

CIRCUIT MODELS OF SENSORY TRANSDUCTION IN THE COCHLEA

John Lazzaro and Carver Mead
Department of Computer Science
California Institute of Technology
Pasadena, California, 91125

Nonlinear signal processing is an integral part of sensory transduction in the nervous system. Sensory inputs are analog, continuous-time signals with a large dynamic range, whereas central neurons encode information with limited dynamic range and temporal specificity, using fixed-width, fixed-height pulses. Sensory transduction uses nonlinear signal processing to reduce real-world input to a neural representation, with a minimal loss of information.

An excellent example of nonlinear processing in sensory transduction occurs in the cochlea, the organ that converts the sound energy present at the eardrum into the first neural representation of the auditory system, the auditory nerve. Humans can process sound input over a 120-dB dynamic range, yet the firing rate of an auditory-nerve fiber can encode only about 25 dB of sound intensity. Humans can sense binaural time differences of the order of ten microseconds, yet an auditory-nerve fiber can fire at most once per millisecond. Using limited neural resources, the cochlea creates a representation that preserves the information essential for sound localization and understanding. Moreover, this neural code expresses auditory information in a way that facilitates feature extraction by higher neural structures.

We are building silicon integrated circuits that model sensory transduction in the cochlea, both to explore the general computational principles of the cochlea, and to create potentially useful devices for sound understanding, for sound localization, and for cochlear prostheses. In this paper, we describe the architecture and operation of an integrated circuit that models, to a limited degree, the evoked responses of the auditory nerve. The chip receives as input a time-varying voltage corresponding to sound input, and computes outputs that correspond to the responses of individual auditory-nerve fibers. The chip models the structure as well as the function of the cochlea; all subcircuits in the chip have anatomical correlates. The chip computes all outputs in real time, using analog continuous-time processing.

NEURAL ARCHITECTURE OF THE COCHLEA

Both mechanical and electrical processing occur in biological cochleas. The sound energy present at the eardrum is coupled into a mechanical traveling-wave structure, the basilar membrane, which converts time-domain information into spatially encoded information by spreading out signals in space according to their time scale (or frequency). Over much of its length, the velocity of propagation along the basilar membrane decreases exponentially with distance. The structure also contains active electromechanical elements; outer hair cells have motile properties, acting to reduce the damping of the passive basilar membrane and thus allowing weaker signals to be heard. Axons from higher brain centers innervate the outer hair cells; these centers may dynamically vary the local damping of the cochlea, providing frequency-specific automatic gain control (Kim, 1984).

Inner hair cells occur at regular intervals along the basilar membrane. Each inner hair cell acts as an electromechanical transducer, converting basilar-membrane vibration into a graded electrical signal. Several signal-processing operations occur during transduction. Inner hair cells half-wave rectify the mechanical signal, responding to motion in only one direction. Inner hair cells primarily respond to the velocity of basilar-membrane motion, implicitly computing the time derivative of basilar-membrane displacement (Dallos, 1985). Inner hair cells also compress the mechanical signal nonlinearly, reducing a large range of input sound intensities to a manageable excursion of signal level.

Spiral-ganglion neurons connect to each inner hair cell, and produce fixed-width, fixed-height pulses in response to inner-hair-cell electrical activity. The synaptic connection between the inner hair cell and the spiral-ganglion neuron may implement a stage of automatic gain control, exploiting the dynamics of synaptic-transmitter release (Geisler and Greenberg, 1986). Auditory-nerve fibers are axons from spiral-ganglion neurons; these fibers present a neural representation of audition to the brain.

When pure tones are presented as stimuli, an auditory-nerve fiber is most sensitive to tones of a specific frequency. This characteristic frequency corresponds to maximum basilar-membrane velocity at the location of the inner hair cell associated with the nerve fiber. The spiral trunk of the auditory nerve preserves this ordering; the nerve fibers are mapped cochleotopically and tonotopically. The mean firing rate of an auditory fiber encodes sound intensity, over about 25 dB of dynamic range. The temporal pattern of nerve firings reflects the shape of the filtered and rectified sound waveform; this phase locking does not diminish at high intensity levels (Evans, 1982).

SILICON MODELS OF THE COCHLEA

Both mechanical and electrical processing occur in biological cochleas. In the chip, however, we model both types of computation using electronic processing. A silicon model of the mechanical processing of the cochlea has been previously described (Lyon and Mead, 1988a; Mead, 1989). The circuit is a one-dimensional physical model of the traveling-wave structure formed by the basilar membrane. In this viewpoint of cochlear function, the exponentially tapered stiffness of the basilar membrane and the motility of the outer hair cells combine to produce a pseudoresonant structure.

