

An Ultra-Low-Power Programmable Analog Bionic Ear Processor

Rahul Sarpeshkar*, Christopher Salthouse, Ji-Jon Sit, Michael W. Baker, Serhii M. Zhak, Timothy K.-T. Lu, Lorenzo Turicchia, and Stephanie Balster

Abstract—We report a programmable analog bionic ear (cochlear implant) processor in a 1.5- μm BiCMOS technology with a power consumption of 211 μW and 77-dB dynamic range of operation. The 9.58 mm \times 9.23 mm processor chip runs on a 2.8 V supply and has a power consumption that is lower than state-of-the-art analog-to-digital (A/D)-then-DSP designs by a factor of 25. It is suitable for use in fully implanted cochlear-implant systems of the future which require decades of operation on a 100-mAh rechargeable battery with a finite number of charge-discharge cycles. It may also be used as an ultra-low-power spectrum-analysis front end in portable speech-recognition systems. The power consumption of the processor includes the 100 μW power consumption of a JFET-buffered electret microphone and an associated on-chip microphone front end. An automatic gain control circuit compresses the 77-dB input dynamic range into a narrower internal dynamic range (IDR) of 57 dB at which each of the 16 spectral channels of the processor operate. The output bits of the processor are scanned and reported off chip in a format suitable for continuous-interleaved-sampling stimulation of electrodes. Power-supply-immune biasing circuits ensure robust operation of the processor in the high-RF-noise environment typical of cochlear implant systems.

Index Terms—Bionic ear, cochlear implant, hearing aids, low power, spectrum analysis, speech recognition.

I. INTRODUCTION

COCHLEAR implants or bionic ears (BEs) restore hearing in profoundly deaf patients. They function by transforming frequency patterns in sound into corresponding spatial electrode-stimulation patterns for the auditory nerve. Over the past 20 years, improvements in sound-processing strategies, in the number of electrodes and channels, and in the rate of stimulation have yielded improved sentence and word recognition scores in patients [1]. Next-generation implants will be fully implanted inside the body of the patient and consequently have very stringent requirements on the power consumption used for signal processing. Our processor is intended for use in such next-generation implants. It can operate on a 100-mAh battery with a 1000 charge-and-discharge-cycle lifetime for 30 years. The digital outputs of the processor, its immunity to power-supply noise, and its programmability ensure ease of use

with the other parts of an implant system such as its wireless communication link and programming interface.

In addition to BEs, low-power processors such as the one described in this paper are useful in portable speech-recognition front ends of the future. Such systems will need programmable front-ends that take a microphone input, and output bits that represent spectral information.

A common speech-processing strategy, used in implants and in speech-recognition systems, employs a mel cepstrum filter bank with 8–20 channels. The mel scale maps frequencies to a perceptually linear scale [2]. Filter banks based on the mel scale use linearly spaced filter center frequencies up to 1 kHz and logarithmically spaced center frequencies above 1 kHz. In the ubiquitous cepstral techniques, a logarithmic measure of the spectral energy in each filter bank channel is used for further processing. In implants, the number of functioning electrodes used for stimulation is also between 8–20 and having more electrodes is often not useful due to spatial interactions amongst the electrodes.

This paper describes a programmable processor for BEs and speech-recognition front ends. The signal-processing blocks of the processor are shown in Fig. 1. A microphone pre-amplifier, audio front end (AFE) in Fig. 1, converts sound over a 77-dB dynamic range into an electrical input that is fed to an automatic-gain-control circuit (AGC). The AGC compresses this input dynamic range to a narrower 57-dB internal dynamic range (IDR) at which each of the 16 channels of the processor operate. The 57-dB IDR fulfills the 50- to 60-dB IDR requirement for good patient performance [3]. Each channel is made up of a programmable Gm-C bandpass filter, an envelope detector with programmable attack and release time constants, and a logarithmic dual-slope A/D converter with programmable calibration, sampling, and precision parameters. The output bits of the A/D converters are sequentially scanned via a shift register, which is clocked by an input electrode-stimulation clock or continuous interleaved sampling (CIS) clock. These bits are fed to a bus whose data constitutes the final output of the processor. Under normal operation, the output bits of the processor would be latched into digital-to-analog converters (DACs) that control electrode stimulation currents.

Current reference circuits provide constant- g_m and power-supply-noise-immune biasing throughout the chip via a current-distribution network. Current-mode DAC converters that are necessary for altering the various parameters of the chip are also biased with these reference currents and implement the programmability of the system. Each channel of the chip is independently programmed via a shift-register-based multiplexing system that operates on a programmability-clock during patient programming. Several crucial voltage and current variables in the channel may be observed through visibility follower circuits

Manuscript received February 26, 2004; revised September 12, 2004. *Asterisk indicates corresponding author.*

*R. Sarpeshkar is with the Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: rahuls@avns1.mit.edu).

C. Salthouse, J.-J. Sit, M. W. Baker, S. M. Zhak, T. K.-T. Lu, L. Turicchia, and S. Balster are with the Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, MA 02139 USA.

Digital Object Identifier 10.1109/TBME.2005.844043

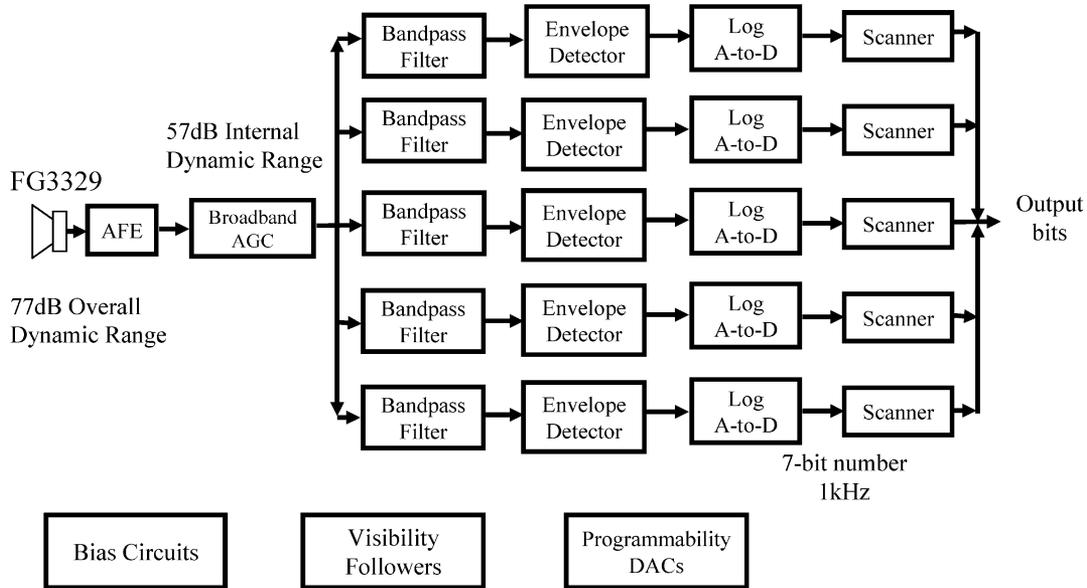


Fig. 1. Overall architecture of the analog BE processor.

that are controlled by the same shift-register multiplexing bits. The visibility of the processor is important for debugging any issues having to do with poor subject performance in a highly variable patient population.

The ultra-low-power specifications of our processor are in large part due to simple but powerful innovations in the design of its building-block analog circuits, some of which have been recently described in detail. These circuits include the microphone front end [4], the bandpass filters [5], the envelope detector [6], and the logarithmic analog-to-digital converter (ADC) [7]. Although some parts of the processor have been reported on, this paper describes in detail how these parts and other new circuit building-blocks all fit together to create a robust, programmable, ultra-low-power, mixed-signal system. We present experimental data from the overall system that verifies its operation. Several simple-but-important issues in the design of a large-scale mixed-signal chip such as this one that are crucial for maintaining robust operation without sacrificing efficiency are discussed.

Subthreshold-MOS, silicon cochlea, and analog circuits for cochlear-implant processing have been previously proposed [8]–[14] as a means for implementing complex signal processing with very low power. This paper proves the promise of such prior work by achieving numbers that make an analog processing system the most power-efficient option in the short and long term, while still preserving the very desirable features of high IDR/channel, programmability, and robustness present in digital systems. This processor presents an example of how deferring digitization by first doing analog preprocessing can lead to tremendous gains in power efficiency (the processor is effectively operating at $5 \mu\text{W}/\text{MIP}$ versus a DSP at $250 \mu\text{W}/\text{MIP}$) because of the exploitation of analog primitives for computation [15].

The organization of this paper follows the signal-flow in the chip: Section II describes the microphone front end of the processor. Section III describes the analog gain-control circuit. Sections IV-A-C describe the signal processing of a single channel, namely, the bandpass filtering, the envelope detection,

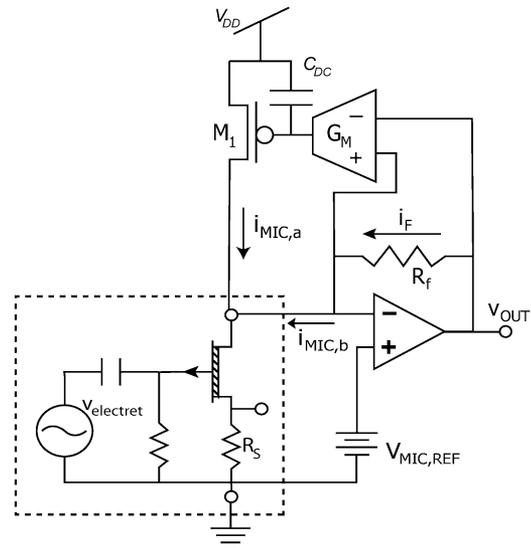


Fig. 2. Analog front end (AFE) with a microphone-current-sense topology and low-frequency feedback.

and the logarithmic A/D conversion. Section V describes the current and voltage biasing scheme used in the processor. Section VI discusses the digital control, programming, and I/O interface of the processor. Section VII presents measured experimental results for the overall operation of the processor. Section VIII compares our processor to more traditional A/D-then-DSP implementations and provides a discussion of why preliminary digitization of the data is not a wise strategy for power-efficient architectures. Section IX presents further improvements that are possible in the processor. Section X concludes by summarizing the key contributions of the paper.

II. THE MICROPHONE FRONT END

Fig. 2 shows the built-in microphone front-end circuit (the AFE circuit) that is used in our processor. A commercial JFET-buffered FG3329 electret microphone transduces sound into an

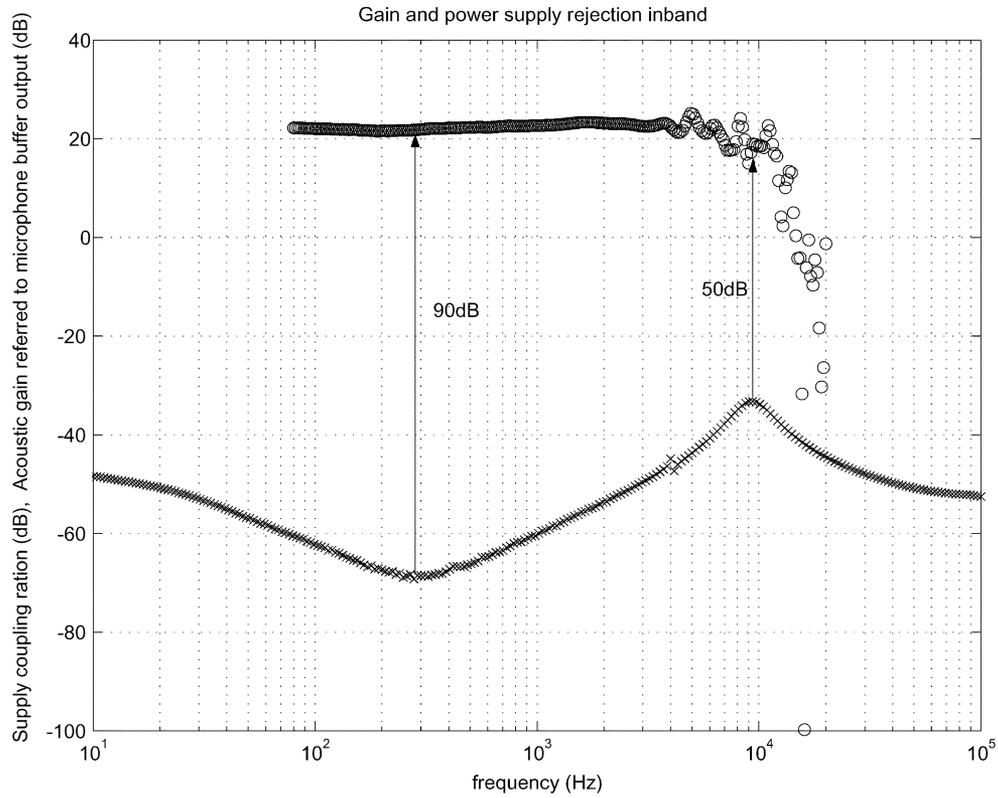


Fig. 4. AFE frequency response and supply rejection performance to 5 MHz.

