Biophysics of Computation
(1999)
Koch

1. The Membrane Equation.

1.6 Recapitulation

In this introductory chapter, we meet some of the key actors underlying neuronal information processing. Basic to all cells is the capacitance inherent in the bi-lipid layer, limiting how quickly the membrane potential can respond to a fixed input current. The simplest of all neuronal models is that of a single compartment that includes a resistance, in series with a battery, and the capacitance. It can be completely described by a linear, low-pass filter. As we will see in a later chapter, such an RC circuit, augmented by a simple voltage threshold, constitutes one of the simplest yet also most powerful models of a spiking neuron: the leaky integrate-and-fire unit.

We introduced fast chemical synapses, the stuff out of which computations arise. Chemical synapses convert the presynaptic voltage signal - via a chemical process - into a postsynaptic electrical signal, via a change in the membrane conductance specific to certain ions. Such a synapse can be described by a time-dependent synaptic conductance $g_{syn}(t)$ and a synaptic battery E_{syn}. In general, synapses cannot be treated as constant current sources.

Chemical synapses, similar to an operational amplifier wired up as a follower, isolate the electrical properties of the postsynaptic site from the presynaptic one. This allows synapses to link neurons with very different electrical impedances. Furthermore, the amplitude, duration, and sign of the postsynaptic signal can be quite different from those of the presynaptic one. Electrical synapses, discussed in Sec. 4.10, share none of these properties.

The fact that synapses act by changing, usually increasing, the postsynaptic membrane conductance has a number of important consequences. It allows for the natural expression of several non-linear operations, in particular saturation and gain normalization. As an example, we saw how shunting inhibition, mediated by a type of synapse whose synaptic reversal potential is close to the cell's resting potential, acts similar to division. We also studied how synaptic input that increases the postsynaptic membrane conductance for some combinations of ions, no matter what its reversal potential, acts to decrease the cell's input resistance and thus its membrane time constant. As postulated by Carandini and Heeger (1994), the effect of massive feedback synaptic input might, in a very physiological manner, implement gain normalization in cortical areas.

2. Linear Cable Theory.

2.5 Recapitulation

One-dimensional cable theory is based on several approximations. (1) The magnetic field due to the movement of charge can be neglected. (2) Changes in the concentration of the charged carriers, Na^+, K^+, and other ions, is slight so that the current can be expressed by Ohm's law, and the intracellular cytoplasm can be mimicked by an ohmic resistance. (3) Due to the wirelike geometry of dendrites and axons and the high resistivity of the neuronal membrane, the radial and angular components of voltage can be neglected, reducing the complexity of the solution from three spatial dimensions to a single one. (4) The extracellular space is reduced to a homogeneous resistive milieu whose resistivity is usually set to zero. This allows us to solve for the potential V (x, t) across the neuronal membrane on the basis of a single equation.

Linear cable theory further assumes that for a limited range of voltage excursions around the resting potential, the membrane properties are independent of the membrane potential, reducing the electrical description of the membrane to resistances and capacitances, greatly simplifying analysis.

Starting in the late 1950s and early 1960s, the linear cable equation was solved by Rall and others to study the dynamics of the membrane potential in dendritic trees. Several key concepts associated with the linear cable equation for a single finite or infinite cylinder are the space constant λ, determining the distance over which a steady-state potential in an infinite cylinder decays e-fold, the neuronal time constant τ_m, determining the charging and decharging times of V (x, t) in response to current steps, and the input resistance R_{in}, determining the amplitude of the voltage in response to slowly varying current injections.

The voltage in response to a current input, whether delivered by an electrode or by synapses, can be expressed by convolving the input with an appropriate Green's function. For passive cables, this always amounts to filtering the input by a low-pass filter function.

While the class of parabolic differential equations (to which the cable equation belongs) does not admit to any wave solutions but only shows dissipative

behavior, one can define input, transfer, and propagation delays by computing and tracking the centroid or center of mass of V (x, t) relative to the centroid of the current input. In the following chapter, we will apply these concepts to realistic dendritic trees.

3. Passive Dendritic Trees

3.7 Recapitulation

While Chap. 2 is concerned with solutions to the cable equation in single cables, this one concentrates on studying the membrane potential in realistic dendritic trees with the aim of understanding the relationship between dendritic architecture and function. A number of exact, recursive techniques exist to derive the membrane potential at any point in the dendritic tree in response to an arbitrary current injection at some other point. Because these techniques do not generalize well to n simultaneous inputs and do not work in the presence of membrane non-linearities, today's method of choice is the numerical solution of the appropriately spatio-temporal discretized cable equation (compartmental method). Several well-documented, graphics-oriented, public software packages are available that implement the appropriate algorithms.

An alternative approach for gaining intuition about the events occurring in a dendritic tree is to "observe" the system at two points i and j, where the input is applied at i and the output recorded at j. The entire system can then be characterized in terms of three frequency-dependent functions that, in general, take on complex values: the input impedances $K_{ii}(f)$ and $K_{jj}(f)$ and the transfer impedance $K_{ij}(f) = K_{ji}(f)$. For sustained current input, these reduce to three real numbers, corresponding to the conventional input resistances K_{ii} and K_{jj} and transfer resistance K_{ii} (with the dimensions of a resistance). This approach allows us to define in a straightforward manner the voltage attenuation A_{is} between any dendritic site i and the soma s and the outgoing voltage attenuation A_{si} from the soma to sites in the dendritic tree. This last measure is identical to the charge attenuation, the ratio of the total charge injected at the synaptic site i to the total charge reaching the soma.

The chapter also introduces the logarithmic transform of the attenuation, $L^{v}_{ij} = \ln|A_{ij}|$. As compared to a true distance metric, this measure is not symmetric: the logarithm of the voltage attenuation between i and j is different from the logarithm of the voltage decrement between j and i. Indeed, for passive trees, the voltage attenuation from any dendritac site to the soma is always considerably larger than in the reverse direction. The reason for this profound asymmetry is that dendritic sites have high impedance while the cell body has a low impedance. A further consequence is that, if integrated over sufficiently long times (in practice several time constants), about half of the total charge injected into distal dendritic sites can reach the soma. We discussed a frequently used measure of synaptic efficiency, the total electronic distance LIs. Due to the fact that real trees are quite distinct from infinite cables, this measure sharply underestimates the true voltage attenuation. All of this can be visualized graphically, using the morphoelectrotonic transforms. METs are a useful dynamic tool for portraying different features of the electroanatomy of a cell.

We also considered delays in the dendritic tree and between dendritic sites and the soma. Using the definition of delays in terms of differences in the centroids of voltages and currents, we can show that tm represents an upper bound on how slowly the voltage decays in response to fast inputs. At the soma, the local delay is on the order of the membrane time constant, while for distal sites the local delay can be much faster, as small as 5 or 10% of τ_m, implying that events in the dendritic tree can be very rapid. This is particularly true for rapid synaptic events in thin dendrites, where the rise and early decay times of the dendritic membrane potential are independent of R_m, allowing for the possibility of carrying out precise timing relationships in these structures that are not limited by τ_m.

The take-home message is that one has to be careful in applying the concepts that characterize an infinite cable (λ, τ_m, and R_{in}) to realistic neurons. These have finite arbors, with multiple dendrites terminating at different distances and branching patterns that do not obey the necessary geometrical constraints for reduction to an equivalent cylinder. These considerations have significant effects on how the voltage spreads and attenuates in dendrites and hence in how synaptic inputs are integrated.

