

1 FILTER CASCADES AS ANALOGS OF THE COCHLEA

Richard F. Lyon

Foveonics Inc., 10131-B Bubb Rd., Cupertino CA 95014

dicklyon@acm.org

1.1 MODELS OF COCHLEAR WAVE PROPAGATION

Wave propagation in the cochlea can be modeled at various levels and for various purposes. We are interested in making models of cochlear signal processing, in analog or digital VLSI or in software, suitable for supporting improved hearing aids, speech-recognition systems, and other engineered hearing machines. We are also interested in developing models that can contribute to a deeper understanding of how hearing works. Hence, a neuromorphic approach, in which the functionality of the model emerges from a form that is loosely copied from the nervous system, is appropriate.

The filter-cascade approach to modeling the cochlea is based on the observation that small segments of the cochlea act as local filters on waves propagating through them. Thus, a cascade of filters can emulate the whole complex distributed hydrodynamic system. This modeling approach can include compressive and adaptive aspects of the peripheral auditory nervous system as well, using analogs of cochlear nonlinear distortion and efferent feedback. We summarize the underpinnings, advantages, and limitations of this approach in this paper, so that readers can more readily understand other papers on filter-cascade approaches and implementations.

Figure 1.1 shows the filter-cascade structure that we discuss in this paper.

1.2 COCHLEAR HYDRODYNAMICS IN THE LIOUVILLE-GREEN APPROXIMATION

Imagine the cochlea as a three-dimensional (3D) hydrodynamic system with a linear or one-dimensional (1D) array of sensors attached to it. In the real




Figure 1.1 A Filter Cascade. This simple structure of cascaded filter stages is a useful analog to the hydrodynamic wave propagation system of the cochlea.

cochlea, the 3D and 1D structures follow a complicated helical path, with the sensors being the inner hair cells (IHCs) of the organ of Corti. Abstractly, we refer to the one dimension that indexes the sensors as the cochlear *place* dimension. From a functional point of view, we care about the response only as a function of the input signal and of the 1D place, so there is only one relevant spatial dimension at the model output.

Wave propagation in the cochlea depends on fluid displacement in three dimensions, on membrane bending and stretching, and on related 2D and 3D and micromechanical effects within the organ of Corti, which have mostly unknown physical parameters. There have been many arguments in the cochlea-modeling business about whether 1D, 2D, or 3D models are good enough to capture the essence of the physics. Independent of these arguments, if the results of the model are needed at only a sequence of places along one dimension (such as at the inner-hair-cell locations), then we can represent the results of the 2D or 3D hydrodynamics by a 1D model system, and can do so by using transfer functions, more economically than by modeling the fluid motion directly.

One key property of cochlear physics on which we must rely for this approach is the unidirectionality of energy flow. Under normal conditions, sound energy enters the ear, propagates, and is absorbed without causing significant energy to reflect and propagate back out [2]. This condition is the one that we shall model; we discuss exceptions in Section 1.9.

The method known as Liouville–Green (LG) or Wentzel–Kramers–Brillouin (WKB, or WKBJ with Jeffreys) give us easy insight into wave propagation in nonuniform media such as the cochlea. This method says that, if a wave is propagating from the input along one dimension, then we can get the response from the input to any point by composing the relative responses from each point to the next along that dimension, using local parameters as though the medium were uniform.

The mathematics is most easily expressed in terms of a wave description in which the local (i.e., at any particular place) wave-propagation properties are characterized by a complex *wavenumber*. To make life simple, we consider

only one frequency at a time. To simplify the analysis further, we use complex numbers as the values of waves, realizing that we can easily constrain wave values to real numbers later by adding pairs of complex waves in a complex-conjugate relationship.

In a uniform medium, a wave propagating toward increasing values of the place dimension is given by

$$W(x) = A \exp(i\omega t - ikx)$$

where A is the amplitude, ω is the frequency, k is the wavenumber, t is time, and x is place. The wavenumber depends on frequency via the physics of the medium; we can write it as the function $k(\omega)$ —the solution of the physical constraints known as the *eikonal* or the *dispersion relation* of the system.

We can think of the wavenumber as the spatial frequency of the wave, in radians per meter in MKS units. If k is real, then the wave described simply propagates with no change in amplitude, with a wavelength of $2\pi/k$, at a velocity ω/k . If k has a nonzero imaginary component, however, then the wave can decay in amplitude as it propagates (i.e., in a passive or attenuating medium, for a negative imaginary part) or increase as it propagates (i.e., an active amplifying medium, for a positive imaginary part). Zweig and colleagues [29] presented an analysis of the 1D long-wave approximation to the cochlea with resonant basilar membrane (BM), using the complex radian wavelength λ (*lambda-bar*, the reciprocal of the wavenumber).

