Log-Domain Circuits for Auditory Signal Processing

R. Timothy Edwards[1] and **Gert Cauwenberghs**[2]

[1] Johns Hopkins University
Applied Physics Laboratory
11100 Johns Hopkins Road
Laurel, MD 20723-6099

[2] Johns Hopkins University
Electrical and Computer Engineering
3400 N. Charles St.
Baltimore, MD 21218-2686

August 6, 1999

Abstract

The theory and practice of log-domain filter design has reached the point where it is possible to incorporate log-domain filter structures into large current-mode VLSI systems. We report on interface circuits used to implement a current-mode frontend filterbank and feature extractor for acoustic pattern recognition. These circuits maintain a log-domain structure, acting on the unexpanded filter output to execute such functions as peak-to-peak voltage measurement, full-wave rectification, smoothing, and normalization. A fabricated VLSI sixteen-channel filterbank feature extractor exhibits 40 dB resolution for short-term energy envelope measurements of class A log-domain second-order bandpass filter outputs.

1 Introduction

Most of the development of log-domain principles and theory have been directed toward the purpose of analyzing and synthesizing filters [1, 2, 3, 4, 5]. This is the natural course of action for developing a solid foundation for log-domain methodology. This paper takes a different tack in outlining the development of current-mode circuits with log-domain properties targeting applications other than filtering.

This research arose out of a project requiring a bandpass filterbank for use in an acoustic signal processing system. We have previously published our results on the log-domain filterbank core of the system, concentrating on methods addressing issues related to filter design in the audio frequency range and the problem of compensation for mismatch of filter parameters across the bank of parallel filters [6, 7]. Mismatch can be extremely poor in a VLSI fabrication process, even one nominally geared toward analog designs. For purposes of signal processing (specifically, acoustic event classification), the quantities of interest are short-term spectral band energies, although it is also possible to design systems around zero-crossing measurements [8]. There exist multiple ways to approach the design of a circuit to extract the short-term energy envelope of a filter output as a directly measurable value. Two obvious approaches, which will be considered here, are rectification followed by smoothing, and peak detection and tracking. Each of these methods is fairly simple to implement in a well-controlled environment. In keeping with our system design goal of minimizing mismatch across channels in a VLSI filterbank system, we will consider designs which are robust to offset errors originating at the filter outputs and other consequences of mismatch.

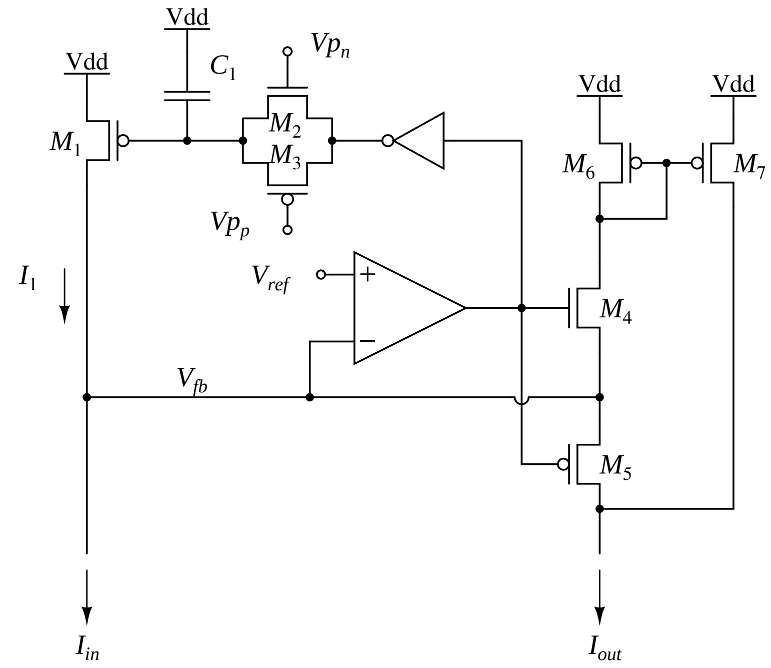

Figure 1: Simplified schematic of the adaptive current full-wave rectifier.

2 Methods for Determining a Short-Term Energy Envelope

The application for which the bandpass filterbank was designed calls for rectification and smoothing of each bandpass filter channel output for the purpose of evaluating the short-term average energy over the frequency band of the channel. The rectified signal is smoothed (integrated) with a time constant of 1 to 2 ms, a time scale appropriate for analysis of audio-frequency events.

We evaluate two ways to compute the needed short-term energy envelope: The first way is to implement exactly the method described above by performing a half-wave or full-wave rectification followed by a log-domain lowpass filter for smoothing. Rectification requires that the midpoint of the signal be known, and from prior experience we have observed that offsets introduced by mismatch and circuit nonidealities can severely complicate the task. To address this problem, we also investigated novel circuits for determining the true peak-to-peak height of a filter output, avoiding the necessity of knowing the signal midpoint.

3 Method 1: Signal Rectification and Smoothing

Figure 1 shows a simplified version of the circuit used for the current full-wave signal rectifier. This circuit employs long-term adaptation through feedback, using a *charge pump* to achieve the long time scale required (about 1/10 s to 1 s would be typical with respect to the audio input signal). The charge pump (including the inverter, transistors M_2 and M_3, and capacitor C_1) acts to keep the current I_1 through transistor M_1 equal to the long-term average input current $\langle I_{in} \rangle$ using feedback from the operational amplifier: If the input current exceeds the current through M_1, then V_{fb} drops, passing V_{ref}, as the system driving the input is driven towards cutoff. Consequently, the op-amp output rises, driving the inverter output to ground. When that side of the charge pump is grounded, the nMOS transistor in the charge pump pair is activated and sinks current, drawing charge off of the capacitor node. The lowering voltage on the gate of M_1 increases its drain-to-source current I_1. This negative feedback process continues until I_1 matches I_{in}. The biases Vp_n and Vp_p on the gates of transistors M_2 and M_3, respectively, determine the rate at which charge can be placed on or removed from C_1. The transistors have fixed gate voltages but are switched at the source, a configuration which minimizes charge injection onto the sensitive capacitor node. The charge pump system acts like a crude, nonlinear lowpass filter with the biases determining the time constant of the filter. Time constants are so large in this system that the nonlinearity of the system is irrelevant. I_1 changes very slowly compared to I_{in}, and thus will eventually drift until it settles at the DC average of I_{in}. It will adapt to any DC drift in the input which is longer than the time constant of the charge pump adaptation.

