

Received April 5, 2022, accepted May 2, 2022, date of publication May 20, 2022, date of current version May 26, 2022.

Digital Object Identifier 10.1109/ACCESS.2022.3176368

Real-Time Multirate Multiband Amplification for Hearing Aids

ALICE SOKOLOVA^{1,2}, (Graduate Student Member, IEEE), DHIMAN SENGUPTA³,
MARTIN HUNT¹, RAJESH GUPTA^{3,4}, (Fellow, IEEE), BARIS AKSANLI^{1,2}, (Member, IEEE),
FREDRIC HARRIS¹, (Senior Member, IEEE), AND HARINATH GARUDADRI⁵, (Member, IEEE)

¹Department of Electrical and Computer Engineering, UC San Diego, La Jolla, CA 92093, USA

²Department of Electrical and Computer Engineering, San Diego State University, San Diego, CA 92182, USA

³Department of Computer Science and Engineering, UC San Diego, La Jolla, CA 92093, USA

⁴Halıcıoğlu Data Science Institute, La Jolla, CA 92093, USA

⁵Qualcomm Institute of Calit2, UC San Diego, La Jolla, CA 92093, USA

Corresponding author: Alice Sokolova (aasokolo@ucsd.edu)

This work was supported in part by the National Institute of Health, NIH/National Institute on Deafness and Other Communication Disorders (NIDCD), under Grant R21DC015046 and Grant R33DC015046 (Self-fitting of Amplification: Methodology and Candidacy); in part by the National Institute of Health, NIH/NIDCD, under Grant R01DC015436 (A Real-time, Open, Portable, Extensible Speech Laboratory); in part by the Division of Information and Intelligent Systems, University of California San Diego, under Grant IIS-1838897 (A Framework for Optimizing Hearing Aids In Situ Based on Patient Feedback, Auditory Context, and Audiologist Input); in part by the Qualcomm Institute, University of California San Diego (UCSD); in part by the Halıcıoğlu Data Science Institute, UCSD; and in part by the Wrethinking, the Foundation.

ABSTRACT Hearing loss is a common problem affecting the quality of life for thousands of people. However, many individuals with hearing loss are dissatisfied with the quality of modern hearing aids. Amplification is the main method of compensating for hearing loss in modern hearing aids. One common amplification technique is dynamic range compression, which maps audio signals onto a person's hearing range using an amplification curve. However, due to the frequency dependent nature of the human cochlea, compression is often performed independently in different frequency bands. This paper presents a real-time multirate multiband amplification system for hearing aids, which includes a multirate channelizer for separating an audio signal into eleven standard audiometric frequency bands, and an automatic gain control system for accurate control of the steady state and dynamic behavior of audio compression as specified by ANSI standards. The spectral channelizer offers high frequency resolution with low latency of 5.4 ms and about 14 \times improvement in complexity over a baseline design. Our automatic gain control includes a closed-form solution for satisfying any designated attack and release times for any desired compression parameters. The increased frequency resolution and precise gain adjustment allow our system to more accurately fulfill audiometric hearing aid prescriptions.

INDEX TERMS Hearing aids, digital signal processing, auditory system, channelization, wearable computers, speech processing, open source hardware, real-time systems, embedded software, research initiatives.

I. INTRODUCTION

Studies have shown that only about one-third of individuals who have hearing loss utilize a hearing aid. Among those individuals, around one-third do not use their hearing aids regularly. The main reason for this disuse is often the dissatisfaction with the speech quality offered by modern hearing aids, especially in noisy environments where hearing-impaired individuals need them the most [1].

The associate editor coordinating the review of this manuscript and approving it for publication was Tianhua Xu .

Achieving music appreciation with hearing aids is an even greater challenge [2].

One highly effective approach for improving the audibility of sound for hearing impaired users is called Wide Dynamic Range Compression (WDRC), which is the amplification and reduction of the dynamic range, or volume swing, of an audio signal. WDRC involves amplifying quiet signals to improve audibility, and simultaneously decreasing the volume of loud signals to reduce discomfort to a hearing-impaired user.

Human hearing, however, is inherently frequency-dependent. The human cochlea perceives finer pitch variation at lower frequencies than at higher frequencies. Additionally,

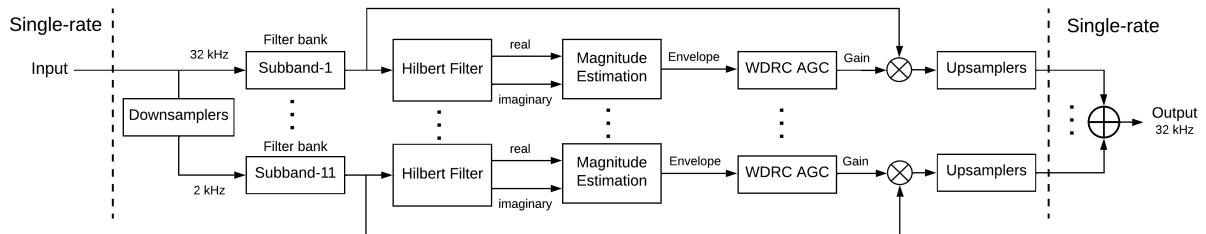


FIGURE 1. A block diagram of the multirate multiband hearing aid amplification system. A 32 kHz input signal is separated into eleven frequency channels with different sampling rates. Each channel is then individually processed. Lastly, all bands are brought back to the original sampling rate and combined to create the output signal for real-time playback.

hearing loss is also typically frequency dependent, affecting certain frequency ranges more than others. For this reason, the compression gains needed to compensate for hearing loss vary across different frequency bands, necessitating a multiband approach to WDRC. Studies have shown that a greater number of frequency bins increases researchers' flexibility, especially for unusual hearing loss patterns [3].

In this paper, we present a Real-time Multirate Multiband Amplification system, which addresses the need for finer, more precise gain control in a hearing aid device. Our design provides the audiology research community with tools which offer higher flexibility and accuracy than currently available on open-source platforms. The system consists of:

- 1) A Multirate Audiometric Filter Bank, offering highly accurate low-latency subband decomposition which can be used for a variety of hearing enhancement algorithms. In this paper, we present a half-octave realization, centered at the standard audiometric frequencies of 250, 375, 500, 1000, 1500, 2000, 3000, 4000, 6000, and 8000 Hz.
- 2) A Multirate Automatic Gain Control system for WDRC that accurately fulfills the static and dynamic properties specified by audiologists, which include steady state gain, as well as attack and release times.

The block diagram in Figure 1 shows an overview of the proposed subband amplification system. This system accepts an audio signal sampled at 32 kHz, performs frequency decomposition on the signal, and transitions from single to multirate processing. The system then computes the gains necessary for Wide Dynamic Range Compression in each band. The final stage converts all multirate outputs back to the original sampling rate and combines the bands into a final output. Multirate processing is a key feature of our design, and is instrumental in ensuring real-time operation of the system and reducing power consumption.

The multirate amplification system is implemented and tested on the Open Speech Platform (OSP) – an open source suite of software and hardware tools for performing research on emerging hearing aids and hearables. The OSP suite includes a wearable hearing aid, a wireless interface, and a set of hearing enhancement algorithms [4]–[7].

