

# Multichannel Dynamic-Range Compression Using Digital Frequency Warping

**James M. Kates**

*GN ReSound, c/o Department of Speech, Language & Hearing Sciences, University of Colorado at Boulder, 409 UCB, Boulder, CO 80309, USA  
Email: jkates@gnresound.dk*

**Kathryn Hoberg Arehart**

*Department of Speech, Language & Hearing Sciences, University of Colorado at Boulder, 409 UCB, Boulder, CO 80309, USA  
Email: kathryn.arehart@colorado.edu*

*Received 22 March 2004; Revised 13 September 2004*

A multichannel dynamic-range compressor system using digital frequency warping is described. A frequency-warped filter is realized by replacing the filter unit delays with all-pass filters. The appropriate design of the frequency warping gives a nonuniform frequency representation very close to the auditory Bark scale. The warped compressor is shown to have substantially reduced group delay in comparison with a conventional design having comparable frequency resolution. The warped compressor, however, has more delay at low than at high frequencies, which can lead to perceptible changes in the signal. The detection threshold for the compressor group delay was determined as a function of the number of all-pass filter sections in cascade needed for a detectible change in signal quality. The test signals included clicks, vowels, and speech, and results are presented for both normal-hearing and hearing-impaired subjects. Thresholds for clicks are lower than thresholds for vowels, and hearing-impaired subjects have higher thresholds than normal-hearing listeners. A frequency-warped compressor using a cascade of 31 all-pass filter sections offers a combination of low overall delay, good frequency resolution, and imperceptible frequency-dependent delay effects for most listening conditions.

**Keywords and phrases:** hearing aids, dynamic-range compression, frequency warping, delay perception.

## 1. INTRODUCTION

Multichannel dynamic-range compression is an important feature in hearing aids (Kates [1]). Compared to normal listeners, hearing-impaired listeners typically have elevated auditory thresholds which interfere with the perception of low-intensity sounds. However, the perception of loudness for intense sounds is often similar to that of normal-hearing listeners. Thus, an objective of a hearing aid is to fit the dynamic range of speech and everyday sounds into the restricted dynamic range of the impaired ear. Furthermore, hearing losses are typically frequency-dependent, so the compressor should provide different amounts of dynamic-range compression in different frequency regions. The solution to this design problem is generally a multichannel system, such as a filter

bank, with different degrees of compression in each channel.

The design of a multichannel compressor involves a fundamental trade-off between frequency resolution and time delay. For any given processing approach, increased frequency resolution comes at the price of increased processing delay. In this paper, a new compression algorithm based on digital frequency warping is introduced. Compared to conventional digital processing algorithms, the use of digital frequency warping inherently gives frequency resolution on an auditory frequency scale, and also reduces the amount of processing delay for a specified degree of low-frequency resolution. However, the processing delay of the frequency-warped compressor is frequency-dependent, with greater delay at low frequencies than at high frequencies. The frequency-warped compressor design must therefore take into account the frequency resolution, overall system processing delay, and delay variation across frequency. The objective is to design a compression system that has good

frequency resolution while avoiding audible artifacts caused by the processing delay.

### 1.1. Frequency resolution

One concern in designing a multichannel compressor is to match the frequency resolution of the digital system to the resolution of the human auditory system. For example, several hearing-aid fitting procedures are based on loudness scaling in the impaired ear (Dillon et al. [2]), and the estimation of loudness presupposes an auditory frequency analysis. Digital frequency analysis, such as the discrete Fourier transform, typically provides constant-bandwidth frequency resolution. The frequency resolution of the human auditory system, however, is more accurately modeled by a filter bank having a nearly constant bandwidth at low frequencies but with bandwidth becoming proportional to frequency as the frequency increases (Moore and Glasberg [3]; Zwicker and Terhardt [4]).

The mismatch between digital and auditory frequency analyses can be greatly reduced if the conventional uniform frequency analysis is replaced by a warped frequency analysis. Frequency warping uses a conformal mapping to give a nonuniform spacing of frequency samples around the unit circle in the complex- $z$  plane (Oppenheim et al. [5]; Oppenheim and Johnson [6]). With an appropriate choice of the parameters governing the conformal mapping (Smith and Abel [7]), the reallocation of frequency samples comes very close to the Bark frequency scale (Zwicker and Terhardt [4]) used to describe the auditory frequency representation. Frequency warping therefore allows the design of digital audio systems that have uniform time sampling but which have a frequency representation similar to that of the human auditory system (Härmä et al. [8]; Karjalainen et al. [9]).

Frequency warping can be used to design both finite-impulse response (FIR) and infinite-impulse response (IIR) filters (Karjalainen et al. [10]). A frequency-warped FIR filter, for example, can be designed by replacing the unit delays in the conventional FIR filter with all-pass filter sections. Improved frequency resolution in a conventional FIR filter requires increasing the filter length, which leads to an increase in the filter group delay. Similarly, improved frequency resolution in a warped FIR filter requires an increase in the number of all-pass filter sections that comprise the filter, which also leads to an increase in the filter delay. Thus, there is a trade-off between frequency resolution and group delay for both conventional and warped filters, although the warped filter has less delay at low frequencies than a conventional filter for the same low-frequency resolution.

Frequency warping has been shown to be effective in a number of audio applications. Linear predictive coding (LPC) of speech using frequency warping (Strube [11]) was found to give higher speech intelligibility and quality than conventional LPC for predictor orders below eight, although for higher predictor orders, the frequency warping offered little benefit. Loudspeaker equalization filters designed directly in the warped frequency domain (Karjalainen et al. [9]) were found to require a lower filter order than their conven-

tional counterparts to achieve comparable perceptual benefits. Warped filters have also proven effective in modeling the acoustic properties of musical instruments, reducing the necessary filter orders by a factor of five to ten, and in reducing the filter order needed to model the head-related transfer function (HRTF) used in synthesizing 3D sound localization cues (Karjalainen et al. [12]).

### 1.2. Delay effects in speech production

A second concern in designing a compression system for a hearing aid is the overall processing delay. These time delays can cause coloration effects to occur when the hearing-aid user is talking. When talking, the talker's own voice reaches the cochlea with minimal delay via bone conduction and through the hearing-aid vent. This signal interacts with the delayed and amplified signal produced by the hearing aid to produce a comb-filtered spectrum at the cochlear. Delays as short as 3 to 6 milliseconds that are constant across frequency are detectable (Agnew and Thornton [13]; Stone and Moore [14]), and overall delays in the range of 15 to 20 milliseconds can be judged as disturbing or objectionable (Stone and Moore [14, 15]).

Stone and Moore [16] also studied the effects of frequency-dependent group delay on the production of speech by ten listeners with bilateral cochlear hearing loss. A delay ranging from 0 to 24 milliseconds was added to the low frequencies relative to the fixed broadband system delay of 2.5 milliseconds. The low-frequency delay thus ranged from 2.5 to 26.5 milliseconds, while the high-frequency delay remained at 2.5 milliseconds. The frequency-dependent delays did not significantly affect the subjects' word production rates. However, the subjects' perception of their own voices was significantly affected by the group delay; mean ratings of speech processed with a 20-millisecond additional low-frequency delay were "disturbing" and an additional low-frequency delay of 9 milliseconds was significantly more disturbing than no additional low-frequency delay.

