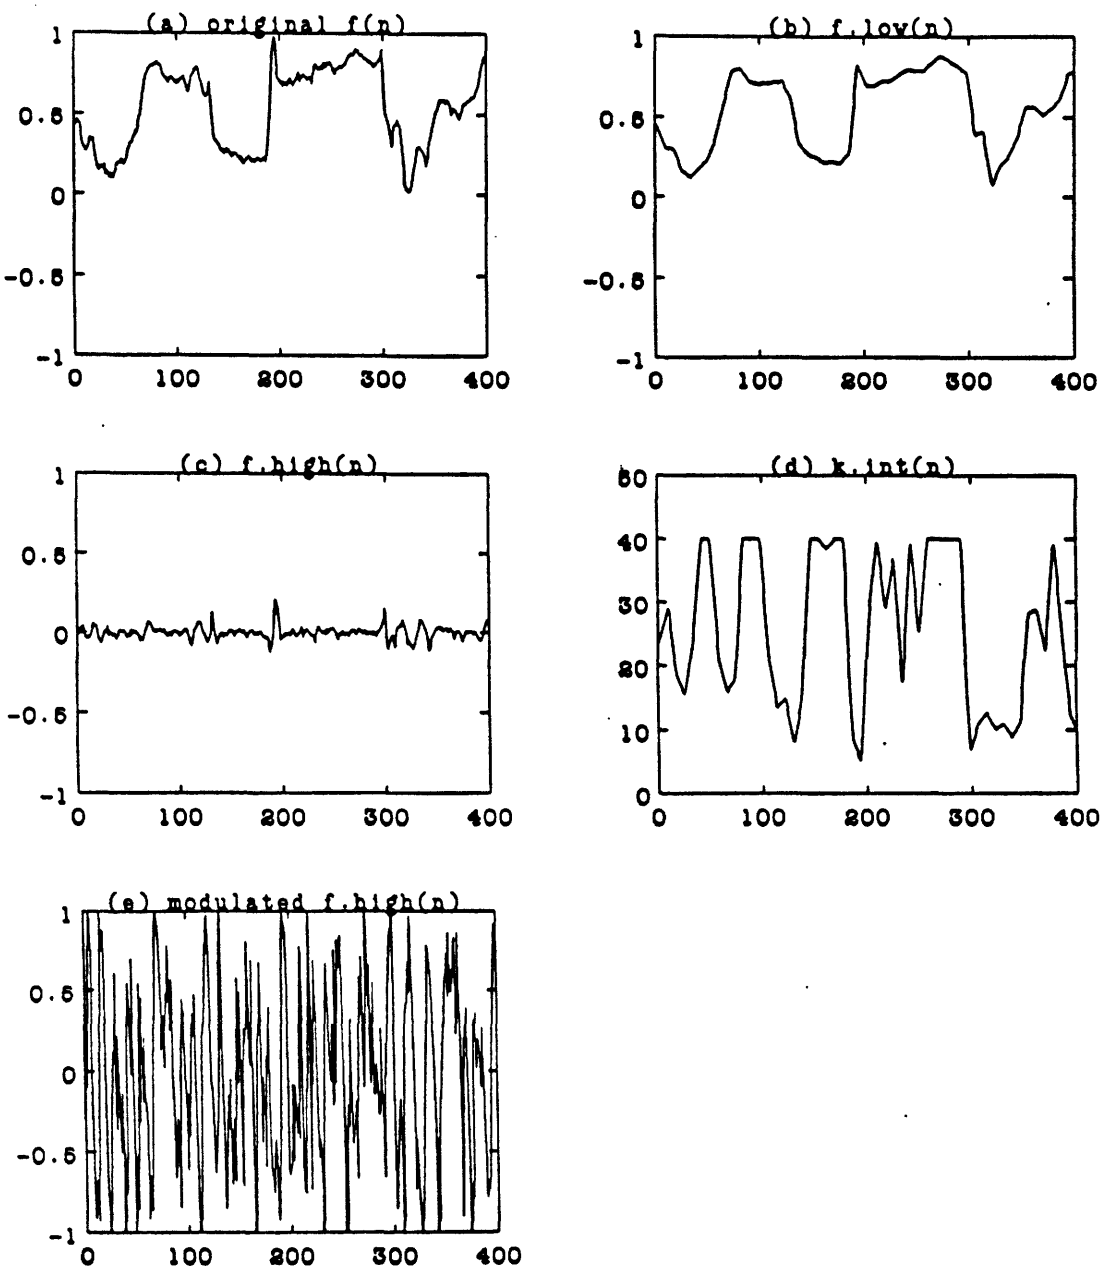

Figure 3.4: Two-channel processing of a one-dimensional signal. (a) The original signal $f(n)$. (b) The interpolated low-frequency portion of the signal $f_{\mathrm{L}}(n)$. (c) The high-frequency portion of the signal $f_{\mathrm{H}}(n)$. (d) The subsampled and interpolated adaptation factors $k_{\mathrm{L}}^{\mathrm{int}}(n)$. (e) The high-frequency portion of the signal after adaptive modulation $f_{\mathrm{H}}^{\mathrm{am}}(n)$.

high-frequency signal $f_{\mathrm{H}}^{\mathrm{am}}(n) + w(n)$, substantially reducing the noise content. The recovered low- and high-frequency component signals are then added together to obtain the recovered signal $\hat{f}(n)$.

Two-channel processing does require that a digital side channel be present to transmit the low-frequency samples and the adaptation factors. The necessary bandwidth of this channel, however, is very small compared to the bandwidth of the original channel. Typical monochrome images have roughly 8 bits of amplitude resolution. Using any simple data compression scheme (*eg.* DPCM with a one-step predictor), the low frequency samples may be encoded with 3–5 bits per sample, while the adaptation factors generally require no more than 2–3 bits of amplitude resolution. Typical block sizes range from 4×4 to 8×8. Using these parameters, one can see that the extra information may be transmitted with a data rate of about 0.078–0.500 bits per pel.

3.3 Receiver-Compatible Two-Channel Processing

Clearly, two-channel processing and adaptive modulation can only improve the signal quality. However, introducing these techniques into a mass-media environment such as television or radio can cause problems with compatibility which may outweigh the benefits. The compatibility issue is particularly troublesome when the baseband of the input signal is perceptually important, since two-channel processing requires that most of the baseband signal be subtracted out and transmitted on a digital side channel. Television is such an environment, since typical broadcast images have very little energy in the high spatial frequencies.

Thus, it would be desirable to have a processing technique which retains the benefits of two-channel processing without significantly altering the original signal, especially in the lower spatial frequency bands, since it is known that those bands are

the most perceptually important. The obvious solution is to add the interpolated low-frequency signal back to the modulated high-frequency prior to transmission to form a "receiver-compatible" signal $f^{RC}(n)$ suitable for transmission over the available analog channel. Since the receiver compatible signal contains the low-frequency portion of the original signal as well as the high-frequency (although in a somewhat altered form), one would expect it to fairly compatible with existing receivers.

The application of this processing to television to go from the original NTSC signal to the receiver-compatible signal is similar to the contrast enhancement algorithm of Lim and Peli [16]. Naturally, one would expect that the receiver-compatible signal would differ from the original NTSC signal in the content of higher spatial frequency bands. In essence, the receiver-compatible two-channel processing deliberately attempts to "hide" these differences in the high spatial frequencies, which are not as perceptually important. This same basic principle was used when color television was introduced. The color components were "hidden" by modulating them on a subcarrier into the frequency region corresponding to fine diagonal detail, which is exactly the spatial frequency band that is most difficult for the human visual system to discern. Thus, while one would expect the receiver-compatible signal to appear somewhat contrasty when it is viewed on a standard receiver, it should be possible to prevent this effect from becoming objectionable.

Several adjustments must be made to the standard two-channel processing algorithm to produce the receiver-compatible signal. In order to facilitate the adaptive modulation of the high-frequency portion of the signal, the low-frequency portion is compressed in dynamic range. If we assume that the allotted dynamic range is the interval $[A_0, B_0]$, then the low-frequency portion of the signal may be compressed to within the interval $[A_L, B_L]$ according to

$$f_L^{com}(n) = A_L + \frac{A_L - B_L}{A_0 - B_0}(f_L^{int}(n) - A_0) \tag{3.3}$$

Of course, if the signal does not occupy all of the allotted dynamic range, then

it is not necessary to completely compress the signal as indicated above. This compression creates margins at the top and bottom of the dynamic range that will be filled when the amplified high-frequency portion is added back in.

The process of choosing the appropriate adaptation factor is slightly more complicated in the receiver-compatible version of two-channel processing. It is now necessary to consider whether the combined signal $f_{\mathrm{L}}^{\mathrm{com}}(n) + f_{\mathrm{H}}^{\mathrm{am}}(n)$ will lie within the overall dynamic range $[A_0, B_0]$, so both the low- and high-frequency portions of the signal must be considered. The new rule for determining the optimal value of the $k(n)$ at each pel is

$$k(n) = \begin{cases} \left| \dfrac{f_{\mathrm{L}}^{\mathrm{com}}(n) - B_0}{f_{\mathrm{H}}(n)} \right|, & f_{\mathrm{H}}(n) > 0 \\ \left| \dfrac{f_{\mathrm{L}}^{\mathrm{com}}(n) - A_0}{f_{\mathrm{H}}(n)} \right|, & f_{\mathrm{H}}(n) < 0 \end{cases} \tag{3.4}$$

One can easily verify that with this choice of $k(n)$, the receiver-compatible signal $f^{\mathrm{RC}}(n) = f_{\mathrm{L}}^{\mathrm{com}}(n) + f_{\mathrm{H}}(n)$ will lie within the allotted dynamic range.

At the receiver, the low-frequency portion of the signal $f_{\mathrm{L}}^{\mathrm{int}}(n)$ is created by interpolating low-frequncy samples $f_{\mathrm{L}}^{\mathrm{sub}}(m)$. It is then compressed to reconstruct $f_{\mathrm{L}}^{\mathrm{com}}(n)$and subtracted from the received analog signal $f^{\mathrm{RC}}(n) + w(n)$ to obtain a noisy version of the adaptively modulated high-frequency signal $f_{\mathrm{H}}^{\mathrm{am}}(n)+w(n)$. The adaptation factors $k_{\mathrm{L}}^{\mathrm{int}}(n)$, which are reconstructed by interpolating $k_{\mathrm{L}}^{\mathrm{sub}}(m)$, are then used to perform inverse adaptive modulation on the received high-frequency signal. This process produces the recovered high-frequency portion of the signal $\hat{f}_{\mathrm{H}}(n)$, which has a dramatically reduced noise content. Finally, the complete recovered signal $\hat{f}(n)$ is assembled by adding adding $\hat{f}_{\mathrm{H}}(n)$ to $f_{\mathrm{L}}^{\mathrm{int}}(n)$.

It should also be pointed out that it is possible to operate a receiver-compatible version of two-channel processing system without transmitting the digital low-frequency samples. These values may simply be estimated at the receiver by low-pass filtering the analog receiver-compatible signal $f^{\mathrm{RC}}(n)$, making use of the fact that the mean of the high-frequency component $f_{\mathrm{H}}^{\mathrm{am}}(n)$of the reciever-compatible

signal should be approximately zero. Such a system would perform well in the presence of additive noise alone, since additive noise is typically a high-frequency degradation. There would be little protection, however, against the effects of echo or interchannel interference, which degrade the low frequencies as well as the high frequencies.

The complete-receiver compatible system is given in Figures 3.5 and 3.6. Figure 3.5 shows the processing that must be done to a signal at the transmitter, while Figure 3.6 illustrates the required receiver processing. Note that while the transmitter processing is fairly involved, the processing at the receiver consists of only a few simple steps. This receiver-end simplicity is especially important when the number of receivers is larger than the number of transmitters by several orders of magnitude, as is the case with AM and FM radio and television. Figure 3.7 shows the transformations that a single sample point of input (*i.e.* a pel) would undergo during the receiver-compatible encoding process and Figure 3.8 demonstrates the effect of applying the receiver-compatible adaptive modulation process to a one-dimensional signal. Note that the dynamic range is used much more completely in the receiver-compatible signal than in the original.

3.4 Two-Channel Processing in Two Dimensions

Because low-pass filtering, subsampling, and interpolation all have natural extensions from one to two dimensions, adaptive modulation and two-channel processing may also be easily extended. Two-dimensional low-pass filtering may be accomplished by forming a separable rectangular filter from a one-dimensional filter. Linear interpolation between two points in one dimension may be replaced by bilinear interpolation between four points in two dimensions.

Suppose, for instance, that the original two-dimensional signal (*i.e.*image) is $f(n_1, n_2)$. The first step in the transmitter processing for the receiver-compatible

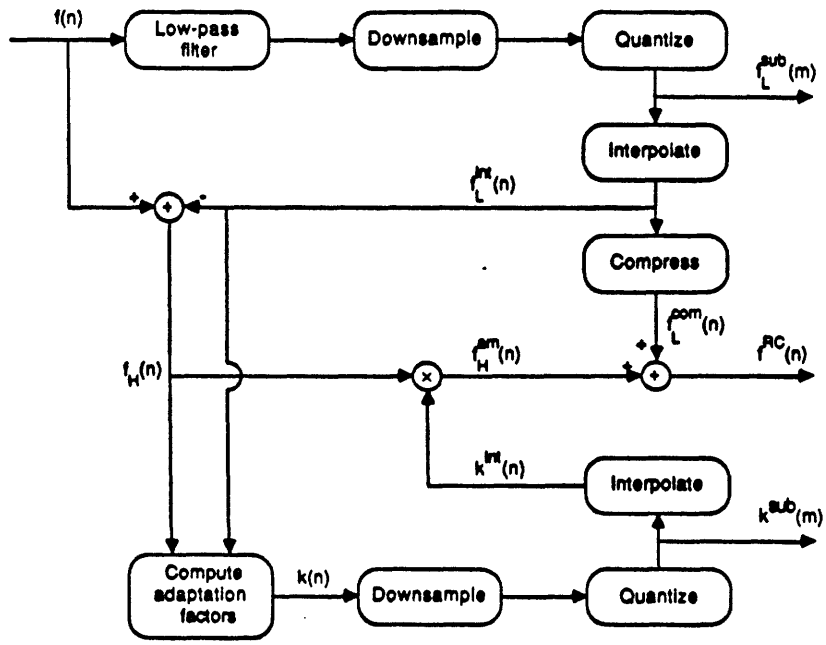

Figure 3.5: Encoder processing for receiver-compatible adaptive modulation.

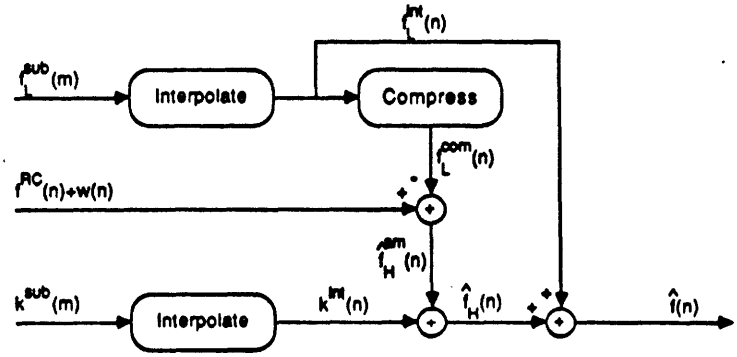

Figure 3.6: Decoder processing for receiver-compatible adaptive modulation.

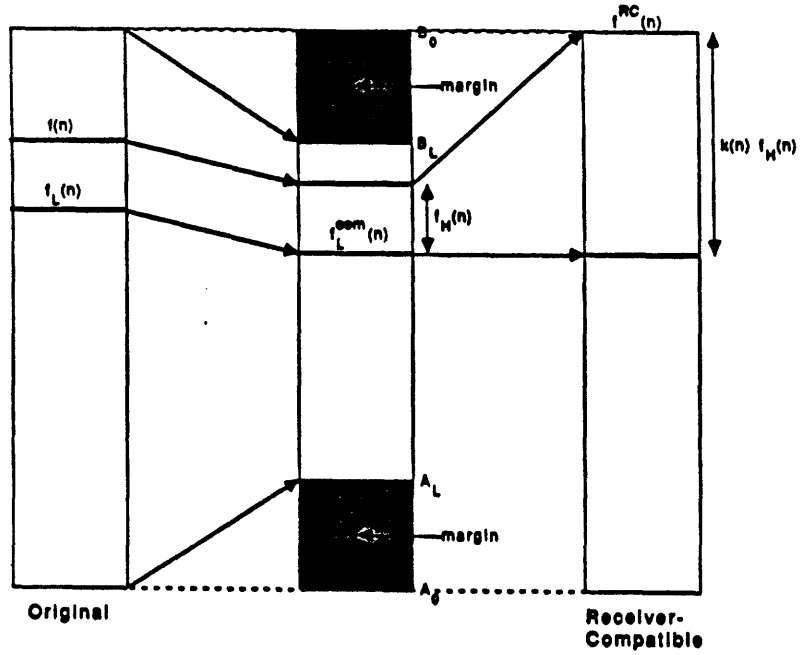

Figure 3.7: Mapping of low-frequency and high-frequency portions of a signal undergoing receiver-compatible adaptive modulation.

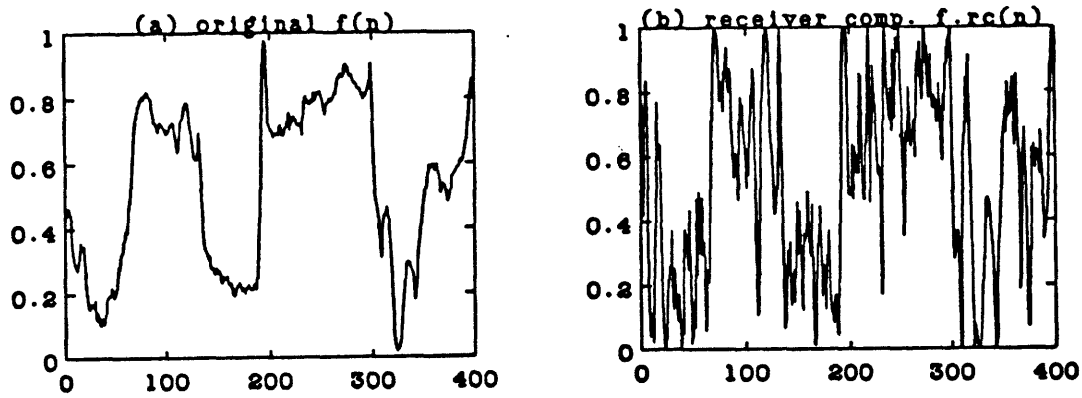

Figure 3.8: Result of performing receiver-compatible adaptive modulation on a one-dimensional signal: (a) Original signal $f(n)$. (b) Receiver-compatible signal $f^{RC}(n)$.

system is to low-pass filter $f(n_1, n_2)$ (usually with a simple boxcar filter) to get the "lows" signal $f_{\mathrm{L}}(n_1, n_2)$. To save on transmission bandwidth, the lows signal is then subsampled to get $f_{\mathrm{L}}^{\mathrm{sub}}(m_1, m_2)$, according to

$$f_{\mathrm{L}}^{\mathrm{sub}}(m_1, m_2) = f_{\mathrm{L}}(n_1, n_2)|_{n_1 = L_1 m_1,\ n_2 = L_2 m_2}, \tag{3.5}$$

where $L_1 \times L_2$ is the two-dimensional block sampling period. A new low-pass version of the picture $f_{\mathrm{L}}^{\mathrm{int}}(n_1, n_2)$ is obtained by interpolating $f_{\mathrm{L}}^{\mathrm{sub}}(m_1, m_2)$ (usually by bilinear interpolation) to get $f_{\mathrm{L}}^{\mathrm{int}}(n_1, n_2)$. The high-frequency portion of the picture $f_{\mathrm{H}}(n_1, n_2)$ is then obtained by subtracting $f_{\mathrm{L}}^{\mathrm{int}}(n_1, n_2)$ from $f(n_1, n_2)$.

