Veterans Administration

Journal of Rehabilitation Research
and Development Vol. 24 No. 4
Pages 21–28

Spectrum shaping with a hardware digital filter

GERALD A. STUDEBAKER,* ROBERT L. SHERBECOE, and JOSEPH S. MATESICH
Department of Audiology and Speech Pathology, Memphis State University, Memphis, Tennessee 38152

Abstract—Digital filters have certain characteristics which make possible some things that are either difficult or essentially impossible using analog equipment. Among the most useful characteristics are steep transition band slopes and the capability to shape the passband transfer function in great detail. Unlike most digital filters implemented in software, hardware digital filters can be used to filter signals in real-time. Also, they can be designed to provide a linear phase response. They have some disadvantages as well, however, including moderately high cost, the need for supporting equipment, and at the present time, the need for a moderate degree of computer programming and interfacing skill. We have used hardware digital filters in a number of applications. These applications and the relevant filter performance characteristics are described.

INTRODUCTION

Filters have been used in auditory research for a very long time (2, 4). Until relatively recent times, most were analog devices that operated on an electrical version of the auditory signal. Because analog circuit systems are quite limited in their filtering performance, it is necessary to use a number of filter sections strung together in series to achieve even the moderate transition band[1] slopes of 18 to 24 dB/octave typical of such units. Performance considered acceptable today (50 to 100 dB/octave, or more) requires many filter sections, which add to the complexity, cost, and internal noise of these systems. Further, analog filters typically introduce unwanted phase shifts into the signal.

In addition to simple bandpass filtering, the ability to specify the relative levels of different frequency regions within the pass band is often important. To do this using analog filters, with the degree of precision and to the extent needed for many applications, normally requires a bank of filters. A bank of filters has one bandpass filter for each frequency band that is to be controlled. A one-third octave band filter set, for example, may have as many as 30 separate bandpass filters, each of which must be adjusted individually to produce the desired response. Thus, such analog filter sets have very many components, are cumbersome and slow in use, and are relatively expensive.

Recently, practical digital filters have become available as a means of overcoming these limitations. Digital filtering is a powerful and flexible technique which can do virtually all of the things done by analog filters, and some things outside the practical limits of analog filters. In addition, digital filters can be designed to eliminate unwanted phase shifts that are an inherent part of analog filtering. However, digital filtration does have its own set of problems. Most have to do with the undesirable effects of sampling. A principal problem is the generation of aliases. Analog filters may be required in order to reduce or eliminate this problem.

*Address correspondence to: Gerald A. Studebaker, 807 Jefferson Avenue, Memphis, Tennessee 38105.

Journal of Rehabilitation Research and Development Vol. 24 No. 4 Fall 1987

Digital filtration can be accomplished either by computer software or by using hardware devices programmed specifically for this task. Software implementations are relatively slow; filtration cannot normally be effected in real-time with this method. Hardware digital filters, however, are fast enough so that they can be used on-line in the same way as an analog filter. Two such filters are used in our laboratory. They are integrated into a system that permits the user to rapidly specify the transfer function the filters will produce. This sytem, its performance, and a number of applications are described in this paper.

SYSTEM DESCRIPTION

Overview

There are many ways to accomplish the functions described below. Thus, our system is an example of only one way to perform such operations. **Figure 1** provides a block diagram of the system, including some additional parts necessary to monitor the system and report results. The actual system has two signal pathways, including two digital filters, two sets of antialiasing filters, two attenuators, etc. In order to reduce complexity, only one channel is shown in **Figure 1**. For the same reason, certain components (e.g., video monitor, numeric keypad, etc.), included when the system is used to test human subjects, are not shown. This system is built using general purpose laboratory equipment. It can be expected that one specially designed would exceed the speed and capabilities of our system in any particular application at considerably less cost.

THE DIGITAL FILTER

The component in **Figure 1** of central interest is the digital filter. It has unique characteristics as compared to analog filters and so we will provide a relatively detailed description of it.

Digital filtration can be accomplished in a variety of ways, as illustrated in the writings of Oppenheim and Schafer (7), Levitt, Neuman, Mills, & Schwander (6), Otnes and Enochson (8), Rabiner and Shafer (9), and others. These references should be consulted for a discussion of the characteristics of these other approaches and for more information on the approach described here.

