CHAPTER

9

FOLLOWER–INTEGRATOR CIRCUIT

In Chapter 4, we noted that neural processes are insulated from the extracellular fluid by a membrane only approximately 50 angstroms thick. The capacitance of this nerve membrane serves to integrate charge injected into the dendritic tree by synaptic inputs. Much of the real-time nature of neural computation is vastly simplified because this integrating capability is used as a way of storing information for short time periods—from less than 1 millisecond to more than 1 second. There is an important lesson to be learned here, an insight that would not follow naturally from the standard lore of either computer science or electrical engineering. Like the spatial smoothing performed by resistive networks in Chapter 7, temporal smoothing is an essential and generally useful form of computation.

In CMOS technology, the elementary temporal-smoothing circuit is the **follower–integrator circuit** shown in Figure 9.1. It is the most universally useful of all time-dependent circuits.

The circuit consists of a transconductance amplifier connected as a follower, with its output driving a capacitor. As usual, we set the transconductance G with the bias voltage V_{b}. The rate at which the capacitor charges is proportional to the output current of the follower:

$$C\frac{dV_{\mathrm{out}}}{dt} = I_{\mathrm{b}} \tanh\left(\frac{\kappa(V_{\mathrm{in}} - V_{\mathrm{out}})}{2}\right) \tag{9.1}$$

where the voltages are measured in units of kT/q.

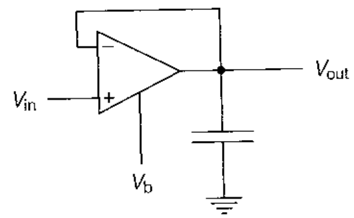

FIGURE 9.1 Schematic of the follower–integrator circuit. The current into the capacitor is proportional to the difference between V_{in} and V_{out}. The rate at which V_{out} is able to respond to changes in V_{in} is set by the transconductance of the amplifier.

First, we will examine the behavior of this circuit for signals that have small time-dependent deviations from some steady (quiescent) value. We will restrict the quiescent input voltage to the range within which the amplifier is well behaved, as described in Chapter 5. Under these conditions, we can analyze the circuit using the linear-systems-theory approach of Chapter 8. We will then compare the results with those we obtained for the RC integrator of Chapter 8. We will consider the large-signal behavior of the circuit at the end of the chapter.

SMALL-SIGNAL BEHAVIOR

For small signals, the tanh can be approximated by its argument, and Equation 9.1 becomes

$$C\frac{dV_{\text{out}}}{dt} = G(V_{\text{in}} - V_{\text{out}}) \tag{9.2}$$

Equation 9.2 can be rewritten in s-notation as

$$\frac{V_{\text{out}}}{V_{\text{in}}} = \frac{1}{\tau s + 1} \quad \text{where} \quad \tau = \frac{C}{G} \tag{9.3}$$

The response of the circuit of Figure 9.1 to a step input is shown in Figure 9.2. It can be compared with that of the RC integrator, which is shown in Figure 9.3. The two responses are not distinguishably different for the small (approximately 40-millivolt) signals used.

Using the principle of linear superposition of Equation 8.28 (p. 143), we can define precisely the temporal-smoothing properties of a single time-constant integration such as Equation 9.3

$$V_{\text{out}}(t) = \int_0^{\infty} V_{\text{in}}(t-\Delta)e^{-\Delta/\tau}\, d\Delta$$

The output at any time t is made up of the input for all previous times; the contribution of the input to the present output decreases exponentially with time into the past. In other words, the output is a moving average of the input, exponentially weighted by its timeliness. One could hardly ask for a more biologically relevant single measure of history—which is, no doubt, why this is the most ubiquitous computation in neural systems.

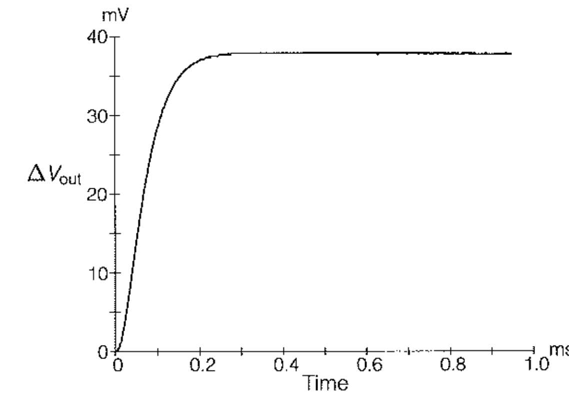

FIGURE 9.2 Measured response of the follower–integrator circuit to an approximately 40-millivolt step change in input voltage. The quiescent input voltage was approximately $V_{DD}/2$. The output is a temporally smoothed version of the input. Because the circuit is not linear, the form of the response will be different for large inputs (see Figure 9.12). We can adjust the time constant of the response over many orders of magnitude by setting the transconductance of the amplifier.

