



Localized Bumps of Activity Sustained by Inhibition in a Two-Layer Thalamic Network

JONATHAN RUBIN

The Ohio State University; University of Pittsburgh

rubin@math.pitt.edu

DAVID TERMAN

The Ohio State University

terman@math.ohio-state.edu

CARSON CHOW

University of Pittsburgh

ccc@math.pitt.edu

Received June 29, 2000 ; Revised December 19, 2000 ; Accepted March 15, 2001

Action Editor: Misha Tsodyks

Abstract. Based on head direction experiments in rats, the existence of localized bumps of thalamic activity has been proposed. We computationally demonstrate the existence of a novel class of localized bump solutions in a two-layer conductance-based thalamic network and analyze the mechanisms behind these stable patterns. In contrast to previous models of bump activity, here inhibition plays a crucial role in initially spreading neuronal firing and in subsequently sustaining it. In our model, we incorporate local strong, fast GABA_A inhibition and diffuse weak, slow GABA_B inhibition, based on previous biophysical experiments. These forms of inhibition contribute in different, yet complementary, ways to the observed pattern formation.

Keywords: localized activity, head direction cells, thalamus, conductance-based model, synaptic coupling

1. Introduction

Spatially localized, sustained neuronal activity in the form of a bump attractor has been proposed for the head-direction system in mammals (Skaggs et al., 1995; Blair, 1996; Redish et al., 1996; Zhang, 1996), working memory (Wilson and Cowan, 1973; Amit and Brunel, 1997; Rao et al., 1999), and orientation selectivity in models of the visual system (for example, see Somers et al., 1995; Hansel and Sompolinsky, 1998, and the references therein). In these previous models, stable bumps are sustained by local recurrent excita-

tions and localized by lateral inhibition (Amari, 1977; Ermentrout, 1998).

Experiments on rats have found sets of neurons that respond only when the animal faces a given direction. It has been proposed that these cells fire collectively within a bump that moves in concert with the head around a network with the topology of a ring. These head-direction (HD) cells are found in the postsubiculum (PoS) (Taube et al., 1990) and in the anterior thalamic nuclei (ATN) (Taube, 1995), which are at least partially recurrently connected (Van Groen and Wyss, 1990, 1995). In particular, it appears that head-direction

activity from the anterior dorsal thalamic nucleus (AD) is transmitted to the PoS, which in turn refines AD activity (Goodridge and Taube, 1997). However, it is not understood how the ATN, which lacks recurrent excitatory connections, could support localized activity. Here, we demonstrate that a localized bump may be sustained by the combined inhibitory and excitatory connections in a thalamic network.

This work is done in a two-layer thalamic network model composed of sets of conductance-based ordinary differential equations for a number of thalamocortical relay (TC) and thalamic reticular (RE) cells and the synaptic connections between them. Thanks to the properties of RE and TC cells and their connections, the network can exhibit temporally periodic, spatially nonlocalized activity after transient stimulation, even when no individual cells are spontaneously oscillatory, as seen previously in simulations and analysis (Destexhe et al., 1993; Golomb et al., 1994; Wang et al., 1995; Destexhe and Sejnowski, 1996; Terman et al., 1996, 1998; Rubin and Terman, 2000; Rubin and Terman, in press). In the resultant rhythms, activity is maintained through a repeated sequence in which RE cells fire in response to excitation from TC cells, TC cells are inhibited by RE cells, and then TC cells fire through postinhibitory rebound (PIR). Without inhibition, activity could not spread through the network, which includes no recurrent excitation (e.g., Steriade et al., 1993, 1997).

In our model networks, we include two types of inhibition from RE cells to TC cells—namely, fast, local GABA_A inhibition and slow, long-range GABA_B inhibition (Sohal and Huguenard, 1998); we explore a variety of architectures of GABA_B connections. Unlike previously analyzed thalamic sleep rhythms, the bump solutions reported here involve activity only over a restricted, connectively localized portion of the cells in the network. Moreover, the firing patterns of cells within the bumps can be quite disorganized, as seen in the wake of some propagating activity by Rinzel et al. (1998), which contrasts with previously observed synchronized or clustered thalamic rhythms.