The results for delays independent of frequency indicate that overall processing delays of 10 to 15 milliseconds will be acceptable under most speaking conditions. However, the frequency-dependent delay results indicate that low-frequency delays of 9 to 15 milliseconds, when added to the overall system delay of 2.5 milliseconds, can cause significant subjective disturbance. Thus total low-frequency delays of less than 11.5 milliseconds, and across-frequency delays of less than 9 milliseconds, are necessary to ensure that objectionable delays will be avoided for most talking conditions.

### 1.3. Frequency-dependent delay effects in listening

The compression system described in this paper uses frequency warping to reduce the system delay while still providing good frequency resolution on a critical-band frequency scale. However, the delay in the warped system is frequency-dependent, with a greater delay at low than at high frequencies. This frequency-dependent delay can also introduce audible artifacts when listening to speech even when the user of the hearing aid is not talking. For example, a click is

converted into a descending chirp when passed through a cascade of all-pass filters having a group delay that increases with decreasing frequency.

Relatively short delays can be detected for click stimuli when the group delay varies across frequency. Blauert and Laws [17] passed clicks through all-pass filters giving increased delay in narrow-frequency regions, and found that normal-hearing subjects can detect delays as short as 1 millisecond at 2 kHz, with the detection threshold increasing to 2 milliseconds at 8 kHz or 1 kHz. In experiments using Huffman sequences and normal-hearing subjects, Green [18] found a delay detection threshold of 2 milliseconds independent of frequency. In experiments using click-like stimuli and normal-hearing subjects, Banno et al. [19] found a detection threshold of 2 milliseconds for group-delay variations that spanned more than one equivalent rectangular bandwidth (ERB). However, they found that the threshold of detection was approximately 8 to 10 milliseconds for group-delay variations that were constrained to lie within an ERB; these results suggest that in normal-hearing listeners, cross-channel phase effects may be more important than within-channel effects for short stimuli.

Group-delay detection thresholds for speech are greater than for clicks. Based on results using one normal-hearing subject, Greer [20] reported detection thresholds between 0.0625 and 0.125 millisecond for dispersed impulses when passed through all-pass filters that give increased delay in narrow frequency regions. For all-pass filters having group-delay effects occurring over a frequency region that corresponded to approximately 20 percent of the filter's center frequency, detection thresholds for speech sounds were 4 to 8 milliseconds for a plosive, 8 to 16 milliseconds for a vowel, and 16 to 32 milliseconds for a fricative. For all-pass filters having group-delay effects occurring over a broader frequency region corresponding to approximately the filter's center frequency, the detection thresholds were 2 to 4 milliseconds for a plosive, 2 to 4 milliseconds for a vowel, and 4 to 8 milliseconds for a fricative. The results of Greer [20] are consistent with those of Banno et al. [19] in that the detection threshold is lower for all-pass filters spanning more than one ERB.

The frequency-dependent group delay also can interfere with speech intelligibility, but at delays that greatly exceed the detection thresholds. Stone and Moore [16] found that hearing-impaired listeners' identification of nonsense syllables decreased by a small but significant amount as the low-frequency delay was increased (from 72.3 percent at no delay to 68.1 percent at 24 milliseconds additional low-frequency delay). Arai and Greenberg [21] introduced delay variations as a function of frequency in sentence materials, and found that word identification accuracy for normal-hearing subjects decreased as the delay variations increased. However, listeners maintained good word identification (75%) with across-band delay variations of 140 milliseconds; this delay duration is long enough to encompass two or more phonemes, and therefore represents a scrambling of the order of portions of the speech sounds.

## 1.4. Objectives

The purpose of this paper is to describe a dynamic-range compression system based on digital frequency warping, and to determine the detection threshold for the frequency-dependent group delay inherent in the warped compression system. The paper begins with a description of frequency warping and the benefits of using symmetric warped filters. The warped compressor algorithm is then described, and its group-delay behavior is illustrated. The warped compressor is shown to have substantially reduced delay in comparison with a conventional design having comparable frequency resolution. Good frequency resolution can be achieved with an overall delay that would be expected to fall below the threshold for audible interference when the hearing-aid user is talking.

The warped compression system, however, introduces a frequency-dependent group delay. The detection threshold for the group delay is then determined as the number of all-pass filter sections in cascade needed for a detectible change in the signal. The test signals include clicks, synthetic vowels, and speech, and results are presented for both normal-hearing and hearing-impaired subjects. The paper concludes with recommendations for warped compressor design considering the trade-offs between frequency resolution, processing delay, and the ability to detect frequency-dependent delay effects while listening.

## 2. COMPRESSION ALGORITHM

### 2.1. Digital frequency warping

Digital frequency warping is achieved by replacing the unit delays in a digital filter with first-order all-pass filters (Oppenheim et al. [5]; Oppenheim and Johnson [6]; Karjalainen et al., [9]; Smith and Abel [7]; Härmä et al. [8]). The all-pass filter is given by

$$A(z) = \frac{z^{-1} - a}{1 - az^{-1}}, \quad (1)$$

where  $a$  is the warping parameter. The frequency warping that results for different choices of the parameter  $a$  is illustrated by Oppenheim and Johnson [6] and Karjalainen et al. [9]. The value for the warping parameter that gives a closest fit to the Bark frequency scale is  $a = 0.5756$  for a 16 kHz sampling rate (Smith and Abel [7]). The group delay for this choice of parameters is illustrated in Figure 1. The delay at low frequencies exceeds one sample, while the delay at high frequencies is less than one sample.

The warped FIR filter transfer function is the weighted sum of the outputs of each all-pass section:

$$B(z) = \sum_{k=0}^K b_k A^k(z), \quad (2)$$

for a filter having  $K + 1$  taps ( $K$  all-pass sections). Forcing the real filter coefficients  $\{b_k\}$  to have even symmetry for an unwarped FIR filter yields a linear-phase filter, in which the filter delay is independent of the coefficients as long as the

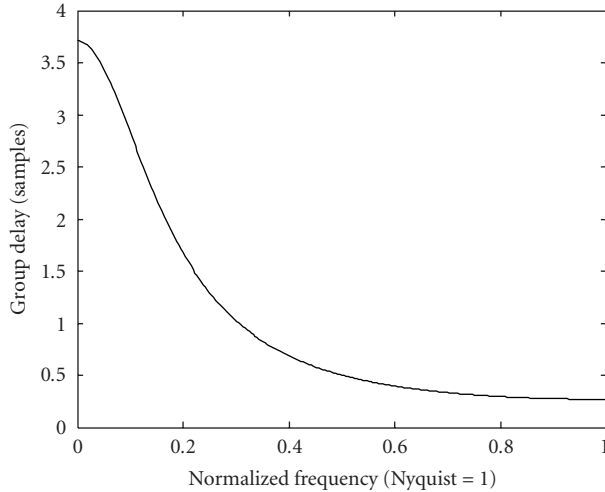


FIGURE 1: Group delay in samples for a single all-pass filter section with  $a = 0.5756$  at a sampling rate of 16 kHz.

symmetry is preserved. If the unwarped FIR filter has  $K + 1$  taps, the delay is  $K/2$  samples. Similarly, forcing even symmetry for the coefficients of a warped FIR filter gives a filter having a fixed frequency-dependent group delay that is independent of the actual filter-coefficient values. As shown in (3) and (4), if the warped FIR filter has  $K + 1$  taps, the group delay is  $K/2$  times that of a single all-pass filter. This filter-coefficient symmetry property guarantees that no phase modulation will occur as the compressor changes gain in response to the incoming signal, thus avoiding the question of the audibility of phase-modulation effects. Furthermore, in a binaural fitting (hearing aids on both ears), the coefficient symmetry ensures that identical amounts of group delay are introduced at the two ears by the hearing-aid processing, thus preserving the interaural phase differences that are used for sound localization.