The digital filters we use were designed by J. P. Trinder of the Institute of Hearing Research (IHR), Nottingham, England (13). Each unit provides a single signal channel. To use the device, one directs an analog signal to its input and receives a filtered analog signal at its output, as with analog filters. A/D (12-bit) and D/A (16-bit) converters are included within the unit, but it is necessary to supply anti-aliasing filters if they are needed.

The IHR filter uses a "direct" convolution method (12) to create an FIR (finite impulse response) digital filter. In a direct convolution the signal is convolved with the impulse response of the filter in the time domain. This is in contrast to what might be described as an "indirect" convolution, wherein Fourier transforms of the signal and the impulse response are multiplied in the frequency domain followed by an inverse Fourier transform of the result back into a time domain signal. For an example of the indirect approach see Levitt, Neuman, Mills, and Schwander (6). In the present case, a discrete finite duration impulse response is derived that produces a transfer function approximately equal to the transfer function called for by the user. The sampled signal is then convolved in the time domain with the obtained coefficients of the sampled impulse response.[2] An advantage of the finite impulse response filter is that it can easily be programmed for a linear phase response.

The IHR filters are supplied with a set of "host files." Each is designed to be used at a maximum sampling rate of 20 k, 30 k, 40 k, 50 k or 60 kHz. Within the limitations set by the number of available coefficients and the sampling rate, a single host file will produce any filter transfer function specified by the user. Higher sampling rates reduce aliasing problems but have the disadvantages that they result in wider transition bands[3] and larger errors.[4] Some host files provide narrower transition bands than others at the same sampling rate but, again, at the cost of larger errors. However, even with the narrowest transition band, the errors are still small enough to be of little practical concern.

The transition band widths associated with the supplied host files range from 78 Hz to 768 Hz. A FORTRAN program is also provided with which the user can create his/her own host files. Thus, within certain limits, users can select combinations

