The Digital Wetgate: A Third-Generation Noise Reducer

By Gerhard Wischermann

This paper describes the development of noise-reduction technology and addresses its applications. First-generation noise reducers using linear filtering techniques such as recursive or transversal filters were designed for removing film grain noise or tape noise of analog VTRs. With the introduction of nonlinear median filtering techniques it became possible to conceal impulse-type noise such as film dust, randomly dispersed film scratches, or FM sparkles.

A third-generation noise reducer now has been developed by BTS. By the use of advanced techniques for error detection and motion adaption it is now possible to conceal large film dirt or fixed vertical film scratches — so-called tramline scratches. A wide range of new features especially designed for film reproduction, such as automatic pulldown and scene-change detection, provide great concealment power in the digital wetgate.

About ten years ago, the first generation of studio noise reducers came to the market. These were designed for removing uniformly spread noise such as film grain noise or tape noise in analog VTRs, using either recursive or transversal filters.

The next step within the development of noise-reduction technology was the introduction of nonlinear filtering methods using median filters. Median filters are able to handle impulse-type noise and can therefore conceal bit errors or film dust. Although median filters were well known and described in many technical articles, they were not used in studio equipment for a long time. The first median noise reducer, the MNR9, was introduced in 1991. Due to a three-dimensional, star-shaped filtering window there was no detail loss for still pictures, but as it did not contain motion adaption, some median artifacts were introduced on critical material. However, the median noise reducer was a breakthrough in noise-reduction technology and started many engineers thinking about how to improve the median principle.

Third-generation noise reducers currently available use very sophisticated methods for motion adaption and error correction to avoid the typical median artifacts. Depending on the nature of the distortion to be removed, special types of self-adaptive median filters are applied to the video signal, so that even large image areas can be restored. Film dirt, randomly dispersed film scratches or dropouts are no longer a problem.

Linear Filtering Methods

The first part of this paper discusses linear filtering methods for removing unwanted parts of the video signal. Some of the typical applications for linear filtering are:

- Camera noise under low-light conditions
- Grain noise in film material
- Tape noise from analog VTRs
- Cross color, cross luminance
- Chroma phase errors (Simple PAL)

Starting at the source, the noise is first introduced by the video camera. The blue channel is very sensitive to noise, which might later lead to chroma-key problems. In the case of film, noise is usually caused by the film grain. If a 16mm film is shot under low-light conditions, grain noise becomes very annoying. During production, analog VTRs will add some tape noise to the video signal, especially after multiple generations of recording. After PAL or NTSC encoding, cross color and cross luminance will further degrade the signal quality. For Simple-PAL decoding, hue errors

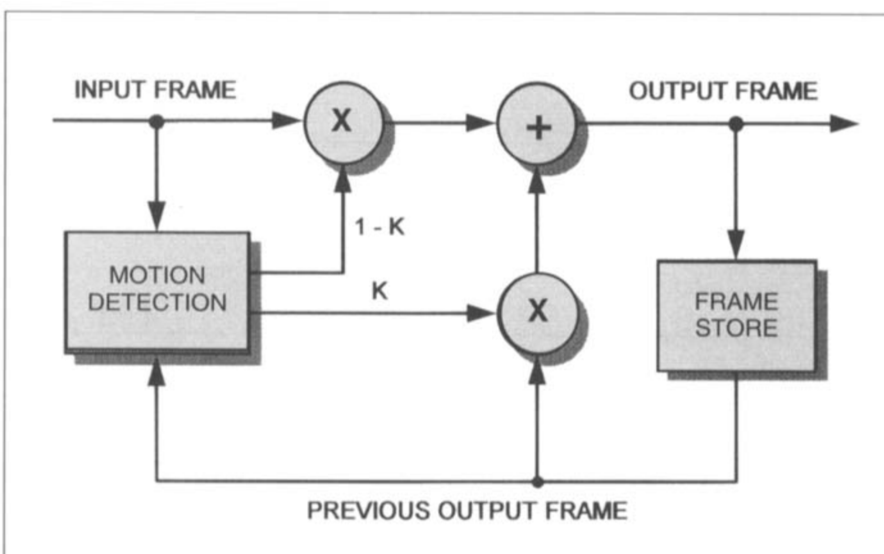

Figure 1. Adaptive recursive filter.