However, the thalamic bumps do share some characteristics with these other rhythms, and these distinguish them from bumps seen previously in other computational models. In particular, PIR is necessary for the existence of these bumps. The absence of lateral excitatory connections in the thalamus implies that excitation alone is completely inadequate for producing or maintaining thalamic bumps. Instead, inhibition

and excitation interact to sustain them, and inhibition largely shapes them. Classic rate models (e.g., Wilson and Cowan, 1973; Amari, 1977) and some more recent integrate-and-fire models (e.g., Somers et al., 1995; Laing and Chow, forthcoming), require lateral excitatory connections for bump formation and do not allow for the rich temporal structure of inhibitory currents that our thalamic bumps require (see Ermentrout, 1998). The bumps in this work also do not require sustained external input to stay active, as needed in other models (Somers et al., 1995; Hansel and Sompolinsky, 1998).

Our aim is to demonstrate how certain basic intrinsic and synaptic features of an RE-TC network interact to produce sustained, localized bumps of activity. Thus, the model network presented in Section 2 consists of a simplified version of an RE-TC network. If one considers the TC cells to represent the HD cells in the rat ATN, which receive RE inputs and project to the PoS (Kultas-Ilinsky et al., 1995; Shibata, 1992; Van Groen and Wyss, 1995), then the mechanisms that generate localized bumps of sustained TC activity in this model may be relevant in explaining experimentally observed firing properties of these HD cells. In Section 3.1.1, we mention specifics about the coupling used in our simulations, while Appendix A is devoted to further numerical issues, including actual parameter values used. Section 3.1 presents results of the simulations. Here we describe the spread of activity, its localization, and the activity within the bumps that arise. The results in Section 3.1.5 show that local transient excitation to cells in either population in the network can initiate bumps of activity, and in some cases can translate established bumps, while a more widespread excitation to either population terminates activity through a synchronization effect.

Biophysical features of our thalamic model network are responsible for the particular characteristics of its activity patterns. In Section 3.2, we analyze how these details sustain and shape this activity. In particular, we compute an estimate for the size of a bump that will develop in a network, depending on intrinsic and synaptic cellular parameters. This involves computing the level of inhibition that must be present at the edges of the bump. We also consider mechanisms for bump termination and translation, as would be necessary in a head direction system. We conclude in Section 4 with a summary of our findings and a further contrast of these results with studies of activity bumps in other models.

2. Methods: Models for Cells and Synaptic Coupling

We performed simulations with a recent version of XPPAUT (Ermentrout, 2000). Certain aspects of the code are discussed in Appendix A. The model network we simulated and analyzed includes two populations of oscillators, an excitatory population corresponding to TC cells and an inhibitory population corresponding to RE cells. We refer to these two populations as E -cells and J -cells, respectively. The model also incorporates synaptic connections between these populations as well as cortical inputs they receive.

Individual cells are represented by conductance-based single-compartment models (Destexhe et al., 1993; Golomb et al., 1994). The equations of each E -cell are

$$\begin{aligned} v' &= -I_T(v, y) - I_L(v) - I_A - I_B - I_{ctx} \\ y' &= \phi(y_\infty(v) - y)/\tau_y(v). \end{aligned} \quad (1)$$

The equations of each J -cell are

$$\begin{aligned} w' &= -I_T(w, z) - I_L(w) - I_{A'} - I_E - I_{ctx'} \\ z' &= \psi(z_\infty(w) - z)/\tau_z(w). \end{aligned} \quad (2)$$

The variables v , w in (1), (2) represent membrane potentials. The terms I_T , I_L are intrinsic currents, namely a low-threshold T -type calcium current and a leak current. They are given in (1) by $I_T(v, y) = g_{Ca}m_\infty^2(v)y(v - v_{Ca})$ and $I_L = g_L(v - v_L)$ with similar forms in (2). The variables y , z denote the inactivation levels of the corresponding T -type currents. Explicit formulas for the nonlinear voltage-dependent functions m_∞ , y_∞ , z_∞ , τ_y , τ_z are given, along with parameter values, in Appendix A. These are chosen such that without coupling, E - and J -cells are silent.

The other currents in the v equation in (1) and the w equation in (2) represent inputs to the cells. The E -cells receive GABA_A and GABA_B inhibition from the J -cells, with the currents generated by these inputs represented in (1) by I_A and I_B , respectively. The J -cells receive GABA_A inhibition $I_{A'}$ from other J -cells and AMPA excitation I_E from E -cells. The currents I_{ctx} , $I_{ctx'}$ denote excitatory input from the cortex. The forms of all of these coupling currents appear in Appendix A.