4. Synaptic Input.

Chemical synapses are far more prevalent and important than electrical synapses.

Synapses consist of a presynaptic axonal terminal and a post synaptic process located on a dendrite spine, or trunk, or on a cell body.

4.1 Neuronal and Synaptic packing Densities. Synapse are either excitatory or inhibitory. Synapse approx 0.5-1.0um diameter, and 100000 below each mm2 of cortex.

4.2 Synaptic transmission is stochastic. An action potential causes a rush of Ca+ ions into the presynaptic terminal which causes neurotransmitter filled vesicles 30-40nm to fuse with the synapse membrane and release their quantum. The neurotransmitter diffuses across the synaptic cleft (20nm) in ?s.

4.2.1 Probability of synaptic release p. The probability of release at an individual synapse is highly variable (0.1-0.9) and there are many contributors to this variability.

4.2.2 What is the synaptic weight? Typically the product of the number of quantal release sites, the probability of release at each site, and a measure of the postsynaptic effect.

4.3 Neurotransmitters. Three major chemical classes: amino acids, biogenic amines and neuropeptides. Fast (1ms rise, 20ms decay) transmission in vertebrate CNS is mainly amino acids - major excitatory being glutamate and aspartate and major inhibitory GABA and glycine. Biogenic amines includes acetylcholine, norepinephrine, dopamine, serotonin, and histamine are typically slower (20ms-2s). An ever growing list of peptides operate over minutes and are probably the closest thing to a global variable.

4.4 Synaptic Receptors. Ionotropic receptors (rapid) are directly coupled to ionic gates. Metabotropic receptors (slower , extended and amplifiers) use internal second messengers to link to the ionic gate by diffusion. These are the modulators.

4.5 Synaptic Input as Conductance Change.

4.5.1 Synaptic reversal potential in Series with an Increase in Conductance. A synapse is non-linear and can saturate.

4.5.2 Conductance Decreasing Synapses.

4.6 Excitatory NMDA and non-NMMDA synaptic input. Glutamate is the predominant fast excitatory neurotransmitter but the corresponding receptors are very diverse. The most important distinction is their response to various pharmacological agents known as agonists. Two such excitatory agonists are NMDA and AMPA. Antagonists specifically block a subclass of glutamate receptors. These tools have allowed the action of the various synaptic mechanisms to be isolated and measured. NDMA channels implement an AND gate (significant depolorisation current only with presynaptic neurotransmitters and postsynaptin depolarisation). Note that the postsynaptic current always leads the postsynaptic voltage. Also note that the current flowing at the soma is a much more meaningful measure of synaptic efficiency than the peak postsynaptic potential.

4.7 Inhibitory GABAergic Synaptic Input. The most common inhibitory neurotransmitter in the CNS of both vertebrates and invertebrates is GABA. There are two major forms of receptors, A & B. Pharmological agents are also used here to activate or block these receptors. A receptors are iontropic, operate with Cl ions, the direction of current reverses at about -70mv and are fast (1ms rise and 10-20ms decay). B receptors are metabotropic, operate with K ions, and are slower (10+ms rise and 100+ms decay). A & B receptors do not appear to be colocalised.

4.8 Postsynaptic potential. The descriptic equation is presented follow by examples of stationary, transients and infinitely fast synaptic inputs.

4.9 Visibility of Synaptic inputs.

4.10 Electrical gap junctions. Found in areas such as the cardiac muscle and glial cells where a signal is rapidly propagated to a large number of cells in a symmetrical and only attenuating manner.

4.11 Recapitulation.

Fast communication among nerve cells occurs at specialized junctions called synapses. They are very compact: between several hundred million and one billion synapses can be packed into one cubic millimeter of neuronal tissue. Of the two types, chemical and electrical, we focus on the former since they are much more frequent and make very specific point-to-point connections.

It is useful to distinguish fast ionotropic chemical synapses, acting on a millisecond time scale, from metabotropic chemical synapses, acting on a time scale of a fraction of a second to minutes. Conceptually and cum grano salis, ionotropic synapses are of the essence in the rapid forms of neuronal communication and computations underlying perception and motor control.

In response to a presynaptic change in membrane potential at a synapse, neurotransmitters are released and diffuses within a fraction of a millisecond across the cleft separating pre- and postsynaptic terminal. At the postsynaptic terminal, neurotransmitter molecules bind to specific receptors, which usually-either directly or indirectly, via involvement of a second-messenger system-open specific ionic channels. Depending on the neurotransmitter-receptor kinetics, these channels

remain open for some time and a synaptic current flows across the membrane. Synaptic transmission at central synapses appears to be stochastic. A presynaptic action potential has a probability p of causing a release of a vesicle and a postsynaptic response, where p can be as small as a few percent and depends on the spiking history of the synapse. The amplitude of the postsynaptic signal is variable as well. These factors need to be taken into consideration when thinking about neural computation.

An electrical engineer would be justified in treating a chemical synapse (at the time scale of tens of milliseconds) as a nonreciprocal, two-port device (see the introduction to Sec. 3.4). A two-port description is necessary, since a pair of equations (for both the pre- and the postsynaptic current and voltage changes) is required to completely characterize its behavior; it is non-reciprocal since changes at the postsynaptic side have no (fast) effect on the presynaptic side. Synapses serve to decouple neuronal elements that can each have very different electrical impedances, rather like a follower-amplifier circuit.

At the macroscopic and phenomenological level, a fast synaptic input induces a time-dependent increase in conductance $g_{syn}(t)$ in series with a battery E^{syn}. The sign of E^{syn} relative to the membrane potential at the postsynaptic terminal determines whether synaptic input causes an EPSP (excitatory synapse), an IPSP (inhibitory synapse), or no change in membrane potential (silent or shunting inhibition). The five dominant types of fast synaptic inputs are (1) non-NMDA or AMPA voltage-independent excitation; (2) ACh-mediated excitation; (3) the voltage-dependent and slower NMDA excitatory input that is thought to be crucially involved in synaptic plasticity; (4) the $GABA_A$ type of silent inhibition; and (5) the slower $GABA_B$ hyperpolarizing inhibition. The fact that synaptic input increases the postsynaptic membrane conductance in series with a battery has important consequences. In particular, synaptic inputs can saturate, will influence the input conductance of the cell, and can interact with each other nonlinearly. Only if the amplitude of the synaptic input conductance change is small relative to the local input impedance can synaptic input be treated as a constant current source.

The immense variety of neurotransmitters and postsynaptic receptors gives rise to a staggering combination of possible pairings that act on all possible time scales, and across different spatial scales, from a single synapse to a single ganglion or an entire neural system, such as the thalamus. Synapses are responsible for the most salient difference between nervous systems and even our most advanced digital computers: while the former adapt and learn - a subject we will cover in depth in Chap. 13 - the latter do not.

Electrical synapses allow for direct current flow among adjacent neurons. A gap junction can usually be modeled by a fixed conductance. Different from chemical synapses where amplification between the pre- and postsynaptic sites can occur, here the signal is always attenuated. One advantage of this mode of cellular communication is speed, since no synaptic delay occurs. Thus, electrical synapses are frequently found in neuronal pathways, which sub-serve information that needs to be communicated very rapidly and faithfully. In the retina, gap junctions among photoreceptors and horizontal cells create vast, electrically interconnected networks that filter the incoming visual signal.