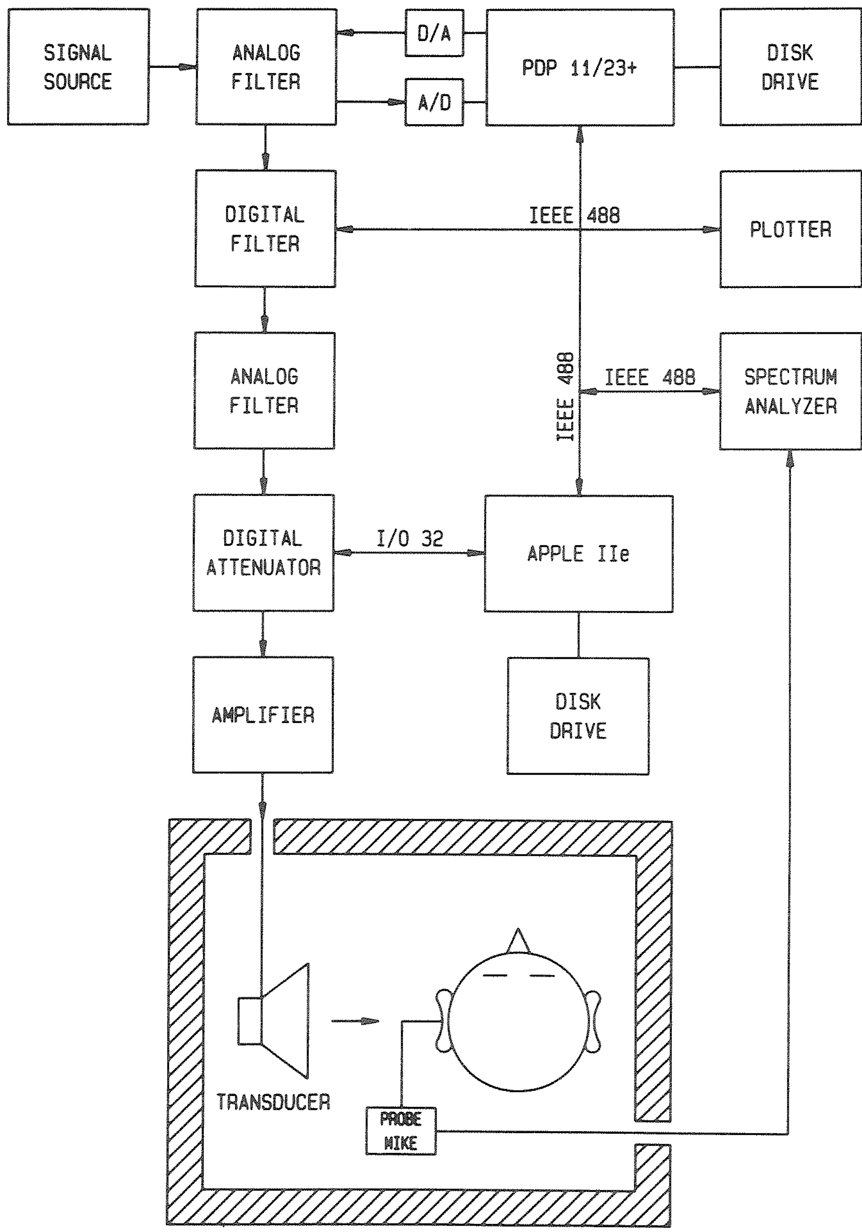

Figure 1.
Simplified block diagram of the system surrounding the digital filters described in this paper. Only one of two signal channels is shown.

of sampling rate, transition band, and error.

The IHR filter can provide an exactly linear phase response by specifying a symmetric or antisymmetric, odd or even impulse response. In these four modes, the maximum number of points in the impulse response is 512. Alternatively, an arbitrary phase response can be specified. In this case, the maximum number of impulse response coefficients is 256. (For a discussion of symmetry and related issues see Ramirez (10) or the other texts on digital signal processing listed in the References.)

OTHER COMPONENTS

Like all devices that sample continuous signals, digital filters may produce aliases of the input signal. Thus, antialiasing filters will be required for the input and the output signals of the digital filters if those signals include frequencies that exceed one-half the sampling frequency. In our applications we have typically used a 20 kHz sampling rate—therefore, antialiasing filtering is necessary. To do so, the input to the digital filters is filtered with an

Journal of Rehabilitation Research and Development Vol. 24 No. 4 Fall 1987

analog filter providing rejection rates of 115 dB/octave. The steep rejection rate permits the cutoff frequency of this filter to be set relatively near the Nyquist folding frequency. The output of the digital filters is also filtered using analog filters producing rejection rates of 48 dB/octave.

The system is controlled by two computers, a Digital Equipment Corporation PDP 11/23+ and an Apple IIe.[5] The computers communicate with each other via an IEEE-488 interface. The same interface is used to send filter coefficients from the PDP 11/23+ computer to the digital filters, to send commands to a digital plotter from either the IIe, the PDP 11/23+, or an RTA (real-time spectrum analyzer), and to send data from the RTA to either computer. The Apple computer also controls a pair of digital attenuators (one for each channel) via an inexpensive 16-bit parallel interface card.

USING THE SYSTEM

The components of the system shown in **Figure 1** can be used to create any spectrum (within the performance limitations of the hardware and the host file) between 0 Hz and a frequency one-half the sampling rate. In early applications, we entered the frequencies and relative intensity levels by hand from a keyboard. In more recent work, this aspect of the procedure has been automated.

The steps necessary to set the digital filter to a particular frequency response characteristic are as follows. First, a data file specifying attenuation values from 0 to −200 dB at user-selected frequencies is created, either by typing the values into the computer using a text editor or by transferring them from another device such as a real-time spectrum analyzer (RTA) or another computer. In the example shown in **Figure 1**, a loudspeaker produces a signal which is monitored in or near the subject's ear canal by a probe microphone. The transduced signal is sent to the RTA which produces a spectrum of the signal. This spectrum, after correction for the microphone system transfer function, is compared with the desired target spectrum and the differences (as a function of frequency) are calculated. The differences and the associated frequencies are then entered into the PDP 11/23+ computer and stored in a disk file in ASCII format. Additional frequency and level entries are allowed depending only on the detail desired in the transfer function. Actual results, however, are limited by system capabilities.