II. FILTER BANK

A. OVERVIEW

Figure 2 shows the Eleven Band Multirate Filter Bank, also known as a channelizer, for subband decomposition. Subband decomposition is the process of separating a signal into multiple frequency bands or channels, and is used in many applications, including hearing aids [8]–[11]. The multirate filter bank possesses the following properties:

a: CENTER FREQUENCIES

The structure of an audiometric filter bank reflects the spectral nature of the human cochlea, which is inherently logarithmic. The American Speech-Language-Hearing Association (ASHA) defines a set of ten audiometric frequencies used for pure-tone audiometry, which are 0.25, 0.5, 1, 1.5, 2, 3, 4, 6, and 8 kHz [12]. These frequencies closely resemble a half-octave logarithmic sequence, and are commonly targeted for audiometric filter banks. However, every other frequency is not a true half-octave frequency, but rather a simplified integer approximation. The audiometric filter bank is a true half-octave channelizer, making it uniformly distributed on the logarithmic scale, as seen from Figure 2a. It spans a range of 0.25 to 8 kHz, which produces eleven bands. Although the true half octave center frequencies diverge from the rounded ASHA approximations, they are functionally the same, and for the sake of simplicity we will be referring to each individual band by its approximate audiometric frequency.

b: ATTENUATION AND RIPPLE

The American National Standards Institute (ANSI) defines specifications for Half-Octave Acoustic filters [13]. The standard includes three classes of filters – class 0, 1, and 2, where class 0 has the strictest tolerances and class 2 has the most lax tolerances. The filter bank meets class 0 standards – the highest of the three. Accordingly, each band of the filter bank has -75 dB sidelobe attenuation, and the in-band ripple is within ± 0.15 dB. The ripple of the composite response of the channelizer is also within ± 0.15 dB.

c: FILTER SHAPE AND COMPOSITE RESPONSE

Figure 2 shows the audiometric filter bank on both the logarithmic and the linear scale. As seen from Figure 2, filters which are symmetrical on the logarithmic scale are

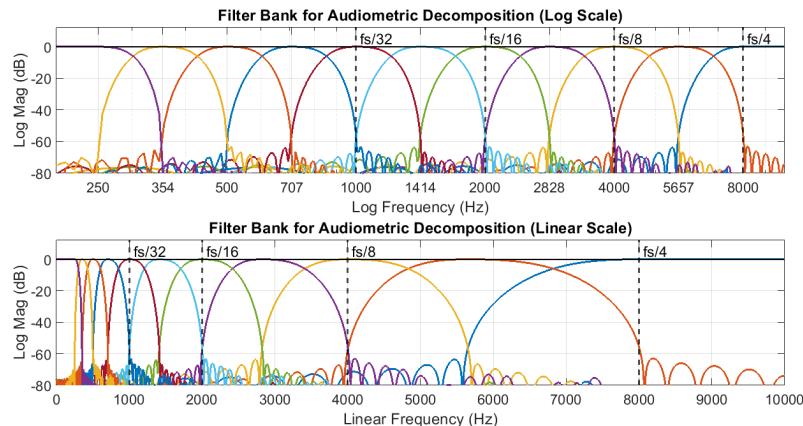


FIGURE 2. The magnitude response and composite response of the proposed multirate filter bank, shown on the logarithmic scale (top) and linear scale (bottom). Vertical dotted lines represent different sampling rates used in the filter bank.

asymmetrical on the linear scale. We designed asymmetrical bandpass filters by convolving a lowpass and a highpass filter for each band.

A more difficult challenge, though, is achieving signal reconstruction. A filter bank has perfect reconstruction if the sum of all outputs is equal to the original input signal. In the frequency domain, this means the composite frequency response of the filter bank is a flat line spanning all frequencies, as shown in Figure 2.

We ensure that our filter bank has perfect reconstruction by employing complementary filter design. Complementary filters are two filters the sum of which is an all-pass filter. For any highpass or lowpass filter, its complement can be found by subtracting it from an all-pass filter, which is simply an impulse in the time domain. We designed all neighboring filter edges to be complements of each other, ensuring that their sum is an all-pass filter, which guarantees signal reconstruction. The channelizer offers perfect reconstruction within ± 0.15 dB.

d: MULTIRATE PROCESSING AND LATENCY

It is well known in the signal processing community that the sharper a digital filter is, the more coefficients it requires. As seen from Figure 2, the audiometric channelizer requires very narrow and sharp filters – the lowest center frequency (0.25 kHz) is 32× smaller than the highest center frequency (8 kHz), and at a 32 kHz sampling rate, the width of the narrowest filter is only 1/64 of the entire signal bandwidth. A conventional implementation of such narrow filters would result in too much latency to meet real-time processing deadlines, and would require excessive processing power.

Our filter bank dramatically reduces both power consumption and latency by employing multirate signal processing. Compared to a single-rate implementation, multirate processing reduces the power consumption by a factor of 13.7×, and reduces latency from 32 ms down to 5.4 ms.

B. MULTIRATE SIGNAL PROCESSING

The motivation behind multirate processing is to decrease the complexity of a filter by reducing the sampling rate. Table 1 lists the number of taps needed to implement the filters shown in Figure 2 at a single sampling rate of 32 kHz. As the filters becomes narrower and sharper, they require an exponentially increasing number of taps, reaching impractical values at the lowest frequencies.

However, the complexity of a filter can be decreased by reducing the sampling rate. For a given bandpass filter, the relative bandwidth is narrower at a higher sampling rate and wider at a lower sampling rate. Thus, a filter spanning a fixed range of frequencies becomes relatively wider as the sampling rate decreases. As the relative filter bandwidth increases, the numbers of taps proportionately decreases. For example, when the sampling rate of a filter is decreased by half, the relative bandwidth of the filter doubles, and the number of taps needed to implement it is also halved.

We exploit the unique structure of the audiometric filter bank to map each frequency octave to a sampling rate. The audiometric channelizer is a half-octave filter bank spanning a frequency range of about 5 octaves, from 250 Hz to 8000 Hz. An octave is a logarithmic unit defined as the difference between two frequencies separated by a factor of two, and a half-octave is the difference between two frequencies separated by a factor of $\sqrt{2}$. Thus, a half-octave filter bank is binary logarithmic, and the bandwidth of any two filters an octave apart differs by a factor of two.

As such, we are able to map each octave of the channelizer to a different sampling rate. We start by designing two bandpass filters at the original sampling rate that span one octave. The next two filters are one octave below, are half as wide, and would require double the number of taps. However, if we lower the sampling rate of the lower octave, the number of taps would decrease by half, resulting in filters of the same length as the ones we started with. Following this pattern, we are able to design all the filters in the audiometric

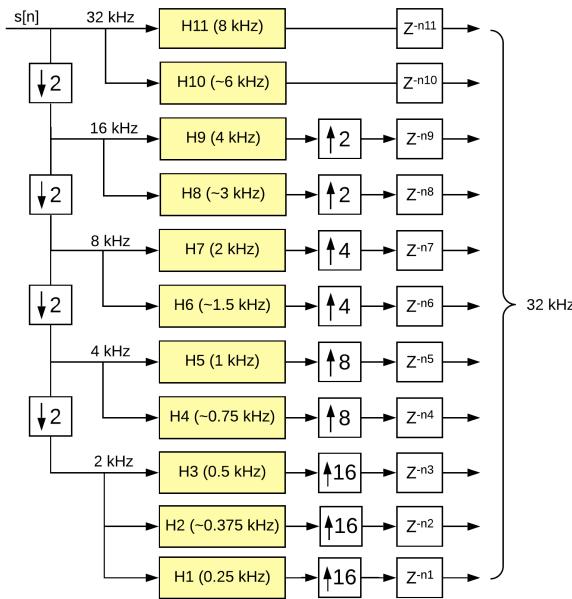


FIGURE 3. Multirate filter bank structure: An input signal is split into five sampling rates and is passed through eleven bandpass filters. Each channel is then restored to the original sampling rate. Delays following the upsampleers are used to compensate for latency disparity among bands.

TABLE 1. A comparison of the number of coefficients in each filter of the Audiometric Filter Bank, with and without multirate processing.

Filter Band:	Filter Taps		Sampling rate
	Single-rate	Multirate	
8 kHz	53	53	1
6 kHz	77	77	1
4 kHz	154	77	1/2
3 kHz	154	77	1/2
2 kHz	308	77	1/4
1.5 kHz	308	77	1/4
1 kHz	616	77	1/8
0.75 kHz	616	77	1/8
0.5 kHz	1232	77	1/16
0.375 kHz	1232	77	1/16
0.25 kHz	1232	77	1/16

channelizer using the same number of coefficients for each filter.

Table 1 compares a single-rate versus a multirate implementation of the channelizer. In the single-rate case, as the bandwidth of the filters is halved for every octave, the number of filter coefficients doubles for every octave. However, in the multirate implementation, we retain constant filter complexity because the decrease in a filter's bandwidth is compensated by a decrease in the sampling rate. (The 8 kHz band is an exception because it is a highpass rather than a bandpass filter.)