Consider a warped FIR filter having even symmetry and an even number of taps. The analysis for an odd number of taps is similar. The  $z$ -transform of a warped FIR filter is given by (2). Assume that the filter coefficients are real and have even symmetry, giving  $b_k = b_{K-k}$ . For  $K$  odd, the filter has an even number of taps and the  $z$ -transform can be rewritten as

$$B(z) = \sum_{k=0}^{(K-1)/2} b_k \left[ A^k(z) + A^{(K-k)}(z) \right]. \quad (3)$$

Rearranging the delay terms leads to

$$\begin{aligned} B(z) &= \sum_{k=0}^{(K-1)/2} b_k A^{K/2}(z) \left[ A^{(k-K/2)}(z) + A^{-(k-K/2)}(z) \right] \\ &= A^{K/2}(z) \sum_{k=0}^{(K-1)/2} b_k \left[ A^{(k-K/2)}(z) + A^{-(k-K/2)}(z) \right]. \end{aligned} \quad (4)$$

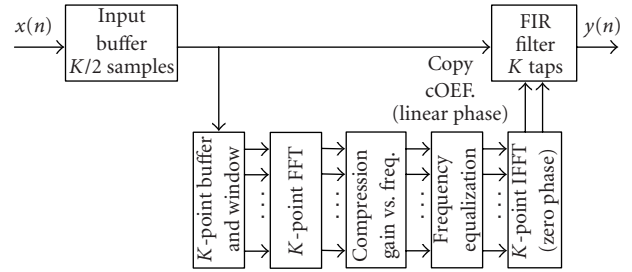


FIGURE 2: Block diagram of a compressor using a side branch for frequency analysis, with compression gains applied to the signal through an FIR filter in the signal path.

The filter delay is determined by evaluating (4) on the unit circle. Because  $A^{-1}(e^{j\omega}) = A^*(e^{j\omega})$ , the term inside the summation is pure real and does not contribute to the filter group delay, while the term outside the summation in (4) represents a fixed frequency-dependent group delay. The symmetric warped FIR filter thus has a fixed group delay that does not depend on the actual filter coefficients as long as the symmetry is maintained. For example, a warped filter using 31 all-pass filter sections and symmetric real coefficients will have a delay equal to that of 15 sections.

## 2.2. Warped compression system

An effective design for a compressor is to use a side branch for the frequency analysis, with the analysis generating the coefficients of a filter placed in the signal path (Williamson et al. [22]). Either a filter bank or an FFT can be used for the frequency analysis; an FFT-based side-branch system is illustrated in Figure 2. The approximation to auditory frequency analysis is provided by using individual frequency FFT bins at low frequencies, and summing groups of adjacent FFT bins at high frequencies. The signal processing delay for this system is the buffer size plus the delay associated with the FIR filter; additional delay is caused by the  $A/D$  and  $D/A$  converters and code execution time.

The resolution of the frequency analysis performed in the side branch is limited by the size of the FFT and its associated input buffer. For example, when a 32-point FFT is computed, the positive frequency samples can be combined to give nine overlapping frequency bands using individual FFT bins at low frequencies and combining bins at frequencies above 1 kHz. A Blackman window was found to give the best combination of frequency resolution and sidelobe suppression for the 32-point FFT. Increasing the FFT size would give better frequency resolution, but would also increase the system's processing delay due to the larger input buffer size and the longer filter length.

A dynamic-range compression system using warped frequency analysis is presented in Figure 3. The basic design is similar to the side-branch compressor shown in Figure 2. The compressor combines a warped FIR filter and a warped FFT. The same tapped delay line is used for both the

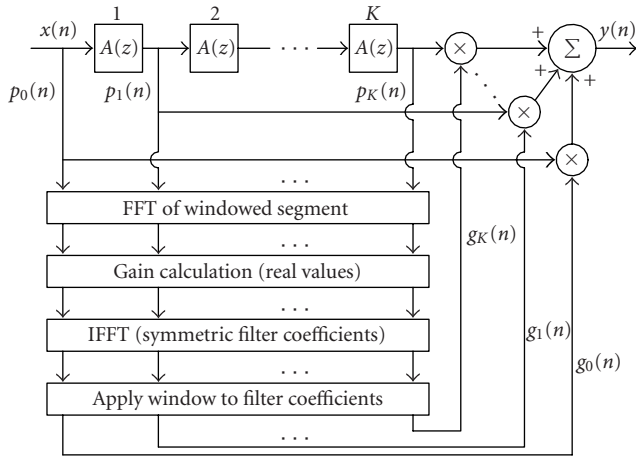


FIGURE 3: Block diagram of a compressor using frequency warping for both frequency analysis and filtered signal synthesis.

frequency analysis and the FIR compression filter. The incoming signal  $x(n)$  is passed through a cascade of first-order all-pass filters of the form given by (1), with the output of the  $k$ th all-pass stage given by  $p_k(n)$ . The sequence of delayed samples  $\{p_k(n)\}$  is then windowed, and an FFT is calculated using the windowed sequence. The result of the FFT is a spectrum sampled at a constant spacing on a Bark frequency scale. The algorithm can be implemented on a sample-by-sample basis or using block data processing. Block processing is typically used with the FFT computed after a block of samples is read in and processed through the cascade of all-pass filters; the compression gains are therefore updated once per block.

Because the data sequence is windowed, the spectrum is smoothed in the warped frequency domain, giving smoothly overlapping frequency bands. The compression gains are then computed from the warped power spectrum for the auditory analysis bands. The compression gains are pure real numbers, so the inverse FFT to give the warped time-domain filter results in a set of filter coefficients that is real and has even symmetry. The system output is then calculated by convolving the delayed samples with the compression gain filter:

$$y(n) = \sum_{k=0}^K g_k(n)p_k(n), \tag{5}$$

where  $\{g_k(n)\}$  are the compression filter coefficients.

In comparison with a conventional FIR system having the same FIR filter length, the warped compression system will require more computational resources because of the all-pass filters in the tapped delay line. However, in many cases the warped FIR filter will be shorter than the conventional FIR filter needed to achieve the same degree of auditory frequency resolution. A nine-band compressor, for example, requires a 31-tap conventional FIR filter but can be realized with a 15-tap warped FIR filter.

### 2.3. Simulation results

Two compression systems were simulated for the performance evaluation. The systems operated at a 16 kHz sampling rate and were simulated in MATLAB using floating-point arithmetic. The first compressor is the side-branch system of Figure 2. For a short system delay, a 16-sample buffer is used for the block time-domain processing, and the signal is processed by a 31-tap FIR filter. The frequency analysis uses a 32-point FFT operating on the present and previous 16-point data segments. A window is used to provide adequate FFT smoothing at low frequencies, and overlapping FFT bins are summed to give the analysis bands at high frequencies. This system has a total of 9 analysis bands, with a low-frequency resolution of 500 Hz. The frequency resolution can be improved by increasing the FFT size, but the system delay will also be increased. The compression gains are calculated in the frequency domain, and the gains inverse transformed to give the symmetric compression filter used to modify the incoming signal.