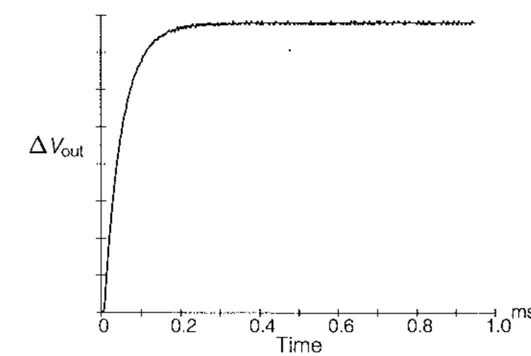

FIGURE 9.3 Measured response of the RC circuit of Figure 8.3 (p. 137) to step change in input voltage. Because the circuit is linear, the form of the response is independent of the size of the input.

COMPOSITION

The transfer function $V_{\text{out}}/V_{\text{in}}$ of the follower–integrator circuit is identical to that of the RC integrator of Chapter 8. You might think, therefore, that we have used a very difficult solution to a simple problem: We have replaced a simple resistor by a complicated transconductance amplifier. From the point of view of the transfer function, that is indeed what we have done; but a system has more properties than just its transfer function. There are important differences between the RC circuit and the follower–integrator circuit, which we can best

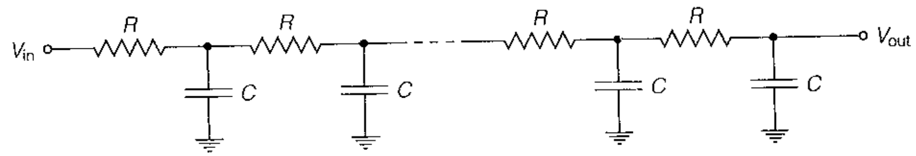

FIGURE 9.4 Delay line made up of a large number of RC-integrator sections in cascade.

appreciate by conducting an experiment. Let us take the RC integrator of Figure 8.3 (p. 137) and say, "If one is good, a lot more are better." We connect the output of one section to the input of another to form the RC delay line shown in Figure 9.4.

The response measured at the output of the first section of the line is shown in Figure 9.5. It is clear that the response of the first RC section has been changed from that of Figure 9.3 by the addition of the rest of the line. The reason is obvious: There is current flowing out of the first output through the second resistor, to charge up the rest of the line. This current must come from the input, and no particular node capacitance can get charged up until the capacitors between that node and the input get charged up. The waveform at the output of the first section of the line is much more sluggish than that out of the same single section with the rest of the line disconnected.

Transfer functions are most useful when their form is not changed by the environment in which they are used. Assume we have two circuit building blocks, the output of the first one feeding the input of the second. If we want to know the transfer function of this combined circuit, we just multiply the two transfer functions of the two individual circuits. The transfer functions are algebraic functions of s, and their product is the transfer function for the whole circuit. For a circuit to be an independent *module*, it must have this composition property.

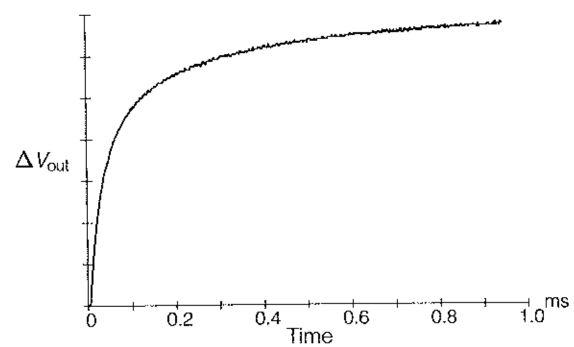

FIGURE 9.5 Measured response of the output of the first section of the RC delay line of Figure 9.4. The addition of subsequent stages has drastically changed the form of the response from that of a single RC integrator (Figure 9.3).

The RC circuit does not have this important property. Although that circuit *is* linear, it is *not* an independent module. The RC circuit looks simple, but in a system context it is actually extremely complicated; whenever we hook up one to another, we change what the first one does. So we cannot just multiply the transfer functions. The transfer function of the *composition* of the two circuits is not the composition of their individual transfer functions. In other words, the RC circuit does not have a simple *abstraction*.