The basilar-membrane circuit model implements this view of cochlear hydrodynamics using a cascade of second-order sections with exponentially scaled time constants. The cascade structure enforces unidirectionality, so a discretization in space does not introduce reflections that could cause instability in an active model. An analog, continuous-time circuit implementation of the model computes the pressure at selected discrete points along the basilar membrane in real time.




Figure 1. Circuit implementation of a second-order section. Input V_i and output V_o are time-varying voltages. The τ and Q control inputs set bias currents on transconductance amplifiers A1, A2, and A3, to control both the characteristic frequency and the peak height of the lowpass-filter response.

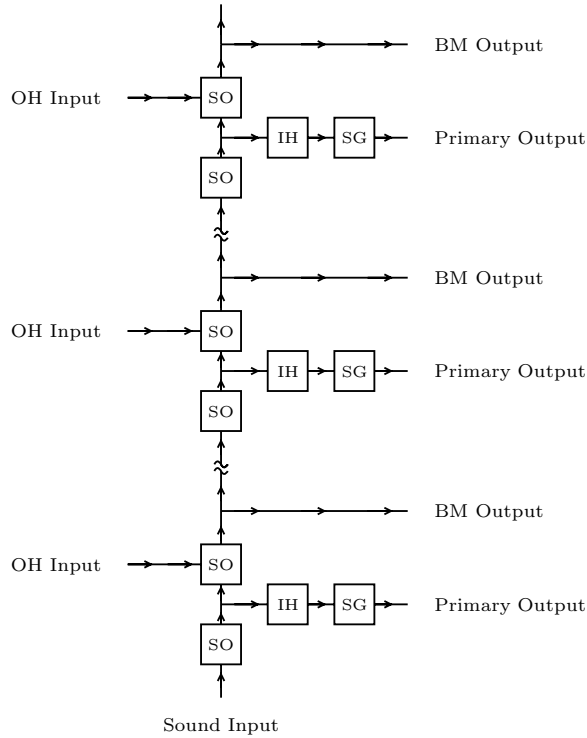


Figure 2. Block diagram of the chip. A time-varying input voltage, representing sound input to the cochlea, travels down the basilar-membrane model, a cascade of second-order sections (SO) with exponentially increasing time constants. Basilar-membrane (BM) circuit outputs show pressure along the membrane, whereas inputs modeling innervation of outer hair cells (OH) control local damping of the membrane circuit. Taps along the basilar membrane connect to a circuit model of inner hair cells (IH); outputs from inner hair cells connect to circuits that model spiral-ganglion neurons (SG). These neurons form the primary output of the chip, thus modeling auditory-fiber response.

Figure 1 shows the CMOS circuit implementation of a second-order section. Input and output signals for the circuit are time-varying voltages. The gain blocks are transconductance amplifiers, operated in the subthreshold regime. Capacitors are formed using the gate capacitance of n -channel and p -channel MOS transistors in parallel. Because of subthreshold amplifier operation, the time constant of the second-order section is an exponential function of the voltage applied to the transconductance control inputs of $A1$ and $A2$, labeled τ in Figure 1. Thus a cascade of second-order circuits, with a linear gradient applied to the τ control inputs, has exponentially scaled time constants. To implement this gradient, we used a polysilicon wire that travels along the length of circuit,

and connects to the τ control input of each second-order section. A voltage difference across this wire, applied from off chip, produces exponentially scaled time constants. The amplifier $A3$ provides active positive feedback to the membrane, modeling the active mechanical feedback provided by the outer hair cells in biological cochleas. A second polysilicon wire is connected to the transconductance inputs of the $A3$ amplifiers in each second-order section (labeled Q in Figure 1); a voltage gradient across this wire similar to that on the τ control inputs sets all the second-order sections to the same response shape.