By examining the ratio of waves at two places separated by a distance Δx , we see that the wave at the farther place is equal to the wave at the nearer place multiplied by $\exp(-ik\Delta x)$, representing a frequency-dependent *filter* characterizing the stretch of length Δx .

In a nonuniform medium, there is no single wavenumber for a frequency, and possibly certain regions amplify a particular frequency while others attenuate it. Under reasonable conditions, however, each point in such a medium (i.e., along the place dimension) can be characterized by a local wavenumber, as though it were part of a uniform medium. The LG approximation then says that a wave propagating an infinitesimal distance dx through that place will be multiplied by $\exp(-ikdx)$ (and possibly also by a real-valued factor near 1, if a constant amplitude does not correspond to a constant power as the parameters of the medium change—but let's neglect that factor for now).

Now consider wavenumber as a function of both frequency and place: $k(\omega, x)$. Within the approximation of the LG method, this function completely characterizes wave propagation in the nonuniform medium along the x dimension. To see what happens between points far apart, we can break the medium into infinitesimal segments of length dx , and can multiply together all the factors for those segments. The factors are exponentials, and the product of exponentials is the exponential of a sum, so the resulting product is the exponential of an integral along the x dimension:

s a useful
with the
stractly,
ar place
nse only
only one
three di-
l 3D and
unknown
odeling
ature the
s of the
(such as
the 2D
transfer
tly.
pproach
d energy
t energy
we shall
Brillouin
gation in
wave is
response
om each
ough the
cription
roperties
consider

$$H = \exp\left(-i \int_{x_a}^{x_b} k dx\right)$$

This complicated-looking frequency-dependent gain and phase factor H is the LG method's representation of the transfer function between points x_a and x_b in a nonuniform medium; it is a generalization of the transfer function $\exp(-ik\Delta x)$ that characterizes a stretch of a uniform medium.

The factor H is still just a linear filter in the usual signal-processing sense. Furthermore, we can factor this filter into a product, or cascade, of several filters by splitting the interval of integration (from x_a to x_b) into N small steps:

$$H = \prod_{j=1}^N \exp\left(-i \int_{x_{j-1}}^{x_j} k(\omega, x) dx\right)$$

Any number and size of steps leads to a factorization, but, if the steps are small enough, then each individual filter will be well approximated from a local wavenumber by $\exp(-ik\Delta x)$, where $\Delta x = (x_b - x_a)/N$ is the step size, making it easier to tie the filters directly to a model of the underlying wave mechanics:

$$H \approx \prod_{j=1}^N \exp(-ik(\omega, x_j)\Delta x)$$

Therefore, independent of the details and dimensionality of the underlying wave mechanics, the responses of the cochlea at a sequence of places are equivalent to the responses at the outputs of a sequence of cascaded filters. The LG method constrains the design of those filters when the underlying physics is known.

How does the filter relate to the wavenumber? For a given value of Δx in a uniform medium, the filter and the wavenumber are in 1-to-1 correspondence via the complex exponential. For a given pair of places x_a and x_b in a nonuniform medium, the filter is determined uniquely by the function $k(\omega, x)$, although the inverse is not necessarily true (i. e., a different $k(\omega, x)$ with the same integral on that interval, such as a spatial reversal of $k(\omega, x)$, also would be a solution).

Even for nonlinear and time-varying wave mechanics, we can reasonably assume that a nonlinear and time-varying filter cascade will be a useful structural analog and a fruitful modeling approach: that of modeling local behavior with local circuits. The approach is neuromorphic in the sense that it is based on the form of wave propagation present in this peripheral part of the sensory nervous system.

If the cochlea's frequency-to-place map is approximately logarithmic, and we model equal place increments with filter stages, then the model stages will have characteristic frequencies (or time constants) in nearly a geometric sequence. We often assume a geometric sequence in model calculations, but the method is more flexible and can be used to match realistic cochlear maps in which the low-frequency region maps nearly linearly to place.

1.3 POWER FLOW AND ACTIVE GAIN

The LG method goes one step further than we just described in providing techniques to compute a slowly space-varying amplitude factor to account for the varying relationship between wave amplitude and power in a system with nonuniform energy-storage mechanisms. For example, if BM stiffness is varying, the proportionality between squared volume displacement and potential energy is changing; the amplitude of a BM displacement wave needs to be adjusted accordingly. For our purposes, we shall typically jump up one level of abstraction by imagining that our wave amplitudes are given in terms of derived variables, such that constant amplitude corresponds to constant power. Therefore, the gain of the filters will be exactly 1 in regions that are passive and lossless, as is typical of extremely low frequencies, relative to the characteristic frequency (CF) of the cochlear place under consideration.