The current-switched architecture of the system ensures that the current difference $I_{in} - I_1$ always has a source or sink, leaving V_{fb} virtually unchanged (high gain on the op-amp, of course, makes the system respond faster and more cleanly). When I_{in} exceeds its midpoint I_1, M_5 is cut off and current is drawn through M_4, mirrored, and produces a positive I_{out}. When I_{in} falls below I_1, M_4 is shut off (and the mirror M_6, M_7 with it), and positive current again flows to the output, this time directly through M_5. Apart from mismatch errors, the output current always equals $|I_{in} - I_1|$.

In addition to the components shown in the simplified schematic, the complete (fabricated) design has 1) a source-degenerated ("diode-connected") pMOS transistor between M_1 and Vdd to raise the required voltage across C_1, 2) a current-conveyor mirror in place of the simple current mirror (M_6 and M_7), and 3) the addition of a small coupling capacitor between the op-amp negative input and the output.

After signal rectification, smoothing is a matter of using a lowpass filter, the quintessential class-A log-domain structure which is known to be very robust (Figure 8). Multiple lowpass filter circuits can be cascaded to form an n-th order filter, as shown in Figure 2. Translinear behavior in any class-A circuit

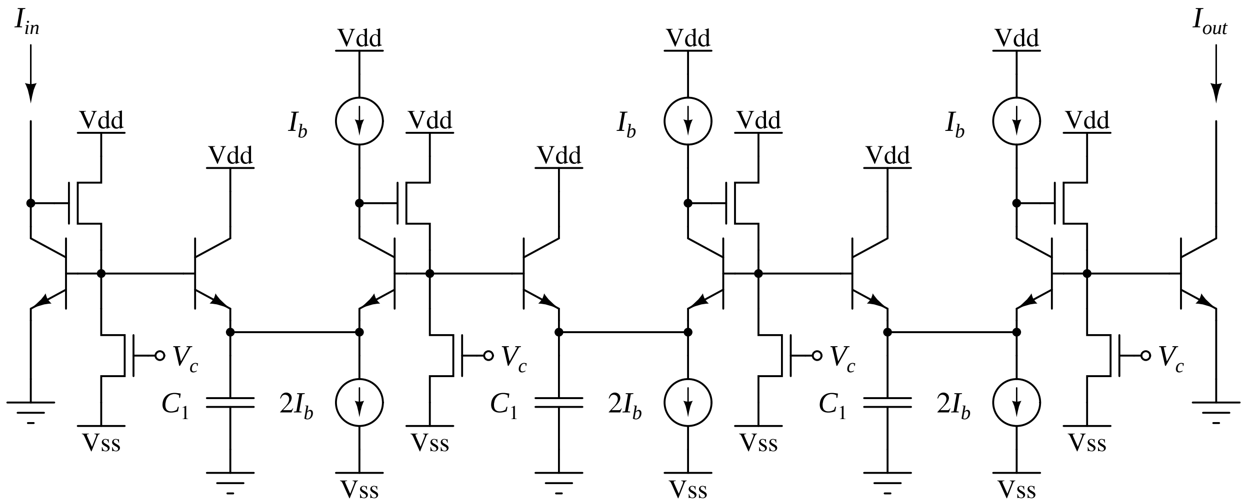

Figure 2: Cascaded log-domain lowpass filters.

breaks down near zero current due to the difference between the exponential approximation for transistor current behavior $I_c \approx I_0 e^{V_{be}/V_t}$ and the slightly better relation $I_c \approx I_0 \left(e^{V_{be}/V_t} - 1\right)$. Departure of the two relations occurs at currents on the order of I_0, or 1 fA for a typical BiCMOS process. This is well below the noise floor of the filter, ensuring that the log-domain lowpass filter can accept input directly from the signal rectifier without requiring any further biasing. The smoothed filter output can reach arbitrarily low current values during quiet periods of the system input.

4 Method 2: Signal Peak-Peak Detector

We developed a novel log-domain circuit for determining the height of a current-mode signal without requiring knowledge of the mean value of the input waveform. The operation of the circuit is based on the

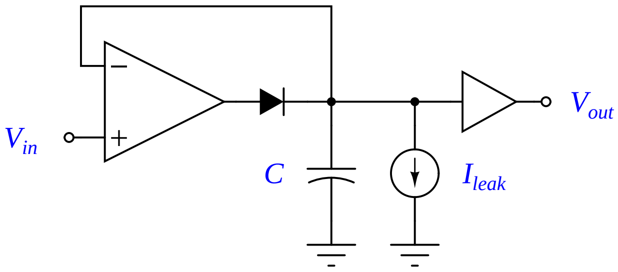

Figure 3: Simple diode-based peak detector.

usual "leaky integrator" peak detector, which consists of a diode used to rectify the input, a capacitor to hold the peak voltage value when the diode is reverse-biased, and a small constant current source leaking charge slowly off of the capacitor (producing a linear voltage drop per unit time) so that the circuit can track the input rather than hold indefinitely at the largest input voltage. Putting the diode in the feedback loop of an amplifier circumvents the voltage drop across the diode, so that the output closely follows the rising edge of the input when the diode is not reverse biased. This circuit is shown in Figure 3.