5. Synaptic Interactions in a Passive Dendritic Tree.

5.1.7 Retinal Directional Selectivity and Synaptic Logic. "It appears likely that even such a simple operation as distinguishing the direction of a moving stimulus is implemented using a plurality of biophysical mechanisms acting at several sites, some requiring inhibition and some not. Such redundancy might be necessary in the face of demands that the circuitry wire itself up during development and retain its specificity in the face of a constantly varying environment." pg 131.

5.4 Recapitulation.

The fact that a synaptic input changes the conductance in the postsynaptic membrane in series with a synaptic battery and does not correspond to a constant current source ultimately implies that synaptic inputs interact with each other via the membrane potential. In particular, the somatic potential in a passive tree in response to two or more inputs is not equal to the sum of the individual synaptic components. We explored the computational consequences of this in two cases.

The interaction between voltage-independent non-NMDA excitatory input and shunting inhibition (that is, when E_i reverses close to the cell's resting membrane potential) in the sub-threshold domain can mediate a veto operation that is specific in space and time. If inhibition is adjacent to the excitatory synapse or on the path between excitation and the cell body and if their time courses overlap, inhibition can effectively suppress the effect of excitation. This implies that specific synaptic arrangements in a dendritic tree can implement

logic-like AND-NOT operations, possibly one of the crucial nonlinearities underlying direction selectivity in retinal ganglion cells.

The specificity of synaptic interaction represents at the same time also its greatest weakness, in the sense that it places great demands on developmental mechanisms to precisely guide synapses and dendrites during development. A more plausible synaptic arrangement could be implemented at the level of synaptic populations: if excitatory and inhibitory synapses are colocalized onto the same part of the dendritic tree, excitatory input can always override any inhibitory influence. Conversely, if inhibition is at a different site, for instance, close to the spike initiating zone, and excitation at more distal sites, then any excitatory input can always be vetoed by inhibition. These two types of synaptic placements might instantiate different kinds of suppressive behaviors. Under certain conditions the threshold for initiating a reflex should be elevated (relative suppression) while under others the behavior needs to be totally abolished (absolute suppression). If GABAergic inhibition has a reversal potential much below the membrane resting potential, as in the fly tangential interneurons, inhibition tends to act akin to a linear subtraction.

A more plausible mechanism to implement multiplicative behavior involves NMDA synapses clustered over the dendritic tree. When clusters of adjacent non-NMDA synapses are randomly sprinkled around the dendritic tree of a pyramidal cell, their activation causes a smaller cellular response than when the synapses are isolated from each other, a consequence of synaptic saturation. A very different behavior is obtained with clusters of voltage-dependent NMDA synapses. Because of their cooperative nature, clusters of 6 to 10 adjacent NMDA synapses are much more effective than the same number of synapses by themselves. A dendritic tree endowed with such synapses can be used to implement a very efficient nonlinear pattern discriminator that is robust to the presence of dendritic nonlinearities, while also serving as a plausible biophysical mechanism for multiplication, which is required for a host of computations, such as motion, binocular disparity, and tabular look-up storage.

Both case studies imply that the dendritic tree, rather than just performing a filtering operation onto the synaptic input, as suggested by linear cable theory presented in Chaps. 2 and 3, can be partitioned into numerous spatial subunits. Within each such subunit, synaptic inputs interact nonlinearly, while the interaction between two or more subunits is approximately linear.

Finally, we mentioned the concept of a synaptic microcircuit, pioneered by Shepherd (1978, 1998). These usually involve a combination of one or several excitatory and inhibitory synapses with a predilection for a specific arrangement among two or three neurons and are common in extracortical structures such as the retina, the olfactory bulb, and the thalamus and in invertebrates.

6. The Hodgkin-Huxley Model of Action Potential Generation.

6.8 Recapitulation

The Hodgkin-Huxley 1952 model of action potential generation and propagation is the single most successful quantitative model in neuroscience. At its heart is the depiction of the time- and voltage-dependent sodium and potassium conductances G_{Na} and G_K in terms of a number of gating particles. The state of G_{Na} is governed by three activation particles m and one inactivating particle h, while the fate of the potassium conductance is regulated by four activating particles n. The dynamics of these particles are governed by first-order differential equations with two voltage-dependent terms, the steady-state activation (or inactivation), and the time constant. The key feature of activating particles is that their amplitude increases with increasing depolarization, while the converse is true for inactivating particles. For rapid input to a patch of squid axonal membrane, spike initiation is exceeded whenever the net inward current becomes negative, that is, when a particular voltage threshold V_{th} is exceeded.

Inclusion of the cable term leads to a four-dimensional system of coupled, nonlinear differential equations with a wave solution that propagates at a constant velocity down the axon. This wave, the action potential, is due to the balance between dispersion and restoration caused by the voltage-dependent membrane. When injecting sustained currents into the axon, the equations predict two important aspects of the squid axon: the abrupt onset of sustained firing with a high spiking frequency and the very limited bandwidth of the firing frequency.

The Hodgkin-Huxley formalism represents the cornerstone of quantitative models of nerve cell excitability, and constitutes a remarkable testimony to the brilliance of these researchers. It should be remembered that their model was formulated at a time when the existence of ionic channels, the binary, microscopic, and stochastic elements underlying the continuous macroscopic, and deterministic ionic currents, was not known.

Wrapping axons in insulating material, such as the many layers of myelin observed in myelinated fibers that are found in all vertebrates, leads to a dramatic speedup over unmyelinated fibers. Conversely, at the same spike propagation speed, myelinated fibers can be up to 50 times thinner than unmyelinated fibers. In mammals, axons above 1μm are usually myelinated, with speeds in the 5-mm/msec range, and rarely exceed 20μm. When axons reach their target zone, they branch profusely, enabling them to make thousands of contacts on postsynaptic processes. As trains of spikes attempt to propagate past these points, they can be slowed down, depending on the exact geometry of the junction. In the more extreme cases, individual spikes can fail to propagate past branch points.

We conclude that pulses can communicate along axons reliably, rapidly (at speeds between 1 and 100 mm/sec) and with little temporal dispersion. The main exception to this appears to be the propagation of trains of spikes past branching points. Here, due to a variety of phenomena, conduction block can occur, which will differentially route information into one of the daughter branches or prevent conduction altogether.

7. Phase Space Analysis of Neuronal Excitability.

7.4 Recapitulation

The theory of nonlinear dynamics represents a powerful tool to characterize the generic mechanisms giving rise to the threshold response, the stereotypical shape of action potentials, or the onset of oscillations. It allows us to understand why these phenomena occur even when we have insufficient information concerning the detailed kinetics or the exact shape of the various rate and time constants. In this chapter, we apply the theory of dynamical systems to analyze two-dimensional systems with the help of the phase portrait. In particular, we focus on the FitzHugh-Nagumo (FitzHugh, 1961, 1969; Nagumo, Arimoto, and Yoshizawa, 1962) and the Morris-Lecar (1981) models of spiking.

The FitzHugh-Nagumo model was based on the observation that the membrane potential and the sodium activation in the four-dimensional Hodgkin-Huxley equations evolve on a similar time scale, while the sodium inactivation and the potassium activation also share similar behavior, albeit on a slower scale. This feature can be exploited by expressing the membrane excitability via the two-dimensional FitzHugh-Nagumo equations with constant coefficients, with V corresponding to the excitability of the system and W its degree of accommodation. Linear stability analysis allows us to understand why the resting state is stable and when spiking first occurs. It is also helpful to explain the all-or-none shape of the action potential as an example of a limit cycle, upon which the system will rapidly converge once threshold is exceeded.