Figure 3 shows the conceptual block diagram of the audiometric filter bank. First the input signal is separated into different sampling rates using downsamplers. Then the inputs are passed through the bandpass filters. Lastly, the outputs are brought back to the original sampling rate using upsampleers.

The five different sampling rates used in the channelizer are represented with dotted vertical lines in Figure 2.

According to the Nyquist Theorem, for any given sampling rate f_s , the only frequencies that can be observed are those lying between $-f_s/2$ and $+f_s/2$. Thus, each line represents the frequency limit of each different sampling rate. For the purposes of space, however, the original sampling rate, spanning $-f_s/2$ to $+f_s/2$, is not explicitly shown in Figure 2.

According to the Nyquist theorem, any frequency band which lies to the left of a dotted line can be processed at that respective sampling rate without aliasing distortion. However, resamplers are not ideal, and require constraints on overlapping transition bandwidths.

C. RESAMPLING

Conventionally, downsampling is performed by passing a signal through an antialiasing filter, and then decimating it. Similarly, conventional upsampling is performed by zero-padding a signal, and then passing it through an interpolating filter. As such, the complexity of conventional resamplers strongly depends on their resampling ratio – a high-ratio downsample would require a sharp antialiasing filter to remove all unwanted frequencies, and a high-ratio upsample would require a sharp interpolating filter to remove spectral signal copies. As before, sharp antialiasing and interpolating filters would require many taps, negating the power and latency benefits of multirate processing.

We combat this issue by performing resampling in multiple stages. Since all of our resamplers are multiples of two, we cascade multiple 1:2 or 2:1 resamplers to achieve the desired resampling ratio. 1:2 and 2:1 resamplers require only a short half-band filter for anti-aliasing and interpolating, which allows us to achieve high reductions of complexity.

Figure 4 compares a single-stage and a cascaded implementation of a 1:8 upsample. A 1/8 band filter suitable for this resampler would require about 261 taps. The number of multiply-and-add operations, equal to the frame size multiplied by the number of filter coefficients, would equal to 8352 operations per 32-sample output frame. However, this upsample can be split into three 1:2 upsample stages, each containing a half-band filter, and after each upsampling stage, the transition bandwidth of the interpolating filter can be increased, which reduces complexity. As such, a cascaded 1:8 upsample requires only 680 multiply-and-add operations.

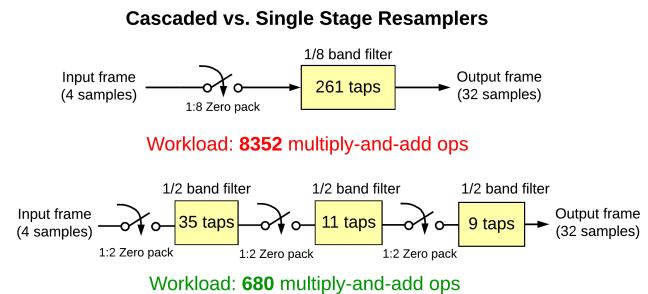


FIGURE 4. A comparison between a conventional single-stage upsample and an equivalent cascaded upsample. It is more efficient to split a single-stage high-ratio resampler into multiple stages.

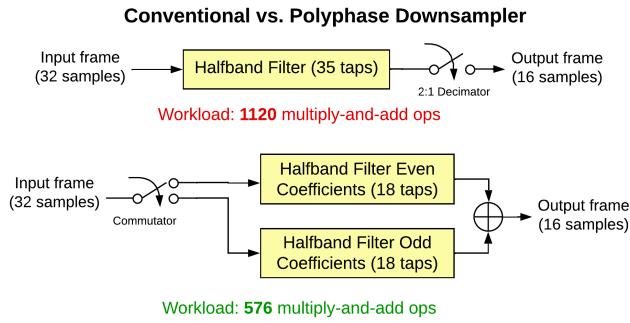


FIGURE 5. A comparison between a conventional 1:2 upsample and an equivalent polyphase implementation. Converting conventional resamplers into polyphase resamplers reduces complexity by about factor of 2.

TABLE 2. The cumulative number of multiply-and-accumulate operations per sample of the audiometric filter bank, with and without multirate processing.

Filter Band:	Operations per sample:		Ratio:
	Single-rate	Multirate	
8 kHz	53	53	1x
6 kHz	77	77	1x
4 kHz	154	74.5	2.07x
3 kHz	154	56.5	2.73x
2 kHz	308	43.25	7.12x
1.5 kHz	308	34.25	8.99x
1 kHz	616	26.63	23.14x
0.75 kHz	616	22.13	27.84x
0.5 kHz	1232	18.31	67.28x
0.375 kHz	1232	16.06	76.7x
0.25 kHz	1232	16.06	76.7x
Total:	5982	437.69	13.67x

We further reduce the complexity of the resamplers by employing polyphase filtering [14]. Conventional resamplers perform many redundant computations, such as computing samples which will be discarded, or computing samples which are known to be zero. Polyphase filtering eliminates these redundant computations by splitting a single filter into multiple paths and employing the Noble identity to rearrange filtering and resampling. Figure 5 compares a conventional and a polyphase 2:1 downampler. Polyphase resamplers always perform filtering at the lower of their input/output rate, and reduce the complexity of resampling by approximately a factor of M , where M is the resampling ratio.

D. POWER

We estimate the cumulative power consumption of the filter bank by computing the total number of multiply-and-accumulate operations per one output sample. For a filter running at a single sampling rate, the number of operations per sample is simply equal to the number of filter taps. However, in a multirate system, samples are continuously removed and added, which makes it impossible to match an input sample to a single output sample. As such, we compute the number of operations per sample of the multirate channelizer by calculating the total number of operations per input frame, and then normalizing by the input frame size. For each stage of the filter bank, we track the current frame size and the

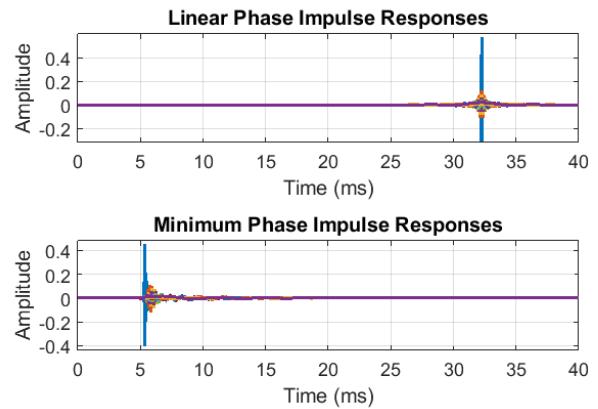


FIGURE 6. A comparison between a linear phase implementation (top) and a minimum phase implementation (bottom) of the Audiometric Filter Bank. The linear phase filter bank has about 32 ms of delay, while the minimum phase implementation has only 5.4 ms of delay.

cumulative operations count. Due to the multirate structure of the channelizer, normalization by frame size results in a fractional number of operations per sample.

Table 2 compares the total number of multiply-and-accumulate operations per sample for a single-rate and multirate implementation of the channelizer. The multirate operations estimate accounts for all filters and resamplers. Our evaluations show that compared to a conventional approach, the multirate filter bank offers 13.7x improvement in complexity. For a wearable battery-operated system, power consumption and processing capabilities are of critical importance. Reducing the number of operations improves battery-life and frees processing power for other tasks.

E. LATENCY

As seen from Figure 3, different frequency bands follow different signal paths and as such, experience varying amounts of delay. Because of the resamplers and lower sampling rates, lower frequency bands incur more delay than higher frequencies. The highest frequency bands (8 kHz and 6 kHz) experience only a few milliseconds of delay. However, the 0.5 kHz, 0.375 kHz, and the 0.25 kHz bands experience over 30 milliseconds of latency. This disparity causes a phase offset among the eleven bands, and causes distortion in the composite frequency response. To certain listeners, this phase disparity sounds like an echo or a distorted sound timbre.