A novel aspect of our circuit is that it operates *differentially* by containing two subcircuits, one which computes the signal peak (maximum), and one which computes the signal trough (minimum). The maximum-detecting circuit leaks toward ground, stopping when its value equals that of the input. The minimum-detecting circuit leaks toward the positive power rail, also stopping when its value equals that of the input. The circuit output is the result of subtracting the minimum-detecting circuit output from the maximum-detecting circuit output, and yields a reasonably accurate measurement of the input signal's instantaneous peak-to-peak amplitude (see Figure 4). Due to the leaky capacitor architecture, the circuit responds quickly

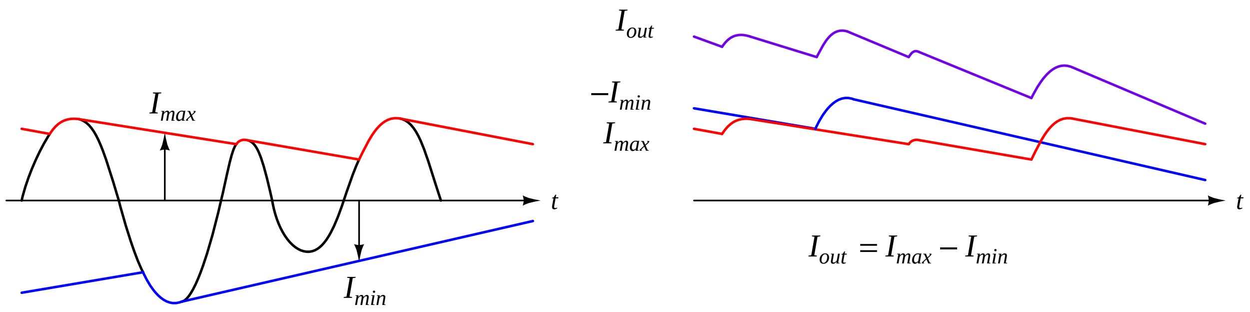

Figure 4: Behavior of the peak-peak detector.

to sharp rising edges in the input but relaxes slowly to zero (with the rate determined by the leak current) on the falling edge of the input.

The use of signals as currents allows a simple implementation of the diode rectifier circuit using a current conveyor structure. This structure is shown in Figure 5. The circuit qualifies as a "log-domain" circuit. The capacitor nodes (V_{max}, V_{min}) hold voltage values which are the logarithm of the maximum (or minimum) input current (I_{in}) and are expanded exponentially when converted from the voltage back into a current (I_{max}, I_{min}). A major consequence of this choice of architecture is that although the leaky current source (M_3, M_6) on each of the capacitors (C_1, C_2) causes a linear change of voltage per unit time, as does the simple voltage-mode peak detector of Figure 3, the effect is that the output current changes *exponentially* with time (or as a square-law, depending on the bias on the MOS transistors (M_9, M_{10}) which do the expansion). This effect is difficult to avoid in a current-mode circuit. However, it is minimized in a class-A architecture by the DC bias of I_{in}, fixed such that the signal amplitude is considerably smaller than its bias.

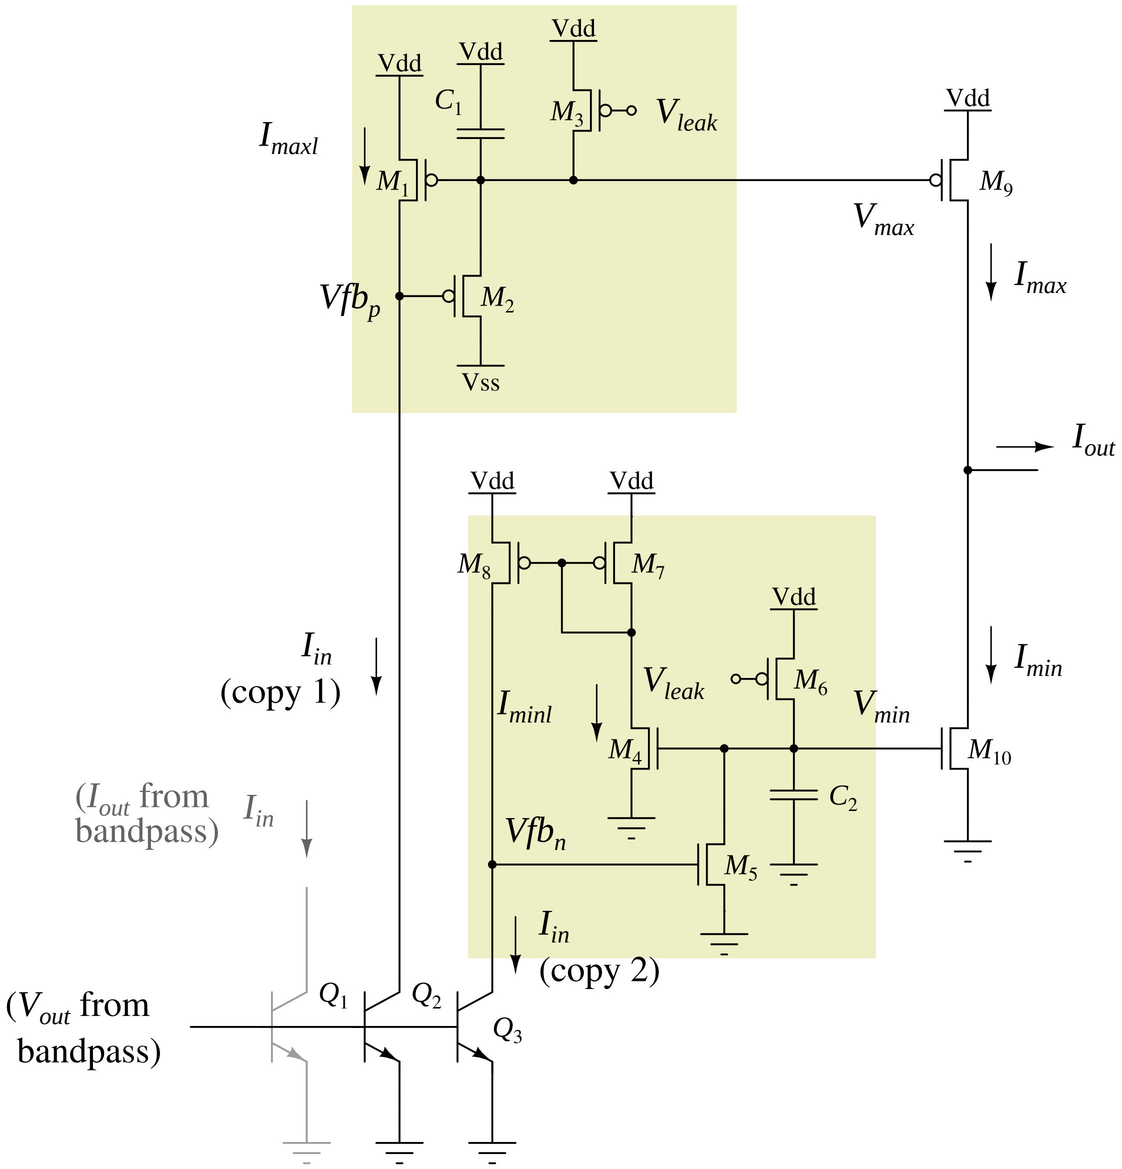

Figure 5: Peak-peak detector circuit, simplified.