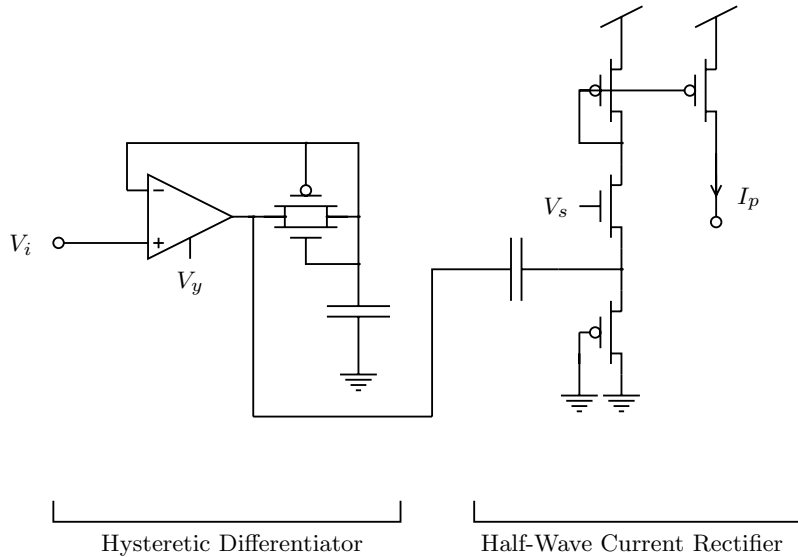


Figure 3. The inner-hair-cell circuit model. Input V_i , from the basilar-membrane circuit, is a time-varying voltage. The hysteretic-differentiator circuit, biased by voltage V_y , performs time differentiation and logarithmic compression. The output of the hysteretic differentiator, a time-varying voltage, connects to the half-wave current-rectifier circuit, which is shown in more detail in Figure 4.

This circuit model of cochlear mechanics is the foundation of our integrated circuit; Figure 2 shows the complete architecture of the chip. A way to model the adjustment of basilar-membrane damping by higher brain centers is to use an automatic-gain-control system that varies the damping of the second-order sections locally. We have not implemented this automatic-gain-control system; however, we have brought off chip several taps from the polysilicon wire that connects to the Q control of the second-order sections, allowing off-chip experiments with automatic gain control. To complete our circuit model of the auditory periphery, we have added circuits that model inner-hair-cell and spiral-ganglion-neuron functions.

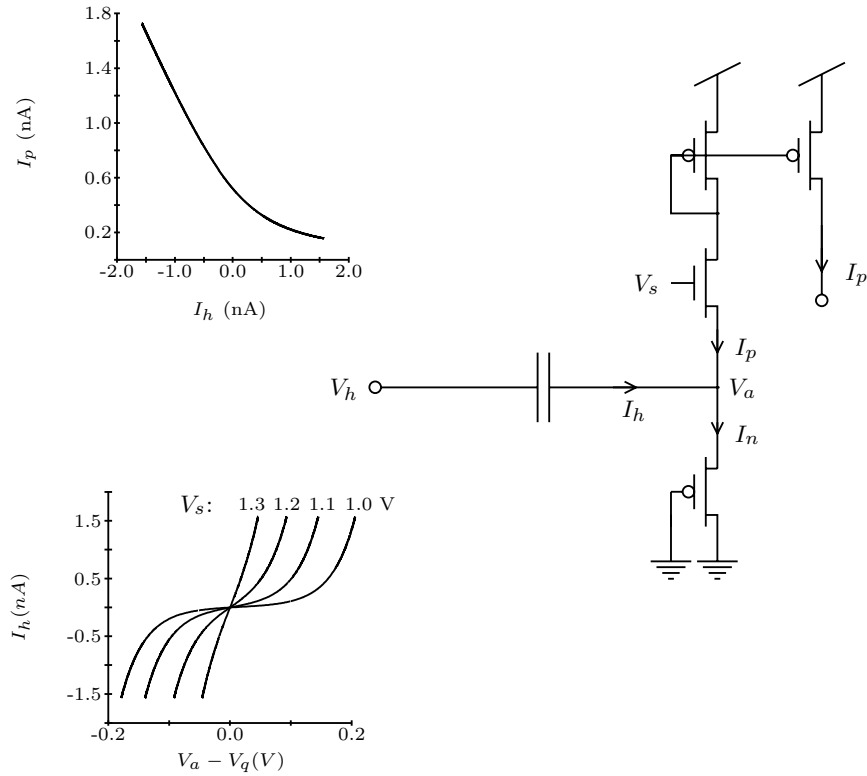


Figure 4. The half-wave current-rectifier circuit. Input V_h , from the hysteretic-differentiator circuit, is a time-varying voltage. A floating capacitor couples V_h into the node associated with V_a , as the bidirectional time-varying current I_h . The bottom graph shows the change in V_a required to sink or source I_h , for several values of bias voltage V_s ; the voltage V_q is the value of V_a when $I_h = 0$. When $V_a = V_q$ and $I_h = 0$, the circuit output, the unidirectional current I_p , is at a quiescent value, I_q , set by V_s . Nonzero values of I_h modulate the output current I_p about I_q ; for large $|I_h|$ relative to I_q , the circuit output I_p is a half-wave-rectified version of I_h , as shown in the top graph. Graphs show theoretical responses.

Figure 3 shows our inner-hair-cell circuit model. A hysteretic-differentiator circuit (Mead, 1989) processes the input-voltage waveform from the basilar-membrane circuit, performing time differentiation and logarithmic compression. The circuit enhances the zero-crossings of the input waveform, accentuating phase information in the signal. The output voltage of the hysteretic differentiator connects to a novel implementation of a half-wave current rectifier.