Presented at the 136th SMPTE Technical Conference in Los Angeles (paper no. 136-21) on October 13, 1994. Gerhard Wischermann is with BTS Broadcast Television Systems GmbH, Weiterstadt, Germany D-64331.

might occur if the burst-to-chroma phase relation has been changed.

Temporal Recursive Filter

Figure 1 shows one way for removing or at least reducing these problems — the temporal recursive filter.

How does this work? The easiest way to understand the operation of the recursive filter is to regard it as an adder that accumulates successive fields or frames of video. By this averaging process the noisy part of the video signal will be suppressed due to its statistical nature. Looking at the block diagram of the recursive filter, the next output frame will be a mixture of the new frame to be processed and the last output frame already being noise reduced. The mixing ratio between the feedback and the direct portion is controlled by a factor called K. The higher the amount of feedback, the more the signal gets noise reduced after some time. To avoid motion blur, however, the K factor has to be set to lower values by a motion detector in those parts where the picture has changed more than a certain amount.

Recursive Noise Reduction

Before addressing motion detection, we will discuss time constants. Figure 2 shows how the recursive noise reduction builds up frame by frame if there is no motion detected.

If 12 dB of noise reduction is selected, it takes more than 20 frames to actually get 12 dB! There is a large time lag, which every operator should keep in mind when using high decibel settings. For normal operation it is better to select only 4 dB or 6 dB so as to get a short time constant. High decibel settings only make sense for still pictures or little movement.

Motion Detection

The parameters of motion detection include:

- Evaluation of the interframe differences
- Distinguishing between noise and slight movement
- Detection of a scene change
- Detection of a fade-in or fade-out
- Choice of a proper K-factor characteristic

Referring to the block diagram of the recursive filter in Fig. 1, motion detection is derived from the difference between the input frame and the previous output frame. But these interframe differences really include a lot more than just motion. Because the output of the recursive filter is noise reduced and the input is not, noise is also present in the difference signal. That is why first there must be some noise reduction in the motion signal itself before it is used to control the recursive filter. This can be done by horizontal and vertical low-pass filtering. To get rid of the DC component arising after noise filtering, an adjustable threshold is used in all noise reducers. The clip level has to be set by the operator or, what is better and more reliable, by an automatic measurement of the noise floor.

After noise reduction, the frame-to-frame difference signal can be used to control the recursive filter via a proper K-factor characteristic. This characteristic has to be chosen very carefully. Noise reduction should not be cut for small movement without introducing smearing effects or motion blur on the other hand. Particularly "afterimages," in conjunction with scene changes, must be strictly avoided.

There is one more thing a good motion detector should include, and that is fading detection. During a fade-in or fade-out, or a fade-to-black transition, the motion detector will normally be triggered by a strong DC component of the motion signal. This will cause noisy pictures during the transition, even if there is no movement. The new BTS noise reducer includes both fading detection and scene change detection.

The scene change detector flushes the store of the recursive filter in order to suppress afterimages of the old scene. Furthermore, the scene change signal is made available externally for controlling other equipment, such as editors or color correctors. As color errors normally occur scene by scene, there is no need for time-consuming moving of the film in order to find the cutting points. The scene change is detected automatically and is frame accurate, so the operator can concentrate on his creative work.