The ASCII data file is read by a FORTRAN program that calculates the filter weighting coefficients necessary to produce the desired result. The output of this program is a binary data file. The binary file is loaded into the filter by a loading program called from the keyboard or from a running application program.

Binary files of filter coefficients are stored on cassette tape or on disk under a user-selected name. Communication between the computer and the storage device can be via an RS-232C or an IEEE-488 interface. We use the IEEE-488 interface in order to minimize filter loading time.

SYSTEM PERFORMANCE

Filter Creation and Loading Times

The coefficient generation program may take from 3 seconds to 3 minutes to generate a 512-point coefficient file, the amount of time depending on the complexity of the transfer function requested. The biggest of the resulting coefficient files consumes less than 3 kilobytes of memory, so all of the files to be used in a test session can be loaded into memory in advance of the session. With the files in memory, the filters can be set or changed to a new setting in 34 ms, using a 512-point impulse response (the largest available) running at 20 kHz. For a 196-point impulse response with the filter running at 60 kHz the filter can be loaded and restarted in just over 12 ms. The signal is turned off during the loading period. When the files are read from a fast Winchester disk, just prior to loading the filters, the time required is increased by about 500 ms depending on the disk speed. The signal "off" time is not increased by a disk access but only delayed by the time needed to fetch the file.

PROCESSING DELAY

The convolution technique results in a delay between the input and output signals that is dependent on the sampling interval and the number of points in the impulse response.[6] Our measurements indicate that the delay is about 12.8 ms with the

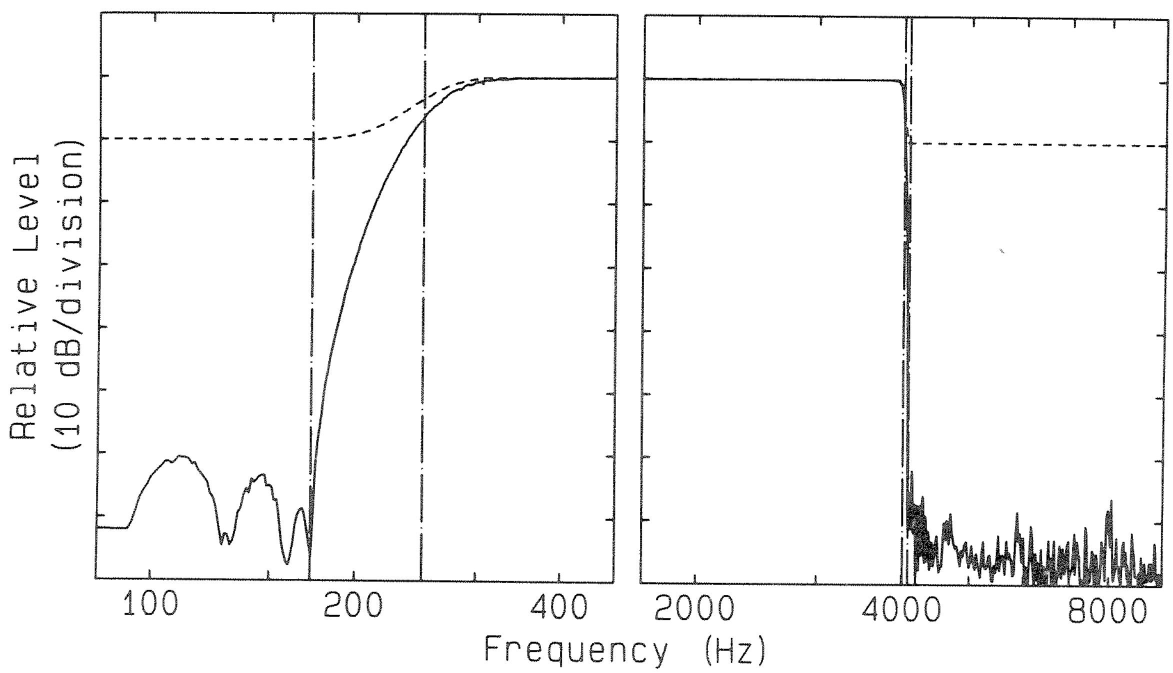

Figure 2.
Transfer function of the IHR digital filter set to pass the band between 250 Hz and 4 kHz. The solid line shows the result with the filter set to maximum attenuation in the stop band (−200) while the dashed line shows the result with stop band attenuations set to −10 dB. The vertical dot-dash lines indicate the nominal position of the 78 Hz transition band.

equipment set up to have a 20 kHz sampling rate and a 512-point symmetrical impulse response. The delay with a 60 kHz sampling rate and a 196-point symmetrical impulse response is about 1.6 ms. These values are essentially equal to those expected from calculation.

