Chapter 3

Current-Mode Filterbank Frontend

3.1 Time-Frequency Representations using Filterbanks

There is a corollary to the observation at the conclusion of the previous chapter that our wavelet transform architecture is perhaps better utilized for things other than time-frequency decomposition: Not every time-frequency decomposition needs to be a wavelet transform. Of course, as the term "wavelet transform" or "wavelet decomposition" gets ever more loosely applied (not unlike the term "neural network"), almost any time-frequency decomposition can be described as a wavelet transform. This tends to blur the lines between what is a wavelet filter, a cochlear filter, or just a bandpass filterbank.

In my mind—and this is clearly a subjective interpretation—the three filterbank types mentioned above *should* be more clearly defined along the following lines:

- A *Bandpass Filterbank* is any architecture which splits its input into multiple, simultaneous outputs by passing the signal through a parallel set of bandpass channels, preferably with minimal overlap. The band centers are often spaced on a logarithmic scale, although that is not a necessity.

- A *Cochlear Filterbank* is any architecture reasonably attempting to model the signal processing properties of the mammalian auditory system, at very least including a model of basilar membrane mechanics. This differs from a bandpass filterbank in that the input passes through a long cascade of lowpass filter sections, from which the output is tapped at various intervals. In keeping with the psychology of auditory perception and known facts about the biological structure of the cochlea and its neural connections, the center frequency of the (highpassed or

bandpassed) taps typically are spaced logarithmically (or more accurately, on a mel scale).

- A *Wavelet Filterbank* is any architecture which implements a wavelet transform by attempting to create a distribution of outputs which covers the time-frequency plane with greatest efficiency. By necessity, the frequency bands are spaced on a logarithmic scale.

Note that the above definitions do not *preclude* any overlap. For example, although the outputs of a cochlear filter do not encode the input with anything approaching maximal efficiency, that does not necessarily mean that the mechanics of the inner ear plus subsequent neural processing do not eventually extract a maximally efficient wavelet-like representation from them. In a way, the cascaded serial processing of the cochlear filter is more faithful to the idea of "dilations" in wavelet transformations, particularly as adopted in DWT implementations, than is the parallel filterbank structure we used for the CWT processor. Cascaded processing avoids redundancy in the filtering and is therefore decidedly more power-efficient than parallel processing. However, problems with cascaded processing, such as cumulative noise and cumulative DC offset error, often make circuit implementations prohibitively difficult.

3.2 Parallel Filterbanks for Transient Classification

About the same time that we were realizing practical circuits for the continuous wavelet transform processor and achieving good results, I began work on a project involving the implementation of an algorithm for acoustic transient classification. The algorithm, initially conceived and developed by Fernando Pineda at the Johns Hopkins University Applied Physics Laboratory, was intended to be most conveniently implemented by analog circuits. The algorithm consisted of two major parts: First, a filterbank system, only roughly specified, from which the bandpassed outputs would be rectified and smoothed to encode the average energy in each channel over time (see Figure 3.1); and second, a classification system based on template correlation. The correlation system is described in detail in Chapter 4. This chapter concerns the design and development of an analog filterbank frontend system targeted for use with the acoustic transient correlator.

Preliminary trials of the classification algorithm used a cochlear filterbank circuit dubbed with the pseudo-acronym "HEEAR" (for the "Hopkins Electronic EAR"), developed in Andreas Andreou's group at Johns Hopkins University by Weimin Liu and later developed into a capable audio signal processing system by Nagendra Kumar and Paul Furth. Once I began work designing the template correlation system, however, an elegant design began to evolve, one which required

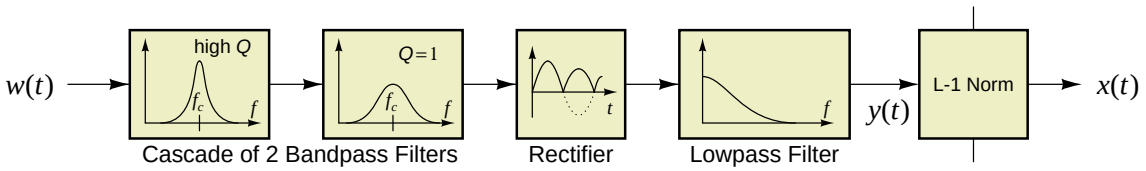

Figure 3.1: Frontend filterbank system—block diagram.

currents at the input. It appeared that it was going to be a difficult task to figure out how to couple a voltage-mode circuit like the HEEAR chip, or even the continuous wavelet transform processor, to the template correlator subsystem. Here is the fundamental problem: Currents are easily turned into voltages by feeding the current through a resistor to ground and buffering the resulting voltage across the resistor. Turning voltages into currents (that is, creating a linear voltage-controlled current source) is a more involved matter involving active feedback. Both methods require that the resistances required are compatible with the process: resistances larger than about 10 kΩ are inconvenient to realize in VLSI, while resistances larger than about 1 MΩ become incompatible with discrete component design.

Subthreshold circuits and, in general, current-mode circuits aimed at ultra-low power applications, typically have currents in the nanoamp range. The correlation processor to be described in detail in Section 4.4.1 and following is such a circuit, with the input currents topping out around 1 µA. Such a system should operate marginally well using an off-chip interface of discrete components between the frontend analog signal processing section and the template correlator. But since one of our goals is to realize single-chip solutions for large systems, we prefer to have the frontend interface directly to the correlator without the need to send signals off-chip, incurring the associated load capacitance, noise pickup, and pad-limited design. This being the goal, either we must devise a voltage-mode correlator or else devise a current-mode frontend.

3.3 Current-Mode Filters for Current-Mode Applications

From as far back as the first few years I began designing CMOS circuits operating in the subthreshold regime, I was aware of the strong duality between voltage-mode and current-mode circuit design. The idea of duality, however, cannot be too rigorously applied and is incomplete in a number of areas. Roughly speaking, voltage-mode circuits typically are based on transconductance amplifiers, whereas current-mode circuits typically are based on current conveyors. Because cur-

rent conveyors have fewer transistors than amplifiers, current-mode circuits can be made extremely compact compared to their voltage-mode equivalents, if such an equivalence exists. For example, adding two currents requires no components at all, just the merging of two current-carrying wires, an operation which is perfectly linear at all scales. Adding voltages requires an amplifier circuit and has a limited range of linearity, especially when the amplifiers are CMOS transconductance amplifiers operating in the subthreshold regime (see Appendix A). The Gilbert multiplier (see Chapter 2), requires at least 17 transistors and has a rather limited linear range due to the same differential input pair configuration as a transconductance amplifier. By contrast, multiplication of currents is easily obtained with a translinear multiplier (see Section 3.6) consisting of four transistors and two current sources, and is linear over many orders of magnitude of the input, especially if implemented with bipolar transistors.

Each "school" of design has its strong and weak points, though. While the strength of current-mode design is in the compactness and elegant simplicity of its circuits and the capability of wide linear dynamic range, its weakness is in filter circuit design. One factor hampering the design of our original analog circuit-based wavelet transform chip was the combination of simple arithmetic functions such as addition and multiplication which would be most conveniently implemented in the current domain, with filtering functions best implemented with voltage-mode transconductance-C filters.

For a long time, very few current-mode filter designs were available, and those that did exist were not amenable or adaptable to CMOS VLSI technology. However, in the early 1990s, interest began to grow around the topic of "log-domain" filters. Researchers discovered simple and elegant circuit topologies which lent themselves readily to VLSI implementation, and suddenly the possibility existed to couple current-mode arithmetic with current-mode filtering to realize all manner of novel systems. Because I was aware of many potential applications for current-mode filtering in analog and mixed-mode VLSI design, I was among the first researchers to begin attempting to build complete systems around log-domain filters and current-mode arithmetic.

Unfortunately, embarking on systems design while the underlying circuit design theory was still in its infancy was a daunting and difficult task, and required research that lead to novel contributions to the field of log-domain circuit theory. The following sections summarize our investigations and contributions to this burgeoning field, directed primarily toward the goal of building practical current-mode filters and related circuits for an acoustic-frequency analog signal processing system.

3.4 High-Level Simulations of the Filterbank Frontend

Preliminary research focused on simulating variations of the filterbank frontend using sampled recordings of acoustic transient events (see Chapter 4, Section 4.3.1). The purpose of the research was to confirm that the parallel, linear filterbank channels we anticipated using would be functionally equivalent to the cascaded, compressively-nonlinear HEEAR chip filters. Proving this in simulation essentially meant converting the continuous-time filter equations into the z-domain and using a computer to apply the filter equations to the original data presented to the HEEAR chip. Then both the simulated parallel filterbank outputs and the HEEAR chip outputs were used as inputs to a simulation of the acoustic transient correlation. Results, which can be found in Chapter 4, Section 4.3.5, showed that at least under the ideal conditions of simulation, the parallel filterbank frontend outputs yield similar results on the classification task, maintaining the classification rates of the HEEAR outputs and often doing better. The improved classification rates almost certainly are due in part to the difference between the noisy outputs of the HEEAR chip with all of its physical limitations versus the clean digital processing of the filterbank simulation.

The filterbank system described in the remainder of this chapter performs the steps shown in Figure 3.1. Input to the system is an acoustic signal. The simulation takes a sampled, digitized version of the acoustic signal. The HEEAR data acquisition system stored its outputs and a sampled version of the original input as 16 bit values sampled at 32 kHz. The input was bandlimited to under 16 kHz to prevent aliasing during sampling. Recording gain was not adjusted to maximize the resolution of the recording so that the output was effectively 8 bits resolution rather than 16. The fact that the HEEAR system copied its input to the output along with its own filtered outputs was useful for making sure that the simulated parallel filterbank system received the same inputs as the HEEAR system.

Figure 3.2 shows the raw sampled input signal for an example transient recording. This and the following figures are screen captures from the X11 interface to our simulation software.

The parallel filterbank frontend processes its inputs through N channels ($n = 1 \ldots N$), where each channel contains two cascaded second-order bandpass filters, as shown in Figure 3.1. The frequency-domain (Laplace domain) transfer function for each filter section is

$$H_n(s) = \frac{(1/Q)\tau_n s}{1 + (1/Q)\tau_n s + \tau_n^2 s^2} \tag{3.1}$$

where $\tau_n = 1/2\pi\omega_{c,n}$ and $\omega_{c,n}$ is the center frequency for channel n. Resonance value Q may be made a function of channel n; however, we enforced the restriction that this would be a *constant* Q

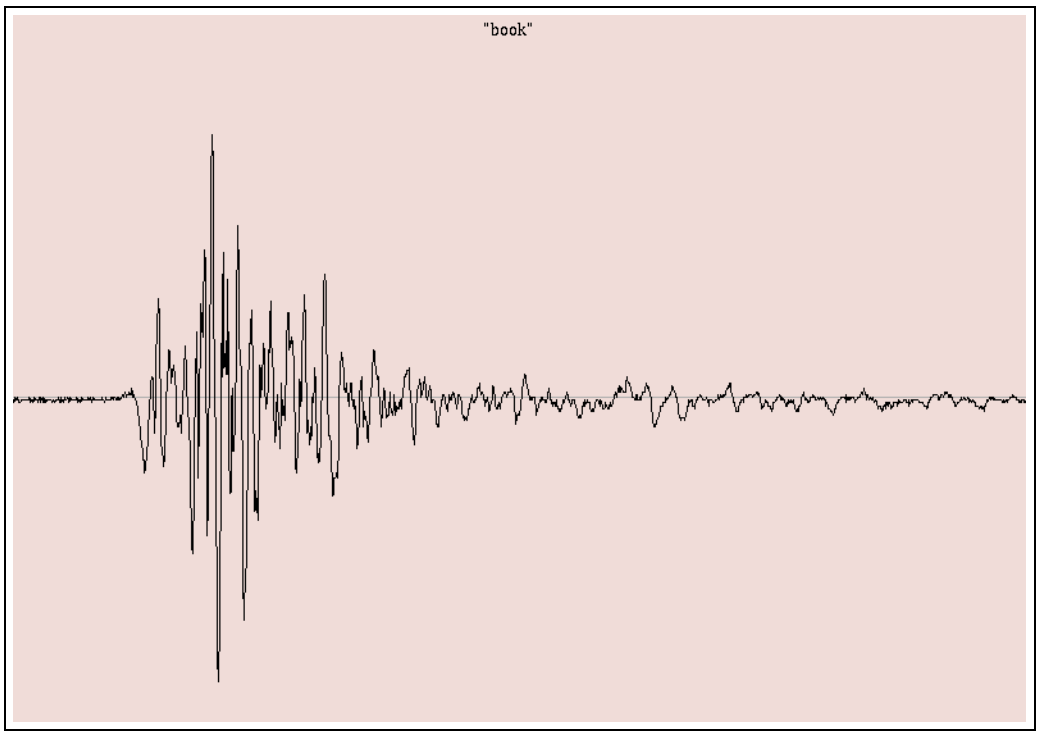

Figure 3.2: Sampled-data input from a recording of an acoustic transient—that of a book being dropped onto a desk.

filterbank.

Simulation of the bandpass filter requires a transformation from the s-domain into the z-domain. A standard way to do this is to use the *bilinear transform*,

$$H(z) = H(s)\Big|_{s=2f_s\frac{(z-1)}{(z+1)}} \tag{3.2}$$

which is valid when the sampling frequency is much larger than the filter bandwidth. The result of applying the bilinear tranform to Equation (3.1) is:

$$H_n(z) = \frac{b_{0,n} + b_{2,n}z^{-2}}{1 - a_{1,n}z^{-1} - a_{2,n}z^{-2}} \tag{3.3}$$

where

$$b_{0,n} = \frac{\tau_n}{1 + (\tau_n/Q) + \tau_n^2} \tag{3.4}$$

$$b_{2,n} = -b_{0,n} \tag{3.5}$$

$$a_{1,n} = \frac{2(\tau_n^2 - 1)}{1 + (\tau_n/Q) + \tau_n^2} \tag{3.6}$$

$$a_{2,n} = \frac{-(1 - (\tau_n/Q) + \tau_n^2)}{1 + (\tau_n/Q) + \tau_n^2} \tag{3.7}$$

and

$$\tau_n = \frac{f_s}{\pi f_{c,n}}. \tag{3.8}$$

with f_s being the sampling frequency (32 kHz), and $f_{c,n}$ being the center frequency of channel n. Thus each channel has its unique transfer function.