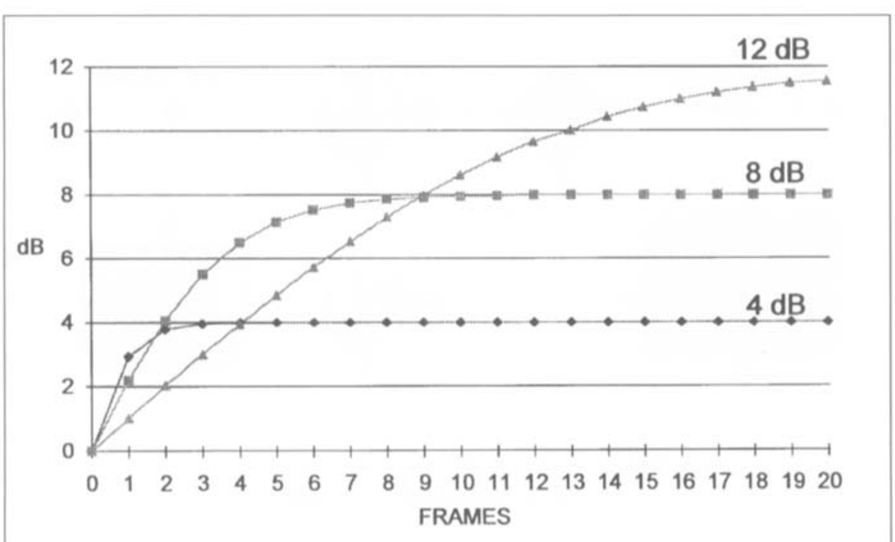

Figure 2. Recursive noise reduction vs. time.

Transversal Filter

Figure 3 outlines the block diagram of a transversal filter that can alternatively be used for the reduction of uniformly spread noise. This configuration requires more hardware in terms of frame stores because it uses three frames or more in parallel for averaging. Also, motion detection is more complex because forward and backward motion, relative to the center frame, has to be calculated. However, at a scene change the transversal filter is superior to the simple recursive filter with only backward motion detection.

This will be illustrated in Fig. 4, where a comparison is shown between

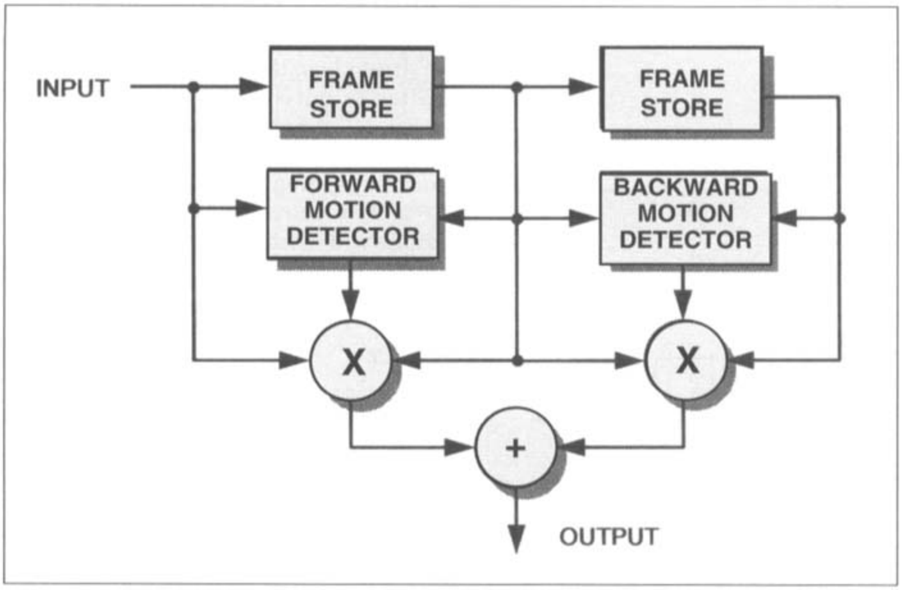

Figure 3. Adaptive three-frame transversal filter.

a 5-tap transversal filter and a recursive filter, both having 7 dB of noise reduction for still pictures. It is clear that the recursive filter has to restart noise reduction from zero after the scene change. The first picture of the new scene is not noise reduced at all. Much better is the dynamic performance of the transversal filter. As it has forward and backward motion detection there are at least three frames of either the old or the new scene that can be used for averaging. Therefore, transversal filtering leads to continuous and symmetrical noise reduction, even at a scene cut.

Another principal advantage of the transversal filter is that it cannot introduce a large motion smear over many frames like the recursive filter. However, for still pictures or little movement, the recursive filter is an excellent companion to the transversal filter, because it has no limitation in the degree of noise reduction.