The high-frequency portion of the signal then undergoes a two-dimensional adaptive modulation process. The process begins by computing $k(n_1, n_2)$, the maximum allowable adaptation factor for each point. Next, the adaptation factors are subsampled by choosing the minimum adaptation factor from each $K_1 \times K_2$ block to produce $k_{\mathrm{L}}^{\mathrm{sub}}(m_1, m_2)$($K_1$ and K_2 are not necessarily equal to L_1 and L_2). Bilinear interpolation is then applied to these samples to construct $k_{\mathrm{L}}^{\mathrm{int}}(n_1, n_2)$, which is used to modulate $f_{\mathrm{H}}(n_1, n_2)$, producing the adaptively modulated signal $f_{\mathrm{H}}^{\mathrm{am}}(n_1, n_2)$.

To create the receiver-compatible signal, the interpolated low-frequeny portion of the signal $f_{\mathrm{L}}^{\mathrm{int}}(n_1, n_2)$ is linearly compressed in dynamic range to produce $f_{\mathrm{L}}^{\mathrm{com}}(n_1, n_2)$, creating margins at the top and bottom of the dynamic range. The receiver-end processing is completed by adding $f_{\mathrm{H}}^{\mathrm{am}}(n_1, n_2)$ to $f_{\mathrm{L}}^{\mathrm{com}}(n_1, n_2)$to construct the receiver-compatible signal $f^{\mathrm{RC}}(n_1, n_2)$.

The receiver decoding process is, of course, also analogous. The received analog signal is $f^{\mathrm{RC}}(n_1, n_2) + w(n_1, n_2)$, where $w(n_1, n_2)$is an additive noise term accounting for all channel degradations. The receiver first recovers the low-frequency portion of the signal $f_{\mathrm{L}}^{\mathrm{int}}(n_1, n_2)$ by interpolating $f_{\mathrm{L}}^{\mathrm{sub}}(m_1, m_2)$. This signal is compressed and subtracted from the received analog signal to obtain $f_{\mathrm{H}}^{\mathrm{am}}(n_1, n_2) + w(n_1, n_2)$. The adaptation factors $k_{\mathrm{L}}^{\mathrm{int}}(n_1, n_2)$ are reconstructed by interpolating $k_{\mathrm{L}}^{\mathrm{sub}}(m_1, m_2)$and are then used to perform inverse adaptive modulation on the noisy high-frequency portion of the signal to produce $\hat{f}_{\mathrm{H}}(n_1, n_2)$, which has a substantially reduced noise

content. Finally, the complete recovered signal $\hat{f}(n_1, n_2)$ is obtained by adding $\hat{f}_{\mathrm{H}}(n_1, n_2)$ to $f_{\mathrm{L}}^{\mathrm{int}}(n_1, n_2)$.

As one might note, when this processing is applied to images, the adaptation factors will not be very large near large-scale luminance transitions, such as edges or detailed regions. In these regions, most of the dynamic range is already consumed by the original signal. Fortunately, however, the characteristics of the human visual system mentioned earlier are such that random noise is least visible in regions with edges or busy detail. Thus, one would expect that the subjective performance of the receiver-compatible adaptive modulation system should be almost identical to that of the two-channel system. Determining the degree to which this assumption is true is the one of the primary subjects of investigation of the next chapter.

Chapter 4

System Design

With the building blocks of adaptive modulation and two-channel processing as a starting point, a study was conducted to determnine the feasibility of a receiver-compatible noise reduction system for television. This investigation was divided in two phases. First, an attempt was made to design the "optimal" system. After the system had been designed, its performance was analyzed. The key performance issues that were studied were the systems's compatiblity with existing NTSC receivers and its ability to produce noise-reduced pictures on advanced receivers.

There were many steps leading up to the design of the receiver-compatible noise reduction system. To begin with, a scheme for embedding a digital side channel within the NTSC signal format had to be devised and an estimate made of its effective digital bandwidth. This bandwidth figure was needed as a guide for determining the feasibility of various design choices. Of course, devising such a digital modulation scheme required some understanding of the relevant features of the existing NTSC television system; thus, a brief introduction to NTSC is given in the next section.

The processing of the luminance was the next item to be considered. There were many decisions to be made regarding the filtering and subsampling of the luminance, the selection of the adaptation factors, and the level of quantization that could be

tolerated. A wide variety of tests were conducted to determine how the various system parameters related to each other and to the overall system performance. After choices for how to process the luminance had been narrowed, the processing of the chrominance signals was considered. Following that, the issue of how to encode the digital side information, particularly the luminance and chrominance samples, was addressed. Finally, attempts were made to correct several of the observed deficiencies of the receiver-compatible noise reduction system. Particular attention was paid to raising the level of performance of the system in neighborhoods of sharp luminance transitions and on full video rate sequences.

In the discussion that follows, there will be repeated references to "advanced" and "standard" receivers. A standard receiver is defined as one that is capable of decoding only signals encoded in the existing NTSC format. An advanced receiver, on the other hand, is capable of using the side information supplied in the encoded receiver-compatible format to perform noise reduction.

4.1 The NTSC Signal

In the NTSC system, moving imagery is temporally sampled into frames at the rate of approximately 30 frames per second. However, since the human visual system is quite sensitive to temporal frequencies of 30 Hz, each frame is scanned in a 2:1 interlace manner. The interlaced scanning divides each frame into two fields and effectively increases the temporal sampling rate to approximately 60 Hz (the actual rate is 59.94 fields per second). Each frame is vertically sampled into 525 scan lines. Of these 525 scan lines, 480 are "active", meaning that they are visible on the television screen. Thus, there are 262.5 scan lines per field, 240 of which are active. Using these numbers, one can quickly calculate that the horizontal line rate is about 15.75 KHz (horizontal line time of roughly 63.5 μsec).

The remaining scan lines (22.5 per field, 45 per frame) make up the vertical

blanking interval (VBI). The VBI was originally inserted to give the electron gun in the receiver's picture tube time to retrace to the top of the screen, but now technology has advanced to the point where the transmission format is independent of the requirements of the display apparatus. Now it is possible that these lines be used to send extra information. Currently the Teletext system makes use of the VBI to transmit alpha-numeric information (used for closed-captioning for the hearing impaired). Also, at least one proposed HDTV system [10] makes use of the VBI to transmit a wide variety of information.

Given that the usable luminance bandwidth is approximately 4 MHz and that the horizontal blanking interval (HBI) occupies roughly 16% of each scan line, the number of distinct samples per line is

$$2 \times 4\,\text{MHz} \times 63.5\,\mu\text{sec} \times 0.84 \approx 426\,\text{samples/line}.$$

It will be assumed later on that digital samples may be sent at this rate using the NTSC format.

In the NTSC signal format, color information is transmitted by a clever form of frequency-division multiplexing. The luminance (Y) and chrominance (I and Q) components are computed from the RGB color coordinates by the linear transformation

$$\begin{bmatrix} Y \\ I \\ Q \end{bmatrix} = \begin{bmatrix} 0.299 & 0.587 & 0.114 \\ 0.596 & -0.274 & -0.322 \\ 0.211 & -0.523 & 0.312 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}. \tag{4.1}$$

The Y component, which is bandlimited to 4.2 Mhz, is transmitted in the baseband. The I and Q components are bandlimited to 1.3 and 0.6 MHz respectively and modulated in quadrature on a subcarrier $f_{SC} = 3.579545$ MHz. This subcarrier frequency f_{SC} is carefully chosen to be an odd multiple of half the horizontal line rate. Because the spectra of the luma and chroma tend to be line spectra with the harmonics spaced at intervals of 30 Hz (*i.e.* the frame rate), this choice of a subcarrier causes the harmonics of the luma and chroma to be interleaved, thus

minimizing the interference between them.

4.2 Digital Data Transmission

Before the particulars of the receiver-compatible noise reduction system are investigated, it is important to get a fairly good estimate of the digital bandwidth available for transmission of the discrete parameters. Since the digital bandwidth is a strict limitation on the system, it can be used to eliminate impractical design choices from consideration.

As in the HDTV system proposed by Zenith [10], the receiver-compatible noise reduction system will use VBI to transmit the digital side information. Zenith proposes that 2 lines of each blanking interval be used for synchronization, leaving an average of 20.5 lines per field available for the digital data. While this system will not require any special synchronization procedure, the 2 line margin is probably a good safety measure to prevent the modulated digital data from appearing on the television screen. According to Zenith, digital data may be transmitted reliably in the VBI using 16-state quadrature amplitude modulation (QAM) with as little as 24 dB SNR. This claim is consistent with the results observed by Seo and Feher [23], who measured bit error probabilities of 10^{-6}–10^{-7} for data transmitted with 16-state QAM in a multi-channel environment with a SNR of 24 dB. Since it is possible to transmit 4 bits per sample with 16-state QAM, the total number of bits that can be transmitted in the VBI of a single field is:

$$(426\,\text{samples/line}) \times (4\,\text{bits/sample}) \times (20.5\,\text{lines/field}) = 27,880\,\text{bits/field}.$$

Clearly, some of these bits must be allocated for error protection. The protection scheme that Zenith recommends is a concatenated interleaved code capable of correcting at least one error in every 10 bits, as well as burst errors. The code carries an overhead of 33% (*i.e.* 1 bit out of 4 is a check bit), so that the number of useful

bits available for transmission of digital data is:

$$0.75 \times (27,880 \text{ bits/field}) = 20,910 \text{ bits/field} = 41,820 \text{ bits/frame.}$$

This is an effective digital data rate of roughly 1.2 Mbps. While all of this available digital bandwidth could be used to encode video information, an allotment of 200–300 Kbps for a digital audio channel and for the transmission of auxiliary information such as teletext makes much more sense. This allocation leaves roughly 900 Kbps to 1 Mbps (slightly more than 30 Kbps per frame) for transmission of video parameters.

There are two major issues to consider when deciding how these bits will be allocated: reduction of channel degradations on advanced receivers and picture quality impairment on standard NTSC receivers. With these issues in mind, there are several ways to trade bits off against performance. Smaller block sizes for both the low-frequency samples and the adaptation factors will yield better subjective results, but will also results in an increased bit rate. Similarly, one may quantize the sample values and adaptation factors coarsely in order to reduce the bit rate, but only at the cost of reducing the performance by adding quantization noise. Predictive or differential coding schemes may be used, but they too may introduce unwanted quantization effects. Where, then, can digital data compression be done with the least adverse effects?

4.3 Luminance Processing

Because the luminance has roughly 8 times the bandwidth of the chrominance, and because the human visual system is much more sensitive to changes in the luminance than to changes in the chroma, it makes sense that providing a good luminance signal on both standard and advanced receivers should be of primary importance. One would also expect that that because of its high bandwidth, the encoding of the luminance signal will require more digital bandwidth than the en-

coding of the chrominance signals. Thus, it is logical to consider the processing of the luminance signal (Y component) first. Once all of the important decisions regarding the processing of the luminance have been made, the processing of the chrominance signals can be considered.

4.3.1 Block Size

There are two types of blocks which must be considered in designing an adaptive modulation system. The first type, the *L-block*, specifies how the low-frequency portion of the signal will be subsampled. Because the vertical and horizontal resolutions of human visual system are not significantly different, the L-blocks are chosen to be square. One sample is taken from each L-block, and is assumed to correspond to the output of the low-pass filter at the center of the L-block. Clearly, if the effect of aliasing are to be avoided, the parameters of the low-pass filter must be directly dependent on the size of the L-block.

The second type of block, the *K-block*, specifies how the adaptation factors will be subsampled. Like the L-blocks, the K-blocks are chosen to be square and are assumed to be centered around samples of the adaptation factors. There is no explicit dependence of the size of the K-blocks on the the size of the L-blocks; however, one would expect that performance considerations would prevent the two sizes from drifting too far apart.

Clearly, both block sizes have a great influence on both bit rate and performance. As the L-block size gets smaller, the interpolated low-frequency portion of the signal will match the original signal more closely. A more accurate interpolant reduces the power (peak power, in particular) in the high-frequency portion of the signal so that adaptive modulation is more effective. The improvement gained by using smaller block sizes is especially apparent in the neighborhood of sharp edges. As the K-block size gets smaller, the spatial adaptivity of the adaptation factor selection algorithm increases. The regions which are dominated by a few large values in the

high-frequency portion of the signal decrease in size, so that on average, the values of the subsampled adaptation factors increase, allowing for more noise reduction.

Several tests were done on 512×512 pixel stationary images to obtain a preliminary estimate of the range of effective block sizes. In all of these tests, the standard two-channel coding scheme was used with no parameter quantization (*i.e.* parameters were transmitted with full floating point precision). The original pictures were encoded, subjected to a variety of degradations, and then decoded. Both subjective and MSE measurements were taken to determine the relative perceptual effects of varying the block sizes. For both the advanced and the standard receivers, the regions of greatest interest were the neighborhoods surrounding sharp edges. Because the low-frequency portion of the signal does not match the complete signal very well in these regions, there will be a large amount of energy left in the high-frequency portion of the signal. Thus, near sharp edges the reduction of channel degradations on advanced receivers will be limited because the adaptation factors will not be very large and the increase in contrast on standard NTSC receivers will be particularly noticeable. In the blank background regions and uniformly detailed regions the effects of changing the block sizes are not quite as noticeable.

In the case of linear one-dimensional interpolation, if the low-frequency samples are chosen optimally, the maximum absolute error between the interpolated low-frequency signal and the original should be $O(N^{-2})$ [2], where N is the number of sample points. In fact, for uniform meshes, the maximum difference between the interpolant and the original satisfies

$$\max_{x \in [a,b]} |f(x) - g(x)| \leq \frac{1}{8} \left(\frac{b-a}{N-1} \right)^2 \max_{x \in [a,b]} |f''(x)| .$$

The error bounds for the two-dimensional case are similar. Assuming that the original two-dimensional function is sampled on a $N \times N$ grid, the maximum absolute error is $O(N^{-2})$. Of course, the performance of the system is not solely dependent on the fit of the interpolant, but these bounds should provide some insight.

For the experiment, block sizes of 2×2, 4×4, 8×8, 16×16, and 32×32 were

used for both the adaptation factors (K-blocks) and the low-frequency samples (L-blocks). All possible combinations of K-block and L-block sizes were tested. The quantitative results are shown in Figures 4.1 and 4.2.

Clearly, it is desirable to choose the system parameters so as to minimize the MSE of both the encoded and the recovered pictures, thereby producing the best possible pictures on both the standard and advanced receivers. Therefore, any figure which is supposed to quantify the overall performance of the system should depend on both the encoded and recovered MSE figures. In accordance with this observation, an *ad hoc* figure of merit was computed for each block size combination by multiplying the MSE of the recovered picture with that of the encoded picture (*i.e.* adding the respective dB MSE figures). Thus, a change in the system parameters which improves the SNR of the decoded picture by making the SNR of the encoded picture proportionately worse (or vice-versa) will cause no change in the figure of merit. A plot of the figure of merit for all of the various L-block and K-block size combinations can be found in Figure 4.3. In all cases, a lower figure-of-merit number indicates better combined overall performance

Several interesting observations may be made. As expected, the MSE of the encoded pictures increases as the L-block size increases or the K-block size decreases. In a corresponding manner, the MSE of the recovered picture increases as the L-block size increases or the K-block size increases. Furthermore, if different block sizes are to be used for the adaptation factors and the low-frequency samples, it is always better to use the smaller blocks for the adaptation factors. In other words, it is better to use 4×4 K-blocks with 8×8 L-blocks than 8×8 K-blocks with 4×4 L-blocks. These quantitative results corresponded exactly to the perceived image quality of both the encoded receiver-compatible picture and the recovered picture.

From a subjective standpoint, excellent results could be obtained even for difficult pictures (*i.e.* picture with many sharp edges and fine details) with 4×4 blocks for both the low-frequency samples and the adaptation factors. Increasing both the

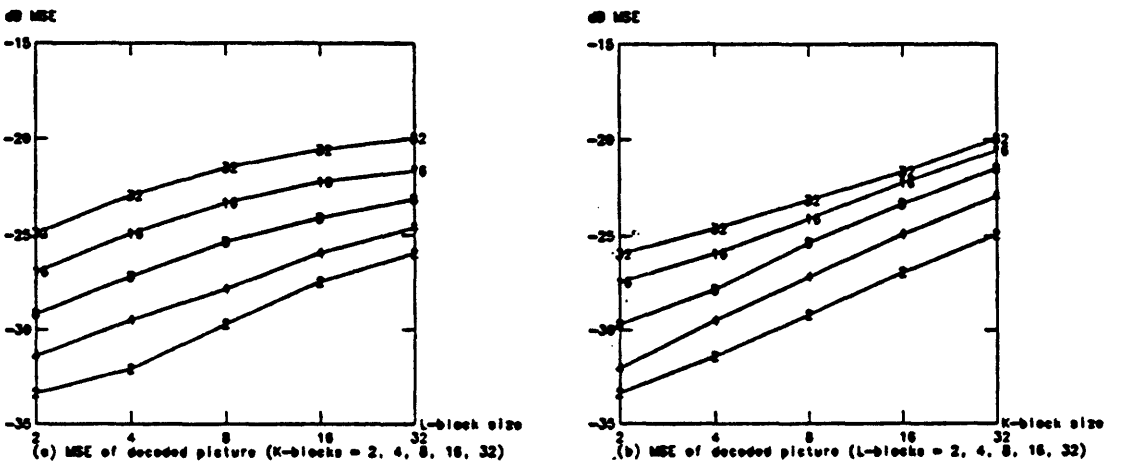

Figure 4.1: Graph of MSE between original picture and picture recovered on advanced receiver for varying block sizes: (a) Each curve is a function of L-block size and corresponds to a fixed K-block size (of 2, 4, 8, 16, or 32). (b) Each curve is a function of K-block size and corresponds to a fixed L-block size (of 2, 4, 8, 16, or 32).

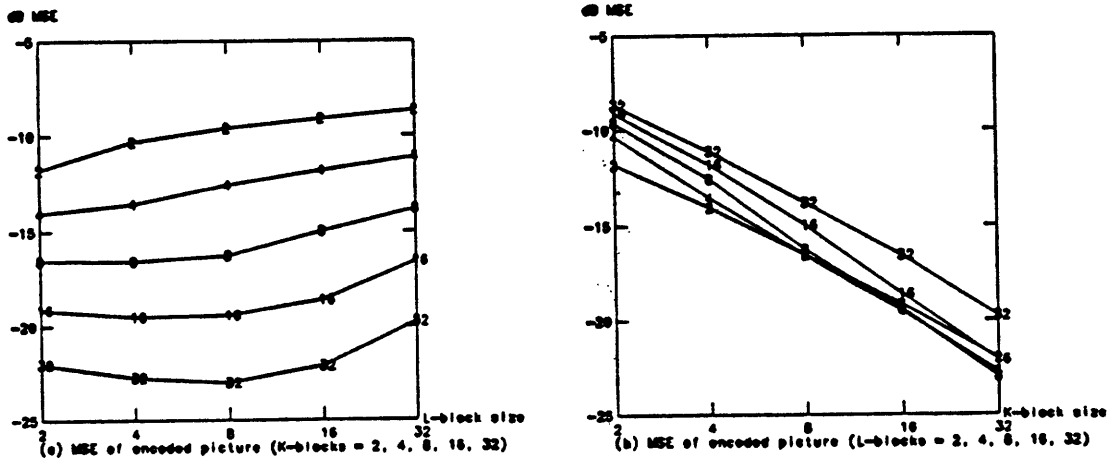

Figure 4.2: Graph of MSE between original picture and encoded picture as viewed on a standard receiver for varying block sizes: (a) Each curve is a function of L-block size and corresponds to a fixed K-block size (of 2, 4, 8, 16, or 32). (b) Each curve is a function of K-block size and corresponds to a fixed L-block size (of 2, 4, 8, 16, or 32).