Center frequency values are rather arbitrarily defined; they are usually defined recursively, using

$$f_{c,n} = \gamma f_{c,n-1} \qquad n = 2 \ldots N \tag{3.9}$$

$$f_{c,1} = f_1 \tag{3.10}$$

for a purely exponential scale, where γ is constant and f_1 is the lower frequency bound of the filterbank, or

$$f_{c,n} = \gamma_0 f_{c,n-1} + \gamma_1 \qquad n = 2 \ldots N \tag{3.11}$$

$$f_{c,1} = f_1 \tag{3.12}$$

for a "mel" scale, where both γ_0 and γ_1 are constant and f_1 is the lower frequency bound of the filterbank. The values γ should ensure that channels overlap in the frequency domain based on some criterion such as -3 dB bandwidth or half-bandwidth.

The two cascaded sections have the combined transfer function

$$G_n(s) = (H_n(s))^2. \tag{3.13}$$

The result of this filtering operation on an example transient signal in simulation is shown in Figure 3.3. The algorithms which generate the filter characteristics and simulate the filter operation are detailed in Appendix D.

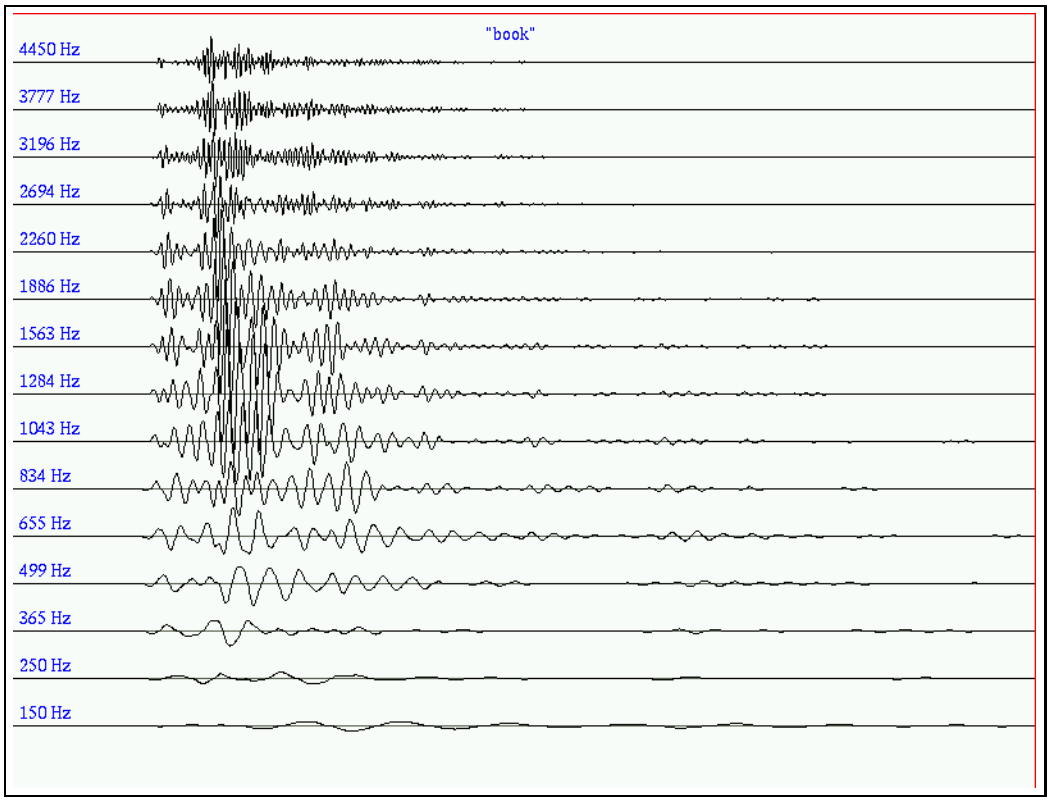

Figure 3.3: Bandpass-filtered acoustic transient input using two cascaded second-order filters on each channel.

The signal rectification process is trivial in simulation since the data are zero-mean after bandpass filtering. The lowpass function is, like the bandpass function, an s-domain filter function describing the analog hardware, transformed into the z domain via bilinear transform. Our choice of smoothing filter is exactly the same as that used for the continuous wavelet transform system: simple first-order lowpass sections in cascade. Each first-order filter has the transfer function

$$H_n(s) = \frac{1}{1 + \alpha_n s} \tag{3.14}$$

for which application of the bilinear transform yields

$$H_n(z) = \frac{d_n + d_n z^{-1}}{1 - c_n z^{-1}} \tag{3.15}$$

having coefficients

$$c_n = \frac{1}{1 + \alpha_n} \tag{3.16}$$

$$d_n = \frac{\alpha_n - 1}{\alpha_n + 1}. \tag{3.17}$$

The value α_n may be simply a constant value across all channels or it may vary across channels as a value proportional to the bandpass constant τ_n. It may be computed from the sampling rate f_s and the desired lowpass cutoff frequency $f_{3\mathrm{dB}}$ by the relationship:

$$\alpha_n = \frac{f_s}{\pi f_{3\mathrm{dB},n}}. \tag{3.18}$$

In cascade, the cumulative lowpass filter transfer function for M cascaded stages is

$$G_n(s) = \left(H_n(s)\right)^M. \tag{3.19}$$

Since this is not a Gaussian filter, usually two or three stages will suffice for the smoothing operation. In simulation, we used three stages. The result of applying rectification and smoothing in simulation can be seen in Figure 3.4.

3.5 Introduction to Translinear Circuits and Log-domain Filtering

Electrical engineering has a long history of attempts to create linear functions from highly nonlinear components. This may be due largely to the fact that engineering began with the (nearly) perfectly linear passive devices: resistors, capacitors, and inductors. The fundamental theories of linear filter design were fleshed out well before the advent of transistors or, for that matter, vacuum tubes. When these "active" devices were invented, designers capitalized on certain obvious nonlinear properties such as transistor switching (for digital circuits) and diode rectification (for AC to DC conversion). For most analog design, and almost exclusively for analog filter design, the goal was to force linear behavior from these devices through the use of high-gain amplification and feedback. The fundamental circuit element of analog integrated filter design became the operational amplifier, made of dozens or more transistors.

Linearization techniques work well, but the nonlinearities complicate circuit analysis considerably. One of the first techniques learned in amplifier and filter design involves the notion of

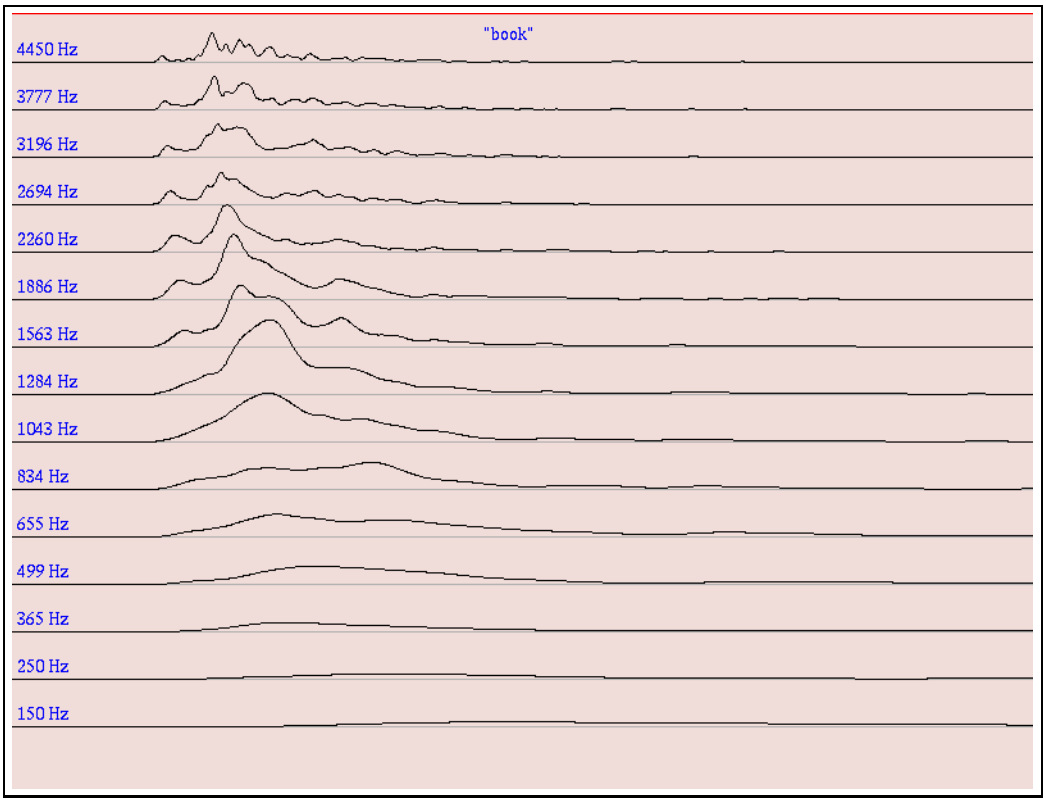

Figure 3.4: Parallel filterbank output after rectification and smoothing of the acoustic transient input signal across all frequency bands.

"small signal" analysis *vs*. "large signal" analysis. Nonlinear circuits act linear around a certain DC operating point, but can achieve only a limited accuracy over a given range. As the signal applied to the circuit departs from the operating point, the circuit departs from linear behavior. Circuit analysis proceeds in two stages by necessity: The nonlinear circuit equations can be used to find the quiescent operating point of the circuit given no input. Unfortunately, the nonlinear circuit equations, although they do describe the complete behavior of the system, are intractable for nonzero input functions. The linear circuit equations by nature are tractable but only approximate the behavior of the system near its operating point. Fortunately the principle of superposition allows these two sets of equations to be solved independently and added together to get symbolic solutions such as transfer function equations.

Log-domain filters are a different beast entirely. They comprise a subclass of circuits having externally linear transfer functions but internally nonlinear components (occasionally referred to by the acronym "ELIN") [40, 36]. This property, of course, could be claimed in reference to any circuit built with operational amplifiers, where every transistor used to build the amplifier is a nonlinear component. However, operational amplifiers are so firmly ingrained into engineering design that they are preferentially regarded as irreducible, fundamental devices without regard to their internal complexities. This view, which hides the internal complexity of the device behind a relatively few parameters like gain, unity-gain bandwidth, slew rate, and output impedance, is quite sensible. It may even be that one day certain subcircuits of log-domain filters, in particular current conveyors, will be regarded in the same way with the same universal scale of acceptability. For now, though, mostly due to the small number of transistors required for simple log-domain filters, they are scrutinized at the transistor level where the nonlinearities cannot be ignored.

As the name implies, log-domain filters are specifically those circuits whose internal state is a logarithmic function of the input and output. The circuit design exploits this particular nonlinearity directly rather than attempting linearization around an operating point. Such manipulations can only be done for specific types of nonlinearity: One is the square law, described by the drain current through a MOS transistor operating in strong inversion relative to its gate-to-source voltage [37]; another is the exponent, described by the collector current into a bipolar junction transistor relative to its base-to-emitter voltage, or the drain current through a MOS transistor operating in weak inversion (subthreshold) relative to its gate-to-source voltage. However, within the range of operation for which the nonlinearity of interest dominates all other second-order corrections to the device behavior, *large-signal linearity* is possible in circuits made from these devices. When a circuit has large-scale linearity, its transfer function uniquely describes the overall behavior of

the system. The equations governing the internal nonlinearity of the system are generally tractable, leading to complete solutions which do not require separate DC and transient analyses.

The foundation of log-domain filter theory was laid more than 20 years ago with the formulation of *translinear loop* circuits by Barry Gilbert [3, 35], which made use of the fact that the I-V equation of the bipolar transistor is an excellent approximation to an exponential over a large number of decades of current; coupled with Kirchoff's Current Law (KCL), it enables a number of mathematical operations to be performed by a small number of transistors. One important circuit based on the translinear principle is the Gilbert multiplier 2.8, whose operation is based on the fact that a multiplication of n numbers can be performed by taking the logarithm of the numbers, adding them all together, and then taking the antilog (exponential) of the result. Conceptually, this seems like a roundabout way to multiply numbers together. But because it relies on one easily-obtained linear operation (summation via KCL) and two easily-obtained nonlinear operations (log and antilog via the bipolar transistor I-V characteristic), the result is that two currents can be multiplied quite simply by this indirect method, while any attempt at a direct multiplication using linearization through feedback is extraordinarily complicated by comparison.

Substantial progress has been made in simplifying the processes of synthesis and analysis of log-domain filters [38, 39, 41, 42, 43, 44] by recognizing that they are by necessity composed of translinear loop circuits. The use of translinear circuit theory circumvents the necessity of explicitly solving transistor I-V equations, instead replacing them with mathematical functions on a more symbolic level, such as log, antilog, sum, and multiply. The translinear loop circuit analysis, as described by Gilbert, pertains only to static circuits, and only goes so far in describing time-dependent circuits such as filters. The crucial step in simplifying log-domain filter synthesis and analysis is to formulate a translinear principle which applies to time-derivative systems. This "dynamic translinear principle" was developed in Mulder *et al.* [38], and is worth paraphrasing here in the context of our method of circuit synthesis, where it allows us to synthesize filter circuits without resorting to state-space manipulation (*c.f.* [41]).

3.6 Principles of log-domain synthesis

A translinear loop is a circuit loop in exactly the same sense as Kirchoff's Voltage Law (KVL): it describes a circuit which originates at some point, passes through a number of circuit elements with associated voltage drops, and returns to the same point with a total sum of zero volts around the whole loop (*i.e.*, potential, and therefore energy, is conserved). In a translinear loop,

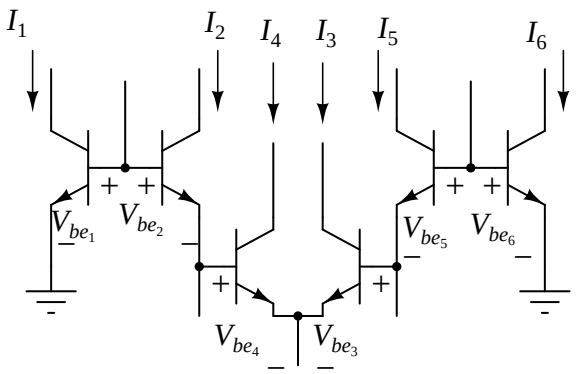

Figure 3.5: Translinear loop with common-base and common-emitter configurations.

however, the elements are nonlinear elements with exponential I-V relationships, namely diodes, bipolar transistors under normal biasing conditions, or MOS transistors in weak inversion. Each voltage drop around the loop is the voltage drop across the p-n junction of the diode or BJT or the voltage drop from the gate to the source in a MOSFET. By applying the I-V relationship of the nonlinear devices to the KVL equation, substituting the currents through the devices (which are not part of the KVL loop) for the voltage drop across the junction, it becomes immediately apparent that the exponential function allows very simple expressions for certain relationships between the currents.

