We Got Rhythm: Dynamical Systems of the Nervous System

Nancy Kopell

Introduction

The nervous system is perpetually active, in waking and sleep, creating its own dynamics as well as reacting to the world around it. Among the dynamics are many periodic rhythms. Some of these are associated with periodic motor behavior, such as walking, chewing, and breathing. Others—the ones I am going to be discussing here—are much more mysterious. They occur in a wide range of frequencies and are associated with mental processing, including sensory activity and cognitive states [1]. These dynamical phenomena raise many mathematical questions, and I shall be focusing on a role for mathematics in the attempt to understand the origin and uses of dynamics in the nervous system.

A previous Gibbs Lecturer, Norbert Wiener, was also intrigued by "brain waves" and wrote about it in his famous book *Cybernetics* (2nd Edition). At that time, almost forty years ago, the technology for data acquisition and analysis was primitive; it was a breakthrough to be able to make measurements amenable to the use of autocorrelations so as to be able to document the existence of rhythms. Now there are powerful computers and very sophisticated technology for measurements of activity of individual neurons and collections of them. However, the wealth of data acts to underscore the limitations of data gathering alone: in addition to the technology for measurements, there must be a way to understand how the observations fit together and a guide to which kinds of data are most important to gather. It is here that mathematics can play a central role in the part of neuroscience that has to do with dynamics.

The mammalian nervous system is extraordinarily complicated, with interconnected neocortex, hippocampus, thalamus, cerebellum, and other structures, each of which has its own special anatomy and many substructures. The dynamics have different features in different substructures and can be highly localized in time and space. Although the ultimate goal is to understand the interconnections and how they organize mental processing, the current state of the art generally focuses on some small subset of the nervous system and the connections within that subset.

The Mathware

The framework for thinking about electrical activity in the brain was introduced in the 1950s by A. L. Hodgkin and A. F. Huxley and won them a Nobel Prize. A good mathematical introduction to these ideas is in [2]. The general theory that they developed plays a role in neurophysiology similar to the one that the Navier-Stokes equations play in fluid mehanics. The equations for each neuron are based on an analogy with electrical circuit theory. The primary equation is for conservation of charge across the membrane of a cell:

$$(1)\quad C\frac{dv}{dt} = -\sum I_{ion} + D\nabla^2 v - \sum I_{synapse} + I_{appl}\,.$$

Nancy Kopell is professor of mathematics and co-director of the Center for BioDynamics at Boston University. Her e-mail address is nk@math.bu.edu.

This article is based on the Josiah Willard Gibbs Lecture given at the Joint Meetings in San Antonio in January 1999.

The author's research was supported by NSF Grant DMS 9706694 and NIH Grant RO1 MH47150.

Here v is the voltage difference across the membrane and I_{ion} denotes an ionic current that passes across the cellular membrane. Each ionic current is like an element in parallel in an electric circuit. Each has its own special "battery", known as the electrochemical driving force, $v - v_R$, where v_R is the "reversal potential" for that current. Each current also has its own "conductance", which is the reciprocal of resistance. This conductance comes about because ionic channels in the membranes can open to allow the passage of particular kinds of ions in particular directions.

It is the dependence of the conductances on the voltage (and sometimes on other quantities as well) that adds more equations to the full set of Hodgkin-Huxley equations and makes them so mathematically challenging. The mathematical representation of the conductance usually involves terms describing up to two molecular "gates", each of which must be open for the current to flow. One of these gates, known as "activation", generally opens more as the voltage increases. The other, known as "inactivation", generally closes more as voltage increases. The full conductance is usually represented as a product $\overline{g}m^j h$, where $\overline{g}$ is the maximal conductance; m and h, both functions of time, are the fraction of the activation and inactivation gates that are open; and j is a positive integer. (For the classical sodium channel involved in producing spikes, $j = 3$.) Each ionic current is the product of the conductance and the driving force, i.e.,

$$I_{ion} = \overline{g}m^j h\,(v - v_R). \tag{2}$$

For each gate there is an auxilliary equation for the kinetics which describes how fast the fraction of open channels changes as membrane voltage changes. The equations for the two kinds of gates have the same form:

$$dx/dt = (x_\infty(v) - x)/\tau_x(v), \tag{3}$$

where $x = m$ or h and $x_\infty(v) = m_\infty(v)$ or $h_\infty(v)$. The difference between the two kinds of gates is in the form of the function $x_\infty(v)$. For most currents the function $m_\infty(v)$ is a monotone increasing function that saturates for low and high values of v, while $h_\infty(v)$ is monotone decreasing with v. These functions give the asymptotic values of the fraction of channels open for a given value of the voltage if the voltage were to be held forever at that value. The functions $\tau_x(v)$ give the voltage-dependent time constants for the gating. The full Hodgkin-Huxley equations are (1–3). Note that they are a family of equations with many possible currents and many possible parameters rather than a single fixed equation. Depending on the type of neuron, a mathematical representation may use as few as three or as many as a dozen currents, differing in reversal potential, maximal conductance, and gating kinetics. An extremely important feature of these equations is that there is a very large spread of time constants, leading to equations that have many simultaneous time scales.

The above description focuses on the change in voltage due to currents across the membrane. The full equation for dv/dt has extra terms as well. One of these reflects the fact that a neuron need not be concentrated near a spatial point, but may have processes (dendrites and axons) that spread out in space in a pattern that is characteristic of that cell. In such a case the voltage need not be the same throughout the cell, and the equation is a partial differential equation with diffusion across the parts of the cell. The other two terms in (1) reflect influences from the rest of the system to which that cell is connected. I_{syn} is the current that flows into the cell due to interactions with other cells. Like the intrinsic currents, they can be inward (depolarizing, or pushing the voltage toward the threshold for firing an action potential) or outward (hyperpolarizing, or pushing the voltage away from the threshold); it is the reversal potential of the current that determines whether it is inward or outward. I_{app} is a current that can be injected by a scientist studying the effects of perturbations.