For more specific information on hydrodynamic modeling and the LG method, relative to analog VLSI implementations and numerical methods, see the dissertation of Lloyd Watts [26].

We have not yet said how the filters in a cascade should look—we have said only that the design can be constrained by models at a lower level. We discuss specific filters in Section 1.5. In general, we expect cochlear filters to be passive and linear for low frequencies, to provide active gain or power amplification for frequencies near CF, and to attenuate high frequencies. Therefore, filters of the class of interest have unity gain at DC, followed by a gain somewhat greater than unity, and a high-frequency gain less than unity.

If the model has many such filters in cascade, then the individual filter stages do not need to have gains far from unity for the cascade to achieve an aggregate pseudoresonant [5] response, with a high peak gain and a large high-frequency attenuation.

The notion of a pseudoresonance differs in a fundamental way from that of a resonance, with which engineers generally are familiar. A pseudoresonance is a broadly tuned gain bump that results from the collective behavior of a large number of broadly tuned (and hence low-precision) stages (or poles), or of a distributed system. A resonance, on the other hand, becomes narrowly tuned and needs high precision to achieve a high gain at its center. Reliance on such a collective computation is another hallmark of the neuromorphic approach.

1.4 WIDE-DYNAMIC-RANGE COMPRESSION VIA FILTER CASCADES

One of the most important nonlinear functions of the cochlea is the compression of a wide range of sound intensities into a narrower range of cochlear-motion intensities at the sensor array, for frequencies near CF. Studies of cochlear mechanical response since about 1970 [19] have repeatedly demonstrated this frequency-dependent compression in live cochleae, and its absence in dead cochleae.

In live cochleae, the overall input-output intensity curves for frequencies near CF have a slope of typically 0.25 to 0.5 on a log-log plot [20]. This reduced slope is known as a *compressive nonlinearity*. The exact slope depends on the quality of the experimental preparation, on the frequency and intensity range, and on whether the response is measured at a fixed frequency or at the frequency of greatest response, which shifts a little with level. A higher compression (slope 0.25, or 4-to-1 compression) is more typical at a fixed frequency at or above the most sensitive frequency (CF), and a lower compression (slope 0.5, or 2-to-1 compression) is more typical at the peak response frequency, as the peak moves to frequencies below CF at higher sound levels.

From our filtering point of view, we need level-dependent filters whose gains decrease as the signal level increases, to model this mechanical compression. Equivalently, we expect that the imaginary part of the wave number will change with level, even changing its sign between damping and amplification for some combinations of frequency and place.

Presumably, the dependence of wavenumber on level comes from nonlinearities in the biomechanics, including the outer hair cells, which are the presumed source of the energy needed to provide active gain. These mechanical changes modify the way that traveling waves pick up energy, and the resulting cascaded filters that model a set of different places are a reflection of the underlying wave mechanics. Therefore, a filter-cascade model can, in principle, exhibit a range of behaviors similar to those of the mechanical system.

As a wave picks up energy in traveling across a range of places, each little increment of place needs to contribute only a small amount of gain. If the filter-cascade model has stages that model small Δx regions, then each filter will need to contribute only a small gain; as the overall gain changes, each filter will have to change only slightly.

A power gain, or a filter gain greater than unity, is correlated in this view with an active process that we think of as providing an active undamping—effectively a negative viscosity. But even if we do not rely on this notion of a literal power gain, the variable-gain filter-cascade structure provides a qualitative functional model of the variable-gain behavior observed in the cochlea—it could be adapted to fit the wave mechanics of a passive model. Indeed, the basis for our first use of the filter-cascade technique [11] was a passive long-wave analysis [29].

In our earlier model [11], we added the gain variation after the filters as a functional afterthought, so the model did not have the right frequency-dependent properties, such as linearity at low frequencies. Because it incorporates gain variation directly into the cascade as filter-parameter (Q) variation, the filter-cascade approach inherently achieves a reasonable constraint on how the overall filter gain can vary with frequency and place: The different places share most of the same cascade filters. That is, it is not possible to have a high peak gain at one place and a low peak gain at a nearby place, even if the cascaded filters vary substantially, because the composite responses at nearby places share most of the same filters. This property arises because we

are modeling a wave propagation directly, again illustrating the benefit of a neuromorphic approach.