The second compressor is the warped FIR side-branch system of Figure 3 in which a 16-sample data buffer and a 32-point FFT are used in conjunction with a 31-tap warped FIR filter. This compressor is essentially the frequency-warped version of the side-branch compressor of Figure 2. The input data segment is windowed with a 32-point Hann window, and no frequency-domain smoothing is applied to the spectrum. The compression gains are smoothed by applying a 31-point Hann window to the compression filter after the gain values are transformed into the time domain. This system is termed the Warp-31 compressor.

The Warp-31 compressor provides frequency analysis with a separation of approximately 1.3 Bark. There are a total of 17 bands covering the positive frequencies, including 0 and  $\pi$  radians. The low-frequency bands are approximately spaced at multiples of 135 Hz, with the spacing increasing to 1800 Hz at the highest frequency. The side-branch compressor using the 32-point FFT, on the other hand, uses the output of the FFT to approximate frequency bands on a Bark scale. The limited resolution of the short FFT with its uniform 500 Hz bin spacing causes a poor match between the side-branch frequency bands and the Bark band spacing at low frequencies. At high frequencies, however, FFT bins can be combined to give a reasonably good match. To achieve the same low-frequency resolution as the Warp-31 system, the side-branch compressor requires an FFT size of 128 points which gives a bin spacing of 125 Hz.

The frequency resolution of the Warp-31 system is illustrated in Figure 4. Each curve in the figure represents the warped FFT magnitude frequency response to a steady-state sinusoid at the indicated frequency. The sinusoids were chosen to lie at the center frequencies of 5 of the 17 warped frequency bands. The shapes of the power spectra for the different excitation frequencies are essentially shifted versions of the same basic response. The response at the adjacent frequency band is about 5 dB below the response at the excitation frequency, and the average slope of the response over the first octave is about 50 dB/oct. Replacing the Hann with

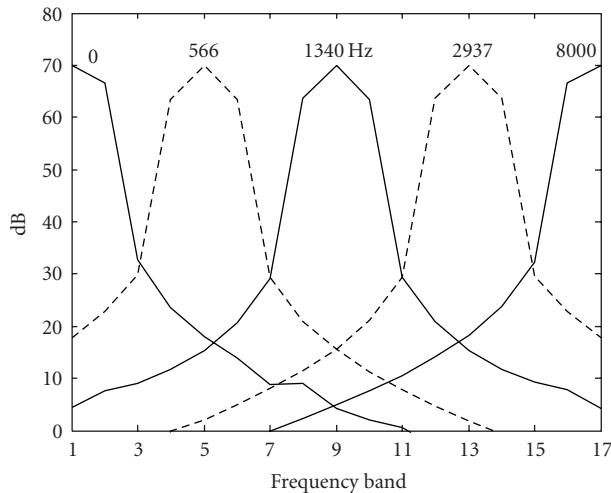


FIGURE 4: Power spectra for the Warp-31 frequency analysis for steady-state sinusoidal excitations at the indicated warped FFT bin center frequencies. The excitation signal is at a level of 70 dB SPL.

a different window shape will modify the spectral response in manner comparable to the effects of the window on a conventional FFT.

The overall system processing group delay is due to several factors. Certain aspects of the overall system delay, such as the *A/D* and *D/A* converter delays, are fixed by the hardware and are not affected by the signal processing. The total software processing delay is the sum of the time required to fill the input buffer, the group delay inherent in the frequency-domain or time-domain filtering operation provided by the compressor, and the time needed to execute the code before the output signal is available.

The side-branch compressor uses a linear-phase FIR filter, so the delay is independent of frequency. The Warp-31 compressor uses all-pass filters to replace the unit delays in the FIR filter implementation, so this system has a frequency-dependent delay. The total delay for the Warp-31 compressor is an estimate assuming that the hardware delays and the time needed for the code execution will be similar to that needed for the side-branch compressor, with an additional allowance for the all-pass filters. The delay values for the 32-point FFT version of the side-branch compressor are based on measurements of an actual hearing aid, and assume 2.5 milliseconds for the hardware and code execution and 1 millisecond for the 16-sample input buffer.

The group delay for the compression systems is plotted in Figure 5. The side-branch system has a constant delay as a function of frequency because of the linear-phase filters used for the processing. The delay is 3.5 milliseconds for an FFT size of 32 points, and increases to 10.5 milliseconds when the FFT size is increased to 128 points. The Warp-31 system has a smooth frequency-dependent delay due to the group-delay characteristics of the all-pass filters used for the warped FIR filtering. The maximum delay for the Warp-31 compressor is 6.1 milliseconds at 0 Hz, with the delay falling to 2.9 mil-

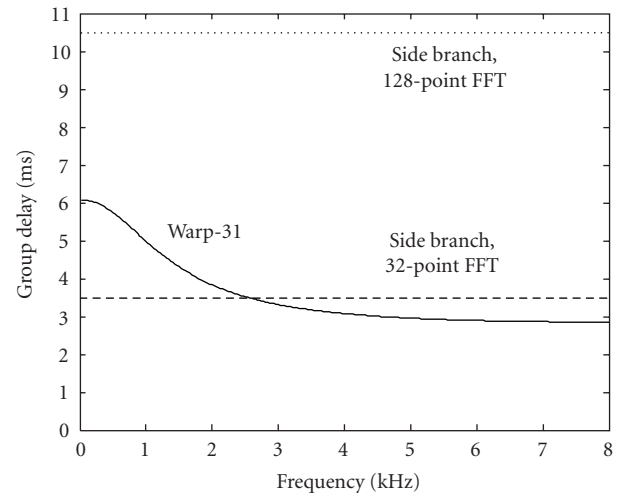


FIGURE 5: Group delay versus frequency for digital compressors based on the side-branch compressor using a 16-sample input buffer and a 32-point FFT (dashed line), 64-sample buffer and a 128-point FFT (dotted line), and the Warp-31 system (solid line).




FIGURE 6: Impulse response for the Warp-31 compressor having a flat frequency response.

liseconds at high frequencies. Thus the Warp-31 compressor has delay characteristics similar to those of the side-branch system with a 32-point FFT, while providing frequency resolution that can only be achieved when a 128-point FFT with its much greater delay is used. The warped compressor thus has substantially reduced delay in comparison with a conventional design having comparable frequency resolution, and the resultant delay in the Warp-31 system would be expected to fall below the threshold of approximately 9 milliseconds for audible interference when the hearing-aid user is talking.

The Warp-31 system has a relative delay of 3.2 milliseconds at low frequencies compared to the delay at high frequencies. The impulse response of the Warp-31 system with a flat frequency response is the same as for a cascade of

TABLE 1: Age, gender, test ear, and audiometric thresholds (dB HL) of listeners with hearing loss (NR means no response).