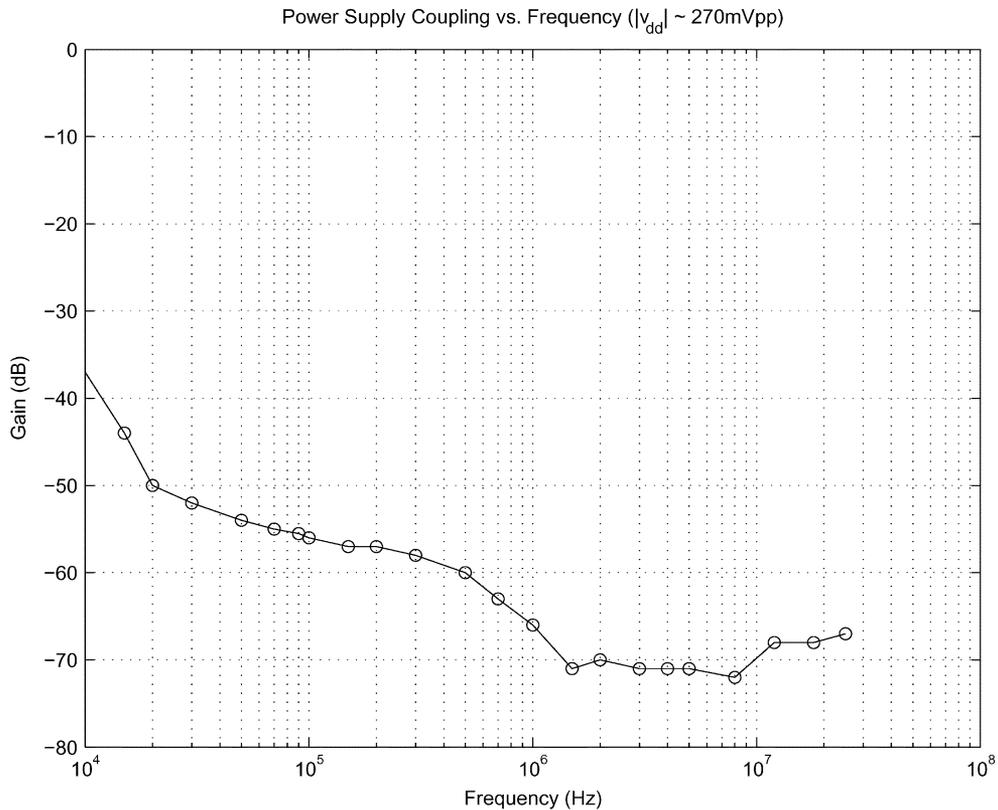


Fig. 5. Microphone front end power-supply Insensitivity to 26 MHz.

shows that the VGA resistance is actually implemented as a transconductance G_{m2} such that the gain of the circuit is given by G_{m1}/G_{m2} . The transconductors G_{m1} and G_{m2} both use

the wide-linear-range (WLR) transconductor circuit described in [16]. The voltage reference inputs were biased very near the microphone reference voltage used in Fig. 2.

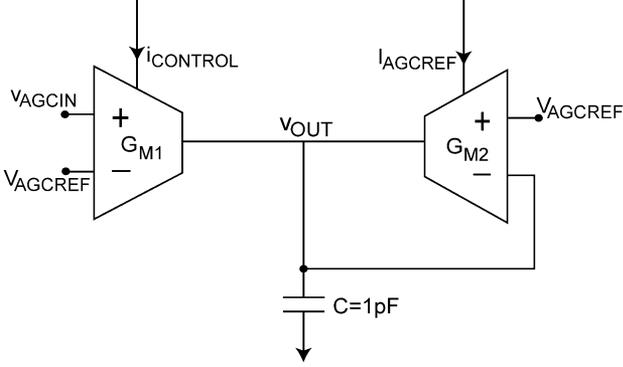


Fig. 6. The VGA.

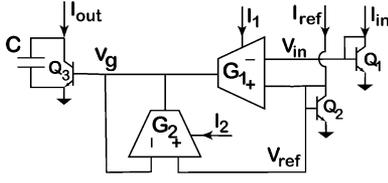


Fig. 7. The translinear compression circuit used for gain control.

The gain regulation of the AGC is performed by the circuit shown in Fig. 7: The input I_{in} is generated by an envelope detector whose output current is linearly proportional to the amplitude of the VGA's *output* voltage. The operation of the envelope detector is identical to that used in the envelope detector described in the single-channel processing circuits of Section IV. Thus, we shall delay the description of its operation to Section IV. A current I_g , proportional to $\ln(I_{in}/I_{ref})$, is output by the transconductance amplifier with bias current I_1 and converted to a voltage V_g on the effective load resistance created by the transconductance amplifier with bias current I_2 . The voltage V_g is exponentiated by the output transistor to create I_{out} , which is conveyed through a cascode transistor and current mirror to create $I_{control}$, the gain-controlling current of the VGA of Fig. 6. The capacitor C in Fig. 7 is necessary for achieving a small value of thermal voltage noise and, consequently, a small value of error in the gain fluctuations of the VGA. Since we desire the output dynamic range of the overall AGC to be around 60 dB, these fluctuations must be significantly less than 1 part in a 1000.

Since the I_1 and I_2 transconductors have the same linear range, the ratio I_1/I_2 is the same as the ratio of their transconductances, and determines the overall compression ratio of the AGC circuit. A simple mathematical analysis reveals that the circuit of Fig. 7 implements the equation

$$I_{out} = I_{ref}(I_{in}/I_{ref})^{-I_1/I_2} \quad (1)$$

which, when applied to the VGA of Fig. 6, yields a power-law relationship between the input and output of the VGA of

$$V_{OUT} \propto V_{AGCin}^{1/(1+(I_1/I_2))}. \quad (2)$$

The knee of the AGC circuit is set by a minimum circuit that limits the control current to the VGA to be the minimum of

that output by the circuit of Fig. 7 and a user-settable current that limits the maximum gain of the VGA. The minimum circuit implemented in our processor is relatively straightforward and is described in [26].

Fig. 8 shows the overall steady-state input-output (I/O) compression curve measured from the AGC circuit in our processor. The knee and compression are clearly visible at low and moderate input amplitudes in the dual log-scaled plot.

Fig. 9(A) and (B) shows the attack and release characteristics of the AGC as it adapts to an abrupt increase or abrupt decrease in the input amplitude respectively. The fast attack and slower release time constant in the envelope detector is inherited by the AGC, which consequently adapts more quickly to increasing amplitudes than to decreasing amplitudes. Such asymmetry in the attack and release time constants of the AGC is necessary to mimic the behavior of the human auditory system.

The power consumption of the overall AGC circuit was measured to be $30 \mu\text{W}$. At its maximum gain setting, the AGC is designed to ensure that its minimum detectable signal is near but below the output noise floor of the preceding AFE ($50\text{-}\mu\text{V}$ rms) and that this noise is amplified by the AGC to be near but below the input-referred noise floor of the succeeding channel circuits. ($310 \mu\text{V}$ rms). The gain of the AGC is then turned down to ensure that, at maximum amplitudes, it does not saturate the filters of the succeeding channel circuits by limiting the output signal strength to 500-mV rms. The power consumption of the AGC circuit is limited by the need to preserve the 10-kHz bandwidth necessary for this application while simultaneously maintaining the low noise floors required for wide-dynamic-range operation of the overall processor.

Programming of the AGC parameters is done with currents from a constant- g_m bias circuit described in Section V. The programming of the translinear circuit may also be implemented by off-chip current biasing.

IV. SINGLE-CHANNEL PROCESSING CIRCUITS

Each channel of the processor consists of a bandpass filter, an envelope detector, and a logarithmic A/D to extract the logarithm of the spectral energy in a given filter band of the spectrum. We shall describe each of these blocks briefly, focusing on key insights and results relevant to processor operation. Further details may be found in [5]–[7].

A. The Bandpass Filter

The bandpass filter is a fourth-order transconductance-capacitor filter and consists of two stages of a capacitive-attenuation second-order filter cascaded via a source-follower to prevent interstage loading as shown in Fig. 10. The upward and downward rolloff slopes of the resulting filter are then both 40 dB/decade. Second-order rolloff slopes are known to be necessary and sufficient for good patient performance. In each stage, C_1 and AC_1 implement capacitive-attenuation to widen linear range in the transconductors [5], G_1 and its attenuating capacitors implement the high-pass portion of the filter, G_2 and its output-loading capacitors implement the low-pass portion of the filter, and G_3 provides low-frequency adaptation to set the floating node of G_2 's capacitive divider. Measurements of

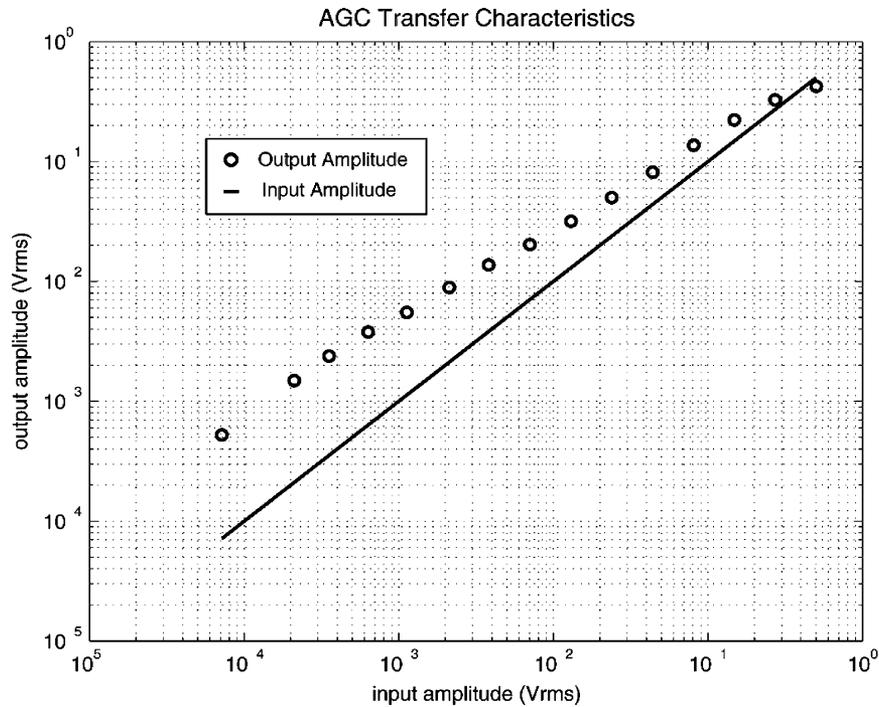


Fig. 8. The AGC input-output compression characteristics.