Strategy and the Spherical Cow

Even without the architectural complexity of the entire brain, the dynamical complexity of even a piece of a brain is daunting. Here we get immediately to an issue of strategy: what sets of equations are appropriate to represent the dynamics? As mathematicians, we are drawn to the simplest descriptions that might be relevant. In the case of periodic dynamics, the simplest description is the one for uniform angular motion, namely,

$$\frac{d\theta}{dt} = \omega. \tag{4}$$

A person steeped in the culture of biology is much more likely to look for a representation that describes her reality: since there can be thousands of cells even in a small slice of tissue, and it may take on the order of fifty "compartments" for each cell to capture the spatial voltage differences, and the currents expressed can differ from cell to cell and between compartments, even when these equations are written as ordinary rather than partial, they can form an immensely large system [1].

Both of these approaches have major difficulties. The problem with the first is the analogue of the well-known "spherical cow": if the cow is not giving milk, its representation as a sphere makes it impossible to address the interesting questions about why that might be. In our case, (4) does not capture enough of the interesting detail to allow us to make relevant distinctions and to answer questions such as why the same group of cells is able to be coherent in multiple frequency ranges and whether the different frequency bands have different dynamical properties. The representa-

tion as "realistic" equations leads to a system that is about as complicated as the original wetware. Even with today's massive computers, which allow such mega systems to be solved numerically, very few people are able to look at the output and figure out what is going on.

The case studies discussed in this article present an intermediate strategy, one that focuses on the structure of the equations. Once some behavior has been identified in both the experiments and the large-scale simulations, the mathematics serves to clarify the behavior and identify the key dynamical features underlying that behavior. There are three case studies, two of which are short. Together they give some feel for the variety of questions and methods in this area of research. More details about the mathematical issues in these case studies can be found in chapters by N. Kopell and G. B. Ermentrout, and by J. Rubin and D. Terman in [3]. These chapters also contain many further references.

Synchrony and Cognition

The rhythms discussed in this section are usually referred to as gamma (30–80 Hz) and beta (12–30 Hz). These rhythms are mysterious and controversial, partly because they are technically difficult to induce and to spot; they can appear for epochs under a second and require appropriate data filtering to see. The stakes in the controversy are high, since some investigators have been marshaling evidence that these rhythms are associated with key cognitive states (attention, perception) as well as with memory. There is also tantalizing new evidence that pathologies in these rhythms are associated with thought disorders such as schizophrenia.

It remains an open question how the brain makes use of these rhythms. The most general, and controversial, idea is that synchrony facilitates coordination between distant parts of the nervous system that need to work together to create a movement or a percept [4]. Though rhythms are not needed to produce synchrony, there are data from P. König and collaborators (1995) that suggest that synchrony across a distance of more than a few millimeters does not happen in the absence of rhythms. An important concrete consequence of synchronous rhythms is plasticity: cells that fire within some window of time of one another are known to be able to increase the strength of the synaptic connections between them. This gives a way to encode the effects of past firing experience in future behavior. Synchronous activity also affects the gating of incoming signals and the strength of the downstream effects.

Though no one is yet sure of the uses of synchronous rhythms in the nervous system, there are many clues suggesting that a further study can provide insights about function. The clues come from the reproducible nature of the circumstances that evoke the rhythms, the spatial locations in which they are found, and the synchronization properties of the different rhythms. For example, the 1989 work by W. Singer and colleagues that led to a huge burst of scientific activity showed that exposing cats to moving bars of light can lead to gamma rhythms in the visual cortex. These rhythms synchronize over a few millimeters of tissue, and the synchronization properties have been reported to be related to the global geometric properties of the stimulus, e.g., whether the stimulus is one bar or two and whether the two bars are moved in the same direction. There is another scientific literature, reviewed in [5], dealing with sensory-motor tasks; here the investigators have mainly found rhythms in the lower-frequency beta range. The rhythms appear in the parts of the cortex expected to be participating in those sensory-motor tasks. The appearance of the rhythms is task dependent: e.g., P. R. Roelfsema and colleagues showed (1997) that at the end of a task, when the animal is rewarded, there was a shift from the beta rhythm to the still lower-frequency alpha rhythm (8–12 Hz).

The synchronization properties of the rhythms appear to be related to the frequency band. For example, in the work of Roelfsema and colleagues, one sees that both gamma and beta rhythms display very precise (within a millisecond or two) synchronization across distances of at least a few millimeters. The alpha rhythm, however, has some coherence but very sloppy synchronization, with changeable phase lags between pairs of sites.

These data raise a number of questions: what determines frequency? That is, why does the same collection of cells sometimes display a 50 Hz rhythm and sometimes an 18 Hz rhythm? Even more fundamental, what causes activity in some subsets of the brain to be, at least temporarily, coherent? Note that the rhythms we are considering here involve self-organization, not the result of outside coordination as by an orchestra conductor, making coherence a puzzle. Are the rhythms associated with distinct frequencies different in *structural* ways, not just in their time scale? Can the structure of rhythms tell us anything about how synchronization is accomplished across significant distances?

The above questions are very far from straightforward because of the complexity of the underlying equations and the large number of interacting dynamic processes with a large range of time scales. It is not only the properties of the individual cells that can affect the outcome but also the dynamics of the synapses that connect the cells. Thus the objective of the mathematics is to tease out how the biophysics of the cells and synapses work together to create coherent synchronous rhythms. An overview of related experimental

results and large-scale numerical modeling can be found in [1].

Slicing the Problem: Clues from the Hippocampus

To be able to think carefully about the origin of rhythms, it is helpful to slice away at the biological and mathematical complexity. For the biological complexity this is done quite literally by preparing slices of living tissue, usually from a rat; this allows an experimenter to deal with intrinsic dynamics without the confounding input from the rest of the brain. In the context of rhythms, one of the most studied slice preparations is from the hippocampus, believed to be important for the formation of memories.