1.5 FITTING A FILTER CASCADE TO THE COCHLEA

For most models of the mechanics of cochlear wave propagation, the qualitative behavior of a stage is just this: the filtering provides a gain bump for frequencies “near” CF, and provides attenuation at higher frequencies. What are the simplest lumped-parameter filters (i.e., small sets of poles and zeros in Laplace or Z-transform spaces) that can model this qualitative behavior? How significantly do the details of the stage filter affect the overall pseudoresonant response of the cascade? We answer these questions using examples.

The simplest stage is a two-pole filter, which we refer to as a *second-order section (SOS)*. The SOS as commonly used in digital signal processing might have either just two poles and no zeros in the simplest case, or might include 1 or 2 zeros in a higher-complexity alternative. We have focused on the simpler all-pole version in recent years [27, 18], whereas our earlier work used both poles and zeros [15, 11, 12, 16, 23].

An active-undamping approach to a physical basis for a wavenumber solution [18] led us to believe that simple two-pole filters may be not sharp enough (have narrow enough relative bandwidth) to be realistic, and that a three-pole filter would be a closer match. But two poles and two zeros can accomplish the same sharpening if that is what we need to fit experimental or theoretical data. In spite of this sharpness discrepancy relative to our particular mechanical model, we see the two-pole filter as a good and useful model of cochlear function. We should not rule out this simplest approach without a compelling reason.

Figure 1.2 shows a composite pole-zero diagram representing four alternative designs for a single filter stage: two-pole, three-pole, two-pole/two-zero, and a sharper two-pole/two-zero designs. The corresponding stage transfer-function gains and group delays are shown in Figures 1.3 and 1.4, respectively. The two-pole/two-zero designs have sharper drops just beyond CF than the all-pole designs, but then level out at some gain less than 1, rather than continuing to drop.

Figure 1.5 shows the composite gains of long cascades of geometrically spaced stages, and Figure 1.6 shows the corresponding total group delays. Note that we can sharpen the two-pole response by adding either an extra pole or a pair of zeros; the resulting cascade gains can be made similar, but adding poles adds to the delay, whereas adding zeros reduces the delay. This dimension of flexibility may be useful if we wish to match the model to a delay or phase measurement.

Moving the zeros closer to the poles and closer to the imaginary axis in the s plane results in more sharpening, especially of the high side, of the overall filter. This configuration fits the *notch* transfer function of a long-wave mechanical approximation, and was the basis of our original cochlea model [11]. We now




Figure 1.2 Composite pole-zero diagram. Four different filter designs are specified within this composite diagram, so that the pole and zero positions can be compared. Since complex poles and zeros are inherently in a complex-conjugate relationship, we label several in the top half-plane and several in the bottom half-plane to reduce clutter.




Figure 1.3 Stage transfer-functions gains. For each of the four filter designs of Figure 1.2, the magnitude of the stage transfer function is plotted.

believe that that model was too sharp, due to the unsuitability of the long-wave approximation for modeling the real cochlea near CF. At the time, we used too-sharp filter models by trying to match transfer functions to iso-rate

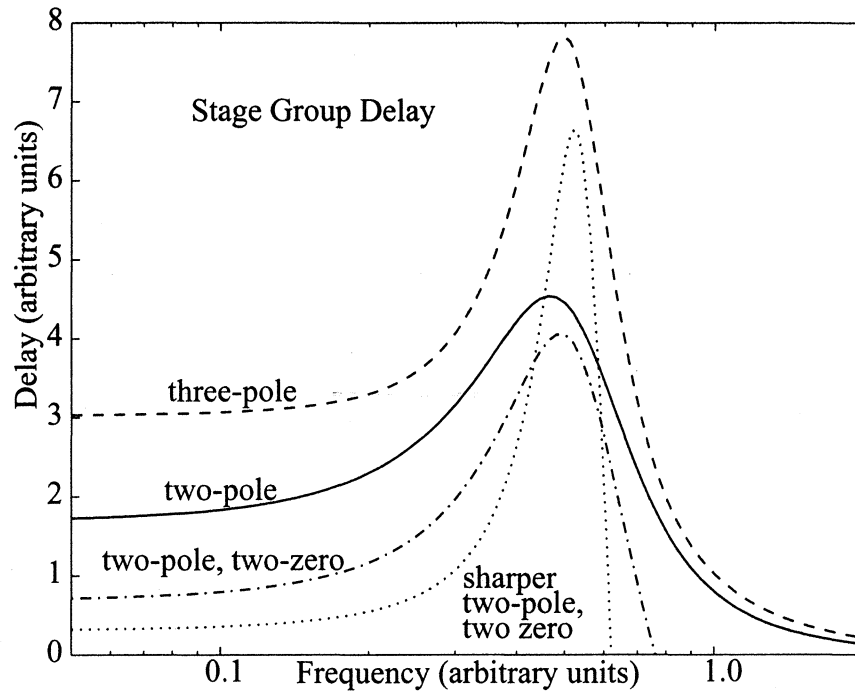


Figure 1.4 Stage group delays. For each of the four filter designs, the group delay of the filter stage is plotted in arbitrary units.