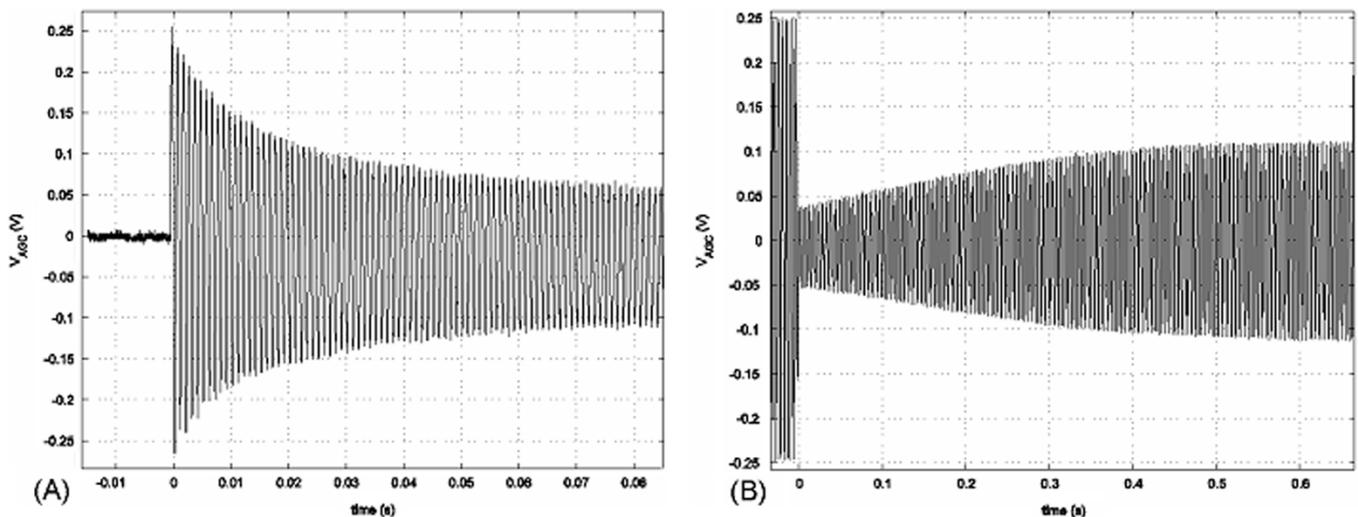


Fig. 9. Experimental AGC output waveforms showing: (A) the dynamics of attack with $T_a = 10$ ms; (B) the dynamics of release with $T_r = 300$ ms.

these filters reveal a 65-dB operational dynamic range from $400 \mu\text{V}$ rms at the input-referred noise floor to 750 mV rms at 5% total harmonic distortion. The power consumption is $5.4 \mu\text{W}$ for a 5- to 10-kHz filter and 112 nW for a 100- to 200-Hz filter. The filter's performance is in good accord with a theoretical noise-and-dynamic-range analysis, which is performed using the procedure outlined in [5]. Each channel of the processor has a 5-bit DAC whose output determines the low-frequency corner of the filter by setting the bias current of G_1 , and a 5-bit DAC whose output determines the high-frequency corner of the filter by setting the bias current of G_2 . The bias current of G_3 , which is less critical in determining the filter characteristics as long as it is small enough, is set to a constant current of approximately 10 pA. To provide good G_m/C corner-frequency

matching, the DAC's reference currents are generated from an on-chip resistor-based constant- g_m current reference described in Section V. Experimental measurements from the array of bandpass filters are presented in Section VII.

B. The Envelope Detector

The envelope detector circuit is shown in Fig. 11(A). It consists of a transconducting rectifier followed by a current-mode peak detector that estimates the dc component of the rectified signal. The details of the rectifier circuit are shown in Fig. 11(A) and (B), and the details of the peak-detector circuit are shown in Fig. 12. Fig. 11(A) shows that the rectifier is implemented by inserting a class AB current conveyor inside a G_m -C filter in a negative-feedback configuration. At frequencies above the

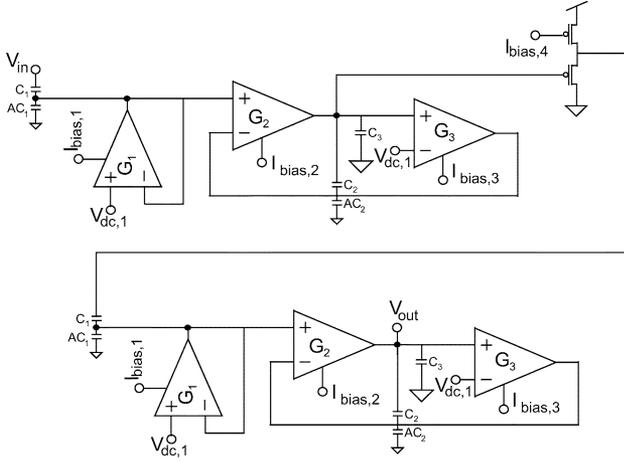


Fig. 10. The capacitive-attenuation bandpass filter. Two second-order filter stages are cascaded.

corner frequency of the filter, the output current of the transconductor, I_{in} in Fig. 11(A)], is proportional to the input voltage V_{in} since the node voltage on the capacitor is very near the dc voltage of the input, i.e., $I_{out} = -I_{in} = G_m V_{in}$. This current is split into two halves with positive current traveling through M_p and its associated mirror, and negative current traveling through M_n and its associated mirror. For small input signals, M_p and M_n are operated in weak inversion and function like MOS versions of an NPN and PNP transistor respectively with exponential current dependence. The positive and negative half currents are added to create a full-wave-rectified signal, and fed to the peak detector. However, Fig. 11 only shows a half-wave implementation for clarity. The local feedback amplifier, shown in symbolic form in Fig. 11(A), and as a transistor schematic in Fig. 11(B), serves to reduce the two diode-drop “dead zone” of the M_n and M_p exponential elements by prebiasing $V_{out,TOP}$ and $V_{out,BOT}$ to enable Class AB operation. It also reduces the effect of the node capacitance by the gain of the amplifier. The corner frequency of the G_m - C filter is set below 100 Hz since we are interested in sensing the envelopes of signals in the 100-Hz to 10-kHz audio frequency range.

The minimum detectable signal of the rectifier is limited by two considerations. First, the magnitude the current I_{in} must be large enough to charge the effective capacitance at the input node of the conveyor to a voltage that has sufficient amplitude to exceed the dead-zone voltage; this effect argues for having as small a dead zone as possible if one wants to have a tiny minimum detectable signal. However, there is a second effect: A small dead zone is detrimental to detecting small signals because it causes large transmission of the thermal noise currents at the input of the conveyor to the peak detector and creates a large rectified-thermal-noise current floor; a small dead zone effectively reduces the natural noise-reduction benefit of a thresholding nonlinearity by reducing the threshold. Thus, to get as small a minimum detectable signal as possible, there is an optimal dead-zone width at which both effects yield the same limiting noise floor. The analysis in [6] and [18] establishes the optimal width of the dead zone as a function of the bias current of the transconductor, the parasitic gate-to-source capacitances of the transistors, and the maximum desired frequency range of

operation (10 kHz). At this optimal width, one can compute that the maximum dynamic range $D_{optimum} = I_B/I_{in,ms}$ achievable in a transconductor with bias current I_B , with n effective devices of shot noise, with an output current mirror ratio of N , and a minimum detectable current of $I_{in,ms} = G_m V_{in,ms}$ is given by

$$D_{optimum} = \sqrt[4]{\frac{2}{\pi}} \cdot \sqrt{\frac{N I_B}{n \cdot q \cdot f_{MAX}}} \quad (3)$$

where q is the charge on the electron. In [6], we showed that we achieved this optimum dynamic range by presenting results from a 75 dB (0.3 mV pp–1.7 V pp), 2.8 μ W, 100-Hz–10-kHz envelope detector with input dc-insensitive operation as in Fig. 11(A). The use of minimum-size subthreshold MOS devices for M_n and M_p with low values of gate-to-source capacitance was an important factor in achieving this dynamic range.

The peak detector shown in Fig. 12 is a current-mode low-pass filter [19], [20] modified to have asymmetric attack and release time constants [6], [18]. If transistor M_5 is shorted, then the topology behaves like a lowpass filter with a time constant $\tau_a = [C_a(kT/q)]/(\kappa I_a)$. When the input current is rising, M_5 and I_r , which form a source follower, do effectively behave like a short since M_5 is strongly turned on in this regime. Thus, for rising inputs, the attack time constant is given by the above equation. For falling inputs, M_5 is strongly turned off and the source follower relies on I_r to provide a slow return to equilibrium by charging C_r . The net effect is that we have a release time constant $\tau_r = [C_r(kT/q)]/(\kappa I_r)$, where κ is the subthreshold exponential constant [17].

The cascade of a rectifier and a peak detector yields an envelope detector. The detector quickly follows the rising portion of the full-wave-rectified input with a fast attack and holds this value with little decay in between cycles during a slow release. Fig. 13(A) shows data from a low-power (0.87 μ W) version of the envelope detector with 63-dB dynamic range used in each of the 16 channels. Fig. 13(B) shows data from a higher-power wider-dynamic-range version (75 dB, 2.8 μ W) used in the AGC circuit of the processor. The outputs of both envelope detectors are relatively invariant with input frequency over the 100 Hz–10 kHz operating range [6].

A 2-bit DAC in each channel allows I_a to be altered such that the attack time constant can vary from 1 ms to 4 ms. A 3-bit DAC in each channel allows I_r to be varied such that the release time constant can vary from 3 ms to 24 ms. Fig. 14 shows how this programmability affects the envelope tracking of a square-wave amplitude modulated, sinusoidal input. The measurements were taken by converting the current to a voltage across an instrumentation transconductor.

C. The Logarithmic A/D

Fig. 15 shows a schematic of the logarithmic A/D with offset and temperature compensation: The WLR circuit is a wide-linear-range transconductor with gate inputs, bump-linearization, and source-degeneration, which extend its linear range to 0.24 V [16]. The comparator is a standard 5-transistor operational transconductance amplifier [17].

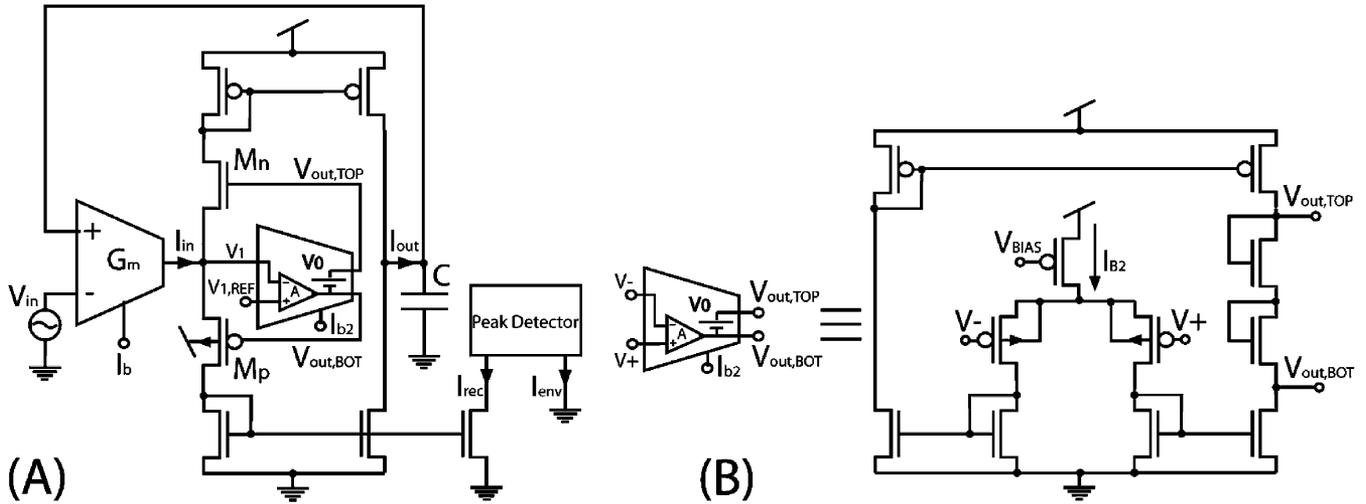


Fig. 11. (A) The envelope-detector topology. (B) The feedback amplifier implementation.

