A Low-Power Magnitude Detector for Analysis of Transient-Rich Signals

Brandon Rumberg, *Student Member, IEEE*, and David W. Graham, *Member, IEEE*

Abstract—**Magnitude detection, such as envelope detection or RMS estimation, is needed for many low-power signal-analysis applications. In such applications, the temporal accuracy of the magnitude detector is as important as its amplitude accuracy. We present a low-power audio-frequency magnitude detector that simultaneously achieves both high temporal accuracy and high amplitude accuracy. This performance is achieved by rectifying the signal with a high-ripple peak detector and then averaging this rectified signal with an adaptive-time-constant filter. The time constant of this filter decreases with increasing amplitude, enabling the filter to quickly respond on a short time scale to transients, while steady-state ripple is averaged on a longer time scale. The circuit has been fabricated in a 0.18 μm CMOS process and consumes only 1.1 nW–1.08 μW when tuned for operation from 20 Hz–20 kHz. It exhibits a dynamic range of 70 dB across typical speech frequencies.**

Index Terms—**Analog processing circuits, CMOS integrated circuits, continuous-time circuits, envelope detectors, magnitude detectors, nonlinear dynamic circuits, peak detectors, ultra-low power.**

I. INTRODUCTION

MAGNITUDE-DETECTION circuits—such as envelope detectors, peak detectors, and RMS-to-DC converters—produce an estimate of a signal's magnitude and are thus important elements in communications transceivers [1], automatic gain control systems [2], and analog spectral analyzers [3]–[7]. Since the magnitude is a time-varying quantity, the accuracy of a magnitude detector has two components: amplitude accuracy and temporal accuracy. Traditionally, the design emphasis of magnitude circuits has been on amplitude accuracy; however, temporal accuracy is crucial when the magnitude changes quickly relative to the frequencies of the underlying carrier signal, such as in speech signals [8]. Thus, temporal accuracy is important in audio-processing systems, including ultra-low-power (ULP) applications such as bionic ears [7], [9] and event detectors for wakeup applications [10]. Existing digital signal processing techniques for temporal accuracy use non-physical, non-causal filters that require powerful processing and non-stop data conversion, thereby limiting

Fig. 1. (a) Schematic of our magnitude detector. (b) Block diagram of an analog spectral-analysis system.

their use in ULP systems. Discrete-time analog circuits have also been explored to achieve good temporal accuracy [11]; however, the power level and the sampled-data representation are inappropriate for many ULP signal-analysis systems. In this paper, we present an ULP continuous-time magnitude detector that has been designed with an emphasis on temporal accuracy, while still achieving high amplitude accuracy.

Fig. 1(a) shows our magnitude detector, wherein a rectifying nonlinearity provides an initial estimate of the signal's magnitude and then a low-pass filter averages this estimate to obtain the final smooth magnitude estimate. For the rectifying nonlinearity, we have developed a voltage-mode asymmetric integrator. This circuit's asymmetry causes the average level of its output to shift in proportion to the input magnitude, thereby providing a magnitude estimate that is superimposed with a ripple. This ripple is then smoothed by a nonlinear low-pass filter with an adaptive time constant. We have designed this filter's nonlinearity such that its time constant shrinks in response to large input-output differential signals, thereby reducing the integration window in order to follow transients more closely, while maintaining a long time constant for small signals in order to retain good ripple suppression.
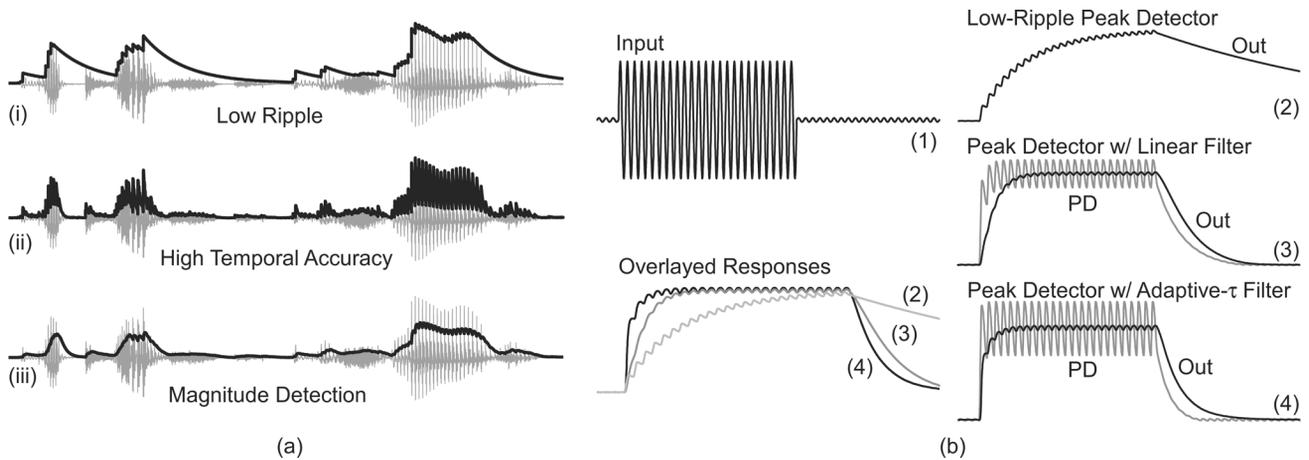
Fig. 2. (a) Tradeoff between a response with low ripple and a response with high temporal accuracy. (b) Comparing the temporal accuracy of three magnitude-detection architectures, each with 1% ripple. "Overlayed Responses" shows the improved temporal response achieved by our adaptive-time-constant filter. All plots in this figure are from numerical simulations of the equations discussed in Sections III and IV.

The remainder of this paper is organized as follows. In Section II, we discuss the application-space of our circuit and describe the high-level design approach. In Sections III and IV, we present the asymmetric integrator and the adaptive-time-constant filter, respectively. Then, in Section V, we combine the subcircuits into the complete magnitude detector and present our experimental results. We have fabricated this circuit in standard 0.18 μm, 0.35 μm, and 0.5 μm CMOS processes. Unless otherwise noted, all plots are measurements from the 0.18 μm circuit.

II. MAGNITUDE DETECTOR ARCHITECTURE

A common application environment for magnitude detectors is within audio- and vibration-processing systems. For example, a standard first step in such systems is spectral analysis, which can be implemented in low-power analog circuits to make an efficient real-time sensor-processing front-end [12]–[14]. Such a spectral analysis front-end can be combined with other analog processing circuits to create an entire ULP system (such as for implantable electronics), or the front-end can be used as an event detector to wake up a higher-power back-end, thus reducing system-level power consumption [10]. Analog spectral-analysis systems typically consist of a bank of filters that decompose the signal into frequency components [4]–[7], followed by subband processing blocks, such as magnitude detectors for extracting the magnitude of the spectrum [15]–[17], as illustrated in Fig. 1(b). Adaptation is also often used to increase the dynamic range of the system, either by using automatic gain control on the pre-filtered signal [7] or by adapting the gain/Q in individual subbands. Such adaptation is typically based on the signal's magnitude, and therefore requires a magnitude detector. In addition to needing to respond quickly to changes in the signal, magnitude detectors for these applications require a smooth/low-ripple magnitude estimate, since ripple adds uncertainty to the estimate.