In order to eliminate this latency disparity, we realign the bands by inserting delays into the signals paths, as seen in Figure 3, such that higher frequency bands are delayed until the lowest frequency bands arrive. Figure 6 (top) shows the aligned impulse responses of the filter bank. Although the solution above preserves perfect reconstruction, the latency far exceeds real-time operation requirements. Conventionally, the latency limit for a real-time hearing aid is considered to be 10 milliseconds [15]. As seen from Figure 6 (top), the latency of the aligned channelizer is about 32 milliseconds. We resolve this issue by converting the filters from linear phase to minimum phase. A minimum phase filter has the

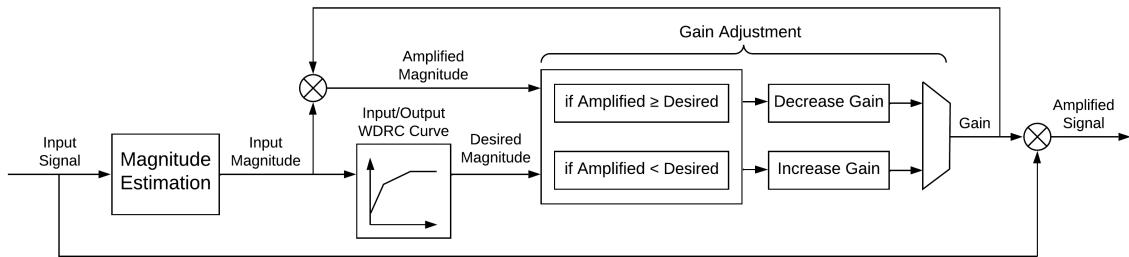


FIGURE 7. Wide Dynamic Range Compression estimates the envelope of a signal and applies positive or negative gain according to the WDRC input-output curve. The speed at which gain is increased or decreased is determined by attack and release times.

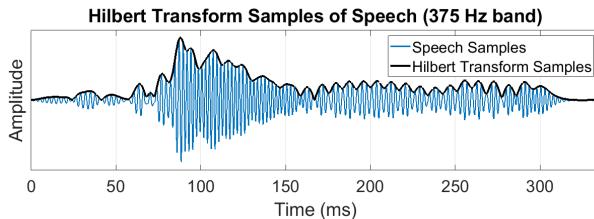


FIGURE 8. The waveform and computed envelope of the word “please” in the 375 Hz band, spoken by a female voice.

same magnitude response as a linear phase filter, but the lowest possible delay. A filter can be converted from linear phase to minimum phase by reflecting all roots which lie outside the unit circle.

Figure 6 (bottom) shows the aligned impulse responses of the minimum phase filter bank. As seen from Figure 6, converting the filters from linear to minimum phase dramatically decreases the delay of each band. While retaining the same functionality as a linear phase filter bank, the minimum phase filter bank has a latency of only 5.4 ms, compared to 32 ms, which makes it suitable for real-time applications.

III. WIDE DYNAMIC RANGE COMPRESSION

A. OVERVIEW

WDRC is a type of automatic gain control (AGC) system which reduces the dynamic range of audio by applying varying gain to a signal depending on the input magnitude. For any instantaneous input magnitude, the WDRC curve, shown in Figure 9 (left), determines the desired instantaneous output magnitude. The WDRC curve is defined by a combination of parameters, which change the gain, the maximum power output, the “knee low” and “knee up” points, and the slope of the compression region. The reciprocal of the slope of the compression region is called the “compression ratio” (CR).

It is insufficient, however, to set the gain of each audio sample independently. Studies in acoustics and speech intelligibility have shown that the rate of change of WDRC gain has a strong effect on speech clarity and legibility [16], [17]. The rate of change of gain is measured using the attack and release times, which play a key role in the performance of WDRC. However, to the best of our knowledge, currently available open-source hearing aids do not have an accurate mechanism

for setting attack and release times independently of other parameters. For example, the attack and release times of the Kates system [18] depend on the user-defined compression ratio, which gives rise to major inaccuracies.

In this paper, we explore the complex relationship between the attack and release times of WDRC and the parameters defining a WDRC curve. We also propose a multirate compression algorithm which yields precise response times in accordance with ANSI standards for any user-defined WDRC parameters.

B. MAGNITUDE ESTIMATION

Wide Dynamic Range Compression calculates compression gains based on input magnitude. However, sound is a modulating signal, meaning the magnitude of the signal is contained in the envelope. Common approaches to finding the envelope of a modulating signal include peak detection [18], per-frame total power [19], sliding RMS windows, and more. However, all of these approaches introduce inaccuracies into the envelope estimate, such as ripple or excessive smoothing.

We estimate the signal envelope by employing the Hilbert Transform. The Hilbert Transform accepts a real signal, and computes a 90-degree phase shifted imaginary component. The magnitude of the input signal is then found as the absolute value of the real and imaginary components.

The accuracy of the Hilbert Transform depends on the accuracy of the underlying Hilbert Filter, which is a filter that cuts off the negative frequencies of the signal spectrum. If the transition bandwidth of the Hilbert Filter overlaps with signal content, then the computed envelope becomes distorted.

As seen from Figure 2, many of the channels are very close to DC, and preserving these frequencies would require an unrealistically sharp Hilbert Filter. However, we prevent distortion in the low-frequency bands by performing magnitude estimation and amplification in the multirate domain, as shown in Figure 1. As we discussed earlier, reducing the sampling rate of a filter increases its relative width. However, for a given center frequency, reducing the sampling rate of the signal also moves said center frequency relatively farther from DC. As such, the channel is no longer affected by the Hilbert Filter’s transition bandwidth.

The multirate Hilbert Transform produces highly accurate signal envelopes for all frequency channels of the filter bank.

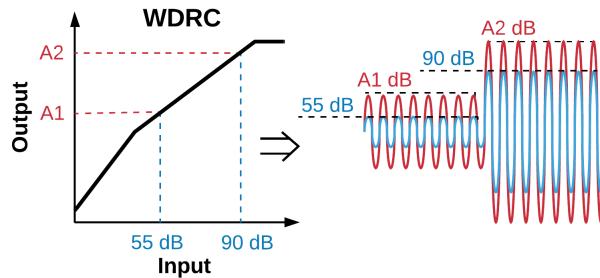


FIGURE 9. ANSI attack and release times of hearing aids are measured using a sinusoidal step input changing from 55 dB to 90 dB. The WDRC curve determines the desired output magnitudes.

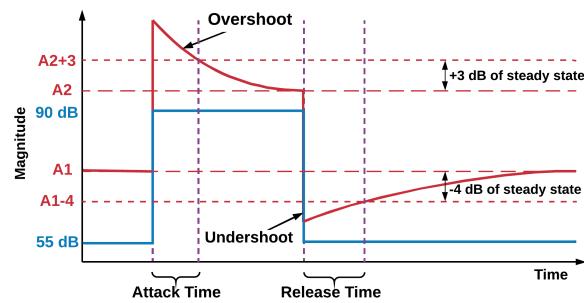


FIGURE 10. ANSI standard attack time is measured as the time it takes for the overshoot to settle within 3 dB of steady state. Release time is measured as the time it takes for the undershoot to settle within 4 dB of steady state.

Figure 8 shows the 0.375 kHz band of the word “please” spoken by a female voice from the TIMIT database [20], as well as the envelope of the waveform computed using the Hilbert Transform.

C. PROPOSED AUTOMATIC GAIN CONTROL

The ANSI Specification of Hearing Aid Characteristics defines the attack and release times for hearing aid devices [21]. Given a step input which changes magnitude from 55 dB to 90 dB, as shown in Figure 10, the attack time is defined as the time elapsed between the step change and the time the output remains within 3 dB of its steady state value, notated as A_2 in Figure 10. Release time is similarly defined as the time elapsed between a step change from 90 dB to 55 dB, and the time the output remains within 4 dB of steady state, notated as A_1 . The steady-state values are obtained from the WDRC curve, shown in Figure 9, and as such, depend on compression parameters.