tuning curves; that approach is clearly inappropriate, as we have explained subsequently [13].

1.6 CASCADE-PARALLEL FILTERBANKS

Our original filter model [11] was based on the longwave approximation to cochlear mechanics, in which a significant membrane mass leads to a true local resonance. We used a cascade of notch filters (two-pole/two-zero) to model pressure-wave propagation, plus a resonator at each tap of the cascade to convert pressure to BM displacement or velocity locally. This cascade-parallel architecture may still be a useful way to separate the propagated variable from the sensed variable, and possibly to simplify the required filters. For the purposes of that model, however, we noticed that, by constraining the parameters and rearranging some poles, we could easily convert the structure to a roughly equivalent pure cascade version, saving complexity and computation [12].

Figure 1.7 shows the cascade-parallel structure; contrast it with Figure 1.1. Notice that the output taps of the cascade-parallel structure are still always related by a relative transfer function, such that a pure cascade equivalent version exists, although stability of the exactly equivalent pure cascade is not ensured unless the parallel filters have stable inverses.

d within
complex
al in the

gure 1.2,

ie long-
ime, we
iso-rate




Figure 1.5 Aggregate transfer-functions gains. For each of the four filter designs, the gain of a long geometrically-spaced cascade is plotted. The scale is arbitrary, because it depends on the density of stages per octave.

1.7 NONLINEAR EFFECTS

A filter cascade can have an overall strongly compressive nonlinear input-output behavior, if the stages are weakly compressive. There are two general forms of nonlinearity that are important to consider, and it is likely that both operate in the real cochlea: instantaneous nonlinear distortion, and feedback control of a peak-gain parameter. For example, the nonlinear model of Kim and his colleagues [8] is just a cascade of two-pole filters with a compressive nonlinearity in each stage; Kim's later suggestions [7] are of the parameter-feedback form, and are motivated as a functional role for the auditory efferent system.

An instantaneous distortion nonlinearity, such as a hyperbolic tangent that puts a firm saturation limit on the amplitude out of each stage, leads to the generation of intermodulation products, such as the cubic-difference tone $2f_1 - f_2$. Distortion products are mostly generated where the primary components (f_1 and f_2 in this case) are large, and distortion products with frequencies below the primaries are then free to propagate farther to their own place. Frequencies above and below the primaries can sometimes be detected propagating back out of a real cochlea, but unidirectional cascades cannot model that effect.

As discussed in Section 1.4, we can use feedback of a detected output level to affect the filter parameters, implementing a less distorting amplitude com-

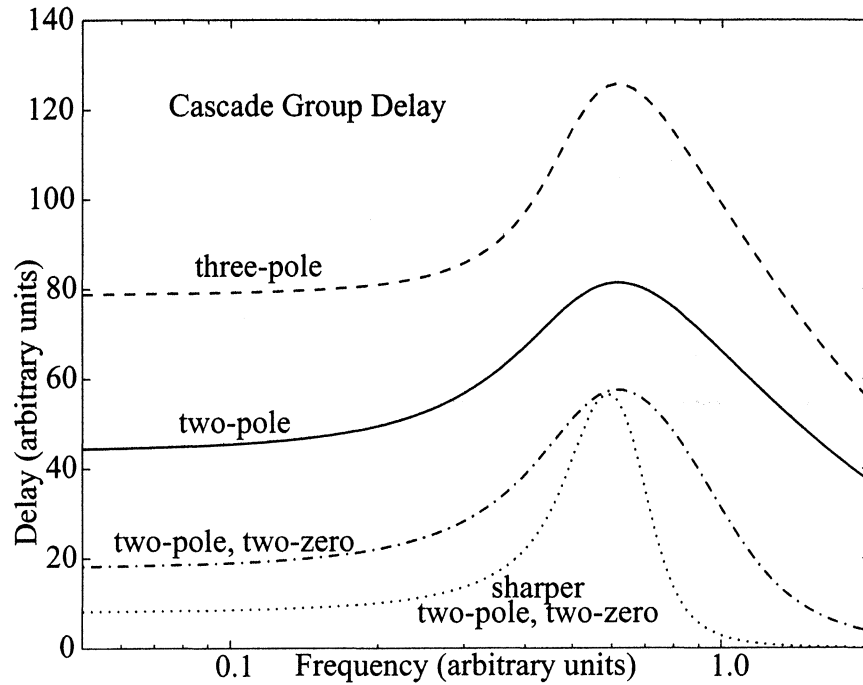


Figure 1.6 Aggregate group delays. For each of the four filter designs, the group delay of the cascade is plotted in arbitrary units.