Figure 5. The spiral-ganglion-neuron circuit. Circuit input, from the half-wave-rectification circuit, is the unidirectional current I_i . The circuit converts this current into fixed-width, fixed-height voltage pulses, at output V_o . The bias voltage V_p sets pulse width; the output voltage V_o pulses between V_{dd} and ground.

Figure 4 shows our half-wave current-rectifier circuit. To understand its operation, consider the state of this circuit when the input voltage V_h is constant. If V_h is constant, $I_h = 0$, and V_a adapts such that $I_p = I_n$. For $I_h = 0$, we define the quiescent conditions $I_q \equiv I_p = I_n$ and $V_q \equiv V_a$. The value of I_q depends on the circuit bias voltage, V_s . A current mirror reflects this quiescent current to the circuit output. Thus, the output of the half-wave current-rectifier circuit in response to a constant voltage input is an adjustable bias current.

Now consider the circuit state when the input voltage V_h is a time-varying waveform. During the positive-going phase of the waveform, the current I_h is positive, and $I_n = I_h + I_p$. As I_n increases, V_a must also increase; the amount of increase depends on the circuit bias voltage, V_s , as shown in the bottom graph in Figure 4. However, if V_a increases, I_p must decrease. So, during the positive-going phase of the waveform, the output current I_p decreases from the quiescent current I_q .

During the negative-going phase of the waveform, the current I_h is negative, $I_p = |I_h| + I_n$, and the output current of the circuit increases from the quiescent current I_q . Thus, the circuit converts the input time-varying voltage waveform V_h into a unidirectional current waveform I_p . For large $|I_h|$ relative to I_q , the current waveform I_p is not symmetrical about I_q , and the average value of I_p is greater than that of I_q ; thus, the circuit performs the rectification function, as shown in the top graph in Figure 4.

The current I_p is the output of the inner-hair-cell circuit. The spiral-ganglion neuron circuit model, shown in Figure 5, converts this current into fixed-width, fixed-height pulses. The circuit — a slightly modified version of the neuron circuit in (Mead, 1989) — creates a pulse rate that is linear in input current, for sufficiently low pulse rates. Thus, the average pulse rate of the circuit reflects the average value of I_p , whereas the temporal placement of each pulse reflects the shape of the current waveform I_p .

SILICON BASILAR-MEMBRANE RESPONSE

To test the tuning properties of the silicon auditory-nerve fibers, we duplicated a variety of classical auditory-nerve measurements. In these experiments, we tuned the basilar-membrane circuit to span about seven octaves, from 50 Hz to 10,000 Hz. We set the maximum firing rates of the auditory-fiber outputs at 150 to 300 spikes per second, with spike widths of 5 to 20 μ s.




Figure 6. a: The response of the basilar-membrane circuit at a single point, to pure tones at a fixed input amplitude (0 dB = 3 mV peak). **b:** Transfer function of a single position on the basilar membrane of the squirrel monkey (Rhode, 1971). The curves show amplitude of vibration for constant malleus displacement.

In this configuration, without an input signal, the auditory-fiber outputs fire at less than 0.1 spike per second. At the characteristic frequency of a fiber, pure tones of a few millivolts peak amplitude produce responses significantly above this spontaneous rate. The chip can process tones up to about 1 V of peak amplitude, yielding approximately 60 dB of usable dynamic range.

Adding a preprocessor to basilar-membrane circuit, to limit intense input signals, would extend the upper limit of the dynamic range. A biological cochlea has a mechanical limiter as a preprocessor — the stapedial reflex. Designing more sensitive inner-hair-cell circuits would extend the lower limit of dynamic range. Both dynamic-range enhancements are currently under development.

Figure 6(a) shows a frequency-response plot for the basilar-membrane circuit, at a position with a best frequency of about 1900 Hz. The plot shows a flat response for frequencies significantly below the best frequency, a 12-dB response peak at the best frequency, and a sharp dropoff to the noise floor for frequencies significantly above the best frequency. This response is qualitatively similar to the frequency-response curve taken from the basilar membrane of the squirrel monkey using the Mossbauer effect, shown in Figure 6(b) (Rhode, 1971). Near the best frequency, basilar-membrane pressure, computed by the chip, is approximately equal to basilar-membrane displacement, measured by Rhode. Quantitatively, the bandwidth of the resonance peak of the chip response is wider than that of the physiological data; a cascade of second-order sections does not yield an optimal model of cochlear hydrodynamics (Lyon and Mead, 1988b).