The general concept of Automatic Gain Control for WDRC, illustrated in Figure 7, is to decrease the gain when the output overshoots, and increase the gain when the output undershoots. However, since the steady state values A_1 and A_2 shown in Figure 10 depend on user parameters, the overshoot and undershoot also depend on user compression parameters. Thus, there is a relationship between user input parameters and the response speed of an AGC loop which is not well explored in modern hearing aids and leads to

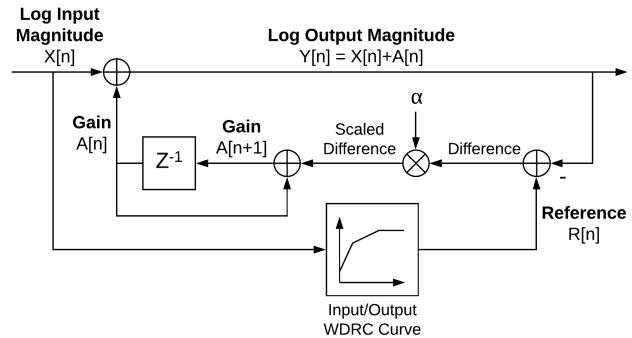


FIGURE 11. The proposed algorithm for automatic gain control, offering precise control over the dynamic response times. The attack and release times of the loop are controlled by the parameters α .

significant error in actual attack and release times compared to desired values.

We derived a closed-form relationship between user compression parameters (compression ratio) and the attack and release times of a hearing aid, and designed an Automatic Gain Control (AGC) loop which yields exact attack and release values for any user-defined compression parameters. Our design builds upon work in [22] by adapting radio AGC to Wide Dynamic Range Compression. The block diagram of the proposed AGC algorithm is shown in Figure 11. For each input sample, the gain of the previous sample is added to the current sample. The sum is then compared to the desired output level based on the WDRC curve. The scaled difference between the desired and the actual output levels is then used to modify the gain of the next sample. In the AGC loop, alpha (α) is an important scaling parameter which determines how quickly the system reacts to changes. As such, α is the only parameter determining the attack and release times of the AGC loop. Since WDRC must respond differently to rising and falling input levels, the AGC loop requires two distinct values of α – one for attack time, one for release time.

In this section, we derive the relationship between α and WDRC parameters such that the system yields exact attack and release times in any configuration. The behavior of the system above is described by the equation below.

$$\begin{aligned} A[n+1] &= A[n] + \alpha \times (R[n] - Y[n]) \\ &= A[n] + \alpha \times (R[n] - (X[n] + A[n])) \\ &= A[n] \times (1 - \alpha) + \alpha \times (R[n] - X[n]) \end{aligned} \quad (1)$$

Consider the ANSI test signal, which is a step input which changes magnitude from 55 dB to 90 dB at time $n = 0$. Let us define G_0 as the initial steady state gain before the step change. For $n < 0$, $R[n] = A_1$, $X[n] = 55$, so $G_0 = R[n] - X[n] = A_1 - 55$.

Let us define G_∞ as the final steady state gain after the step change. For all times $n \geq 0$, $R[n] = A_2$, $X[n] = 90$, so $G_\infty = R[n] - X[n] = A_2 - 90$. Using these definitions, for all $n \geq 0$, equation 1 can be rewritten as:

$$A[n+1] = A[n] \times (1 - \alpha) + \alpha \times G_\infty \quad (2)$$

TABLE 3. Maximum and average wide dynamic range compression steady state error for seven standard hearing loss profiles.

No.	Category	Max (dB)	Average (dB)										
			250	354	500	707	1000	1414	2000	2828	4000	5657	8000
N1	Very Mild	0.5	0.0	-0.2	0.0	-0.2	0.0	0.0	-0.1	-0.1	-0.3	0.0	-0.2
N2	Mild	1.0	-0.1	-0.4	0.1	0.0	-0.1	0.0	-0.1	-0.1	-0.3	0.2	0.1
N3	Moderate	1.5	-0.1	-0.3	0.1	0.0	0.1	0.1	0.0	0.0	-0.1	0.0	0.3
N4	Mod/Severe	2.0	0.0	0.5	0.1	0.5	0.5	0.5	0.4	-0.1	0.0	0.0	0.5
N5	Severe	1.5	-0.1	-0.1	0.1	0.0	0.0	0.3	0.1	-0.1	0.0	-0.1	0.2
N6	Severe	2.0	-0.1	-0.1	0.2	0.0	0.2	0.3	0.2	-0.2	-0.1	-0.2	0.4
N7	Profound	3.0	0.2	-0.3	0.0	0.1	0.0	0.2	-0.1	-0.4	0.2	-0.1	0.3

In order to gain insight into the behavior of the system, let us write out the gains of the first few samples:

$$A[0] = G_0 \quad (3a)$$

$$A[1] = G_0 \cdot (1 - \alpha) + \alpha \cdot G_\infty \quad (3b)$$

$$A[2] = G_0 \cdot (1 - \alpha)^2 + \alpha \cdot G_\infty \cdot (1 - \alpha) + \alpha \cdot G_\infty \quad (3c)$$

$$A[3] = G_0 \cdot (1 - \alpha)^3 + \alpha \cdot G_\infty \cdot (1 - \alpha)^2 + \alpha \cdot G_\infty \cdot (1 - \alpha) + \alpha \cdot G_\infty \quad (3d)$$

As seen from the pattern formed in equation 3, the gain of the n 'th sample is found as a geometric series, shown in equation 4a and simplified in equation 4b.

$$A[n] = G_0 \times (1 - \alpha)^n + \alpha \times G_\infty \times (1 + (1 - \alpha) + (1 - \alpha)^2 + \dots + (1 - \alpha)^{n-1}) \quad (4a)$$

$$A[n] = G_0 \times (1 - \alpha)^n + \alpha \times G_\infty \times \left(\frac{1 - (1 - \alpha)^n}{\alpha} \right) = G_0 \times (1 - \alpha)^n + G_\infty \times (1 - (1 - \alpha)^n) = (G_0 - G_\infty) \times (1 - \alpha)^n + G_\infty \quad (4b)$$

This important result provides us with an equation for gain as a function of time and α . As expected, at time $n = 0$ the gain is equal to G_0 , and as n reaches infinity the gain approaches G_∞ .

Using the equation above, we can use known values of n to solve for α . As explained earlier, α is the only parameter which sets the attack and release times of the AGC system. Let AT represent the attack time. From the ANSI definition of attack time, we know that at time $n = AT$, the gain needs to be within 3 dB of steady state, which is $G_\infty + 3$. Substituting these values into equation 4b yields:

$$G_\infty + 3 = (G_0 - G_\infty) \times (1 - \alpha)^{AT} + G_\infty \quad (5)$$

The equation above contains only one unknown variable, allowing us to solve for α_{attack} :

$$\alpha_{attack} = 1 - \left(\frac{3}{G_0 - G_\infty} \right)^{\frac{1}{AT}} = 1 - \left(\frac{3}{A1 - A2 + 35} \right)^{\frac{1}{AT}} \quad (6)$$

Following similar steps and using the ANSI definition for release time, we can find a similar expression for $\alpha_{release}$:

$$\alpha_{release} = 1 - \left(\frac{4}{G_0 - G_\infty} \right)^{\frac{1}{RT}} = 1 - \left(\frac{4}{A1 - A2 + 35} \right)^{\frac{1}{RT}} \quad (7)$$

Equations 6 and 7 provide us with values for α_{attack} and $\alpha_{release}$ that guarantee exact attack and release times for the

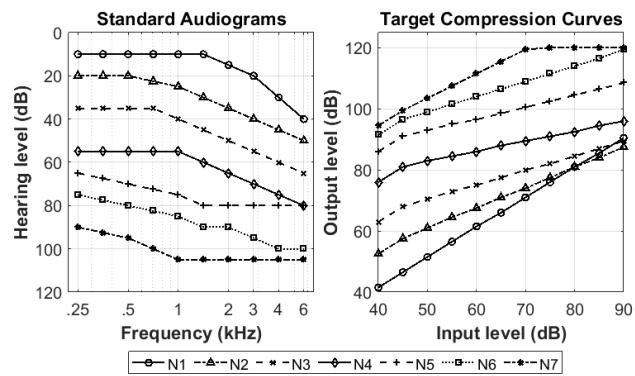


FIGURE 12. Standard Audiograms for hearing aid testing developed by the ISMADHA group [23] (left); The corresponding target compression curves at 1000 Hz (right).