Figure 1.7 The cascade-parallel structure. This cascade-parallel filter configuration can model the response of the cochlea with more flexibility than the pure cascade structure has.

pression known as *automatic gain control* (AGC). A small reduction in the pole Q of each stage in response to output energy can lead to a highly compressive overall response. This AGC is one purported function of the auditory efferent innervation: to tell the outer hair cells to turn down their level of activity [7]. Sound in one ear is even known to drive the efferent neurons to the contralateral cochlea, perhaps to keep the gains of the two ears more synchronized than they would be otherwise, and thus to aid the brain in binaural level comparisons [1].

1.8 SILICON VERSUS COMPUTER MODELS, AND PRACTICAL PROBLEMS

We would need a high-power programmable processor to implement a filter-cascade model of the cochlea in real time. Dedicated silicon implementations, on the other hand, can be made with less silicon and much lower power consumption [27, 12, 27]. An important cost factor in analog or digital sampled-data implementations is that avoidance of aliasing in the nonlinear operations that follow the cochlea model requires a substantial oversampling factor. By avoiding clocks, sampling, and high-speed circuits, the continuous-time analog approach yields by far the lowest power consumption, but requires novel solutions to noise, matching, tuning, stability, and communication problems [3, 6, 11, 25, 24, 25, 32, 90]; see also the other papers on neuromorphic analog cochleae in this book.

1.9 LIMITATIONS

As we mentioned in Section 1.7, distortion products can propagate backward out of a real cochlea, but not out of a unidirectional filter-cascade model. This limitation applies to other otoacoustic emissions as well, both stimulated and spontaneous. So the filter cascade is not a suitable modeling substrate for such effects.

The filter-cascade model is based on looking at a set of points along only one dimension, and as such provides no direct help in our understanding the motion of other parts of the cochlea or in the cochlea's fluid. Micromechanical models, 2D and 3D models, and other modeling approaches can help to inform the design of the filter cascade, but the filter cascade then captures only a slice of the more detailed models.

Any small stretch of cochlear transmission line acts approximately as the filter $\exp(-ik\Delta x)$; but does this filter, or an approximation to it, have useful properties, such as stability or causality? The filter specification derived from the wavenumber tells us what happens at all sine frequencies; to address stability, however, we need to consider approximate filter models with poles and zeros. We believe, but have not proved, that, if the mechanical wave system is stable, then stable rational filters exist that are reasonable approximations to the system's frequency response.

Causality is more complicated, because the response at a point is not physically caused by only the action at a different point upstream, even under the unidirectional assumption; rather, it is caused by the combined actions of nearby points in the whole 2D or 3D motion of the medium. The resulting filter, or approximations to it, could conceivably show *precursors* in response to an impulse. So, if we design a filter with the right magnitude frequency response, causality may force the filter to have excess delay if the Δx value is short compared to the wavelength. Thus, fine division of the place dimension may make it increasingly difficult to get the phase right in low-order filter approximations—especially in the case of causal all-pole filters. Adding zeros

helps us to cancel some of the delay of the poles, thus making it easier for us to develop a model with reasonable phase. Kates [6] has explored one class of filter cascades using zeros to arrive at lower overall delay.

The LG method breaks down in the cochlea in the *cutoff region*, where the eikonal has multiple complex solutions for k . In this region, the wavenumber changes so rapidly that there is effectively a *mode coupling* phenomenon that allows energy to couple into several of these different wave modes, which interfere with each other in complex ways, and which decay more slowly with x than does the original mode [26]. The resulting high-frequency irregularity, or plateau, in the response gain, which is found in numerical 2D solutions and is sometimes observed in real cochleae, is not easily modeled by cascades of simple filters. This discrepancy is an obstacle not to the concept of a filter cascade, but rather to the modeling details and to the desire to use simple stages. It seems likely, however, that the high-frequency plateau has no functional importance in normal hearing. Complicated response patterns due to cochlear micromechanics may lead to similar considerations, depending on one's goals in modeling.