We also acquainted the reader with a qualitatively similar system of equations that describes muscle fiber excitability in terms of a leak, a calcium, and a potassium current (Morris and Lecar, 1981). Following the lead of Rinzel and Ermentrout (1998), we assume that the calcium current is so rapid that it is always in the steady-state with respect to the membrane potential and the potassium activation, and we discuss two variants of these equations. The first behaves similar to the squid axon and to the FitzHugh-Nagumo equations, generating oscillations with a nonzero oscillation frequency (via a sub-critical Hopf bifurcation). Modifying the potassium conductance to be a steeper function of V_m allows the system to spike at very low frequencies via a saddle-node bifurcation, similar to what occurs in cells that have a marked delay between the onset of current injection and the first spike (due to the presence of IA).

The advantage of using models based on a very small number of variables-rather than relying on biophysical very detailed models with an exuberance of variables-is that they offer us a qualitative, geometrical way to understand why a cell spikes and why it switches between two modes of firing without necessarily having to know every single detail of the system.

Given the enormous numerical load involved in simulating the dynamics of hundreds or thousands of neurons, such a simplified single-cell model will allow us to study neural networks containing realistic numbers of neurons. The price one pays for this reduced complexity is a lack of quantitative predictions.

8.4 Recapitulation

Underlying the entire gamut of electrochemical events in the nervous system are proteins inserted into the bilipid membrane, so-called ionic channels, which allow specific ions passage across the bilipid membrane. Given their small electrical conductance (between 5 and 200 pS), it required the development of the giga-seal technique by Neher and Sakmann to readily observe their behavior. From a functional point of view, the key properties of channels are that they possess one or a very small number of open, conductive states and that the transitions among

closed, open, and inactivated states are governed in a probabilistic manner by the amplitude of the applied membrane potential (for voltage-dependent channels) or the presence of various agonists (for ligand-gated channels).

A very large effort in the field is directed toward identifying and characterizing the molecular sequence of these channels and in relating specific structural features of the channel protein to its voltage and ionic selectivity and, ultimately, to its function. Such a detailed molecular understanding goes hand in hand with the construction of ever more complex kinetic models, describing the transitions of a single channel among a large number of internal states in terms of probabilistic Markov models.

Numerical studies have related the microscopic, stochastic chatter of individual ionic channels opening and closing to the observance of macroscopic currents changing in a highly deterministic and graded manner. A membrane at rest and studded with a few thousand discrete channels behaves in general little different from the deterministic Hodgkin-Huxley equations. Deviations are only expected to occur when the membrane potential is close to threshold or when the channels are embedded into a small patch of neuronal membrane with a low leak conductance. In the latter case, the single-channel conductance can have the same magnitude as the passive leak conductance, and the opening of one or a few channels can depolarize the membrane beyond Vth. The resultant spikes are Poisson distributed. In other words, the microscopic behavior of individual molecules is amplified and causes a macroscopic event, an action potential. Such a mechanism, known to occur in certain experimental preparations, could be of functional relevance in very small cells or in electrically decoupled, distal parts of the dendritic tree as a "random event" generator.

In the remainder of this book we will typically deal with large enough membrane areas - and therefore channel numbers - that we are usually justified in treating electrical events in terms of deterministic, continuous currents, rather than in terms of probabilistic, all-or-none ionic channels.

9.6 Recapitulation

In this chapter, we summarized some of what is known about ionic currents populating the nervous system other than I_{Na} and I_{DR} of squid axon fame. In particular, we introduced calcium currents. Ca^{2+} ions are crucial for the life and death of neurons; they link rapid changes in the membrane potential, used for computation, with action, such as neurotransmitter release at a synapse, muscle contraction, or biochemical changes underlying cellular plasticity. Three broad classes of calcium currents are known: L, T, and N currents, although it is unclear whether these are just three samples along a continuum of such currents. Because of the substantial imbalance between the concentration of Ca^{2+} inside and outside the cell, electrical rectification needs to be taken account of. Typically, this requires the use of the Goldman-Hodgkin-Katz equation to describe calcium currents, rather than the linear Ohm's law. The accumulation of calcium ions following action potential activity has the useful function that the intracellular calcium concentration in the cell represents an index of the recent spiking history of the cell. Various proteins and enzymes, such as calcium-dependent potassium channels, might be able to read out this variable and use it to normalize excitability or for gain control.

While calcium ions themselves play almost no role in generating conventional action potentials, calcium currents can support a slower all-or-none electrical event, the LT spike, leading to a burst of fast sodium spikes. It may be possible that these and other forms of bursts represent events of special significance in the nervous system.

We introduced some of the many potassium currents that have been identified, focusing on the transient inactivating potassium current IA. It serves to linearize the discharge curve around threshold and implements a delay element between the onset of depolarizing input and spike initiation. It does so via a voltage-dependent current whose time constant is considerably faster than the duration of the delay.

Also important are the calcium-dependent potassium currents, collectively referred to as $I_{K(Ca)}$. Calcium ions that move into the cell during the depolarizing phase of the spike turn on these hyperpolarizing outward currents, rendering subsequent spiking that much more difficult. That this might have profound significance for the nervous system can be observed in hippocampus, neocortex, and elsewhere. The release of noradrenaline by brainstem fibers abolishes spike frequency accommodation, greatly increasing the excitability of these cells. Interestingly, the effect of blocking only IAHp appears to be specific to increasing the slope of the cell's f -I curve, that is, increasing its gain.

We keep on drawing analogies between the operations the nervous system carries out to process information and digital computers. There are, of course, very deep, conceptual differences between the two. One is the amazing diversity seen in biology. An I_A current in a neocortical pyramidal

cell is not the same as the I_A current in hippocampal pyramid cells or in thalamic relay cells. Each has its own unique activation and inactivation range, kinetics, and pharmacology, optimized by evolution for its particular role in the survival of the organism. A case in point appears to be calcium channels. So far, molecular biological techniques have revealed about a dozen different genes coding for two of the five subunits of the calcium channel, a number that is likely to increase over time. This provides the nervous system with the wherewithal to generate a very large number of calcium channels with different functional properties (Hofmann, Biel, and Flockerzi, 1994). This tremendous diversity could be necessary to allow each organism to optimize its complement of ionic currents for its particular operations depending on its particular and unique developmental history. This is, indeed, very far removed from the way we design and build digital integrated circuits using a very small library of canonical circuit elements.

10.5 Recapitulation

Although it can be argued that a linear analysis of a nonlinear phenomenon does not do justice to it, it will certainly help us to understand certain aspects of the mechanism underlying the phenomenon. This is true when considering certain resonant or oscillatory behaviors evident in nerve cells.

As Detwiler, Hodgkin and McNaughton (1980) pointed out, a time- and voltage-dependent potassium current that is activated by depolarization or an inward current that is activated by hyperpolarization (as the h part of the fast sodium current) does behave, under certain restricted conditions, like an inductance in series with a conductance would. That is, one can mimic the small-signal behavior of such a current by an electrical circuit that includes such an inductance. This explains why EPSPs or the depolarization in response to even very small current injections will be followed by a hyper-polarizing overshoot (provided the experimental setup is sensitive enough to record such small voltages superimposed onto the background noise). The membrane impedance of a membrane that includes resistances, capacitances, as well as phenomenological inductances (RLC circuit) peaks at some nonzero resonant frequency f_{max}, Such a quasi-active membrane with a bandpass response will cause the membrane potential to show damped oscillations at or close to the resonant frequency fmax in response to a current step. This has been confirmed empirically when recording the response of the squid giant axon to small, sub-threshold current steps.