Second-Generation Noise Reducers

This part of the paper will discuss second-generation noise reducers using nonlinear filtering techniques. Some "video nasties," which cannot be removed by linear filtering, are:

- Tape dropouts and bit errors
- FM sparkles from satellite links
- Film dust
- Film dirt
- Film scratches

All these examples have one thing in common: the video signal is distorted by high-energy pulses in isolated parts of the picture. Therefore, filtering this kind of errors cannot be done by averaging the bad pixels, but by replacing them. This is what a median filter does.

Median Filters

The function of a median filter can be characterized as follows:

- A median filter is a switch, selecting an output pixel within a group of input pixels.
- The group of input pixels to be filter-fed is the so-called filtering window or cluster.
- The median filter does not calculate the average value like the transversal filter.
- It selects the input pixel with the median amplitude value within the window.
- The median filter drops bad pixels and repeats good pixels (hopefully).
- It suppresses minimum or maximum values within the filtering window.

Simple Median Filter

An example of a simple median filter is shown in Fig. 5. The filtering window consists of three horizontal pixels, if it is assumed the signal to be filtered is a video signal. As the median filter suppresses local minimum and maximum amplitudes, the two isolated pulses shown in the diagram, which might be bit errors, are concealed at the output. On the other hand, the edges are not affected by the median filter, so there is no loss of sharpness during transitions. The two pixel wide pulse at the end of the sequence is not removed (it could be removed with a five-tap median filter). This leads to the following rule applying for all median filters: If N (odd) is the number of pixels to be

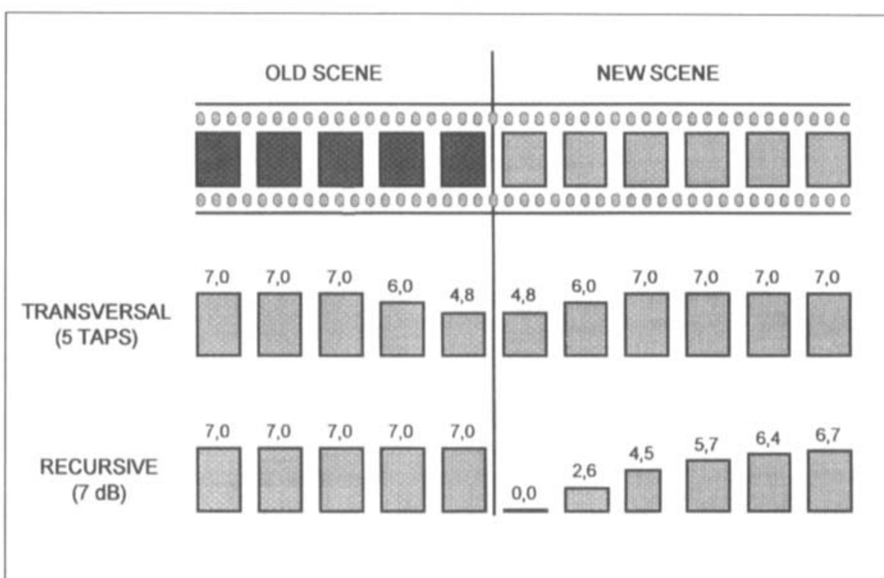

Figure 4. Noise reduction at a scene change.