Figure 3.5 shows a simple example of a translinear loop formed by the base-emitter junctions on bipolar transistors. For purposes of analysis and synthesis, we will consider all transistors (whether BJT or, equivalently, MOSFET operating in weak inversion[1]) to be ideal, perfectly matched devices implementing the simplified exponential function $I_c = I_s\, e^{V_{be}/V_t}$, or inversely, $V_{be} = V_t\, \ln\left(I_c/I_s\right)$, where V_{be} is the BJT base-to-emitter voltage, or the gate-to-source voltage of a corresponding subthreshold MOS transistor, V_t is the thermal voltage (0.025 V at room temperature) and I_s is the reverse saturation current of the transistor. Circuit modifications necessary to deal with the nonidealities of real devices will be addressed in Section 3.9. By Kirchoff's voltage law, the sum of the V_{be} voltages around the loop is zero:

$$\sum_{i\,\mathrm{odd}} V_{be_i} = \sum_{i\,\mathrm{even}} V_{be_i}. \tag{3.20}$$

[1]I do not intend to suggest that the physical mechanisms by which BJTs and MOSFETs in weak inversion operate are equivalent, which they are not [34], only that the I-V characteristic of both can be closely approximated by an exponential function which allows either transistor type to become the fundamental translinear circuit component for realizing equivalent (to first order) linear systems.

Substituting the idealized transistor equation results in an expression for the currents in which the logarithm function can be eliminated due to the property of logarithms relating the sum of logs to the log of products. This is the "static translinear principle":

$$\sum_{i\,\mathrm{odd}} V_t \, \ln\left(I_i/I_s\right) = \sum_{i\,\mathrm{even}} V_t \, \ln\left(I_i/I_s\right); \tag{3.21}$$

$$\prod_{i\,\mathrm{odd}} I_i = \prod_{i\,\mathrm{even}} I_i. \tag{3.22}$$

As Equation (3.22) shows, the static translinear principle gives a simple rule for computing products of currents, revealing a minimalist way of multiplying the value of two currents (always scaled by another current, so that the result is a current, not an unmeasurable or impractically-measurable unit like amperes squared). To design filters, or, for that matter, any dynamic system, we need an equivalent, simple rule to generate equations involving time-derivatives of currents, *i.e.*,

$$I_{out} \propto \dot{I}_{in}. \tag{3.23}$$

Fortunately, such a simple rule exists, one which again hinges on the use of an arithmetic property of logarithms/exponentials to simplify resulting expressions. The key to the problem is the property of the derivative of an exponential function. Applying the time derivative to the simplified transistor equation results in the "dynamic translinear principle":

$$I_{out} = I_s \, e^{(V_{be})/V_t} \tag{3.24}$$

$$\dot{I}_{out} = \frac{d}{dt}\left(I_s \, e^{(V_b - V_e)/V_t}\right) \tag{3.25}$$

$$= \frac{1}{V_t} I_{out} \frac{d}{dt}\left(V_b - V_e\right). \tag{3.26}$$

Grounding the BJT emitter ($V_e = 0$) and adding a capacitor C to the system at node V_b ($I_C = C\dot{V}_b$) yields an equation composed entirely of current-mode variables. Figure 3.6 shows such a system, a basic building-block of log-domain filters (*c.f.* [39], Fig. 1, and [38], Fig. 1). It should be noted at this time that a constant voltage V_{shift} (as shown in Figure 3.6) may be inserted between the capacitor node Z and the transistor base. This constant is canceled on both sides of the equation and therefore does not affect the solution:

$$\dot{I}_{out} = \frac{1}{V_t} I_{out} \frac{I_C}{C}. \tag{3.27}$$

From this equation we can compute the derivative of a current by multiplying two currents (I_{out} and I_C) together. As mentioned earlier, multiplication of two currents can be easily accomplished with a translinear loop circuit.

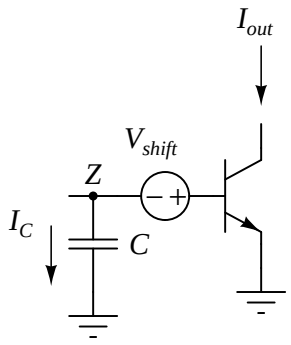

Figure 3.6: Filter pole formed using a transconductor.

Note that while neither equations (3.22) and (3.27) describe any circuit behavior that cannot be derived by directly solving the (simplified) I-V equations, they both can be used as methods to quickly analyze translinear loop circuits without explicitly referring to either exponent and logarithm functions or to device parameters (*e.g.*, I_s), and also to synthesize filters from building blocks based on simple translinear loops.

3.7 First-Order Circuit synthesis

To show how these circuit concepts can be used in practice to synthesize complete filter circuits, consider for instance a generic first-order system with the (current) transfer function:

$$\frac{I_{out}(s)}{I_{in}(s)} = \frac{1}{A + \tau s} \tag{3.28}$$

Equation (3.28) is not in a canonical form, but is presented here as a useful building block for generating higher-order circuits, as we will show in Section 3.8. Working in the time domain and substituting, from Equation (3.27), $\dot{I}_{out}$ for sI_{out}, gives:

$$AI_{out} + \left(\frac{\tau\, I_C}{V_t\, C}\right) I_{out} = I_{in}, \tag{3.29}$$

$$I_{out}\left(A + \frac{\tau\, I_C}{V_t\, C}\right) = I_{in}. \tag{3.30}$$

We then can define the time constant τ in terms of some (constant) bias current I_b:

$$\tau = \frac{V_t\, C}{I_b}. \tag{3.31}$$

This may seem like an arbitrary step, but the units of a time constant are of course equal to RC, and $R = V/I$, and V_t is the thermal voltage with units of volts ($kT/q \approx 0.025\,\mathrm{V}$), so this expression

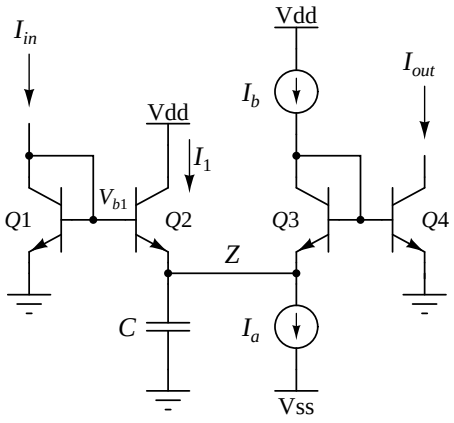

Figure 3.7: First-order log-domain filter circuit.

does indeed have the correct units of time. Substituting τ into Equation (3.30) and multiplying through by I_b, we get

$$I_{out}\,(AI_b + I_C) = I_{in} I_b. \tag{3.32}$$

Equation (3.32) has a familiar form, that of the four-component translinear loop, Equation (3.22), with the constraint that I_C and I_{out} must have the relationship shown in Figure 3.6. One of several possible implementations is the "up-down" (common-base) configuration shown in Figure 3.7, in which transistors $Q1$ through $Q4$ form the translinear loop $I_{out}\, I_1 = I_{in}\, I_b$. $Q3$ is a voltage level-shifter (*c.f.* [39], Fig. 2), and Equation (3.28) is satisfied if $I_a = I_b\,(1 + A)$. Thus, the value A is tunable and controlled by the ratio of I_a to I_b, both of which are unconstrained constant bias voltages. Useful values of A are $A = 1$ at $I_a = 2I_b$ (a lowpass filter) and $A = 0$ at $I_a = I_b$ (an integrator, useful in systems with feedback such as second-order filters).

The negative power supply, $\mathrm{V_{ss}}$, is a somewhat arbitrary value which needs to be kept far enough below ground that any transistors used as current sources which draw current off of the capacitor node Z can operate correctly. For a simple, single nMOS transistor used as the current source I_b, $\mathrm{V_{ss}}$ should be about $-0.5\,V$.

Translinear loop equations break down as voltages across the translinear element approach zero because the diode equation, which holds (with different premultipliers) for all three types of translinear elements, differs from a true exponential by a single "-1", which is the manifestation of the impossibility of having currents move against the direction of the applied voltage across the junction:

$$I_c = I_s\,(\exp\left((V_b - V_e)\,/V_t\right) - 1)\,. \tag{3.33}$$

As an approximation to a true exponential, this equation breaks down at currents in the range of I_s, which for most applications can be considered equal to zero compared to the nominal operating current through the device. Translinear devices cannot behave properly at all for negative input currents. An exception is the use of carefully constructed *class A-B* log-domain circuits, in which the input is centered around zero current, and two symmetric log-domain circuits handle the signal alternately in the positive and negative regions [40]. Class A-B circuits are particularly interesting for reasons of their ability to act as instantaneous companding circuits to reduce internally-generated noise as seen at the filter output. However, they are beyond the scope of this thesis, which is concerned mainly with efficient implementations and only covers single-ended, class A filters. Correct operation of a class A first-order section requires that the entire circuit be biased by adding a positive DC current to the input large enough to keep the input positive and therefore keep the transistor base-emitter (or gate-source) junctions forward biased at all times.

One noticeable consequence of class A operation is that while noise on the signal generated at or before the filter input is logarithmically compressed and expanded along with the signal, internally generated noise only experiences the expansion. Thus a portion of the noise seen at the filter output increases with the instantaneous signal level. Fortunately, there are few sources of noise internal to the log-domain circuit, and according to the equipartition theorem [6], they are all referred to the capacitor node. Thus the magnitude of the exponentially-rectified portion of the noise at the output can be calculated, and the capacitor size can be increased to make the value as small as required by the filter specifications.

3.8 Designing second-order sections

Construction of higher-order log-domain circuits from simple component cells has been described in various ways, such as the Bernoulli cell of [41] and the E+ and E− cells of [44]. We generate higher-order functions in a similar way, factoring the desired current transfer function into equations which can be directly implemented by the first-order subcircuit of Figure 3.7. Consider, for example, a second-order bandpass equation (Frey [39] uses a more general form of the same function for state-space synthesis of a second-order section):

$$I_{out}(s) = I_{in}(s)\frac{\tau s}{1 + (1/Q)\tau s + \tau^2 s^2} + I_{DC}(s). \tag{3.34}$$

This describes a bandpass function with a gain and resonance of Q, except that the output is biased to a positive DC operating point. Since the bandpass function itself eliminates DC components

from the input, the output bias must be provided separately, denoted here as a constant current I_{DC}, independent of the input signal. The class A filter circuit requires only that the input be strictly positive. The output can potentially be bidirectional. However, there are several reasons for biasing the output to make it strictly positive: The filter section is then in a structural form simple to cascade, and the circuit implementations are less prone to output errors caused by device mismatch.

An elegant way to implement the above equation is to multiply the DC current term by the denominator of the bandpass function, $1/1 + (1/Q)\tau s + \tau^2 s^2$, which by itself is a second-order lowpass filter function. A lowpass operation on the DC term retains the DC term unaffected (except for possible offsets as a result of nonidealities in the actual circuit). On the other hand, the additional term makes the equation much easier to factor.

$$\begin{aligned} I_{out}(s) &= I_{in}(s)\left(\frac{\tau s}{1+(1/Q)\tau s+\tau^2 s^2}\right) \\ &+ I_{DC}(s)\left(\frac{1}{1+(1/Q)\tau s+\tau^2 s^2}\right). \end{aligned} \tag{3.35}$$

Equation (3.35) can be easily factored into two simple first-order current transfer functions by introducing an intermediate current term I_x:

$$\frac{I_{out}(s)}{I_{in}(s)-I_x(s)} = \frac{1}{1/Q+\tau s}; \tag{3.36}$$

$$\frac{I_x(s)}{I_{out}(s)-I_{DC}(s)} = \frac{1}{\tau s}. \tag{3.37}$$

The right-hand side of (3.36) and (3.37) takes the form of the first-order function (3.28): for the first section, a lossy integrator, set $I_a = I_b(1 + 1/Q)$, and for the second section, a lossless integrator, set $I_a = I_b$. Replacing τ by the expression of Equation (3.31) shows that the bandpass filter's center frequency is

$$f_c = 2\pi\omega = \frac{1}{\tau} = \frac{I_b}{2\pi V_t C}. \tag{3.38}$$

The input to each section has a negative feedback term from the other section. As presented in the equations above, and shown in Figure 3.8A, the positive and negative feedback terms combine to generate an input which can be a positive or negative value. But this presents a problem: it violates the class A log-domain circuit structure in which input currents must be strictly positive, driving the transistors out of the active mode and invalidating the translinear model. In this case of carefully orchestrated feedback which is the bandpass filter, the inputs to the log-domain sections cannot simply be biased upward: The lossless integrator, in particular, would not work at all, since it would continuously integrate the positive input until its output saturated. There is a solution,

however: The negative feedback term can be separated from the positive term, and with the proper transformation, it can be referred to the capacitor node as a dependent current source. This solution is shown in Figure 3.8B. To find the equation for this dependent current source, the translinear loop

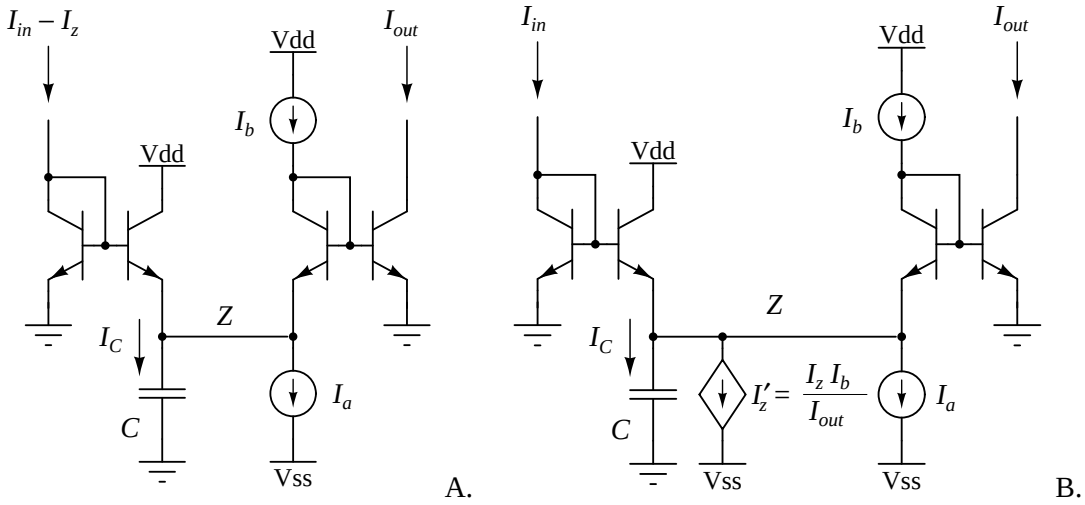

Figure 3.8: Computing a current difference at a log-domain filter input. A) The underlying idea, which is physically unrealizable. B) An equivalent working implementation.

equations for circuits A and B can be written, respectively:

$$(I_{in} - I_z)I_b = (I_a + I_C - I_b)I_{out}; \tag{3.39}$$

$$I_{in}I_b = (I_z' + I_a + I_C - I_b)I_{out}. \tag{3.40}$$

Solving for I_z', we find that the dependent current source must have the value $I_z' = (I_z\, I_b)/I_{out}$. It should be immediately apparent that this familiar form can be conveniently implemented by a simple four-transistor translinear loop circuit.