The resonance peak of the chip response decreases for large-amplitude sinusoids, because the feedback amplifier *A3* in the second-order sections saturates. The resonance peak in a physiological cochlea also decreases for large-amplitude inputs (Rhode, 1971). The silicon and physiological cochleas may show decreased resonance for similar reasons; for high sound intensities, outer hair cells in the physiological cochlea may not be capable of a linear response to basilar-membrane motion. Alternatively, an automatic-gain-control system may increase basilar-membrane damping locally for high-intensity sounds, by modulating the mechanical effect of the outer hair cells (Kim, 1984).

TUNING PROPERTIES OF THE SILICON AUDITORY NERVE

We characterized the tuning properties of the auditory-nerve-fiber circuit model, using pure tones as input. In response to a pure tone of sufficient intensity and appropriate frequency, the silicon auditory fiber produces spikes at a constant mean rate, as shown in Figure 7. The mean spike rate of a silicon fiber, in response to a constant tone, does not decrease over time, unlike that of a physiological auditory fiber; this lack of adaptation indicates the absence of dynamic automatic gain control in our model.

Figure 8(a) shows the mean spike rate of a silicon auditory fiber as a function of pure tone frequency. For low-amplitude tones, the fiber responds to a narrow range of frequencies; for higher-intensity tones, the fiber responds to a wider range of frequencies. The saturating nonlinearities of the basilar-membrane circuit and of the inner-hair-cell circuit cause the bandwidth of the fiber to increase with sound intensity. Qualitatively, this behavior matches the

iso-intensity plots from an auditory-nerve fiber in the squirrel monkey (Rose et al., 1971), shown in Figure 8(b). Quantitatively, the saturation of the amplifiers in the forward path (A1 and A2) produce a detuning that is not a proper model of basilar-membrane mechanics.




Figure 7. Output of a silicon auditory fiber (bottom trace) in response to a sinusoidal input (top trace). The frequency of the input is the characteristic frequency of the fiber.

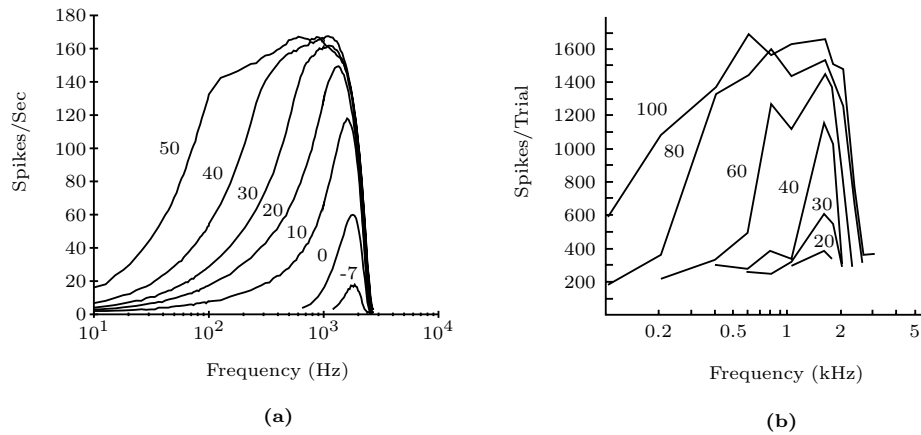


Figure 8. a: Plots showing the mean spike rate of a silicon auditory fiber as a function of pure tone frequency. Legend numbers indicate tone amplitude, in dB. **b:** Plots showing the number of discharges of an auditory fiber in the squirrel monkey, in response to a 10-s pure tone (Rose et al., 1971). Legend numbers indicate tone amplitude, in dB.

Figure 9(a) shows the mean spike rate of a silicon auditory fiber as a function of pure tone amplitude, at frequencies below, at, and above the best frequency of the fiber. In response to its characteristic frequency, 2100 Hz, the fiber encodes about 25 dB of tone amplitude before saturation. Figure 9(b) shows rate-intensity curves from an auditory fiber in the cat (Sachs and Abbas, 1974). At its characteristic frequency, the physiological fiber also encodes about 25 dB of tone amplitude before saturation. The shape of the biological and silicon curves at the characteristic frequency is remarkably similar, giving us some confidence in the validity of this modeling paradigm. In response to frequencies below and above the characteristic frequency, the functional forms of the silicon fiber responses are different from those of the physiological data. Most notably, the saturation rate of a silicon fiber for frequencies below the fiber's characteristic frequency exceeds the saturation rate of the silicon fiber at the fiber's characteristic frequency. This behavior is also a direct result of the undesired saturation at high input intensities of second-order-section amplifiers A1 and A2, shown in Figure 1, which model the stiffness of the basilar membrane. Above its best frequency, the response of the model decreases in a manner that is reminiscent of its biological counterpart.