AGC loop. It is important to note that in equation 6 and 7, the units for AT and RT are samples. Samples and milliseconds are related to each other through sampling rates which, as described earlier, varies between the different subbands.

It can be noted that the difference $G_0 - G_\infty$ is none other than the Overshoot pictured in Figure 10. The Overshoot is a variable which depends on the parameters setting the WDRC curve. By deriving the relationship between α and Overshoot, we account for all WDRC parameters, including compression ratio, in our calculations for attack and release times.

Another feature of the AGC loop, shown in Figure 11, is that the reference signal $R[n]$ needs not be a piecewise curve, as shown in Figure 9. The piecewise input-output WDRC curve benefits from simplicity, but our system can accept any function for the input-output curve, including smooth continuous functions and 'S' curves. This flexibility allows the user to employ other input-output curves, which may be more appropriate for the user.

IV. EXPERIMENTAL RESULTS

A. IMPLEMENTATION TESTBED

We have integrated the audiometric filter bank into the Open Speech Platform (OSP) [4]–[6], which is an open source suite of hardware and software tools for conducting research into many aspects of hearing loss both in the lab and the field. The hardware system consists of a battery operated wearable device running a Qualcomm 410c processor, similar to those in cellphones, with two ear-level assemblies attached – one for each ear. More details about the hardware systems can be found in [5].

At the core of OSP software is the real-time Master Hearing Aid (RT-MHA) reference design. Initially, the incoming

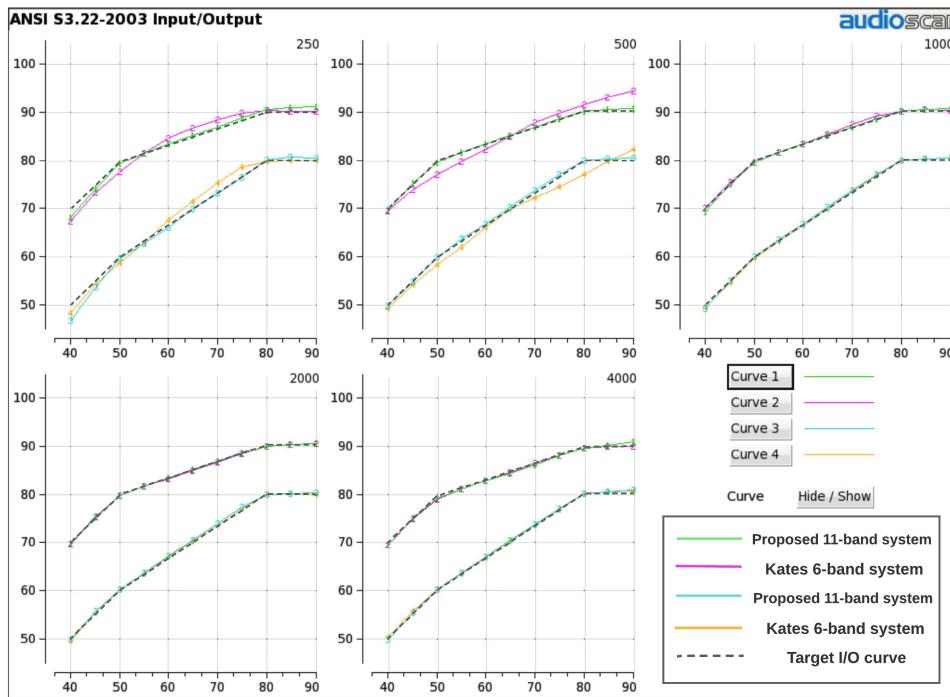


FIGURE 13. Verifit Verification Toolbox measurements comparing the steady state behavior of the multirate 11-band system and the Kates 6-band system. The Verifit tool generates tones of increasing magnitude, indicated on the x-axis in units of dB SPL. It then records and plots the steady state output on the y-axis, forming the WDRC input-output curves for both systems at 250, 500, 1000, 2000, and 4000 Hz.

audio signal from the microphones is sampled at 48 kHz, and is then downsampled to 32 kHz (not to be confused with the resamplers present in the channelizer). The audio signal is then routed to the channelizer.

The outputs of the channelizer then pass through the WDRC unit to compensate for the user's hearing loss. Then the amplified outputs are recombined and passed through a Global Maximum Power Output (MPO) controller in order to limit the power outputted by the speaker. Finally, the audio is upsampled from 32 kHz back to 48 kHz and outputted through the speakers. Additionally, the RT-MHA reference design contains Adaptive Feedback Cancellation (AFC) in order to compensate for the feedback arising from the close proximity of the microphone and the speaker. More detailed explanations of the RT-MHA components can be found in [5], [6].

B. VERIFIT VERIFICATION TOOLBOX

We evaluated the design using the widely accepted Audioscan Verifit 2 Professional Verification system [24]. Verifit 2 is a verification tool consisting of a soundproof binaural audio chamber, a display unit, and a set of powerful testing procedures, such as speech map, ANSI tests, and distortion.

We conducted steady state input-output measurements to evaluate the multirate amplification system running on Open Speech Platform hardware. The purpose of this test is to compare the experimentally measured input-output curve of our device to the ideal target curve specified by a hearing loss prescription. In this experiment, the hearing aid device

is placed into the soundproof audio chamber. The Verifit's reference speaker plays calibrated audio signals with known acoustical properties into the hearing aid microphone, which becomes the input signal for the hearing aid. The processed output signal of the hearing aid is then collected by the Verifit's coupler microphone and is compared to the input signal to identify the measured gain.

We verified our system using seven standard pure tone audiograms developed by the International Standard for Measuring Advanced Digital Hearing Aids (ISMADHA) group [23], which represent a broad class of hearing loss patterns, from very mild to profound. We obtained compression parameters by passing a subset of ISMADHA profiles through the NAL-NL2 Prescription Procedure [25], which is a widely accepted algorithm for generating hearing aid prescriptions from pure tone audiograms. Figure 12 shows the ISMADHA standard pure tone audiograms, and an example of the obtained target input/output amplification curves for each audiogram at 1 kHz.

We performed steady state measurements at the eleven half-octave frequencies offered by the audiometric filter bank. For each frequency, we obtained the target compression curves, such as the ones shown in Figure 12. We then took measurements for each combination of audiogram, frequency, and input level, resulting in 847 data points. Table 3 shows the maximum and average errors we obtained for each audiogram as a function of frequency. Our results show that the compressed output values closely match the desired target values, often with 0 dB average error. The maximum error

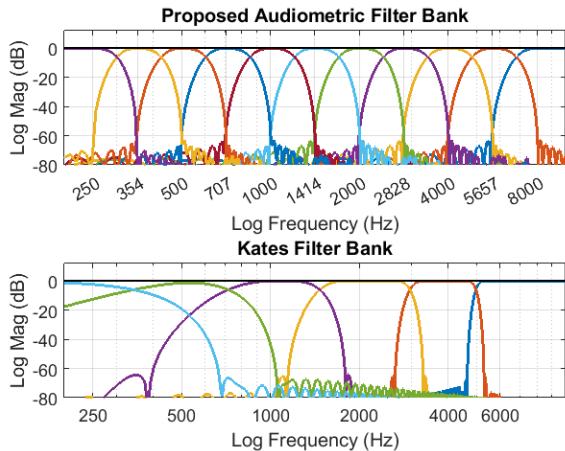


FIGURE 14. The magnitude responses of the multirate audiometric filter bank versus the widely used Kates Filter Bank. The multirate filter bank offers more bands and sharper band edges with lower processing complexity than the Kates system.

TABLE 4. Complexity and latency comparison between the Multirate Audiometric Filter Bank and Kates Filter Bank.

Filter Bank	Bands	Operations per sample	Latency
Proposed OSP Filter Bank	11	437.69	5.43 ms
Kates Filter Bank	6	1542	4.03 ms

(usually found in the MPO region) is also small, and never exceeds 3 dB, which was shown to be the threshold of just noticeable difference in speech-to-noise ratio [26].

V. COMPARISON WITH PRIOR WORK

We compared the (i) Multirate Audiometric Filter Bank and (ii) Multirate Wide Dynamic Range Compression System with the Kates Digital Hearing Aid [4], [5], [18], one of the most popular open-source tools for hearing aid research.