Subject	Age	Gender	Ear	Frequency (Hz)					
				250	500	1000	2000	4000	8000
I1	50	F	L	70	70	85	80	80	85
I2	61	M	R	25	10	35	60	55	60
I3	64	M	R	30	50	65	70	65	70
I4	38	F	R	25	30	50	60	75	60
I5	75	F	L	50	50	65	60	60	75
I6	78	M	L	50	45	50	65	70	95
I7	44	F	R	15	20	35	55	95	NR
I8	26	F	L	25	30	45	50	50	45
I9	69	F	L	80	75	60	55	40	55
I10	22	F	R	30	35	50	55	55	60
I11	76	M	R	15	30	40	40	60	90

15 all-pass filter sections; the impulse response, shown in Figure 6, illustrates the relative delay between the initial high-frequency output of the filter and the later low-frequency output. In processed speech, the effects of the group delay will be to delay the onset of the first formant relative to the second and third formants. The low-frequency content of bursts, as in stops and plosives, will also be delayed relative to the high-frequency content. Stone and Moore [16] found no significant effect of low-frequency delay on voicing information for hearing-impaired subjects. They found that low-frequency delays greater than 15 milliseconds were significant for manner information (Miller and Nicely [23]), which consists of nasality, affrication, and duration, and delays greater than 9 milliseconds were significant for place information. The 3.2-millisecond relative low-frequency delay of the Warp-31 system should therefore have no measurable effect on speech intelligibility.

### 3. DETECTION OF FREQUENCY-DEPENDENT GROUP DELAY

An important design objective of the frequency-warping system is to determine the optimal number of filter sections that would give effective frequency resolution while minimizing audible delay. As shown in Section 1.2, the overall group delay is short enough that there should be minimal audible interaction between the user's own voice and the delayed sound from the hearing aid. The data of Stone and Moore [14, 15, 16] indicate that there is a threshold of approximately 9 milliseconds for audible interference when the hearing-aid user is talking, and their experiments do not need to be duplicated for the frequency-warped system. However, there is still a question whether the frequency-dependent delay will cause audible timbre or transient effects when listening to speech.

This issue is addressed in the present experiment by determining the minimum boundary for detectible group delay for impulsive sounds, for steady-state sounds, and for continuous speech. Specifically, detection thresholds for frequency-dependent group delay for these stimuli were measured in a

group of normal-hearing listeners and a group of hearing-impaired listeners. The conditions included in this study were ones in which processed sounds are perceived alone, not combined with unprocessed sounds. As such, the conditions studied here are applicable to situations in which a hearing-aid wearer is listening but not talking.

#### 3.1. Method

##### 3.1.1. Listeners

Ten listeners with normal hearing and 11 listeners with hearing loss participated in this study. Listeners with normal hearing had thresholds of 20 dB HL or better at octave frequencies from 250 to 8000 Hz, inclusive. Listeners with hearing loss demonstrated test results consistent with cochlear pathology: normal tympanometry, absence of otoacoustic emissions in regions of threshold loss, and absence of an air-bone gap exceeding 10 dB at two or more frequencies. Listeners with hearing loss had a mild-to-severe hearing loss. Table 1 provides a summary of the characteristics of the listeners with hearing loss, including the audiometric thresholds of the test ear. All listeners were tested monaurally. The right ear was tested in normal-hearing listeners and in hearing-impaired listeners with symmetrical hearing loss. The left ear was tested in some hearing-impaired listeners when the threshold configuration of the left ear allowed for more optimal digital filter design for linear amplification (see below). Listeners were tested individually in a double-walled sound proof booth. Daily test sessions typically lasted one hour but did not exceed two hours. Listeners were compensated 9 dollars per hour for their participation.

##### 3.1.2. Stimuli

Test stimuli included clicks, sentences, and vowels processed to duplicate the delay effects of the frequency-warped system. These stimuli were included in order to assess the perceptual effects of group delay on stimuli ranging from impulsive (clicks) to steady-state (vowels).

Vowels were synthesized using Sensimetrics cascade formant software (Klatt [24]) with a 16,000 Hz sampling rate

TABLE 2: Formant bandwidths and frequencies (Hz) for the vowels /i/ and /a/ for fundamental frequencies of 125 Hz and 200 Hz.

Formant	Bandwidth	/i/	/a/
F1	90 Hz	270 Hz	730 Hz
F2	110 Hz	2290 Hz	1090 Hz
F3	170 Hz	3010 Hz	2440 Hz
F4	250 Hz	3350 Hz	3350 Hz
F5	300 Hz	3850 Hz	3850 Hz




FIGURE 7: Group delay (in milliseconds) is shown as a function of frequency for 8, 15, 30, 60, and 120 filter sections.

and a duration of 1000 milliseconds. Two different vowels (/i/ and /a/) were each generated with two different fundamental frequencies ( $F_0 = 125$  Hz and 200 Hz). Formant frequencies of the vowels, based in part on those published by Peterson and Barney [25], are listed in Table 2. The sentence stimuli were selected from the TIMIT corpus of digitally recorded speech. Specifically, two versions of the same sentence (“Don’t ask me to carry an oily rag like that.”) were included in the stimulus set. The first version was spoken by a male talker and the second version was spoken by a female talker.

### 3.1.3. Signal processing and presentation

The stimuli were processed using a cascade of frequency-warping all-pass filter sections at a 16 kHz sampling rate and with delays ranging from 0 to 150 filter sections in one-section increments. A subset of these delays is shown in Figure 7. After processing with the delay filters, the middle 500-millisecond portions of the vowel stimuli were excerpted. The 500-millisecond excerpts were then used for stimulus presentation in order to assess the steady-state effects of the group delay. Dynamic-range compression was not used in this study because the objective was to determine the threshold of detection for the group delay, and compression would have audibly altered the signal envelopes.

For listener presentation, the digitally stored stimuli went through a digital-to-analog converter (TDT AP2,DD1), a 10 000 Hz antialiasing filter (TDT FT3), an attenuator (TDT PA4), and a headphone buffer (TDT HB6). Finally, the stimuli were presented monaurally to the test ear of each listener through a TDH-49 earphone.

The system for signal presentation described above has its own inherent frequency-dependent group delay due to the antialiasing filter. This inherent group delay will be constant from trial to trial and will not affect the primary objective of this study, namely to determine a listener’s sensitivity to the delay versus frequency characteristics of the warped delay line. Nevertheless, we sought to minimize the absolute delays in the system. After the stimuli were processed with the delay filters, the stimuli were upsampled from 16 000 Hz to 24 000 Hz using linear-phase filters and then passed through a lowpass (cutoff 10 000 Hz) antialiasing filter. Therefore, the dominant system group delay will be due to the TDH-49 headphone cutoff frequency of 6 kHz, which is consistent with the cutoff frequency of hearing aids.

### 3.1.4. Signal levels

#### Normal-hearing listeners

The click stimuli all had the same total power, independent of the number of delay sections. The clicks with no delay had a peak level of 80 dB SPL. Detection thresholds for the click (with no delay) were approximately 26 dB peak SPL for the normal-hearing listeners. Vowels and sentences were all presented at an equalized RMS level corresponding to 65 dB SPL.

#### Hearing-impaired listeners

Stimuli were amplified (through digital linear-phase filtering) for each individual hearing-impaired listener, approximating the linear gain prescribed by the NAL-R fitting procedure (Byrne and Dillon [26]). The input levels to this amplification were as follows: a peak dB SPL of 80 dB for the no-delay click stimuli and 65 dB SPL for the speech stimuli.