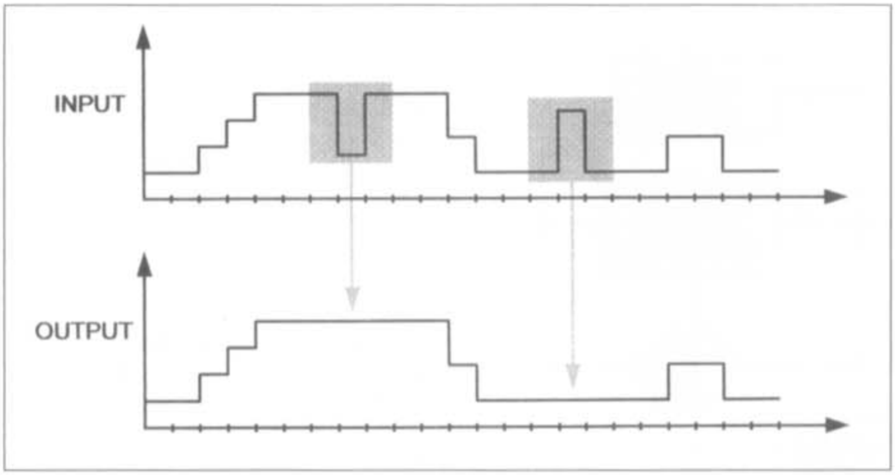

Figure 5. Effect of median filtering.

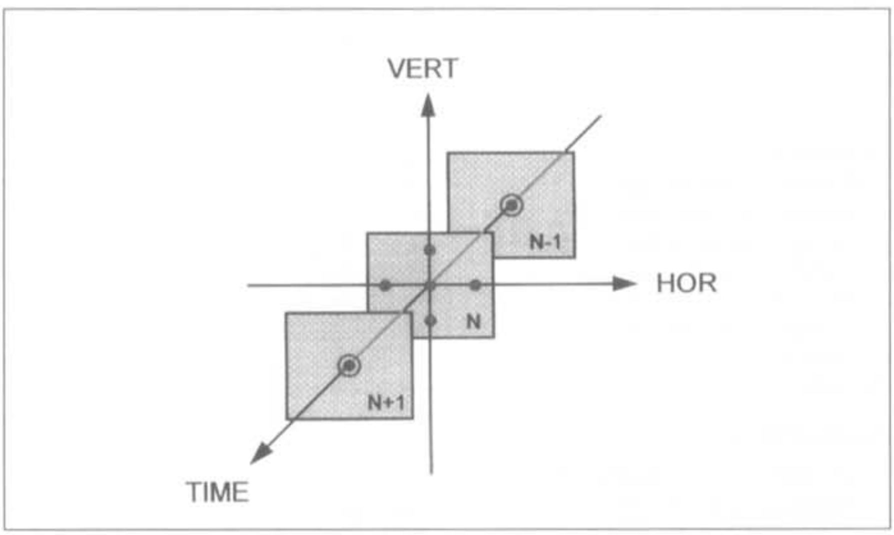

Figure 6. Three-dimensional median filter.

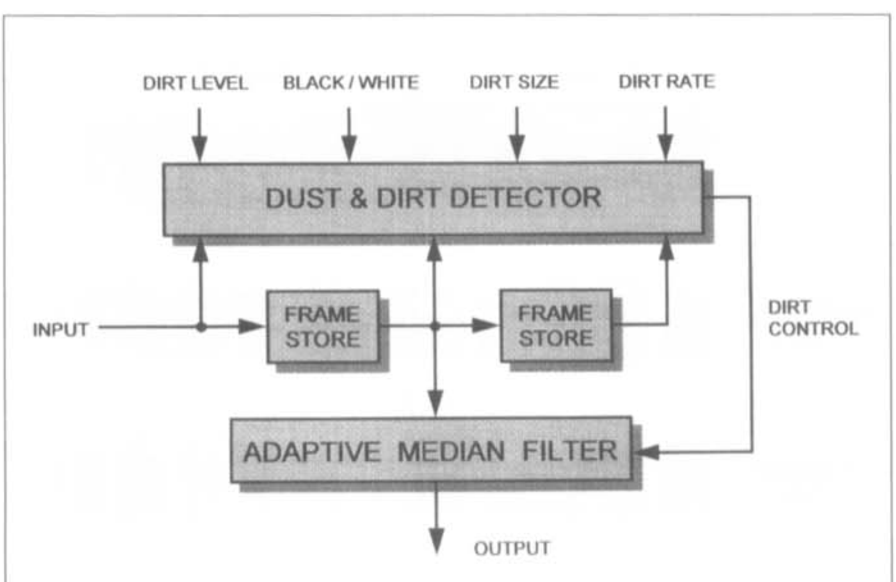

Figure 7. Dust and dirt concealment.

filtered, then (N-1) bad pixels can be concealed.