Magnitude detectors for these types of applications are typically peak/envelope detectors which extract the envelope of the waveform by finding local maxima and then providing a slow decay between individual peaks. This operation is illustrated in Fig. 2(a), which shows numerical simulations of a peak detector responding to a speech input. As can be seen in Trace (i), a peak detector with a slow decay provides a smooth envelope of the waveform. However, the slow decay rate causes the peak detector to respond too slowly to decreases in the input's magnitude, thereby masking low-amplitude content that follows high-amplitude content. This temporal masking loses information about the signal, which is unacceptable for analysis applications. Setting a faster decay rate, as shown in Trace (ii), enables the peak detector to track the decreases in amplitude, but the output has too much ripple to be useful as a magnitude estimate. Consequently, a tradeoff exists between a smooth/low-ripple response and being able to quickly track signals so that information is not lost. To address this tradeoff, we have developed an adaptive-time-constant low-pass filter that is able to operate on a high-ripple output of a peak detector (similar to Trace (ii)) to provide a low-ripple output with good temporal accuracy, as shown in Trace (iii). This filtered output is a scaled version of the envelope, prompting the term magnitude detector instead of peak/envelope detector.

To illustrate that an adaptive-time-constant filter is needed to simultaneously achieve a low-ripple output and a good temporal response, we provide the numerical simulations of Fig. 2(b). This figure demonstrates the "acquire times" for different magnitude architectures, each of which provide 1% ripple. Low-ripple operation can be achieved with a standalone peak detector by slowing down its operation and using its inherent integration; however, the resulting response has a long acquire time and responds slowly to changes in the envelope of the input, as shown in Trace (2). The acquire time can be improved by cascading a low-pass filter with a peak detector, where the peak detector is biased to respond quickly to rising/falling signals and the low-pass filter is used to smooth the ripple; this response is shown in Trace (3). While the acquire time is improved for the downward step, the upward step response is limited by the filter's time constant, which is longer than the peak detector's attack time constant in order to achieve low ripple. To improve the acquisition time while still achieving low ripple, we need a filter that adjusts its time constant based on the amplitude of the
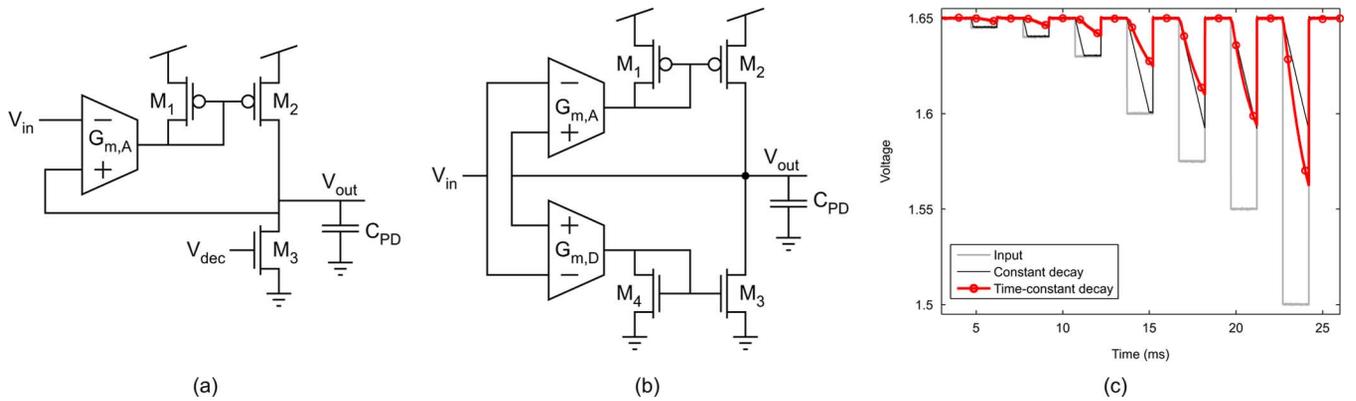
Fig. 3. (a) Peak detector with a constant decay-rate. (b) Our peak detector with a time-constant decay. (c) Comparing the constant decay-rate and time-constant decay peak detectors. For a fair comparison, the measurements were taken from circuits fabricated on a 0.5 μm CMOS process, since that is the only process in which we have fabricated the constant decay-rate peak detector.

signal. To accomplish this, we have developed a nonlinear filter with an adaptive time constant. When the amplitude changes, the integration window shrinks to track more quickly; when in the steady state, the time constant returns to a larger value to suppress the ripple. This operation is demonstrated in Trace (4) where the increasing amplitude is followed more quickly than for the linear filter in Trace (3). This nonlinearity in the filter helps the magnitude circuit to achieve better temporal responsiveness and still achieve low ripple in the steady state.

In summary, our magnitude detector [Fig. 1(a)] consists of an asymmetric integrator followed by an adaptive-time-constant filter. In the remainder of this paper, we describe these two sub-circuits in detail and show the results of the complete magnitude detector circuit.

III. PEAK DETECTOR

In this Section, we present a voltage-mode peak detector that provides an initial estimate of a signal's magnitude. This circuit has tunable attack and decay integration rates, allowing it to be used as an asymmetric integrator. By setting the attack rate faster than the decay rate, the average level of the output shifts in proportion to the amplitude of the input signal, thereby providing a measure of the signal's magnitude. Due to its tunability, this circuit can be biased to extract any fraction of the input envelope (such as the full envelope or $1/\sqrt{2}$ for RMS detection) and with any amount of ripple. In this Section, we describe the development of this peak detector circuit, analyze its operation, and provide a design procedure to allow it to be used in the larger magnitude-detector circuit.

A. Overview of the Peak Detector Circuit

Fig. 3(a) shows a common voltage-mode CMOS peak detector topology [18]–[20] that is based on a peak-detect-and-hold circuit [21] but with the reset transistor replaced by the constant-current sink, M_3. In this circuit, the current mirror half-wave rectifies the output current of the operational transconductance amplifier (OTA) onto the capacitor, causing the circuit to act as a follower when $V_{in} > V_{out}$ (i.e., the attack phase), assuming that the transconductance $G_{m,A}$ is large; reducing $G_{m,A}$ causes the circuit to act as a follower-integrator during the attack phase with an attack time constant $\tau_A = C_{PD}/G_{m,A}$. When

$V_{in} < V_{out}$ (i.e., the decay phase), no current flows through M_2, and the capacitor discharges through M_3, causing the output voltage to decrease at a constant rate. The response of the circuit to downward steps of varying sizes is shown in Fig. 3(c), which illustrates that the circuit has a constant decay-rate regardless of the step size. Unfortunately, this circuit cannot be biased for both good dynamic range and good temporal performance. For example, a very slow decay rate must be used in order to detect small signals and provide some amount of "holding" instead of simply following the small downward steps [see the first downward step in Fig. 3(c)]. However, using this same slow rate for large signals results in very long acquire times, thereby limiting the circuit's temporal response.

A precision magnitude detector, however, should produce a magnitude estimate that is amplitude invariant and has low temporal error. To accomplish these requirements, we have altered the peak detector of Fig. 3(a) to decay with a time constant, resulting in the circuit of Fig. 3(b). The operation of the lower half of the circuit mirrors that of the top half, causing the circuit to follow downward-going signals with time constant $\tau_D = C_{PD}/G_{m,D}$. The amplitude-dependent decay is demonstrated in Fig. 3(c), showing that it provides an appropriate decay over a much larger range of steps than the constant decay-rate version. The circuit is typically biased with $G_{m,A} > G_{m,D}$ to extract the positive magnitude, and it can be thought of as an asymmetric integrator due to its two different time constants, as described by the piecewise differential equation

$$C_{PD}\frac{dV_{out}}{dt} = \begin{cases} G_{m,A}(V_{in} - V_{out}), & V_{in} > V_{out} \quad \text{(attack)} \\ G_{m,D}(V_{in} - V_{out}), & \text{else} \quad\quad\;\; \text{(decay)}. \end{cases}$$
$$(1)$$

B. Peak Detector Analysis

This peak detector can be biased for different tracking levels and ripple levels through the choice of the transconductance values $G_{m,A}$ and $G_{m,D}$. We define the tracking level, $A_t = V_{out,DC}/V_{in,pk}$, as the DC output level normalized by the input amplitude. The tracking level is the magnitude metric; for example, $A_t = 1/\sqrt{2}$ is used for RMS tracking. We define the ripple ratio, $R_O = V_{out,pk}/V_{in,pk}$, as the amplitude of the

output ripple normalized by the amplitude of the input. Typically, the peak detector is tuned for 10%–30% ripple to achieve 1% ripple from the complete magnitude circuit, as discussed further in Section V. To design and bias this peak detector for use in the complete magnitude detector, we need to know how to choose the biases to achieve the specified tracking levels and ripple amounts.