A. SPECTRAL DECOMPOSITION

The main motivation for this work was to improve the spectral resolution of hearing aids. Figure 14 compares the magnitude responses of the proposed audiometric filter bank and the Kates 6-band filter bank. In addition to offering more bands, the multirate filter bank also offers better filter sharpness. Although most of Kates's filter satisfy ANSI class 0 requirements, the filters lose their sharpness at lower frequencies, and the 500 Hz filter does not satisfy the requirements for any of the ANSI classes.

We also used the Verifit's input-output curve feature to compare the prescription accuracy of the multirate eleven-band system versus the Kates system. Figure 13 shows two target compression curves and the six band versus eleven band realizations. At higher frequencies, both realizations accurately fulfill the target prescription. However, starting from 1000 Hz, the Kates implementation begins to diverge from the target curve, and both the 250 and 500 Hz bands lose their shape integrity. This is due to the high side lobes of the low frequency bands seen in Figure 14.

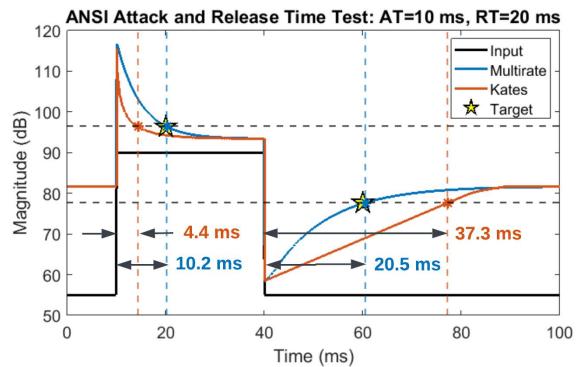


FIGURE 15. ANSI attack and release time test for the Proposed Multirate and Kates's automatic gain control. The input is a 2 kHz sine wave alternating between 55 dB and 90 dB. Asterisks mark the measured attack and release times, and stars mark the ANSI S3.22 target values.

Table 4 compares the complexity and latency of the Kates filter bank and the eleven band filter bank. In addition to offering almost twice the number of bands compared to Kates's filter bank, the proposed filter bank achieves about 3.5× improvement in complexity, with a comparable algorithmic latency of 5.43 ms.

B. AUTOMATIC GAIN CONTROL

We also compared our Multirate Multiband Automatic Gain Control with Kates's approach. As described Section III-C, the relationship between WDRC parameters and AGC response times are not explored in previous works. In the Kates approach, the AGC response times are controlled by the coefficients of the peak detector used to estimate the signal magnitude. The resulting coefficients are approximated to meet ANSI attack and release time standards, but diverge from target values significantly.

As a test case, Figure 15 compares the dynamic responses of the proposed multirate system and the Kates system. The input is a gated sinusoid test signal stepping between 55 and 90 dB, as defined by the ANSI S3.22 standard [21], centered at 2000 Hz. Both systems were configured to have a compression ratio of 3:1, and the attack and release times were set to 10 ms and 20 ms respectively. The dynamic responses of the two systems are shown in Figure 15.

In this experiment, the measured attack and release times of the proposed Multirate system are 10.2 ms and 20.5 ms respectively, which deviate from the target values by 0.2 ms (2%) and 0.5 ms (2.5%). On the other hand, the measured attack and release times of the Kates system are 4.4 ms and 37.3 ms respectively, which is a 5.6 ms (45%) and 17.3 ms (87%) deviation from the target values. This experiment shows that the proposed Multirate system satisfies attack and release times within 0.5 ms of the target value. However, the Kates system yields attack and release time values that significantly diverge from the target. Furthermore, this error is unpredictable because the internal coefficients responsible for attack and release times of the Kates system are designed to be “fudge” factors.

The proposed Multirate systems offers very accurate fulfillment of user designated attack and release times. However, neither the current standards [21] nor popular HA prescription tools, e.g., [25] provide guidance for the dynamic aspects of dynamic range compression. There is a need for the signal processing and audiology research community to address this important gap and investigate the role that response times play in speech legibility and perception.

VI. CONCLUSION

In this paper, we presented a real-time multirate, multiband amplification system for hearing aids. Our system improves upon the prescription accuracy of hearing aids and provides an open source tool for hearing loss research.

We designed a channelizer offering eleven frequency sub-bands centered at the standard frequencies used in pure-tone audiometry, with high side-lobe attenuation and low ripple. This high frequency allows our hearing aid system to accurately satisfy hearing aid prescriptions, even for complex and unusual hearing loss patterns. The channelizer uses multirate processing to reduce the complexity by about 14 \times compared to a single-rate implementation. By employing minimum-phase filters, we decreased the latency of our filter bank to 5.43 ms, which is within the conventional threshold for modern hearing aids.

We also designed an automatic gain control (AGC) system which provides accurate control of the steady state and dynamic behavior of dynamic range compression. We use the Hilbert Transform to find the instantaneous signal magnitude, which provides higher accuracy than conventional instantaneous power estimation methods. Furthermore, we derived the closed-form relationship between the compression parameters of our AGC loop, and the attack and release times at the output. The accurate fulfillment of attack and release times in dynamic range compression opens new opportunities for exploring the relationship between response times and hearing impaired users' satisfaction.

We implemented the Multirate Multiband Amplification System on Open Speech Platform - an open source suite of hardware and software tools for hearing loss research. The system runs in real-time on a wearable device, and is suited for hearing loss research both in the lab and in the field.