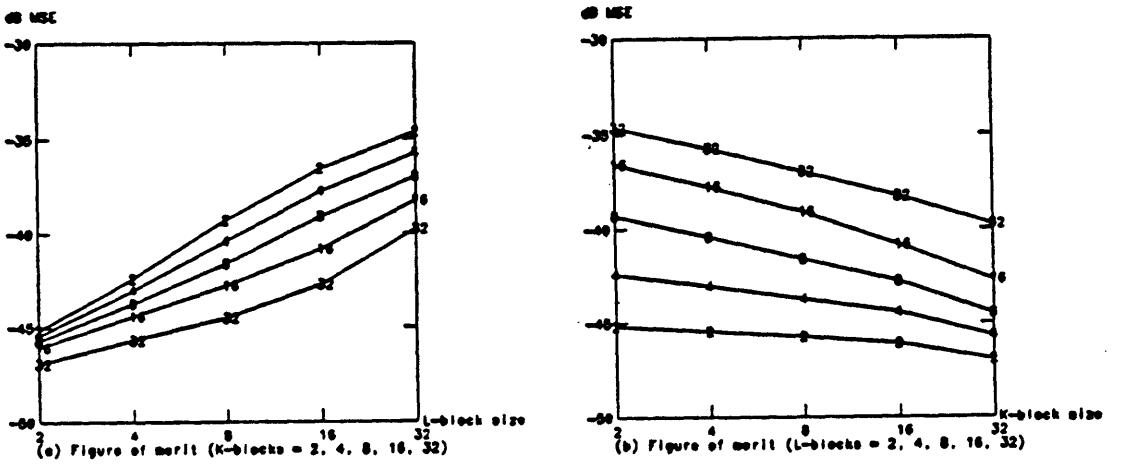

Figure 4.3: Graph of figure of merit for varying block sizes: (a) Each curve is a function of L-block size and corresponds to a fixed K-block size (of 2, 4, 8, 16, or 32). (b) Each curve is a function of K-block size and corresponds to a fixed L-block size (of 2, 4, 8, 16, or 32).

L-block size and K-block size to 8×8 yielded acceptable results, although performance was noticeably worse than in the 4×4 block case. When the block sizes were both increased to 16×16, significant objectionable artifacts became apparent. Figure 4.4 shows how varying the L-block size from 4×4 to 16×16 with a fixed K-block size of 4×4 affects the quality of both the encoded and the recovered pictures. Figure 4.5 demonstrates how picture quality is affected as the K-block size is varied from 4×4 to 16×16 for a fixed L-block size of 4×4. Based on these perceptual results and in consideration of the constraint imposed by the digital bandwidth, it was determined that the best choices for block sizes were 8×8 for the L-blocks and 4×4 for the K-blocks.

In converting these results to television, several factors had to be considered. Although, the dimensions of the NTSC picture (480 rows $\times$ 426 pels) are very close to the dimensions of the images used in the testing (512×512 pixels), the aspect ratios were significantly different. An NTSC picture has an aspect ratio of 4:3, while the stationary images were all 1:1. Thus, a square block on a television screen is composed of roughly $3N$ rows of $2N$ pels each for some number N, whereas a square block in one of the stationary images had as many rows as columns. An additional consideration is the fact that rows are not the correct unit of vertical resolution for a television system, particularly for a system which uses interlaced scanning, such as NTSC. The number of effective vertical samples is obtained by multiplying the number rows by the Kell factor, which is a number between 0 and 1 (roughly 0.70 for NTSC).

One would expect, then, that results for blocks of size 4×4 and 8×8 on stationary images would apply approximately for blocks of size 6×4 and 12×8 respectively on television pictures. Using these figures, one can compute the total number of blocks per frame that would be required for a 6×4 block:

$$\frac{480 \text{ lines}}{6} \times \frac{424 \text{ pixels}}{4} = 8,480 \text{ blocks/frame}$$

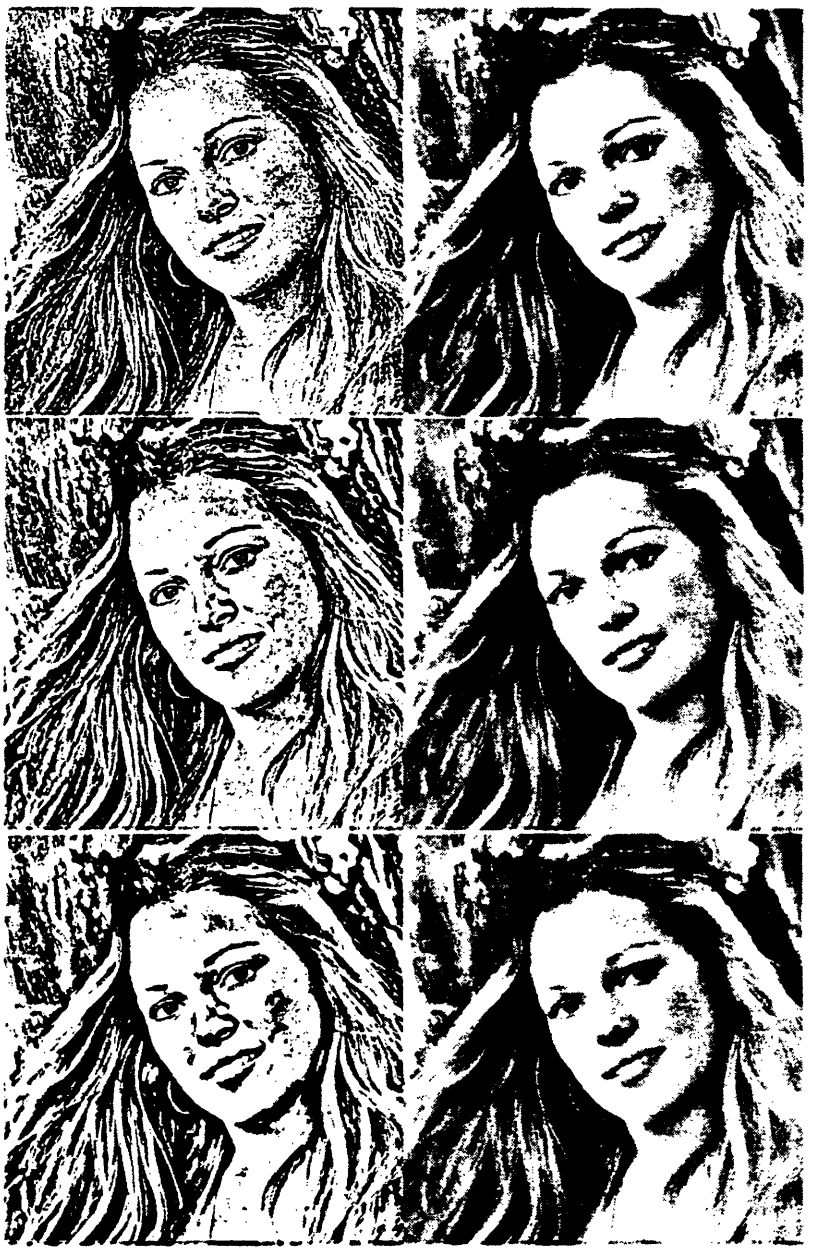

Figure 4.4: Effects of varying the L-block size with a fixed K-block size of 4×4: 4×4 L-blocks (a) encoded, (b) recovered. 8×8 L-blocks: (c) encoded, (d) recovered. 16×16 L-blocks (e) encoded, (f) recovered.

(a) (b)

(c) (d)

(e) (f)

Figure 4.5: Effects of varying the K-block size with a fixed L-block size of 4×4: 4×4 K-blocks (a) encoded, (b) recovered. 8×8 K-blocks: (c) encoded, (d) recovered. 16×16 K-blocks (e) encoded, (f) recovered.

and for a 12×8 block:

$$\frac{480 \text{ lines}}{12} \times \frac{424 \text{ pixels}}{8} = 2,120 \text{ blocks/frame}$$

Assuming that 1.2 Mbps of digital bandwidth are available, the system can use up to 4.93 bits per 6×4 block or up to 19.73 bits per 12×8 block. Keeping in mind, however, that since it would be desirable to have a digital audio channel, 950 Kbps is a more realistic figure for the digital bandwidth available for video information. Using this number, there are 3.95 bits per 6×4 block or 15.80 bits per 12×8 block available. These bits-per-block figures will become important later when the issue of parameter quantization is addressed.

4.3.2 Adaptation Factors

The choice of an appropriate allowed dynamic range for the adaptation factors was another issue requiring some investigation. As one might expect, there is a trade-off involved in selecting the maximum allowed adaptation factor K_{max}. Clearly, one would like to make the adaptation factors as large as possible to achieve the greatest possible noise reduction on the advanced receivers. There is, however, a price to pay for allowing the adaptation factors to be arbitrarily large. Large adaptation factors will yield regions of extremely high contrast on standard NTSC receivers, thus impairing the receiver compatibility of the system.

To determine the optimal trade-off point for the allowed dynamic range of the adaptation factors, several still images were processed with block sizes (L-blocks and K-blocks) of 4×4 and 8×8 and K_{max} ranging from 2 to 64. As in the block size experiment, the original pictures were encoded, subjected to a variety of channel degradations (producing the picture as it would be seen on a standard NTSC receiver), and then decoded (producing the recovered picture as it would be seen on an advanced receiver). The encoded pictures were viewed both with and without added channel degradations, to determine how impairments due to large adaptation factors compare perceptually with those due to the channel degradations.

The MSE between the processed picture and the original was computed for all encoded and recovered pictures in order to get a quantitative measure for the effects of varying K_{max}. These MSE figures are shown in Figures 4.6 and 4.7. From these figures, it can be seen that, in general, very little overall improvement is achieved by increasing K_{max} beyond 16. In fact, all but about 0.5 dB of the maximum possible noise reduction is retained by setting K_{max} to 8. More importantly, however, it should be noted how dependent both the compatibility and noise-reduction performance of the system are on the choice of K_{max}. As will be seen, K_{max} is the single most important parameter of the system.

As with the block size comparison, a combined figure of merit was computed by multiplying the recovered MSE by the encoded MSE. A plot of this figure of merit for varying values of the maximum adaptation factor can be found in Figure 4.8. It is interesting to note that even for relatively small values of K_{max} (*eg.* $K_{max} \leq 8$), changing K_{max} has little effect on the figure of merit. However, within the range $2 \leq K_{max} \leq 8$, varying K_{max} has a substantial effect on both the encoded and decoded MSE figures. Changing K_{max} to reduce the encoded MSE results in an increase in the decoded MSE, and vice-versa, so that the overall figure of merit changes only slightly. Thus, K_{max} is somewhat of a free parameter that may decreased to improve compatibility with standard receivers or increased to improve the noise reduction on advanced receivers.

Perceptually, it was found that once the adaptation factors were allowed to be as large as about 8, the impairments of the picture quality on the standard NTSC receiver due to the encoding process dominated those due to channel degradations. It was certainly clear that an adaptation factor of 16 or greater would not be acceptable. Impairments due to adaptation factors of 4 or less, however, appeared to lie well within the acceptable perceptual limits. The effects of varying K_{max} from 2 to 8 on a typical still image can be seen in Figure 4.9. As a compromise between compatibility and noise reduction, the working value of K_{max} was set at 4.

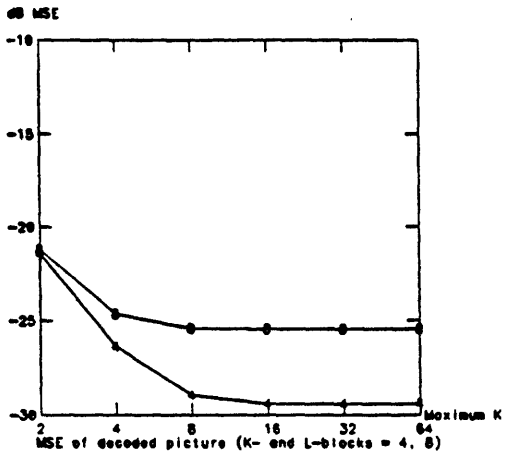

Figure 4.6: Graph of MSE between original picture and picture recovered on advanced receiver for varying values of the maximum adaptation factor K_{max} (blocks sizes of 4×4 and 8×8).

4.3.3 Low-frequency Compression

The issue of how the low-frequency portion of the signal should be compressed is closely related to the allowed dynamic range of the adaptation factors. One would expect that increasing the dynamic range of adaptation factors would require increasing the level of low-frequency signal compression, since larger margins would have to be created to accommodate the adaptively modulated high-frequency portion of the signal. There is no reason to allow large adaptation factors if these margins are going to be the restricting factor in actually selecting the adaptation factors for each block.

After running several examples on typical images, it was determined that low levels of low-frequency signal compression restricted the selection of the adaptation factors only when the maximum allowed adaptation factor was very large (*i.e.* greater than 16). Even in these cases, though, there was no significant loss of perfor-

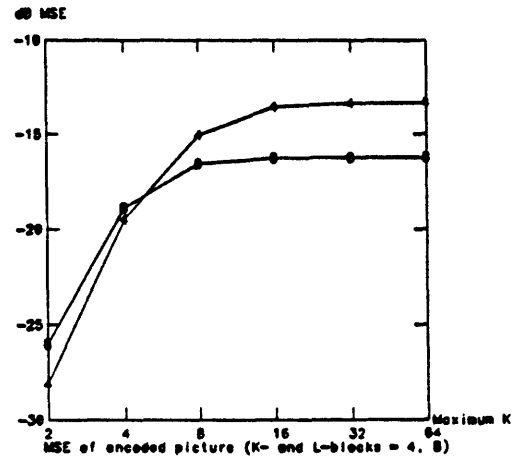

Figure 4.7: Graph of MSE between original picture and encoded picture as viewed on a standard receiver for varying values of the maximum adaptation factor K_{max} (K- and L-blocks of 4×4 and 8×8).

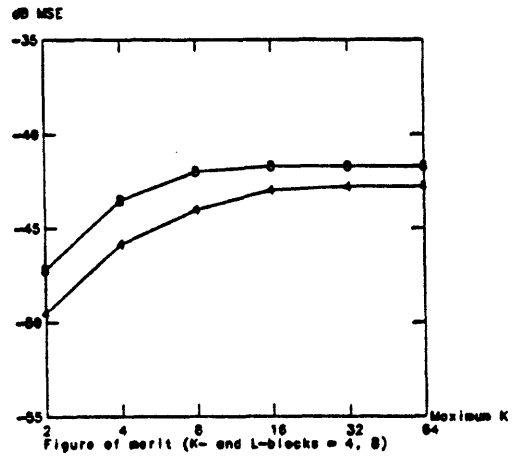

Figure 4.8: Graph of figure of merit for varying values of the maximum adaptation factor K_{max} (K- and L-blocks of 4×4 and 8×8).

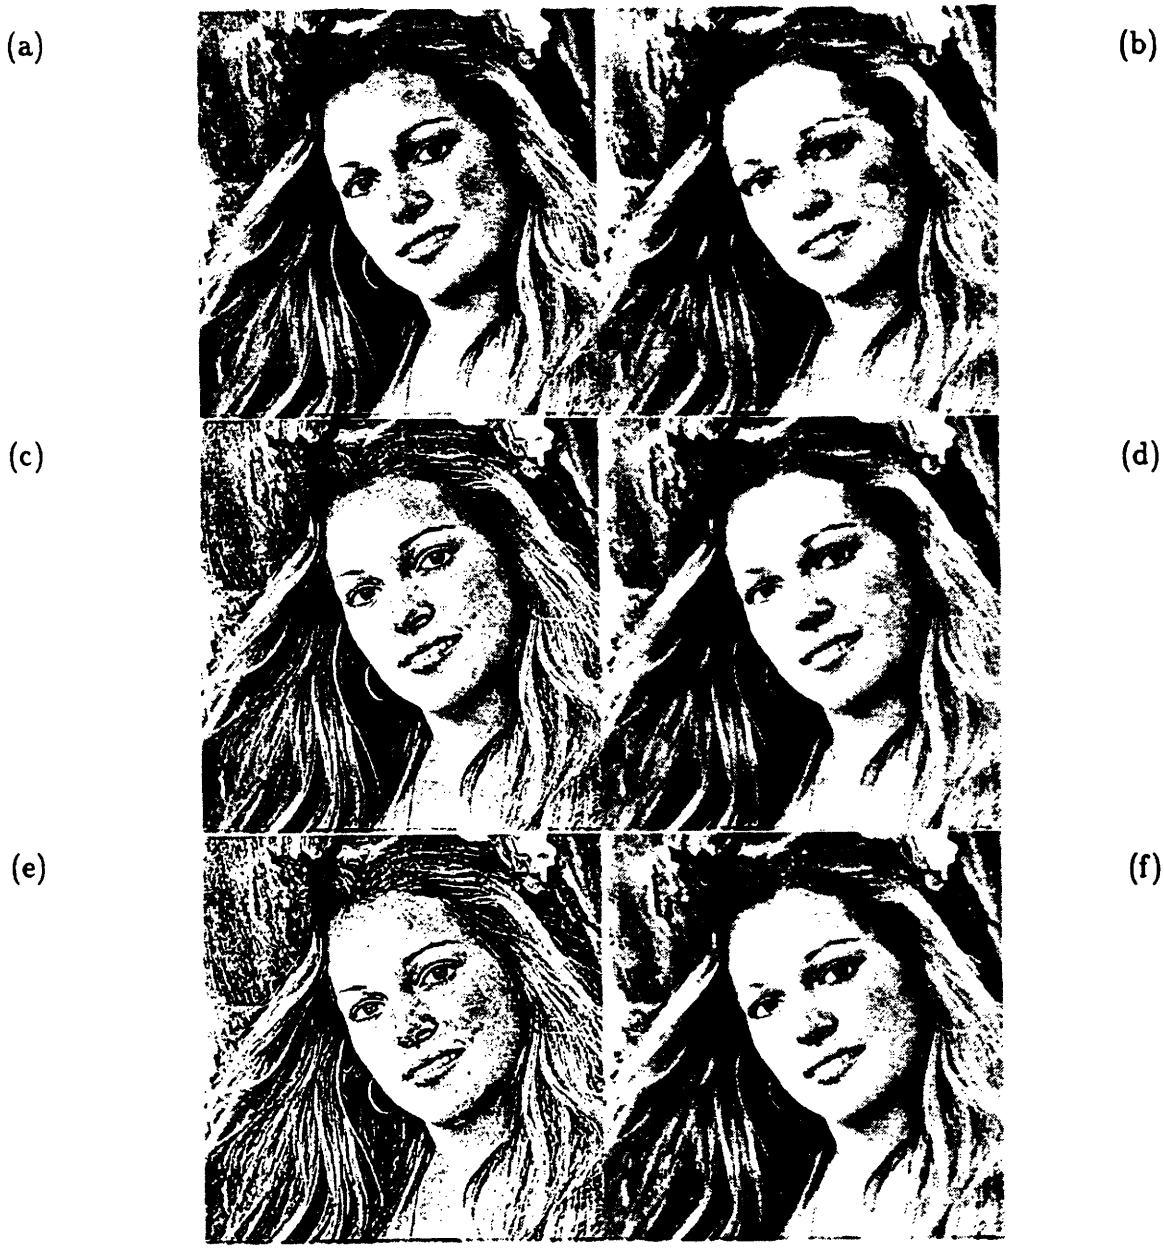

Figure 4.9: Effects of varying the maximum adaptation factor K_{max} (K- and L-blocks of 4×4): $K_{max} = 2$ (a) encoded, (b) recovered. $K_{max} = 4$ (c) encoded, (d) recovered. $K_{max} = 8$ (e) encoded, (f) recovered.

mance with low levels of compression, because there were very few blocks in which large adaptation factors would be selected even if a high level of compression were used. More importantly, it was found that when high levels of compression were used (*i.e.* a reduction in the dynamic range of the low-frequency signal by a factor of more than 2), the resulting receiver-compatible signal was very objectionable because of the overall gray-scale modification. Furthermore, since adaptation factors of larger than 8 proved to be impractical, there was no reason to require a large amount of low-frequency signal compression. In fact, good results could be obtained using no compression at all. However, as one would expect, there was some loss of performance in the regions where an original image had luminance levels near the limits of the dynamic range. For the practical ranges of adaptation factors, it was found that compressing the low-frequency signal to $\frac{7}{8}$ of the original dynamic range (thereby creating upper and lower margins of $\frac{1}{16}$ the total dynamic range) was sufficient. Thus, the low-frequency signal was compressed from an initial dynamic range of [0,1] to a reduced dynamic range of $[\frac{1}{16}, \frac{15}{16}]$.