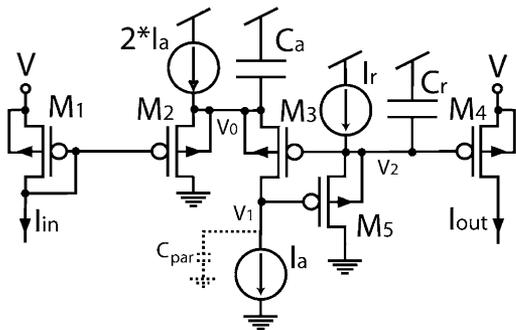


Fig. 12. The current-mode peak-detector topology with adjustable attack and release time constants.

The output current from the envelope detector is converted to a logarithmic voltage on a diode. The proportional-to-absolute temperature (PTAT) dependence of this diode voltage is cancelled by the inverse PTAT dependence of the transconductance of the WLR transconductor [16], which linearly reconverts the diode voltage to a current. The ratio of this current to the bias current of the transconductor is then digitized in a dual-slope analog-to-digital (A/D) conversion scheme described below. Any temperature dependence of the bias current of the transconductor is cancelled out in the ratiometric computation. The switches and clocking controls in Fig. 15 implement offset compensation via the automatic incorporation of auto zeroing in the conversion cycle.

The clocking sequence employed for dual-slope A/D conversion involves 3 phases: During the first auto-zeroing phase (AZ), the input to the WLR is switched to a reference current I_{ref} , the WLR's offset is sampled and stored on C_{az} and C_{int} via negative-feedback action on the follower-configured WLR circuit, and the comparator's offset is similarly sampled and stored on $C_{coupling}$. During the AZ phase, the offset current I_{os} serves to auto-zero the WLR to operate from a starting negative-offset voltage. This strategy ensures that a substantial fraction of the full differential linear range of the transconductor is available for use in the circuit rather than just the positive half of the differential linear range if I_{os} is zero. On the second

or following integration phase, the AZ and AZ_{delayed} feedback switches to the WLR and comparator are opened, and the input of the WLR is switched to the envelope-detector current I_{in} . The delayed opening of the AZ_{delayed} switch prevents charge injection from the AZ switch from affecting the comparator output. During the integration phase, a current proportional to the logarithmic value of I_{in} is integrated for a fixed period of time to generate a voltage on the integration capacitor C_{int} . On the third or following de-integration phase, the input to the WLR is switched to I_{ref} , the current I_{os} is disconnected from C_{int} , and the negative-offset voltage on the WLR causes it to de-integrate the voltage developed on C_{int} during integration. If a counter, initialized at 0, counts the number of clock cycles until the voltage on C_{int} de-integrates to the auto-zeroed trip point of the comparator, then, the frozen value of the counter after the comparator trips represents a digitization of $(\phi_t/V_L) \ln(I_{in}/I_{ref})(I_b/I_{os})$, where $\phi_t = kT/q$ is the PTAT thermal voltage and V_L is the inverse-PTAT linear range of the transconductor. To ensure wide dynamic range, V_L needs to be at least as large as the 180 mV of input diode-voltage variation seen over a 60-dB input-current variation, but it should not be too much larger, or the sampled thermal noise on C_{az} will reduce the precision of the converter [7]. A small value of V_L also reduces the power required to maintain the bandwidth necessary for settling during the auto-zeroing phase.

The use of two local-feedback loops to auto zero the WLR and comparator separately instead of one global-feedback loop, yields a 2.5-bit improvement in precision for the same power [7]. The improvement in precision results because a larger value of C_{az} and C_{int} can be used to reduce sampled thermal noise, while the power used in the more complex biasing-and-compensation circuitry of global feedback, can be diverted into maintaining the bandwidth of the WLR.

Since the envelope detectors across various channels are likely to have gain mismatches and differing noise floors, it is advantageous to adjust I_{ref} in each channel to compensate by providing calibration bits for each channel. Each channel, therefore, has a 3-bit DAC to alter the value of I_{ref} to assume one of 8 possible values from approximately 50 pA to 400 pA.

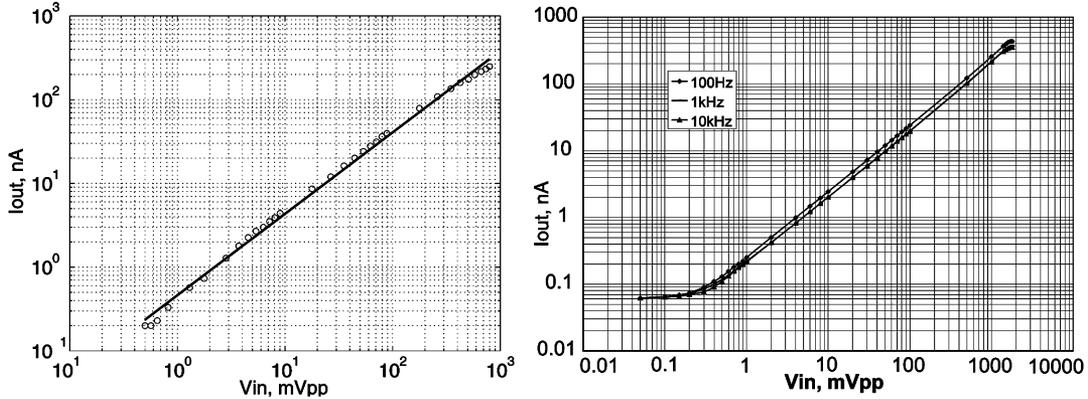


Fig. 13. Experimentally measured envelope detector characteristics: (A) for the 63-dB dynamic range used in the channels, $f = 1$ kHz; (B) for the 75-dB dynamic range used in the AGC, $f = 100$ Hz, 1 kHz, 10, kHz.

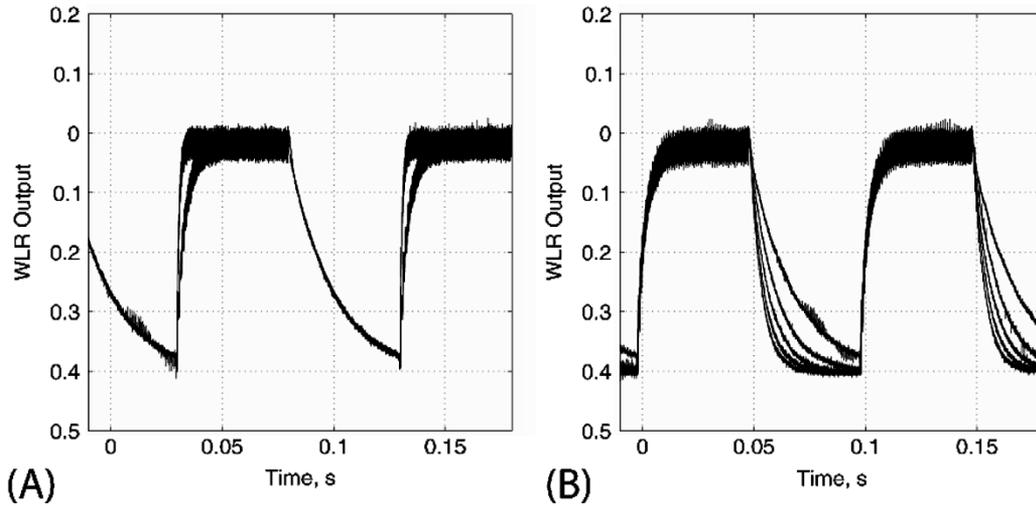


Fig. 14. Envelope detector: (A) attack-time programmability ($T_a = 1$ ms and $T_a = 4$ ms); (B) release-time programmability ($T_r = 4$ ms, $T_r = 6$ ms, $T_r = 8$ ms, $T_r = 12$ ms, $T_r = 24$ ms).

Sampling rate and power are easily traded in this converter by adjusting the value of I_b and I_{os} . A $2\times$ increase in sampling rate at the same precision requires a $2\times$ increase in I_b and I_{os} and also doubles the digital switching power consumption. At constant power, speed and precision may also be traded off against each other, albeit in a more speed-advantageous fashion: Since the converter's precision is thermal-noise limited [7], a reduction in capacitance by $4\times$ worsens the precision by 1 bit. However, the converter can now run at a sampling rate that is four times faster.

The converter described in [7] achieved 8-bit precision at 300 Hz sampling rate with $3 \mu\text{W}$ of power ($1 \mu\text{W}$ of analog power and $2 \mu\text{W}$ of switching digital power in a $1.5\text{-}\mu\text{m}$ process) for appropriate choices of capacitances and bias currents. While these numbers are perfectly suitable for generating spectrograms of speech signals for a speech-recognition front end (the envelope of speech varies slowly at a 100-Hz rate and there is much variability in the signal), they are more precise than required for cochlear-implant processors. Patients can discriminate 8–50 electrode-stimulation steps [21], which implies that a 5-bit quantizer may be perfectly adequate. It appears that a higher stimulation rate may be more beneficial than precise stimulation. Thus, in this processor, while keeping

the power constant at $3 \mu\text{W}/\text{channel}$, we reduced the values of capacitances in the circuit to lower the precision to 7-bits and increased the sampling rate to 1 kHz. As we discuss in Section VI, our processor allows the user of the chip to globally increase the sampling rate at the cost of increased power by altering DAC bits or by inputting external bias currents for I_b and I_{os} .

V. CURRENT-AND-VOLTAGE BIASING

Fig. 16 shows our PTAT current-reference circuit that is designed to generate a current of $(kT/q) \ln(9)/R$ where R is near $1 \text{ M}\Omega$, and the transistors are in subthreshold MOS operation. A description of how standard PTAT circuits work may be found in [22].