A slice is still a complicated piece of wetware; it has many substructures and different kinds of cells. For the purpose of this case study, I am focusing on one substructure (the CA1 region) and lumping together the different kinds of cells into "excitatory" (E-cells) and "inhibitory" (I-cells). In general, excitatory cells use a transmitter whose downstream effect is to push cells toward threshold; i.e., make it easier for them to fire. Inhibitory cells use transmitters that have the opposite effect.The signals from these transmitters create currents that are mathematically similar to those in (2), but with the gates dependent on the voltage of the "presynaptic" cell, or cell giving input to the cell in question. Whether the presynaptic cell is inhibitory or excitatory is encoded in the mathematics by the value of the reversal potential of the driving force associated with the synaptic current in the postsynaptic cell.

In 1995 M. A. Whittington and collaborators showed that it is possible to induce a rat hippocampal slice to display a gamma rhythm. The most surprising thing about this work was that the rhythms were obtained when the signals from the excitatory cells were entirely blocked, leaving a network containing only inhibitory cells. This led to a pair of related mysteries: how can inhibition alone lead to synchronization, and why does the rhythm appear in the gamma frequency range? The latter question is especially puzzling, because the inhibitory neurons participating in the rhythm are individually capable of firing at any frequency between 0 and 200+ Hz.

Inhibition-Based Rhythms

An essential clue about the frequency came from pharmacology: it was found that substances applied to the slice that changed the frequency significantly were ones (e.g., barbiturates) known to change the time course of the decay of the inhibition. It is important to mention here that signals that come from synapses are not point events; they have time courses that turn out to be very important in the behavior of the system.

These and related data led various groups to look at networks of inhibitory neurons to try to understand both coherence and frequency. Some of the authors are C. van Vreeswick, L. Abbott, and G. B. Ermentrout; W. Gerstner, L. van Hemmen, and J. Cowan. Several different mathematical techniques were used; these are described in the chapter by Ermentrout and Kopell in [3]. The ideas can be most easily illustrated by the "spike response method" originated by Gerstner. The cells involved are relatively simple, almost linear between spikes. Thus they are reasonable candidates to be described by one of the simplest reductions of the Hodgkin-Huxley equations, the so-called "integrate and fire" model. These equations, which combine the gating variables with the voltage, have the form

$$C\frac{dv}{dt} = I - g_m(v - v_R) - I_{syn}; \tag{5}$$

the voltage is reset to some specified point when it reaches a specified threshold. The spike-response method works with an integrated version of this or a somewhat more general equation, using kernels to keep track of effects of spikes from a given cell and from the other cells in the network on the voltage of that cell.

From this work and related work it could be seen that spiking cells can form coherent rhythms with signals that are only inhibitory. The mathematics showed that, if the time course of the onset and offset of the inhibitory signal is long enough, inhibition alone can stabilize a synchronous solution. However, this made the mystery of gamma even deeper, since the range of frequencies for which stable synchrony is possible is much broader than the gamma range, even for a fixed, physiologically reasonable inhibitory signal.

The central clue turned out to be a major difference between networks of identical cells and networks in which the cells are heterogeneous, e.g., in drive, modeled by the parameter I in (5). X.-J. Wang and G. Buzsaki gave simulations to show that the coherence is fragile if there is even mild heterogeneity in the network. J. White, C. Chow, J. Ritt, and I analyzed the breakdown and found that it occurs in different ways in different parameter regimes. Furthermore, those parameter regimes are closely related to control of frequency: in different parts of parameter space, the frequency of the coupled network can be primarily governed by different parameters, including the decay rate of the inhibition, the membrane time constant, or the drive to the cell. The conclusion of the analysis was that coherence is most robust to disruption from heterogeneity when the system is in a parameter regime in which the period of the network is proportional to the time constant for the decay of the inhibition. Since a real slice network that exhibits this interneuron gamma rhythm is very likely to have some heterogeneity, this strongly suggests that one should expect the period to be proportional to the inhibition decay time. This

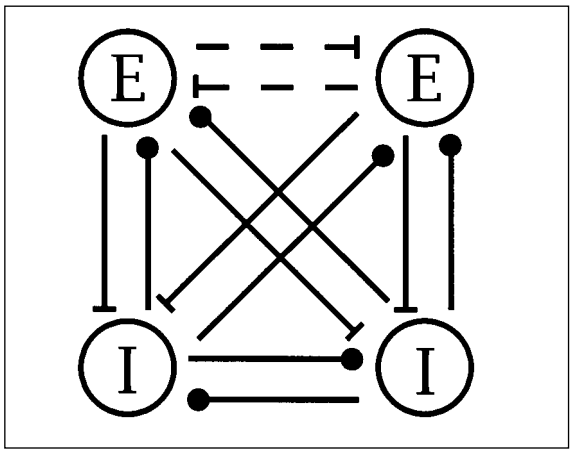

Figure 1. A simple circuit suitable for investigating synchronization of E- and I-cells. For the gamma rhythm, each I-cell receives input from both E-cells and the other I-cell. The E-cells get input from the I-cells, but not from the other E-cell (circuit without the dashed line); during the time that the beta rhythm is displayed, there are extra connections between the E-cells (dashed lines).

was verified using data from the lab of J. G. R. Jefferys. So the mathematics was able to reveal where the frequency of the gamma rhythm comes from and that it takes some heterogeneity to see it.

Transition between Two Rhythms

The same piece of tissue may be capable of multiple rhythms, with transitions between them, and the hippocampal slice provides a striking example, as shown by the 1997 work of Whittington and collaborators. If the slice is stimulated, provoking activity in both excitatory and inhibitory cells, it displays a transient gamma rhythm. The more surprising observation is that, at a higher level of stimulation, the rhythm starts off at gamma, undergoes a period of 150–200 ms (milliseconds) of incoherence, and then switches to a slower beta frequency rhythm. Such transitions have been seen using EEGs (electroencephalograms) to measure neural activity in humans and have been recently shown by the group of J. Gruzelier to correlate with the presentation of novel auditory stimuli.