4.3.4 Parameter Quantization

Initially, the studio quality luminance and chrominance signals have about 8 bits each of amplitude resolution. The adaptation factors don't have any fundamental amplitude resolution, but from the results presented in the previous section, one might guess that about 4 bits would be required to prevent any loss of performance due to quantization. If the L-blocks were chosen to be twice as large in each dimension as the K-blocks, as proposed earlier, roughly 40 bits per L-block would be required to transmit all of the parameters using straight PCM. This initial bit allocation is shown in Table 4.1. Thus, even if 12×8 L-blocks were used (with a maximum of 19.73 bits available per block), a data compression by a factor of slightly more than 2.0 would be required to get the digital data to fit within the allotted bandwidth. Furthermore, if digital bandwidth were to be left available for

Sample	Bits/Block
Y low	8
I low	8
Q low	8
K factor	4×4
Total	40

Table 4.1: Original bit allocation

an audio channel and other auxiliary information, even larger compression factors (roughly 2.5 to 3.0) would be required.

Clearly, the low-frequency samples and adaptation factors must be quantized in some manner. It is important to know, then, exactly how performance is related to the relative coarseness of the quantization. Several tests were conducted to determine how coarsely the low-frequency samples and the adaptation factors would be quantized before the effects of the quantization became objectionable.

The effects of quantizing the low-frequency samples were analyzed first. A test was conducted to determine how quantization of the low-frequency samples effects the close of the fit between the original signal and the interpolant constructed from the samples. The full-precision low-frequency samples were computed and then uniformly quantized to a varying number of reconstruction levels (from 3 to 6 bits of quantization were used). The MSE between the full-precision and each of the quantized versions of the samples was computed. These MSE figures are shown in Figure 4.10. As one would expect, the MSE is almost exactly inversely proportional to the number of reconstruction levels.

Next, the full-precision and quantized low-frequency samples were interpolated to construct approximations to the low-frequency portion of the signal. These interpolants were then compared to the original signal. The MSE between the

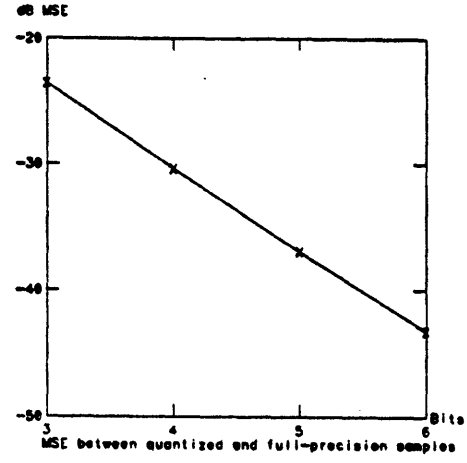

Figure 4.10: Graph of the MSE between the full-precision and quantized low-frequency samples.

original signal and each interpolant was then computed. These MSE figures are shown in Figure 4.11. It is clear that the number of quantization levels has very little effect on the MSE between the interpolant and the original. Comparing these results to those of Section 4.3.1, one can see that the L-block size has a much more direct influence on performance than the number of quantization levels.

When quantization was applied to the low-frequency samples in the encoding and decoding processes, however, blocking effects in both the encoded and recovered pictures became perceptible when fewer than 5 bits of uniform quantization were used. Fortunately, it is possible to exploit the typically high level of statistical redundancy exhibited by most images to lower the required bit rate. Making use of DPCM and nonuniform quantization schemes, 5–6 bits of uniform quantization may be effectively achieved with only 3 bits per sample. This compression will be explained in detail when coding and data compression are discussed.

Next, the issue of quantizing the adaptation factors was considered. Because of

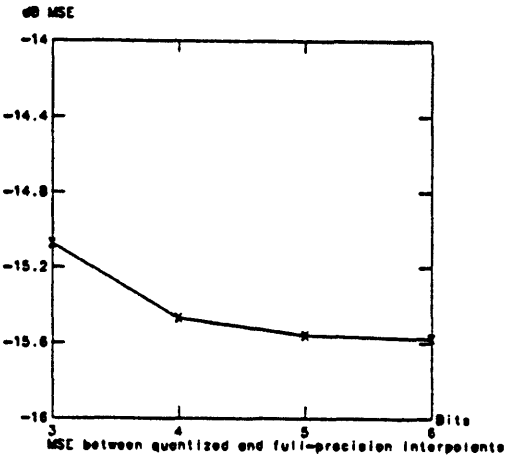

Figure 4.11: Graph of the MSE between the original signal and the interpolants obtained from quantized low-frequency samples.

the nature of the adaptation factors, it was reasoned that logarithmically spaced quantization levels would yield better results than uniformly spaced levels. The idea here is that as the adaptation factor k increases, tightly spaced levels become more and more redundant. For instance, the difference between using $k = 1$ and $k = 2$ is much greater than the difference between using $k = 15$ and $k = 16$, both perceptually and quantitatively (in dB MSE). Assuming, then, that the adaptation factors were limited in dynamic range to the interval $[1, K_{max}]$, the N_k quantization levels for the adaptation factors were

$$K_i = K_{max}^{(i-1)/(N-1)} \text{ for } i = 1, 2, \ldots, N_k. \tag{4.2}$$

Using this quantization scheme, an experiment was conducted to determine how performance would be affected by quantizing the adaptation factors to a varying number of bits. Using L-blocks and K-blocks of sizes 4×4 and 8×8, several still images were encoded, subjected to several channel degradations, and decoded as in the tests mentioned earlier. The maximum allowed adaptation factor was set at 8

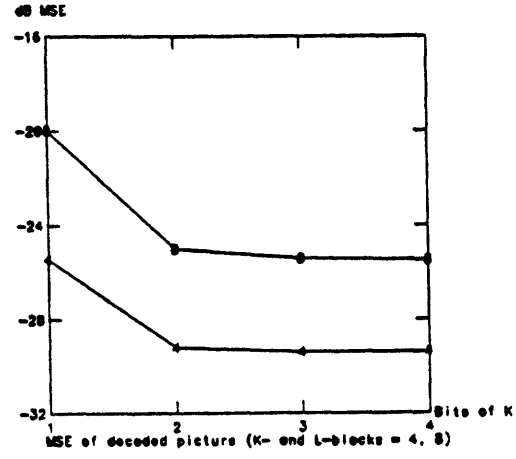

Figure 4.12: Graph of MSE between original picture and picture recovered on advanced receiver for varying amounts of adaptation factor quantization (blocks sizes of 4×4 and 8×8.)

for these trials and from 1 to 4 bits of quantization were used.

As before, a quantitative measure of performance was obtained by computing the MSE between the original picture and each of the encoded and recovered pictures. These MSE figures, as well as the combined figure of merit described earlier, are shown in Figures 4.12, 4.13, and 4.14. Since all of these graphs level off abruptly as the number of bits is increased beyond 2, it is readily apparent that there is very little to be gained by quantizing the adaptation factors to any more than 2 bits. Subjective analysis confirmed this finding. In fact, a perceptually significant, but slightly less than adequate, amount of noise reduction could be achieved even with 1 bit of quantization for the adaptation factors.

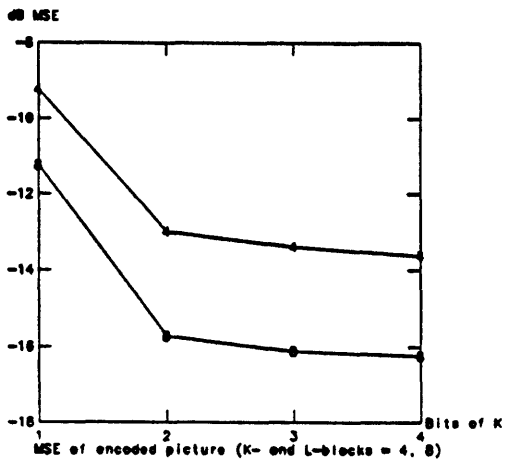

Figure 4.13: Graph of MSE between original picture and encoded picture as viewed on a standard receiver for varying amounts of adaptation factor quantization (K- and L-blocks of 4×4 and 8×8.)

4.4 Chrominance Processing

The addition of color to the picture adds several new problems to the two-channel system. According to Schreiber [19], the chroma are not suitable for two-channel processing because of their typically low-pass nature. Tests were carried out several color still images to determine if this was indeed the case. The pictures were first divided into Y, I, and Q components. Then, each component was encoded by applying the receiver-compatible encoding designed for the luminance. These encoded components were combined to produce the color receiver-compatible signal. The components were again separated, subjected to channel degradations. Finally, each component was decoded using the receiver-end processing for the luminance component described earlier. Finally, the three recovered components were matrixed together to produce the recovered color picture. Trials were run with the L-block and K-block sizes ranging from 4×4 to 8×8 and with the maximum adaptation

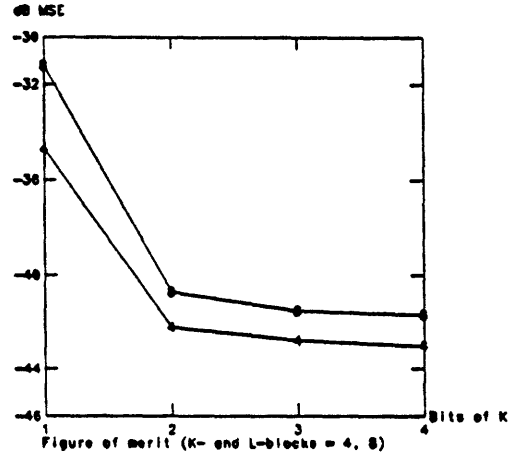

Figure 4.14: Graph of figure of merit between original picture and encoded picture as viewed on a standard receiver for varying amounts of adaptation factor quantization (K- and L-blocks of 4×4 and 8×8.)

factor ranging from 4 to 8. No parameter quantization was done.

What was found was that the effects of the various channel degradations, especially multipath, on the chrominance signal were too severe to make two-channel processing and adaptive modulation practical. In the presence of strong multipath, the colors tended to "leak" beyond edges in the recovered pictures, causing an effect that was far more disturbing than the multipath of the luminance alone. Furthermore, the encoded pictures demonstrated strange color distortions, especially in the vicinity of edges. These color distortions were so perceptually disturbing that the encoded signal could no longer be considered compatible with the standard NTSC receivers. Apparently, because the chrominance signals contain only a very small amount of high-frequency information, the process of adaptive modulation actually makes the overall subjective performance worse near edges.

In addition to the subjective performance problems, performing adaptive modulation on the chrominance signals would also cause a problem with side information. Transmitting the adaptation factors for the I and Q signals would require triple the digital bandwidth needed for transmitting just the adaptation factors for the luminance. Of course, it would be possible to use a single adaptation factor to modulate corresponding blocks of the Y, I, and Q signals, but this strategy also has several problems. First, it would require that the luminance and the chrominance be processed with the same K-block size, which would not be efficient because of the typically low-pass character of the chrominance signals. This property allows that larger blocks be used to process the chrominance than those that are use to process the luminance. More importantly, however, with all three components placing restrictions on the adaptation factor to be selected from each block, the overall performance for any particular block would be limited by the worst of three components in that block. Thus, overall noise reduction would be reduced.

As a result of these observations, it was decided instead that the chrominance should simply be low-pass filtered and subsampled. The samples could then be

interpolated at the receiver to reconstruct fairly good approximations to the original chrominance signals. In order to determine how the sampling and quantization of the chrominance would affect the quality of the recovered picture, a subjective experiment was carried out. Several 512×512 color still images were separated into Y, I, and Q components. The luminance was processed as in the receiver-compatible system described earlier. The I and Q components were low-pass filtered and subsampled, with the chrominance L-block sizes ranging from 4×4 to 16×16. The I and Q components were then quantized using the DPCM scheme of Netravali coupled with mu-law compression of the residuals, using from 3 to 5 bits of quantization. At the receiver end, the luminance component was recovered from the analog receiver compatible signal and the digital side information as in the system for the luminance already described. The chrominance components were recovered by bilinearly interpolating the corresponding digitally encoded samples.

Trials were done both with and without channel degradations being added to the analog receiver-compatible luminance signal. It was found that chrominance L-block sizes of 16×16 or more lead to noticeable blurring of colors near sharp color transitions. Acceptable pictures were recovered when 8×8 blocks were used, and only a marginal improvement in performance was noticeable with 4×4 blocks. In all cases, it was found that 3 bits of quantization was more than sufficient. Based on these results and on the constraints imposed by the digital bandwidth limitation, it was decided that the chrominance components would be processed by subsampling them by a factor of 8×8 and quantizing them to 3 bits.

It should be noted that techniques to sharpen the interpolated chrominance, increasing its effective bandwidth, have been proposed [3]. These techniques make use of the fact that sharp transitions in luminance are usually accompanied by corresponding transitions chrominance. However, because acceptable performance could be obtained without any sharpening, the effects of applying such a technique were not investigated.

After it was determined that the best performance could be obtained simply by reconstructing the low-frequency portion of the chrominance signals from low-frequency samples, the issue of how to process the analog chroma that would be transmitted as part of the receiver-compatible signal remained. There were basically two ways to proceed. The chrominance components could simply be left alone and transmitted along with the receiver-compatible luminance signal. In doing so, one would expect that the colors in the encoded picture would differ from those the original in brightness and saturation, but not in hue. However, in consideration of the fact that the transformation from RGB color coordinates to YIQ color coordinates is linear and that hue should remain constant if R, G, and B are all scaled by the same amount, it was thought that better results could be obtained by appropriately scaling of the chrominance. For instance, suppose $Y(n_1, n_2)$ and $Y^{\mathrm{RC}}(n_1, n_2)$ are the values of the luminance in the original and receiver-compatible signals. The value of the chrominance signals in the receiver-compatible signal would be obtained according to

$$\begin{bmatrix} I^{\mathrm{RC}}(n_1, n_2) \\ Q^{\mathrm{RC}}(n_1, n_2) \end{bmatrix} = \left(\frac{Y^{\mathrm{RC}}(n_1, n_2)}{Y(n_1, n_2)} \right) \begin{bmatrix} I(n_1, n_2) \\ Q(n_1, n_2) \end{bmatrix} \tag{4.3}$$

Because of the fundamental nonlinearity and non-orthogonality of the three-dimensional space of colors, it was unclear which of these two approaches would be better. Thus, both of these schemes were implemented and tested on several color still images. The difference in subjective quality of the resulting color receiver-compatible pictures was judged to be very slight. Both schemes appeared to result in barely noticeable distortions in the colors. These color distortions were expected, even in the cases where the chrominance signals were not processed. They are a result of the fact that the hue of color is not completely independent of its computed Y component. Thus, when the Y component is processed in a nonlinear manner, some small changes in the hue may result. These distortions were not significant enough, however, to destroy the receiver-compatibility of the encoded color signal. Because the scaling the chrominance signals offered no appreciable improvement

over simply transmitting them without any alteration, it was decided on the basis of computational simplicity that the I and Q components of the receiver-compatible signal would require no processing.

4.5 Parameter Coding

The next issue that was addressed was choosing an appropriate coding scheme for the system parameters. Consider first the problem of coding the low-frequency luminance and chrominance samples. Initially, each luminance and chrominance component has about 8 bits of amplitude resolution, which means there are 2^{24} possible color/intensity combinations for each picture element. Of course, as noted earlier, only about 5 bits of uniform quantization are needed for either the luminance or chrominance low-frequency samples in order for the receiver-compatible noise reduction system to achieve adequate overall performance. Suppose that some block coding scheme were used to code these low-frequency samples. Even if a block as small as 4×4 samples were used, there would still be $(2^5)^{(4\times4)} = 2^{80} \approx 10^{24}$ possible combinations for each component. Even if only 1% of those combinations were used as codebook entries (an extremely optimistic amount of compression), a codebook of 10^{22} entries would be required. Storing and searching through such a codebook is simply impossible with current technology. Thus, block coding of the low-frequency samples was ruled out as a feasible option.

However, if the same size L-blocks are used for both the luminance and chrominance, one might consider trying vector quantization (VQ). For each L-block, a three-element vector (Y, I, Q) could be coded. Since each component requires 8 bits, there would be a total of 2^{24} elements in this vector space. Searching through this space to locate the best codeword for a particular Y-I-Q combination would be a difficult task. There do exist sub-optimal coding schemes that divide the three-dimensional space of colors into as many as 2^{22} cells in a manner that facilitates

efficient codebook lookups [19]; however, storing a codebook of that size at the receiver would require more memory than is practical. Even if the luminance and chrominance samples were coarsely quantized to 5 bits (the minimum possible without perceptible quantization noise), there would still be 2^{15} possible color vectors. While it would be reasonable to store a codebook of this size at the receiver, the codebook would be much too large to be transmitted, and it is unlikely that adequate performance could be achieved with a permanently fixed codebook. Thus, vector quantization was also ruled out as a practical alternative.

It was observed, however, that even for relatively large block sizes there is a significant amount of correlation between spatially adjacent low-frequency samples. Because the high degree of spatial correlation between neighboring pixels in a typical image, a substantial compression of the bit rate rate is possible by using a differential PCM (DPCM) coding technique, as opposed to PCM. Since the computational complexity of a DPCM algorithm depends only on the order of the predictive filter that is used, the implementation of a DPCM algorithm is typically much simpler than that of a VQ coding algorithm. Also, there is no need for a codebook with a DPCM algorithm (although one could encode the error residuals to get further bit rate compression). Thus, for the case of coding the the low-frequency samples, DPCM techniques offer savings in both computation and storage, with only a small performance penalty.

With any DPCM technique, the amount of compression that is possible depends on the accuracy of the predictor that is being used. When the input signal can be modeled as a stationary stochastic process (or a stochastic process with slowly varying statistics), the most common choice for a predictor is a linear predictor of the form

$$\hat{f}(n_1, n_2) = \sum_{(k_1,k_2)\in R_a} a(k_1, k_2) f(n_1 - k_1, n_2 - k_2), \tag{4.4}$$

where R_a is the region of support of the predictive filter. The predictor coefficients $a(n_1, n_2)$ are optimized by minimizing the expected squared error $\mathrm{E}[(f(n_1, n_2) -$

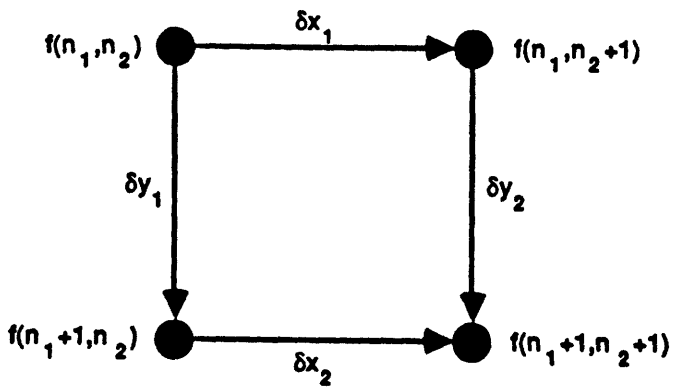

Figure 4.15: Prediction method of Netravali.

$\hat{f}(n_1, n_2))^2]$. The problem with attempting to use this technique to code images, however, is that most images are far from stationary.

Fortunately, due to the high degree of spatial correlation in most images, excellent results can be obtained with simple *ad hoc* predictors. For image coding applications, it is desirable that a predictor respond quickly to sudden changes in intensity and that it be computationally simple. Often, using the value of an adjacent vertical or horizontal pixel, or some simple linear or nonlinear function of the two, is sufficient to achieve a significant compression. One particularly effective nonlinear predictor proposed in [13] attempts to determine the direction in which the spatial correlation is greatest and then uses the value of the appropriate adjacent pixel as a predictor. Suppose that one wishes to predict the value of $f(n_1+1, n_2+1)$ given the (already quantized) values of $f(n_1, n_2)$, $f(n_1, n_2+1)$, and $f(n_1+1, n_2)$, as indicated in Figure 4.15. The first step is to compute $\delta x_1 = f(n_1, n_2+1) - f(n_1, n_2)$ and $\delta y_1 = f(n_1+1, n_2) - f(n_1, n_2)$. Then, if $\delta x_1 > \delta y_1$, $f(n_1, n_2+1)$ is the predictor output and δy_2 is the prediction error. Similarly, if $\delta y_1 > \delta x_1$, $f(n_1+1, n_2)$ is the predicted output and δx_2 is the prediction error.