The differential current change of the maximum and minimum currents is small, and so traverses a shorter range of the nonlinear expansion function. To a first-order approximation, the change in output current is nearly linear with respect to time.

Here is a brief description of how the circuit operates: On the maximum current-finding side, when current I_{in} tries to exceed I_{maxl}, voltage Vfb_p drops to near ground to reduce the collector current of Q_2 to match I_{maxl}. The lowering of Vfb_p turns on transistor M_2, causing the circuit formed by transistors M_1 and M_2 to become a current conveyor. V_{max} drops immediately until I_{maxl} reaches the same value as I_{in}. This current-conveyor feedback ensures that as long as the input current is equal to I_{maxl} and rising, I_{max} will continue to track it. On the other hand, when I_{in} tries to drop below I_{maxl}, Vfb_p rises toward Vdd so that M_1 is driven toward cutoff and I_{maxl} drops until it becomes equal to I_{in}. When Vfb_p is close to Vdd, transistor M_2 is cut off, and leaves V_{max} floating save for the small leakage through M_3. So although I_{maxl} is forced to track the input current, the output I_{max} remains at its maximum level, only drifting slowly toward zero until either it falls below the level of the input current, or the input current catches up with its value.

The minimum-tracking circuit is symmetric to the maximum-tracking circuit. To keep the rate of charge leakage equal in the two subcircuits, both circuits use the same size capacitor and pMOS transistor gated by the same bias voltage. The capacitor voltage always leaks toward Vdd. The capacitor node voltage drives the output transistor gate directly. Where it drives a pMOS transistor, the (maximum) current drifts toward lower values; where it drives an nMOS transistor, the (minimum) current drifts toward higher values.

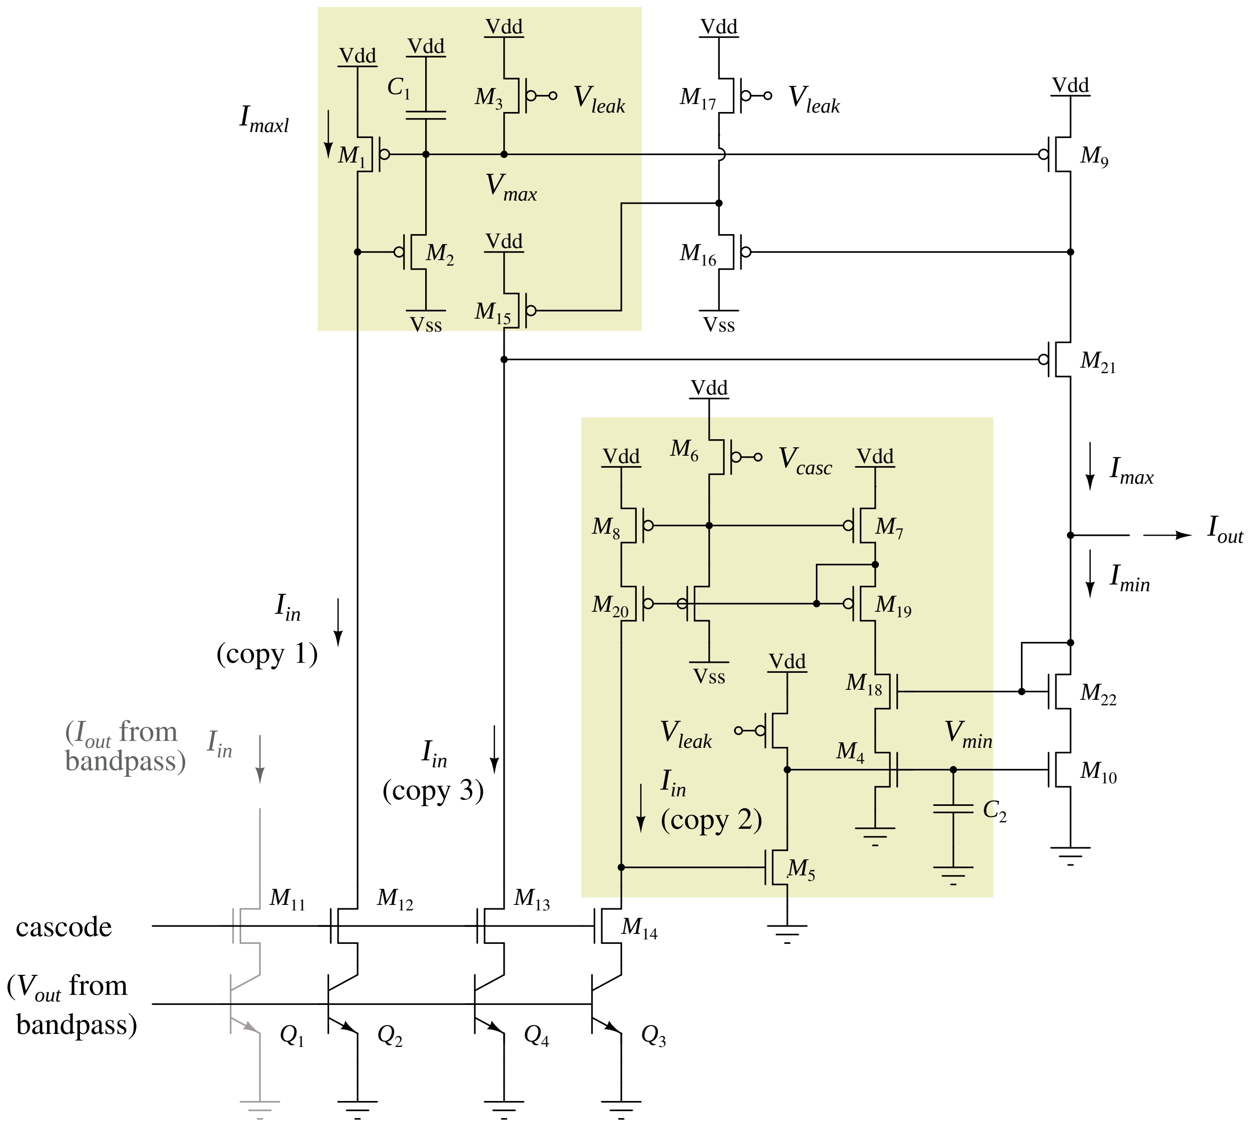

Figure 6: Peak-peak detector with cascodes, as fabricated.