In the slice, where it is possible to make detailed measurements from single cells, it can be seen that the beta rhythm is not just a slower version of the gamma rhythm. Indeed, it is a nested rhythm, with the inhibitory cells continuing to fire at a gamma frequency and the excitatory cells missing cycles. Furthermore, on the whole, they miss the same cycles, so that population averages reveal the slower rhythm.

The most significant thing about the transition of rhythms in the slice is that it has semipermanent traces. After stimulations of intensity high enough to produce the gamma to beta switch, later stimulation that would only have induced the gamma rhythm now is capable of inducing the entire gamma-to-beta sequence. Thus, beta gives a way to read back at a later time that some experience has happened.

The experimental manipulations yielded some evidence of what physiological changes are involved in this transition. One of these concerns is which ionic currents are expressed during each rhythm. Mathematically, this means which terms should be included in I_{ion} in (1). Some currents in the excitatory cells that are suppressed by the consequences of the stimulation return and participate during the beta phase. These currents are outward currents that pull the voltage away from the firing threshold and hence tend to slow down those cells. The more surprising and more permanent change is the growth of new functional connections among the cells. These connections are between pairs of excitatory cells and are fostered by the synchrony between those cells during the gamma phase. A much-used phrase, attributed to work of D. O. Hebb, is that "cells that fire together wire together."

One mathematical question is whether the above changes in ionic currents and synaptic connections are sufficient to produce the transition between the rhythms. This was recently addressed by R. D. Traub, Whittington, Ermentrout, and me. Understanding this required having a working definition of beta; indeed, one critical role of mathematics in science is the creation of idealizations that allows one to conceptualize what is in fact much more messy. In the case of beta, the working definition is that the I-cells synchronize at a gamma rhythm controlled essentially by the inhibition, while the E-cells synchronize at a subharmonic.

The simplest network that embodies this idea and allows one to address how synchrony comes about has two E- and two I-cells (Figure 1). It is easy to see that if the E-cells are unable to fire quickly enough to keep up with the I-cells, each will skip cycles; what is not clear is why the E-cells skip the *same* cycles and synchronize with one another. The added recurrent excitation (new connections between E-cells) are critical for this, but not in a straightforward way: as Gerstner, van Vreeswijk, and their collaborators have shown, connections between excitatory cells can actively foster antisynchrony or asynchrony between those cells. In this case, the extra excitation comes in a context in which the E-cells have the slow currents described above and the network has inhibition; these factors turn out to change the synchronizing properties of connections between the E-cells.

Long Distance Coordination and Feedback Loops

A central question associated with both beta and gamma rhythms is how synchronization is possible between sites with a conduction delay between

them that is a significant fraction (e.g., 20%–40%) of the period of the cycle. The long-distance coordination can be mimicked in the hippocampal slice by stimulating at two different sites in the slice, separated by a distance yielding a conduction delay of 6–8 ms. In simulations and experiments on stimulation-induced gamma rhythms in the hippocampal slice, Traub and collaborators (1996) noted that there was a strikingly regular structure in the order of the spikes when the two sites synchronized: the I-cells almost always displayed "doublets", i.e., a pair of spikes in each cycle; any parameter range in the simulations in which those doublets were absent led to nonsynchronous dynamics. The mathematical questions addressed by Ermentrout and me were how the synchronization comes about and what the role of the doublets is in this. There had been earlier (1991) simulations by König and T. B. Schillen showing that synchronization can take place in the presence of conduction delays. However, the models in those simulations used "activity" (essentially firing rate) as a variable instead of voltage; this kind of description is valid when one can average many spikes over a relevant period of time. In the current case, with the E-cells firing at most one spike per cycle, a rate model is not an appropriate description, and a different mathematical representation is needed. As discussed below, it is possible to develop low-dimensional maps (i.e., maps with a low-dimensional domain) that capture the essential parts of the physiology and screen out the rest.

A minimal network for analyzing long-distance connections lumps together all the cells at each site that fire synchronously. Such a network has one E- and one I-cell per site, a total of four cells. Each E-I pair oscillates: a spike from the E-cell induces a spike in the I cell, which inhibits the E-cell. The latter fires again when the inhibition wears off and, in the gamma rhythm, the I-cell waits for this excitation to spike again. The essential long-distance connections are the ones from the E- to I-cells. (Figure 2).

Analyzing networks of spiking cells: treating high-dimensional systems as a collection of low-dimensional maps. The simple four-cell network can be written as Hodgkin-Huxley equations, with extra equations for the synaptic currents. To include the minimal number of currents to get action potentials, plus the slow outward current in the E-cells, plus the equations for the synaptic gating variables requires a system of at least twenty highly nonlinear differential equations. Since it is almost impossible to analyze systems of that size, it is very helpful to find an analytical technique that captures the essence of the high-dimensional model in a faithful enough way that one can ask about effects of biophysical parameters. This cannot be done by reducing the number of cells or equations, since

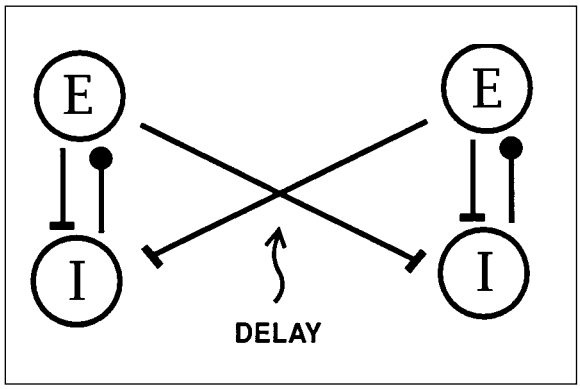

Figure 2. A simple circuit for investigating synchronization across distances. Each pair of E-I cells is an oscillator, with long connection from each E-cell to the distant I-cell, with signals arriving after a conduction delay.

the network and resulting equations already constitute a "minimal" model. The central idea for analyzing stability of a periodic network behavior is to construct a low-dimensional approximation to the Poincaré map, valid only in a neighborhood of a particular periodic trajectory.