1.10 RELATION TO OTHER APPROACHES

The most common functional approach to computational models of the cochlea is the bandpass filterbank. In this approach, every place to be modeled has its own filter, which is designed to match experimental data. Because there is usually no good basis for constraining a filter design using poles, an all-zero (transversal or finite impulse response) filter is often employed. Both of these features—independent filters and lack of poles—make the implementation of this approach computationally expensive.

Filterbanks that use poles—such as the gammatone filterbank (GTF) and its all-pole variant (APGF) [14, 23]—are becoming more widely used, because of their efficiency and simpler parameterization. The GTF is popular, but has an inappropriate symmetric passband; the APGF is closely related to a cascade of two-pole filters, and is therefore much more realistic in terms of transfer function and of the possibility of parametric nonlinearity.

An analog silicon model of the cochlea that can propagate waves bidirectionally has been reported by Watts [27]. It uses a 2D resistive grid as a substrate for directly solving Laplace's equation for wave propagation in a 2D fluid model of the cochlea, with second-order filters along one edge modeling the BM-fluid interaction. This approach needs to be further developed to see whether it leads to an overall advantage in implementing an effective cochlea model. A potential problem is that irregularities or the inherent spatial discretization may lead to reflections that cause instability, as has sometimes been a problem in 2D numerical solutions of active cochlea models.

1.11 CONCLUSIONS

The filter-cascade structure for an cochlea model inherits two key advantages from its neuromorphic roots: efficiency of implementation, and potential realism. Both the filter transfer functions, in terms of magnitude and delay dispersion, and the nonlinear behaviors of the cochlea, in terms of distortion and adaptation, are modeled realistically under the constraints imposed by the cascade. Minor problems, such as excess total delay in the finely discretized all-pole version, are tolerable in practical applications.

Analog VLSI implementations of the filter-cascade cochlea model are currently being explored at a number of laboratories around the world. The ideal of a practical micropower real-time artificial cochlea circuit is rapidly coming closer to reality.

Acknowledgments

I thank Malcolm Slaney, Lloyd Watts, John Lazzaro, Rahul Sarpeshkar, and Carver Mead for their contributions to the work discussed. And I thank Tor Sverre Lande for encouraging me to write this paper, and Lyn Dupré for skillfully editing several drafts of the manuscript.

References

- [1] J. F. Brugge. An overview of central auditory processing. In A.N. Popper and R.R. Fay, editors, *The Mammalian Auditory Pathway: Neurophysiology*, pages 1–33. Springer-Verlag, 1992.
- [2] E. de Boer and R. MacKay. Reflections on reflections. *J. Acoust. Soc. Am.*, 57:882–890, 1980.
- [3] P. Furth and A. B. Andreou. A design framework for low power analog filter banks. *IEEE Transactions on Circuits and Systems, Part I: Fundamental Theory and Applications*, 42:966–971, 1995.
- [4] P. Furth and A. B. Andreou. Linearized differential transconductors in sub-threshold CMOS. *IEEE Electronics Letters*, 31(7):545–547, March 1995.
- [5] M. Holmes and J. D. Cole. Pseudo-resonance in the cochlea. In Boer E. de and Viergever M.A., editors, *Mechanics of Hearing*. Martinus Nijhoff Publishers, the Hague, 1983.
- [6] J. M. Kates. A time-domain digital cochlear model. *IEEE Trans. Signal Processing*, 39:2573–2592, December 1991.
- [7] D. O. Kim. Functional roles of the inner- and outer-hair-cell subsystems in the cochlea and brainstem. In Berlin C., editor, *Hearing Science*, pages 241–261. College-Hill Press, San Diego, 1984.
- [8] D. O. Kim, C. E. Molnar, and R. R. Pfeiffer. A system of non-linear differential equations modeling basilar membrane motion. *J. Acoust. Soc. Am.*, 54:1517–1529, 1983.