Each input of the follower–integrator circuit is the gate of a transistor. The gate is isolated from the rest of the circuit by the gate oxide, which for all practical purposes is a perfect insulator. We can hook the input of such an amplifier to the output of any other circuit without drawing any current. We can obtain the transfer function of the composition by multiplying the two transfer functions, because we have not disturbed the first one by hooking up the second one. Using the amplifier instead of a resistor has bought us a clean abstraction of the smoothing function.

By throwing away the amplifier voltage gain A, we have obtained unity gain at very low frequencies to a high accuracy. We also have got much better control over the value of G, because G is directly related to the bias current in the amplifier, which we can control with a current mirror. The time constants are useful up to about 10 seconds. The follower–integrator first-order section works just like the RC integrator would have done had we put unity-gain amplifiers between every stage.

IMPLEMENTATION

The layout of a typical follower–integrator circuit is shown in Plate 8(a). The circuit consists of a wide-range amplifier driving a capacitor structure. In CMOS technology, the only excellent capacitor material we have available is the gate oxide. Unfortunately, the p- and n-type diffused areas do not extend under the polysilicon gate material, and hence we do not have a structure with good conductors on both sides of the thin gate oxide. Instead, we use a p-type transistor with its source tied to V_{DD}, and an n-type transistor with its source tied to ground. Because each transistor is biased above its threshold voltage, its capacitance is very nearly equal to the *oxide capacitance*—the same value as that of an ideal capacitor employing the gate oxide as its dielectric. If the gate voltage falls below the transistor threshold voltage, the capacitance falls rapidly. For the structure shown in Plate 8(a) to maintain a relatively constant capacitance, the voltage on the common polysilicon gate area should be kept away from both rails by at least the threshold voltage of the relevant transistor. This limitation is not much more severe than are the voltage limitations imposed by other circuits in a system.

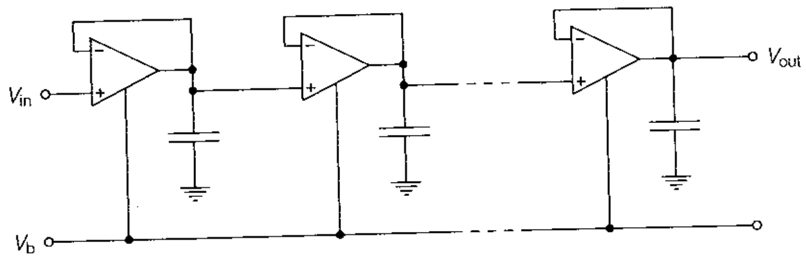

FIGURE 9.6 Delay line formed by connection of a large number of follower–integrator circuits in cascade.

DELAY LINES

We can create a delay line using follower–integrator first-order sections, as shown in Figure 9.6. The step response at the first few taps is shown in Figure 9.7. We can compare this figure with a similar plot for the RC line of Figure 9.4, shown in Figure 9.8. It is clear that the signal decays much faster in the RC line, because current must flow all the way from the input to the point at which the capacitor is being charged. In the follower–integrator line, the current required for charging each capacitor is supplied from the power supply, mediated by the transconductance amplifier.

Follower–Integrator Delay Line

Because of the modular nature of the sections, we can write the transfer function of the line up to the nth section as the product of the transfer functions of each of the individual sections:

$$\frac{V_{out}}{V_{in}} = \left(\frac{1}{\tau s + 1}\right)^n \tag{9.4}$$

To understand how Equation 9.4 represents a signal propagating along the line, we will pick a particular form for the input—a sine wave. We represent a sinusoidal signal of angular frequency ω by setting s equal to $j\omega$ in the transfer function. We need to understand what the nth power of a complex number represents, and how to take the inverse of a complex number. We can delay worrying about the inverse by considering V_{in}/V_{out} instead of V_{out}/V_{in}. We therefore write the inverse of Equation 9.4 for $s = j\omega$:

$$\frac{V_{in}}{V_{out}} = (j\omega\tau + 1)^n \tag{9.5}$$

We will be interested in lines with many sections. For such lines, the response at high frequency will decrease very steeply, falling with increasing frequency as $(1/\omega\tau)^n$. We are concerned with frequencies for which the line has measurable output, so we can safely assume that $\omega\tau$ is much less than 1. In Chapter 8, we