Highpass and Lowpass Filtering

In many experiments, very steep filter slopes are needed. For example, steep slopes were useful in several of our recent speech intelligibility studies (3, 11) because they greatly reduced the amount of information contributed by speech energy in the filter skirts beyond the nominal cutoff frequencies. We achieved such slopes by using a filter sampling rate of 20 k samples/second and a host file with a nominal transition bandwidth of 78 Hz, the steepest available. (This host file also produced the largest error of any available to us that operated at this sampling rate.)

Figure 2 shows four examples of system performance in the form of highpass and lowpass filters for each of two attenuation settings in the rejection regions. The low-frequency cutoff was 250 Hz and the high-frequency cutoff was 4 kHz. Attenuation in the pass band was set at 0 dB in both cases. Attenuation was set at the maximum value (nominally −200 dB) outside that band in the case of the solid line and at −10 dB in the case of the dashed line. The nominal transition bands for both the highpass and lowpass filters (78 Hz) are indicated by the dash-dot lines. The results shown in the figure reveal that the filters produce transition bands very close to the nominal value and that exactly 10 dB of attenuation was obtained when called for. In the case of the solid line, the depth of the rejection region is limited by stop band ripple and noise which were between 60 dB and 75 dB down from the pass band. Stop band ripple may be made somewhat lower with a different host file at the cost of a wider transition band.

Figure 2 shows that the transfer function slope within the transition band varies with frequency and depth of the rejection region. When maximum attenuation was specified, slopes of approximately 107 dB/octave and 3436 dB/octave for the highpass

Journal of Rehabilitation Research and Development Vol. 24 No. 4 Fall 1987

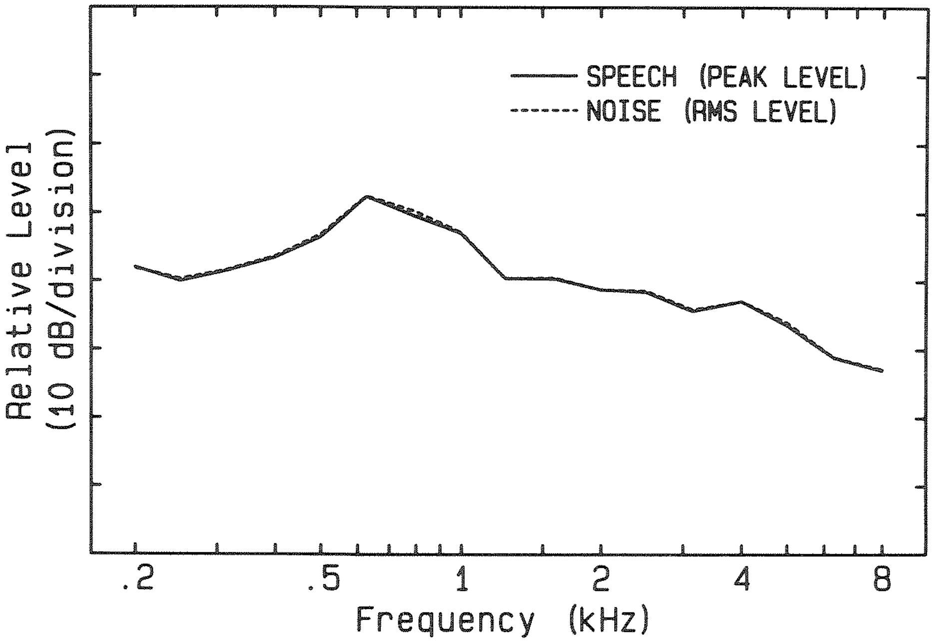

Figure 3.
The one-third octave band levels of a target speech peaks spectrum (solid line) and the one-third octave band levels of a thermal noise produced by a filter designed to match those peak levels (dashed line).

and lowpass filters were produced. In contrast, 10 dB of attenuation resulted in highpass and lowpass slopes of only 16 dB/octave and 512 dB/octave. Thus, very steep slopes are achieved in the high frequencies while in the low-frequencies slopes in dB/octave become increasingly shallow. Nevertheless, in the frequency regions important to most auditory research (>100 Hz), slopes equal to or greater than those obtainable with most analog filters are possible.