One consequence of a bandpass-like membrane impedance is that the transfer impedance and the frequency-dependent space constant associated with an infinite cable covered by such a quasi-active membrane can also show resonant behavior. This implies that input signals with frequencies around fmax are treated preferentially in terms of a smaller voltage attenuation to the soma (or other measures of synaptic efficiency) than inputs at faster or slower frequencies: the cable is spatio-temporally tuned to some frequency range. If the transfer impedance of a cable with a quasi-active membrane can be approximated by a bandpass, injecting a current whose dominant signal content lies in the spectrum between dc and fmax induces a voltage change that approximates the continuous temporal derivative of the injected current. In other words, a small piece of quasi-active dendritic cable can implement a temporal derivative operation.

Hudspeth and Lewis (1988b) used a nonlinear, squid axon-like membrane description of a calcium and a calcium-dependent potassium current (in series with leak and capacitive currents) to model the electrical behavior of bullfrog hair cells in the cochlea. Confirming earlier linear RLC analyses, each hair cell by itself acts as an electrical resonant element best tuned to respond to sounds in the fmax frequency band. The kinetics and density of $I_{K(Ca)}$ control fmax as well as the tuning of the bandpass. The morale is that individual neurons can implement a variety of different linear and nonlinear computational operations on the basis of the cornucopia of known membrane currents.

11.9 Recapitulation

Diffusion is a fundamental fact of life for molecules in the intracellular or extracellular cytoplasm. Through its random action, it acts to move substances throughout the cell.

From a computational point of view, the most important fact about diffusion is that it places strong constraints on how rapid calcium or other second messenger molecules can affect things far away. The distance over which some concentration increase diffuses is proportional to the square root of the time that has passed. In the absence of any calcium nonlinearities and active transport processes, this square-root law fundamentally limits the ability of the calcium signal to implement the fast type of information processing operations required for many perceptual, cognitive, or motor tasks. Recognizing a friend's face, shifting visual attention from one location to a neighboring one, or raising one's hand to catch a ball can all be accomplished within a few hundred milliseconds.

The calcium that rushes into the cell via ionic channels is tightly regulated. The vast majority is bound to a host of intracellular buffers, such that only one out of 20 Ca^{2+} ions is free to interact with other molecules, severely limiting the effective diffusion coefficient of calcium.

While calcium ions diffuse along some process, their concentration rapidly decreases. This is especially true for a substance that diffuses in the three-dimensional extracellular tissue; its spatial concentration profile decreases sharply with distance from the source (as e^{-r2}). A relevant case are certain unconventional neuroactive substances, such as nitric oxide, that can diffuse across the membrane cytoskeleton (Secs. 20.2 and 20.3).

Of course, these constraints do not argue against the use of the local intracellular calcium concentration for computing and for short-term memory storage (Sobel and Tank, 1994). When calcium and its protein targets are in close spatial proximity, the rate at which calcium can bind to this protein limits the speed of the computational operation being implemented
(Sec. 20.1). This allows chemical switching to proceed in the sub-millisecond domain. Using concentration changes for implementing rapid operations does impose stringent conditions on a fast local input and a fast local read-out mechanism.

In general, movements of ions due to the inhomogeneous distribution of the various relevant ions (Ca^{2+}, Na^{+}, K^{+} and Cl^{-}) must be incorporated into the cable equation, leading to the Nernst-Planck electrodiffusion equation. However, as long as the diameter of the neuronal process is above a fraction of a micrometer, this equation is well approximated by the cable equation. Only when studying very small processes, such as dendritic spines or very thin dendrites, does the longitudinal diffusion of the carriers need to be taken into account.

If the buffering reaction is substantially faster than diffusion and if the calcium concentration is small (technically, if $[Ca^{2+}]_i < K_d$ of the pump and of the buffer), the coupled system of reaction-diffusion equations can be reduced to a single linear partial differential equation, which is formally equivalent to the cable equation. This allows us to define space and time constants and input resistances in analogy to these parameters in passive dendrites. One important insight is that $\lambda_{r-d} \ll \lambda$, implying that from the point of view of spatial compartmentalization, the presence of reasonable amounts of calcium pumps and buffers in the dendritic tree will fractionate the tree into a series of small and relatively independent compartments. In each of these subunits, independent calcium-initiated chemical computations could be carried out. This is in contrast to the relatively smaller attenuation experienced by the membrane potential in a dendritic tree. It may well be possible that the architecture and morphology of the dendritic tree reflects less the need for electrical computations but more its role in isolating and amplifying chemical signals. We will study a beautiful instance of this in the following chapter on dendritic spines.

Because of the mathematical equivalence between electrical, chemical, and even biochemical networks, (Busse and Hess. 1973: Eigen, 1974: Hielmfelt and Ross. 1992; Barkai and Leibler, 1997) that derives from their common underlying mathematical structure, appropriate sets of reaction-diffusion systems can be devised that emulate specific electrical circuits. In principle, computations can be carried out using either membrane potential as the crucial variable-controlled by the cable equation-or concentration of calcium or some other substances-controlled by reaction-diffusion equations (for examples of this, see Poggio and Koch, 1985). The principal differences are the relevant spatial and temporal scales, dictated by the different physical parameters, as well as the dynamical range of the two sets of variables. Given neuronal noise levels, the membrane potential can be considered to vary by a factor of hundred or less, while the concentration of calcium or other substances can vary by three or more orders of magnitude during physiological events.

12.7 Recapitulation

The small extent of dendritic spines, somewhat smaller than the average-sized Escherichia coli colon bacterium, precluded until very recently direct experimental access, making the function and properties of dendritic spines a favorite subject among modelers. Hypotheses concerning their function can be divided into two major categories: spines as devices for the induction and/or the expression of synaptic plasticity and spines as devices sub-serving specific computations. Other proposals have been advanced, such as that spines primarily serve to connect axons with dendrites (Swindale, 1981), that spines serve to increase the effective membrane capacity of a dendrite (Jaslove, 1992), or that spines serve to protect the dendrite from high, and therefore possibly toxic, overdoses of calcium during synaptic activation (Segal, 1995b). While each of these ideas may contain some grain of truth, we here emphasize those properties of spines of direct relevance to information processing and storage.

That changes in the electrical resistance of the spine neck can modulate the EPSP amplitude of a synapse on this spine has been a popular idea since it was discussed in detail by Rall (1970, 1974). It now appears, at least in the case of hippocampal CAI pyramidal cells (but most likely also for neocortical spines), that the spine neck conductance is too large relative to the synaptic-induced conductance change to be able to effectively modulate the dendritic EPSP for passive spines. Furthermore, as covered in the following chapter, the weight of the evidence leans in favor of a pre- rather than a postsynaptic site as the locus where the long-term changes in synaptic weight are affected. It is most likely in the induction phase of Hebbian long-term changes in synaptic plasticity that spines play a critical role.