Once an adequate prediction scheme had been chosen, the next issue was to determine how the error residuals should be quantized. There are basically two

approaches to designing a quantizer. In the first approach, a simple quantization scheme (usually uniform quantization) is used. Such a simple quantizer cannot take full advantage of the redundancy in the error residuals, so some type of variable-length entropy coding (*eg.* Huffman coding or block coding) must be applied to achieve the minimum bit rate. For example, DPCM error residuals typically have either gamma or Laplacian probability distributions, corresponding to the fact that with a good predictor, small error values are much more probable than large ones. If a uniform quantizer is used to code DPCM errors, most of the samples will be quantized to one of the smaller reconstruction levels, while the larger reconstruction levels will be used only infrequently. Thus, the fundamental inefficiency in the quantizer is remedied by entropy coding.

The second approach attempts to eliminate this inefficiency by allowing the quantizer to have nonuniform step sizes. The quantizer decision and reconstruction levels are then tailored to match the probability distribution of the signal being quantized. The problem of designing such a quantizer to minimize the mean square error has been addressed by Max [12] and Lloyd [9]. Suppose that a random variable x with probability distribution $p_x(x)$ is to be quantized to L reconstruction levels $y_0, \cdots, y_{L-1}$ with the $L+1$ decision levels $x_0, \cdots, x_L$. Max and Lloyd have shown that the optimal reconstruction and decision levels must satisfy

$$\begin{aligned} x_0^* &= -\infty \\ x_k^* &= \frac{1}{2}(y_{k-1}^* + y_k^*), \quad k = 1, 2, \cdots, L-1 \\ x_L^* &= \infty \end{aligned} \tag{4.5}$$

and

$$y_k^* = \frac{\int_{x_k^*}^{x_{k+1}^*} x p_x(x)\, dx}{\int_{x_k^*}^{x_{k+1}^*} p_x(x)\, dx}, \quad k = 0, 1, \cdots, L-1 \tag{4.6}$$

In general, it is impossible to directly solve for the x_k^* and y_k^*, so complicated iterative methods must be used instead. However, for most common probability distributions

(particularly for gamma and Laplacian distributions [14]), these values have been computed and tabulated.

The problem with using either of these schemes to quantize images is that the dynamic range of most images, particularly television pictures, is strictly limited, which means that DPCM quantization errors are highly dependent upon the signal value. For instance, if the signal is near the bottom of the dynamic range, then it is known that the error cannot extend very far in the negative direction. Similarly, it is impossible to have a large positive error when the signal is near the top of the dynamic range. The probability distribution of the error is conditioned upon the value of the signal. Thus, to get the maximum efficiency, a different quantizer must be used for every possible signal value, since the probability distribution of the error will be different for each.

While it is certainly possible to estimate each of the conditional error probability distributions, an extremely large sample space of typical images would be required to get any accuracy, especially in the tails of the distributions. Because of this difficulty in estimating the tails of distributions and because worst-case performance is more important than average-case performance for coding television pictures, it is possible to obtain performance close to that of the Max-Lloyd quantizer with a simple suboptimal nonuniform quantization scheme, such as mu-law quantization.

Mu-law quantization provides a good compromise between the efficiency of the Max-Lloyd quantizer and the simplicity of uniform quantization. The basic idea behind mu-law quantization is to use a uniform quantizer, but to apply it to a nonlinearly distorted version of the source signal. A mu-law encoder/decoder system is shown in Figure 4.16. The nonlinear compression function $c(x)$ is given by

$$c(x) = x_{max} \frac{\log_e \left(1 + \frac{\mu|x|}{x_{max}}\right)}{\log_e (1 + \mu)} \text{sgn}(x). \tag{4.7}$$

This compression characteristic is linear for small signal values ($|x| \ll x_{max}$) and logarithmic for large signal values ($|x| \approx x_{max}$). Traditionally, the parameter μ is chosen to be a number of the form $2^m - 1$, so that the above expression reduces to

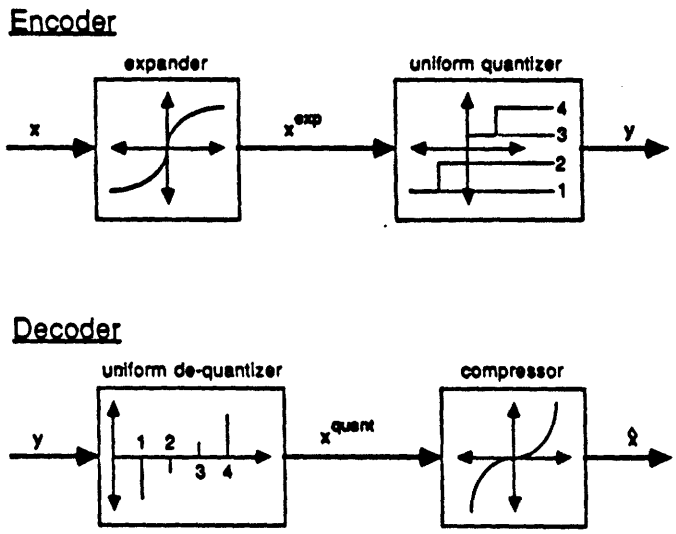

Figure 4.16: A mu-law encoder and decoder.

a logarithm base 2^m, which is generally easy to compute. It was determined after several brief tests that $\mu = 63$ yields a good compression characteristic for coding television pictures.

A test was done to determine how much improvement there was to be gained in going to the DPCM technique described above. The full-precision low-frequency samples were computed for several still images. These samples were then quantized to from 3 to 6 bits using several different quantizers. The first quantizer was a simple uniform PCM quantizer. Next, non-adaptive one-step horizontal and vertical predictors, with mu-law quantization of the error residuals, were tried. Finally, the combination of Netravali's prediction scheme and mu-law quantization was tested. In each case, the error between the quantized version of the samples and the full-precision samples was computed. These MSE figures are shown in Figure 4.18(a). For each case, a second error measure was obtained by interpolating the sample values (using bilinear interpolation) and computing the error between the resulting interpolant and the original image. These MSE figures are shown in Figure 4.18(b).

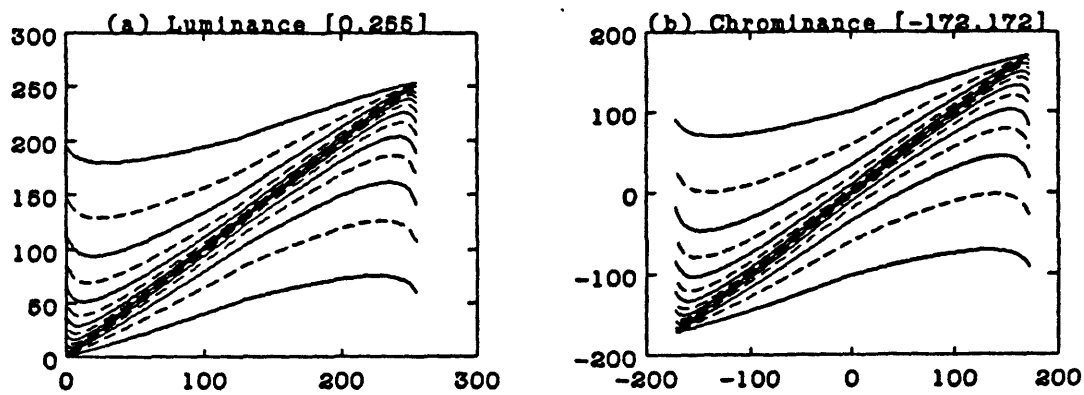

Figure 4.17: Reconstruction levels (solid) and decision levels (dashed) as a function of the input pixel value for (a) luminance and (b) chrominance.

As one would expect, in almost every case the magnitude of the error is inversely related to the number of quantization levels. One can see that the quantization scheme based on Netravali's predictor offers roughly a 3 dB improvement over straight PCM and about a 1.5 dB improvement over either of other nonadaptive one-step predictors. It is interesting to note, however, that these improvements are almost completely eliminated when the samples are interpolated. The interpolant generated from the quantized samples have almost identical errors. These quantitative results seem to indicate that many quantization levels are not necessary and that the underlying sampling structure is the greater factor in determining how well the interpolant will fit the original. Subjective results, however, indicate that using fewer than 3 bits of quantization leads to noticeable blocking effects. Thus, due to digital bandwidth limitations, it was decided that 3 bits would be used to quantize the low-frequency samples.

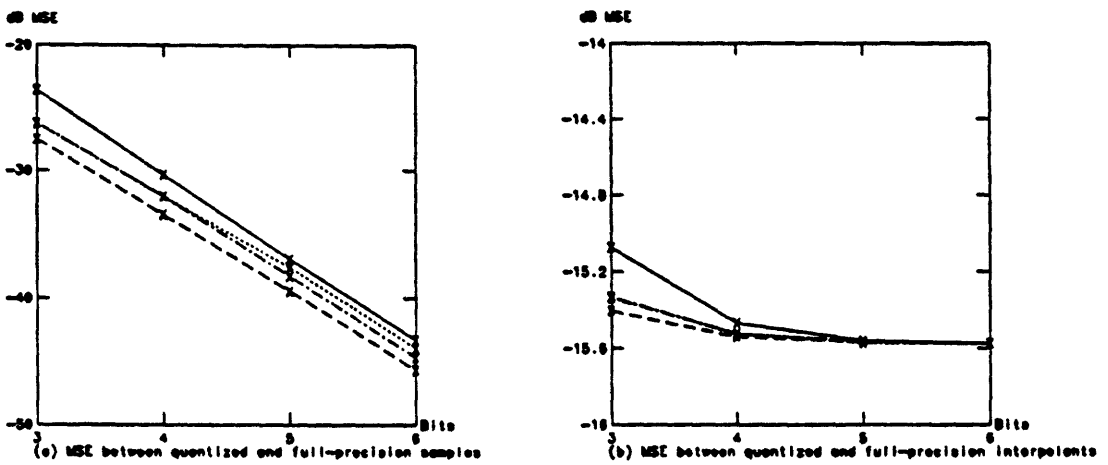

Figure 4.18: Comparison of Netravali's DPCM (dashed) to standard PCM (solid), horizontal one-step prediction (dotted), and vertical one-step prediction (dash-dotted). These results are for a typical 512×512 still image with 8×8 L-blocks for a varying number of bits per sample. (a) MSE between quantized and full-precision low-frequency samples. (b) MSE between quantized and full-precision interpolants.

4.6 Performance Improvement

After all of the above mentioned parameters had been optimized, there were still several problems with the receiver-compatible noise reduction system. The most noticeable of these problems was the break-down in performance in the vicinity of sharp edges. As one would expect, the low-frequency portion of the signal is generally not an adequate approximation of the original signal in these regions. The result is that the high-frequency portion of the signal has large amplitudes near sharp edges. Since the maximum amplitude of the high-frequency signal limits the size of the adaptation factor, the performance of the noise reduction system will be impaired in these regions.

Another problem that had to be addressed was the performance trade-off involved in the selection of an appropriate value for K_{max}. As mentioned earlier, the parameter K_{max} may be varied over a range in order to improve either compatibility with standard receivers or noise reduction on advanced receivers, but not both. Ideally, one would want a high level of both compatibility and noise reduction, instead of a compromise at some intermediate value of K_{max}. This problem, too, is due in part to a poor fit between the original signal and its low-frequency portion. If there existed a way to improve this fit, it would be possible to increase K_{max} without increasing level of distortion of the encoded picture. Thus, the trade-off would be eased, since more noise would be removed from the decoded picture.

In an attempt to reduce these problems, several different methods were investigated. First, a smoother interpolation technique was implemented and tested. Next, a variety a methods for choosing the optimal sample values were investigated. Finally, experiments were conducted with a sophisticated nonuniform sampling technique. The underlying idea behind all of these techniques was to improve the fit of the low-frequency interpolant to the original signal (particularly near edges), thereby reducing the amplitude of the high-frequency residual portion of the signal.

4.6.1 Interpolation Scheme

The first attempt to improve the performance of the system was to alter the interpolation scheme. It was thought that one of the major factors contributing to the lack of performance near sharp edges was the relative inaccuracy of the bilinear interpolation scheme. Thus, in an attempt to remedy the problem, smoother interpolating filters were tried.

Consider the interpolation problem in one dimension. The filter $h_L(n)$ that corresponds to linear interpolation is the triangular filter given by

$$h_L(n) = \begin{cases} 1 - \frac{1}{B}\left|n - \frac{1}{2}\right|, & n = -B+1, -B+2, \ldots, B-1, B \\ 0, & \text{otherwise} \end{cases} \tag{4.8}$$

where the block size B is assumed to be even (so that the filter has a half-sample delay). Performing bilinear interpolation in two dimensions corresponds to convolving the zero-padded subsampled signal with the separable filter $h_{L2}(n_1, n_2) = h_L(n_1)h_L(n_2)$.

One might expect that better results could be obtained by performing ideal low-pass filtering and subsampling to get the low-frequency, and then using the same ideal low-pass filter to do the interpolation. Of course, it is impractical to use ideal low-pass filters because of their infinite extent, but a windowed version can be used instead with very little change in performance. The filters $h_{S2}(n_1, n_2)$ that were tried were obtained from the ideal low-pass filter $h_I(n_1, n_2)$ by windowing it with separable Hamming windows of various sizes.

Tests were done using window sizes of 4, 6, and 8 times the L-block size, so that the resulting windowed filters had 4×4, 6×6, and 8×8 taps respectively. The L-block size was also varied from 4×4 to 8×8. The original still pictures were encoded, subjected to channel degradations, and then decoded to produce the recovered pictures. The MSE between the original picture and each of the encoded and recovered pictures was computed to provide a means of quantitative performance comparison. These MSE figures are given in Tables 4.2 and 4.3. As

		dB MSE	
filter type	taps	recovered	encoded
bilinear	2	-26.97	-18.12
windowed ideal LPF	4	-27.68	-17.28
	6	-27.86	-17.05
	8	-27.70	-17.41

Table 4.2: Comparison of bilinear interpolation to (windowed) ideal low-pass interpolation for various window sizes. In each case, the filter size is the number of taps times the L-block size, which is 4×4 for all cases.

		dB MSE	
filter type	taps	recovered	encoded
bilinear	2	-26.12	-16.17
windowed ideal LPF	6	-27.18	-15.07

Table 4.3: Comparison of bilinear interpolation to (windowed) ideal low-pass interpolation for 8×8 L-blocks.

one can see from this data, using the windowed ideal low-pass filters appears to offer about a 1 dB improvement in the recovered pictures MSE at the cost of making the encoded pictures MSE about 1 dB worse. These results appear to hold regardless of block size or number of filter taps. Thus, there is no advantage in going to smoother interpolation from a MSE standpoint, since this same trade-off could be achieved by varying other parameters.

Based on subjective perceptual criteria, however, the simple bilinear interpolation did a much better job. The improvement in the recovered pictures due to using the windowed ideal low-pass filters was not perceptible. However, the pictures encoded using the windowed ideal low-pass filtering and interpolation appeared to

be far worse than those encoded with simple block averaging and bilinear interpolation. In the blank areas, there was a significant increase in the grainyness, and near sharp edges, there was a disturbing ringing effect. These effects appeared to get worse as the block size was increased. Thus, it was determined that smoother interpolation was not an effective means to improve performance.

4.6.2 Fixed Least-Squares Fit

When changing the interpolation filters didn't yield much improvement in the performance, an investigation was done to determine whether there was a better way to choose the low-frequency samples that were being interpolated. The hope was that the fit of the low-frequency interpolant could be improved by applying a least-squares criterion to choose the "optimal" sample values, rather than relying on simple boxcar or triangular filters to get them.

Given that linear interpolation is to be used and that the locations of the sample points are fixed (although not necessarily uniform), this least-squares optimization problem may be set up rather easily. Suppose that the objective is to approximate the function $f(x)$, which is twice differentiable on $[a, b]$, by a piecewise-linear interpolant $g(x)$. Assume furthermore that the breakpoints between which the interpolation will be done are

$$a = x_0 < x_1 < x_2 < \cdots < x_{N-1} < x_N = b. \tag{4.9}$$

For the sake of convenience, let $x_{-1} = x_0$ and $x_{N+1} = x_N$. It is then possible to express the interpolant $g(x)$ as a weighted sum of basis functions

$$g(x) = \sum_{i=0}^{N} \alpha_i h_i(x) \tag{4.10}$$

where $\alpha_i = g(x_i)$ and the $h_i(x)$ are a series of shifted triangular window functions

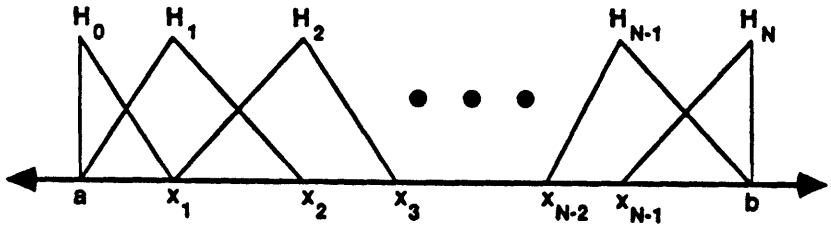

Figure 4.19: The basis functions $h_i(x)$.

given by

$$h_i(x) = \begin{cases} \frac{x - x_{i-1}}{x_i - x_{i-1}}, & x_{i-1} \leq x < x_i \\ \frac{x_{i+1} - x}{x_{i+1} - x_i}, & x_i \leq x < x_{i+1} \\ 0, & \text{otherwise} \end{cases} \tag{4.11}$$

as shown in Figure 4.19.

The objective, then, is to determine the values of the parameters α_i such that the expression

$$e_{int} = \int_a^b |f(x) - g(x)|^2 \, dx \tag{4.12}$$

is minimized. The form of the solution is independent of the basis functions and is given by [2]. The α_i must satisfy the set of linear equations given by

$$\sum_{j=0}^{N} \left[\int_a^b h_i(x) h_j(x) \, dx \right] \hat{\alpha}_j = \int_a^b h_i(x) f(x) \, dx. \tag{4.13}$$

In the case of uniformly spaced with a gap of λ and linear interpolation, these equations may be reduced to the tridiagonal Toeplitz system

$$\left(\frac{\lambda}{6}\right) \hat{\alpha}_{i-1} + \left(\frac{2\lambda}{3}\right) \hat{\alpha}_i + \left(\frac{\lambda}{6}\right) \hat{\alpha}_{i+1} = \beta_i \triangleq \int_a^b h_i(x) f(x) \, dx. \tag{4.14}$$

The two-dimensional extension of this derivation is straightforward. Assume that the goal is to approximate a two-dimensional function $f(x, y)$ on the mesh

$[a_x, b_x] \times [a_y, b_y]$ by an interpolant

$$g(x,y) = \sum_{i=0}^{N_x} \sum_{j=0}^{N_y} \alpha_{i,j} h_{i,j}(x,y)$$

where the basis functions $h_{i,j}(x,y)$ are a series of shifted pyramid-shaped windows (obtained by multiplying a triangular window in one dimension with a triangular window in the other). Let the sample spacing be λ_x in the x-direction and λ_y in the y-direction. Solving for the $\alpha_{i,j}$ that minimize the mean square error between the interpolant $g(x)$ and the original function $f(x)$ yields a highly structured block tridiagonal Toeplitz system defined by

$$\begin{aligned} \beta_{i,j} &\triangleq \int_{a_x}^{b_x} \int_{a_y}^{b_y} h_{i,j}(x,y) f(x,y)\, dx\, dy \\ &= \left(\tfrac{\lambda_x \lambda_y}{36}\right) \hat{\alpha}_{i-1,j-1} + \left(\tfrac{\lambda_x \lambda_y}{9}\right) \hat{\alpha}_{i-1,j} + \left(\tfrac{\lambda_x \lambda_y}{36}\right) \hat{\alpha}_{i-1,j+1} + \\ &\quad \left(\tfrac{\lambda_x \lambda_y}{9}\right) \hat{\alpha}_{i,j-1} + \left(\tfrac{4\lambda_x \lambda_y}{9}\right) \hat{\alpha}_{i,j} + \left(\tfrac{\lambda_x \lambda_y}{9}\right) \hat{\alpha}_{i,j+1} + \\ &\quad \left(\tfrac{\lambda_x \lambda_y}{36}\right) \hat{\alpha}_{i+1,j-1} + \left(\tfrac{\lambda_x \lambda_y}{9}\right) \hat{\alpha}_{i+1,j} + \left(\tfrac{\lambda_x \lambda_y}{36}\right) \hat{\alpha}_{i+1,j+1} \end{aligned} \tag{4.15}$$

Despite the fact that this system of equations is fairly large (it is $N_x N_y \times N_x N_y$, where N_x and N_y are generally in the range 30–120), it requires relatively little computational effort to solve because of its extremely regular structure. In fact, the most computionally intense part of solving the system is computing the $\beta_{i,j}$ by multipying $f(x,y)$ by all of the windows.