To derive the dependence of both the tracking level and the ripple ratio upon the peak detector's biases, we have used the harmonic balance method [22], [23]. This procedure uses (1) modeled in the form of Fig. 4(a), and it assumes a sinusoidal input. First, the output equation is written in terms of both the tracking level and the output ripple, yielding $V_{\text{out}} = V_{\text{in,pk}} (R_O \sin(\omega t + \phi) + A_t)$, where $V_{\text{in,pk}}$ is the input amplitude and ϕ is the phase shift from the input to the output. Next, this equation is applied to (1), and the terms are balanced both at 0 Hz and at the fundamental frequency to obtain equations for the tracking level and the ripple ratio, respectively. Since the peak detector has a low-pass form, we have neglected the harmonics of the fundamental to derive an approximation that is sufficient for choosing biases and predicting the operation of the circuit.

By solving the loop at 0 Hz, the following equations were obtained for the tracking level, A_t:

$$\frac{1}{2} \ln\left(\frac{G_{m,A}}{G_{m,D}} \right) = \coth^{-1}\left(\frac{2}{\pi} \left(\frac{R_e}{A_t} \sqrt{1 - \frac{A_t^2}{R_e^2}} + \sin^{-1}\left(\frac{A_t}{R_e} \right) \right) \right) \tag{2}$$

$$R_e = \sqrt{1 + R_O^2 - R_O \cos(\phi)} \tag{3}$$

R_e is the amplitude at node e normalized by the input amplitude. As shown by (2), the tracking level depends on the ratio of transconductances and also on the output ripple; this was verified experimentally in Fig. 4(b), which shows how the tracking level varies with the attack-to-decay ratio for a fixed output ripple ratio ($R_O = 3\%$). As expected, the tracking level is zero/centered when $G_{m,A} = G_{m,D}$, and the tracking level increases as $G_{m,A}$ is increased above $G_{m,D}$, saturating as the tracking level approaches 100%.

By solving the loop at the fundamental frequency, ω, the following equation was obtained for the output ripple, R_O:

$$\frac{\omega C_{\text{PD}} R_O}{G_{m,D} R_e} = \frac{R_g + 1}{2} + \frac{1 - R_g}{\pi} \sin^{-1}\left(\frac{A_t}{R_e} \right)$$
$$+ \frac{1 - R_g}{\pi} \frac{A_t}{R_e} \sqrt{1 - \frac{A_t^2}{R_e^2}} \tag{4}$$

where $R_g = G_{m,A}/G_{m,D}$. In normal operation, the peak detector is tuned for a particular operating frequency by using the transconductances to obtain the desired ripple at that frequency. This procedure is demonstrated in Fig. 4(c), wherein (2) and (4) are used to bias the circuit for $R_O = 0.1$ (10% ripple) at three different frequencies: 500 Hz, 5 kHz, and 50 kHz. The ripple has a first-order low-pass dependence on frequency, since the circuit is a first-order asymmetric integrator. The ripple increases as the frequency decreases, until the frequency is below the corner
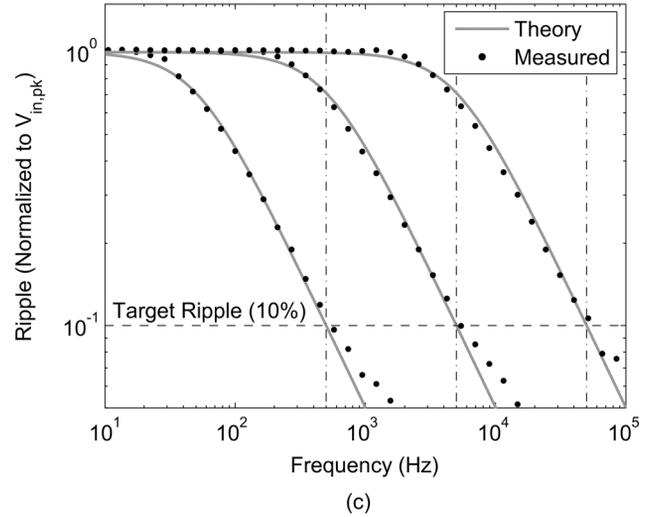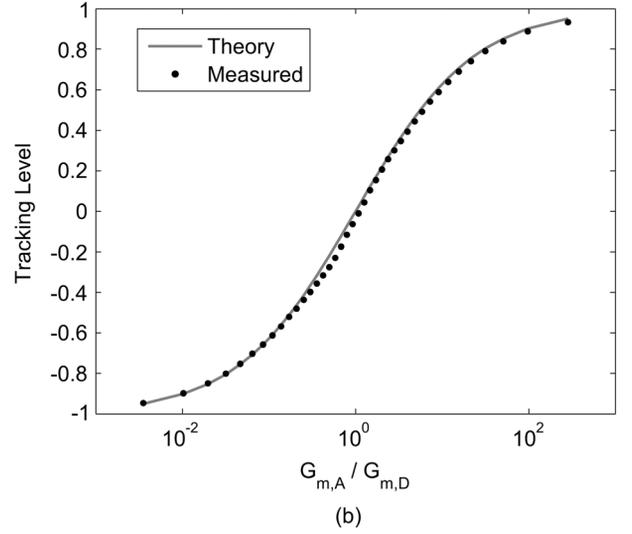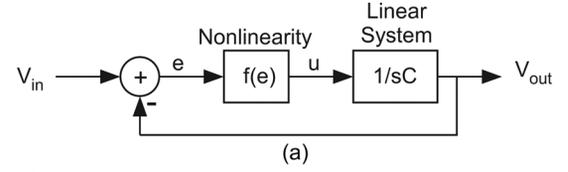
Fig. 4. (a) Model that is used for analyzing the peak detector. (b) Peak detector tracking level as a function of the attack-to-decay ratio. (c) Ripple as a function of frequency for three different biases: targeting 10% ripple at 500 Hz, 5 kHz, and 50 kHz. The data in (b) and (c) were measured with a peak detector fabricated on a 0.35 μm CMOS process, since we did not have direct access to the peak detector output with the 0.18 μm circuit.