Both theoretical and experimental evidence is accumulating that spines may create an isolated biochemical micro-environment for inducing changes in synaptic strength. In the induction of associative (or Hebbian) UP, for example, the spine could restrict changes in postsynaptic calcium concentration to precisely those synapses that met the criteria for potentiation. Furthermore, changes in spine shape could control the peak calcium concentration induced by synaptic input: all other factors being equal, long and skinny spines would have higher peak calcium levels than short and stubby spines. This could provide an alternative explanation for some of the environmental effects on spine morphology. The biochemical compartmentalization provided by dendritic spines could, of course, be equally crucial for a number of other diffusible second messengers, such as IP3, calmodulin, cyclic AMP, and others. We conclude that while spines may not play any role in the expression of synaptic plasticity, they may be crucial in the induction phase by offering a protected micro-environment for calcium and other messenger molecules.

Spines can provide a very rich substrate for computation if they are endowed with regenerative, all-or-none electrical properties as experimental data are now suggesting. The elevated spine input resistance, as compared to the input resistance of the parent dendrite, and the partial electrical isolation of the spine from the rest of the cell, could make spines a favorable site for the initiation of action potentials. The proposal that spines contain voltage-dependent membrane conductances, first advanced by theoreticians and now experimentally supported by evidence for calcium spikes, implies that a dendritic spine is a basic computational gate with two states, on or off. Modeling studies, treated in Sec. 19.3.2, demonstrate that a population of strategically placed active spines can instantiate various classes of quasi-Boolean logic. The combination of excitatory and inhibitory synapses on a single spine-whether passive or active-observed in certain brain areas can only enhance this computational power. With up to 200,000 spines on certain cell types, it remains to be seen whether the nervous system makes use of this vast computational capacity.

13.7 Recapitulation

Behavioral plasticity or adaptation is critical to an organism's survival. Adaptation occurs throughout the nervous system and on many different time scales, from milliseconds to days and even longer. Changes in synaptic strength are widely postulated to be the primary biophysical substrate for many forms of behavioral plasticity, including learning and memory, although our understanding of the link remains far from complete. Use-dependent forms of synaptic plasticity have been characterized in many preparations.

Synaptic strength or weight can be characterized by the triplet (n, p, q), where n is the number of release -sites, p the probability of synaptic release, and q the amplitude of the postsynaptic response following the docking of a single vesicle. The induction of most short-term forms of plasticity depends only upon the history of activity in the presynaptic terminal, while longer term forms require that appropriate conditions be met at both the pre- and the postsynaptic sites. It is therefore only these longer term forms of plasticity that can implement Hebbian type of learning. Biophysical modulation of p underlies most short-term forms, and appears to account for at least a component of some long-term forms as well.

By far the best studied biophysical model of long-term synaptic change is UP The induction of UP requires a conjunction of presynaptic neurotransmitter release, combined with postsynaptic depolarization of the postsynaptic site. There appears to be quite a requirement for the presynaptic input to precede firing activity in the postsynaptic cell. These experimentally observed forms of plasticity can be described at a more abstract level in terms of synaptic algorithms, in particular by temporally asymmetric Hebbian learning rules. Such formulations are useful because a great deal is known from the literature on artificial neural networks about the computational possibilities of Hebbian synapses.

It is important to realize the prevalence of usage-dependent forms of synaptic plasticity. While digital transistors have been designed to be as constant as possible at switching speeds of hundreds of megahertz over the lifetime of the processor, a

single synapse will vary its weight considerably in response to two or more consecutive spikes. These short-term changes can take many forms, including depression and facilitation. In at least one case, the three different types of synaptic input to layer 4 spiny cells show the entire gamut of short-term plasticity: none, short-term depression and short-term enhancement (Stratford et a1.,1996). In summary, synaptic properties show a complex dependency on their previous history of usage and on the postsynaptic activity. We are only beginning to understand the computational significance of such dynamic switching elements.

14.5 Recapitulation

We started off this chapter by defining the instantaneous firing frequency f(t). It is a fictive variable that can be obtained by averaging the spiking response of a single neuron to multiple presentations of the same stimulus. Due to the stochastic nature of the neuronal response the exact micro-structures of spike trains are rarely reproducible from trial to trial. This is why the temporal average (f (t)) of the firing rate is the most common variable measured during neurophysiological experiments. In a handful of experiments J),), evaluated over fractions of a second or longer, has been directly related to the behavior of the animal.

There is no question that firing rate codes that preserve temporal information at the 5-10 msec level are used in the nervous system. To what extent more complex correlation codes-exploiting information encoded among n-tuples of spikes from one neuron or spiking information across multiple neurons-exist remains a subject of considerable debate. What is the simplest model of a single neuron that captures some of its key operations? Two families of models are in common use today: integrate-and-fire and continuous firing rate models. While the former retains the timing information of individual action potentials, the latter assumes that it is only the average or mean firing rate of a neuron that matters to its postsynaptic targets. Both neglect the dendritic tree and both eliminate the complex time course of the sodium and potassium membrane conductances underlying spiking.

The key insight behind the various guises of the integrate-and-fire model is that from a phenomenological point of view the neuron possesses two domains of operation, a subthreshold and a suprathreshold one. In the subthreshold domain, synaptic inputs are integrated and decay away; their temporal evolution is governed by the time constant t. Once the voltage threshold is reached, a pulse is generated and the membrane potential is reset. Different versions of integrate-and-fire models, incorporating various mechanisms to account for adaptation, can be well fitted to the discharge curves of cortical and other cells. It will be argued in Sec. 17.3 that firing in response to fast synaptic input in complex, conductance-based single-cell models is, indeed, initiated whenever a voltage threshold is exceeded.

In response to a supra-threshold stimulus, these units, in accordance with their biological counterparts, can spike in a time Tth « t. A network of integrate-and-fire units can respond almost instantaneously to a stimulus. The take home lesson is that the dynamics of the sub-threshold domain do not carry over into the supra-threshold domain.

In firing rate neurons the continuous output variable is an instantaneous function of the voltage. Since the evolution of the voltage is dictated by a time constant, the firing rate will always be low-pass filtered with respect to the input current, distinct from the response of real neurons, and different from integrate-and-fire units. If one would like to retain the continuous nature of the firing rate model, a more physiologically correct way to achieve this would be to make the steady-state firing rate a function of the total current (synaptic, dendritic, or otherwise) at the cell body. The output of such a neuron can be interpreted as the firing rate associated with a population of spiking cells.

At the heart of the vast majority of neural networks lies the assumption that synaptic inputs interact in a linear manner. The nonlinearity that is necessary for computation is relegated to the firing mechanism at the output. A biophysically more faithful and more complex model that incorporates multiplicative interactions among synaptic inputs is the polynomial or sigma-pi unit.

Multiplication is a key operation underlying many neuronal operations. Chapter 5 treated the evidence in favor of the view that a dendritic tree endowed with NMDA synapses and voltage-dependent membrane conductances (see Chap. 19 as well) can implement a robust version of such a polynomial unit. The nonlinear operations underlying the polynomial interactions do not depend on the threshold occurring at the cell body but precede it. The computational power of such neurons is considerably beyond that of their feeble-minded linear threshold counterparts.

15.4 Recapitulation

In this chapter we tried to address and quantify the

stochastic, seemingly random nature of neuronal firing. The degree of randomness speaks to the nature of the neural code used to transmit information between cells. Very regularly firing cells are obviously not very good at encoding information in their timing patterns; yet this will not prevent information from being encoded in their mean firing rates, albeit at a lower rate (in terms of bits per spike) but in a robust manner.

The spiking behavior of cells has traditionally been described as a random point process, in particular as a renewal process with independent and identically distributed interspike intervals. Cortical cells firing at high rates are at least as variable as expected from a simple Poisson process. Variability is usually quantified using two measures: Cv to assess the intersDike interval variability and FM for spike count variability, with both on unity for a Poisson process. If spike trains are generated by a renewal process, F = C2v in the limit of large observational intervals.