Unfortunately, this least-squares optimization of the sample values offers little improvement over simple averaging or low-pass filtering. Analysis of the situation revealed that the major reason for large differences between the interpolant and the original function was not because the sample values were being chosen poorly, but because the uniform sampling grid made it difficult to find an interpolant that matched the original function well in the vicinity of sharp edges.

4.6.3 Optimization by Linear Programming

Given the failure of the least-squares signal approximation approach to choosing the low-frequency sample values, one might contend that the correct thing to do is to

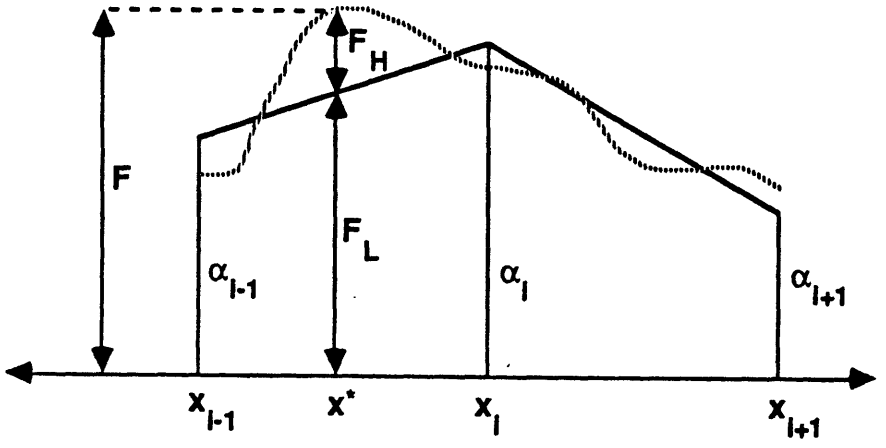

Figure 4.20: One-dimensional linear programming problem.

choose the samples so that the adaptation factors are maximized, rather than trying to match the original signal as closely as possible. Of course, these two criteria are highly interdependent, but they are not equivalent. The following derivation of the algorithm for selecting the sample values based on this criterion should make the differences clear.

Consider the continuous interpolation problem in one dimension. As in the previous section, the sample points, or interval breakpoints, are x_i and the values of the low-frequency samples at those points are α_i for $i = 0, 1, \ldots, N$. Linear interpolation is used, so that the value of the interpolant between two consecutive breakpoints x_i and x_{i+1} is

$$g(x) = \left(\frac{x - x_i}{x_{i+1} - x_i}\right) \alpha_{i+1} + \left(\frac{x - x_{i+1}}{x_i - x_{i+1}}\right) \alpha_i. \tag{4.16}$$

Now, assume that α_{i-1} and α_{i+1} are fixed and that α_i is a free parameter. It is possible to view the value of the interpolant $g(x)$ on the interval $[x_{i-1}, x_{i+1}]$ as a bilinear function of the coordinate x and the sample value α_i. In particular, one

may express the interpolant as

$$g(x) = \begin{cases} p_i(x) + q_i(x)\alpha_i & x_{i-1} \le x \le x_i \\ r_i(x) + s_i(x)\alpha_i & x_i \le x \le x_{i+1} \end{cases} \tag{4.17}$$

where

$$\begin{aligned} p_i(x) &= \left(\tfrac{x-x_i}{x_{i-1}-x_i}\right)\alpha_{i-1} & q_i(x) &= \left(\tfrac{x-x_{i-1}}{x_i-x_{i-1}}\right) \\ r_i(x) &= \left(\tfrac{x-x_i}{x_{i+1}-x_i}\right)\alpha_{i+1} & s_i(x) &= \left(\tfrac{x-x_{i+1}}{x_i-x_{i+1}}\right). \end{aligned} \tag{4.18}$$

Now consider a particular point $x^* \in [x_{i-1}, x_1]$. For simplicity, make the following definitions

$$F = f(x^*) \tag{4.19}$$

$$P = p_i(x^*) \tag{4.20}$$

$$Q = q_i(x^*), \tag{4.21}$$

where $f(x^*)$ is the value of the original signal at x^*. With these definitions, the low-frequency and high frequency portions of the signal at x^* may be expressed as

$$F_L = P + Q\alpha_i \tag{4.22}$$

$$F_H = F - P - Q\alpha_i. \tag{4.23}$$

Now, assume that receiver-compatible adaptive modulation is used, so that the low-frequency portion of the signal undergoes a linear compression according to

$$f^{\text{com}} = af + b. \tag{4.24}$$

where a and b are fixed constants that depend on the original dynamic range $[A_0, B_0]$ and the compressed range $[A_L, B_L]$. In accordance with this compression rule, define

$$F^{\text{com}} = aF + b \tag{4.25}$$

$$F_L^{\text{com}} = aP + aQ\alpha_i + b. \tag{4.26}$$

Now (3.4) may be applied to determine how varying α_i affects the value of the adaptation factor $K = k(x^*)$ at x^*. If F_H is positive ($\alpha_i < \frac{F-P}{Q}$), then the adaptation

factor K is limited by the inequality

$$\begin{aligned} K &\leq \frac{B_0 - F_L^{\text{com}}}{F_H} \\ &= \frac{B_0 - aP - AQ\alpha_i - b}{F - P - Q\alpha_i} \\ &= a + \frac{B_0 - (aF + b)}{F - P - Q\alpha_i} \\ &= a + \frac{B_0 - F_L^{\text{com}}}{F - P - Q\alpha_i}. \end{aligned} \tag{4.27}$$

Similarly, if F_H is negative ($\alpha_i > \frac{F-P}{Q}$), the adaptation factor K must satisfy

$$\begin{aligned} K &\leq \frac{F_L^{\text{com}} - A_0}{-F_H} \\ &= a + \frac{A_0 - F_L^{\text{com}}}{F - P - Q\alpha_i}. \end{aligned} \tag{4.28}$$

Thus, given α_{i-1} and $f(x^*)$, it is possible to get a constraint on the maximum adaptation factor as a function of α_i for any point of interest $x^* \in [x_{i-1}, x_i]$. A similar constraint may be derived given α_{i+1} and $x^* \in [x_i, x_{i+1}]$. An example of such a constraint function is shown in Figure 4.21.

Unfortunately, this constraint is not linearly dependent on the free parameter α_i. This problem may, however, by overcome by making the change of variables

$$\kappa = \frac{1}{K - a} \geq \begin{cases} \frac{F-P-Q\alpha_i}{B_0 - F_L^{\text{com}}}, & \alpha_i < \frac{F-P}{Q} \\ \frac{F-P-Q\alpha_i}{A_0 - F_L^{\text{com}}}, & \alpha_i > \frac{F-P}{Q}. \end{cases} \tag{4.29}$$

Since K is always greater than a, maximizing K corresponds to minimizing κ. An example of a constraint function for κ is shown in Figure 4.22. It is interesting to note the very special structure of this constraint on κ. Regardless of the value of x^*, $f(x^*)$, $p_i(x^*)$, or $q_i(x^*)$, the constraint is a "V"-shaped curve with its vertex sitting on the horizontal axis. Furthermore, it can be shown that for a fixed x^* (and therefore for fixed $p_i(x^*)$ and $q_i(x^*)$ as well), both the left and right branches of each "V-constraint" rotate clockwise as $f(x^*)$ increases.

In the case of a discretely (although not necessarily uniformly) sampled signal, there are many sample points y_j, $j = 0, 1, \ldots, M$ within the interval $[x_{i-1}, x_i]$. If it

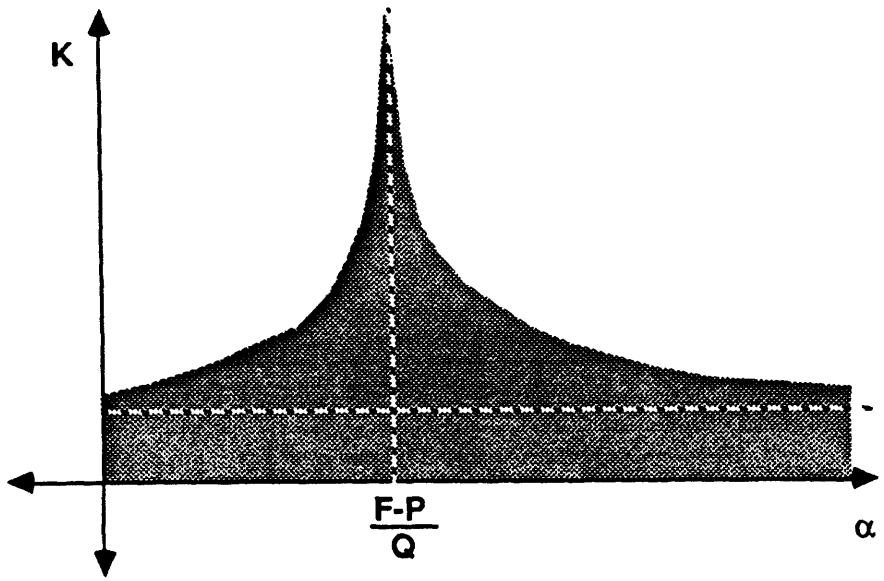

Figure 4.21: Maximum adaptation factor at a particular point $x^* \in [x_{i-1}, x_i]$ as a function of α_i. Feasible adaptation factors are in the shaded region.

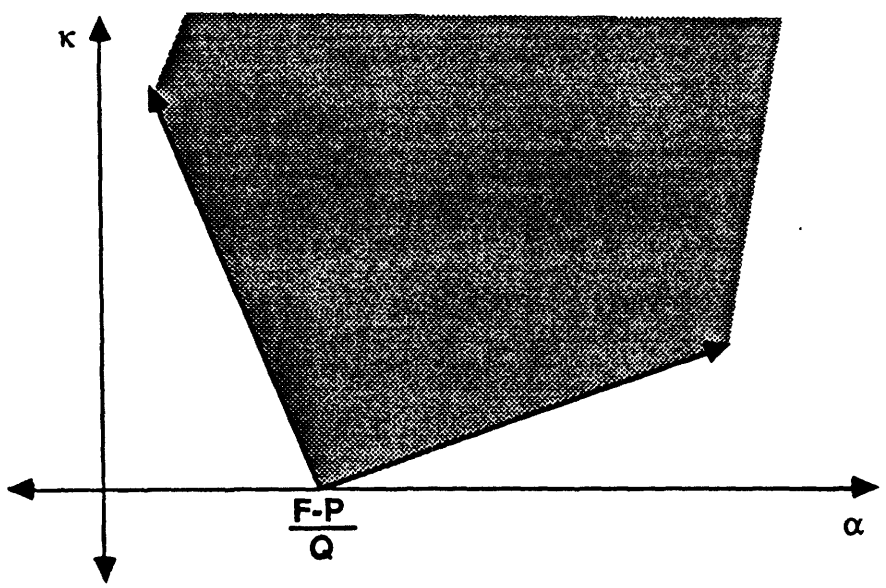

Figure 4.22: Feasible region for κ at a particular point $x^* \in [x_{i-1}, x_i]$ as a function of α_i.

is assumed that α_{i-1} and α_{i+1} are fixed, then at each point y_j it is possible to derive a constraint on the range of possible adaptation factors that depends linearly on α_i. If these nonlinear constraints on $k(x)$ for $x = y_j$, $j = 0, 1, \ldots, M$ are mapped to linear constraints on $\kappa(x)$, then the resulting set of constraints is as shown in Figure 4.23.

Next, it may be noted that the magnitude of the adaptation factor for any particular block is set to be the maximum value such that the dynamic range limitations at each point in the block is satisfied. For each value of α_i, there is a feasible range of values of the adaptation factor K which satisfy all of the constraints. Clearly, there must also be a feasible range of κ values for each α_i. Thus, the problem of choosing the optimal value for α_i,

$$\begin{aligned} \alpha_i^* &= \arg\max[\min_{x \in [x_{i-1}, x_{i+1}]} k(x)] \\ &= \arg\min[\max_{x \in [x_{i-1}, x_{i+1}]} \kappa(x)], \end{aligned} \tag{4.30}$$

reduces to a minimax problem. A problem with this structure may be solved by a variety of well-known linear programming algorithms. In fact, because of the special "V" shape of all of the constraints, a very simple algorithm may be used.

Clearly, it would be impossible to optimize all of the sample values simultaneously taking this approach, since to optimize a sample value requires that the neighboring sample values be fixed. However, it is possible to apply this linear programming technique on an iterative block-by-block basis in a manner similar to Gauss-Seidel iteration. To compute a new value of α_i, the best current estimates of the optimal α_{i-1} and α_{i+1} are used. Several passes over all of the sample points are required to get the sample values to converge to within an acceptable tolerance of their optimal values.

Clearly, this algorithm may be easily extended to handle a two-dimensional signal. However, since the value of an "interior" point of the interpolant depends on the four (as opposed to two) surrounding sample values, more sample values must be fixed in order to optimize a particular sample value. In particular, to solve

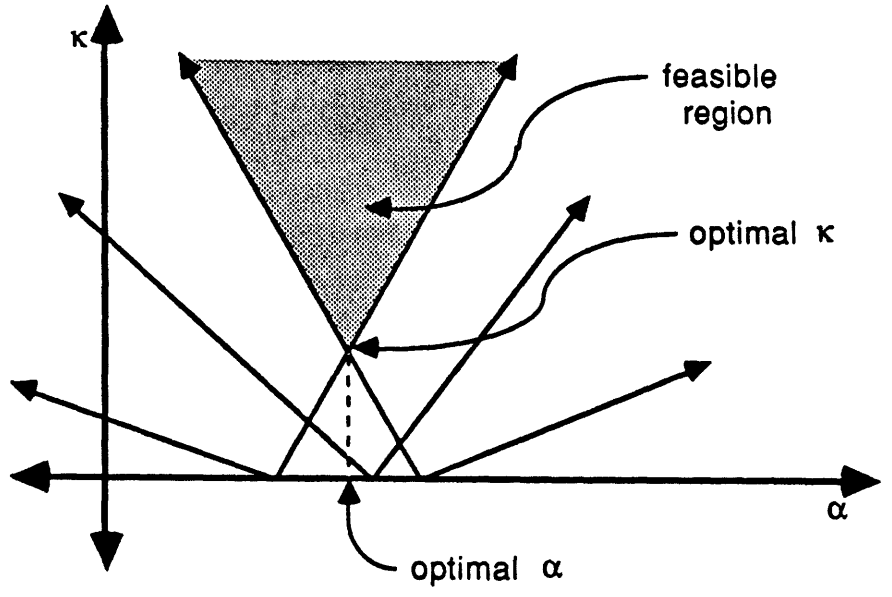

Figure 4.23: Constraints on the adaptation factor and the feasible region.

for a particular sample point $\alpha_{i,j}$, the two vertically adjacent ($\alpha_{i-1,j}$ and $\alpha_{i+1,j}$, two horizontally adjacent ($\alpha_{i,j-1}$ and $\alpha_{i,j+1}$, and four diagonally adjacent ($\alpha_{i-1,j-1}$, $\alpha_{i-1,j+1}$, $\alpha_{i+1,j-1}$, and $\alpha_{i+1,j+1}$) must be specified, as shown in Figure 4.24.

Unfortunately, when this sample value optimization algorithm was implemented, there was no significant improvement in performance over that achieved by simple block averaging. This new optimization algorithm also had the drawback that its computational requirement was enormous. Even though the number of iterations required to converge to a reasonably stable solution was few, the number of computations per iteration were enormous, since a complete linear programming problem had to be set up and solved for each block. As a result, this method was ruled out as a practical alternative.

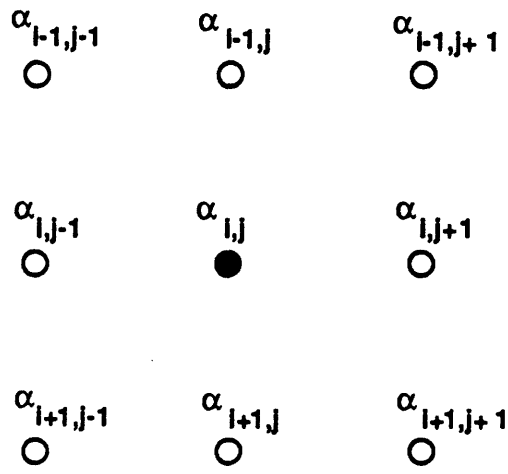

Figure 4.24: Optimization of a single sample value in the two-dimensional case requires the specification of the eight surrounding sample values.

4.6.4 Nonuniform Sampling

After the experiments both with varying interpolation filters and varying sample selection algorithms failed to produce significant performance improvements, it was realized that choosing the proper spacing between sample points is much more important for improving the fit of the low-frequency interpolant than the methods used to select the sample values or interpolate between them.

If a sharp edge falls exactly half-way between two sample points, the interpolant is not going to approximate the original signal very well in that interval. Clearly, this problem may be reduced by going to smaller block sizes. This approach, however, requires the availability of a very high bandwidth digital channel to send all of the block means and adaptation factors. A more practical solution is to use nonuniform sampling and nonuniform interpolation. Ideally, one would like to sample the original image densely near sharp edges and sparsely in slowly varying background

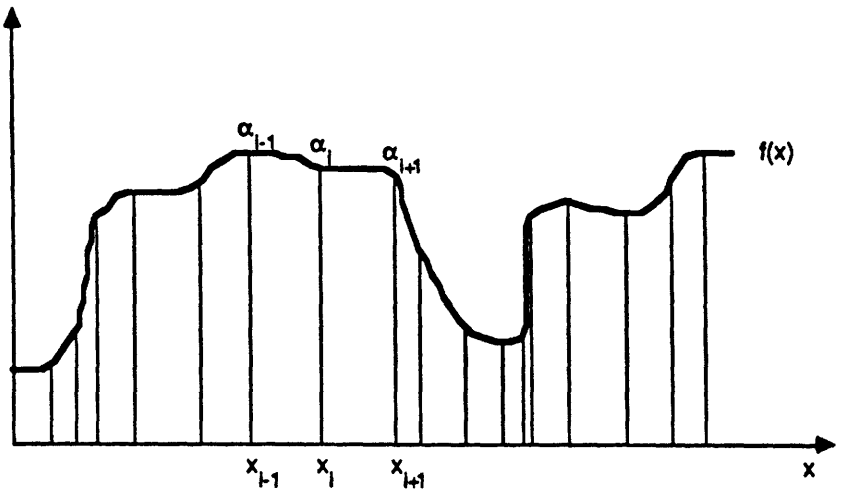

Figure 4.25: Non-uniform sampling of a one-dimensional function.

regions, as in Figure 4.25.