Figure 6 is the final version of the circuit, as implemented in a BiCMOS fabrication process. Extra circuitry minimizes offsets seen at the output due to mismatch between the maximum-detecting circuit and the minimum-detecting circuit. The design takes into account all the major systematic sources of mismatch, primarily by cascoding devices to reduce the Early effect. Most of the additional transistors (Q_4, M_{15}, M_{16}, M_{17}) exist merely to compute the proper cascode bias for transistor M_{21} to keep I_{max} consistent with I_{maxl}.

We employ a short cascade of log-domain first-order lowpass filters at the end of the circuit to smooth out the bumps in the peak detector output, as was done for the signal rectification and smoothing described in the previous section. This final operation presents a much smoother signal to the system output at the expense of only a small overhead in circuitry.

5 Current-Mode L-1 Normalization

Our acoustic signal classifier model calls for a transformation of the energy envelope by a normalization across all channels. The "L-1" normalization, in the current domain, follows the following equation:

$$I_{out}(i) = I_{norm} \frac{I_{in}(i)}{\sum_{j=1}^{N} I_{in}(j)} \tag{1}$$

In the application, the L-1 norm forces the outputs to maintain the same total value as a form of automatic gain control.

$$\sum_{j=1}^{N} I_{out}(j) = I_{norm} \tag{2}$$

A circuit implementing an L-1 normalization in the current domain has been known for many years: it is one of the "original" translinear circuits. It has a simple and elegant form and as such makes a voltage-mode implementation almost unthinkable. The circuit (Figure 7), invented by Barry Gilbert [9], is known

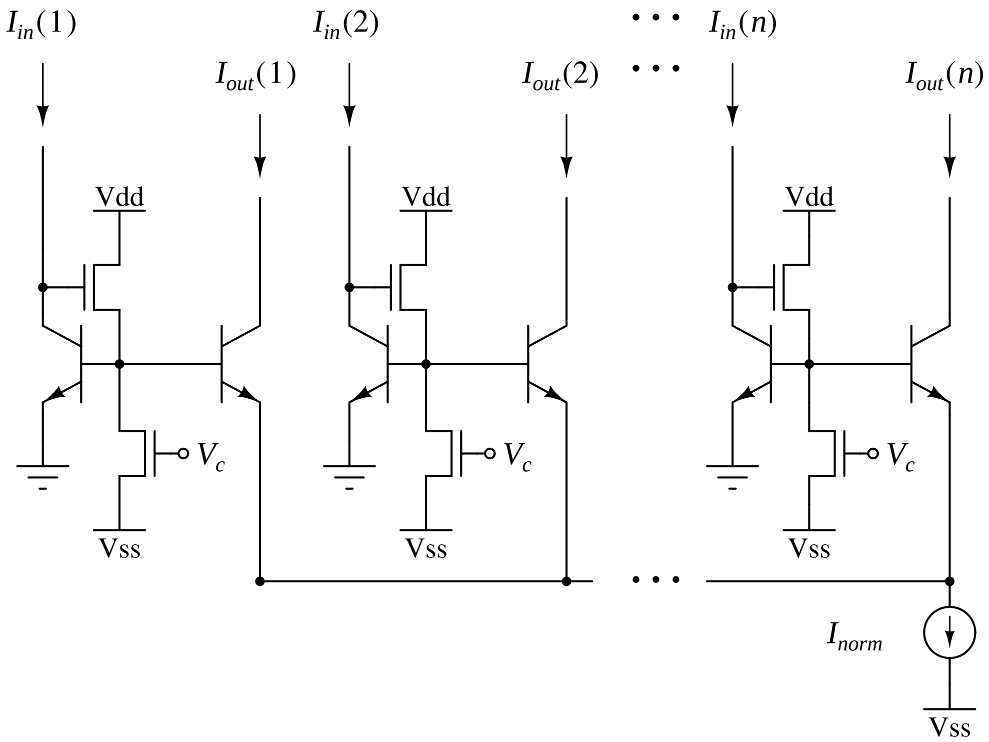

Figure 7: L-1 normalization circuit, after Gilbert [9].

(at least in analog VLSI circles) as the "Gilbert normalizer" [10]. The bias current I_{norm} is a simple way to assure the relationship of Equation (2) (excepting the slight difference between the desired I_{norm} and the true I_{norm} due to the BJT β values). Otherwise, the circuit is simply made up of many translinear loops, for which

$$I_{in}(i)I_{out}(j) = I_{in}(j)I_{out}(i) \qquad \forall\, i, j \tag{3}$$

Solve by adding together all of the loop equations for input/output pair i:

$$\sum_{j=1}^{N} I_{out}(j) = \sum_{j=1}^{N} I_{in}(j)\frac{I_{in}(i)}{I_{out}(i)} \tag{4}$$

$$I_{norm} = \sum_{j=1}^{N} I_{in}(j)\frac{I_{in}(i)}{I_{out}(i)} \tag{5}$$

$$I_{out}(i) = I_{norm}\frac{I_{in}(i)}{\sum_{j=1}^{N} I_{in}(j)} \tag{6}$$

6 Experimental Results

We have fabricated and tested several different VLSI architectures implementing a complete signal-processing system [11] containing a filterbank frontend of log-domain bandpass filters. The filterbank has 16 parallel channels each consisting of two cascaded log-domain second-order bandpass filters followed by the log-domain peak-detecting circuit described in Section 4, a smoothing (lowpass) filter, and the normalizer circuit described in Section 5. The acoustic signal classifier for which this filterbank was designed is also a current-mode system which accepts as input the unipolar current outputs of the frontend system, without any conversions.