Now the bandpass filter can be drawn as a cascade of two first-order circuits, with the output of the first section shared with the input of the second, a method which is described briefly in [39]. The connected sections are shown in Figure 3.9.

As a further simplification, by noting the translinear loop $Q2$-$Q3$-$Q4$-$Q5$ which gives the equation $I_x I_y = I_b I_{out}$, the dependent current sources I'_{DC} and I'_x can be written in terms of I_y rather than I_x, and the entire subcircuit which generates I_x (dotted box in Figure 3.9) can be eliminated, avoiding the need for one set of matched current sources.

It remains to generate the translinear loops implementing the dependent current sources. Generally, two different solutions are possible for any simple translinear loop of this type: A

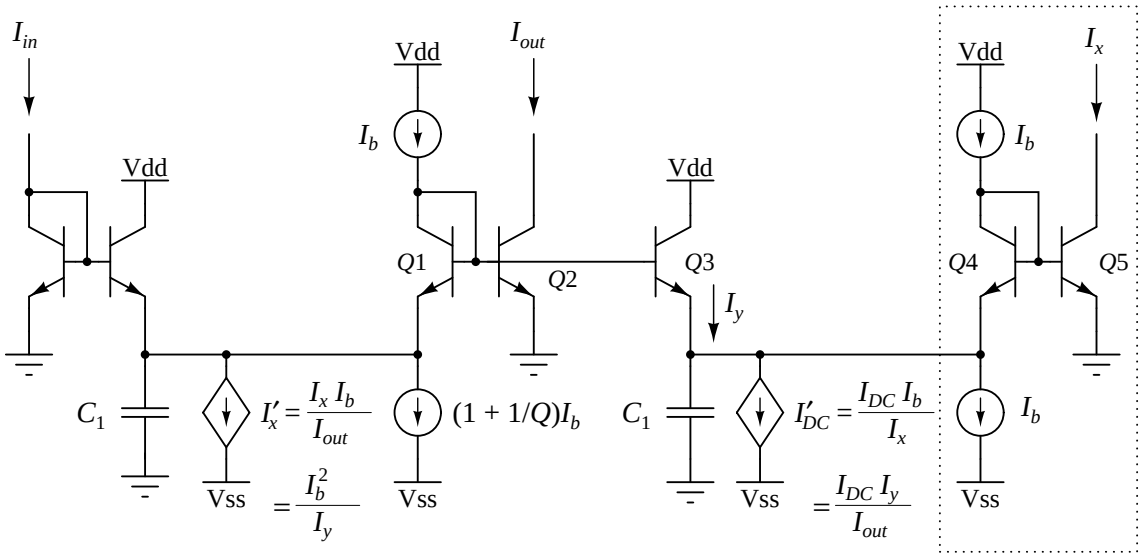

Figure 3.9: Bandpass structure formed from first-order sections.

common-base configuration, and a common-emitter configuration. Thus far we have used common-base configurations exclusively to design first-order log-domain sections. Naturally, connecting transistors together at the emitter works just as well to create KVL loops, with the preferred structure depending on the circumstance. Both structures are illustrated in the bandpass circuit schematic of Figure 3.11: In the first solution, $I'_x = I_b^2/I_y$ is generated by the internal translinear loop $Q3$-$Q7$-$Q6$-$Q1$ using the common-emitter configuration of transistors $Q6$ and $Q7$. This configuration has been used by Frey [39] and others. The other dependent current source, $I'_{DC} = I_{DC}I_y/I_{out}$, is implemented through the translinear loop $Q8$-$Q9$-$Q3$-$Q2$ using a common-base instead of a common-emitter configuration, in which the sign of the current is reversed, as required in Figure 3.9, by mirroring, doubling, and subtracting from the same node. This configuration has been described previously by Fox [44]. Notice that both translinear loops implement the same function with different circuits.

The author of each of the works mentioned above used the same circuit configuration (common-base or common-emitter) twice. It is possible to mix both types of structures to take advantage of the strengths of each in the appropriate context. This is the approach we used for our bandpass filter after encountering some stability problems with our first prototype, which used two common-base structures. On the left-hand side of the filter, a common-base configuration is generally less desirable due to the required extra pair of matched I_b current sources. It is critically

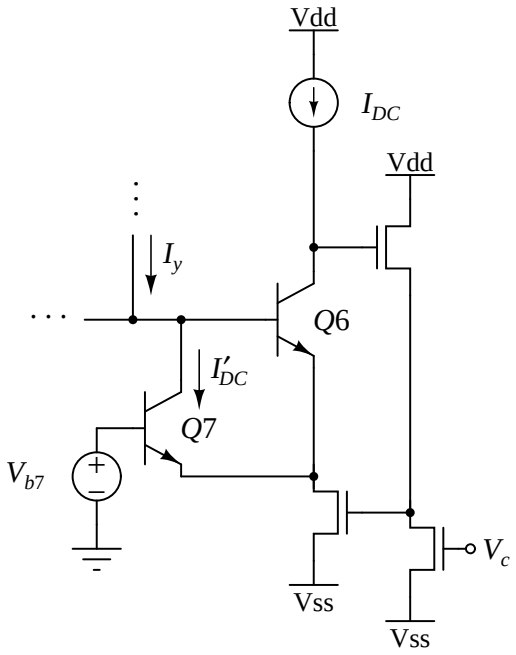

Figure 3.10: An alternative common-emitter circuit generating I'_{DC} (see text for discussion).

important to reduce the number of matched current sources in the filter circuit, which is the primary source of gain and offset errors and inter-channel mismatch in a filterbank. The common-base circuit also incurs a signal delay around the current mirror loop, and has been observed to latch up when the MOSFETs in the mirrors are made too large, increasing the signal delay to the point at which the circuit becomes unstable. A full analysis of this circuit's stability has not been attempted. On the right-hand side of the filter, use of a common-emitter circuit to replace the $Q8 - Q9$ common-base pair (Figure 3.10, with V_{b7} = 0 V) results in current draw through the base of $Q6$, turning the lossless integrator into a lossy integrator. At audio frequencies, where the bias current I_b can be smaller than the base current I_{DC}/β of 10 to 100 nA, the circuit will fail. Note, however, that I'_{DC} is proportional to $I_{DC} \cdot e^{V_{b7}/V_t}$. I_{DC} can be reduced to the order of magnitude of I_b if V_{b7} is raised to compensate. This will allow the common-emitter circuit to work, though with the drawback that the output bias level is no longer equal to I_{DC}, but is a derived function of I_{DC} and V_{b7}. The derived function can in turn be generated from the real values of I_{DC} and V_{b7} using yet more translinear circuits 3.12. The resulting circuit has a minimum number of matched currents throughout, and therefore potentially has the best behavior in terms of mismatch. We took this approach in our latest prototype of the bandpass filterbank. Unfortunately, this approach merely shifted the mismatch problem from one

place to another in the circuit, resulting in improved matching of Q but not in gains.

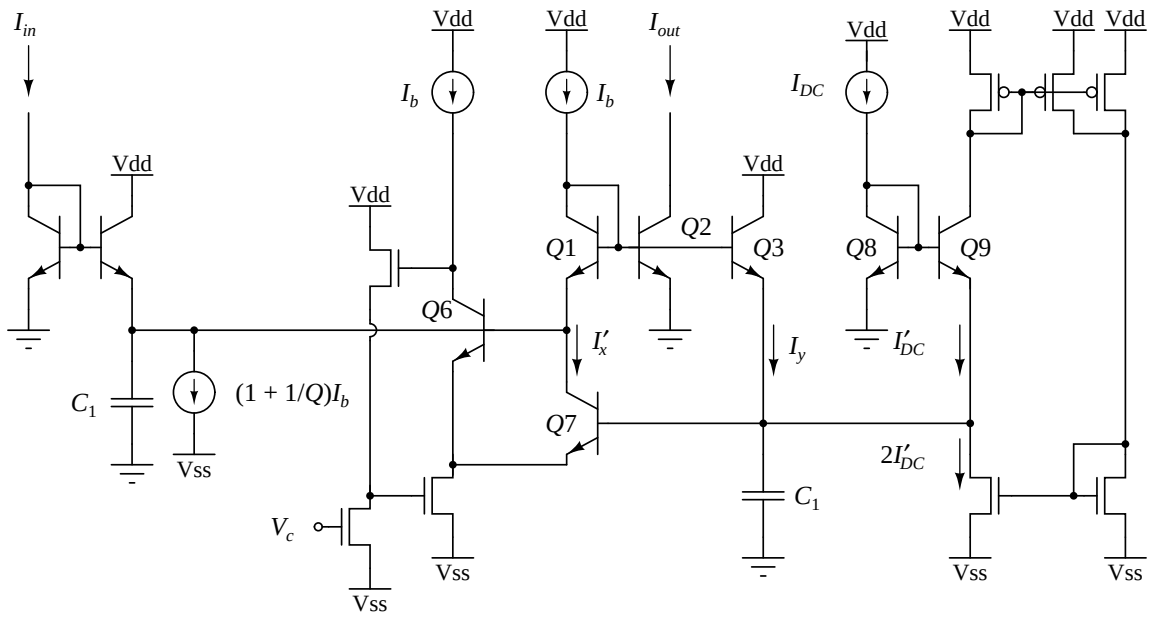

Figure 3.11: Complete circuit schematic for the second-order bandpass filter.

3.9 Technology limitations for low-frequency filter design

Much previously published work on log-domain filters has extolled their potential for replacing conventional filters in high-frequency filter design. Our work instead concentrates on the use of log-domain filters for audio-frequency applications [44, 45, 46], the use of current-mode filtering for current-mode applications, and the application of log-domain filters in system design [57]. The low frequency range of audio requires a large RC time constant, which for the log-domain filters described here is inversely proportional to the bias current I_b. For large-scale integrated systems where capacitors cannot reasonably be made larger than a few picofarads, the bias current can be as low as several tens of picoamps on the low end of the audio frequency range, which places some important restrictions on circuit technology. Problems arising from established designs are:

1. Traditional bipolar designs fail due to base current draw and can suffer from β mismatch at low emitter current values.

2. MOSFET designs have poor matching of currents in mirrors, and require that the input bias

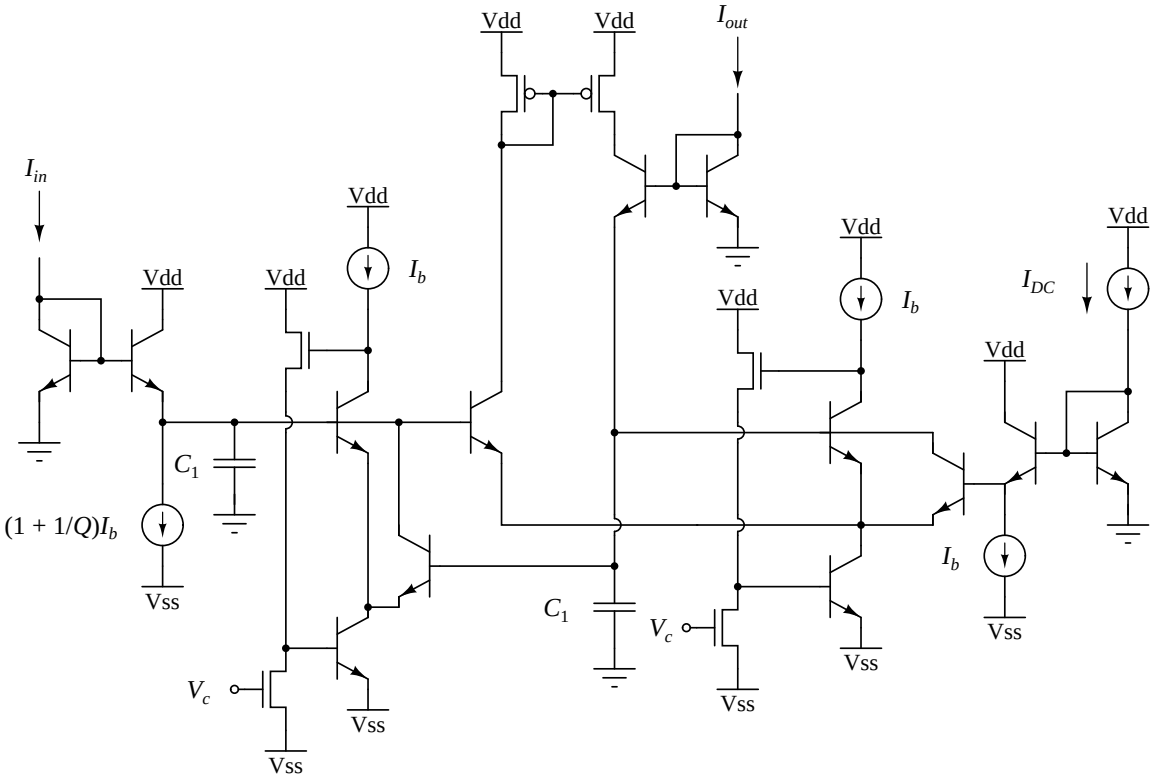

Figure 3.12: Another BiCMOS Log-domain bandpass filter, using the common-emitter configuration for all feedback circuits.