If it is known that linear interpolation is to be used, then there is an optimal least squares solution to the problem of choosing the breakpoints between blocks. As before, it is assumed that the objective is to approximate the function $f(x)$, which is twice differentiable on $[a,b]$ by a piecewise-linear interpolant $g(x)$. The difference now is that the breakpoints are allowed to vary instead of being fixed at pre-determined (typically uniformly spaced) locations.

Intuitively, one would guess that the best way to choose these breakpoints is to sample most densely in the regions where there is the most "activity" in the function $f(x)$. Ideally, one would like to have some measure of activity per unit length, so that it would be possible to segment the interval into regions of equal activity. As it turns out, the optimal way to choose the breakpoints x_i is such that [2]

$$\int_a^{x_i} \sqrt{|f''(x)|}dx = \frac{i}{N}\int_a^b \sqrt{|f''(x)|}dx, \quad \text{for all } i \tag{4.31}$$

which indicates that the square root of the second derivative of $f(x)$ is the best

measure of activity. In other words, the optimal solution is to divide the interval into regions such that the integral of the square root of the second derivative of $f(x)$ is constant across all regions, *i.e.*

$$\int_{x_i}^{x_{i+1}} \sqrt{|f''(x)|}dx = C_0 \tag{4.32}$$

for some constant C_0. This results makes good intuitive sense. The second derivative of a function is closely related to its curvature

$$K = \frac{f''(x)}{\sqrt{1 + (f'(x))^2}}. \tag{4.33}$$

Since curvature is a measure of how rapidly the function is changing direction (essentially it is the inverse of the instantaneous radius of curvature), it is desirable to place more samples in the regions where the curvature is high.

While these results apply to continuous signals, it is clear that a digitized signal can be (low-pass) filtered so that a reasonable estimate of the second derivative of the sampled analog signal can be obtained. The underlying sampling grid limits the granularity of the breakpoint selection, but that should not be a problem, since the only points at which the value of the interpolant matters are the grid points.

A test was conducted to determine how much the fit of the low-frequency interpolant could be improved by using nonuniform sampling. Several 512-point lines were extracted from a two-dimensional image. Each line was then sampled and interpolated using both the uniform and nonuniform sampling strategies, with the number of samples per line varying from 16 to 128.

The interpolants obtained by uniform and nonuniform sampling for the case of 64 sample points are shown in Figure 4.26. Note that the nonuniform interpolant fits the original signal better at the sharp transitions than the uniform interpolant does. The price that must be paid for this improved performance near edges, however, is a slightly worse fit in the more smoothly varying regions.

The MSE between the interpolant and the original signal was computed for each case. The results are shown in Figure 4.27. From these results, one can see that

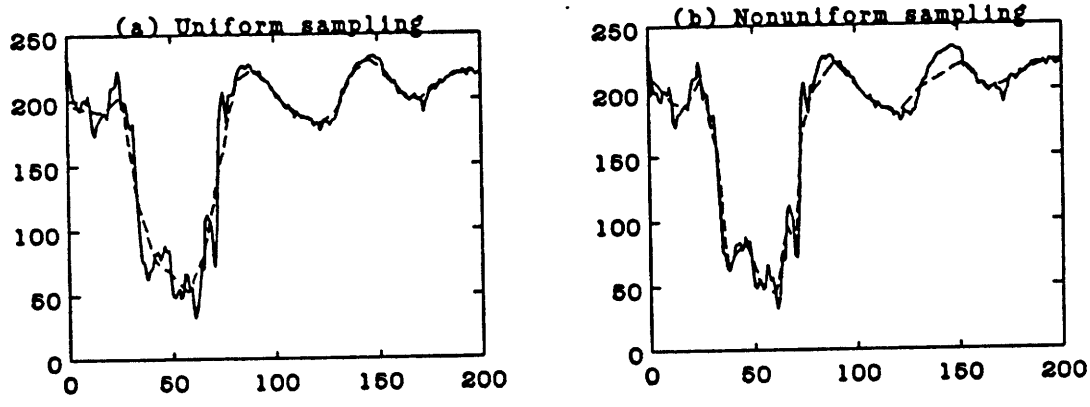

Figure 4.26: Interpolants obtained by (a) uniform sampling and (b) nonuniform sampling. In each case the interpolant (dashed line) was was constructed from 64 samples of the 512-point original (solid line).

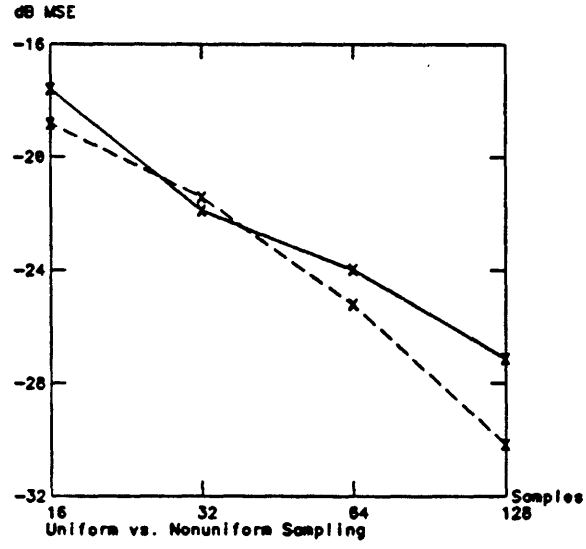

Figure 4.27: Graph of the MSE between the low-frequency interpolant and the original signal, which was one line of a typical image, for uniform (solid line) and nonuniform (dashed line) sampling.

nonuiform sampling does offer some improvement over uniform sampling in most cases. However, it is interesting to note that in all cases, better performance may be achieved by using uniform sampling with $2N$ sample points than nonuniform sampling with N sample points.

There are two-dimensional analogs for these results [15], but both the theory and its application are much more complicated. In one dimension, a block is entirely specified by its two endpoints. In two dimensions, specification of a block is much more difficult. The boundaries between blocks may be arbitrary one-dimensional contours in two-dimensional space. The representation of even discretized versions of the contours requires special data structures. More importantly, the implementation of any algorithm to compute these contours would require an enormous amount of computational effort, even after several simplifying assumptions have been made.

The best that can realistically be done, then, is to do non-uniform sampling

in one spatial dimension and normal uniform sampling in the other. Because the horizontal dimension of the television picture is larger than the vertical and because most of the motion in a typical television sequence is in the horizontal direction, it is likely that more edges will be crossed in scanning a typical horizontal line than in scanning a typical vertical line. Thus, it makes more sense to do the non-uniform sampling in the horizontal direction. Using this technique, one would expect to get a gain in performance similar to that achieved in the previous one-dimensional example.

This basic idea may be improved upon by performing directional interpolation between horizontal lines. The basic system described above will perform ideally only if all of the edges are oriented in the vertical direction (so that the luminance gradient will be in the horizontal direction), but as the edges become closer to horizontal in orientation, the system will perform no better than one which uses standard uniform sampling. Performance would be significantly improved if it were possible to interpolate along lines oriented in the same direction as the edges in the luminance. In this manner, sharp edges could be preserved all directions except those that are very nearly horizontal. Algorithms for locating edges in images and determining their orientation are well-documented in the literature.

Regardless of the sophistication of the inter-line interpolation technique, however, non-uniform sampling incurs a significant cost in the amount of additional side information that must be sent. Since there are roughly 426 pels per line, 9 bits would need needed to transmit the location of each sample if absolute addressing were used. Of course, more sophisticated coding schemes could be used to reduce this figure somewhat. For instance, one could simply code the lengths of the gaps between sample points. Depending on the number of samples per line, such a coding scheme would probably require 5–6 bits per sample. If entropy coding or non-uniform quantization were used, it is conceivable that adequate performance could be achieved by coding the positions to as few as 3 bits per sample.

Even at this optimistically low bit rate, transmitting the sample positions would require at least the same number of bits as are used to transmit the values of the low-frequency samples. Thus, in order to keep the required amount of digital bandwidth constant, the number of samples would have to be reduced by a factor of 2. Clearly, such a scheme would wasteful, since, as the above experiment demonstrates, better performance could be obtained simply by doubling the number of luminance samples and using uniform sampling. The digital bandwidth problem would be further compounded if the directional interpolation scheme were used. Extra bits would be required to transmit the interpolation directions, leaving fewer bits available for the transmission of the low-frequency sample values.

Ultimately, non-uniform sampling was abandoned as a useful alternative for two reasons. First, as mentioned above, it incurred a large additional cost in required bit rate. More importantly, however, non-uniform sampling did not achieve a perceptually significant increase in performance. There was some increase in the amount of noise reduction near sharp edges; however, objectionable amounts of noise still remained. The major effect of non-uniform sampling seemed to be to limit the extent of the noisy region surrounding the edges. In those regions very close to the edges where noise remained, there was no noticeable decrease in its magnitude, so it was still very much visible.

4.7 Video

After the parameters of the system had been optimized for still images, the system was tried on several monochrome video sequences. Each original sequence was encoded, using the scheme described above, on a frame by frame basis. This encoding process produced a receiver-compatible video sequence and a stream of digital side information. The receiver-compatible sequence was then subjected to various channel degradations. Finally, the advanced-receiver decoding process was applied to

the noisy receiver-compatible sequence, producing the recovered video sequence. As before, it was assumed that the stream of digital side information was transmitted without error.

While the quantitative results were almost identical to those for the still images, the subjective results were somewhat disappointing. They seemed to indicate that processing in three dimensions is fundamentally different from processing in two dimensions. The major problem was that the underlying block structure had become visible in the encoded receiver-compatible sequence. Because the location of the L-blocks and K-blocks did not change from one frame to the next, it apparently was possible for the human visual system to "lock on" to the block boundaries and track them through time.

If one considers what happens to the receiver-compatible as a detailed region in the original sequence moves across stationary sample points, the reason for this blocking effect becomes clear. If the value of the original signal changes abruptly from one frame to the next, then corresponding changes in the low-frequency sample values would result. These fluctuations in the sample values would lead to fluctuations in the value of the low-frequency interpolant in the neighborhoods of the affected sample points, which in turn would lead to large-scale changes in the value of the high-frequency potion of the signal. Because the high-frequency portion of the signal is being amplified by the adaptation factors, these changes are extremely visible.

This blocking effect appeared to be most noticeable when a low-contrast edge moved across a sample point. Edges of this type appeared to move in a jerky manner, jumping from one sample point to the next, rather than in the smooth manner in which they moved in the original sequence. This effect is demonstrated in Figure 4.28. What is shown in this figure is the receiver-compatible signal as a function of time. Note how the values at all of the points in the trailing block appear to converge simultaneously, while all of the values in the leading block diverge

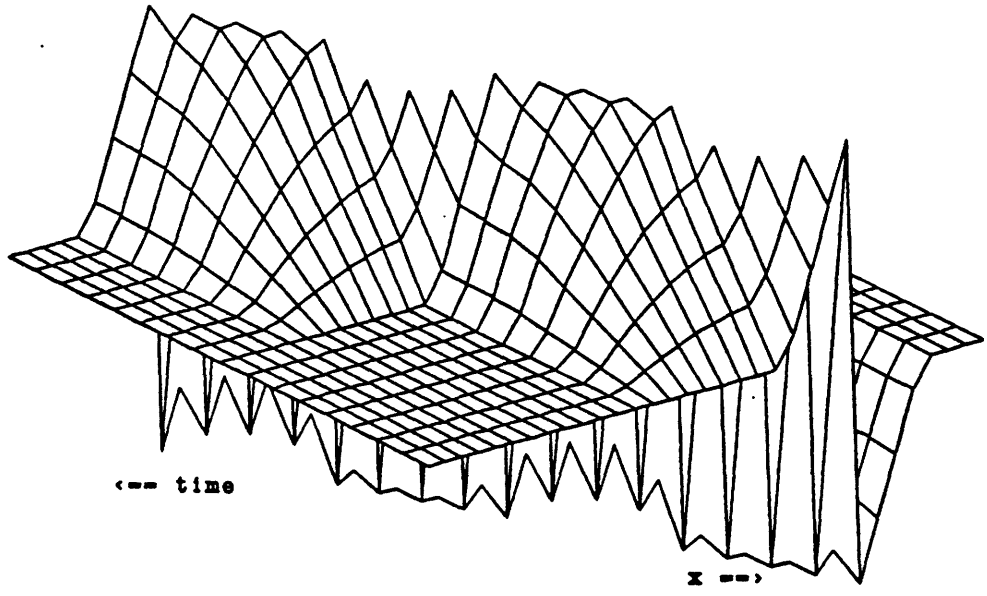

Figure 4.28: Evolution of the receiver-compatible signal as a low-contrast edge (in the original) moves across a sample point.

simultaneously.

In an attempt to remedy this problem, several different adjustments were made to the system. The first adjustment that was tried was to use smoother, approximately bandlimited interpolation filters instead of the triangular filters that correspond to bilinear interpolation. It was thought that some of the blocking effect might be due to the abrupt changes in slope that are present in the interpolant when bilinear interpolation is used. Thus, by using an interpolation filter with continuous first derivatives, the abrupt slope discontinuities in the interpolant might be eliminated, thereby reducing the blocking effects. The windowed ideal low-pass filters which were described earlier were used. Unfortunately, these smoother filters failed to produce the desired effect. As in the case of improving the performance of the system near sharp edges, it appeared that sample point position was more important than sample values or the interpolation scheme used to construct the interpolant from them.

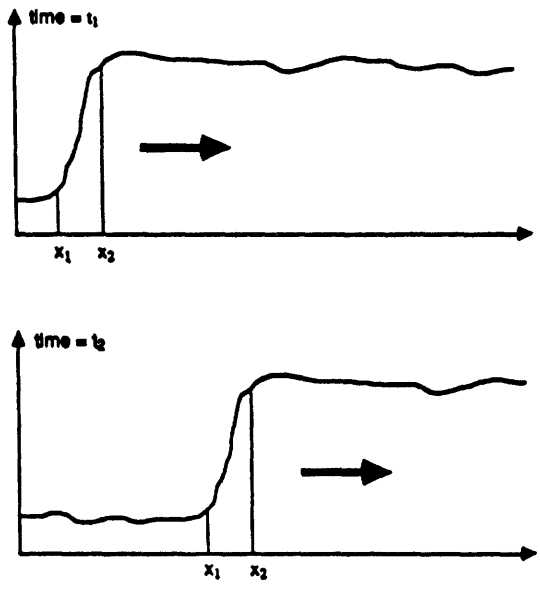

Figure 4.29: Tracking of edge from one frame to the next using nonuniform sampling.

When these smoother interpolation filters failed to produce any improvement, the next idea that was tried was nonuniform sampling. Even though nonuniform sampling was ruled out as an efficient method of sampling still images, it was thought that it might have some merit in the case of sequences. Since the nonuniform sampling algorithm places the sample points at the positions in the signal where the curvature is the greatest, one would expect that sample points would be placed at the beginning and end of any sharp transition, as in Figure 4.29. As the edge moves from one frame to the next, these sample points should move with it, as shown in Figure 4.29. Thus, since the number of occurrences of edges moving over sample points would be small, one would expect that nonuniform sampling would reduce the blocking effect by a large amount. Unfortunately, that was not the case. Nonuniform sampling did offer some improvement, but not enough to warrant the additional digital bandwidth it would consume.

It was finally decided that the dynamic range of the adaptation factors might

have to be limited even further. Instead of using the range [1,4], as with the still images, an adaptation factor range of [1,2] was tried. This reduction in the magnitude of the adaptation factors did hide much of the blocking effects in the receiver-compatible picture, but this improvement was obtained at the cost of roughly 6 dB of SNR in the pictures recovered on the advanced receivers.

This simple adjustment, however, is not really an acceptable way to solve the blocking problem. Limiting the maximum adaptation factor to 2 limits the gain in SNR on the advanced receivers to roughly 6 dB. Such an increase in SNR is certainly noticeable, but the expense of the additional hardware needed to achieve this modest improvement would make the system impractical.

It is possible, though, that this blocking problem may be overcome by some type of temporal processing. By comparing successive frames at the transmitter, it would be possible to limit the amount of change in either the high-frequency component or the adaptation factors (or both) from one frame to the next. Such processing would eliminate the blocking problem, but would result in a loss of performance in the regions containing abrupt luminance transitions. A better idea would be for transmitters to detect moving edges and somehow pass the positions, orientations, and velocities of these edges to the receivers. The receivers could then buffer several frames and use this extra side information to interpolate temporally. The difficulty with this approach, however, is that even though algorithms do exist for performing all of these functions [11], the additional storage and digital bandwidth they would require make them impractical at this point in time.

Chapter 5

Performance

Throughout the design phase, only a very limited set of channel degradations were used to evaluate the trade-offs involved in varying the parameters of the system. Therefore, once the system configuration and parameters had been set, several tests were run to evaluate the performance of the system with varying channel degradations. The parameters of the system that were used during this testing are shown in Table 5.1. The coding of the low-frequency samples was done by Netravali's DPCM technique in conjunction with mu-law quantization of the error residuals.

The three types of channel degradations that were simulated were random additive noise, multipath interference, and interchannel interference. To measure the performance of the system in the presence of these degradations, several monochrome still images were encoded into receiver-compatible format. The receiver-compatible encoded signals were subjected to each of these three channel degradations individually and to all three at once. To provide a standard for comparison, the corresponding NTSC-encoded signals were subjected to the same set of channel degradations. The subjective performance of the system in the presence of 20 dB random additive noise may be seen in Figure 5.1. Similarly, Figures 5.2 and 5.3 demonstrate the performance of the system in the presence of 20 dB multipath

L-blocks	Y	8×8
	I	8×8
	Q	8×8
K-blocks	Y	4×4
Bits of signal samples	Y	3
	I	3
	Q	3
Range of adaptation factors	Y	[1,4]
Bits of adaptation factor	Y	2
Low-frequency compression factor	Y	.875

Table 5.1: Parameters of the Receiver-Compatible Noise Reduction System.

and 20 dB interchannel interference, respectively. Figure 5.4 demonstrates the system's performance in the presence of all three types of degradations simultaneously. In all cases, one can see that the quality of the picture resulting from receiver-compatible encoding and decoding is far superior to that of the picture resulting from NTSC encoding and decoding. Furthermore, one can see that in most cases, the picture obtained by applying NTSC decoding to a degraded receiver-compatible signal has roughly the same quality as the degraded NTSC picture.

In typical environments, however, one would expect a SNR of quite a bit more than 20 dB. Thus, a second experiment was carried out to determine how the quality of both the pictures on the standard receiver (*i.e.* the encoded pictures) and the pictures on the advanced receiver (*i.e.* the decoded pictures) varies as a function of the noise level. The encoded receiver-compatible signal was subjected to noise levels ranging from 10 dB SNR to infinite SNR (no degradation at all). For subjective comparison, an NTSC-encoded signal was subjected to the same set of noise levels. These pictures are shown in Figure 5.5. A graph of the MSE between the original and the degraded receiver-compatible signal is shown in Figure 5.6(a). A graph of

(a) (b)

(c) (d)

Figure 5.1: Performance of receiver-compatible system in the presence of additive random noise (20 dB SNR): (a) Original picture. (b) NTSC encoding and decoding. (c) Receiver-compatible encoding and NTSC decoding. (d) Receiver-compatible encoding and decoding.

(a) (b)

(c) (d)

Figure 5.2: Performance of receiver-compatible system in the presence of multipath (20 dB SNR): (a) Original picture. (b) NTSC encoding and decoding. (c) Receiver-compatible encoding and NTSC decoding. (d) Receiver-compatible encoding and decoding.

(a) (b)

(c) (d)

Figure 5.3: Performance of receiver-compatible system in the presence of interchannel interference (20 dB SNR): (a) Original picture. (b) NTSC encoding and decoding. (c) Receiver-compatible encoding and NTSC decoding. (d) Receiver-compatible encoding and decoding.

(a) (b)

(c) (d)

Figure 5.4: Performance of receiver-compatible system in the presence of additive random noise (20 dB SNR), multipath (20 dB SNR), and interchannel interference (20 dB SNR): (a) Original picture. (b) NTSC encoding and decoding. (c) Receiver-compatible encoding and NTSC decoding. (d) Receiver-compatible encoding and decoding.