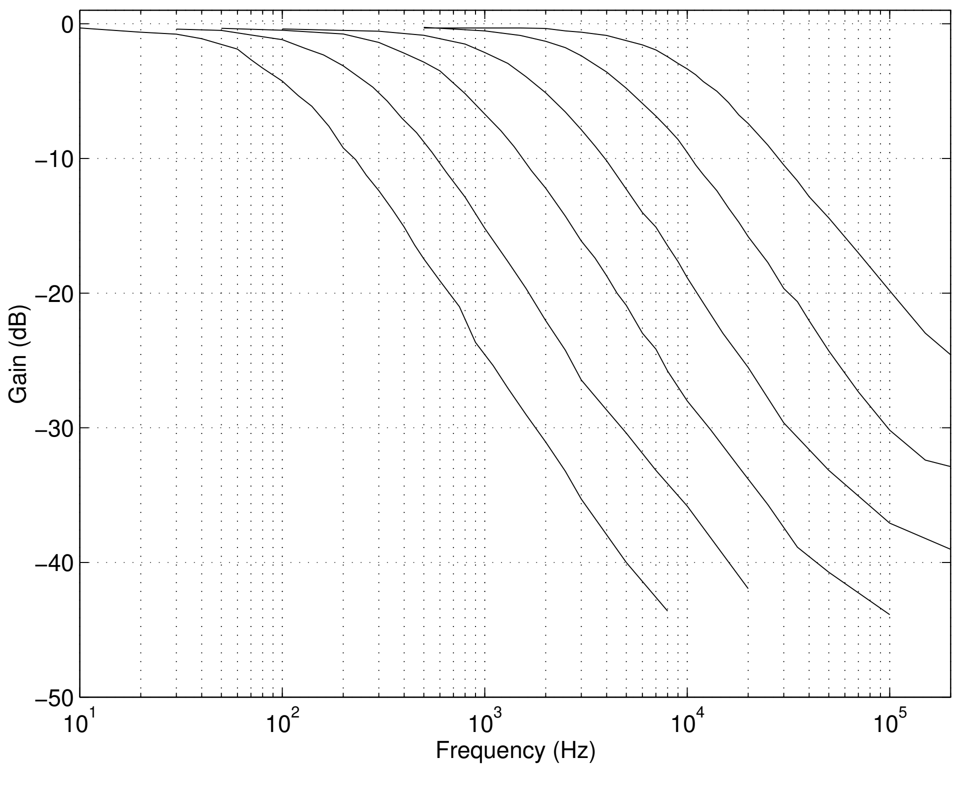

Figure 8: Measured magnitude response of the log-domain first-order lowpass filter from the test chip.

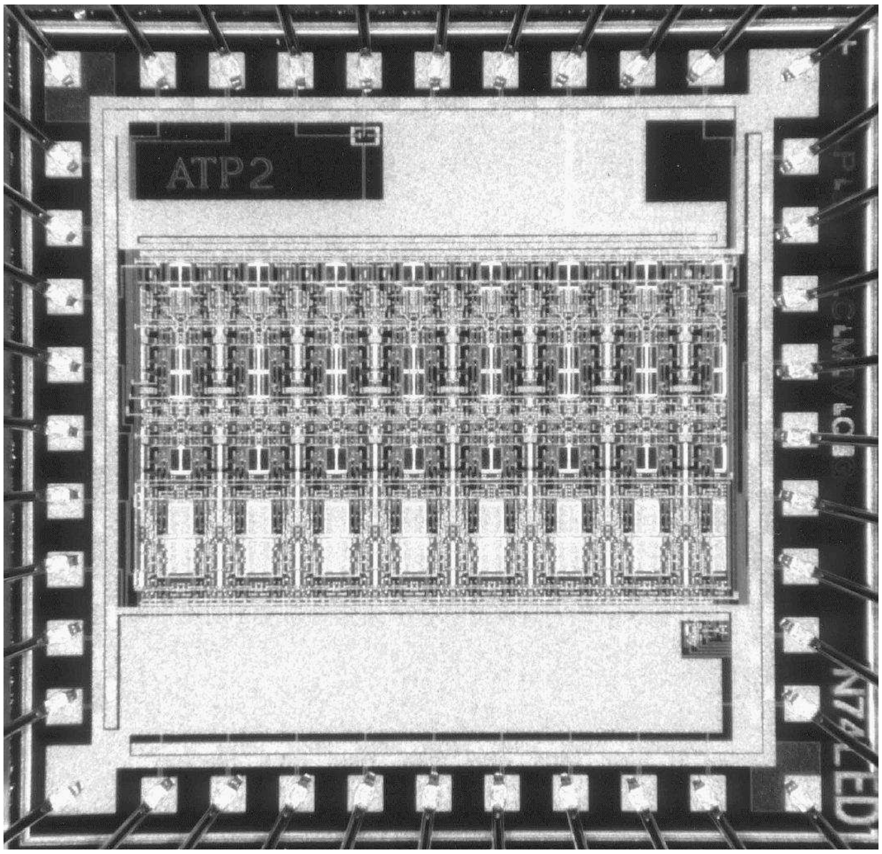

Figure 9: Photograph of the fifteen-channel bandpass filterbank fabricated in 1.2 μm technology inside a 2.2 mm × 2.2 mm padframe.

The chips were fabricated on 2 mm × 2 mm die in a standard 1.2 μm analog n-well BiCMOS technology, with double metal routing, double poly for linear capacitors, and a p-base implant for creating vertical NPN bipolar transistors. For high-density integration, we limited capacitors to about 2 pF, used single-ended (class-A) circuit architectures throughout, and made all bipolar transistors minimum size.

Figure 8 shows measurements taken from a cascade of first-order lowpass filters on a test chip. This is the filter used as the smoothing filter at the output of the rectifier or peak detector. The circuit is that shown in Figure 2, with a capacitor value of about 1 pF. The circuit is quite robust and the response maps well over the entire audio frequency range. The dynamic range is limited by an approximately 50 dB noise floor, which is at least 15 dB better than any of the circuits which precede it in the signal processing flow.

The remaining data were measured from the filterbank system chip, shown in the photograph of Figure 9. All measurements shown were made with a power supply of 5 V, though the system showed similar characteristics from 4 V to 6 V. Cascoded MOS transistors implementing the current sources and mirrors plus the p-n junction voltage drops in the translinear loops limit circuit operation to a power supply of at least 4 V. Power consumption of the circuit is dominated by the DC bias added to the filter input and output (typical for class-A designs), which affects the SNR at the filter output. With a 5 V power supply and an input/output bias of 10 μA, the filterbank dissipates about 200 μW and has a dynamic range of about 35 dB.