frequency of the peak detector, at which point the peak detector acts as a follower. In order to maintain peak detector operation, the signal frequency must remain within the -20 dB/decade slope region; if the signal frequency drops too low, then the circuit no longer performs rectification. Thus, if the circuit is going to be used for broadband operation, it should be biased such that the lowest frequencies of interest remain above the corner frequency. For example, in the broadband speech demonstration of Fig. 11, the peak detector was biased for 30% ripple at 200 Hz, which is within the range of fundamental frequencies for speech [8]. If the circuit is used for subband operation, such as the filter
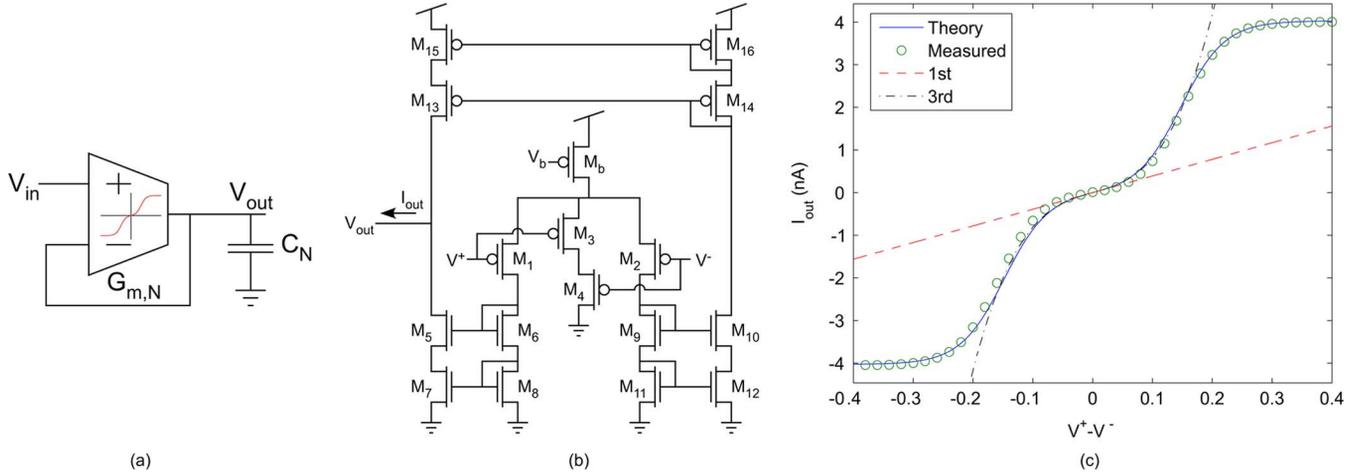
Fig. 5. (a) The adaptive-τ filter is a follower-integrator where the transconductance element has an expansive nonlinearity. (b) The nonlinear transconductance element is a pFET-based OTA with bump de-linearization. (c) I–V curve for the nonlinear OTA, shown with first- and third-order Taylor series expansions.

bank in Fig. 1(b), then each band's detector is biased to have the desired ripple at that subband frequency.

C. Peak Detector Biasing

The following is a procedure for using (2) and (4) to choose the attack and decay rates required to operate the circuit at a specified tracking level, ripple level, and input frequency.

1) Specify the tracking level A_t, ripple level R_O, and operating frequency ω. If biasing for envelope-detector operation (i.e., tracking to the top of each peak), use $A_t = 1 - R_O$;
2) Initialize $\phi = -\pi/2$;
3) Use (2) to solve for the attack-to-decay ratio $G_{m,A}/G_{m,D}$;
4) Use (4) to solve for the decay rate $G_{m,D}$;
5) Refine $\phi = \tan^{-1}(-\omega C_{\mathrm{PD}}/G_{m,D})$ and repeat steps 3–4;
6) (Optional) To bias for envelope-detector operation, use $G_{m,A}/C_{\mathrm{PD}} = 10\omega$ to ensure the output reaches the peaks with aligned phase.

IV. ADAPTIVE-TIME-CONSTANT FILTER

As discussed in Section I, the second stage of the magnitude detector integrates the first stage's initial magnitude estimate, removing the ripple that couples in from the carrier signal to produce a smooth magnitude estimate. To obtain a response with low ripple and high temporal accuracy, we have developed a low-pass filter with an expansive nonlinearity to achieve an amplitude-dependent time constant. We call this filter an adaptive-time-constant filter, or adaptive-τ filter. The expansive nonlinearity is achieved by using a nonlinear transconductance element in a follower-integrator filter topology, as shown in Fig. 5(a). The transconductance element has a sinh-shaped voltage-to-current relationship. Thus, for small differential voltages, it has a low and essentially linear transconductance, resulting in a long time constant for suppressing ripple; for large differential voltages, the transconductance increases, resulting in a shorter time constant to provide a better temporal response. In this Section, we present this adaptive-τ filter.

A. Nonlinear Transconductor

We have formed the expansive nonlinearity by using a "bump circuit" within a standard OTA [see Fig. 5(b)] [24]. Such "bump-OTAs" have been used to create linearized transconductors through appropriate sizing of the "bump" transistors, M_3 and M_4 [25], [26]. Here, we have used the bump transistors to design the cubic nonlinearity in Fig. 5(c) that is used in our adaptive-time-constant filter; a similar nonlinear transconductor was used for circuits implementing Hebbian learning [27]. In the bump-OTA, the current through the tail transistor (M_b) is shared by the input transistors (M_1 and M_2) and the bump transistors (M_3 and M_4). The current through the bump transistors is greatest when $V^+ = V^-$. By making the bump transistors have a large W/L ratio, they steal a significant amount of current from the input pair, creating a low-transconductance region for small differential voltages and generating the expansive nonlinearity. The voltage-current relationship for this circuit is described by

$$I_{\mathrm{out}} = I_b \frac{\sinh\left(\frac{\kappa}{U_T}(V^+ - V^-)\right)}{1 + S/2 + \cosh\left(\frac{\kappa}{U_T}(V^+ - V^-)\right)} \qquad (5)$$

where κ is the subthreshold slope, U_T is the thermal voltage, and the strength parameter $S = (W/L)_{3,4}/(W/L)_{1,2}$ is the relation between the aspect ratio of the bump transistors and the input pair [25]. The voltage-to-current relationship for our OTA is shown in Fig. 5(c), and the first two nonzero Taylor series coefficients are

$$a_1 = I_b \frac{\kappa}{U_T} \frac{1}{2 + S/2}, \quad a_3 = I_b \left(\frac{\kappa}{U_T}\right)^3 \frac{S/2 - 1}{6(S/2 + 2)^2}. \qquad (6)$$

These coefficients are shown with the V-I sweep in Fig. 5(c). For small differential voltages, the nonlinear-OTA acts as a linear transconductor with transconductance a_1. Increasing the input-output differential voltage increases the effective transconductance according to a_3.
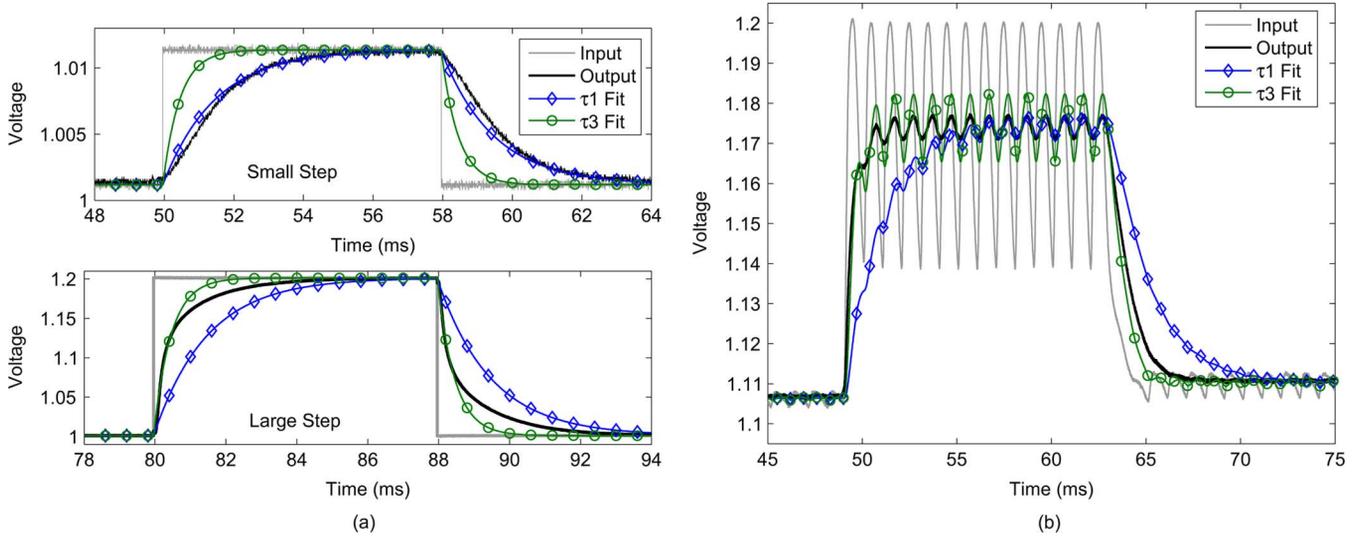
Fig. 6. (a) Measured step response of the adaptive-τ filter, shown with simulated first-order linear filters which correspond to the adaptive-τ filter's effective time constants for small and large amplitudes. (b) Measured output of the adaptive-τ filter in response to the peak detector's output, shown with the same simulated first-order linear filters as in (a).