MSE between the original and the picture recovered from the degraded encoded picture is shown in Figure 5.6(b). There is some question as to the value of the MSE figure as a means of comparison, since as Schreiber [19] and Netravali [13] point out, it is really the noise visibility, rather than the absolute noise level that matters. However, these graphs seem to correspond fairly well with the observed picture quality of the images under consideration.

It is interesting to note that as the quality of the channel decreases, the impairments due to the receiver-compatible encoding become less objectionable. In other words, the process of going from original picture to the receiver-compatible encoded one appears to add a constant amount of degradation, which becomes insignificant as the level of the actual channel degradations increases. Also, as the quality of the channel decreases, the difference in subjective performance between the NTSC system and the receiver-compatible system becomes greater, as the noise reduction of the receiver-compatible system offers increasingly more subjective improvement.

A similar set of tests was carried out on color still images. The MSE figures followed a pattern similar to the one observed for the monochrome images, but there were some subjective differences. Under close scrutiny, one could find some loss of color resolution in the pictures received by the advanced receiver. However, this loss of resolution was confined primarily to those regions in which a sharp transition between extremely saturated colors occurred. The addition of the slowly varying color components did, however, have at least one beneficial effect. The color tended to hide some of the contrast in the luminance of the encoded pictures, thereby increasing their subjective compatibility.

Tests were also performed on several video sequences to determine the subjective performance level of the system. The quality of the encoded pictures was given special consideration, to make sure that the blocking effect mentioned earlier did not produce any objectionable effects. After running several trials with varying noise levels and with the value of the K_{max} parameter set to 4 (as in the previous

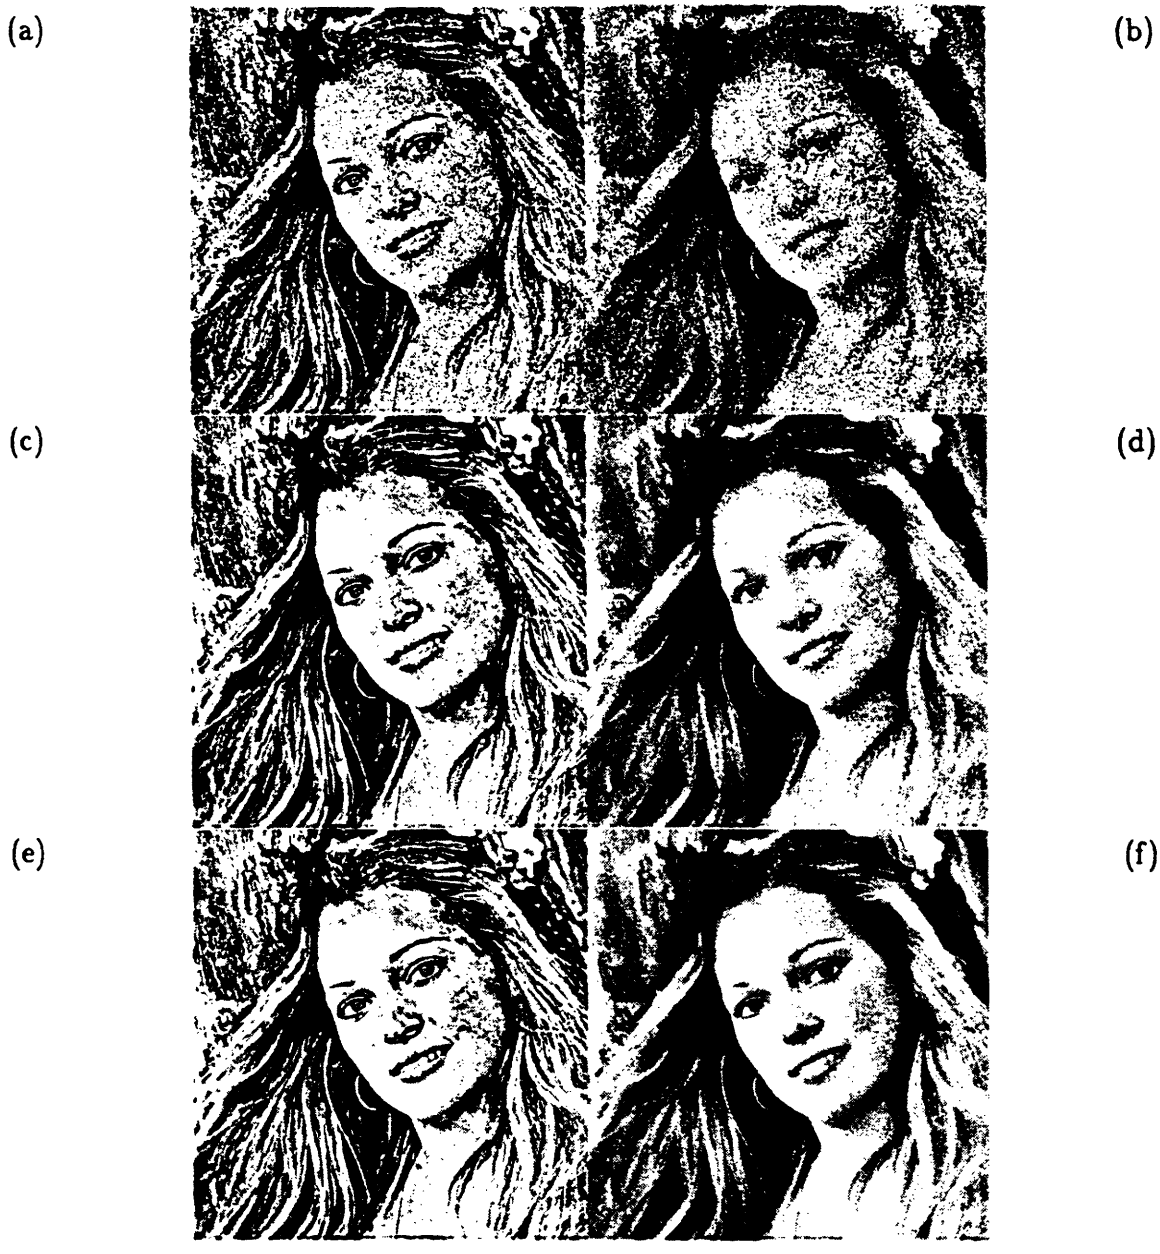

Figure 5.5: Performance of the receiver-compatible system in the presence of varying levels of additive random noise. 15 dB AWGN: (a) degraded receiver-compatible-encoded picture (b) degraded NTSC-encoded picture. 25 dB AWGN: (c) degraded receiver-compatible-encoded picture, (d) degraded NTSC-encoded picture. 35 dB AWGN: (e) degraded receiver-compatible-encoded picture, (f) degraded NTSC-encoded picture.

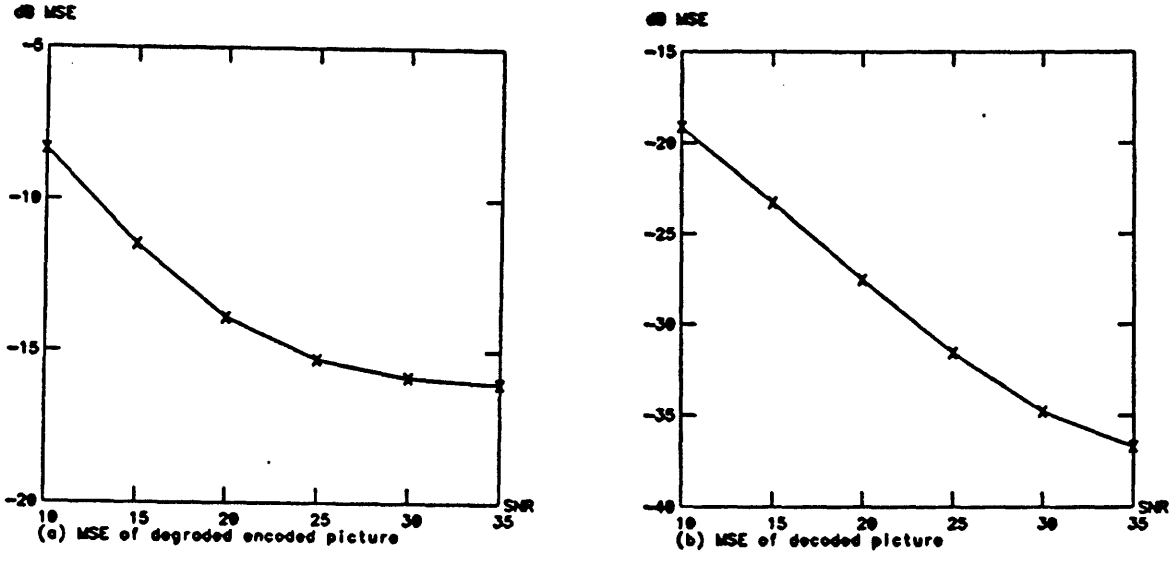

Figure 5.6: Graph of MSE for varying levels of additive random noise: (a) MSE between original picture and receiver-compatible encoded picture. (b) MSE between original and receiver-compatible decoded picture.

trials), the quality of the encoded picture sequences was found to be acceptable, but a few of the blocking artifacts discussed earlier were noticeable. Reducing K_{max} to between 2 and 3 resulted in noticeably better encoded sequences; however, it also caused a noticeable loss in noise reduction in the decoded sequences.

The one underlying fact that became evident during the course of these trials was that there is no good way to fix the system parameters to balance the trade-off between compatibility with standard receivers and noise reduction on advanced receivers. As mentioned earlier, the parameter K_{max} may be varied over a limited range to trade compatibility for noise reduction, or vice-versa, but the results of these trials indicate that there is no fixed set of parameters which achieves a truly significant amount of noise reduction while maintaining complete subjective compatibility.

While this result is somewhat discouraging, it does not imply that there is no use for a system of this type. Since it is possible to vary the parameters of the system to emphasize either compatibility or noise reduction, the system may be tailored to the particular application. In low SNR areas, it would be possible to allow more distortion in the encoded pictures, since the low quality of the channel would mask the distortion. In high SNR areas, compatibility could be emphasized by limiting this distortion. A typical terrestrial television channel usually has an SNR of 35-40 dB. Thus, one would expect that in most cases compatiblity would be emphasized over noise reduction. However, as standard recievers are phased out, it would be possible to increase the level of noise reduction even in these high SNR areas.

Chapter 6

Conclusions

An NTSC receiver-compatible system for the reduction of channel degradations has been presented. The system has been shown to work very well on still images, maintaining a high level of compatibility while offering a theortical improvement in SNR of up to 12 dB (based on $K_{max} = 4$). In fact, SNR improvements of slightly more than 12 dB were observed during several of the performance tests. Video sequences caused some problems for the receiver-compatible noise reduction system, as blocking effects became much more noticeable, especially in the presence of moving of regions of sharp detail.

What was found is that it is very difficult to achieve compatibility with existing NTSC receivers and high levels of noise reduction on advanced receivers simultaneously. Indeed, the two are almost inversely related to one another, since the primary factor governing both receiver compatibility and the level of noise reduction is the dynamic range of the adaptation factors. As the dynamic range is restricted, compatibility is increased at the cost of reduced noise reduction; as the dynamic range is increased, noise reduction is improved at the cost of impaired compatibility. Of course, if techniques were developed to improve the fit of the low-frequency interpolant to the original signal, this trade-off would be eased somewhat. Better interpolants result in smaller amplitudes in the high-frequency portion of the signal,

which in turn allow for larger adaptation factors for the same level of compatibility. This idea is discussed below.

One of the attractive features of the receiver-compatible noise reduction system, though, is that even if the compatibility with existing NTSC receivers and high performance on advanced receivers cannot be achieved simultaneously, it is possible to make a smooth transition from one end of the spectrum to the other by slowly varying the parameters of the system. In particular, one could initially implement the system with a very small dynamic range for the adaptation factors and a very low level of low-frequency compression. Under these conditions, compatibility with existing receivers would be excellent, while there would be only a small, but appreciable, amount of noise reduction on the advanced receivers. Then, as standard NTSC receivers are phased out in favor of advanced receivers, the parameters could be smoothly adjusted to increase the level of noise reduction possible with advanced receivers at the expense of impairing the compatibility of the receiver-compatible signal with existing NTSC receivers. The dynamic range of the adaptation factors and the level of low-frequency compression would both gradually be increased. Eventually, it would be possible to make the low-frequency compression infinite, so that no low-frequency information at all would be transmitted in the analog signal. This format would allow for extremely large adaptation factors, thereby making it possible to eliminate even high-power interferences from impairing the quality of the received picture.

6.1 Further Improvement

There is still some room for improvement of this system. As mentioned earlier, the key to making any adaptive modulation system work well is to make the portion of the signal that is to be adaptively modulated as small as possible. In the case of the receiver-compatible noise reduction system, this goal may be achieved by making

the "low-frequency" portion of the signal fit the original signal as closely as possible, thereby reducing the amplitude of the "high-frequency" portion of the signal. As the high-frequency signal decreases in amplitude, the level of signal distortion my be decreased simply by holding the adaptation factors constant. Conversely, it would also be possible to increase the adaptation factors, thereby increasing the level of noise reduction, while keeping the level of signal distortion constant.

There are many ways in which the fit of the low-frequency interpolant might be improved. It is possible that some additional digital bandwidth might be made available either through an improved digital modulation scheme or through more sophisticated coding techniques (perhaps exploiting temporal correlation). These additional bits could be used to reduce the block sizes or increase the number of quantization levels of the low-frequency samples, both of which would improve the fit of the interpolant to the original.

During the course of this research, several different techniques were investigated in an attempt to get a better fit. Least-squares optimization and linear programming were used in an attempt to obtain a better fit by selecting the "optimal" sample values, rather than computing the sample values simply as the outputs of a low-pass decimation filter. When these techniques failed to produce significant improvements, the traditional idea of the "low-frequency" portion of the signal was abandoned and nonuniform sampling was tried. However, that technique also failed to significantly improve the closeness of the fit between the interpolant and the original signal.

Despite these discouraging results, it is likely that there does exist some technique which will achieve a sufficiently close fit to the original so that performance is significantly improved. What is needed is a way to represent the original luminance signal that would minimize the absolute errors between the original signal and its representation. The problem of finding such a representation may be stated as follows: Given the characteristics of the luminance component, such as its dynamic

range and bandwidth, an appropriate error criterion, and limitations on computational complexity and storage, what is the best way to represent the original signal with a given amount of digital information? Determining an appropriate error criterion is a problem in itself. Although the least-squares criterion was used throughout most of this research, a minimax error criterion probably would have been better suited to the problem. A minimax error criterion would tend to favor schemes whose performance does not degrade significantly around sharp discontinuities. It is possible, however unlikely, that low-pass filtering and sampling at the Nyquist rate may be the best way to proceed. To determine an optimal or near-optimal solution to this problem, however, is an extremely difficult problem in rate distortion theory, well beyond the scope of this work.

The receiver-compatible noise reduction system might also be improved by introducing temporal processing. Temporal processing was excluded from consideration during the course of this research because of the increased storage requirement it would entail, but it is possible that the additional expense might be worth the improvement to be gained.

For instance, the high level of temporal redundancy in a typical video sequence could be exploited to get some coding gain. A three-dimensional prediction scheme could be used to estimate the low-frequency sample values. One would expect that such a predictor would be much better than the two-dimensional spatial predictor that is currently being used, thus reducing the magnitude of the error residuals and therefore the number of required quantization levels. The bits freed by a more efficient coding scheme could be used elsewhere to improve the performance of the system.

More importantly, however, it is possible that temporal processing could be used to eliminate the disturbing blocking problem that became apparent when the receiver-compatible encoding scheme was applied to video sequences. The low- and high-frequency components and adaptation factors of successive frames could

be compared at the transmitter to detect and correct potential problems, such as a sudden shift in the sign of the high-frequency component. It is possible that processing of this type might not even require any extra storage in the advanced receivers, making it extremely practical.

Bibliography

[1] M. J. J. C. Annegarn, J. P. Arragon, G. de Haan, J. H. C. van Heuven, and R. N. Jackson. HD-MAC: a step forward in the evolution of television technology. *Philips Technical Review*, 43(8):197–212, August 1987.

[2] Carl de Boor. *A Practical Guide to Splines.* Springer-Verlag, New York, 1978.

[3] Yves Faroudja and Joseph Rozien. Improving NTSC to achieve near-RGB performance. *SMPTE Journal*, 750–761, August 1987.

[4] Takahiko Fukinuki and Yasuhiro Hirano. Extended definition TV with existing standards. *IEEE Transactions on Communications*, COM-32(8):948–953, August 1984.

[5] W. E. Glenn, K. G. Glenn, J. Marcinka, R. Dhein, and I. C. Abrahams. Reduced bandwidth requirements for compatible high definition television transmission. In *National Association of Broadcasters Conference Proceedings*, 1984.

[6] William E. Glenn and Karen G. Glenn. High definition television compatible transmission system. *IEEE Transactions on Broadcasting*, BC-33(4):107–115, December 1987.

[7] Michael A. Isnardi, Terrence R. Smith, and Barbara J. Roeder. Decoding issues in the ACTV system. *IEEE Transactions on Consumer Electronics*, 34(1):111–120, February 1988.

[8] Jae S. Lim. *Two-Dimensional Signal and Image Processing.* Prentice-Hall, Englewood Cliffs, New Jersey, 1990.

[9] S. P. Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 129–136, March 1982.

[10] Wayne Luplow. *Zenith's Spectrum Compatible High Definition Television System.* Executive Summary, Zenith Electronics Corporation, Glenview, Illinois, September 1988.

[11] Dennis M. Martinez. *Model-Based Motion Estimation and its Application to Restoration and Interpolation of Motion Pictures.* PhD thesis, Massachusetts Institute of Technology, 1987.

[12] J. Max. Quantizing for mimimum distortion. *IRE Transactions on Information Theory*, 7–12, March 1960.

[13] Arun C. Netravali. Interpolative picture coding using a subjective criterion. *IEEE Transactions on Communications*, COM-25(5):503–508, May 1977.

[14] M. D. Paez and T. H. Glisson. Minimum mean-squared-error quantization in speech PCM and DPCM systems. *IEEE Transactions on Communications*, COM-20(4):225–230, April 1972.

[15] Theodosios Pavlidis. Optimal piecewise polynomial L_2 approximation of functions of one and two variables. *IEEE Transactions on Computers*, 98–102, January 1975.

[16] Tamar Peli and Jae S. Lim. Adaptive filtering for image enhancement. *Optical Engineering*, 21(1):108–112, January/February 1982.

[17] William F. Schreiber. Improved television systems: NTSC and beyond. *SMPTE Journal*, 66(8):717–725, August 1987.

[18] William F. Schreiber. Psychophysics and the improvement of television image quality. *SMPTE Journal*, 93(8):717–725, August 1984.

[19] William F. Schreiber and Robert R. Buckley. A two-channel picture coding system: II—Adaptive companding and color coding. *IEEE Transactions on Communications*, COM-29(12):1849–1858, December 1981.

[20] William F. Schreiber and Andrew B. Lippman. *Reliable EDTV/HDTV Transmission in Low-Quality Analog Channels.* Technical Report ATRP-T-96R, The Media Laboratory, Massachusetts Institute of Technology, October 1988.

[21] William F. Schreiber and Andrew B. Lippman. Single-channel high definition television systems, compatible and noncompatible. In *SPIE Symposium*, November 1988.

[22] William F. Schreiber and Andrew B. Lippman. *Single-Channel High-Definition Television Systems, Compatible and Noncompatible.* Technical Report ATRP-T-95, The Media Laboratory, Massachusetts Institute of Technology, 1988.

[23] Jong-Soo Seo and Kamilio Feher. Performance of 16-state SQAM in a nonlinearly amplified multichannel interference environment. *IEEE Transactions on Communications*, 36(11):1263–1267, November 1988.

[24] Yoshio Yasumoto, Sadashi Kageyama, Syuhji Inouye, Hideyo Uwabata, and Yoshio Abe. A wide aspect ratio television system with full NTSC compatiblity. *IEEE Transactions on Consumer Electronics*, 34(1):121–127, February 1988.