Generating Arbitrary Master Spectra

One application in which we used the IHR filters was in the creation of masking noises with one-third octave band spectra matched to the one-third octave band speech spectra of talkers of various test materials. The filters can carry out this function at the same time as they produce highpass and/or lowpass filtering such as that shown in **Figure 2**.

A masking noise ipsilateral to the test signal was essential in some of our recent work (3, 11) because without it the test material would have been too well understood to produce useful results. Further, the masking noise spectrum had to be matched to the spectrum of each talker in order to avoid masking some parts of the speech spectrum more than others. In the study by Studebaker Pavlovic and Sherbecoe (11), for example, three different masking spectra were used, each matched individually to the spectrum produced by the speech of one of three different talkers of the test materials. When the testing protocol required a particular talker, the appropriate masking spectrum filter file was automatically loaded into the digital filter.

The first step in creating the masking noise was to estimate the long-term average levels of speech in one-third octave bands. A filter was then created to produce these levels. The resulting levels were compared to the target levels and the filter file was adjusted as necessary. **Figure 3** shows the result of this procedure from a study by Duggirala et al. (3). The rms difference between the two sets of band levels was 0.26 dB. This is a typical result using this method.

Creating a Spectrum with an RTA

In some applications there is a need to specify the spectrum in more detail than one-third octave bands. For that purpose, we developed a system wherein the levels at each of the 400 frequency points provided by our RTA are transferred to the computers via the IEEE-488 interface. To illustrate the success of this technique for this report, we chose to emulate a hearing aid frequency response under four conditions of sampling rate. A thermal noise served as the test signal. The solid line in each