B. Demonstration of Performance

Fig. 6(a) demonstrates the step response of the adaptive-τ filter. The response is shown for two steps of different sizes: a small step for which the linear term dominates (top pane) and a large step for which the higher-order terms dominate (bottom pane). Shown with the measured response of the adaptive-τ filter are the responses of two simulated first-order *linear* filters: the τ_1 filter has a time constant corresponding to the linearized transconductance for small signals (i.e., $\tau_1 = C_N/a_1$), and the τ_3 filter uses a shorter time constant corresponding to the linearized transconductance for the large step (i.e., the adapted time constant, $\tau_3 = C_N/(a_1 + a_3 3V_{\text{step}}^2/4)$). For the small step, the adaptive-τ filter's response follows the τ_1 filter's response since the first-order term dominates. For the large step, the adaptive-τ filter initially follows the τ_3 filter's response, but it reverts to the longer time constant of the τ_1 filter as it gets close to the final value of the step. This changing time constant helps the filter achieve a faster response for large transients.

To motivate the choice of the adaptive-τ filter for the magnitude detector, we exhibit the experiment in Fig. 6(b), which compares the performance of the adaptive-τ filter with the two simulated first-order linear filters used in the experiment of Fig. 6(a). The input to the filters is the response of the peak detector to a sine wave stepped from 5 mV$_{\text{pk}}$ to 100 mV$_{\text{pk}}$ and then to 10 mV$_{\text{pk}}$. The τ_1 filter yields the same ripple as the adaptive-τ filter but cannot follow the steps closely in time. The τ_3 filter follows the steps but has more ripple than the adaptive-τ filter. These results show that the adaptive-τ filter achieves a good tradeoff between ripple suppression and temporal response, while also being compact and low-power.

C. Design

To design the adaptive-τ filter, we need to know how its time constant depends on the input amplitude, bias current, and strength parameter (S). Here we develop an approximation to relate those parameters and then show how to use this

approximation to design and bias the circuit. Our approximation is based on the describing function [22] of the nonlinear transconductance element. The filter has the form of Fig. 4(a) with $f(V_{\text{in}} - V_{\text{out}}) = a_1(V_{\text{in}} - V_{\text{out}}) + a_3(V_{\text{in}} - V_{\text{out}})^3$, where a_1 and a_3 are the Taylor series coefficients given by (6) and are controlled by the bias current and the strength parameter. The sinusoidal-input-describing-function (i.e., an amplitude-dependent gain term) for this nonlinearity is $a_1 + a_3 3|V_{\text{in}} - V_{\text{out}}|^2/4$ [28]. Knowing that the input-output differential is related to the input amplitude $V_{\text{in,pk}}$, the transfer function can be approximated as

$$H(s, V_{\text{in,pk}}) \approx \frac{1}{1 + sC_N \Big/ \Big(a_1 + a_3 \frac{3}{4}V_{\text{in,pk}}^2\Big)} \qquad (7)$$

where the corner frequency has a quadratic relation to the input amplitude. Equation (7) gives an approximate transfer function for the circuit and is verified with real data in Fig. 7. Fig. 7(a) shows the frequency response measured at different input amplitudes and demonstrates an increasing corner frequency for increasing amplitude. Fig. 7(b) shows the variation in corner frequency as a function of input amplitude for two different filter biases. The circles show the measured corner frequencies and the solid lines show the predicted values using $2\pi f_c = (a_1 + a_3 3V_{\text{in,pk}}^2/4)/C_N$ [i.e., the corner frequency in (7)]. This experiment verifies that the corner frequency is a function of the square of the input amplitude.

Using (7) and the definitions of a_1 and a_3, the circuit can be designed to exhibit the desired ripple suppression and transient response. The procedure to design and bias the filter is as follows:

1) Specify (a) the ripple-suppression time constant, τ_{rip}, (b) the transient-response time constant, τ_{tran}, and (c) the amplitude that is considered a transient (and should be followed with the transient-response time constant τ_{tran}), A_{tran}.

2) Use the ripple-suppression time constant to compute $a_1 = C_N/\tau_{\text{rip}}$.

(a)   (b)
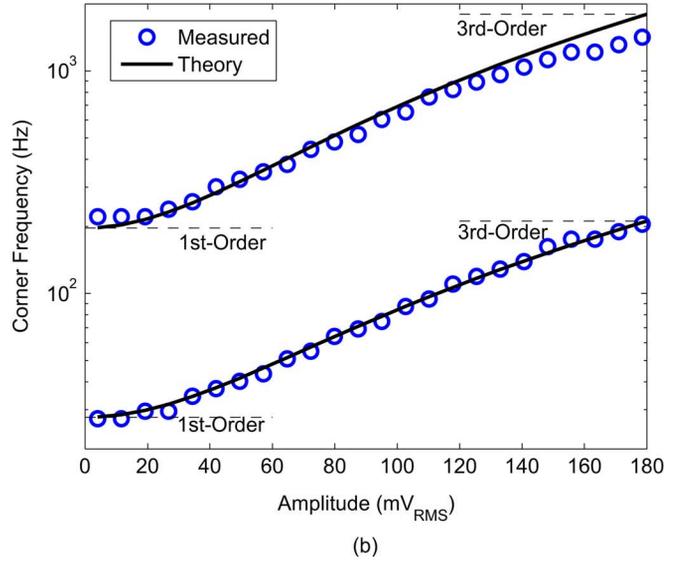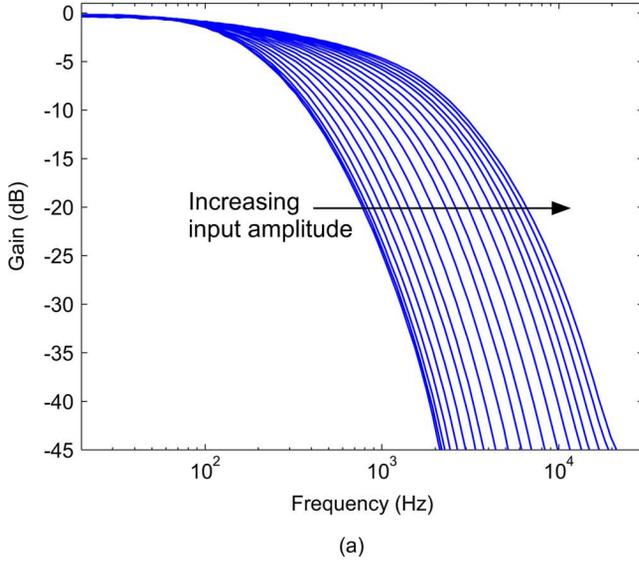
Fig. 7. (a) Measured frequency response of the adaptive-τ filter. Each line is a frequency response for a different input amplitude. (b) Corner frequency of the adaptive-τ filter as a function of the input amplitude, shown for two different OTA biases.
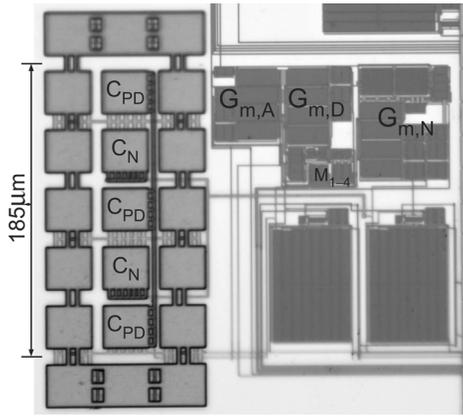
Fig. 8. Micrograph of the magnitude circuit, which was fabricated in a standard 0.18 μm CMOS process.