The construction starts with a fixed order of the cells firing in the periodic trajectory, allowing some cells to be simultaneous. The neighborhood of allowable initial conditions and allowable parameter ranges is constrained so that the firing order is not changed. In principle the Poincaré map around a periodic trajectory has dimension only one less than that of the system. However, the spread of time scales now enters as a simplifying factor. In the relevant neighborhood of parameters and intitial conditions, most of the degrees of freedom of the system do not affect the timing of the spikes in a significant way—in general, they involve processes fast enough to relax away before another slower process determines a spike time. The construction of the map requires identifying the relevant degrees of freedom and ignoring the others. This will be illustrated below. So far it has not been proved that this method gives correct approximations, though predictions of the method have been tested successfully on the minimal biophysical networks and on ones that are large and biologically realistic. The conjecture is that the method works because it corresponds to identification of a center manifold in the dynamics, one containing the periodic trajectory in question.

Synchronization of gamma and beta: different feedback loops. Using the above minimal model (Figure 2), we look in the parameter range in which the first spike in the doublet of each I-cell is induced by excitation from the local E-cell; the second is a consequence of the distant excitation. With more excitable I-cells or stronger synapses, a single pulse of excitation can produce the doublet; however, such a doublet does not help with synchronization for reasons to be seen below.

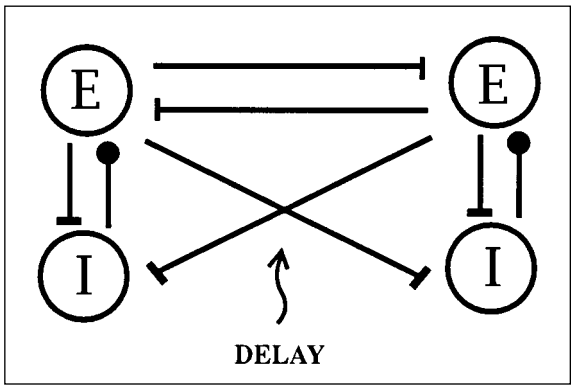

Figure 3. The same circuit as in Figure 2, with added long connections between the E-cells.

It turns out that, with these assumptions, there is only one important degree of freedom in the system: the time between the two E-cells on a given cycle. Nevertheless, it requires insight into the full network of Figure 2 to see where in the dynamics the synchronization signal is. The key feature turns out to be a property of many kinds of neurons, including the fast-firing I-cells of this network. Hodgkin-Huxley models of these cells have a property related to what is sometimes known as "relative refractory period". If the cell receives excitation shortly after it has fired, it may wait to fire again, with the delay to firing decreasing with the time between the first firing and the new excitation. In the doublets configuration the first spike in a cycle of the I-cell results from local excitation, and the second from the distant one. Thus the delay in the time to fire the second spike encodes information about the relative phases of the two E-cells, and the decreasing nature of that function turns out to synchronize the E-cells. The delay function, which depends on other inputs to the I-cells as well as on the intrinsic properties of those cells, can be computed numerically from the network. This gives a way to understand conceptually how changes in physiological parameters change the synchronization properties of the network by understanding how they affect the critical synchronization signal. More mathematical details and generalizations are in [6].

The minimal network described above is rigged, both in the architecture and parameter range, so that the doublet configuration is a solution; the mathematics then addresses the issue of stability. However, in a large, distributed network the I-cells can get excitation from many cells at many times, and even the existence of a solution with the doublet configuration is at issue. J. Karbowski and I recently addressed this question in the context of a large lattice of oscillators, with signals arriving at each I-cell from many neighbors at many times. Use of methods like those in statistical physics showed that the doublet configuration is a stable one for a significant parameter range. Furthermore, though there is a parameter range in which other solutions (e.g., one with a triplet configuration) are stable, the doublet configuration is stable over a larger parameter regime. The work suggested that the messiness of a real network might actually be helpful to synchronization: in the model distributed network some disorder in the arrival times of the signals increased the parameter range over which synchrony is stable.

When the local rhythms at the two sites are beat-skipping beta rhythms as described above, there are more dynamical processes that can be called into play to create the synchronization. For the gamma rhythm the essential synchronizing signal comes from the timing between the spikes of the doublets. For the beta rhythm the extra time scale associated to the slow outward current creates new nonlinear interactions with the potential to help the synchronization.

In recent work the contrast between the synchronization properties of gamma and beta was investigated by Ermentrout and me. Since the E-cells all synchronize in the idealized version of the beta rhythm as well as the gamma rhythm, the same anatomical network as above (Figure 2) can be used to analyze the long-distance synchronization. Once again the relevant approximation to the Poincaré map is one-dimensional, but the map changes, with the timing information coming from the distant cell interacting nonlinearly with timing information carried by the slow outward current. The result is that the beta rhythm can tolerate significantly longer conduction delays than can gamma and still synchronize robustly. This conclusion was tested by Traub in very large-scale models (as in [1]) and was shown to hold. Current work by S. R. Jones and collaborators is showing that other rhythms with still lower frequency need not be able to synchronize over distances, underscoring the need to look at how detailed biophysical structure—specifically which ionic currents are involved, what kinetics they have—translates into mathematical structure in the associated maps.

Within this framework Ermentrout and I also looked at the role of the long-range connections between the E-cells in the synchronizing process (Figure 3) for the beta rhythm. The analysis and simulations of the biophysically based networks show that these connections are not needed to create stable synchrony and indeed do not much change the dynamical behavior in a neighborhood of that trajectory. However, without those connections there are other stable solutions, and a critical role for the long E-E connections in the beta regime is to make other competing solutions unstable or nonexistent.