Our circuit uses cascode mirrors to reduce the drain-voltage dependencies of the currents, and capacitive bypassing to attenuate power-supply effects on the reference output. The capacitive bypassing keeps gate-to-source voltages of all transistors nearly invariant with high-frequency noise on the supply, and is based on some of our prior work on designing high-power supply rejection ratio (PSRR) amplifiers [23]. Buffered versions of the four gate voltages labeled V_n, V_{nc}, V_p, V_{pc} are used as

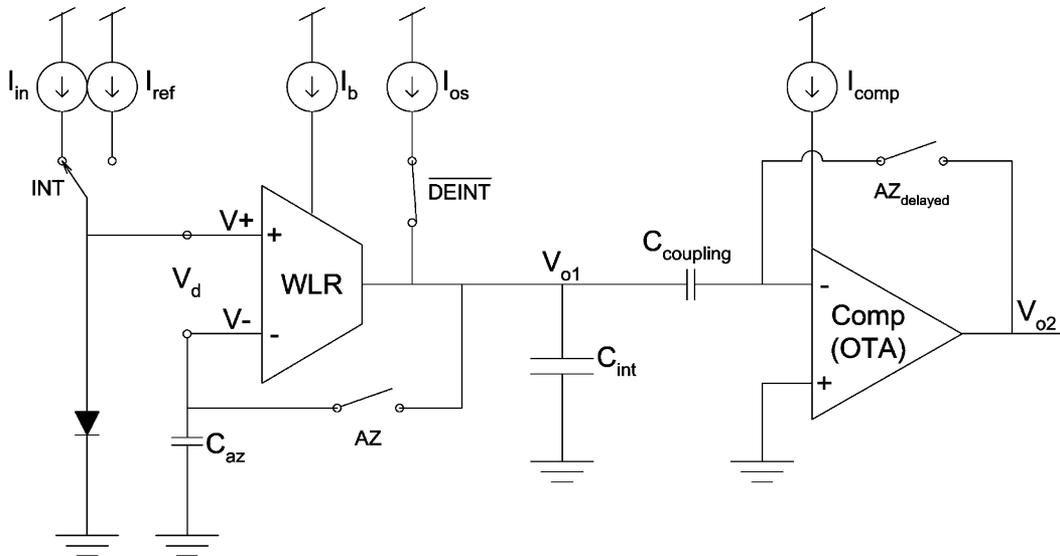


Fig. 15. The Logarithmic A/D topology.

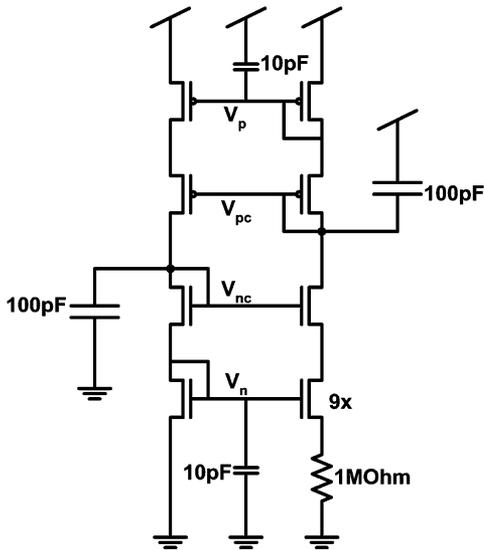


Fig. 16. The on-chip reference current and voltage bias circuit.

reference voltages in various parts of the processor, primarily for setting cascode voltages. With a 2.8-V supply, in our circuit, they are nearly at 0.45, 1.05, 2.0, and 0.95 V, respectively.

The resulting PTAT reference current of 45 nA is used in a current-distribution network to distribute scaled-up and scaled-down current copies to bias the various DAC reference currents, amplifier bias currents, and comparator bias currents all over our chip. The PTAT property results in constant- g_m behavior with temperature such that the corner frequencies of the bandpass filters and the peak-detector time constants are invariant with temperature. The small PTAT variation with temperature has little effect on the operation of any other circuits in the processor, where it is either cancelled out (e.g., in the logarithmic A/D), or causes a small percentage change in the bandwidth, noise, or power that does not affect operation significantly (e.g., in the envelope detector). This processor is expected to function in an environment inside the body of a patient, where the temperature

is typically well regulated, making such effects more inconsequential. Nevertheless, just as in bipolar biasing circuits, it is extremely important in subthreshold MOS operation to pay careful attention to temperature effects because of the exponential dependence of the current on temperature. Our processor simply would not operate robustly if the gate voltages of biasing transistors were set with potentiometer voltages, a common practice during quick characterization of a chip, but unsuitable for medical use. Furthermore, for portable speech-recognition front ends, where our processor could also be used, robust operation with temperature variations is important.

For simplicity, an early version of our processor distributed voltages created by reference currents on a diode to create current copies elsewhere on the chip. Since transistor threshold voltages can vary substantially across the chip, and subthreshold circuits are exponentially sensitive to these variations, we obtained relatively poor matching and consequently a huge power overhead of almost two. The power overhead arises because poor matching causes the bias-voltage to have to be set at a worstcase value to get all circuits to work well. In contrast, by always distributing current and doing local mirroring, the matching is greatly improved. However, care must be taken to ensure that there are not too many stages of mirroring between the initial reference current and the final current copy that gets used, and that there is little power overhead because large currents are always used at the local site of generation and not distributed. The power overhead due to our biasing scheme was small and measured to be 3 μ W. Experimental data on matching and immunity to power-supply noise are presented in Section VII where we discuss the overall operation of the processor.

VI. DIGITAL CONTROL, PROGRAMMABILITY, AND THE I/O INTERFACE

The processor interacts with the outside digital world by outputting bits that represent the result of its computation, and by receiving digital bits that allow its various DAC parameters to

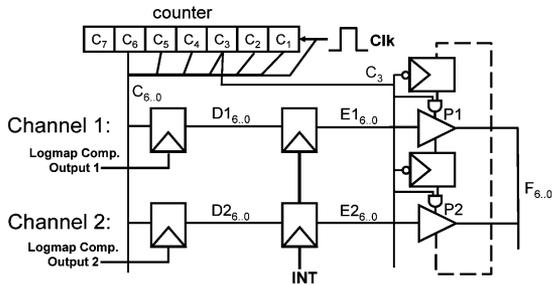


Fig. 17. The digital I/O interface for Continuous Interleaved Sampling (CIS) output stimulation.

be programmed. In addition, the processor allows a select set of analog waveforms from each channel to be visible. These analog waveforms could be potentially useful for debugging and testing subjects in a highly variable patient population. We shall discuss the output-bit interactions in Section VI-A and discuss the programmability-and-visibility interactions in Section VI-B.

A. The CIS Output Interface

Fig. 17 reveals how a 7-bit digital number from each of the 16 channels is created, synchronously latched, and serially multiplexed onto a common output bus for use in electrode stimulation. The final output of the chip is the digital data on the F bus, which contains 7-bit numbers from each of the 16 channels, one at a time, in cyclic-and-sequential fashion. If 16 electrodes are correspondingly synchronously stimulated with current pulses whose charge magnitude is proportional to the 7-bit number, then the method of stimulation is referred to as continuous interleaved sampling (CIS) stimulation. This method of stimulation is the most common method of stimulation today because it minimizes current spreading and interaction amongst the electrodes. The figure only shows the digital circuitry in two adjacent channels, namely channel 1 and channel 2, but all channels have identical circuitry.

An 8-bit global counter cyclically counts from 0 to 255 over each sampling period. In Fig. 17, the clock is the C_0 bit, while the counter stages go from C_1 – C_7 . An 8-bit counter is needed to cover the 128 cycles of de-integration (to produce a 7-bit output) plus 1 additional bit for the overhead of auto-zeroing and integration.

The dual-slope converters in each channel are configured such that de-integration is performed when the counter has values in the set $\{0$ to $127\}$, auto-zeroing for the next conversion is performed from the end of de-integration to counter value 191, and integration for the next conversion is performed when the counter has values in the set $\{192$ to $255\}$. There is only one counter for the whole chip and all converters are performing conversion in parallel in accordance with this scheme. Depending on the value of their input current magnitude sampled during integration, somewhere during the de-integration phase, each of the comparator outputs of the logarithmic converters will asynchronously go high. The exact value of the counter at the time a particular comparator output goes high is a digital measure of the log spectral energy in that comparator's channel. Thus, at the time a comparator goes high, the value of the counter at this instant is latched into the first

set of registers to create the 7-bit D numbers for each channel. Note that since the counter only goes from 0 to 127 during de-integration, we need only latch the C_0 through C_6 bits. The D numbers become valid at some asynchronous time which is different for each channel. To synchronize these numbers, they are all latched into a further set of registers to create the E numbers at the start of the integration phase (when the counter hits 192, or equivalently when C_6 and C_7 go high). The 16 E numbers are now multiplexed onto the F bus by clocking a shift-register-activated bank of tri-state P buffers with bit C_3 of the counter. The shift register is configured to shift an active token bit through each channel in a cyclical fashion on the falling edge of C_3 . To summarize and clarify the description, we note that a 1-kHz sampling rate/channel on C_7 would cause CIS data to be multiplexed onto the output F bus at 16 kHz, while the input clock would operate at 128 kHz.

To prevent the tri-states from fighting each other and being simultaneously on during transitions, a hazardous and power-wasting situation, the tri-state control signal from the shift register is ANDed with C_3 : Whenever C_3 goes low, all tristates are inactivated. A half clock cycle later, when C_3 goes high again, only one tri-state with an active shift-register bit is activated. Thus, there is a half clock cycle gap between the activation of successive tri-states. The output data of the F bus is guaranteed to be valid after the falling edge of C_3 , hence, C_3 is outputted as a "CIS clock" for latching the data in subsequent electrode stimulation circuits.

In Figs. 15 and 17, the scheme described above is implemented by using C_6 AND C_7 to generate INT. All other timing signals [such as NOT(DEINT) and AZ] are generated asynchronously by a state-machine in each logarithmic A/D. Our choice of using a dual-slope conversion scheme is convenient because the sharing of the counter amongst all channels saves digital power, and because the counter conveniently provides various timing control signals.

In addition to outputting digital bits, the de-integration pulses from each channel are also directly reported off chip such that external A/D conversion may be performed on these pulses. Alternatively, external pulse-processing circuitry can directly transform these pulses to stimulate electrodes.

B. Programmability and Visibility

A programmability-and-visibility clock sequentially shifts an active token in a shift-register-based scanner on the chip [24]. When this token activates a particular channel, that channel can be programmed and its visible outputs can be examined. The active token can also activate no channel, the default mode when the processor is in use and not being programmed. If a programming control line is enabled, any of the 18 programmable bits of this channel can be altered and stored in on-chip latches. The 18 bits consist of 5 bits for setting a DAC that controls the low-frequency bandpass filter corner, 5 bits for setting a DAC that controls the high-frequency bandpass filter corner, 3 bits for setting the DAC release current in the envelope detector, 2 bits for setting the DAC attack current in the envelope detector, and 3 bits for setting the DAC reference current in the logarithmic A/D. In addition, 3 global bits can be configured to allow the sampling

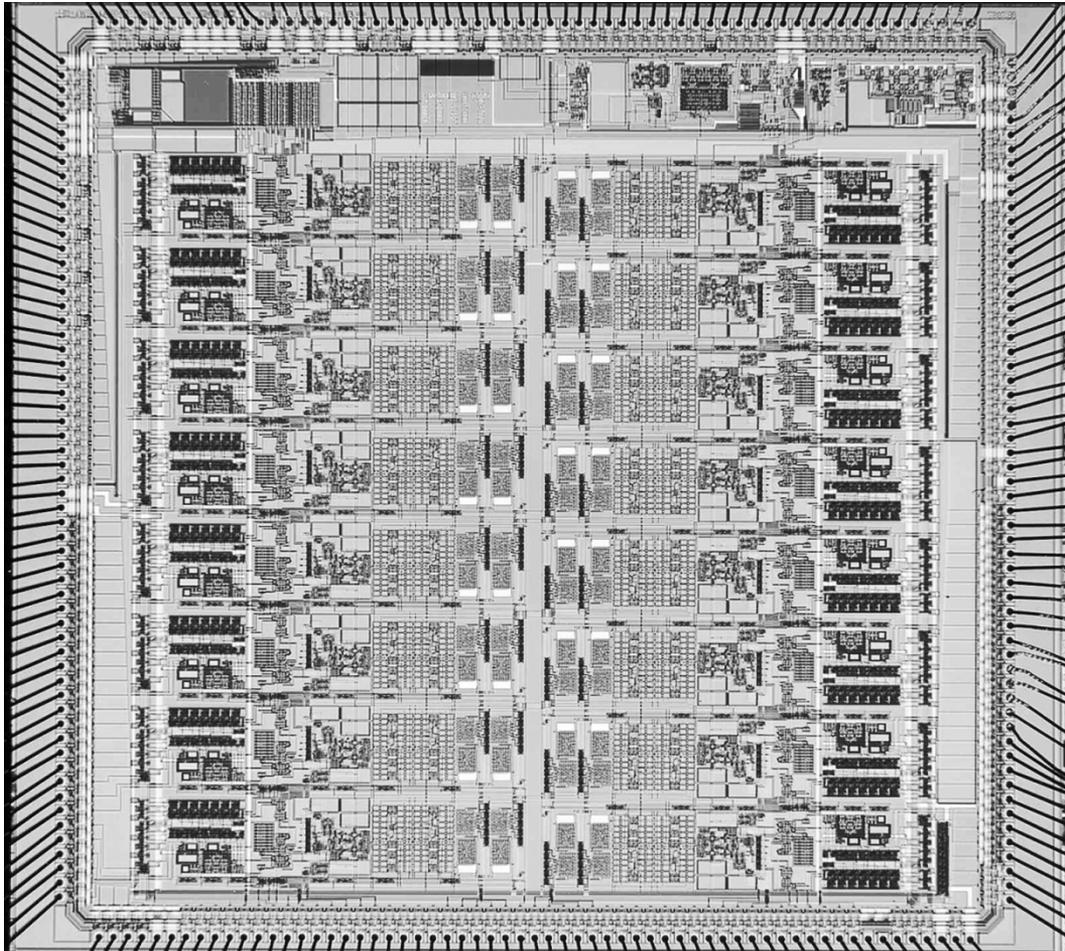


Fig. 18. A die photo of the bionic-ear processor chip.