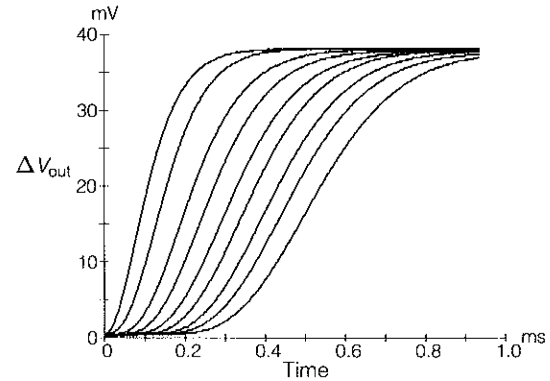

FIGURE 9.7 Measured response at the outputs of the first nine taps of the delay line of Figure 9.6. The spacing between any two curves is the delay of a particular section. The delay varies randomly due to transistor mismatch, as discussed in Chapter 3.

noted that the polar form of a complex number was the representation in which multiplication is a simple operation. In this regime, each term in Equation 9.5 can be approximated by the polar form

$$1 + j\omega\tau \approx \left(1 + \tfrac{1}{2}(\omega\tau)^2\right)e^{j\omega\tau} \tag{9.6}$$

Using Equation 9.6 and the rule for multiplying complex numbers given in Chapter 8, we can approximate Equation 9.5 for $j\omega\tau$ much less than 1:

$$\frac{V_{in}}{V_{out}} = \left(1 + \tfrac{n}{2}(\omega\tau)^2\right)e^{nj\omega\tau} \tag{9.7}$$

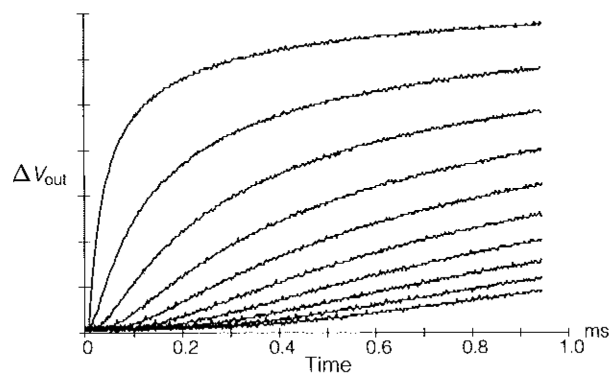

FIGURE 9.8 Measured response at the outputs of the first 10 sections of the RC delay line of Figure 9.4. The rise time of the waveforms increases rapidly with distance through the line.

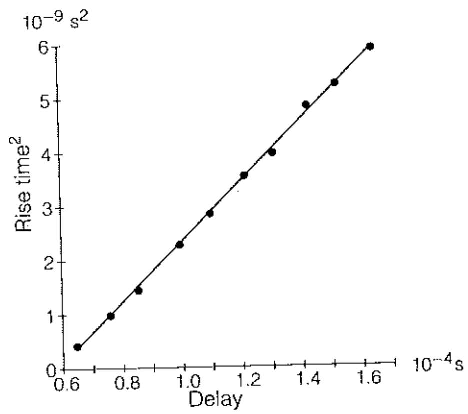

FIGURE 9.9 Linear relationship between the delay at several taps along a follower-integrator line and the square of the rise time at the same taps. The data were obtained from Figure 9.7 by the procedure described in the text.

Equation 9.7 leads to the following approximate form for the transfer function:

$$\frac{V_{\text{out}}}{V_{\text{in}}} \approx \frac{1}{1+\frac{n}{2}(\omega\tau)^2}e^{-jn\omega\tau} \tag{9.8}$$

The line acts as a phase delay of $n\omega\tau$ radians, corresponding to a time delay of $n\tau$ seconds—τ seconds per section. The signal magnitude is attenuated by $\frac{1}{2}(\omega\tau)^2$ per section. We can estimate the **bandwidth** of the line by finding the frequency where the response has decreased to one-half of the input amplitude. To the level of approximation used in Equation 9.8, the cutoff frequency ω_c is given by

$$\omega_c\tau = \sqrt{\frac{2}{n}} \tag{9.9}$$

The bandwidth of the line thus decreases as the square root of the number of sections, whereas the delay is linear with the number of sections.