Fig. 9. Dynamic range measurement of the complete magnitude circuit. The response remains linear to within 1 dB across a range of 70 dB.

3) Use a_1 to find the value of a_3 that yields the desired transient-response time constant at an amplitude of A_{tran} by using $a_3 = 4(C_N/\tau_{\mathrm{tran}} - a_1)/(3A_{\mathrm{tran}}^2)$.

4) Use (6) to find S in terms of a_1 and a_3 by using

$$S = 2\frac{(\kappa/U_T)^2 + 12a_3/a_1}{(\kappa/U_T)^2 - 6a_3/a_1}.$$

5) Use (6) to find the bias current $I_b = a_1(2 + S/2)U_T/\kappa$.

V. COMPLETE MAGNITUDE CIRCUIT

In this section, we present the complete magnitude detector circuit, shown schematically in Fig. 1(a). This circuit is formed by combining the peak detector of Section III with the adaptive-τ filter of Section IV. Fig. 8 shows a micrograph of the circuit, which was fabricated in a standard 0.18 μm CMOS process. The values used for capacitors C_{PD} and C_N are 4.5 pF and 3 pF, respectively. The total area of the circuit is 0.019 mm^2.

As discussed in Section II, our objective is to develop a magnitude detector with improved temporal accuracy, while maintaining a low-ripple response. This tradeoff between ripple and

temporal accuracy can be understood graphically from Fig. 4(c). For example, a standalone peak detector can be tuned for low ripple by biasing the circuit's corner frequency, ω_c, to be much less than the operational frequency, ω, which is the frequency for the desired target ripple. This biasing results in a long acquire time since the acquire time constant, τ_{aq}, is given by $\tau_{\mathrm{aq}} = 1/\omega_c$ and is long compared to the input signal's period at the operational frequency.

The $\omega\tau_{\mathrm{aq}}$ product is related to the number of cycles required to acquire transients and is thus a good metric for temporal responsiveness. In Table I, we compare $\omega\tau_{\mathrm{aq}}$ for different magnitude-detection architectures for a given total ripple at the output, R_T. The third column states the factor of improvement for each architecture with 1% total ripple over the baseline case of the standalone peak detector (PD). For example, a PD combined with a low-pass filter (LPF) is 7.1 times faster than a PD by itself. This improvement is because adding an LPF to the PD
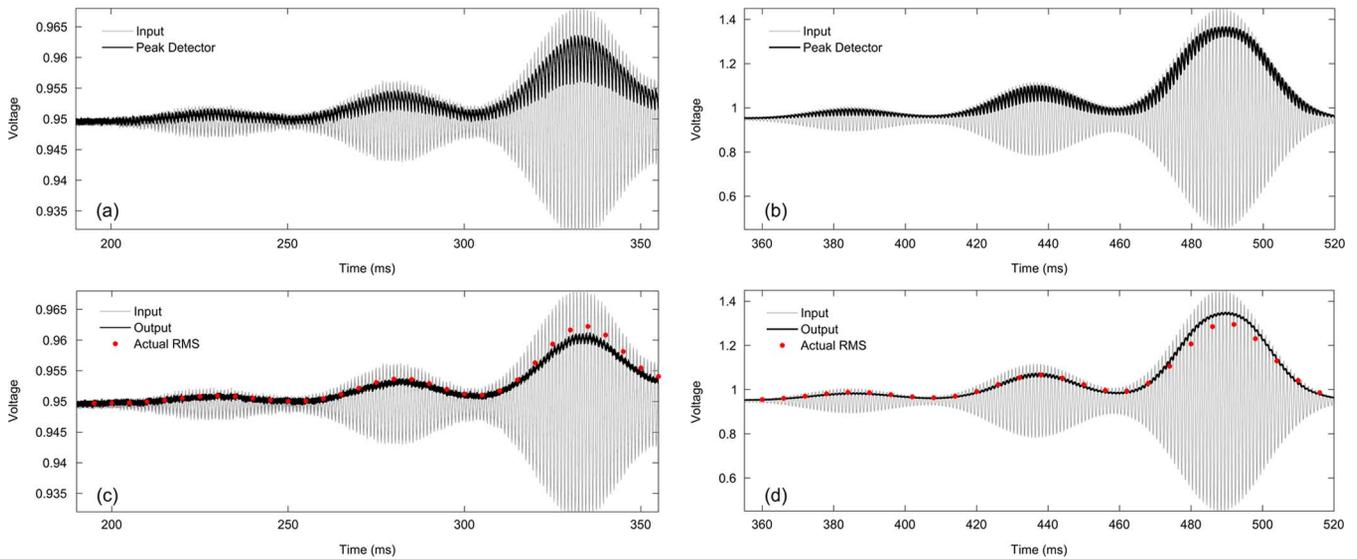
Fig. 10. Measured transient response of the magnitude detector for logarithmically increasing amplitude from 2 mV$_{pk}$ to 500 mV$_{pk}$.
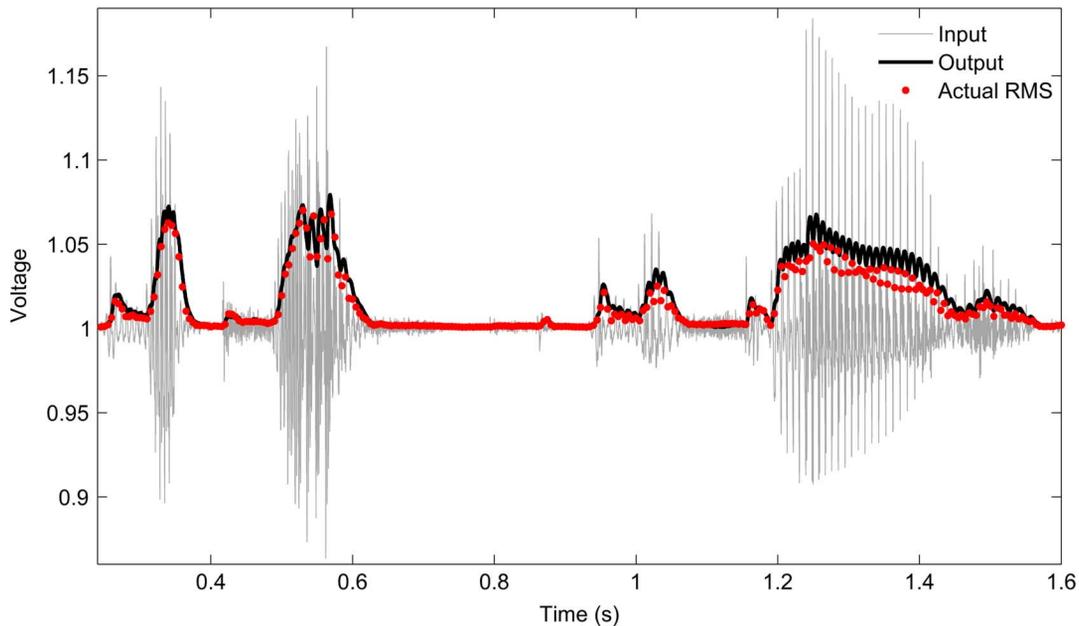
Fig. 11. Measured response of the magnitude circuit to a speech waveform. For comparison, the response is shown with a computer calculated RMS.