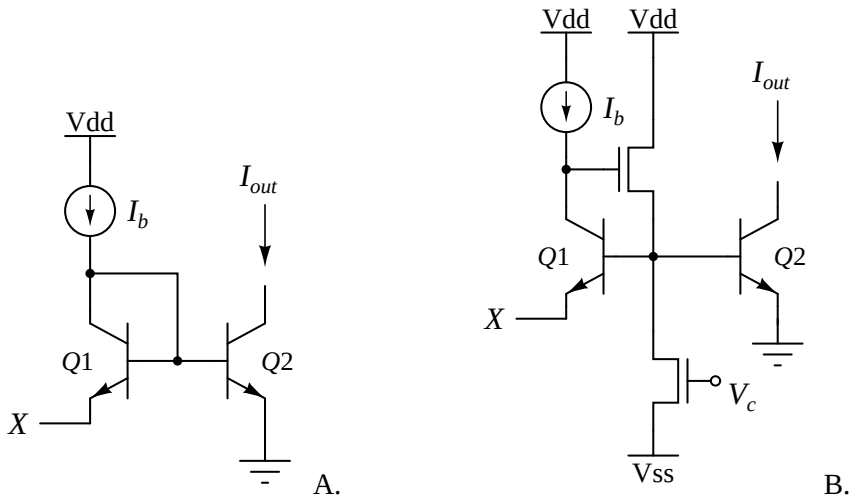

Figure 3.13: Base compensation (B) eliminates undesirable behavior due to significant base current which occurs in (A).

current be very low to keep the MOSFET transistors in weak inversion. The input and output are then close to the noise floor and it becomes difficult to retrieve a clean output signal.

The first of these problems deserves some explanation: The failure of the circuit stems from the design assumption that base current is negligible in translinear-loop circuits. This assumption becomes invalid in situations where the transistor $Q1$ in Figure 3.13A attempts to drive $Q2$ when the collector current through $Q2$ is many orders of magnitude higher than that through $Q1$, due to the emitter voltage of $Q2$ being significantly lower than that of $Q1$. The result is that $Q2$ requires more current through its base than the current source above $Q1$ can provide. This situation occurs in a number of critical places in the bandpass filter circuit; for instance, at the output, where the (DC current biased) output is many orders of magnitude higher than bias current I_b which is driving its base.

We have fabricated and tested both MOSFET and BiCMOS designs. Results from the MOSFET versions indicate that the large noise floor is pervasive and difficult to impossible to eliminate by design. On the other hand, there exist a number of established circuit techniques collectively known as *base-current compensation* for dealing with unwanted base current draw, making the first problem more likely to be overcome. We devised a technique by which base current loss can be eliminated using base-current compensation on all diode-connected transistors in the translinear loops as shown in Figure 3.13. Voltage V_c is a constant bias which generates a bias current that reduces the impedence of the base node. In the second-order filter, the V_c bias on the

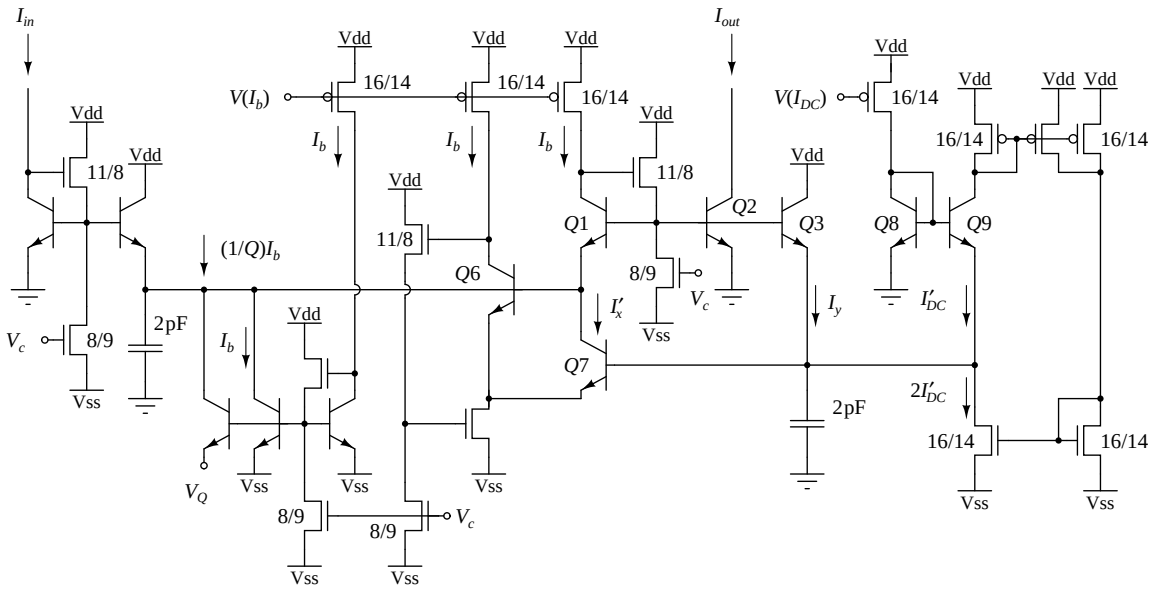

Figure 3.14: Complete circuit schematic for the second-order bandpass filter. Bipolar transistors are minimum size, and MOS dimensions are indicated as W/L in units of λ = 0.6 µm.

base-current compensation circuit can be shared with the V_c bias needed for the common-emitter translinear circuit of Figure 3.10.

Figure 3.14 shows the complete bandpass filter, including all base compensation circuits and current sources, as it was fabricated and from which the measurements in Section 3.16 were taken. Note the circuit implementation of the current $(1 + 1/Q)I_b$: in this configuration, $Q = e^{(V_Q - \mathrm{V_{ss}})/V_t}$, making tuning of Q independent of I_b, the center frequency tuning control, and giving Q a natural range from 1 to the practical maximum allowed by the circuit.

Figure 3.15 shows the alternative bandpass circuit from Figure 3.12 with all the circuits used to compensate for nonidealities. The use of single MOSFET transistors as the current sources requires that cascode connections be added to compensate for the differing source voltages on otherwise matched transistor pairs.

3.10 Layout Considerations for VLSI Log-Domain Circuits

The log-domain filter circuits are so sensitive to mismatch error that it was deemed necessary to invent a regular layout structure to use throughout the chip. Conventional wisdom is that the most obvious measures to take against potentially bad matching of paired transistors are: use the same layout length and width for each device, keep the devices facing the same direction, keep the

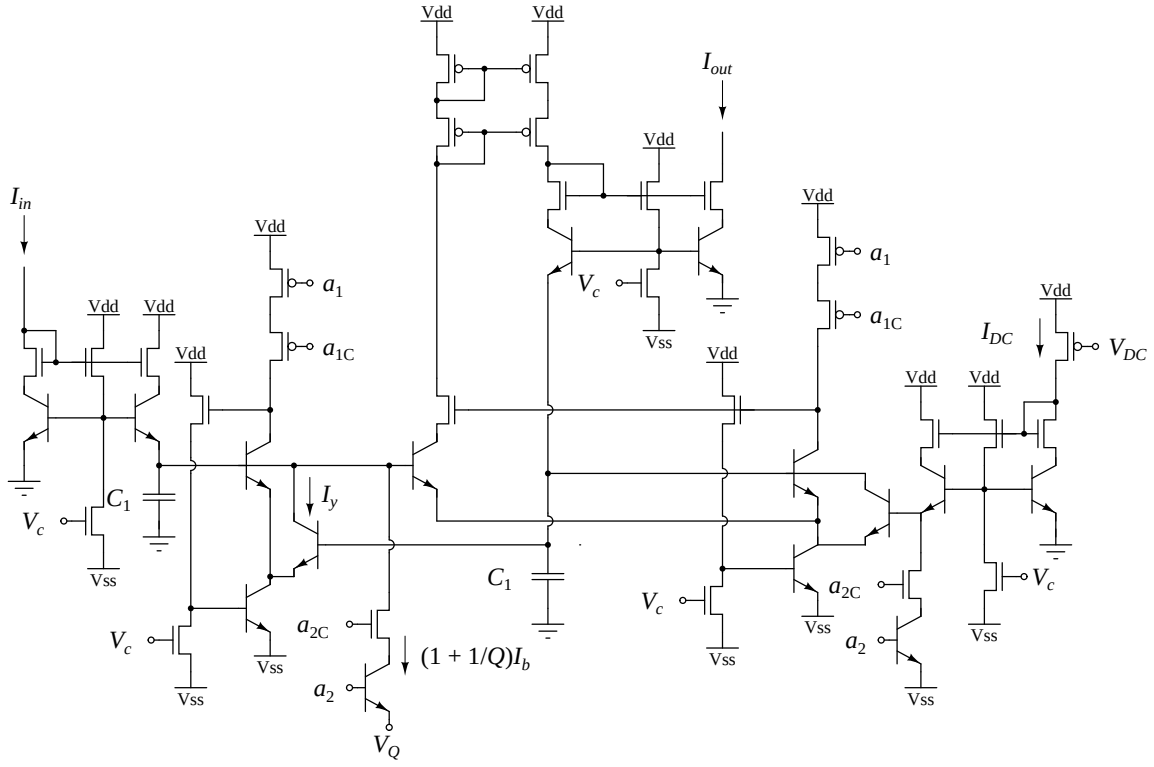

Figure 3.15: Circuit of Figure 3.12, showing base compensation circuits, cascode connections, and the Q-generating circuit

.

devices close together, and use quadrature symmetry when practical. After considerable experience in chip design and testing, I have concluded that the greatest influence on transistor matching, after ensuring the same length and width of each device, is the effect of "same surround." This appears to be more critical even than having transistors facing the same direction (at least this is true for anti-symmetry; I have not looked at 90-degree rotations as the opportunity for this kind of layout while maintaining same surround is rare) and having transistors close together. Indeed, transistors which are displaced by half the length of a chip (say, 1 mm) but which are surrounded by exactly the same layout to a radius of about 50 to 100 μm in all directions tend to be much better matched than two transistors placed next to each other but with different circuits to the right and left of the pair.

Generally speaking, one would expect that the more regular and symmetric the layout, the better the matching between transistors throughout. The parallel nature of the filterbank aids considerably in the matching between channels. Edge effects which might cause increased mismatch in the highest- and lowest-frequency channels can be nearly eliminated by placing the elements most sensitive to mismatch as far away from the edge as possible. Double polysilicon capacitors are the least sensitive to mismatch caused by differences in the surrounding layout. Also, the capacitors are large compared to the size of transistors and most of them have one side grounded. The width of the capacitors usually satisfies the distance of 50 to 100 μm over which circuit differences can affect transistor matching. If the capacitors face outward to the edges of the chip, MOS transistors placed behind the capacitors will not suffer any mismatch due to edge effects. The structure of this design is shown in Figure 3.16.

The edge-effect phenomenon was noticed between two versions of the frontend filterbank. One version contained fifteen channels and the other, sixteen. Each filterbank was designed with pairs of channels stacked antisymmetrically, facing each other, as shown in Figure 3.16. The fifteen-channel filterbank, however, had one unpaired channel on the high-frequency end. The lack of pairing was the only factor influencing mismatch that was not present in the sixteen-channel filter. In the sixteen-channel filter, all measured center frequencies were within 2% of their nominal values. In the fifteen-channel filter, all measured center frequencies were within 2% of their nominal values *except for the unpaired channel*, whose center frequency was over 30% too high! By contrast, no systematic difference occurred between even and odd channels, revealing that antisymmetry has no apparent effect on transistor mismatch.

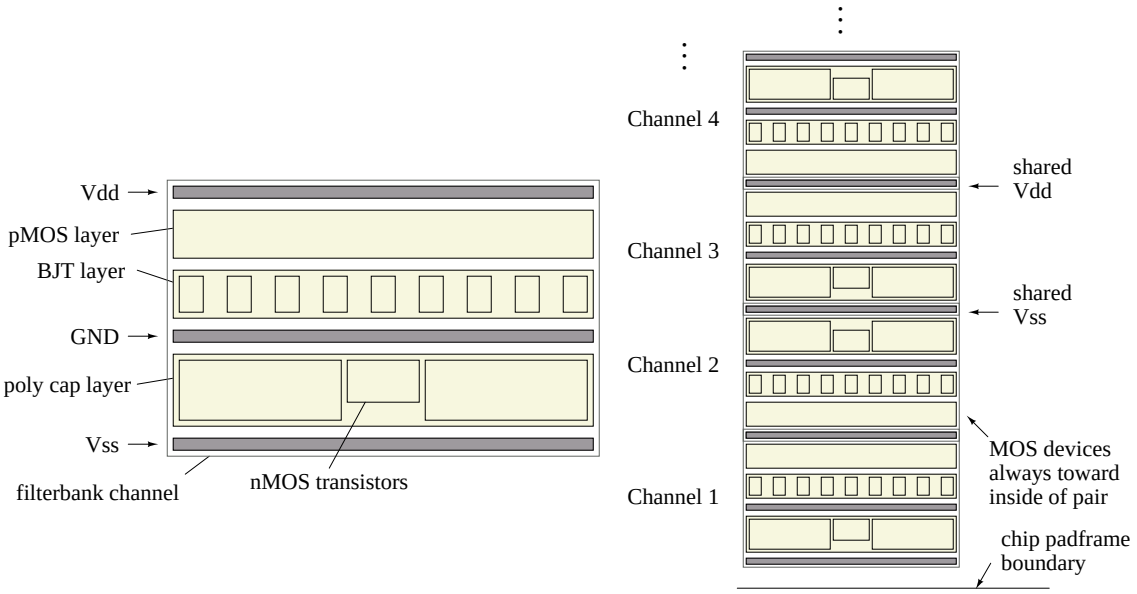

Figure 3.16: Structure of a bandpass filterbank channel and stacking of channels to form the whole filterbank.

3.11 Current-Mode Circuits for Non-Filtering Applications

So far in this chapter, we have developed log-domain theory for the purpose of analyzing and synthesizing filters. Much of the novel aspect of the research work involved in developing the current-mode filterbank, however, was concerned with the other critical circuits necessary to extend the simple parallel filterbank into a useful signal processing frontend.

In keeping with the desire to reduce transistor mismatch through the use of regular and symmetric layout, every attempt was made to maintain that structure not only for the filter, but for all the peripheral circuits as well.

3.12 Signal Rectification and Smoothing

The application for which the frontend filterbank was designed called 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 signal was smoothed (integrated) with a time constant of 1 to 2 ms.