The **rise time** to a step input is approximately the reciprocal of the bandwidth. We can measure the rise time of the curves of Figure 9.7 by extrapolating the steep part of the signal upward to its upper steady-state value, and downward to its initial value. The delay is the time when the output crosses the level midway between the two limiting values. The square of the rise time is plotted as a function of the delay in Figure 9.9. The result is an excellent straight line, as predicted by Equation 9.9. Note that the delay of the individual sections varies considerably, due to the mismatch among the bias transistors in the individual amplifiers. In spite of this variation, the relationship between delay and bandwidth is in excellent qualitative agreement with our analysis. We will see in Chapter 16 that the form of the behavior of much more sophisticated delay lines also is preserved despite a wide variation among the individual delay elements.

RC Delay Line

In Chapter 7, we analyzed the propagation of steady (DC) signals in a passive dendritic process. The signals were not varying with time, and hence we could neglect the membrane capacitance. We will now derive the behavior of such a process, including the effects of membrane capacitance on time-dependent signals.

We model the process by adding capacitors at every node, as shown in Figure 9.10. We can analyze the line by using the same approach by which we derived Equation C.5 (p. 340). The differential equations for current and voltage are

$$-\frac{\partial V}{\partial x} = IR \tag{9.10}$$

and

$$-\frac{\partial I}{\partial x} = C\frac{\partial V}{\partial t} + VG \tag{9.11}$$

where R is the axial resistance, G is the conductance to ground, and C is the capacitance, all given per unit length of line. Both x and I are taken to be positive pointing to the right. Differentiating Equation 9.10 with respect to x, and substituting into Equation 9.11, we can eliminate I, and thus we obtain the equation for $V(x,t)$:

$$\frac{\partial^2 V}{\partial x^2} = RGV + RC\frac{\partial V}{\partial t} \tag{9.12}$$

We recognize our old friend $RG = \alpha^2 = 1/L^2$, where α is the *space constant* and L is the *diffusion length* of the line, as defined in Equation 7.2 (p. 108).

Equation 9.12 is called the **diffusion equation**; it governs the time course of signal propagation in a dissipative passive medium where the stuff out of which the signal is made is stored in the medium, and the fraction lost as the signal propagates is proportional to the amount present. The passive dendrites of a neuron obey this equation: V is the voltage across the membrane, C is the capacitance of the membrane, R is the axial resistance of the cytoplasm, and G is the conductance to the extracellular fluid, all given per unit length. The flow of

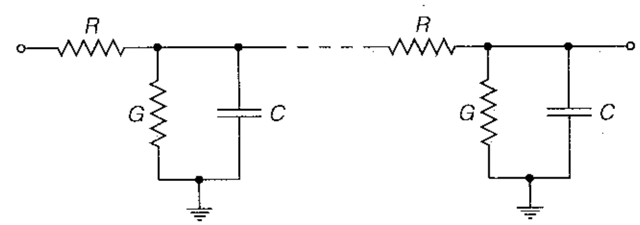

FIGURE 9.10 Network model of a dendritic process. The Rs represent the axial resistance of the cytoplasm, the Cs represent the membrane capacitance, and the Gs represent the membrane conductance.

heat in a medium also is governed by Equation 9.12: R is the thermal resistance, C is the heat capacity, and V is the temperature. The diffusion of minority carriers in a semiconductor follow Equation 9.12 as well; in that medium, G models the recombination process with the majority carriers. All these applications are rich sources both of mathematical treatments of the problem and of intuition concerning the nature of particular solutions. Treatments of the diffusion equation in one, two, and three dimensions, under various boundary conditions in time and space, are the subject of entire books [Carslaw et al., 1959]. Time-dependent electrotonic spread in neural processes was first discussed by Wilfrid Rall [Rall, 1960], and is treated at length in Jack, Noble, and Tsien [Jack et al., 1983]. We will derive only certain solutions of the one-dimensional problem of Equation 9.12, and will compare them with those for the follower–integrator line.

We use the method of separation of variables, assuming $V(x,t)$ can be expressed as

$$V(x,t) = X(x)T(t) \tag{9.13}$$

where X is a function of x alone, and T is a function of t alone. Substituting Equation 9.13 into Equation 9.12, we obtain

$$T\frac{\partial^2 X}{\partial x^2} = RGXT + RCX\frac{\partial T}{\partial t}$$

Dividing by XT, we obtain

$$\frac{1}{X}\frac{\partial^2 X}{\partial x^2} = RG + \frac{RC}{T}\frac{\partial T}{\partial t} = -\lambda^2$$