Figure 9. a: Plots showing the mean spike rate of a silicon auditory fiber as a function of pure tone amplitude. Legend numbers indicate tone frequency, in Hz. **b:** Plots showing the mean spike rate of an auditory fiber in the cat, as a function of pure tone amplitude (Sachs and Abbas, 1974). Legend numbers indicate tone frequency, in Hz.

Figure 10(a) shows iso-response curves for four silicon auditory-nerve fibers. These plots represent an iso-rate section through the iso-intensity curves of Figure 8(a), at a spike rate for each fiber that was comfortably above the spontaneous rate. The chip response accurately models the steep high-frequency tail of tuning curves from cat auditory fibers (Kiang, 1980), shown in Figure 10(b); the shapes of physiological and chip tuning curves are qualitatively similar.

The bandwidth of the chip fibers for low sound intensities, however, is significantly wider than that of the physiological response. This problem stems from the wider bandwidth of the basilar-membrane circuit model, relative to that of the physiological data, as well as from the lack of a dynamic automatic-gain-control system for modulating the damping of the basilar-membrane circuit. The high-frequency cutoff of the iso-response curves, shown in Figure 10(a), is much steeper than is the cutoff of the iso-input curves shown in Figure 8(a). In a linear system, these two measurements would give identical results. The difference reflects the presence of a saturating nonlinearity in the system; the inner-hair-cell circuit and the basilar-membrane circuit provide this saturation function.




Figure 10. a: Plots showing iso-response curves for four silicon auditory fibers. The plots represent an iso-rate section through the iso-intensity curves of each fiber. Constant rates for each curve are, from the highest-frequency curve downward, 21.5, 16, 61, 59 spikes/s. **b:** Plots showing tuning curves from auditory fibers in the cat (Kiang, 1980). Fifty-ms tone bursts were presented at 10/s. Each tuning curve shows the sound pressure level (SPL) at the tympanic membrane (eardrum) that generates 10 spikes/s more activity during the tone bursts than during the silent interval.

TIMING PROPERTIES OF THE SILICON AUDITORY NERVE

The temporal firing patterns of the silicon auditory-nerve fibers encode information. Figure 11(a) shows period histograms of a chip fiber, in response to -5 - to 50 -dB pure tones at the fiber's characteristic frequency (0 dB = 3 mV peak); these histograms show the probability of a spike output occurring within a particular time interval during a single cycle of the input sinusoid. The fiber preserves the shape of the input sinusoid throughout this intensity

range; this behavior matches data from an auditory fiber in the cat (Rose et al., 1971), shown in Figure 11(b). Unlike the cat fiber, however, the silicon fiber does not preserve absolute phase at higher intensities; this deficiency results from the saturation of the amplifiers $A1$ and $A2$ that model basilar-membrane stiffness. The temporal firing patterns of the silicon auditory-nerve fiber are, however, a good representation of signal periodicity; the synchronization ratios (normalized magnitude of the first Fourier coefficient) of the period histograms in Figure 11(a) are 0.5 to 0.6, comparable to those of physiological data at the same frequency (Rhode et al., 1978).