We considered two ways to compute the short-term energy envelope: The first way was to implement the method of the simulation exactly; that is, to perform a half-wave or full-wave

rectification followed by a log-domain lowpass filter for smoothing. However, we were aware from prior experience that rectification requires that the midpoint of the signal be known, and offsets introduced by mismatch and circuit nonidealities can severely complicate the task. So in addition to developing circuits for the direct implementation of the simulation method, we also invented novel circuits which would determine the true peak-to-peak height of each bandpass output, avoiding the necessity of knowing the signal midpoint.

3.13 Signal Rectifier

Figure 3.17 shows a simplified version of the circuit used for the current full-wave signal rectifier. This circuit is adaptive on a fairly long time scale (about 1/10 s to 1 s would be typical), using a *charge pump* for the adaptation. The charge pump (including the inverter, transistors Q_2 and Q_3, and capacitor C_1) acts to keep the current I_1 through transistor Q_1 equal to the input current I_{in} using feedback from the operational amplifier: If the input current exceeds the current through Q_1, then V_{fb} drops, passing V_{ref}, as the system driving the input is driven towards cutoff. Consequently, the op-amp output rises, causing the inverter output to be driven 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 Q_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 Q_2 and Q_3, respectively, determine the rate at which charge can be placed on or removed from C_1. Thus, the 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 architecture of the system ensures that the current difference between I_1 and I_{in} always has somewhere to go, 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, Q_5 is cut off and current is drawn through Q_4, mirrored, and produces a positive I_{out}. When I_{in} falls below I_1, Q_4 is shut off (and the mirror Q_6, Q_7 with it), and positive current again flows to the output, this time directly through Q_5. Apart from mismatch errors, the output current always equals $|I_{in} - I_1|$.

The main differences between the simplified schematic and the one used on the fabricated IC are a source-degenerated ("diode-connected") pMOS transistor between Q_1 and Vdd to raise

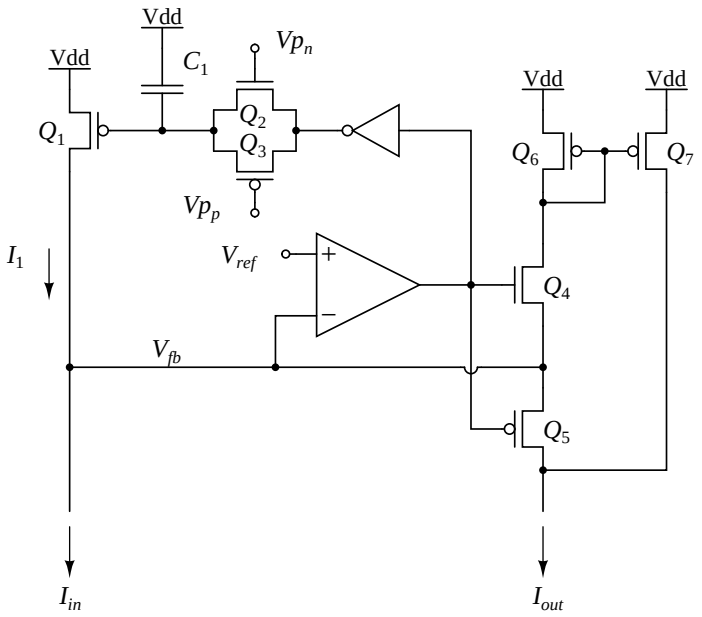

Figure 3.17: Simplified schematic of the adaptive current full-wave rectifier.

the required voltage across C_1, the replacement of the current mirror formed by Q_6 and Q_7 by a current-conveyor mirror of the type shown in Figure 3.13 (B), and 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 log-domain lowpass filter, a structure which has been developed and is known to work well (Figure 3.24). The first-order lowpass is, in fact, the most robust of all the log-domain filter circuits. It can be easily cascaded to form an n-th order filter, as shown in Figure 3.18. Because the translinear circuits operate down to practically zero input current, the log-domain lowpass filter can operate directly on the output of the signal rectifier without requiring any further biasing of the input. The smoothed output can reach arbitrarily low current values during quiet periods of the system input.

3.14 Signal Peak-Peak Detector

We invented and developed a novel analog circuit for determining the height of a signal without requiring any knowledge of the mean value of the input waveform. The operation of the circuit is based on the 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

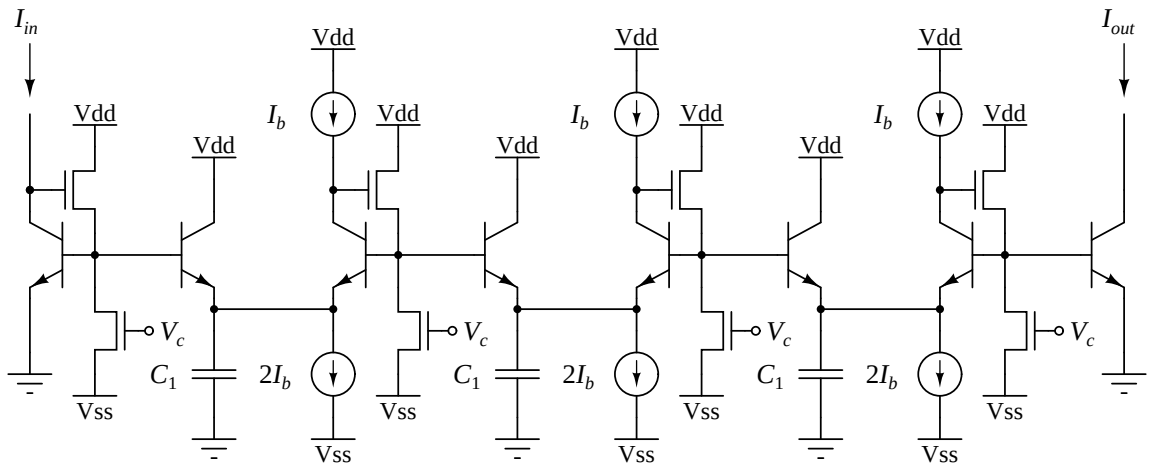

Figure 3.18: Cascaded log-domain lowpass filters.

small current source to cause the capacitor to leak slowly to ground so that the circuit will track the input rather than remain stuck at the largest input voltage. Putting the diode in the feedback loop of an amplifier causes the circuit to be unaffected by the voltage drop across the diode, so that the output very closely follows the rising edge of the input when the diode is not reverse biased. This circuit is shown in Figure 3.19.

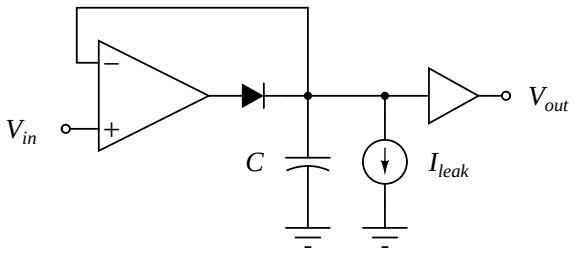

Figure 3.19: Simple diode-based peak detector.

One novel aspect of our circuit is that it operates *differentially* by containing two sub-circuits, one which computes the maximum peak, and one which computes the minimum peak (technically speaking, the "trough"). The maximum peak-detecting circuit leaks toward ground, stopping when its value equals that of the input. The minimum peak-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 peak-detecting circuit output from the maximum peak-detecting circuit output, which gives a reasonably accurate measurement of the input signal's instantaneous

peak-to-peak amplitude (see Figure 3.20). Due to the leaky capacitor architecture, the circuit responds quickly to sharp rising edges in the input but relaxes slowly to zero on the falling edge of the input. This should be a good model for establishing the energy envelope of acoustic transients, which contain much of the detailed information required for classification in the shape and frequency distribution of the rising edge.

Another novel aspect of our circuit is that it operates in the current domain, in keeping with the mode of the log-domain filters preceding it and the normalizer and correlation processor following it. The use of currents allows us to implement the diode rectifier very simply with a current conveyor structure. These structures are shown in Figure 3.21.

This 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 3.19, 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. We have minimized the effect by raising the DC bias of I_{in}, so that the signal amplitude is considerably smaller than its bias. The differential current change of the maximum and minimum currents is small, and so traverses a shorter range of the nonlinear expansion function, rendering the output nearly linear as a function of 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

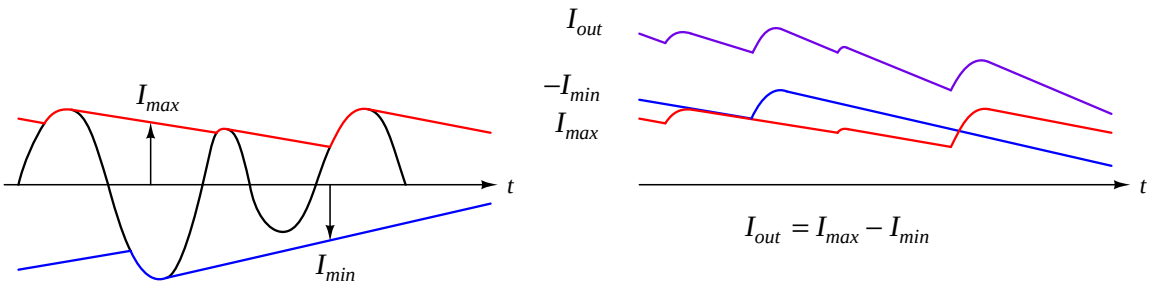

Figure 3.20: Behavior of the peak-peak detector.

rate of charge leakage equal in the two subcircuits, the same size capacitor and pMOS transistor gated by the same bias voltage was used for both. 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 3.22 is the final version of the circuit, as implemented on the chip. Extra circuitry has been added to minimize any offset 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. A significant part of the extra circuitry (Q_4, M_{15}, M_{16}, M_{17}) exists 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, just as was done for the signal rectification and smoothing described in the previous section, to smooth out the bumps in the peak detector output. This presents a much smoother signal to the system output at the expense of only a very small amount of circuitry.

3.15 L-1 Normalization Array

The simulation model calls for a final transformation of the energy envelope: a normalization across all channels. The normalization is the "L-1 norm," which follows the following equation (written in units relevant to the current-mode system):

$$I_{out}(i) = I_{norm} \frac{I_{in}(i)}{\sum_{j=1}^{N} I_{in}(j)} \tag{3.41}$$

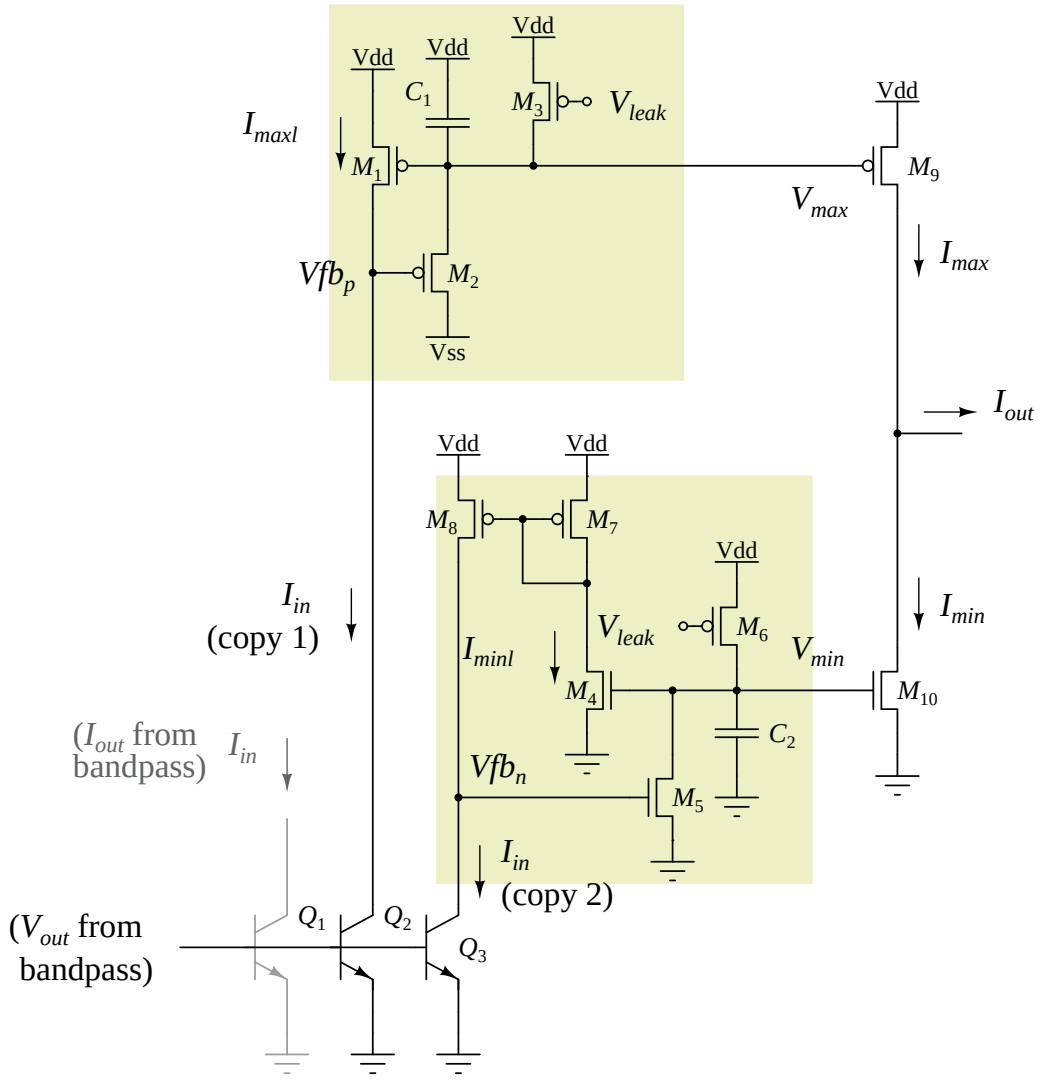

Figure 3.21: Peak-peak detector circuit, simplified.