It is clear that this simple median filter is not suitable for broadcast applications, because a video signal contains a lot of local minimum and maximum amplitudes that are not necessarily bit errors. Better results can be achieved with two-dimensional or three-dimensional median filters.

Three-Dimensional Median Filter

Figure 6 shows the three-dimensional median filter, which was used for the first time in the MNR9 noise reducer. In the MNR9, the filtering window consisted of a total of nine pixels, five spatial pixels within the center frame, and four pixels on the temporal axis coming from the preceding and subsequent frames. With this nine-tap filter configuration, up to four pixel wide errors could be concealed. This was quite sufficient for removing a lot of film dust. In principle, this three-dimensional, star-shaped window did not lead to any loss of spatial resolution for still pictures, because the center pixel to be filtered repeated four times on the temporal axes. This is an important rule to remember.

For moving pictures with low detail the results still were very good, because in this case the center pixel repeats four times on the horizontal and vertical axes. However, for critical material such as scrolling titles or fast camera pans, the median filter had to be switched off, because no motion detector was built in at this time.

The Digital Wetgate

This device was specially designed for film applications, including concealment of film dust and dirt and removal of randomly dispersed scratches (not from the film, but from the video). Concealing fixed vertical film scratches is a unique task, because they are nonstatistical, occurring for several frames at the same place. By means of an additional digital filter, now for the first time even these tramline scratches can be suppressed. Therefore, the digital wetgate contains two filters: the dirt filter for statistical errors, and the scratch filter for nonstatistical errors.

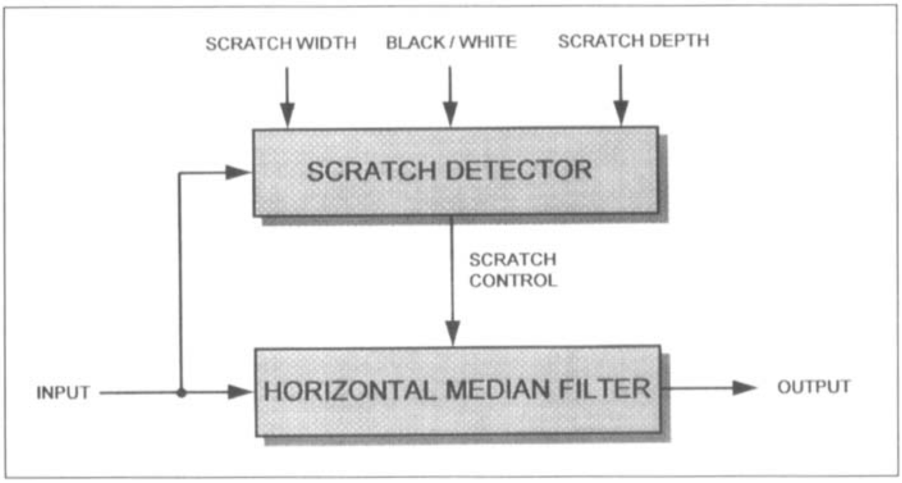

Figure 8. Tramline scratch concealment.

Dirt Filter

Figure 7 is a block diagram showing the dirt filter implemented in the third-generation MNR 11. The filter consists of two parts, one for detection and a second for concealment. The concealment is normally switched off and will only be activated in those parts of the picture where the detector indicates the presence of dirt. Detection can be controlled within seven levels of sensitivity, according to the dirt contrast. It is adaptable to positive or negative film, as there are either black or white spots to be removed. Of course, black and white detection at the same time is no problem. In most applications the operator wants to adapt the concealment to the size and the shape of the error. This is also no problem — eight different settings are available, optimized for individual treatment of film dust, film dirt, random scratches, FM glitches, or dropouts. Even an "Extra Large" setting for unlimited error size is available.