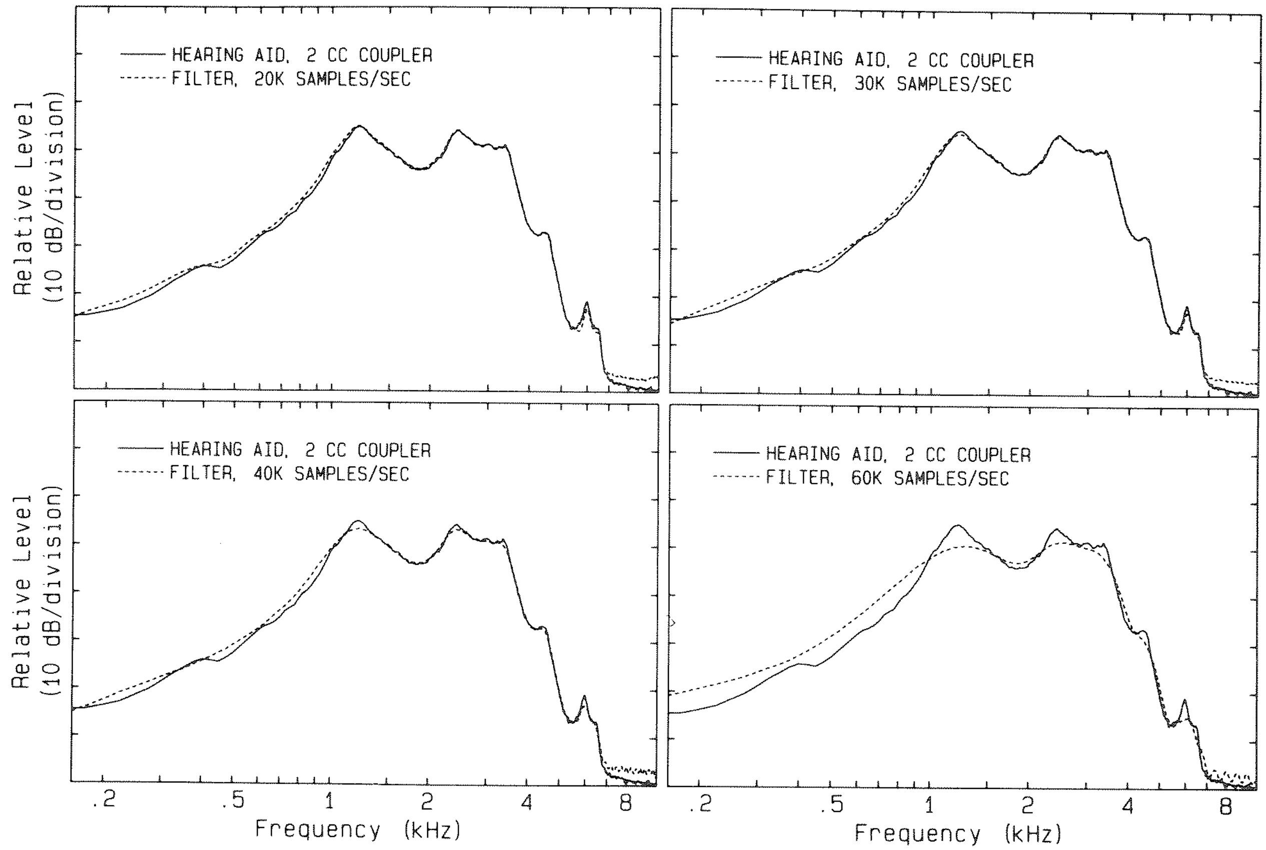

Figure 4.
The output spectra of a hearing aid (solid line) and of a digital filter designed to emulate that hearing aid (dashed lines) operated at four sampling rates. A thermal noise was the test signal and the output was observed on a 400 line RTA. The dashed lines were adjusted to produce the least square difference between the two lines in each case.

of the panels of **Figure 4** is the averaged 2-cm^3 coupler response of the hearing aid as seen on the RTA. This spectrum was sent to the PDP 11/23+ via the IEEE-488 interface. The filter coefficients were calculated from the resulting file and sent to the digital filter. Finally, the thermal noise was directed to the digital filter and its response observed on the same RTA.

The dashed lines in **Figure 4** represent the response of the digital filter at each sampling rate. Before plotting, the digital filter's response curve was adjusted iteratively to produce the minimum rms difference between it and the hearing aid response curve below 6700 Hz. For the lowest sampling rate, 20 k samples/second, the rms difference between the two lines was about 0.7 dB. When the rate was increased to 30 k and 40 k samples/second the accuracy of the emulation changed only slightly, the rms difference always remaining less than 0.8 dB. Apparently, the transition band remains sufficiently narrow through this sampling frequency to provide a good emulation of a hearing aid with a response like the one shown.

At 60 k samples/second substantial differences can be seen between the two curves. This is further reflected by an rms difference of about 2 dB under this condition. This condition is probably not adequate to emulate the precise details of a particular hearing aid transfer function. However, the ability to emulate such details is often not necessary or even desirable. In addition, a high sampling rate provides the advantages of a decreased or nonexistent need for antialiasing filters and of a relatively short time delay between input and output signals.

The host files used in all four panels were those that produce relatively large amounts of error and narrow transition bands. Error does not change significantly with sampling rate. This suggests that transition band width, and not error, was the important factor limiting the accuracy of emulation at the 60 kHz sampling rate.

Direct quantitative comparisons with the results reported by Levitt, Sullivan and Hwang (5) in the same kind of emulation are not possible because their numerical data were not available to us. However, for the lowest three sampling rates, the relationships between the pairs of lines seen in **Figure 4** appear to be essentially the same as they reported.