The power spectrum of cortical spike trains is flat with a dip around the origin, as expected from a Poisson process modified by a refractory period. (This refractory period only comes into play at very high firing rates, serving to regularize them.) The rate of the spiking process is usually not constant in time, but is up or down regulated at the 5-10 msec level. The principal deviation from Poisson statistics is the fact that adjacent inter-spike intervals are not independent of each other (even when neglecting bursting cells), that is, spike trains cannot really be described by a renewal process.

As always in science, this conclusion gives birth to intertwining considerations. What type of models of synaptic integration give rise to the high degree of randomness apparent in neuronal firing and how do these constrain the nature of the neuronal code? The standard Geiger counter model predicts that when an integrate-and-fire unit needs to integrate over a large number of small synaptic inputs, it should fire very regularly. Since this is patently not true in the cortex, it needs to be abandoned. Out of the many alternatives proposed, two divergent views crystallize. One school retains the idea that neurons integrate over large number of excitatory inputs with little regard for their exact timing by invoking a large degree of inhibitory inputs (following the random walk model advocated by Gerstein and Mandelbrot, 1964), a depolarizing reset, or correlated synaptic input. The other school sees cortical cells as coincidence detectors, firing if small numbers of excitatory events arrive simultaneously at the millisecond (or even sub-millisecond) scale (via powerful and fast dendritic nonlinearities). Under these circumstances, detailed timing information can be used to transmit information in a manner much more efficient; yet also more demanding, than in a mean rate code. Only additional experimental evidence can resolve this issue.

Of course, at small enough time scales or for a handful of spikes, the debate loses its significance, since it becomes meaningless to define a rate for a 20 msec long segment of a spike train with just two action potentials.

What this dispute shows is that integrate-and-fire units serve as gold standard against which models of variability are evaluated. Given their relative simplicity-compared against the much more complex conductance-based models of firing-this is quite remarkable.

The highly variable character of cortical firing allows neurons to potentially pack one or more bits of information per action potential into spike trains (as done in sensory neurons closer to the periphery). Information theory as applied to a band-limited communication channel has taught us that the optimal code - optimal in terms of using the entire bandwidth available - looks completely random, since every redundancy has been removed to increase the efficiency. One could infer from this that neurons make optimal use of the limited bandwidth of axons using a sophisticated multiplexed inter-spike interval code from which all redundancies have been removed, and that neurons, properly decoded, maximize the existing channel bandwidth. To what extent they actually do for physiologically relevant stimuli remains an open issue.

16.4 Recapitulation

Bursts, that is, two to five closely spaced, sodium-dependent fast action potentials riding on top of a much slower depolarization, are a dominant feature of a number of cell classes, not only in the cortex but also among thalamic relay neurons and elsewhere. In the cortex, intrinsically bursting cells have a unique morphology and are confined to layer 5, where they constitute the dominant output cell class.

The biophysical mechanisms underlying bursting are diverse. A low-threshold calcium current at the soma is the principal agent for bursting in thalamic cells (Sec. 9.1.4). Bursts in pyramidal cells can originate when sodium spikes propagate back into the dendritic tree, causing parts of the arbor to be depolarized. Under the right circumstances, for instance,
when amplified by dendritic sodium and/or calcium currents, this signals returns to the soma in the form of an after depolarization that triggers several fast

sodium spikes.

It has been argued that bursting represents a special code, quite distinct from the firing of an isolated action potential. This hypothesis has some experimental support from the electrosensory system of the electric fish. Spikes taken from bursts signal the presence of particular features in the input in a much more reliable manner than isolated spikes.

This implies that spiking cells are not restricted to using an asynchronous binary code but might employ a two-channel multiplexing strategy, with the decoding being carried out by rapidly adapting synapses.

A non-exclusive alternative is that bursts serve to synchronize distant neuronal populations.

17.6 Recapitulation

This chapter treated the characterization of a nonlinear and non-stationary neuron from the point of view of the soma on the basis of two measures: the relationship between the current flowing across the somatic membrane and the voltage and the local slope of such an I-V curve. The inverse of this last variable is the slope or input resistance. The traditional I-V curve Istatic8(Vm) is defined under voltage clamp conditions: the voltage is fixed and the current flowing at this voltage is recorded. The local peak around Vrest defines the minimal current It, necessary to initiate spiking. Under the assumption that the activation and inactivation variables of the voltage-dependent currents change much more slowly than the activation m(t) of the sodium current, it is possible to define an instantaneous current-voltage curve Io(Vm). (Conceptually, this can be measured by yanking the membrane potential from Vrest to V„t and instantaneously recording the resulting membrane current.) Stability considerations of the sort explored in Chap. 7 dictate that the zero crossing of Io, if the slope is negative, corresponds to the voltage Vh that needs to be exceeded for a spike to occur.

We conclude that fast and powerful versus sustained but peri-threshold stimuli define two different threshold conditions: in the first case Vn must exceed Vth while in the latter case Isoma must exceed Ith. The larger the current delivered to the soma by a synaptic input, the faster Vth can be reached and the sooner the cell spikes. This process is not limited by t
but by the amplitude of the upstroke of the action potential (that is, by the amount of sodium current and the somatic capacitance). For cells with no or little distributed capacitances, which a fixed amount of charge needs to be placed onto the somatic capacitance in order to reach Vn.

Introducing (Vm) as the temporal average of the somatic membrane potential (including spikes) allows us to extend most of the above concepts into the spiking regime. Biophysically, (V.) approximates the average membrane potential experienced by a distal dendritic site, where the high-frequency components of the back-propagating somatic action potential
have been filtered out by the capacitances distributed throughout the cable.

The dynamic current-voltage curve Idynamic8 ((Vm)) is defined as the inverse of the relationship between the sustained current injected via a micro-electrode and the average membrane potential (Koch, Bernander, and Douglas, 1995). It can be thought of as a single, effective spike current that serves to stabilize the membrane potential in the neighborhood
of Vth and that reverses at Vrest. This current can be described rather well by an exponential relationship: increasing (V m> by 3.7 mV roughly doubles 1. The fact that this current behaves similarly to a diode in the forward direction makes such a relationship particularly easy to implement using electronic circuits. A similar exponential relationship also exists for a patch of squid axon described by the Hodgkin-Huxley equations (not shown).

One caveat. These dynamical measures are based on time-averaged quantities; they cannot be used under conditions where the detailed time course of spiking is thought to be relevant for postsynaptic processing.

18.6 Recapitulation

The past few pages treated synaptic input arriving at the passive dendritic tree of a spiking cell. Only the cell body is assumed to contain voltage-dependent membrane conductances. We first discussed unitary EPSPs and EPSCs, concluding that the former can be used to predict quite well the voltage threshold Vh of the cell. We then investigated the exact relationship between the temporal jitter in a normally distributed set of synapses 6;„ and the jitter in the output spike Qo„t. Surprisingly, we discovered that oo„t a 6;„, with the constant of proportionality being much less than unity. Once again (e.g., Chap. 14) for strong synaptic input (that is, in the above threshold domain), the passive membrane time constant plays only a minor role in determining the temporal behavior of the cell. This has important implications,

in particular for the faithful conservation of timing relationships across many layers of neurons (Abeles, 1990; Marsalek, Koch, and Maunsell, 1997).