These results have implications for differential uses of the gamma and beta rhythms in the nervous system. They are consistent with data from electrophysiology on animals and human EEG recordings, suggesting that coherence in the

gamma range is used in the nervous system for relatively local processing (e.g., within the primary visual cortex), while the beta range is used for coordination among sites that have a longer conduction delay between them. Relevant data include work of A. von Stein and collaborators and Gruzelier and collaborators. The analysis shows why the ionic currents and network connectivity associated with the beta rhythm have the nonlinear properties appropriate to this behavior and why other rhythms can lack those properties.

Sleep Rhythms

We now focus on other parts of the nervous system to discuss rhythms that are associated with different stages of sleep. The thalamus is deep in the brain and has traditionally been considered the gateway between the sensory organs and the neocortex. Unlike the earlier descriptions of the thalamus as a relay station during wakefulness and a wall during sleep, it has been discovered that the thalamus is the origin of many complex dynamical behaviors that are state dependent.

From a mathematical point of view, the study of thalamic rhythms creates new issues, because, at least during sleep rhythms, many of the cells undergo "bursting". That is, they have a succession of spikes followed by a quiet period. In bursting cells there are more nonlinearities associated with the ionic currents, and the mathematical techniques that work for spiking cells are not in general usable. The major cellular players in the sleep rhythms are thalamocortical cells (TC) that are excitatory and reticular cells (RE) that are inhibitory; both are bursting neurons.

Two of the main rhythms in sleep are known as "spindling" and "delta". Spindling occurs in early and light sleep and has a population rhythm of 8–14 Hz. J. Rinzel and A. Destexhe have written extensively about models of this rhythm. Delta occurs in deep sleep. It is more regular and has a population rhythm of 1–4 Hz. These rhythms have very different dynamical stuctures, and the focus of this case study is the insights that mathematics provides in understanding the origins of those structures and the transitions between them.

If we temporarily bracket the physiology and just think about structure, there is a useful way to characterize the main differences between these two rhythms. The delta rhythm corresponds to synchrony of the TC cells, while spindling corresponds to "clustering" of the TC cells. In the latter the population of TC cells breaks down into subsets called clusters. Within a cluster the cells synchronize while the different clusters fire at different phases. In the clustered state the frequency of the population rhythm is several times that of any given cell; it is the frequency of a cell times the number of clusters. In the delta rhythm the TC cells fire on each cycle in a synchronous manner.

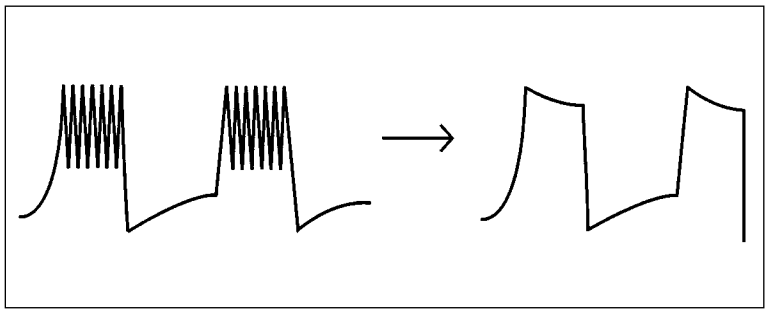

Figure 4. Schematic drawing of a bursting neuron with a set of spikes. If the spikes are removed, the resulting envelope has the shape of a relaxation oscillator.

In spindling each TC cell fires once per several population cycles, and the population rhythm is higher. Unlike the beta rhythm described above, in the spindling rhythm the excitatory cells miss different cycles, not the same cycle.

The network that produces these two rhythms is anatomically the same. The mathematical question is what kind of change gives rise to a transition between these rhythms or, equivalently, a transition between clustering and synchrony. As in the other case studies in this lecture, the physiology gives clues to the mathematics. In this case the clue is in what is happening to the inhibition on a cycle-by-cycle basis. For spindling it was shown by M. Steriade and his collaborators that the RE cells fire on each cycle. In the delta rhythm the RE cells fire only at the beginning of an episode of this rhythm and then are silent until another episode is induced by excitation from the cortex.

We know from the work in spiking cells that inhibition can be important to synchronization. But it can also be the enemy of synchronization. In the context of bursting cells, one can understand the complexities of inhibition in geometrical ways, especially by methods associated with singularly perturbed systems. Such methods have been developed by many people over the last decade. A recent review is the chapter by Rubin and Terman in [1]. Here I discuss a simple version of some of these ideas to give the flavor of these techniques.

When dealing with bursting neurons, one commonly used simplification is to airbrush out the individual spikes and to work with the envelope of the voltages of the burst. In such an envelope the voltage trace has some slowly changing periods, interspersed with rapid transitions (Figure 4). The geometrical methods used to study these systems exploit the difference in the time scales.

One simple idea that is useful even in large, complicated networks can be seen most easily in two dimensions. It deals with the effects of long-lasting inhibition. In singularly perturbed systems, parts of the trajectory hug the nullsurface (nullcline in two dimensions) of the fast variable. A constant source of inhibition changes the effective phase space by moving the voltage nullclines, which moves the trajectory to a different rest point

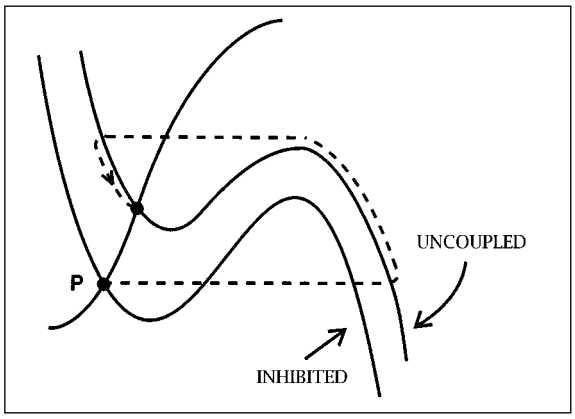

Figure 5. Post-inhibitory rebound. Inhibition lowers the nullcline of the voltage variable. If the inhibition is held long enough, the system goes toward the new stable rest point P. When the inhibition is released, the trajectory from the initial condition P moves around the original nullcline before returning to the old rest point.