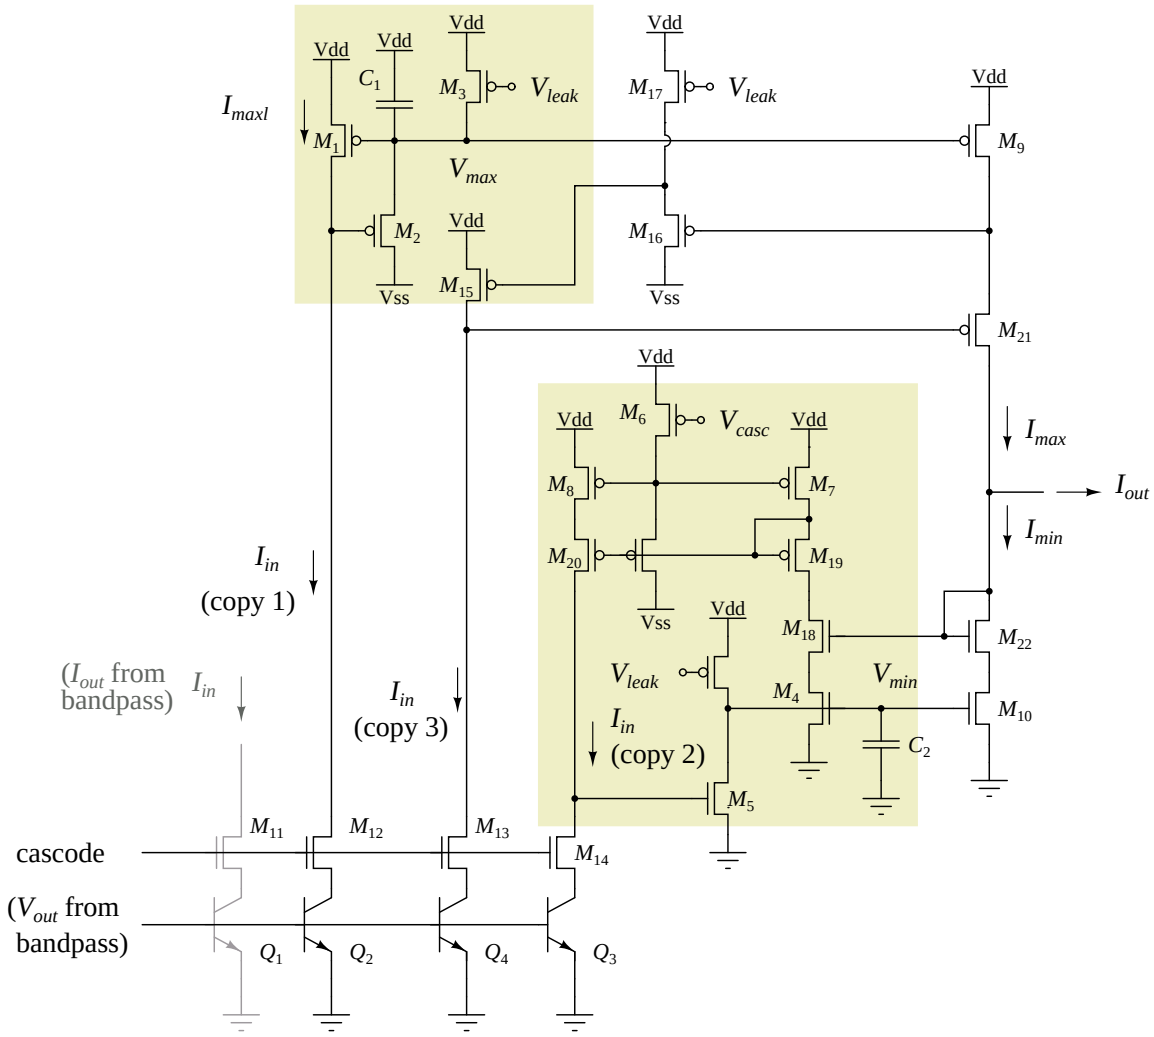

Figure 3.22: Peak-peak detector with cascodes, as fabricated.

The point of the L-1 norm is to force the outputs to maintain the same total value as a form of automatic gain control.

$$\sum_{j=1}^{N} I_{out}(j) = I_{norm} \tag{3.42}$$

Such a circuit 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. It was, in fact, the initial impetus for developing the current-mode, log-domain bandpass filters. The circuit was invented by Barry Gilbert [50] and is known, at least in analog VLSI circles, as the "Gilbert normalizer" [1]. Figure 3.23 shows this circuit. The bias current I_{norm} is a simple way to assure the relationship of Equation (3.42) (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

$$I_{in}(i)I_{out}(j) = I_{in}(j)I_{out}(i) \qquad \forall\, i, j \tag{3.43}$$

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

$$I_{out}(1) = I_{in}(1)\frac{I_{out}(i)}{I_{in}(i)} \tag{3.44}$$

$$I_{out}(2) = I_{in}(2)\frac{I_{out}(i)}{I_{in}(i)} \tag{3.45}$$

$$\vdots$$

$$I_{out}(N) = I_{in}(N)\frac{I_{out}(i)}{I_{in}(i)} \tag{3.46}$$

$$\sum_{j=1}^{N} I_{out}(j) = \sum_{j=1}^{N} I_{in}(j)\frac{I_{in}(i)}{I_{out}(i)} \tag{3.47}$$

$$I_{norm} = \sum_{j=1}^{N} I_{in}(j)\frac{I_{in}(i)}{I_{out}(i)} \tag{3.48}$$

$$I_{out}(i) = I_{norm}\frac{I_{in}(i)}{\sum_{j=1}^{N} I_{in}(j)} \tag{3.49}$$

3.16 Experimental Results

We have fabricated and tested several different chips, one of which was a test chip with individual filter structures on it, and one a complete signal-processing system [57] containing a

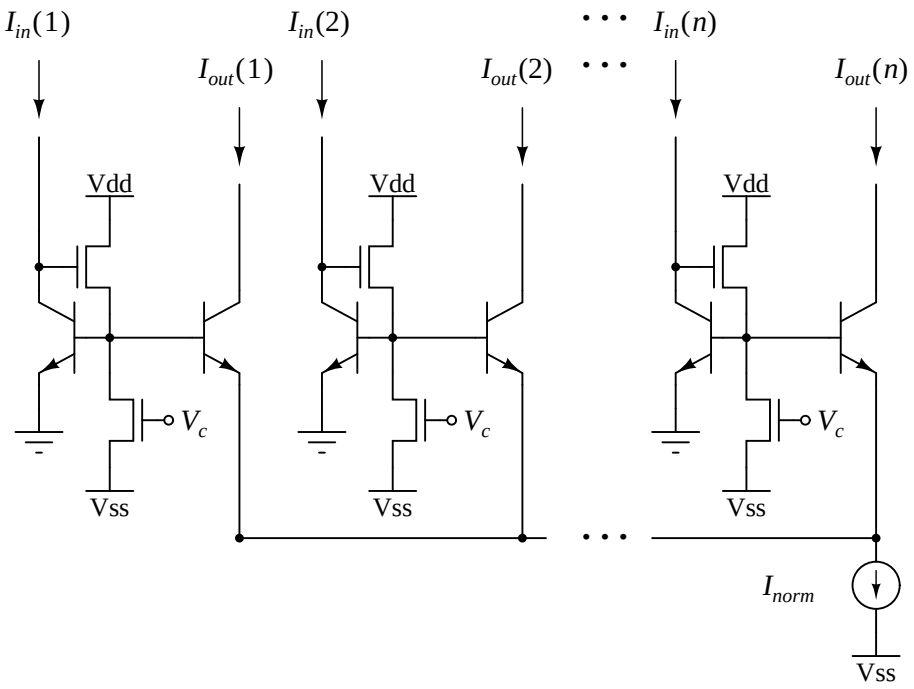

Figure 3.23: L-1 normalization circuit, after Gilbert [50].

filterbank of log-domain bandpass filters (Figure 3.28). The filterbank consists of 15 channels each consisting of two cascaded bandpass filters followed by rectification and smoothing of the signal at the output; The classifier circuit for which this filterbank was designed (Chapter 4) is a current-mode system which takes the current outputs of the frontend system directly, without converting to voltage. The chips were fabricated on 2 mm dies in 2 µm (test chip) and 1.2 µm (filterbank) analog n-well BiCMOS processes, with double metal, double poly, and a p-base layer for creating vertical NPN bipolar transistors. For high-density integration, we limited capacitors to about 2 pF and made all BJTs with a minimum size well. The size of the capacitors implies (from Equation (3.31)) bias currents of 30 pA to 3 nA for center frequencies spaced from 100 Hz to 10000 Hz, respectively.

Figure 3.24 shows measurements taken from a first-order lowpass section on the test chip. The circuit is that shown in Figure 3.7, with a capacitor value of about 1 pF, a bias current ratio $I_a = 2I_b$ and bias currents generated by applying a voltage to the gate of a (cascoded) nMOS transistor in weak inversion. The gate voltage was stepped in linearly-spaced increments, which ideally should yield exponentially-spaced corner frequencies. The data show a falloff of approximately 22 dB/decade and confirm the exponential spacing of corner frequencies. The response maps well

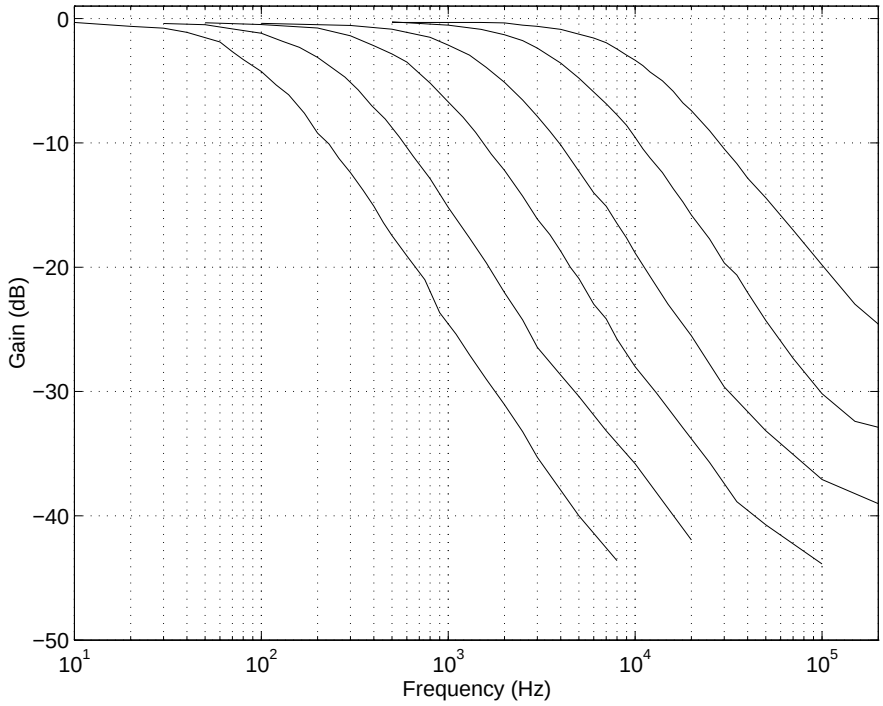

Figure 3.24: Measured magnitude response of the log-domain first-order lowpass filter from the test chip.

Table 3.1: Measured filterbank characteristics.

Die Size	2.2 mm $\times$ 2.2 mm
Process	1.2 µm n-well BiCMOS double-poly, double-metal
Power supply $V_{dd} - V_{ss}$	4.0–6.0 V
V_{ss}	-0.5 V
Number of channels	15
Dimensions of bandpass filter	75 µm $\times$ 234 µm
Capacitor size	2.0 pF
Input DC bias	10 µA
Output bias I_{DC}	10 µA
Max. input AC signal	8 µA peak-peak.
Power consumption	200 µW at 5.0 V
Dynamic range	35 dB at listed conditions
Bias current I_b	30 pA to 3 nA
Frequency tuning range	50 Hz to 15 kHz
Q tuning range	1 to 10
Gain nonuniformity across channels	18% at Q = 6

over the entire audio frequency range, although the dynamic range is limited by an approximately 50 dB noise floor. Note that the transfer function at 100 Hz corner frequency implies that with the use of base compensation, the bipolar transistors maintain a well-defined exponential I-V relationship even at collector currents as low as 30 pA.

The remaining data were measured from the filterbank system chip, the characteristics of which are summarized in Table 3.16. All measurements shown were made with a power supply of 5 V, though the system showed similar characteristics at 6 V and 4 V. Cascoded MOS transistors implementing the current sources and mirrors plus the p-n junction voltage drops in the translinear loops limit circuit operation at less than approximately 4 V. Power consumption of the circuit is dominated by the DC bias added to the filter input and output, which affects the SNR at the filter output. For the 5 V power supply and an input/output bias of 10 µA, the filter circuit of Figure 3.14 dissipates about 200 µW and has a dynamic range of about 35 dB.

Figure 3.25 shows characteristics of the bandpass filter at three different frequencies in the audio band and three different Q values. The filter consists of two second-order bandpass sections of the type shown in Figure 3.14, connected in cascade. Both filters in the cascade have adjustable Q values; for our measurements we varied the Q of the first filter while keeping the Q of the second filter fixed. The Q of the second filter has a value approximately equal to but strictly less than one and is used primarily to give the response a 40 dB/decade drop on the skirts of the passband.

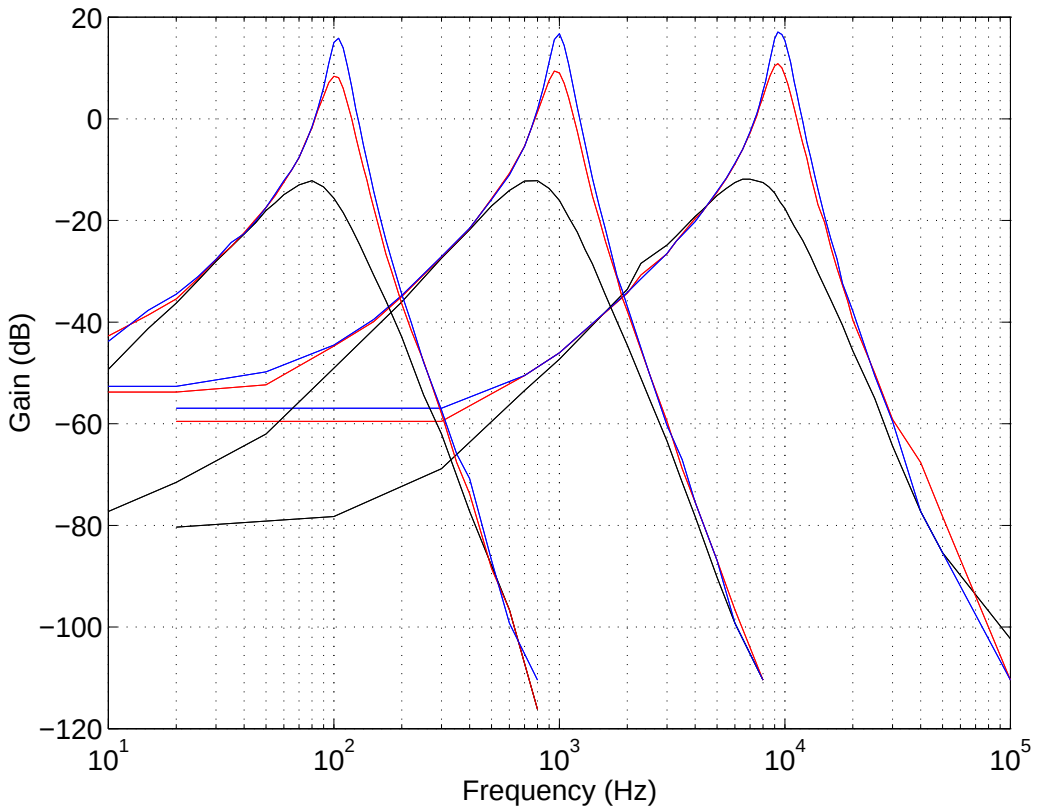

Figure 3.25: Measured magnitude response of one bandpass channel in the filterbank system, made of two cascaded second-order log-domain bandpass filters, over three tunings of the center frequency.