### 3.1.5. Test procedure

The just-noticeable delay (JND) was obtained in listeners using a three-interval three-down one-up adaptive procedure (Levitt [27]). Each trial consisted of three 500-millisecond observation intervals with an interstimulus interval of 400 milliseconds. Two of the three intervals on each trial contained a standard stimulus with no frequency-warped group delay (0 cascaded filter sections). One of the three intervals on each trial contained a comparison stimulus with a frequency-warped delay, described in terms of the number of cascaded filter sections. On each trial, the order of presentation of the standard and comparison stimuli was randomized among the three intervals. The listener’s task was to identify the interval with the frequency-warped group delay. Instructions presented to the listeners are provided in the appendix. Feedback was provided after each trial.






FIGURE 8: Thresholds for detection of group delay (expressed as the just-noticeable delay (JND), in terms of number of filter sections) are shown for each of the seven stimulus conditions for NH listeners (open circles) and for HI listeners (filled symbols).

The initial group delay of the comparison stimulus was chosen so as to be above a listener’s delay threshold, as determined from an initial practice session. A large step size (5 filter sections) was used for the initial two turnarounds. A small step size (2 filter sections) was used for the final ten turnarounds. The JND for each adaptive test run was based on the arithmetic mean of the final ten turnarounds. Four estimates of the discrimination threshold were obtained for each listener in each condition. The first estimate in each condition was part of initial practice session. The discrimination thresholds reported below are based on the average of the threshold estimates obtained in the final three test runs. The discrimination threshold is referred to here as the just-noticeable difference (JND) for group delay and is described in terms of number of filter sections.

**3.2. Results and discussion**

Figure 8 shows the just-noticeable delays (JNDs) for normal-hearing listeners and hearing-impaired listeners for each of the seven conditions. Large variability is evident across listeners and across stimulus conditions. For most listeners, detection thresholds were well below 150 filter sections across the seven conditions. However, it is important to note that reliable threshold estimates from the adaptive procedure were not obtainable from some listeners in some of the steady-state vowel conditions: listeners I2 and I6 in the /i/ 125 Hz condition; listeners N5, I1, I2, I6 in the /i/ 200 Hz condition; listener I2 in the /a/ 125 Hz condition, and listeners N5, I4, I6, and I9 in the /a/ 200 Hz condition. Since the adaptive pro-

TABLE 3: Median JND values (in terms of filter sections) for normal-hearing (NH) listeners and hearing-impaired (HI) listeners. The *p* values for between-group comparisons obtained using the Wilcoxon-Mann-Whitney tests are also shown.

Condition	Group		<i>p</i> value
	NH	HI	
Click	9	19	0.072
Sentence (M)	17	23	0.324
Sentence (F)	19	34	0.041
/a/ 125 Hz	34	87	0.003
/a/ 200 Hz	100	107	0.306
/i/ 125 Hz	39	114	0.016
/i/ 200 Hz	132	133	0.273

cedure had a stimulus range of 0 to 150 sections, these listeners were unable to consistently detect group-delay effects of 150 sections or below. Presumably, their thresholds were greater than 150 sections. For the presentation of the data here, these censored thresholds have been assigned a value of 150. Note that the nonparametric analyses described here are robust to the exact censoring value. For example, the same results would be obtained had the censored thresholds been set to a value of 200. Similarly, the medians that we present in the tables do not depend on the censoring value.

Table 3 shows the median just-noticeable delays (JNDs, described in terms of number of filter sections) for normal-hearing (NH) and hearing-impaired (HI) listeners. The median threshold values are the lowest for the click condition, intermediate for the sentences, and the greatest for the steady-state vowels. The right column of Table 3 also shows the *p* values for between-group comparisons obtained using the Wilcoxon-Mann-Whitney tests. While the median JND values are greater in the HI group than in the NH group across all conditions, significant between-group differences were observed in only three of the conditions (/i/ 125 Hz, /a/ 125 Hz, and the sentence spoken by the female talker). The lack of significance between groups in some conditions is consistent with the large intersubject variability and the overlap among listeners across subject groups.

The large variability in the HI group might be due in part to degree of hearing loss. However, while detection thresholds in the group of hearing-impaired listeners were significantly correlated with the pure tone average for click stimulus (0.81, *p* < 0.05), they were not significantly correlated with any of the other stimuli.

The Wilcoxon signed-rank test was used to evaluate whether JND values were significantly different between conditions. All pairs of conditions were included. Table 4 shows the *p* values for this analysis, with adjustments for multiple comparisons (Holm’s method). JND values are significantly different in all but three pairs of stimulus comparisons: click versus male sentence; /i/ 125 Hz versus /a/ 125 Hz; and male sentence versus female sentence. The fact that statistically significant differences are observed between most of pairs of conditions is consistent with the idea that the detection thresholds are stimulus dependent.

TABLE 4:  $p$  values for Wilcoxon signed-rank test for all pairs of conditions, adjusted for multiple comparisons using Holm's method.

	Click	/i/ 125 Hz	/i/ 200 Hz	/a/ 125 Hz	/a/ 200 Hz	Sent. (F)
Click	—	—	—	—	—	—
/i/ 125 Hz	0.0012	—	—	—	—	—
/i/ 200 Hz	0.0012	0.0012	—	—	—	—
/a/ 125 Hz	0.0012	0.1567	0.0012	—	—	—
/a/ 200 Hz	0.0012	0.0287	0.0037	0.0056	—	—
Sent. (F)	0.0425	0.0012	0.0012	0.0056	0.0012	—
Sent. (M)	0.1567	0.0012	0.0012	0.0012	0.0012	0.7677

The differences in detection thresholds across stimulus conditions and across listeners may be partially explained by the differential use of auditory cues in the detection task. For example, listeners might use the delayed onset of low frequencies relative to the high frequencies in the detection of group delay for the click stimulus. Such across-channel onset cues would be unavailable for the vowels, since the vowel stimuli were limited to a steady-state portion of the processed signal. Listeners might be limited to within-channel cues (e.g., changes in the envelope structure with an auditory filter). That is, the group delay may cause different phase relations among harmonics falling in the same auditory filter, resulting in potentially audible changes in frequency modulation and/or amplitude modulation. These within-channel cues would be expected to be most evident for vowels with lower fundamental frequencies, since the more closely spaced harmonics are more likely to interact within an auditory filter. This idea is consistent with the results showing that detection thresholds for vowels with a 125 Hz fundamental frequency are significantly better than for the vowels with a 200 Hz fundamental frequency. The use of within- and across-channel cues would also be expected to differ in hearing-impaired listeners, given the assumption of broader auditory filters in listeners with cochlear hearing loss (Moore [28]). Further speculation regarding the possible mechanisms underlying the detection thresholds is limited by the experimental design. Specific auditory cues were not parametrically varied as a function of the number of all-pass filters. As such, listeners may have used multiple auditory cues in the detection task, some of which may not have been consistently available (e.g., the phase relations among the vowel harmonics may not have changed monotonically as a function of the number of all-pass filters).

The primary goal of the perceptual study was to determine the minimum detectable boundary for a wide range of stimuli processed with the warped compression algorithm in order to guide its implementation. The objective is to determine the optimal number of filter sections that would give effective frequency resolution while minimizing audible delay.