The experimental procedure used to capture the data presented in this paper requires a system able to present a signal as a current and read the output as a current. The input is the most difficult, due to the small size of the input signal and the lack of precision current-mode signal generators. The circuit adopted for the system input is shown in Figure 11. The use of large resistances and multiple amplifier circuits to convert

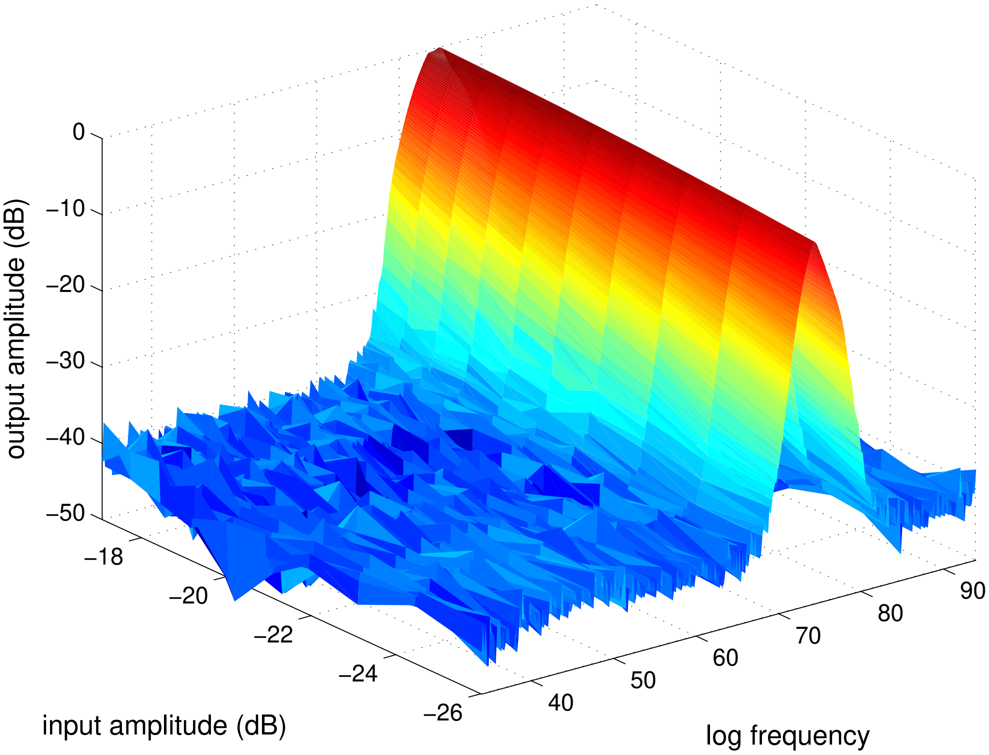

Figure 10: Measured current response of one bandpass channel in the filterbank system as a function of frequency and input amplitude.

signal and bias voltages into the input current limits noise performance of the entire system. For transfer function measurements like those presented below, output was measured at steady-state on a Keithley 486 picoammeter, with a large time constant on the smoothing filter.

Figure 10 shows the transfer function of a log-domain bandpass filter in the filterbank system, measured from the peak detector output as a DC current over a range of input sinusoid magnitudes and frequencies. The peak detector output departs from the filter transfer function on the lower edge of the passband, when the maximum rate of change of the filtered signal becomes less than the drop rate of the leaky capacitor. This condition prevents the peak detector from tracking, and its output falls to zero, as shown in Figure 12. In this particular application, the peak detector actually corrects for the shallow roll-off of the bandpass filter (an effect due to transistor mismatch), creating a sharper filter.

7 Conclusions

In this paper, we have outlined the design and use of current-mode circuits based on class-A log-domain architectures. These circuits perform signal processing functions complementary to the filters and oscillators primarily found in the log-domain literature; these functions include signal rectification and smoothing, peak detection and tracking, and normalization. Experimental data captured from a filterbank system demonstrates the operation of a peak-to-peak tracking circuit performing an energy envelope measurement on the output of a bandpass filter. This circuit tracks signal current peaks independently of DC bias and achieves 35 dB performance at 200 μW power dissipation. Additionally, it exhibits the property of sharpening the bandpass response, compensating for the shallow rolloff at low frequency typical of class-A log-domain bandpass filter designs.

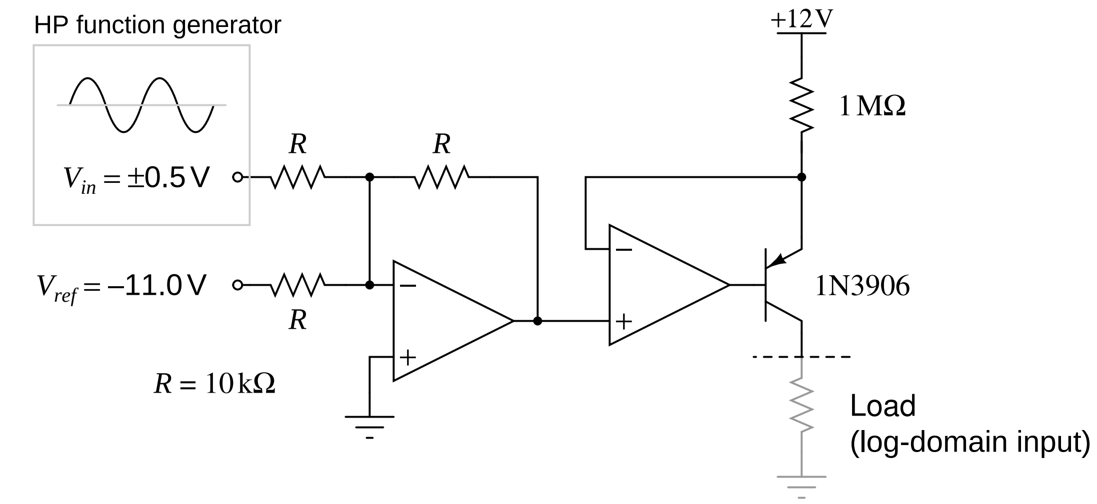

Figure 11: Test setup for driving the log-domain circuit input (signal and bias) from a standard function generator.