rate/channel to be 500 Hz, 1 kHz, 1.5 kHz, 2 kHz, or to an externally settable value with a user-supplied bias current. Thus, there are $16 \times 18 + 3 = 291$ programmable bits in the system.

When a channel is visible, the input of the first highpass stage in the bandpass filter, the output of the bandpass filter, the rectifier output current, the peak-detector output current, the diode input to the logarithmic A/D, the dual-slope waveform of the logarithmic A/D, and the comparator output of the logarithmic A/D may be observed off chip. Internal current waveforms are transduced to voltages with an instrumentation transconductor whose bias current is electronically adjustable.

VII. EXPERIMENTAL PERFORMANCE OF THE PROCESSOR

Fig. 18 shows a die photo of the processor. The chip was fabricated in a $1.5 \mu\text{m}$ AMI BiCMOS process available through MOSIS, the prototyping chip-fabrication service. The $9.58 \text{ mm} \times 9.28 \text{ mm}$ chip has 11 640 NMOS transistors, 11 537 PMOS transistors, 287 capacitors, and a handful of bipolar transistors.

Fig. 19 shows the output frequency response of all channels when all filters are configured with the same DAC bits. The data reveal that there is good time-constant matching across all channels since all bandpass filters peak at the same frequency. In addition, the second-order roll-off slopes of a bandpass filter ought to be linear on a log-log plot. Since the output bits of the chip

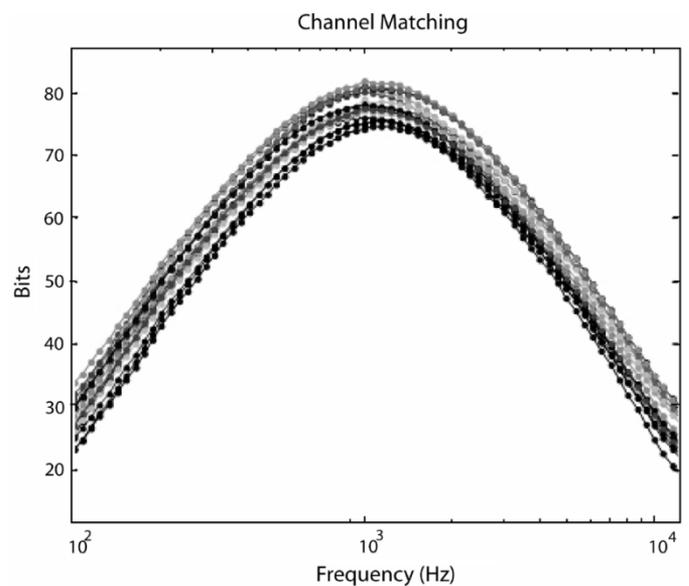


Fig. 19. Matching time constant data across all 16 channels of the chip at the same bandpass filter DAC bits. The bit output of the chip on the y axis corresponds to 1.5 bits/dB of input amplitude.

represent the logarithm of the spectral energy in each channel, they behave like they are on a logarithmic scale, and exhibit a nice linear change with frequency. This data reveals that each

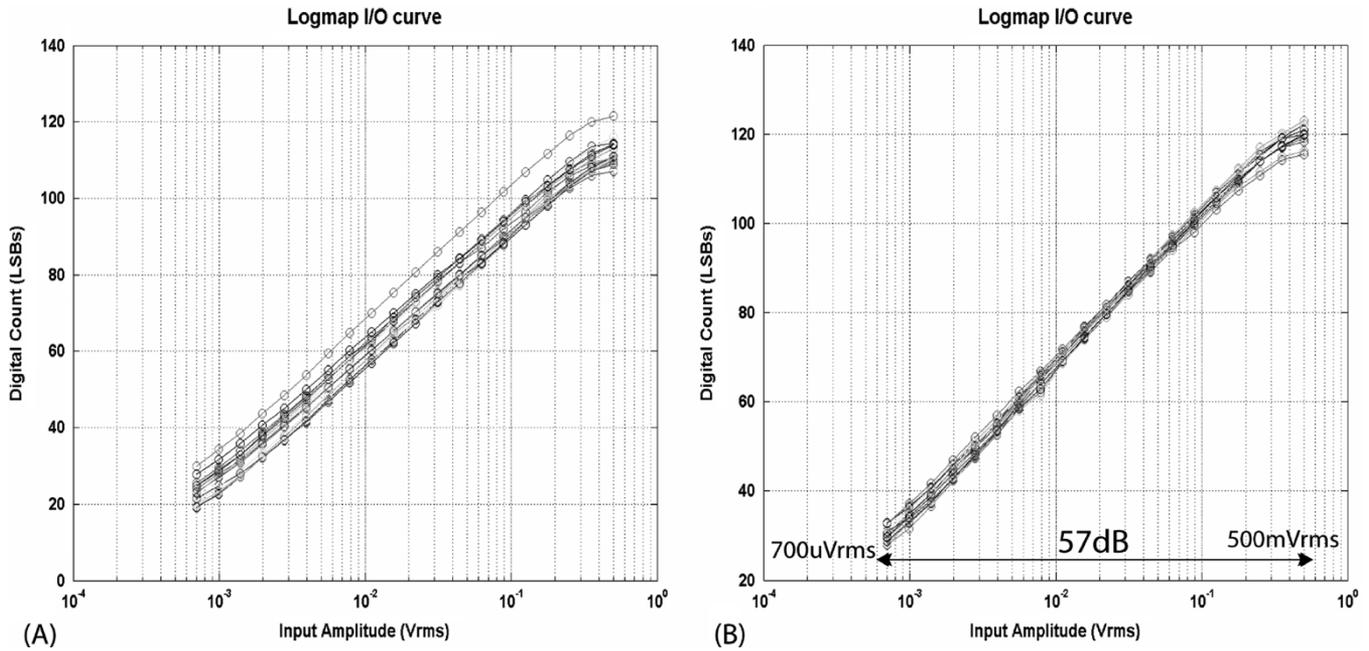


Fig. 20. Logarithmic A/D offsets: (A) before calibration; (B) after calibration. All 16 channels have an IDR of 57 dB.

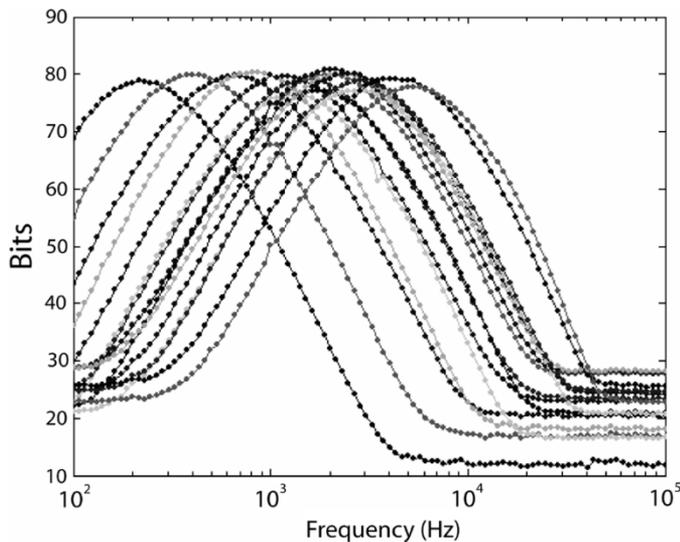


Fig. 21. The bits-versus-frequency curve for all 16 channels. The filters span the frequency range from 200 Hz to 6 kHz. The bit output of the chip on the y axis corresponds to 1.5 bits/dB of input amplitude.

channel is behaving as expected except that the gains of each channel are not the same, resulting in offset shifts in logarithmic bit space. Fig. 20(A) and (B) shows that the offset calibration bits of the logarithmic A/D may be configured to attenuate these offsets: Fig. 20 shows that the logarithmic response of the bits to varying amplitudes of a sine tone at the peak frequency of the channel have much less offset after calibration [see Fig. 20(B)] compared with before calibration [see Fig. 20(A)]. We also see that there is a 57-dB IDR of operation from 700- μ V to 500-mV rms. Fig. 21 shows that the bandpass filters in the array may be configured to span the entire audio frequency range from 200 Hz to 6 kHz with approximately 50 dB of stop band rejection (1.5 decades with 40 dB/decade).

Fig. 22 shows that a Matlab simulation of our chip generates an output similar to that from the chip: The input to the chip is a sine tone with a frequency that ramps up and down as revealed by its fast Fourier transform (FFT) in the lowest plot. The Matlab plot in the center and the chip output at the very top generate broad responses across the channels due to the relatively broad filters in both implementations.

Fig. 23 shows photographs of real-time spectrograms generated by plotting the output bits from the 16 channels as a function of time with dark values representing high bit numbers and lighter values representing low bit numbers. By talking into the FG3329 microphone and watching the computer, a running spectrogram of one's voice can be constructed as one speaks. The spectrograms were taken while the speaker was saying the word "bit." The word is said twice in a row with a gap. On the spectrogram shown on the left, there is no RF noise in the supply. On the spectrogram shown on the right, we intentionally capacitively couple a 49-MHz carrier modulated at 1 kHz onto the power supply of the chip. The amplitude of the carrier as measured directly on the power supply was 430 mVpp and the amplitude modulation depth was 100%. We see that there is barely any change in the two spectrograms even with large interference. Such robust operation is critical to the performance of the chip in real implant environments and confirms that the attention paid to power-supply-immune design was worthwhile. In addition, tests of RF interference performed with a large RF coil surrounding the test board reveal that very high levels of radiated RF (1 Vrms on a 10-turn coil of dimensions 30 cm \times 30 cm encircling the chip, and inducing a carrier amplitude of 400 mVpp on the power supply) do not interfere with the operation of the chip.