Figure 9 shows the cumulative distribution functions for the group of normal-hearing listeners (Figure 9b) and the group of hearing-impaired listeners (Figure 9a). The cumulative distribution functions show estimates of the probability ( $F_n(x)$ ) of a given JND value being detectible.  $F_n(x) = 0.5$

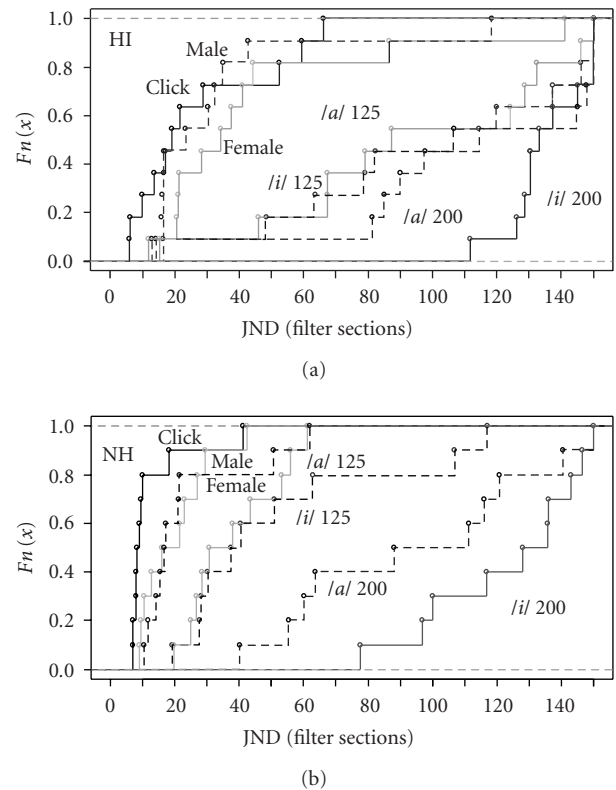


FIGURE 9: The cumulative distribution function for the group of normal-hearing listeners (b) and the group of hearing-impaired listeners (a) is shown for each of the seven stimulus conditions. The cumulative distribution function shows the estimate of the probability ( $F_n(x)$ ) that a particular JND value is detectible.  $F_n(x) = 0.5$  is the median JND threshold for each listener group in each condition.

is the median JND threshold for each group (NH and HI) in each of the seven conditions.

The design of the warped compressor involves a trade-off between the frequency resolution and the group delay. An additional practical concern is the computational load for implementing the warped compressor in the hearing-aid digital processor. The number of multiply adds per second scales directly as the number of all-pass filter sections, and a practical maximum is 31 all-pass filter sections combined with a 32-point FFT for the frequency analysis.

The warped compressor using 31 all-pass filter sections has a group delay equivalent to 15 sections because of the filter-coefficient symmetry. To address the issue of the audibility of the delay associated with a warped compressor with 15 sections, it is helpful to consider an estimate of the probability of the JND = 15 sections being at or above the detection threshold. In both listener groups, the estimate of the probability of listeners being able to detect a frequency-dependent group delay of 15 sections is greater than 0.4 for the click stimulus. However, the probability that listeners in either group will be able to detect the group delay of 15 sections decreases substantially for all other stimulus conditions. Thus, a warped compressor using 31 all-pass sections should give a system with inaudible delay under nearly all listening conditions.

#### 4. CONCLUSIONS

Frequency warping offers definite advantages in designing a digital compressor. The warped frequency scale gives a much better match to auditory perception than the uniform frequency scale inherent in conventional digital signal processing. With the appropriate choice of the warping parameter, the warped FFT bins closely approximate a Bark frequency scale. Furthermore, using a compression filter having even symmetry guarantees that the group delay does not depend on the instantaneous compression gains, and thus removes any phase modulation that could occur as the gains change in response to the incoming signal and ensures that localization phase cues are preserved in a binaural fitting.

The simulation results showed that a 31-tap warped FIR filter, using a 32-point warped FFT, gave frequency analysis results comparable to those from a 128-point conventional FFT. Frequency warping allows a substantial reduction in the system order when compared to a conventional FIR filter giving similar frequency resolution, requiring less than half compression-filter length. The primary disadvantage in implementing frequency warping is the computational cost of replacing unit delays with first-order all-pass filters. The all-pass filters could conceivably double the computational time needed to implement an FIR filter. However, since a filter only half as long is needed for performance equivalent to a conventional FIR compression filter, the net cost should be minimal.

The frequency-warped compressor introduces a frequency-dependent group-delay characteristic. Detection thresholds for the delay variation with frequency were obtained for a group of normal-hearing subjects and for a group of hearing-impaired subjects. The subject tests indicated that the median detection threshold for the frequency-dependent group delay is the lowest for click stimuli, the highest for steady-state vowels, and intermediate for speech. Normal-hearing subjects had lower thresholds on average than the hearing-impaired subjects, but there was a large intersubject variability.

The detection thresholds obtained in this study provide insights into the optimal number of filter sections in

a frequency-warping system that would give effective frequency resolution while minimizing audible delay. A warped compressor design using 31 all-pass filter sections gives a delay equivalent to 15 sections in cascade when symmetric filter-coefficients are used. The maximum delay in a practical system is just over 6 milliseconds, which is comfortably below the threshold of approximately 9 milliseconds found for audible interference when the hearing-aid user is talking. The results reported in this paper show that the frequency-dependent group delay produced by 15 sections is inaudible for most listeners for the click stimuli and inaudible for all listeners for steady-state speech sounds. Thus, a warped compressor using 31 all-pass sections should give a system with inaudible delay under nearly all listening conditions.

US, European, and other patent applications have been filed on the signal processing described in this paper (Kates [29]).

#### APPENDIX

##### INSTRUCTIONS TO LISTENERS

Listeners were presented with the following instructions. Signal processing in hearing aids can help improve what we listen to. Sometimes the signal processing can also make speech sound different. We are studying how perceptible these differences are. Throughout this study, you will be hearing different kinds of sounds. These sounds include (1) clicks, (2) the vowel sound “ah” as in “hot,” (3) the vowel sound “ee” as in “heed,” and (4) the sentence “Don’t ask me to carry an oily rag like that.” During any given listening set, you will hear the same kind of sounds. On each trial, you will hear three sounds in a row. Your task is to pick the one sound (1, 2, or 3) that sounds different from the other two sounds. You will need to wait until all three sounds have played out before pressing the appropriate button.

#### ACKNOWLEDGMENTS

The second author is supported by a grant from the Whitaker Foundation. The authors thank Jessica Rossi-Katz, Kim Howry, Michelle Itano, and Ajay Natarajan for their assistance in data collection and analysis. The authors extend thanks to Reinier Kortekaas for providing helpful comments on an earlier draft of this paper.