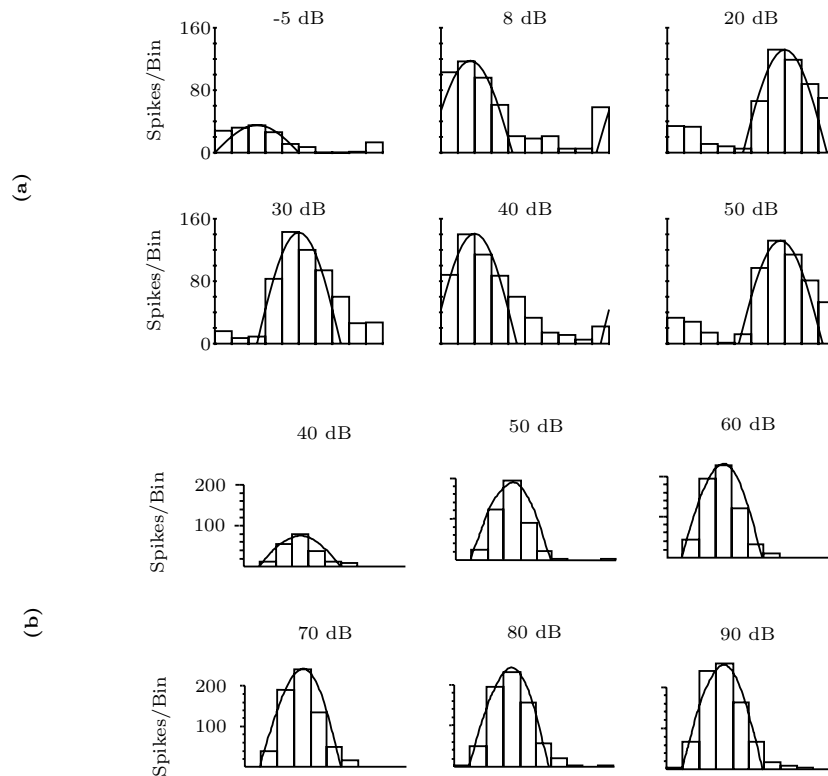


Figure 11. a: Period histograms of the silicon auditory-fiber response to a pure tone of 1840 Hz, near the fiber’s best frequency. Amplitude of tone is shown above each plot. Histogram width is $54 \mu s$. Each histogram begins at a constant position, relative to the input sinusoid; each is fitted to a sinusoid of best amplitude and phase. **b:** Period histograms of the response of an auditory fiber in the cat, to a low-frequency tone (Rose et al., 1971). Amplitude of pure tone is shown above each plot. Each histogram is fitted to a sinusoid of best amplitude but fixed phase.




Figure 12. a: PST histogram of the rarefaction click response of a silicon auditory-nerve fiber. Click amplitude is 60 mV (26 dB peak); click width is 100 μ s. Histogram is for 2000 click presentations; the width of each bin is 58 μ s. **b:** Compound PST histogram of the click response of a silicon auditory-nerve fiber. Rarefaction click response is plotted as positive values; condensation click response is plotted as negative values. Conditions are identical to those of Figure 3(a). **c:** Compound PST histogram of the click response of an auditory fiber in the cat (Kiang et al., 1965). Click level is 30 dB relative to threshold response level; clicks width is 100 μ s. Rarefaction click response is plotted as positive values; condensation click response is plotted as negative values. **d:** Compound PST histogram of the click response of a silicon auditory-nerve fiber, for a 200-mV click (36-dB click). All other conditions are identical to those of Figure 3(a).

The timing properties of silicon auditory-nerve fibers encode the click response of the basilar-membrane circuit. In response to a click of medium intensity, a silicon auditory-nerve fiber produces one or several spikes. To extract the click response from these spikes, we present the click stimulus to the chip

many times, and record the responses of a silicon auditory-nerve fiber. These data are reduced to a poststimulus-time (PST) histogram, in which the height of each bin of the histogram indicates the number of spikes occurring within a particular time interval after the presentation of the click.

A PST histogram of the response of a silicon auditory-nerve fiber to a repetitive rarefaction click stimulus shows a half-wave-rectified version of a damped sinusoidal oscillation (Figure 12a). The frequency of this oscillation, 1724 Hz, is approximately the best frequency of the basilar-membrane position associated with this silicon nerve fiber. The half-wave rectification of the inner-hair-cell circuit removes the negative polarity of oscillatory waveform from the PST histogram of the click response. Repeating this experiment using a condensation click recovers the negative polarity of oscillation; a compound PST histogram, shown in Figure 12(b), combines data from both experiments to recreate the ringing waveform produced by the basilar-membrane circuit. Figure 12(c) shows a compound PST histogram of the click response of an auditory fiber in the cat (Kiang et al., 1965). Qualitatively, the circuit response matches the physiological response.

Figures 12(a) and 12(b) are chip responses to a 60-mV click stimulus (26 dB, 0 dB = 3 mV peak). Higher-intensity clicks produce oscillatory responses with increased damping; a compound PST histogram of chip auditory-nerve response to a 36-dB click shows reduced ringing (Figure 12d). This effect is a direct result of the nonlinear response of the basilar-membrane model; physiological basilar-membrane click responses also show reduced ringing at high click-intensity levels (Robles et al., 1976).

DISCUSSION

Our integrated circuit model captures many essential features of data representation in the auditory nerve; moreover, it computes the representation in real time. There are many traditional engineering representations of audition, however, that are also amenable to analog implementation. What advantages does a silicon auditory-nerve representation offer to a designer of artificial sensory systems?

As shown in Figures 11 and 12, an auditory-nerve fiber encodes a filtered, half-wave-rectified version of the input waveform, over a wide dynamic range, using the temporal patterning of fixed-width, fixed-height pulses. This representation supports the efficient, massively parallel computation of signal properties, using autocorrelations in time and cross-correlations between auditory fibers. In this representation, a correlation is simply a logical AND operation, performed by a few synapses in neural systems, or by a few transistors in silicon systems. Axonal delays in neural systems provide the time parameter for computing autocorrelations; in silicon systems, we model this delay with compact

monostable circuits (Mead, 1989). We have used these techniques in a 220,000-transistor chip that models the auditory-localization system of the barn owl (Lazzaro and Mead, 1989).