The power of various portions of the processor was measured by measuring the current consumption of that portion with a Keithley electrometer (separable supply rails for each portion

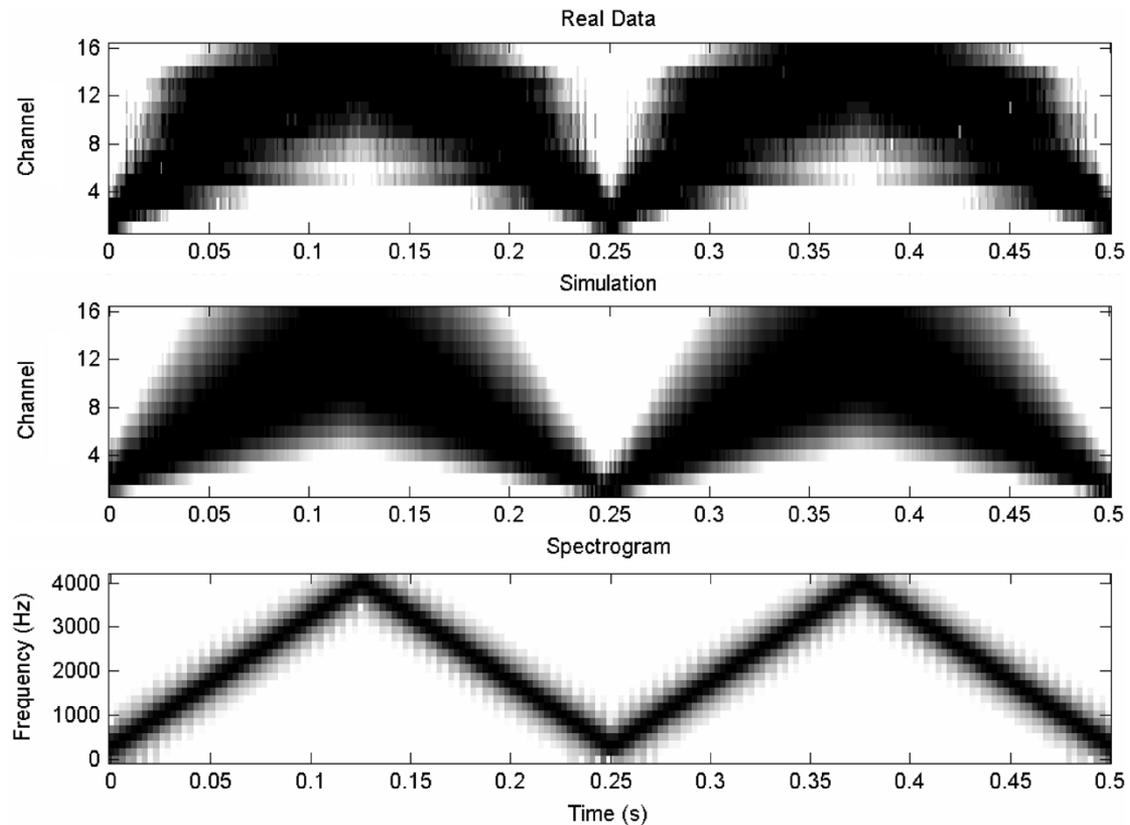


Fig. 22. Experimental chrip spectrogram (top figure), Matlab-simulation chrip spectrogram (middle figure), and FFT chrip spectrogram (bottom figure) to compare the system against.

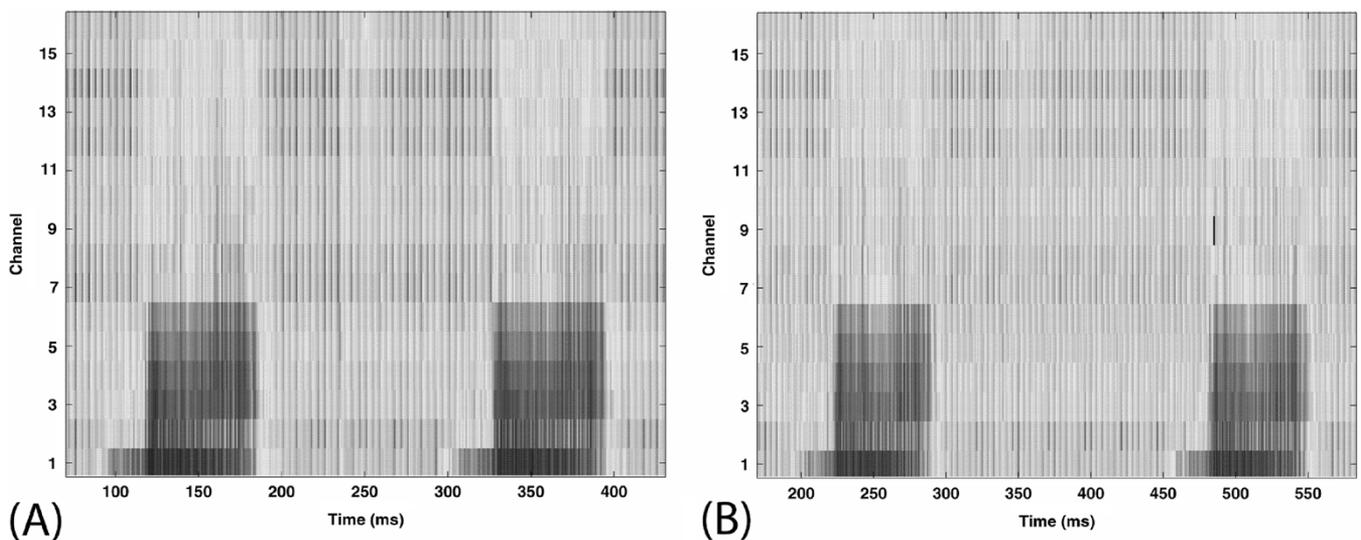


Fig. 23. Spectrogram of the word *bit/put* through the actual system: (A) The spectrogram in the absence of the RF noise; (B) The spectrogram in the presence of a 49-MHz parasitic sinusoid modulated with an in-band 1 kHz signal with 100% AM modulation. The RF noise is capacitively coupled to the power supply causing a signal as large as 430 mVpp on its bus. Nevertheless, we see that the spectrogram barely changes. Note that the time axes in the two plots have different origins.

made this measurement possible) and then multiplying by 2.8 V, the power supply voltage used for the processor. These power measurements are listed in Table I, with a total power consumption of $211 \mu\text{W}$. The numbers were measured at a 1-kHz sampling rate/channel with the lowest-frequency filter configured to have corners at 200 and 400 Hz, and the highest-frequency filter configured to have corners at 5 and 10 kHz.

VIII. ANALOG VERSUS DIGITAL

We estimate that a traditional A/D-then-DSP implementation of our processing would use about 0.25 to 0.5 mW for the microphone front end and A/D, and $250 \mu\text{W}/\text{MIP} \times 20 \text{ MIPS} = 5 \text{ mW}$ for the other processing, yielding a total power consumption of about 5.5 mW. These numbers are representative of

TABLE I
THE POWER OF VARIOUS PORTIONS OF THE PROCESSOR

| Portion of System | Power Consumption |
|--|-------------------|
| Microphone Front End and Microphone | 100 μ W |
| Automatic Gain Control Circuit | 30 μ W |
| Bandpass Filters | 19 μ W |
| Envelope Detectors | 14 μ W |
| Logarithmic A/D (analog section) | 13 μ W |
| Digital Control and CIS Output Interface | 32 μ W |
| Bias Circuits | 3 μ W |
| Entire Chip Processor | 211 μ W |

state-of-the-art cochlear-implant processing. Thus, our 211 μ W implementation is at least 25 times more power efficient even though it is in a 1.5 μ m process, not in an advanced submicron process. If we subtract the 100 μ W power consumption of our microphone and microphone front end, we are doing 20 MIPS of processing in 100 μ W of power, i.e., we are operating at around 5 μ W/MIP.

Needless to say, the digital μ W/MIP number will constantly improve with Moore's law. If we generously assume that it's actually 0 at the end of Moore's law, the power consumption of a very low power microphone front end, anti-alias filter, and A/D would likely still exceed 211 μ W. It must be remembered that A/D scaling in speed, power, and precision is much slower than Moore's law and that some circuits in our analog implementation could also benefit from these improvements. For example, the envelope detector and logarithmic A/D would both have reduced power consumption in a more advanced technology with reduced parasitic capacitance. Reducing the IDR in each channel and doing more gain control would cut analog power consumption further. For example a 40-dB IDR system could operate with 140 μ W–150 μ W of power. The use of very narrowband filters or higher-order filters could increase our power consumption to 300 μ W. However, we would still be an order of magnitude more efficient than an A/D-then-DSP implementation. It should be noted that very narrowband filters and higher-order filters have poor timing resolution and usually do not improve patient performance or speech-recognition performance. A custom digital solution would certainly reduce our power advantage but the high cost of the A/D and microphone would still give us a significant advantage.

It is useful to understand why we are operating more efficiently than an A/D-then-DSP implementation: An A/D immediately creates a representation of the incoming information as a series of relatively high-precision and high-speed numbers (16

bits at 44 kHz is typical in this application) that by themselves carry very little meaningful information. This digitization costs a lot of power because it is expensive to do any task at high speed and high precision. The high precision is necessary if we want wide-dynamic-range operation and are doing all computations including gain control in the digital domain, and the high-speed is necessary to avoid aliasing. Then, a DSP takes all of these numbers and crunches them with millions of multiply-accumulate operations/second, burning power in several switching transistors. It finally extracts more meaningful log-spectral-energy information at a much slower 100 Hz–1 kHz rate in 16 parallel bands, and at 8-bit precision, due to the high variability in speech data. In contrast, analog preprocessing allows for efficient compression of the incoming data such that low-speed and low-precision A/Ds at a later stage of the computation quantize the meaningful information. Some of our prior work analyzes the optimal point for digitizing information in more general systems [15]: Too much analog preprocessing before digitization is inefficient because the costs required to maintain precision begin to rise steeply, while too little analog preprocessing before digitization is inefficient because analog degrees of freedom that can be exploited to improve computational efficiency are ignored in the digital system.

Analog systems are more efficient than digital systems at low output precision while digital systems are more efficient than analog systems at high output precision [15]. In this processor, the output precision in each channel is 7 bits, and the output bandwidth is a few kilohertz at most. The IDR is near 60 dB with gain control allowing 77 dB of input dynamic range. An analog solution can, therefore, compete with a digital solution if we are careful in maintaining the necessary precision throughout the system. If the task required 14 bits of output precision, 72-dB IDR, and 100-kHz bandwidth at each channel, the A/D-then-DSP strategy would definitely be more efficient than our solution. An analog solution has to preserve its efficiency advantage by paying a lot of attention to robustness. The robustness is not attained in every device and in every signal as in a digital solution but at important locations in the signal-flow chain, where it matters. In other words, the robustness-efficiency tradeoff is addressed very differently in an analog system than in a digital system. The programmability is certainly not as great in an analog system as it can be in a digital one. However, as in our case, this is less of an issue when an algorithm that is known to work needs to be implemented and excessive programmability is attained for a huge loss in efficiency.

Signal processing algorithms based on zero-crossings save power because they avoid unnecessary switching. They have successfully been implemented in low-power subthreshold circuits [25]. Zero-crossing-based processing is advantageous for certain auditory recognition tasks especially those with periodic and musical signals. The system in [25] relies on matching between the threshold level and the dc component of the input signal that can be hard to control, thereby significantly reducing the dynamic range to an offset-limited regime. In [26], we show how to enhance our speech processor to be able to extract zero crossing information in an offset-free manner in our wide dynamic range envelope detector with essentially no overhead.

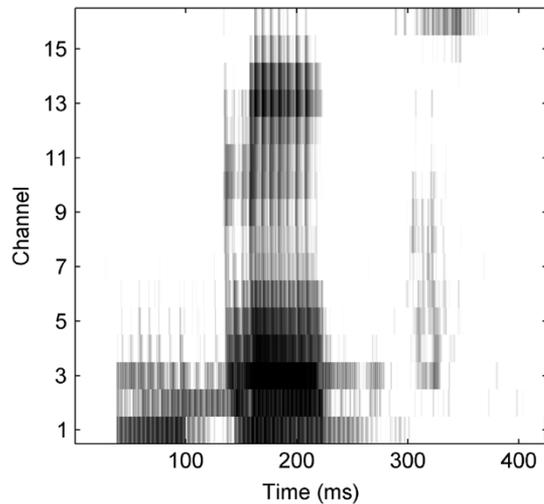


Fig. 24. Spectrogram-like plots for the word /bit/ in a version of the processor with sharper bandpass filters. Amplitude Outputs represented by the bits of the chip are used to construct a spectrogram.