The first filter is responsible for the sharp bandpass response near the center frequency. Both filters are biased to have the same center frequency. Results show that considerable reduction in rolloff occurs on the lower side of the response, due to mismatch of components in the circuit, in particular the matching of bias currents I_b. These data confirm the simulation results of [44], although it is important to note that the effective gain (Q) of the filter appears to be completely independent of the center frequency (bias current I_b). The full tuning range of the circuit is from about 50 Hz to 15 kHz, limited on the low end by $1/f$ and thermal noise and on the top end by voltage drops across the CMOS current sources, which were sized appropriately for audio-frequency operation. Bipolar-based log-domain filters of similar design are capable of high-frequency operation in the tens of MHz.

The distribution of Q values, shown in Figure 3.26 and listed in Table 3.16, shows considerable variation across all channels in the filterbank with the same V_Q (see Figure 3.14) applied

to each. The variation is due primarily to error in matching I_b to $I_b + (1/Q)I_b$ in Figure 3.14. The dependency on $1/Q$ results in increasingly large gain differences between channels as Q becomes larger. A filterbank with sufficient number of channels cannot necessarily utilize precision trimming of current sources separately for each channel. Although much of the mismatch can be attributed to the quality of the fabrication process, better matching should be sought through circuit designs which minimize the number of matched current sources, especially those which must match a pMOS current source to an nMOS current sink.

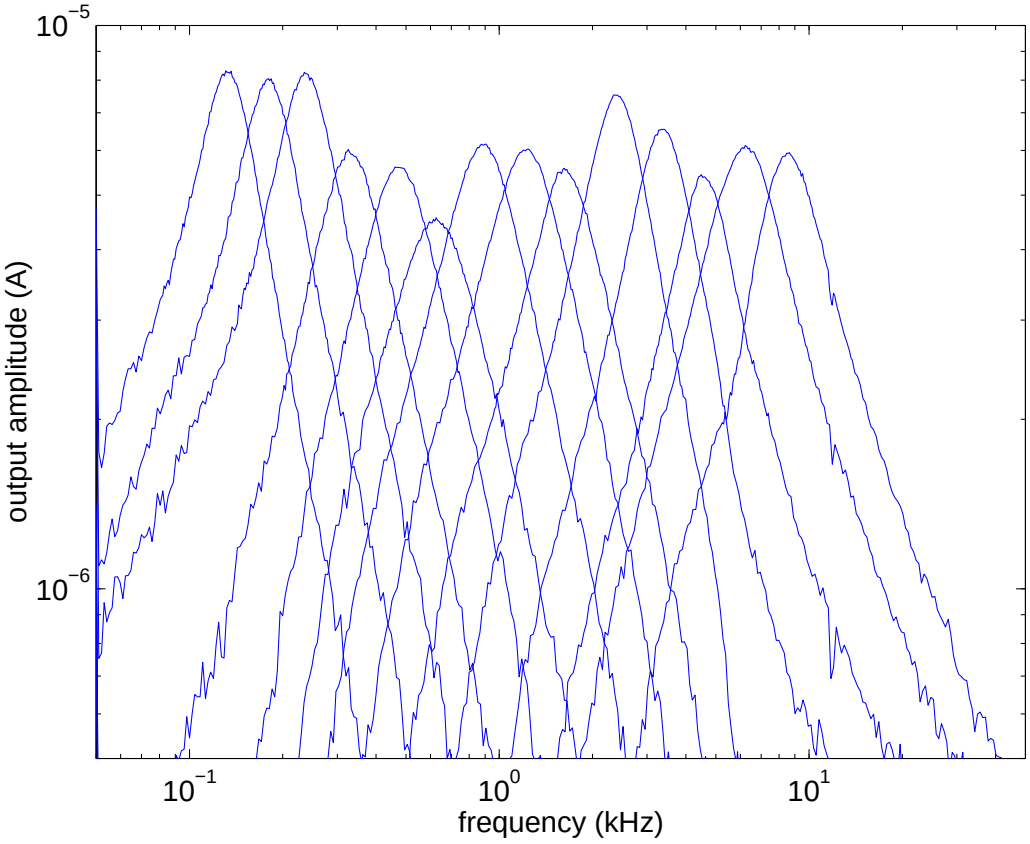

Figure 3.26: Measured magnitude response of all bandpass channels in the filterbank system, as measured at the output after peak-detection and smoothing.

The distribution of center frequencies derives from the same value I_b as the distribution of Q values, and follows the same exponential behavior (as discussed previously in Section 3.9). The main difference is that the filter center frequency is directly proportional to I_b, thus variance in center frequency is directly proportional to variance in I_b. Variations in I_b can come from only two sources: Physical nonidealities in the polysilicon strip which may cause the voltage taps to deviate from linear

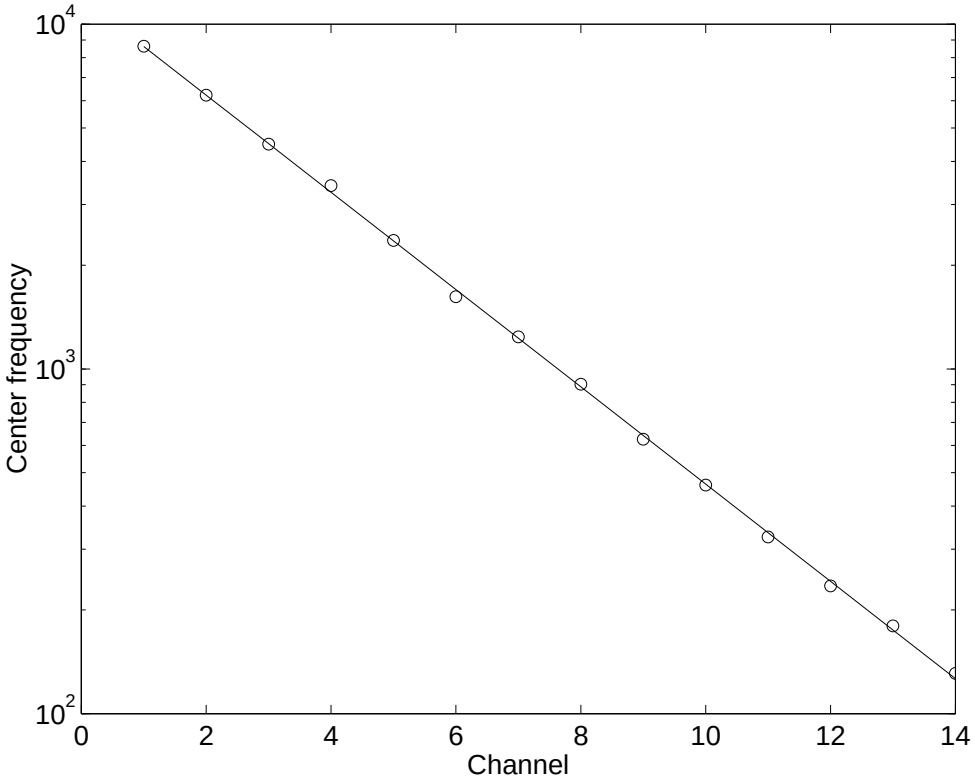

Figure 3.27: Measured center frequencies of all bandpass channels in the filterbank system (circles), compared to the ideal exponential spacing (solid line).

spacing, and mismatch between the pMOS transistors which are gated by the tapped voltages and directly generate I_b. Symmetry in the layout guarantees good matching between transistors at the same position in each channel. This is, in fact, true even between odd and even channels which are antisymmetric (one is flipped across the horizontal axis with respect to the other). The result is nearly perfect exponential spacing seen in Figure 3.27. The poorer matching in Q values derives from two sources: One is that there are two currents which are compared against each other. While good matching between I_b in one channel against I_b in another channel is guaranteed by large-scale symmetry, good local matching of two currents is not so easy to obtain, as the transistors which generate the two matched currents do not have the same physical environs. Particularly in this circuit, the current which generates I_b is sourced from the positive power supply into the capacitor node while the current which generates Q by its ratio to I_b is drawn from the capacitor node to the negative power supply, implying that one current is generated by a pMOS transistor and the other by an nMOS transistor. This situation gives notoriously poor matching properties.

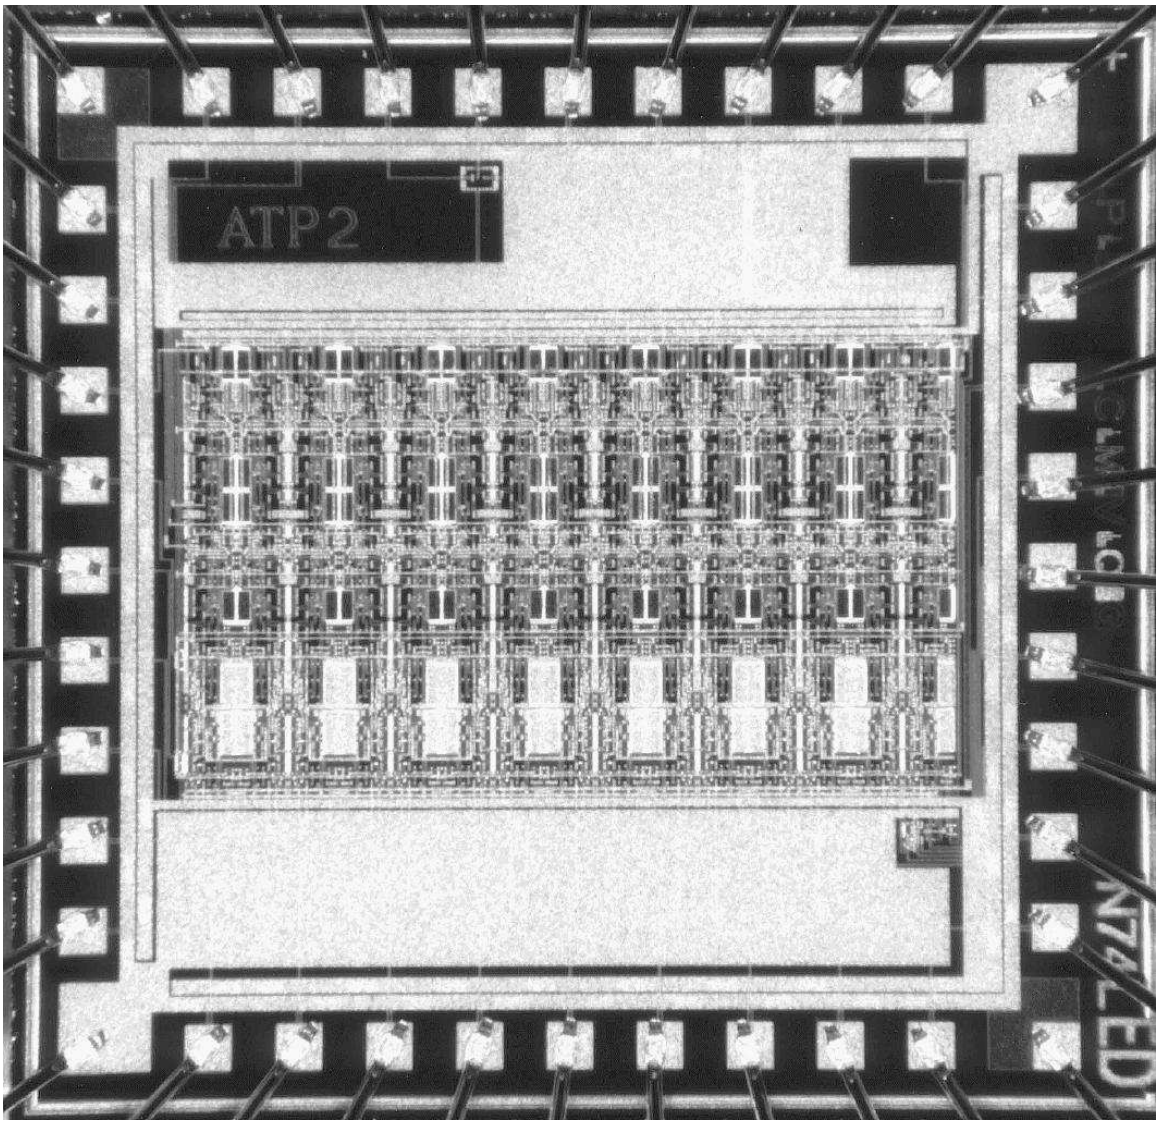

Figure 3.28: Photograph of the fifteen-channel bandpass filterbank fabricated in 1.2 µm technology inside a 2.2 mm $\times$ 2.2 mm padframe.

3.17 Summary

We have described how first-order systems can be synthesized from static and dynamic translinear principles, and how to generate higher-order filter transfer functions from these first-order building blocks. In particular, we have addressed circuit design issues relating to audio-frequency applications. We used these synthesis methods to design and fabricate VLSI analog signal processing systems of log-domain filters with current-domain input and output. Results from a first-order lowpass filter and a second-order bandpass filter fabricated in standard BiCMOS technologies show that these filters have transfer functions which map consistently over the entire audio frequency band. Design considerations for low-frequency operation ensure correct translinear operation of the bipolar transistors even at collector currents lower than 10 pA. We have characterized performance of the bandpass filters in the context of a high layout density 15-channel filterbank system. Performance of these filters is adequate for the applications for which they were designed [57].