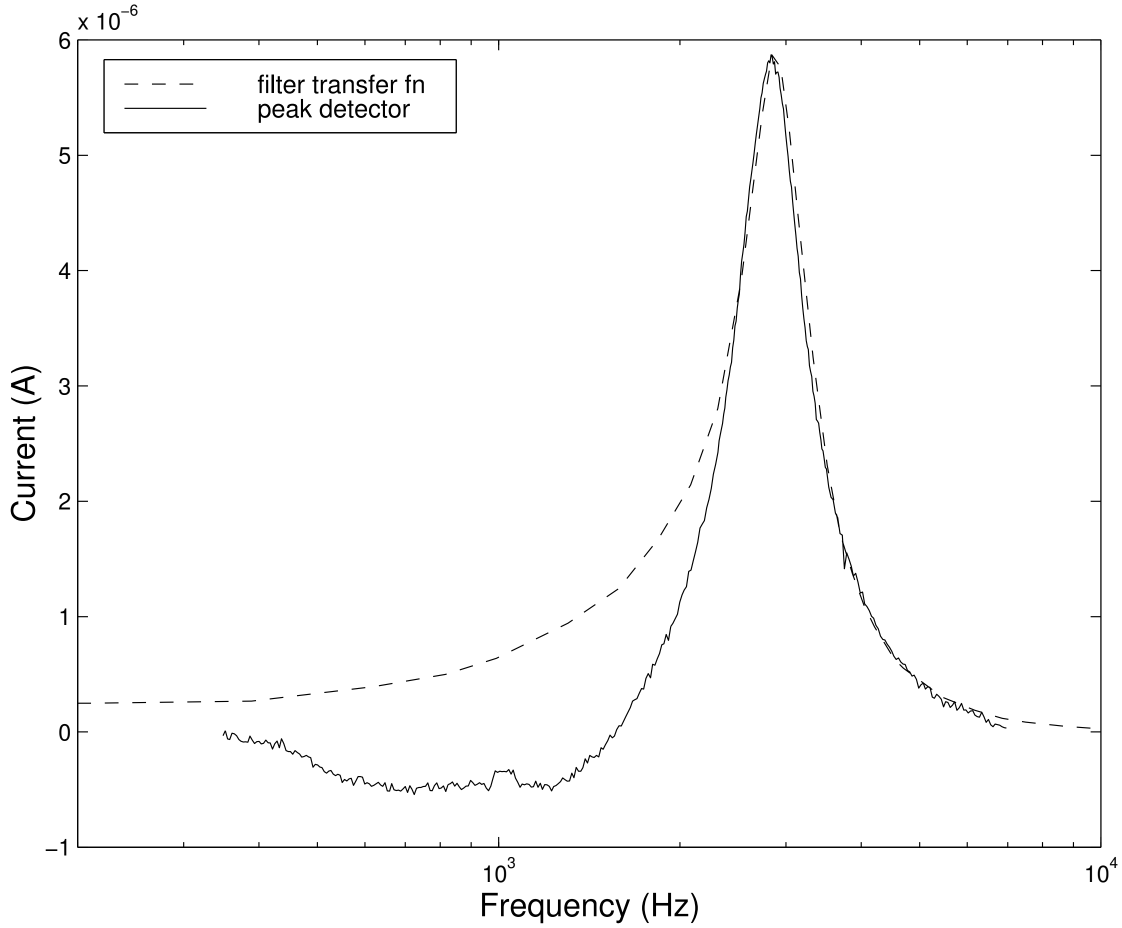

Figure 12: Comparison between the measured transfer function of the bandpass filter of a single channel (dotted line), normalized to the current output of the peak detector and smoothing filter on the same channel (solid line).

References

[1] D. Frey, “Log Domain Filtering for RF Applications,” *IEEE J. Solid-State Circuits*, **31** (10), pp. 1468–1475, Oct. 1996.

[2] D. Perry and G. Roberts, “The Design of Log-Domain Filters Based on the Operational Simulation of *LC* Ladders,” *IEEE Trans. Circuits Syst. II*, **43**, No. 11, pp. 763–774, Nov. 1996.

[3] E. M. Drakakis, A. Payne, and C. Toumazou, “Log-domain filters, translinear circuits and the Bernoulli Cell,” *Proc. ISCAS ’97*, **1**, pp. 501–504, 1997.

[4] R. Fox, M. Nagarajan, and J. Harris, “Practical Design of Single-Ended Log-Domain Filter Circuits,” *Proc. ISCAS ’97*, **1**, pp. 341–344, 1997.

[5] J. Mulder, A. van der Woerd, W. Serdijn, and A. Roermund, “General Current-Mode Analysis Method for Translinear Filters,” *IEEE Trans. Circuits Syst. I*, **44** (3), pp. 193–197, Mar. 1997.

[6] R. T. Edwards and G. Cauwenberghs “Synthesis of Log-Domain Filters from First-Order Building Blocks” (to appear in) *Analog Integrated Circuits and Signal Processing*, special issue on translinear circuits.

[7] R. T. Edwards and G. Cauwenberghs, “A Second-Order Log-Domain Bandpass Filter for Audio Frequency Applications,” *International Symposium on Circuits and Systems (ISCAS)*, Monterey, CA, June 1998.

[8] N. Kumar, W. Himmelbauer, G. Cauwenberghs and A. Andreou, “An Analog VLSI Chip with Asynchronous Interface for Auditory Feature Extraction,” *IEEE Trans. Circuits and Systems II: Analog and Digital Signal Processing*, **45** (5), pp 600–606, 1998.

[9] B. Gilbert, “A 16-channel array normalizer,” *IEEE J. Solid-State Circuits.*, **SC-19**, pp. 956–963, 1984.

[10] C. A. Mead, *Analog VLSI and Neural Systems*, Reading, MA, Addison-Wesley (1989).

[11] R. T. Edwards, G. Cauwenberghs, and F. J. Pineda, “A Mixed-Signal Correlator for Acoustic Transient Classification,” *Proc. ISCAS ’97*, **1**, pp. 621–625, Jun. 1997.