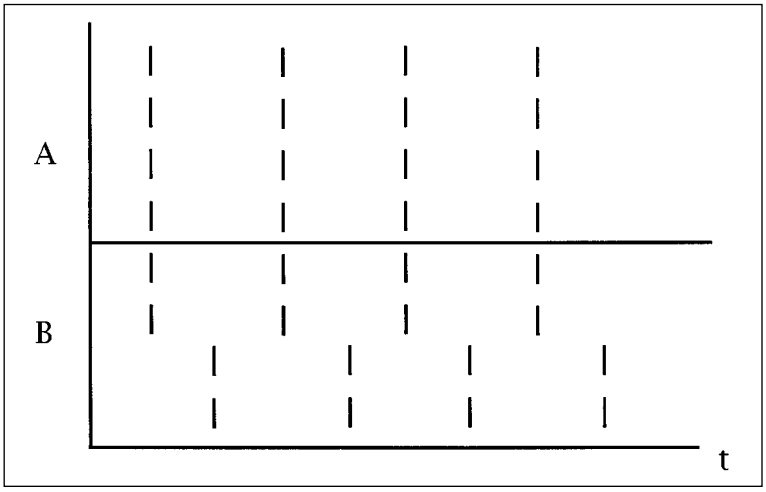

Figure 6. A schematic diagram showing synchrony and clustering in a network of 4 cells. In A the cells are synchronous; i.e., they fire at the same time. In B there are two clusters of 2 cells each, synchronous within a cluster but not between them. The population frequency of B is twice that of A.

(Figure 5). When the inhibition is released, the trajectory makes a long excursion corresponding to a burst before returning to rest. This is known as "post-inhibitory rebound". Now consider many cells, all receiving the same inhibition. They are each pushed to a neighborhood of the same point and, on release, follow nearby trajectories. Estimates, either in phase space or using time-based metrics, show that the trajectories are pulled closer together. Thus, simultaneous release from inhibition can be a synchronizing mechanism.

In some parameter regimes, inhibition need not be synchronizing and can lead to the formation of clusters (Figure 6). Whether there will be synchronization or some other network behavior is a subtle issue depending on interactions of time scales. Geometric methods give a powerful framework for understanding and predicting emergent network behavior, as Terman and Rubin have shown in recent work. These methods help to explain, for example, why fast onset of inhibition works actively to prevent synchrony.

We now use the mathematics to go back to the rhythms in the thalamocortical network. In that neuronal network there are two kinds of inhibition, known as $GABA_A$ and $GABA_B$. The first has a fast onset and a short-lived effect; the second has a much longer time to onset and a longer decay time. Insights from the mathematics suggests that a key feature in the transition from spindling to delta rhythms, i.e., clustering to synchrony, is the removal of the fast-onset inhibition, which prevents the synchrony of the TC cells. Indeed, data show that the cells providing the fast inhibition do not participate in the synchronous delta after the initial burst in some episode, while they do take part, cycle by cycle, in the spindling rhythm. How those cells might get functionally removed from the network is discussed in [7]. The role of the mathematics is to indicate what kinds of changes in the network are sufficient to produce the transition. This kind of "functional reorganization", in which networks switch to very different behavior, appears to be widespread in neural dynamics and could be important in pathology. For example, a model by Destexhe with a similar flavor describes the transition from normal thalamic dynamics to certain kinds of seizure activity.

Dynamics in a Spatially Extended Neuron

The rhythms associated with cognition and sleep are by no means all the rhythms in the nervous system. The last example comes from a different part of the nervous system and illustrates very different mathematical ideas. The neurons involved in this example, from the subcortical structure called the substantia nigra, provide the neurotransmitter dopamine to other parts of the brain; degeneration of these cells leads to Parkinson's disease.

Unlike many cells in the brain, these cells never appear to be in relay mode, faithfully reporting on their inputs. Instead, they have a set of predetermined patterns in their dynamics. The switches among these patterns are mysterious, but so is the formation of the patterns themselves. For example, a pattern of bursting associated with reward signals has not been replicated by injecting any pattern of electrical signals into the cell body of the neuron in a slice. Some current thinking about the origin of the dynamical patterns focuses on the dendrites, the part of the neuron where most of the inputs arrive. So far in this article I have been talking about cells as if the spatial structure does not matter; in this case the spatial structure may be a critical part of the dynamics.

C. J. Wilson and J. C. Callaway have done experiments using a combination of electrophysiology and imaging of calcium concentrations to help illuminate the origin of the cellular dynamics. They found that pharmacologically isolated cells in a slice can oscillate at 1–2 Hz. Furthermore, there is evidence that there are differences between the cell body and the dendrites in the rates at which calcium goes into and out of the cell. This has led them to a model in which the parts of the cell are treated as separate oscillators, with a gradient in natural (uncoupled) frequency along the dendrite. If we think of the dendrite as broken up into spatial compartments, then a rough mathematical description of a single cell is that of a chain of oscillators with a gradient in frequency. The coupling is electrical rather than via chemical synapses; mathematically, this is written as a discretization of a Laplacian.

Simulations of this model were able to reproduce various experimentally produced dynamical behaviors, including the observation that the frequency slows down over time. The latter phenomenon occurs in many neural systems but is usually associated with special ionic currents such as the slow outward current discussed in connection with the beta rhythm. In this model the description of the local oscillators has no feature that could account for the slowdown. The mathematical issue here is to understand the origin of the slowdown, which is believed by some to be related to the origin of the bursting *in vivo*.