Because one side of the equation is a function of x alone, and the other is a function of t alone, both sides of the equation must be independent of x and t. Therefore, λ^2 must be a constant. The left-hand side then becomes

$$\frac{\partial^2 X}{\partial x^2} = -\lambda^2 X \tag{9.14}$$

and the right-hand side becomes

$$\frac{\partial T}{\partial t} = -\left(\frac{1}{\tau} + D\lambda^2\right)T \tag{9.15}$$

where $\tau = C/G$ is the **time constant** of the line, and $D = 1/(RC)$ is called the **diffusion constant** of the line. Because R and C are the values per unit length, the units of D are length^2/time. We have encountered Equation 9.15 before; it is of the same form as Equation 8.3 (p. 130). We know that its solutions are

$$T = e^{st} \tag{9.16}$$

Both λ and s are, in general, complex. Substituting Equation 9.16 into Equation 9.15, we obtain

$$s = -\left(\frac{1}{\tau} + \lambda^2 D\right) \tag{9.17}$$

Similarly, the solutions to Equation 9.14 can be written

$$X = e^{j\lambda x} \tag{9.18}$$

There are two special cases in which physically meaningful solutions can be obtained trivially: one when X is constant, and the other when T is constant. The first is obtained by setting λ equal to 0, in which case

$$T = e^{-t/\tau}$$

As we expect, when the voltage on the line is independent of x, it dies away exponentially with time constant τ.

The second case is obtained by setting s equal to 0 in Equation 9.17:

$$\lambda^2 = -\frac{1}{D\tau} \quad \text{or} \quad \lambda = \pm j\alpha \tag{9.19}$$

where $\alpha = \sqrt{RG} = 1/\sqrt{D\tau}$ is the space constant of the line, as defined in Equation 7.2 (p. 108). As we found in that analysis, when the line is driven by a DC source at the origin ($x = 0$), the voltage dies away exponentially with space constant α:

$$X = e^{-\alpha|x|} \tag{9.20}$$

The $+j$ root in Equation 9.19 corresponds to a signal dying out in the $+x$ direction, and the $-j$ root in Equation 9.19 represents a signal dying out in the $-x$ direction. These two roots give rise to the absolute value of x in Equation 9.20; the signal dies out as it propagates in either direction away from the source.

Now that we have a sanity check on the solutions of Equations 9.18 and 9.16, we can derive the response of the line to a sine-wave input. To compare the results directly with those for the follower–integrator line, we will treat the case where G is 0. Substituting $j\omega$ for s into Equation 9.17, we obtain

$$\lambda^2 = -\frac{j\omega}{D} \tag{9.21}$$

The square root of a complex number N has magnitude $\sqrt{|N|}$ and angle one-half that of N. In Equation 9.21, the magnitude of λ^2 is ω/D and the angle is 270 degrees; λ therefore will have magnitude $\sqrt{\omega/D}$ and angle 135 degrees. In terms of real and imaginary parts,

$$\lambda = (j-1)\sqrt{\frac{\omega}{2D}}$$

The solution for $V(x,t)$ can thus be written[1]

$$V = e^{j\lambda x}e^{j\omega t} = e^{-kx}e^{j\omega(t-x/v)} \tag{9.22}$$

[1] Equation 9.21 also has a $1-j$ solution, it is a wave propagating in the $-x$ direction.

where

$$k = \sqrt{\frac{\omega}{2D}} \qquad \text{and} \qquad v = \sqrt{2D\omega} \tag{9.23}$$

Equation 9.22 is the classic form for a wave traveling at velocity v, attenuated with space constant k. For any given point x along the line, the bandwidth can be defined as the **cutoff frequency** ω_c at which kx is equal to 1. From Equation 9.23,

$$\text{bandwidth} = \omega_c = \frac{2D}{x^2}$$

The delay is x/v, which at the cutoff frequency is $1/kv$.[2] From Equation 9.23,

$$\text{delay} = \frac{1}{kv} = \frac{1}{\omega_c} = \frac{x^2}{2D}$$

We can compare these results directly with those for the follower–integrator line given in Equations 9.9 and 9.8. In that case, the delay was linear in x, and the bandwidth decreased as $1/\sqrt{x}$. The delay of the RC line is quadratic in x, and the bandwidth is the inverse of the delay.

The rise time to a step input is approximately the inverse of the bandwidth at any point x. Thus, for the RC line, *the rise time is equal to the delay.* This behavior can be seen in the plots of Figure 9.8. The response curves become slower, but they still extrapolate back to near the origin; they never develop the delay-line behavior exhibited in Figure 9.7.