The nonlinear filtering properties of the auditory-nerve fibers, shown in Figures 8 and 10, enhance these correlations. In a quiet environment, auditory fibers have narrow bandwidths; each fiber carries independent information, yielding rich correlations. In noisier environments, the tuning of auditory fibers widens, increasing the number of fibers that carry information about the signal. This detuning ensures that some fibers still encode signal properties reliably (Greenberg, 1988).

As shown in Figure 9, auditory fibers encode about 25 dB of signal intensity. Dynamic automatic gain control, present in a physiological cochlea, enhances this range; in addition, different populations of auditory fibers have different thresholds, further enhancing the encoding of signal intensity. Although not sufficient as a primary representation of sound, rate encoding of signal intensity is a valuable secondary cue, particularly for the detection of rapid spectral changes and the encoding of aperiodic sounds. Future versions of our chip will include these enhancements for rate encoding of signal intensity.

In conclusion, we have designed and tested an integrated circuit that computes, in real time, the evoked responses of auditory nerve, using analog, continuous-time processing. The chip offers a robust representation of audition, which can serve as a solid foundation for analog silicon systems that model higher auditory function.

Acknowledgements

We thank R. Lyon for valuable contributions throughout the project. We thank R. Lyon, M. Mahowald, L. Dupre; and D. Gillespie, for critically reading and correcting the manuscript. We thank Hewlett-Packard for computing support, and DARPA and MOSIS for chip fabrication. This work was sponsored by the Office of Naval Research and the System Development Foundation.

References

- Dallos, P. (1985). Response characteristics of mammalian cochlear hair cells. *J. Neurosci.* **5**: 1591–1608.
- Evans, E. F. (1982). Functional anatomy of the auditory system. In Barlow, H. B. and Mollon, J. D. (eds), *The Senses*. Cambridge, England: Cambridge University Press, p. 251.
- Geisler, C.D. and Greenberg, S. (1986). A two-stage nonlinear cochlear model possesses automatic gain control. *J. Acoust. Soc. Am.* **80**: 1359–1363.
- Greenberg, S. (1988). The ear as a speech analyzer. *J. Phonetics* **16**: 139–149.

- Kiang, N. Y.-s, Watanabe, T., Thomas, E.C., and Clark, L.F. (1965). *Discharge Patterns of Single Fibers in the Cat's Auditory Nerve*. Cambridge, MA: M.I.T Press.
- Kiang, N. Y.-s, (1980). Processing of speech by the auditory nervous system. *J. Acoust. Soc. Am.* **68**: 830–835.
- Kim, D. O. (1984). Functional roles of the inner- and outer-haircell subsystems in the cochlea and brainstem. In Berlin, C. I. (ed), *Hearing Science*. San Diego, CA: College-Hill Press, p. 241.
- Lazzaro, J. P. and Mead, C.A. (1989). Silicon models of auditory localization, *Neural Computation* **1**: 41–70.
- Lyon, R. F. and Mead, C. A. (1988a). An analog electronic cochlea. *IEEE Trans. Acoust., Speech, Signal Processing* **36**: 1119–1134.
- Lyon, R. F. and Mead, C. A. (1988b). *Cochlear Hydrodynamics Demystified*. Caltech Computer Science Technical Report Caltech-CS-TR-88-4, Pasadena, CA, February.
- Mead, C. A. (1989). *Analog VLSI and Neural Systems*. Reading, MA: Addison-Wesley.
- Rhode, W. S. (1971) Observations of the vibration of the basilar membrane in squirrel monkeys using the Mossbauer technique. *J. Acoust. Soc. Am.* **49**: 1218–1231.
- Rhode, W. S., Geisler, C.D., and Kennedy, D.T. (1978). Auditory nerve fiber response to wide-band noise and tone combinations. *J. Neurophysiol.* **41**: 692–704.
- Robles, L., Rhode, W. S., and Geisler, C.D. (1976) Transient response of basilar membrane measured in squirrel monkeys using the Mossbauer effect. *J. Acoust. Soc. Am.* **59**: 926–939.
- Rose, J.E., Hind, J.E., Anderson, D. J., and Brugge, J. F. (1971). Some effects of stimulus intensity on response of auditory nerve fibers in the squirrel monkey. *J. Neurophysiol.* **34**: 685–699.
- Sachs, M. B. and Abbas, P. J. (1974) Rate versus level functions for auditory-nerve fibers in cats: Tone-burst stimuli. *J. Acoust. Soc. Am.* **56**: 1835–1847.