For the remainder of the chapter we assume that synaptic input is not correlated at a fine time scale and that, indeed, we can define an average input frequency (f). This allows us to treat massive synaptic input in three separate cases.

Neurons receive between 103 and 105 synapses onto their dendritic tree from other neurons that fire spontaneously between 1 to 10 or more spikes per second in the behaving animal. This massive synaptic background activity has several effects on the spatio-temporal behavior of the cell. As this background firing rate goes up, the electrotonic dimensions
of the cell increase (that is, distal synaptic input becomes more decoupled) and the input resistance and time constant decrease by one or more orders of magnitude. This implies, among other things, that the cell becomes more sensitive to temporal coincidences as the overall network activity increases. Given the rapid changes in firing activity in large parts of the brain, as assessed by EEG and other macroscopic techniques recording large-scale brain activity, this implies that the collective behavior of the network can directly influence the spatio-temporal integrative properties of individual neurons.

We also described a method that characterizes the efficiency of massive synaptic input, involving the current Isyn,s that is delivered by synaptic input to the soma. Deriving Isy„,S for synaptic input to a particular spatial region in the tree firing at a frequency (fl') has the major advantage that the resultant function can be combined with the f -I curve measured at the soma to yield a complete relationship between the sustained presynaptic input firing frequency (f„) and the adapted output firing rate fo„r.

When this method is applied to study synaptic input into the distal part of the apical tree of a pyramidal cell, it is seen that the current from these sites (1) saturates already at very small presynaptic firing rates and (2) influences the maintained firing rate only marginally. This saturation is only to a very minor extent caused by current leaking across the membrane between the synaptic site and the soma. As we will see in the next chapter, introducing voltage-dependent currents into the dendrite can both prevent saturation and amplify these small currents, effectively turning the apical tree into a better "wire."

Finally, we reexamined the effect of shunting inhibition. While it acts in a divisive manner in sub-threshold cable models, it behaves in a linear, subtractive manner when spiking is accounted for. This is due to the fact that the average somatic membrane potential does not exceed the voltage threshold, limiting the maximal inhibitory current. Thus, at least for
models where inhibition acts in a maintained manner at or close to the cell body, shunting inhibition does not implement a divisive normalization operation as postulated by many.

However, at the next level of organization, the small recurrently connected network, linear inhibition can act again in a divisive manner (Douglas et al., 1995). We conclude with the observation that a biophysical mechanism that implements one type of operation in the sub-threshold domain might implement a quite different one in the supra-threshold domain. The morale is that it is dangerous in neurobiology to study any one mechanism at only a single, isolated level of complexity. Phenomena at multiple levels, such as ionic channel, synapse, dendrite, neuron, small network, and so on, interact in highly nonlinear and nonintuitive ways. This is, of course, a characteristic of any evolved systems and makes them so interesting.

19.4 Recapitulation

it was held for a long time that voltage-dependent membrane conductances were restricted to the cell body and axon. Given the overwhelming evidence for the presence of a variety of sodium, calcium, and potassium conductances in the dendritic tree it is imperative that such active membranes be included into biophysically realistic compartmental models. We must also understand the functional role of such conductances.

Because of the great diversity of voltage-dependent currents found in so many different cell types throughout the animal kingdom, it is difficult to summarize the experimental data. The dendritic membrane of many neurons, in particular pyramidal cells, contains a relatively low but homogeneous distribution of ionic channels. This is enough to support the propagation of spikes from the axon and soma back into the dendritic arbor. That the spike is initiated in the axon and not locally in the dendrite can be explained by the fact that the difference in voltage threshold for spike initiation in the dendrite and Vth in the axon (due to much higher densities of active channels and/or intrinsic activation curves that are shifted toward more hyper-polarized values) is less than the voltage attenuation from the dendrite to the site of spike initiation

Under certain conditions, both rapid (1 msec) all-or-

none sodium spikes and slower (10-50 msec) calcium-based action potentials can be observed in the dendritic tree. These may or may not propagate to the soma. The extent to which faster or slower spikes are initiated under physiological conditions in the dendrites and propagate to the soma remains unclear.

Numerous specific functions have been ascribed to active dendritic conductances, some dependent on dendritic spikes while others work in the absence of action potentials. Among the former resides the hypothesis that the timing of the back-propagated spike relative to the timing of any near-simultaneous presynaptic input is a crucial variable controlling whether the weight of the associated synapse decreases or increases. A number of researchers simulated the ability of locally generated spikes in dendrites and spines to instantiate various logical operations as well as temporal coincidence detection.

Two quite plausible proposals for the functional role of active inward currents in the dendrites see them as providing the biophysical substrate either to implement a multiplicative-like operation in a distributed fashion among spatially clustered groups of synapses (the neuron as a sigma-pi element) or to linearize and amplify distal synaptic input selectively so that this input can modulate the cell's discharge more effectively (dendritic amplifier). Any such proposal raises many more questions than it answers, one of the most pertinent one being the question of the necessary developmental and learning rules that can assemble all of this hardware with the required temporal and spatial specificity (that is, inserting a certain number of channels of a specific type into a specific location at some sub-neuronal site). While a passive dendritic tree would require much less specificity to "wire up," its computational abilities are far less than those of an active dendritic arbor.

20.6 Recapitulation

We here dealt with a number of mechanisms that are not conventionally thought of as sub-serving specific neuronal computations. None of them involve the membrane potential, the firing rate, or the intracellular calcium concentration.

The entire realm of biochemical computations has been neglected. Yet at present there are no solid arguments ruling out why specific molecular reactions might not subserve specific computations. As one instance we introduced a molecular flip-flop switch that relies on positive feedback-via autophosphorylation-to implement a long-term memory device with the storage capacity of a few bits that could reside at individual synapses. This is but one example of a realm of computation about which we know almost nothing. Given the extremely large and complex regulatory cascades and networks of proteins and enzymes, the possibilities for nested multilevel computations are staggering. The crucial question is whether the brain avails itself of these possibilities or whether such computations cannot be implemented for reasons having to do with lack of bandwidth, signal-to-noise ratio or specificity.

Two other candidate mechanisms rely on the precise three-dimensional arrangements of neuronal components, either at the sub-cellular or at the cellular level. Whether or not the short-term depletion of extracellular Ca 2+ ions implements a universal presynaptic inhibition that works without any conductance changes is pure speculation at the moment, but is too important to neglect from an experimental point of view.

That puffs of nitric oxide (and possibly carbon monoxide) are released in nervous tissue following local hot spots of synaptic-induced calcium activity opens up new avenues of spreading information in a retrograde manner, back across the synapse. Given the inexorable square-root law of diffusion and the aggressive chemical nature of nitric oxide, NO is unlikely to be effective beyond a small fraction of a millimeter from the site of its synthesis: This sphere of influence does include a potentially very large number of synapses. One of the "unfortunate" consequences of such a diffusing substance is that the specificity of Hebb's synaptic plasticity rule would be significantly reduced. The unit of learning would not be individual synapses, but groups of adjacent synapses.

The last two mechanisms exploit the very large laundry list of neuroactive substances iogenic amines, neuropeptides, hormones) known to be present in any nervous system to implement global variables that act over a fraction of a millimeter and longer distances and i a seconds to minutes and longer time scale. We speculated on the role of neuromodulators
routing information, for reprogramming a particular neural network to change its mode of operation, for adapting the retina or other sensory surfaces to different operating conditions, and the like.