LARGE-SIGNAL BEHAVIOR

The large-signal response of the follower–integrator circuit is either not so nice or quite nice, depending on our point of view. We will look at the large-signal behavior of the circuit both in the time domain and in the frequency domain.

Transient Response

If we put a small step function of amplitude Δv into this circuit (Figure 9.11), the output responds as

$$V_{\text{out}} = \Delta v(1 - e^{-t/\tau})$$

The $e^{-t/\tau}$ term is the homogeneous solution. The "1" occurs because the DC value after the step is different from that before the step.

[2] A rigorous treatment of this problem is beyond the scope of this book. An excellent treatment of time-domain solutions to various forms of the diffusion equation is given in Carslaw and Jaeger [Carslaw et al., 1959].

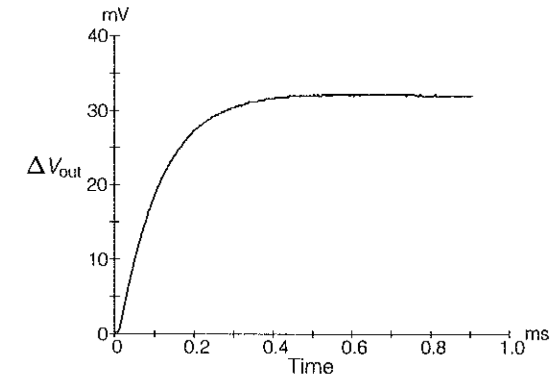

FIGURE 9.11 Response of the follower-integrator circuit to a small (approximately 30-millivolt) step-function input.

If we put in a big step (Figure 9.12), we get *tanh-ed.* Remember, these circuits can supply only a certain amount of current—the bias current I_b. Once the difference in the input voltages is larger than about 100 millivolts, the output is just a current source; the current I charges the capacitor C at a constant rate. Eventually, when the output gets close enough to its final value, the response approaches its final value as $e^{-t/\tau}$.

One way of looking at this behavior is with horror—there is no linear system anymore! If we double the input, we certainly do not get an output that is just scaled up by a factor of two. On the other hand, as the output gets close to its final value, the approach is just like the small signal response. The voltage just does not get there as fast because there is a limit to the maximum rate at which the output can charge its capacitor.

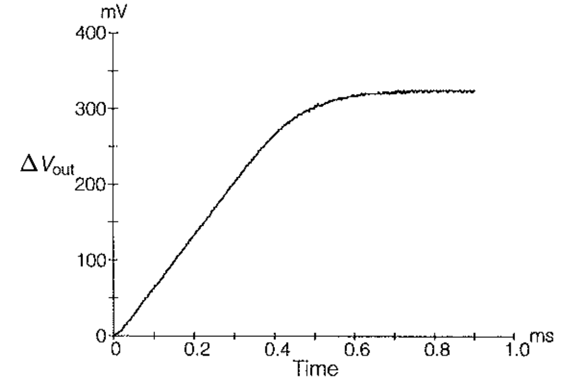

FIGURE 9.12 Response of the follower-integrator circuit to a large (approximately 300-millivolt) step-function input. The constant slope of the initial response is due to the current-limiting behavior of the transconductance amplifier.

If you have ever watched a plotter plot with a pen, you may have noticed that this system has the same property. There is a maximum speed at which the servo system that runs the pen can drive the motors. When the plotter is programmed to draw a shape in one corner of the page, and then to draw the next shape in the opposite corner, it goes zzzZZZZZZ.... There is a maximum tone the plotter makes when it tries to get to the other side; that tone represents the device's **slew-rate limit**.

Our amplifier is slew-rate limited, just like the plotter; in fact, every physical system has a slew-rate limit. There is an inevitable capacitance associated with any electrical node—in particular, with the output of our amplifier—and we can draw only a finite current out of any power supply. There are finite energy resources we can devote to getting from here to there. When a system reaches its slew-rate limit, it cannot accelerate beyond this speed. An automobile has a velocity limit at which the horsepower of the engine matches the drag due to the friction of the air; that is the car's cruising speed on a straight, deserted road. When you start a car from a stop light, it is acceleration limited. But when you are driving it across the desert, it is the slew-rate limit you are up against.

So every physical system has a slew-rate limit, but the follower–integrator circuit has it in spades—at 100 millivolts. That might seem to be a problem. The circuit is not a linear system even in the voltage range in which we are going to use it. Alternately, we can look at the low slew-rate limit as a fortunate factor—because we are going to build VLSI systems. If anything goes wrong somewhere (which it certainly will—somewhere), the amount of damage any one of these amplifiers can do is restricted. If one input gets stuck on, or if something at one spot is driving the system crazy, the magnitude of the damage one amplifier can cause is limited. So the slew-rate limit can be either a blessing or a curse, depending on how you look at it.