#### REFERENCES

- [1] J. M. Kates, “Signal processing for hearing aids,” in *Applications of Digital Signal Processing to Audio and Acoustics*, M. Kahrs and K. Brandenburg, Eds., Kluwer Academic, Norwell, Mass, USA, pp. 235–277, March 1998.
- [2] H. Dillon, R. Katsch, D. Byrne, T. Ching, G. Keidser, and S. Brewer, “The NAL-NL1 prescription procedure for non-linear hearing aids,” Annual Rep. 1997/98, pp. 4–7, National Acoustics Laboratories Research and Development, Chatswood, NSW, Australia, 1998.
- [3] B. C. J. Moore and B. R. Glasberg, “Suggested formulae for calculating auditory-filter bandwidths and excitation

- patterns," *Journal of the Acoustical Society of America*, vol. 74, no. 3, pp. 750–753, 1983.
- [4] E. Zwicker and E. Terhardt, "Analytical expressions for critical-band rate and critical bandwidth as a function of frequency," *Journal of the Acoustical Society of America*, vol. 68, no. 5, pp. 1523–1525, 1980.
- [5] A. V. Oppenheim, D. H. Johnson, and K. Steiglitz, "Computation of spectra with unequal resolution using the fast Fourier transform," *Proc. IEEE*, vol. 59, no. 2, pp. 299–300, 1971.
- [6] A. V. Oppenheim and D. H. Johnson, "Discrete representation of signals," *Proc. IEEE*, vol. 60, no. 6, pp. 681–691, 1972.
- [7] J. O. Smith III and J. S. Abel, "Bark and ERB bilinear transforms," *IEEE Trans. Speech Audio Processing*, vol. 7, no. 6, pp. 697–708, 1999.
- [8] A. Härmä, M. Karjalainen, L. Savioja, V. Välimäki, U. K. Laine, and J. Huopaniemi, "Frequency-warped signal processing for audio applications," *Journal of the Audio Engineering Society*, vol. 48, no. 11, pp. 1011–1031, 2000.
- [9] M. Karjalainen, E. Piirilä, A. Järvinen, and J. Huopaniemi, "Comparison of loudspeaker equalization methods, based on DSP techniques," *Journal of the Audio Engineering Society*, vol. 47, no. 1-2, pp. 14–31, 1999.
- [10] M. Karjalainen, A. Härmä, and U. K. Laine, "Realizable warped IIR filters and their properties," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP '97)*, vol. 3, pp. 2205–2208, Munich, Germany, April 1997.
- [11] H. W. Strube, "Linear prediction on a warped frequency scale," *Journal of the Acoustical Society of America*, vol. 68, no. 4, pp. 1071–1076, 1980.
- [12] M. Karjalainen, A. Härmä, U. K. Laine, and J. Huopaniemi, "Warped filters and their audio applications," in *Proc. IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA '97)*, New Paltz, NY, USA, October 1997, session 11, paper 2, 4 pages.
- [13] J. Agnew and J. M. Thornton, "Just noticeable and objectionable group delays in digital hearing aids," *Journal of the American Academy of Audiology*, vol. 11, no. 6, pp. 330–336, 2000.
- [14] M. A. Stone and B. C. J. Moore, "Tolerable hearing aids delays. II: Estimation of limits imposed during speech production," *Ear and Hearing*, vol. 23, no. 4, pp. 325–338, 2002.
- [15] M. A. Stone and B. C. J. Moore, "Tolerable hearing aids delays. I: Estimation of limits imposed by the auditory path alone using simulated hearing losses," *Ear and Hearing*, vol. 20, no. 3, pp. 182–192, 1999.
- [16] M. A. Stone and B. C. J. Moore, "Tolerable hearing-aid delays. III. Effects on speech production and perception of across-frequency variation in delay," *Ear and Hearing*, vol. 24, no. 2, pp. 175–183, 2003.
- [17] J. Blauert and P. Laws, "Group delay distortions in electroacoustical systems," *Journal of the Acoustical Society of America*, vol. 63, no. 5, pp. 1478–1483, 1978.
- [18] D. M. Green, "Temporal acuity as a function of frequency," *Journal of the Acoustical Society of America*, vol. 54, no. 2, pp. 373–379, 1973.
- [19] H. Banno, K. Takeda, and F. Itakura, "The effect of group delay spectrum on timbre," *Acoustic Science and Technology*, vol. 23, no. 2, pp. 113–116, 2002.
- [20] W. H. Greer, *Monaural sensitivity to dispersion in impulses and speech*, Ph.D. thesis, University of Utah, Salt Lake City, Utah, USA, 1975.
- [21] T. Arai and S. Greenberg, "Speech intelligibility in the presence of cross-channel spectral asynchrony," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP '98)*, vol. 2, pp. 933–936, Seattle, Wash, USA, May 1998.
- [22] M. J. Williamson, K. L. Cummins, and K. E. Hecox, "Adaptive programmable signal processing and filtering for hearing aids," US Patent 5,027,410, June 1991.
- [23] G. A. Miller and P. E. Nicely, "An analysis of perceptual confusions among some English consonants," *Journal of the Acoustical Society of America*, vol. 27, no. 2, pp. 338–352, 1955.
- [24] D. H. Klatt, "Software for a cascade/parallel formant synthesizer," *Journal of the Acoustical Society of America*, vol. 67, no. 3, pp. 971–995, 1980.
- [25] G. E. Peterson and H. L. Barney, "Control methods used in a study of the vowels," *Journal of the Acoustical Society of America*, vol. 24, no. 2, pp. 175–184, 1952.
- [26] D. Byrne and H. Dillon, "The National Acoustic Laboratories' (NAL) new procedure for selecting the gain and frequency response of a hearing aid," *Ear and Hearing*, vol. 7, no. 4, pp. 257–265, 1986.
- [27] H. Levitt, "Transformed up-down methods in psychophysics," *Journal of the Acoustical Society of America*, vol. 49, no. 2, pp. 467–477, 1971.
- [28] B. C. J. Moore, *Cochlear Hearing Loss*, Whurr Publishers, London, UK, 1998.
- [29] J. M. Kates, "Dynamic range compression using digital frequency warping," US Patent Application 20030081804, May 2003.

**James M. Kates** was born in Brookline, Mass, in 1948. He received the degrees of B.S.E.E. and M.S.E.E. from the Massachusetts Institute of Technology (MIT) in 1971, and the professional degree of Electrical Engineer from MIT in 1972. He currently is a Senior Research Engineer in the Algorithm Development Group of GN ReSound. He has an office in Boulder, Colo, where he is involved in research and development of digital signal processing for hearing aids. He is also an Adjunct Professor in the Department of Speech, Language & Hearing Sciences at the University of Colorado at Boulder, where he conducts research in auditory perception, hearing loss, and signal processing for hearing aids. Prior to moving to Boulder, he was with the Center for Research in Speech and Hearing Sciences of the City University of New York, where he conducted research in hearing aids and auditory perception. Previous positions include research and development for hearing aids (Cirrus Logic, AudioLogic, and Siemens Hearing Instruments), signal processing for radar, speech, and hearing applications (SIGNATRON Inc.), and loudspeaker design and signal processing for audio applications (Acoustic Research and CBS Laboratories).



**Kathryn Hoberg Arehart** received her B.S. degree in biological sciences from Stanford University in 1984. She received an M.S. degree in 1987 and a Ph.D. degree in 1992, both in speech and hearing sciences from the University of Washington in Seattle. She also has clinical certification in audiology from the American Speech-Language-Hearing Association. In 1992, she joined the faculty of the Department of Speech, Language & Hearing Sciences at the University of Colorado at Boulder, where she now is an Associate Professor. Her research interests include auditory perception by listeners with cochlear hearing loss and design and evaluation of signal processing algorithms for hearing aids.