The simple, biophysically based model belongs to a class of equations of the form

$$\frac{dv_i}{dt} = f_i(v_i, w_i) + d(v_{i+1} - 2v_i + v_{i+1})$$
$$\frac{dw_i}{dt} = \varepsilon g_i(v_i, w_i).$$

Here the uncoupled equations for the i^{th} compartment describe a relaxation oscillator, with the fast variable v_i denoting the voltage of that compartment and w_i the concentration of calcium in the i^{th} compartment. The uncoupled oscillators have a gradient in frequency along the chain, and d is much greater than 1.

To get a sense of how the spatial structure of the network affects the transient frequency behavior, G. Medvedev and I recently considered a simple version of equations of this type, namely, the van der Pol equations in Lienard form. He and I discovered that the equations had some unexpected mathematical properties. The coupling between the compartments is very strong, so that the voltages are essentially the same throughout the chain. Nevertheless, because the different compartments in the chain have calcium dynamics with different rates, the w_i are not the same in the different compartments. The most unexpected outcome was the observation that, for very large

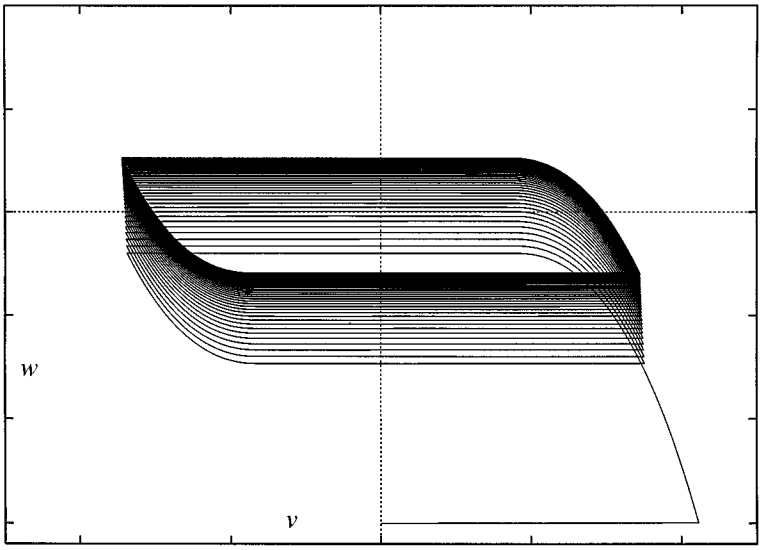

Figure 7. A trajectory for equations describing two compartments coupled strongly with electrical coupling. The plot shows v and w for one of the compartments. Notice that after the initial part of the trajectory, the phase plane portrait displays the trajectory moving slowly along a 1-parameter family of periodic orbits.

electrical coupling ($d > 1$), the system has an invariant manifold on which it behaves like a perturbed conservative system; in the $d \to \infty$ limit, the system has a family of periodic solutions, and for $d < \infty$ there is a drift along the invariant manifold toward a final state (Figure 7), accompanied by a slowly changing frequency, as in the experimental data. Thus the mathematics shows how the strong electrical coupling translates into a Lyapunov function defined on an invariant manifold and how this spatial interaction between the compartments produces the observed frequency modulation. The mathematics also makes the unintuitive prediction that the stronger the coupling, the slower the drift toward the final state.

What Next?

There are several kinds of challenges I see in trying to understand neural dynamics and their functional importance. The first is to build a vocabulary of examples, in the spirit of the ones I discussed, to try to understand in a clear and conceptual way the mechanisms for producing intrinsic dynamics in parts of the nervous system. This is best done in the context of specific examples, especially since the biology gives generous hints about mathematical structure. The biological structure (time scales of synapses and intrinsic currents, spatial effects, slow modulators, etc.) creates mathematical structure (invariant manifolds, Lyapunov functions, singularly perturbed systems, etc.), which can, in turn, illuminate the biology.

The second challenge is to take the ideas understood in idealized situations and explore

whether they work in larger and more realistic but messier circumstances. This may involve statistical and probabilistic notions, a kind of statistical mechanics for neurons. I believe it will be critical for such a theory to respect the structure of the smaller idealized building blocks, not simply mimic the concepts developed for statistical physics. An example in which a simple and rigorous analysis was "scaled up" in this way was given above in the case study on synchronization of the gamma rhythm across distances.

Finally, in a working brain that is listening to a lecture, taking notes, daydreaming, or thinking about the next theorem to be proved, parts of the brain are not isolated from one another and the outside world. A very large challenge is to make use of knowledge about dynamics of parts of the nervous system to understand what role dynamics plays in filtering and processing the inputs from other parts of the brain and the outside world, including inputs that are not periodic. In all of these questions, by exploring the origin and functional implications of dynamics, mathematics can play a significant role in helping us to understand how the brain organizes itself.

References

[1] R. D. Traub, J. G. R. Jefferys, and M. A. Whittington, *Fast Oscillations in Cortical Circuits*, MIT Press, Cambridge, MA, 1999.

[2] J. Rinzel and G. B. Ermentrout, Analysis of neural excitability and oscillations, *Methods of Neuronal Modeling*, 2nd ed. (C. Koch and I. Segev, eds.), 1998, MIT Press, Cambridge, MA, pp. 251–292.

[3] B. Fiedler, G. Iooss, and N. Kopell, eds., *Handbook of Dynamical Systems*, vol. 3: Toward Applications, Elsevier, Amsterdam, 2000.

[4] *Neuron* **24** (1999), Cell Press (special volume devoted to reviews on the binding problem).

[5] S. F. Farmer, Rhythmicity, synchronization and binding in human and primate motor systems, *J. Physiol.* **509** (1998), 3–14.

[6] G. B. Ermentrout and N. Kopell, Fine structure of neural spiking and synchronization in the presence of conduction delays, *Proc. Nat. Acad. Sci. USA* **95** (1998), 1259–1264.

[7] D. Terman, A. Bose, and N. Kopell, Functional reorganization in thalamocortical networks: Transition between spindling and delta sleep rhythms, *Proc. Nat. Acad. Sci. USA* **93** (1996), 15417–15422.