Frequency Response

Of course, the slew-rate limit also affects the time scale of the response. In a strict sense, a bandwidth is defined for only a linear system. On the other hand, if we think about bandwidth in a looser sense, it is reasonable to define one for the real system, but that bandwidth will be a function of the amplitude of the input. For small signals, we saw that the rise time for a step input was inversely proportional to the bandwidth. We can define an amplitude-dependent bandwidth that is the inverse of the rise time, for any input amplitude.

When the circuit is slew-rate limited,

$$C\frac{dV}{dt} = I$$

where I is the current that is set by the transconductance control. The solution under these conditions is just a straight line. The input step has amplitude Δv, so the time t it would take the output to get to Δv if it kept going at its maximum rate is

$$t = \frac{C}{I}\Delta v \tag{9.24}$$

That peculiar rise time depends on the size of the input signal—of course! The larger the signal, the longer it takes the output to reach its final value. If you go from Los Angeles to New York, it takes longer than if you go from Los Angeles to Las Vegas. That is not a big surprise. In the small-signal case, $t = \tau$. We can compute the fraction of the small-signal bandwidth we have available for any size signal. We recall from Figure 5.5 (p. 71) that the tanh function has unity slope at the origin, and the point at which the tangent intersects the asymptote I is $2kT/(q\kappa)$, which is 90 millivolts or so. We can thus relate the slew rate to the small-signal parameters: $G = I/(2kT/(q\kappa))$ and $\tau = C/G$. So, from Equation 9.24, the time for a large-signal response is

$$\frac{t}{\tau} = \Delta v/\big(2kT/(q\kappa)\big)$$

What else could it be? It is the size of the output signal in the natural voltage units of the technology. If we make the output signal many times larger, then we get a frequency response that is that many times lower. Or a time response that is that many times longer. Another way of looking at it is that, for large signals, this device is a perfect integrator; it turns into a single time-constant circuit for small signals.

STAYING LINEAR

We can always make sure the follower–integrator circuit acts like a linear system; we just do not allow the input to change by more than approximately 1 $kT/(q\kappa)$ in 1 time-constant τ. In that way, the dV/dt of the input is always less than the slew-rate limit, and the system is always linear. The signal range can be as large as we want if the rate of change of the input is less than the slope of the large-signal response waveform. Under those conditions, the *difference* between the output and the input of the amplifier never gets very big, even though the signal may be huge. The circuit will stay linear nearly up to V_{DD} and nearly down to ground, provided we use one of the wide-range amplifiers, and do not put large step functions into it. If we increase dV/dt past $\big(kT/(q\kappa)\big)/\tau$, the circuit still works well, but it is slew-rate limited.

That is a graceful way for a computation to degrade; it does not give you a floating exception or an integer overflow or any dumb thing like that, it just follows as fast as it can. Such a nice gentle way to behave itself. In Chapter 16, we will study a system in which we can guarantee, by the way the system is arranged, that the input never varies so quickly that the circuit becomes nonlinear. Because the follower–integrator circuit responds only up to a maximum rate, it automatically limits the rate any subsequent device will have to follow. We use the slew-rate limit of the first to control the rise time of the input to the second.

If staying linear is important to us, we organize the system so that it will not be subjected to rise times that are embarrassingly short, and we thus ensure that it will be a linear system. And because it is a follower, it will try very hard to keep in the linear range if we just give it half a chance.

SUMMARY

We have introduced our first explicitly time-dependent computational metaphor. The follower–integrator circuit allows us to perform the same kind of smoothing in the *time domain* that resistive networks achieved in the *space domain.* Integration in one domain or the other (or in both) is the basis on which a large (but unknown) fraction of neural computation is built. The balance of this part of the book is devoted to the more elaborate computations that can be derived from this humble beginning.

REFERENCES

Carslaw, H.S. and Jaeger, J.C. *Conduction of Heat in Solids.* 2nd ed., Oxford, Oxfordshire: Clarendon Press, 1959.

Jack, J.J.B., Noble, D., and Tsien, R.W. *Electric Current Flow in Excitable Cells.* Oxford, England: Claredon Press, 1983.

Rall, W. Membrane potential transients and membrane time constant of motoneurons. *Experimental Neurology,* 2:503, 1960.