TABLE I
TRADEOFF BETWEEN RIPPLE AND ACQUISITION TIME

Magnitude Structure	$\omega\tau_{aq}$	$\omega\tau_{aq}$ Improvement (1% Ripple)
PD Only	$1/R_T$	1
PD w/ LPF	$\sqrt{2}/R_T$	7.1
PD w/ 2nd-Order LPF	$\sqrt{3}/\sqrt[3]{R_T}$	11.3
PD w/ Adaptive-τ	$\sqrt{\dfrac{1}{4R_T} + \dfrac{4}{R_T\left(1 + A_s^2\frac{3a_3}{4a_1}\right)^2}}$	17.7

increases the slope of the ripple-frequency relationship [i.e., the slope of Fig. 4(c)], thereby moving ω_c closer to ω; the best result is obtained by splitting the ripple suppression evenly between the PD and the linear LPF (e.g., 10% suppression in each stage

to obtain 1% total suppression). Using a second-order filter further improves the temporal accuracy by increasing the slope of the ripple-frequency relationship. However, this technique only achieves 1.6-times improvement beyond the first-order filter, while adding significant increases in power, area, and frequency sensitivity (due to the larger slope). With the adaptive-τ filter, we operate the peak detector faster than with the linear filter (typically twice as fast) which, accordingly, yields more ripple at the PD's output. To achieve the same total ripple at the output of the magnitude detector, the value of a_1/C_N in the adaptive-τ filter is tuned to compensate for the increased PD speed. Large transients then cause the filter's time constant to decrease such that the PD's speed is the main limitation to the magnitude detector's overall speed. As can be seen from the table, the

TABLE II
COMPARISON

	Proposed	Zhak [29]	Steyaert [30]	Zhou [31]	Rodriguez-Villegas [32]	Alegre [33]
Technology	0.18μm	1.5μm	1μm	0.18μm	0.35μm	0.35μm
Supply Voltage	1.8V	2.8V	3V	1.8V	[1V 1.5V]	
Frequency Range	20Hz–20kHz	100Hz–10kHz	12MHz	100Hz–1.6GHz	100Hz	1–10MHz
Dynamic Range	70dB (200Hz) 70dB (2kHz) 64dB (20kHz)	75dB	54dB	50dB	[73dB 83dB]	42dB
Power	1.1nW–1.08μW	2.8μW	10mW	6.3mW	[60nW 90nW]	2.98mW
FOM (x10^{12})	57.9	20.1	0.60	80.3	[7.44 15.7]	0.42
Comments	DR <1dB nonlinearity				Rectifier Only	Simulated; DR <1dB nonlinearity

adaptive-τ filter yields a 2.5-times improvement in temporal response for an amplitude step (A_s) of 300 mV as compared to the linear filter and without the cost of the second-order filter.

Furthermore, we have compared the acquisition times—defined as the number of cycles to reach 99% of a step—of the adaptive-τ ripple suppression case with the first-order linear filtering case, for steps between the minimum (200 μV$_{\rm RMS}$) and maximum (630 mV$_{\rm RMS}$) detectable signals of this circuit. For an upward step, the acquistion time improves from 7.58 cycles with a linear filter to 1.45 cycles with the adaptive-τ filter; and for a downward step, the acquisition time improves from 18.3 cycles with a linear filter to 14.1 cycles with the adaptive-τ filter.

Fig. 9 shows the measured dynamic range of the magnitude circuit for operational frequencies of 200 Hz and 2 kHz (i.e., typical speech frequencies). The biasing routines discussed in Sections III-C and IV-C were used to bias the circuit to track the RMS with 1% ripple. The response remains within 1 dB linearity from 200 μV$_{\rm RMS}$ to 630 mV$_{\rm RMS}$ for both cases, yielding a dynamic range of 70 dB. These two measurements are characteristic of the performance of this circuit for typical speech frequencies, including those covering the telephony frequency range (i.e., $\leq$5 kHz). Additionally, the dynamic range was measured to be 64 dB at the upper end of the audio frequency range (20 kHz).

The minimum detectable signal is limited by the greater of two nonidealities: the noise or the "deadzone" (which is created by offsets in the peak detector's OTAs). Since the magnitude detector can easily be designed to achieve a noise floor below typical offsets for this type of OTA architecture, the designer should focus on the offsets. For example, our complete magnitude detector was measured to have $\leq$115 μV$_{\rm RMS}$ of noise for all biasings across the audio frequency range, which is well below the minimum detectable signal of 200 μV$_{\rm RMS}$, indicating that offsets dominated our minimum detectable signal. Offsets in the peak detector OTAs affect overall circuit operation as follows. Due to the symmetry of the peak detector, its minimum detectable signal is not degraded when $G_{m,A}$ and $G_{m,D}$ have equal offsets with the same sign. The minimum detectable signal is only degraded by a mismatch between the offsets, which will either create a gap between the attack and decay states or will cause an overlap of the attack and decay states; both cases compromise the accuracy of the magnitude estimate for signals

smaller than the difference between the offsets. Thus, the most critical factors for improving the minimum detectable signal are minimizing the offsets and matching the peak detector OTAs to ensure similar offsets with low variance.

To verify the simultaneous achievement of temporal and amplitude accuracy, we performed the experiments of Figs. 10–11. In both experiments, the circuit was biased for RMS tracking and 1% ripple. In Fig. 10, a sine wave with a frequency of 2 kHz had its amplitude modulated by six Gaussian pulses with amplitudes increasing logarithmically from 2 mV$_{\rm pk}$ to 500 mV$_{\rm pk}$. The first three pulses are shown in the left panes and the last three pulses are shown in the right panes with different y-axis limits. The top panes show the peak detector response and the bottom panes show the magnitude response along with the actual RMS. Fig. 11 shows the response of the magnitude circuit to a speech signal. A mathematically calculated RMS is shown alongside the circuit's response. In both Figs. 10 and 11, we see that the magnitude circuit closely follows the same shape as the RMS calculations, and that the magnitude detector accurately follows the quickly changing RMS across a wide dynamic range.

The characteristics of this circuit are summarized in Table II and are compared with relevant state-of-the-art magnitude detectors. The power consumption of our circuit scales linearly with the operational frequency (since transconductance scales linearly with current in subthreshold operation); for typical biasing, the power varies from 1.1 nW–1.08 μW across the 20 Hz–20 kHz audio frequency range. The figure of merit (FOM) in the table is defined as FOM $= 10^{\rm DR/20} f_{\rm max}/P$, where DR is the dynamic range, $f_{\rm max}$ is the maximum operational frequency, and P is the power consumption. Our FOM number is given for a frequency of 200 Hz, since that is the tuning condition used to operate on wideband speech signals.

VI. CONCLUSION

In this paper, we have presented a low-power magnitude detector circuit, which achieves good temporal responsiveness through the use of a novel peak-detector-nonlinear-integrator topology. The circuit was built in a 0.18 μm CMOS process. At 200 Hz, which is a typical operating point for wideband speech signals [8], the circuit achieves a dynamic range of 70 dB with a power consumption of 10.92 nW. The compactness and

low-power operation of the circuit, combined with its flexible biasing, make it a good choice for applications such as spectral analysis.

REFERENCES

[1] F. Yuan, "Design techniques for ASK demodulators of passive wireless microsystems: a state-of-the-art review," *Analog Integr. Circuits Signal Process.*, vol. 63, pp. 33–45, 2010.

[2] M. Baker and R. Sarpeshkar, "Low-power single-loop and dual-loop AGCs for bionic ears," *IEEE J. Solid-State Circuits*, vol. 41, no. 9, pp. 1983–1996, Sep. 2006.

[3] D. Graham, P. Smith, R. Chawla, and P. Hasler, "A programmable bandpass array using floating-gate elements," in *Proc. IEEE Int. Symp. Circuits and Systems (ISCAS)*, Vancouver, BC, Canada, May 2004, vol. 1, pp. I-97–I-100.

[4] E. Fragniére, "A 100-channel analog CMOS auditory filter bank for speech recognition," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2005, vol. 1, pp. 140–141.