Systematic differences between the target and the resulting spectra may be noted in the four panels of **Figure 4**. Specifically, there is a tendency for the digital filter result to be somewhat higher in the rising parts of the spectrum and somewhat lower in the falling parts of the spectrum. This anomaly is caused by the way in which the filter interprets the level and frequency information it receives. That is, the digital filter produces the same intensity at all frequencies from the first frequency for which a level is specified up to the next frequency for which a level is specified. Although the effect of this factor can be seen in **Figure 4**, its extent cannot be gauged precisely from that figure, particularly in the case of the 60 k samples/second condition, because of the absence of a basis for precisely aligning level. In other cases, such as in the emulation of flat-topped bandpass filters, the effect can be readily judged, and in the worst case (steeply rising transition bands in the low frequencies) may result in discrepancies in the range of 2 to 3 dB. These differences can be made smaller by increasing the frequency resolution of the RTA and, to a degree, they can be corrected by software. Thereby, with the expenditure of additional time, this error could be made acceptably small for nearly any application. However, in many of the more common applications, such as the one shown in **Figure 4**, the magnitude of the error, even without any special measures, is too small to be a problem.

FOOTNOTES

[1]For practical purposes, the transition band may be thought of as the difference in Hertz between the nominal cut-off frequency and the frequency at which the specified intensity level of the adjacent band is achieved. However, the precise definition is rather more complicated. For details see Abed and Cain (1).

[2]For an explanation of time-domain convolution see, for example, Ramirez (10). For a discussion of the specific window design technique used in the IHR filter, see Abed and Cain (1).

[3]The increase in the transition band width with sampling rate is a direct result of 1) the increased total band width resulting from the higher sampling rate (thus, each coefficient must account for a broader frequency region) and 2) the number of coefficients must be decreased in order for the computations to be completed within the time available at the higher sampling rate. The IHR Universal filter computational rate is just over 6 million multiplications per second. Thus, at 60 k samples/second, for example, approximately 100 calculations per sample are possible. Ninety-eight coefficients are, in fact, used by the IHR filters at this sampling rate.

[4]Errors include ripple in the transfer function and variations in the transition band width.

[5]The Apple computer is not an essential component in this system. It is included because it makes the system very much easier to use in day-to-day applications.

[6]The delay equals mT in the case of odd length impulse responses and (m-1/2)T in the case of even length impulse responses where m equals one-half the number of points in the impulse response, in the case of symmetry, or equals the number of points in the impulse response in the case of an arbitrary phase response. T equals the sampling interval (12).

REFERENCES

1. **Abed AHM and Cain GD:** Lowpass digital filtering with the host windowing design technique. *The Radio and Electronic Engineer* 48(6):293–300, 1978.
2. **Crandall IB:** The composition of speech. *Phys Rev* 10:74–76, 1917. Reprinted in: Hawley ME (ed), *Speech Intelligibility and Speaker Recognition*. Stroudsburg, PA: Dowden, Hutchinson and Ross, 1977.
3. **Duggirala V, Studebaker GA, Pavlovic CV, and Sherbecoe, RL:** Band importance functions for certain consonant features. Paper presented at the 111th meeting of the Acoustical Society of America, Cleveland, OH, 1986.
4. **Fletcher H and Steinberg JC:** Articulation testing methods. *The Bell System Technical Journal* 8:806–854, 1929.
5. **Levitt H, Sullivan J, and Hwang J:** A computerized hearing aid measurement/simulation system. *Hearing Instruments* 37(2):16–18, 1986.
6. **Levitt H, Neuman A, Mills R, and Schwander T:** A digital master hearing aid. *J Rehabil Res Dev* 23(1):79–87, 1986.
7. **Oppenheim AV and Schafer RW:** Digital signal processing. Englewood Cliffs, NJ: Prentice-Hall, 1975.
8. **Otnes RK and Enochson L:** *Applied time series analysis* (Volume 1). New York: Wiley-Interscience Publication, 1978.
9. **Rabiner LR and Shafer:** *Digital Processing of Speech Signals*. Englewood Cliffs, NJ: Prentice-Hall, 1978.
10. **Ramirez RW:** *The FFT. Fundamentals and Concepts*. Englewood Cliffs, NJ: Prentice-Hall, 1985.
11. **Studebaker GA, Pavlovic CV, and Sherbecoe RL:** A band importance function for continuous discourse. Paper presented at the annual meeting of the American Speech-Language-Hearing Association, Detroit, MI, 1986.
12. **Trinder JP:** Hardware-software configuration for high performance digital filtering in real-time. In: *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*, 687–690. New York: IEEE, 1982.
13. **Trinder JP:** *IHR Universal Filter: Hardware Manual*. Nottingham, England: MRC Institute of Hearing Research, Undated.