REFERENCES

- R. J. Bennett, A. Laplante-Lévesque, C. J. Meyer, and R. H. Eikelboom, "Exploring hearing aid problems: Perspectives of hearing aid owners and clinicians," *Ear Hearing*, vol. 39, no. 1, pp. 172–187, 2018.
- M. Chasin and F. A. Russo, "Hearing aids and music," *Trends Amplification*, vol. 8, no. 2, pp. 35–47, 2004.
- P. E. Souza, "Effects of compression on speech acoustics, intelligibility, and sound quality," *Trends Amplification*, vol. 6, no. 4, pp. 131–165, Dec. 2002.
- H. Garudadri, A. Boothroyd, C.-H. Lee, S. Gadiyaram, J. Bell, D. Sengupta, S. Hamilton, K. C. Vastare, R. Gupta, and B. D. Rao, "A real-time, open-source speech-processing platform for research in hearing loss compensation," in *Proc. 51st Asilomar Conf. Signals, Syst., Comput.*, Oct. 2017, pp. 1900–1904.
- L. Pisha, J. Warchall, T. Zubatiy, S. Hamilton, C.-H. Lee, G. Chockalingam, P. P. Mercier, R. Gupta, B. D. Rao, and H. Garudadri, "A wearable, extensible, open-source platform for hearing healthcare research," *IEEE Access*, vol. 7, pp. 162083–162101, 2019.
- D. Sengupta, T. Zubatiy, S. K. Hamilton, A. Boothroyd, C. Yalcin, D. Hong, R. Gupta, and H. Garudadri, "Open speech platform: Democratizing hearing aid research," in *Proc. 14th EAI Int. Conf. Pervasive Comput. Technol. Healthcare*, May 2020, pp. 223–233.
- THE_Lab at UC San Diego. (2019). *Open Speech Platform*. [Online]. Available: <http://openspeechplatform.ucsd.edu/>
- Y. Hu and P. C. Loizou, "Subjective comparison and evaluation of speech enhancement algorithms," *Speech commun.*, vol. 49, nos. 7–8, pp. 588–601, Jul. 2007.
- Y. Ghanbari, M. Karami, and B. Amelifard, "Improved multi-band spectral subtraction method for speech enhancement," in *Proc. 6th IASTED Internat. Conf. Signal Image Process.*, 2004, pp. 225–230.
- J. M. Kates, "Dynamic range compression using digital frequency warping," U.S. Patent 8 014 549, Sep. 6, 2011.
- L. Lee and R. Rose, "A frequency warping approach to speaker normalization," *IEEE Trans. Speech Audio Process.*, vol. 6, no. 1, pp. 49–60, Jan. 1998.
- Guidelines for Manual Pure-Tone Threshold Audiometry*, American Speech-Language-Hearing Association, Rockville, MD, USA, May 2020.
- Specification for Octave-Band and Fractional-Octave-Band Analog and Digital Filters*, Standard ANSI S1.11-2004.
- R. E. Crochiere and L. R. Rabiner, *Multirate Digital Signal Processing*. Upper Saddle River, NJ, USA: Prentice-Hall, 1983.
- M. A. Stone and B. C. Moore, "Tolerable hearing aid delays. I. Estimation of limits imposed by the auditory path alone using simulated hearing losses," *Ear Hearing*, vol. 20, no. 3, pp. 182–192, 1999.
- L. M. Jenstad and P. E. Souza, "Quantifying the effect of compression hearing aid release time on speech acoustics and intelligibility," *J. Speech, Lang., Hearing Res.*, vol. 48, no. 3, pp. 651–667, Jun. 2005.
- A. C. Neuman, M. H. Bakke, C. Mackersie, S. Hellman, and H. Levitt, "The effect of compression ratio and release time on the categorical rating of sound quality," *J. Acoust. Soc. Amer.*, vol. 103, no. 5, pp. 2273–2281, May 1998.
- J. M. Kates, *Digital Hearing Aids*. San Diego, CA, USA: Plural Publishing, 2008.
- G. Grimm, T. Herzke, S. Ewert, and V. Hohmann, "Implementation and evaluation of an experimental hearing aid dynamic range compressor," *Threshold*, vol. 80, no. 90, p. 100, 2015.
- J. S. Garofolo, L. F. Lamel, W. M. Fisher, J. G. Fiscus, D. S. Pallett, N. L. Dahlgren, and V. Zue, *Timit Acoustic Phonetic Continuous Speech Corpus*. Philadelphia, PA, USA: Linguistic Data Consortium, 1993.
- Specification Hearing Aid Characteristics*, Standard ANSI S3.22, 2009.
- F. Harris and G. Smith, "On the design, implementation, and performance of a microprocessor controlled AGC system for a digital receiver," presented at the IEEE Mil. Commun. Conf. (MILCOM), San Diego, CA, USA, Oct. 23–26, 1988.
- N. Bisgaard, M. S. M. G. Vlaming, and M. Dahlquist, "Standard audiograms for the IEC 60118-15 measurement procedure," *Trends Amplification*, vol. 14, no. 2, pp. 113–120, Jun. 2010.
- Audioscan Verifit User's Guide 4.28*, Audioscan, Dorchester, ON, Canada, 2022.
- G. Keidser, H. Dillon, M. Flax, T. Ching, and S. Brewer, "The NAL-NL2 prescription procedure," *Audiol. Res.*, vol. 1, no. 1, pp. 88–90, 2011.
- D. McShefferty, W. M. Whitmer, and M. A. Akeroyd, "The just-noticeable difference in speech-to-noise ratio," *Trends Hearing*, vol. 19, Dec. 2015, Art. no. 233121651557231.



ALICE SOKOLOVA (Graduate Student Member, IEEE) received the B.S. degree (*summa cum laude*) in electrical engineering from San Diego State University, San Diego, CA, USA. She is currently pursuing the joint Ph.D. degree with San Diego State University and the University of California San Diego.

She was the valedictorian of her graduating class at San Diego State University. Her field of study is digital signal processing, under the mentorship of Prof. Emeritus Fred Harris. In addition to research, she enjoys teaching and has worked as an Instructional Assistant for numerous undergraduate and graduate level classes in electrical and computer engineering, including senior capstone design. Her current research interests include signal processing for hearing aids, healthy aging, and other biomedical applications.



DHIMAN SENGUPTA received the B.S. degree in computer engineering from the University of Michigan, with a focus on DSP and computer architecture. He is currently pursuing the Ph.D. degree with the Microelectronic Embedded Systems Laboratory, University of California San Diego, under the supervision of Prof. Rajesh Gupta.

His background lies in working on interdisciplinary projects with different types of embedded sensors and handling their data flows by processing the data in real-time using novel algorithms. He also worked for three years at the Advanced Space Position, Navigation and Timing Branch, Naval Research Laboratory, where his research was focused primarily on time synchronization and syntonization of clocks across long distances. His current research interest includes developing real-time embedded tools for the next generation of wearable technology.



BARIS AKSANLI (Member, IEEE) is currently an Assistant Professor with the Electrical and Computer Engineering Department, San Diego State University, San Diego, CA, USA. Previously, he was a Postdoctoral Researcher at the Computer Science and Engineering Department, University of California San Diego. As a Researcher, his affiliations include the Multi Scale Systems Center (MuSyC), the TerraSwarm Research Center, and the Center for Networked Systems (CNS); and the collaborators of his projects include Google, Microsoft, Panasonic, Intel, and IBM. His research interests include energy efficiency and peak power management of large-scale systems, such as data centers and smart grids, efficient battery usage in data centers and residential houses, battery lifetime modeling, cost and energy aware automation of residential houses, learning techniques to enhance user behavior modeling and context extraction, house/building/data center, and grid interaction.



MARTIN HUNT received the B.S. degree in computer science from the California Institute of Technology, CA, USA. He is an Experienced Software Engineer with years of embedded and OS kernel work (C/C++ and Linux), with more recent scientific programming projects (python, high performance computing, and visualization). He has led teams developing commercial products and software for scientific research. He currently holds the position of Software Team Lead for the open speech platform with the University of California San Diego.



FREDRIC HARRIS (Senior Member, IEEE) received the bachelor's degree in electrical engineering (EE) from the Polytechnic Institute of Brooklyn, in 1961, the master's degree in electrical engineering (EE) from San Diego State University, in 1967, and the Ph.D. degree from the University of California San Diego, in 1973.

He is currently a Faculty Member with the Electrical and Computer Engineering Department, University of California San Diego. Previously, he taught at the College of Engineering, San Diego State University, for 50 years. He continues to teach courses in digital signal processing and communication systems at UCSD and lectures throughout the world on DSP applications. He has extensive practical experience applying his skills to satellite and cable TV communication systems, underwater acoustics, advanced radar, and high performance laboratory instrumentation. He is well published, holds a number of patents, and has contributed to a number of books on DSP. He is a Former Adjunct Member of the Center for Communications Research, Princeton, and Imperial College. In 1990 and 1991, he was the Technical Chair and then the General Chair of the Asilomar Conference on Signals, Systems, and Computers.



RAJESH GUPTA (Fellow, IEEE) is currently the Founding Director of the Halicioğlu Data Science Institute and also a Distinguished Professor of computer science and engineering with the University of California San Diego. He currently leads the NSF Project MetroInsight and a co-PI on DARPA/SRC Center on Computing on Network Infrastructure (CONIX), with the goal to build new generation of distributed cyber-physical systems that use city-scale sensing data for improved services and autonomy. His research interests include embedded and cyber-physical systems, with a focus on sensor data organization and its use in optimization and analytics. He is a fellow of ACM and the American Association for the Advancement of Science (AAAS). He holds the Qualcomm Endowed Chair in embedded microsystems at UC San Diego and the INRIA International Chair at the French International Research Institute, Rennes, Bretagne Atlantique. He currently served as the Editor-in-Chief for the IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN OF INTEGRATED CIRCUITS AND SYSTEMS, *IEEE Design & Test* and as founding EIC of IEEE EMBEDDED SYSTEMS LETTERS.



HARINATH GARUDADRI (Member, IEEE) received the Ph.D. degree in electrical engineering from The University of British Columbia, Vancouver, BC, Canada, in 1988. In The University of British Columbia, he spent half his time in ECE and the other half at the School of Audiology and Speech Sciences, Faculty of Medicine. He is currently an Associate Research Scientist with the Qualcomm Institute, UC San Diego. He moved to academia, in November 2013, after 26 years in the industry to work on technologies that will improve healthcare delivery beyond hospital walls. His research interests include signal processing applications in diverse fields, such as speech recognition, machine learning, speech, audio and video compression, multimedia delivery in 3G/4G networks, low-power sensing and telemetry of physiological data, reliable body area networks (BAN), noise cancellation, and artifacts mitigation, among other areas. He holds more than 50 granted patents and over 20 pending patents in these areas.