- [9] J. Lazzaro, J. Wawrzynek, M. Mahowald, M. Sivilotti, and D. Gillespie. Silicon auditory processors as computer peripherals. *IEEE Journal of Neural Networks*, 4(3):523–528, 1993.
- [10] W. Liu, A. Andreou, and M. Goldstein. Voiced-speech representation by an analog silicon model of the auditory periphery. *IEEE Transactions of Neural Networks*, 3(3):477–487, 1992.
- [11] R. F. Lyon. A computational model of filtering, detection and compression in the cochlea. In *Proc. IEEE Intl. Conf. on Acoust. Speech and Signal Proc.*, pages 1282–1285, Paris, 1982.
- [12] R. F. Lyon. Computational models of neural auditory processing. In *Proc. IEEE Intl. Conf. on Acoust. Speech and Signal Proc.*, San Diego, 1984.
- [13] R. F. Lyon. Automatic gain control in cochlear mechanics. In P. Dallos et al., editor, *The Mechanics and Biophysics of Hearing*, pages 395–402. Springer-Verlag, 1990.
- [14] R. F. Lyon. All-pole auditory filter models. In E. Lewis, editor, *Diversity in Auditory Mechanics*. World Scientific, In press.
- [15] R. F. Lyon and L. Dyer. Experiments with a computational model of the cochlea. In *Proc. IEEE Intl. Conf. on Acoust. Speech and Signal Proc.*, pages 1975–1978, Tokyo, 1986.
- [16] R. F. Lyon and N. Lauritzen. Processing speech with the multi-serial signal processor. In *Proc. IEEE Intl. Conf. on Acoust. Speech and Signal Proc.*, Tampa, 1985.
- [17] R. F. Lyon and C. Mead. An analog electronic cochlea. *IEEE Trans. Acoust., Speech, Signal Processing*, 36:1119–1134, July 1988.
- [18] R. F. Lyon and C. Mead. Cochlear hydrodynamics demystified. Caltech Computer Science Technical Report Caltech-CS-TR-88-4, Caltech, 1989.
- [19] W. S. Rhode. Observations of the vibration of the basilar membrane in squirrel monkeys using the Mössbauer technique. *J. Acoust. Soc. Am.*, 49:1218–1231, 1971.
- [20] L. Robles, M. A. Ruggero, and N. C. Rich. Basilar membrane mechanics at the base of the chinchilla cochlea. input-output functions, tuning curves and response phases. *J. Acoust. Soc. Am.*, 80:1364–1374, 1986.
- [21] R. Sarpeshkar, Lyon R. F., and C. A. Mead. An analog VLSI cochlea with new transconductance amplifiers and nonlinear gain control. In *Proc. IEEE Intl. Conf. on Circuits and Systems*, volume 3, pages 292–295, Atlanta, May 1996.
- [22] R. Sarpeshkar, Lyon R. F., and C. A. Mead. Nonvolatile correction of q-offsets and instabilities in cochlear filters. In *Proc. IEEE Intl. Conf. on Circuits and Systems*, volume 3, pages 329–332, Atlanta, May 1996.
- [23] M. Slaney. Lyon's cochlear model. Apple Technical Report 13, Apple Computer, Cupertino, CA, 1991.

- [24] C. Summerfield and R. F. Lyon. ASIC implementation of the lyon cochlea model. In *Proc. IEEE Intl. Conf. on Acoust. Speech and Signal Proc.*, San Francisco, 1990.
- [25] A. van Schaik, E. Fragnière, and E. A. Vittoz. Improved silicon cochlea using compatible lateral bipolar transistors. In David S. Touretzky, Michael C. Mozer, and Michael E. Hasselmo, editors, *Advances in Neural Information Processing Systems*, volume 8, pages 671–677. The MIT Press, 1996.
- [26] L. Watts. *Cochlear Mechanics: Analysis and Analog VLSI*. Ph.d. dissertation, California Institute of Technology, 1993.
- [27] L. Watts, Lyon R. F., and Mead C. A bidirectional analog VLSI cochlear model. In C. Sequin, editor, *Advanced Research in VLSI, Proceedings of the 1991 Santa Cruz Conference*, pages 153–163, Cambridge, MA, 1991. MIT Press.
- [28] L. Watts, D. Kerns, R. F. Lyon, and C. Mead. Improved implementation of the silicon cochlea. *IEEE Journal Solid-State Circuits*, 27(5):692–700, May 1992.
- [29] G. Zweig, R. Lipes, and J. R. Pierce. The cochlear compromise. *J. Acoust. Soc. Am.*, 59:975–982, 1976.

NEUROMORPHIC SYSTEMS ENGINEERING

Neural Networks in Silicon

edited by

Tor Sverre Lande
University of Oslo
Norway



KLUWER ACADEMIC PUBLISHERS
Boston / Dordrecht / London

E
G
ty

Harjani,

g, ISBN:

IG A/D

50-9
LS, Ron

KS, Sied

Iaddara,

Serdijn,

ormance

! Analog

CUITS,

0-7923-

H. Alan

entation,

SBN: 0-

, Ronald

Michiel

rence P.

ie, Johan

CUITS,

in Baher,

3-9272-8