[5] B. Wen and K. Boahen, "A 360-channel speech preprocessor that emulates the cochlear amplifier," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2006, pp. 2268–2277.

[6] S.-C. Liu, A. van Schaik, B. Minch, and T. Delbrück, "Event-based 64-channel binaural silicon cochlea with Q enhancement mechanisms," in *Proc. IEEE Int. Symp. Circuits and Systems (ISCAS)*, May 2010, pp. 2027–2030.

[7] R. Sarpeshkar, C. Salthouse, J.-J. Sit, M. Baker, S. Zhak, T.-T. Lu, L. Turicchia, and S. Balster, "An ultra-low-power programmable analog bionic ear processor," *IEEE Trans. Biomed. Eng.*, vol. 52, no. 4, pp. 711–727, Apr. 2005.

[8] J. Picone, "Signal modeling techniques in speech recognition," *Proc. IEEE*, vol. 81, no. 9, pp. 1215–1247, Sep. 1993.

[9] J. Georgiou and C. Toumazou, "A 126-μW cochlear chip for a totally implantable system," *IEEE J. Solid-State Circuits*, vol. 40, no. 2, pp. 430–443, Feb. 2005.

[10] B. Rumberg, D. Graham, V. Kulathumani, and R. Fernandez, "Hibernets: Energy-efficient sensor networks using analog signal processing," *IEEE J. Emerging and Selected Topics in Circuits and Systems*, vol. 1, no. 3, pp. 321–334, Sep. 2011.

[11] J. Alegre, S. Celma, B. Calvo, and J. G. del Pozo, "A novel CMOS envelope detector structure," in *Proc. IEEE Int. Symp. Circuits Syst. (ISCAS)*, 2007, pp. 3538–3541.

[12] P. Hasler, P. Smith, D. Graham, R. Ellis, and D. Anderson, "Analog floating-gate, on-chip auditory sensing system interfaces," *IEEE Sensors J.*, vol. 5, pp. 1027–1034, Oct. 2005.

[13] X. Arreguit and E. Vittoz, "Perception systems implemented in analog VLSI for real-time applications," in *Proc. Perception to Action Conf.*, Lausanne, Switzerland, 1994, pp. 170–180.

[14] P. Hasler and D. Anderson, "Cooperative analog-digital signal processing," in *Proc. IEEE Conf. Acoustics, Speech, and Signal Processing*, Orlando, FL, May 2002, vol. 4, pp. 3972–3975.

[15] R. Edwards and G. Cauwenberghs, "Log-domain circuits for auditory signal processing," in *Proc. IEEE Midwest Symp. Circuits and Systems*, Aug. 1999, vol. 2, pp. 968–971.

[16] S. Haddad and R. H. W. Serdijn, "Analog wavelet transform employing dynamic translinear circuits for cardiac signal characterization," in *Proc. IEEE Int. Symp. Circuits and Systems (ISCAS)*, May 2003, vol. 1, pp. I-121–I-124.

[17] R. Ellis, H. Yoo, D. Graham, P. Hasler, and D. Anderson, "A continuous-time speech enhancement front-end for microphone inputs," in *Proc. IEEE Int. Symp. Circuits and Systems (ISCAS)*, May 2002, vol. 2, pp. II-728–II-731.

[18] S. Ravindran, C. Demiroglu, and D. Anderson, "Speech recognition using filter-bank features," in *Proc. IEEE Asilomar Conf. Signals, Systems, and Computers*, Pacific Grove, CA, Nov. 2003, vol. 2, pp. 1900–1903.

[19] H. Abdalla and T. Horiuchi, "An analog VLSI low-power envelope periodicity detector," *IEEE Trans. Circuits Syst. I*, vol. 52, pp. 1709–1720, 2005.

[20] S.-B. Park, J. Wilson, and M. Ismail, "Peak detectors for multistandard wireless receivers," *IEEE Circuits and Devices Mag.*, vol. 22, no. 6, pp. 6–9, Nov.–Dec. 2006.

[21] M. Kruiskamp and D. Leenaerts, "A CMOS peak detect sample and hold circuit," *IEEE Trans. Nucl. Sci.*, vol. 41, no. 1, pp. 295–298, Feb. 1994.

[22] J. Taylor, "Describing functions," *Electrical and Electronics Engineering Encyclopedia*, pp. 77–98, 2000, Supplement 1.

[23] R. Gilmore and M. Steer, "Nonlinear circuit analysis using the method of harmonic balance—A review of the art. Part I. Introductory concepts," *Int. J. Microw. Millimeter-Wave Computer-Aided Eng.*, vol. 1, no. 1, pp. 22–37, 1991.

[24] T. Delbrück, ""Bump" circuits for computing similarity and dissimilarity of analog voltages," in *Proc. Int. Joint Conf. Neural Networks*, Seattle, WA, Jul. 1991, pp. I-475–I-479.

[25] R. Sarpeshkar, R. Lyon, and C. Mead, "A low-power wide-linear-range transconductance amplifier," *Analog Integr. Circuits Signal Process.*, vol. 13, pp. 123–151, 1997.

[26] P. Furth and H. Ommani, "Low-voltage highly-linear transconductor design in subthreshold CMOS," in *Proc. IEEE Midwest Symp. Circuits and Systems*, Sacramento, CA, Aug. 1997, vol. 1, pp. 156–159.

[27] M. Cohen and A. Andreou, "MOS circuit for nonlinear Hebbian learning," *Electron. Lett.*, vol. 28, no. 9, pp. 809–810, 1992.

[28] A. Gelb and W. Vander Velde, *Multiple-Input Describing Functions and Nonlinear System Design*. New York: McGraw-Hill, 1968.

[29] S. Zhak, M. Baker, and R. Sarpeshkar, "A low-power wide dynamic range envelope detector," *IEEE J. Solid-State Circuits*, vol. 38, no. 10, pp. 1750–1753, Oct. 2003.

[30] M. Steyaert, W. Dehaene, J. Craninckx, M. Walsh, and P. Real, "A CMOS rectifier-integrator for amplitude detection in hard disk servo loops," *IEEE J. Solid-State Circuits*, vol. 30, no. 7, pp. 743–751, Jul. 1995.

[31] Y. Zhou, G. Huang, S. Nam, and B.-S. Kim, "A novel wide-band envelope detector," in *Proc. IEEE Radio Frequency Integrated Circuits Symp.*, Jun. 2008, pp. 219–222.

[32] E. Rodriguez-Villegas, P. Corbishley, C. Lujan-Martinez, and T. Sanchez-Rodriguez, "An ultra-low-power precision rectifier for biomedical sensors interfacing," *Sensors and Actuators A: Physical*, vol. 153, no. 2, pp. 222–229, 2009.

[33] J. Alegre, S. Celma, J. G. del Pozo, and N. Medrano, "Fast-response low-ripple envelope follower," *Integration, the VLSI J.*, vol. 42, no. 2, pp. 169–174, 2009.

Brandon Rumberg (S'08) received the B.S. degrees in electrical engineering and computer engineering in 2007, and the M.S. degree in electrical engineering in 2009, from West Virginia University, Morgantown, WV, where he is currently working toward the Ph.D. degree in electrical engineering.

His research interests include low-power sensing systems, mixed-signal integrated circuits, and extending circuit-level innovations to benefit system-level applications.

David W. Graham (S'00–M'07) received the B.A. degree in natural science from Covenant College, Lookout Mountain, GA, in 2001. He also received the B.S. degree in electrical engineering, and the M.S. and Ph.D. degrees in electrical and computer engineering all from Georgia Institute of Technology, Atlanta, GA, in 2001, 2003, and 2006, respectively.

He is an Assistant Professor in the Lane Department of Computer Science and Electrical Engineering, West Virginia University, Morgantown, WV. His research interests are in developing biologically inspired electronics, cooperative analog and digital signal-processing systems, programmable analog devices, and low-power electronics.