If no special care is taken, concealment errors may occur if the picture content has the same statistical nature as the errors. In that case the noise reducer would act, for example, like a "rain reducer." For minimizing fault triggering, the dirt filter is adaptable to the error rate. To prevent concealment errors during fast camera pans, it is additionally controlled by a global motion detector.

Tramline Scratch Filter

For a complete clean-up of dirty film, the second filter of the digital wetgate is required. The block diagram of the tramline scratch filter (Fig. 8) looks similar to that of the dirt filter. Again it consists of two sections: the detecting part and the filtering part. The filtering part is quite obvious — just a horizontal median filter is needed to do the job. What is more complex is controlling the filter in the right way. For proper scratch detection, at least one frame has to be analyzed in respect to characteristics applicable to tramline scratches. But this does not have to be done by the operator. The only thing he has to do is to set the detector to black, white, or both; select the maximum scratch width to be concealed; and adjust the sensitivity according to the scratch depth.

Performance

The question will arise about the performance of the digital wetgate compared to that of conventional wetgate scanning of film. The first difference, which is very obvious, is that the digital wetgate can be applied at any time, not only during the transfer to video. The main advantage, however, is the concealment power of this digital solution if the film is seriously damaged. Chemical treatment of film material can only be successful as long as scratches are not too deep. If the emulsion layer is damaged, there is no way. And last but not least, digital signal processing is not toxic and has no unpleasant odor.

Main Features of the MNR 11

Some other interesting features of the MNR 11 are:
- Digital wetgate
- Noise reduction
- Contour enhancement

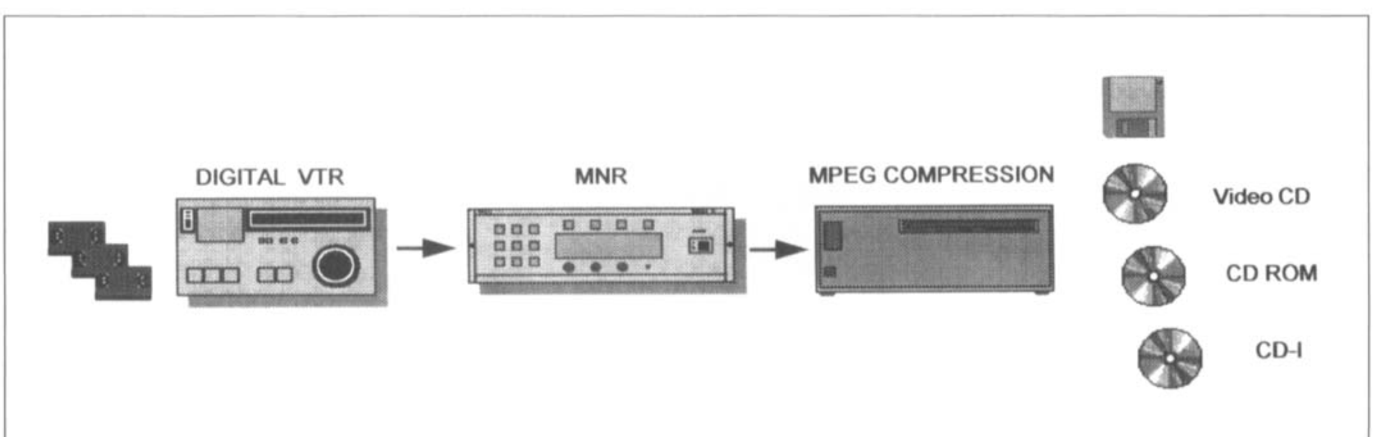

Figure 9. MPEG preprocessing with the MNR.