IX. FURTHER IMPROVEMENTS

We have fabricated an improved 16-channel, 75-dB, 251- μ W version of the analog speech processor with more programmable bits, and described it in a brief conference publication [26]. It provides bandpass filters with programmable quality factors up to 10. The 251 μ W power consumption is for $Q = 4$, a typical value at which the processor is operated. This processor also has zero-crossing outputs for stimulation strategies that use phase information to enhance performance. Fig. 24 shows that spectrograms from this processor have sharper frequency resolution because of the use of sharper filters.

X. CONCLUSION

We described the operation of a 16-channel, 211 μ W, 77-dB dynamic range programmable analog BE processor that cuts the power consumption of state-of-the-art A/D-then-DSP approaches by a factor of 25. The processor is effectively operating at 5 μ W/MIP. Even if significant changes are made to its architecture, the lower-power advantage of our processor is likely to stand at the end of Moore's law because its total power consumption is less than that of a very low power A/D and microphone front end alone. The processor could run for 30 years on a 100 mAh battery capable of 1000 recharges. It is, thus, suited for use in fully implanted BEs of the future and in portable speech-recognition systems, where it could function as an ultra-low-power spectrum analyzer. This design suggests that the current trend of digitizing analog information at high-speed and high precision as soon as possible followed by processing in the digital domain is not an efficient solution if power consumption is of paramount importance. Rather, it is more advantageous to do robust-and-programmable analog preprocessing and digitize higher-level information at lower speed and lower precision.

REFERENCES

- [1] P. Loizou, "Mimicking the human ear," *IEEE Signal Process. Mag.*, vol. 15, no. 5, pp. 101–130, Sep. 1998.
- [2] J. W. Picone, "Signal modeling techniques in speech recognition," in *Proc. IEEE*, vol. 81, Sep. 1993, pp. 1215–1247.
- [3] F.-G. Zeng, G. Grant, J. Niparko, J. Galvin, R. Shannon, J. Opie, and P. Segel, "Speech dynamic range and its effect on cochlear implant performance," *J. Acoust. Soc. Am.*, pt. 1, vol. 111, no. 1, pp. 377–386, Jan. 2002.
- [4] M. Baker and R. Sarpeshkar, "A low-power high-PSRR current-mode microphone preamplifier," *IEEE J. Solid-State Circuits*, vol. 38, no. 10, pp. 1671–1678, Oct. 2003.
- [5] C. Salthouse and R. Sarpeshkar, "A practical micropower programmable bandpass filter for use in bionic ears," *IEEE J. Solid-State Circuits*, vol. 38, no. 1, pp. 63–70, Jan. 2003.
- [6] 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.
- [7] J. J. Sit and R. Sarpeshkar, "A micropower logarithmic A/D with offset and temperature compensation," *IEEE J. Solid-State Circuits*, vol. 39, no. 2, pp. 308–319, Feb. 2004.
- [8] M. Baker, T. Lu, C. Salthouse, J. J. Sit, S. Zhak, and R. Sarpeshkar, "A 16-channel analog VLSI processor for bionic ears and speech-recognition front ends," presented at the Custom Integrated Circuits Conf., San Jose, CA, Jul 2003.
- [9] R. J. W. Wang, R. Sarpeshkar, M. Jabri, and C. Mead, "A low-power analog front-end module for cochlear implants," presented at the XVI World Congr. Otorhinolaryngology, Sydney, NSW, Mar. 1997.
- [10] R. Sarpeshkar, R. F. Lyon, and C. A. Mead, "A low-power wide-dynamic-range analog VLSI cochlea," *Analog Integrated Circuits Signal Process.*, vol. 16, pp. 245–274, 1998.
- [11] W. Germanovix and C. Tourmazou, "Design of a micropower current-mode log-domain analog cochlear implant," *IEEE Trans. Circuits Syst. II: Analog Digit. Signal Process.*, vol. 47, pp. 1023–1046, Oct. 2000.
- [12] T. S. Lande, J. T. Marienborg, and Y. Berg, "Neuromorphic cochlea implants," in *Proc. IEEE Int. Symp. Circuits and Systems (ISCAS 2000)*, vol. 4, 2000, pp. 401–404.
- [13] R. F. Lyon and C. Mead, "An analog electronic cochlea," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 36, pp. 1119–1134, Jul. 1988.
- [14] W. Liu, A. G. Andreou, and M. H. Goldstein, "Voiced-speech representation by an analog silicon model of the auditory periphery," *IEEE Trans. Neural Netw.*, vol. 3, no. 3, pp. 477–487, May 1992.
- [15] R. Sarpeshkar, "Analog versus digital: Extrapolating from electronics to neurobiology," *Neural Computation*, vol. 10, pp. 1601–1638, 1998.
- [16] R. Sarpeshkar, R. F. Lyon, and C. A. Mead, "A low-power wide-linear-range transconductance amplifier," *Analog Integrated Circuits Signal Process.*, vol. 13, no. 1/2, May/June 1997.
- [17] C. Mead, *Analog VLSI and Neural Systems*. Reading, MA: Addison Wesley, 1989.
- [18] S. Zhak, M. Baker, and R. Sarpeshkar, "A low-power wide-dynamic-range envelope detector," Research Lab. Electron., Mass. Inst. Technol., Cambridge, Memo 664, 2003.
- [19] Y. Tsvividis, "Externally linear, time-invariant systems and their application to companding signal processors," *IEEE Trans. Circuits Systems II: Analog Digit. Signal Process.*, vol. 44, no. 2, pp. 65–85, Feb. 1997.
- [20] R. T. Edwards and G. Cauwenberghs, "Synthesis of log-domain filters from first-order building blocks," *Analog Integrated Circuits Signal Process.*, vol. 22, no. 2/3, pp. 177–186, Mar. 2000.
- [21] F. G. Zeng and R. V. Shannon, "Psychophysical laws revealed by electric hearing," *Neuroreport*, vol. 10, no. 9, pp. 1931–1935, Jun. 1999.
- [22] P. Gray, R. Meyer, P. Hurst, and S. Lewis, *Analysis and Design of Analog Integrated Circuits*, 4th ed. New York: Wiley, 2001.
- [23] M. O'Halloran and R. Sarpeshkar, "A low open-loop gain high-PSRR micropower CMOS amplifier for mixed-signal applications," in *Proc. IEEE Int. Symp. Circuits and Systems*, vol. 2, Scottsdale, AZ, May 2002, pp. 424–427.
- [24] C. A. Mead and T. Delbruck, "Scanners for visualizing activity of analog VLSI circuitry," California Inst. Technol., Pasadena, CA, Caltech Computation and Neural Systems Program Memo 11, Jun 1991.
- [25] N. Kumar, G. Cauwenberghs, and A. G. Andreou, "Auditory feature extraction using self-timed, continuous-time discrete-signal processing circuits," *IEEE Trans. Circuits Systems-II: Analog Digit. Signal Process.*, vol. 44, no. 9, pp. 723–728, Sep. 1997.
- [26] R. Sarpeshkar, M. Baker, C. Salthouse, J. J. Sit, L. Turicchia, and S. Zhak, "An analog bionic ear processor with zero-crossing detection," in *Proc. IEEE Int. Solid-State Circuits Conf. (ISSCC)*, San Francisco, CA, Feb. 2005, Paper 4.2, pp. 78–79.



Rahul Sarpeshkar received the Bachelors degrees in electrical engineering and physics at the Massachusetts Institute of Technology (MIT), Cambridge, and the Ph.D. degree from the California Institute of Technology (Caltech), Pasadena.

He was with Bell Labs. as a member of technical staff. Since 1999, he has been on the faculty of MIT's Electrical Engineering and Computer Science Department where he heads a research group on Analog VLSI and Biological Systems, and is currently the Robert J. Shillman Associate Professor. He holds

over a dozen patents and has authored several publications including one that was featured on the cover of NATURE. His research interests include analog and mixed-signal VLSI, ultra-low-power circuits and systems, biologically inspired circuits and systems, and control theory.

He has received several awards including the Packard Fellow award given to outstanding young faculty, the ONR Young Investigator Award, and the NSF Career Award.



Christopher Salthouse received the Bachelors and Masters degrees from the Massachusetts Institute of Technology (MIT) in 2000. Since then he has been working towards the Ph.D. degree at MIT in electrical engineering.

His research focuses on building systems including significant analog components. Currently, this focuses on analog circuit design of micropower filter systems. Past projects have included his master's thesis in the area of power electronics and previous research on software radio.



Ji-Jon Sit received Bachelors degrees in electrical engineering and computer science from Yale University, New Haven, CT, in 2000, and a Masters degree in electrical engineering from the Massachusetts Institute of Technology (MIT), Cambridge, in 2002. He is currently working towards the Ph.D. degree at MIT, on low-power electrode stimulation circuits.



Michael W. Baker received the S.B. degree in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, in 2000. He spent summers with Bell Laboratories, Holmdel, NJ, where he worked with several groups, including the Networked Multimedia Research Department, the Biological Computation Department, and the High-speed Electronics Research Department. He received the M.Eng. Degree from MIT on high-linearity mixers for 5-GHz receivers working with Lucent Technologies'

High-Speed group. He is currently working towards the Ph.D. degree in analog VLSI and biological systems at MIT.

His research interests include biomedical implants for the deaf, low-power/low-noise integrated analog design, and integrated radio-frequency circuits.



Serhii M. Zhak received the S.M. degree in physics from the Moscow Institute of Physics and Technology, Moscow, Russia, in 1999. He is currently working toward the Ph.D. degree in electrical engineering and computer science in the analog VLSI and biological systems at the Massachusetts Institute of Technology (MIT), Cambridge.

His research interests include biomedical implants for the deaf, low-power and low-noise integrated analog design, and integrated radio-frequency circuits.



Timothy K.-T. Lu received the Bachelors and Masters degrees in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, in 2003. He has also worked on software engineering at the IBM Almaden Research Center and the MIT Media Laboratory. His master's thesis focused on the modeling of outer hair cells and feedback loops in the mammalian cochlea. He is currently pursuing the Ph.D. degree in medical engineering/medical physics at the Harvard/MIT Division of Health Sciences and Technology (HST),

Cambridge. His research interests include synthetic biology, biological systems modeling, novel strategies for combating antibiotic resistance, and translating basic research into clinical medicine.

Mr. Lu was awarded a Howard Hughes Medical Institute (HHMI) Predoctoral Fellowship in Biological Sciences in 2003 and received a Siebel Scholar Award in 2003.



Lorenzo Turicchia was born in Bologna, Italy. He received the Laurea degree in electrical engineering from the University of Padova, Padua, Italy, and the Ph.D. degree in computer science from the Department of Mathematics and Computer Science, University of Udine, Udine, Italy. In 2002, he joined the Analog VLSI and Biological Systems group at the Massachusetts Institute of Technology, Cambridge, to work on and finish the Ph.D. degree where he is currently a Postdoctoral Associate.

His main research interest is in nonlinear signal processing, such as audio signal processing and biomedical signal processing.



Stephanie Balster received the S.B. degree in electrical engineering from the Massachusetts Institute of Technology (MIT), Cambridge, in 2004. She is currently working towards the M.Eng. degree at MIT in electrical engineering and computer science.