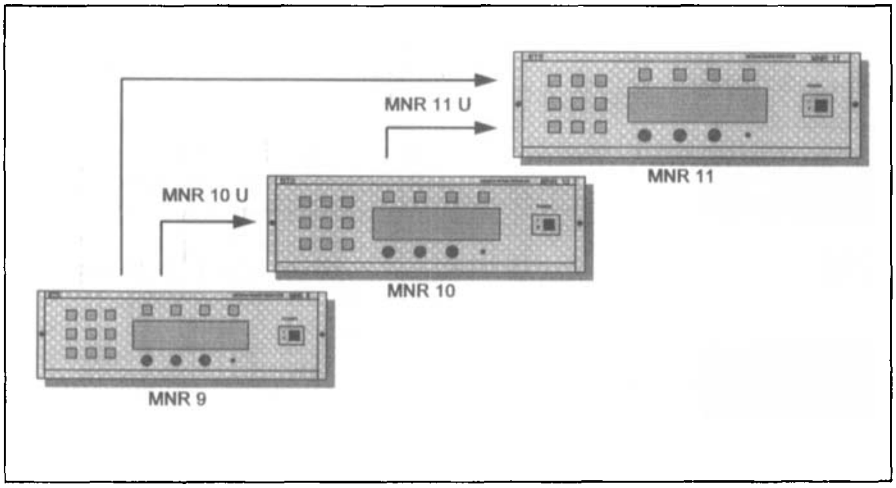

Figure 10. Upgrading the MNR.

- Time-base correction
- Frame synchronization
- Signal format transcoding

For noise reduction and image restoration in NTSC, a 3/2 pulldown detection is built in. This is required for film material and it can even be used for all slow-motion film speeds of the telecine. As any sequence is derived from the video signal itself, there is no need for a special interface. With this feature, even tape-recorded film material can be processed with full performance of the noise-reduction system. For further improving the picture quality, a two-dimensional contour enhancer is implemented with various settings. Images can be provided with a proper sharpness, or even softness, to meet subsequent encoder characteristics. Smoothing the image is often required for captions or MPEG encoding.

Time-base correction and frame synchronizing present no problems. Provision is made for adjusting luminance and chroma gain and correcting the black level. Due to the large number of different input/output (I/O) interfaces, the box can be easily used for signal format transcoding. Special interfaces for PAL and NTSC are available as an option.

Main Applications

The main applications of the MNR11 include:

- Film-to-tape transfer
- Tape-to-tape transfer
- Satellite and cable transmission
- Enhancement of MPEG encoding

Apart from these transfer applications, the MNR can be used after satellite or cable transmission for synchronization and general improvement of the signal-to-noise ratio.

MPEG Encoding

There is a fast-growing market for MPEG encoding. Noise reduction prior to MPEG compression is of major importance in achieving optimum encoding efficiency. In order to get a low data rate at the output, the MPEG encoder tries to remove all nonrelevant information from the picture. This is not so with the noise. As noise is regarded as part of the video signal, the data rate increases, which may result in a poor encoding quality with visible blocking artifacts and so-called "mosquito noise."

The MNR is very helpful in this application (Fig. 9), as noise reduction is really the key to high-quality MPEG compression. THE MNR10 is already a popular noise-reduction system in many MPEG facilities and is a recommended option for encoding hardware and software.

Upgrading the MNR

If you are the owner of an MNR9 or MNR10, you may ask: Do I have to throw away my old noise reducer? The answer is no, because the new system is compatible. All you need is one new board to upgrade your MNR with the latest noise-reduction technology (Fig. 10).

Conclusion

The digital wetgate has become an indispensable tool in the high-end post-production industry, especially for feature film transfers and commercial productions. It effectively removes, in real time, the disturbing effects of film dust, dirt, and scratches with minimum side effects. Even damages to the emulsion layer can be dealt with, not only during the telecine transfer but also later in a tape-to-tape session. The digital wetgate largely avoids time-consuming paintbox work and the hassle of conventional wetgate scanning of film, with regard to the handling and disposal of toxic chemicals.

THE AUTHOR

Gerhard Wischermann is with BTS Broadcast Television Systems GmbH, Weiterstadt, Germany, where he is involved in project management and responsible for the development of the BTS noise-reducer family. Wischermann holds a Diplom-Engineer degree in communications engineering from the Technische Hochschule Darmstadt. In 1977 he joined BTS, where he worked on advanced developments in digital signal processing until taking on his current responsibilities.

Wischermann has presented three papers at SMPTE conferences; two of these have been published in the *Journal*.