The existence of two types of RE cells has been reported (Cox et al., 1996, 1997). One is coupled to TC cells only via weak connections extending

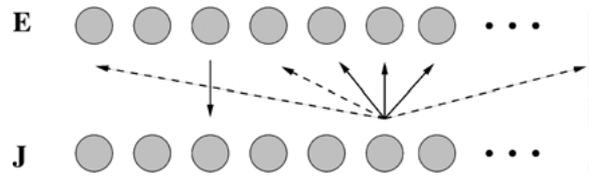


Figure 1. Model thalamic network: example connectivity diagram. Solid lines from J to E -cells represent GABA_A inhibition, while dashed lines represent weaker, more widespread GABA_B inhibition, or ticklers; in this illustration, ticklers are not uniformly distributed. Solid line from E to J represents excitation. Connections between other cells are analogous to the ones shown.

sparingly over a large radius, approximately five times the footprint of the other group's inhibitory connections; these are called ticklers by Sohal and Huguenard (1998). The two types of RE cells have similar intrinsic properties. Hence, following Sohal and Huguenard (1998), we model the corresponding two types of connections to E -cells as coming from a single J -cell population.

The GABA_B connections in (1) in our model represent the ticklers, as shown in Fig. 1. Note that GABA_A turns on and off on much faster time scales than GABA_B (Golomb et al., 1994; Huguenard and Prince, 1994; Destexhe and Sejnowski, 1995). Correspondingly, we omit long-range GABA_A connections in our model, since these will be weak and wear off quickly (Cox et al., 1997). We also omit any stronger local GABA_B connections, since Golomb et al. (1996) found these to be of limited importance for related rhythms.

The curves where $v' = 0$ and $y' = 0$ hold in (1), called the v - and y -nullcline respectively, are shown in Fig. 2; the w - and z -nullclines, also shown, are qualitatively similar. We choose ϕ , ψ as relatively small parameters. Hence, each E -cell's (v, y) values generally lie on the left or right branch of its v -nullcline, except for fast jumps between these branches. When a cell has voltage and inactivation values (v, y) on the left (right) branch of its v -nullcline, we say that it lies in the *silent* (*active*) phase. Jumping up from the silent to the active phase corresponds to action potential generation or firing (Rubin and Terman, forthcoming). Note in Fig. 2 that the v - and y -nullclines intersect on the left branch of the v -nullcline. Their intersection yields a rest state corresponding to the state of uncoupled E -cells without external input. Similar statements hold for J -cells.

As discussed in Appendix A, GABA_A input affects E -cell v -nullclines as shown in Fig. 2; GABA_B input has a similar, albeit weaker, effect. This implies that if

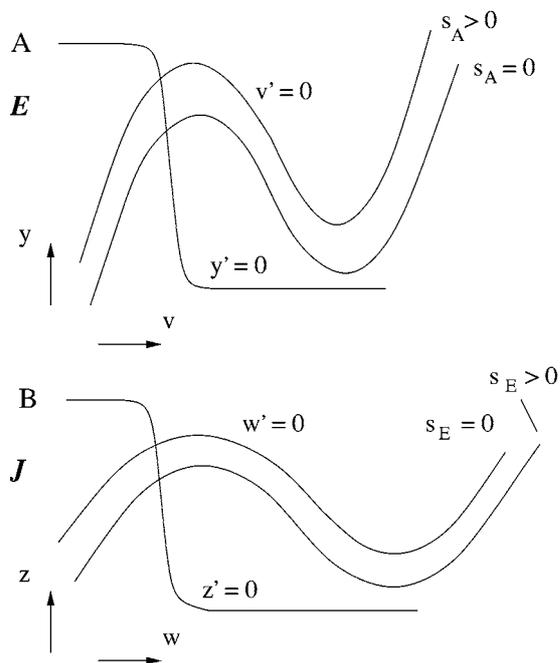


Figure 2. Nullclines for A) E -cells and B) J -cells. Inhibition raises an E -cell's nullcline while excitation lowers a J -cell's nullcline.

an E -cell receives sufficiently strong GABA_A inhibition for a sufficiently long time, then its T -current deactivates sufficiently that it can fire by postinhibitory rebound (PIR) upon release from inhibition, as shown in Fig. 9 below and discussed in Section 3.2.1 (see also Rubin and Terman, forthcoming). Such generation of action potentials by PIR plays a crucial role in producing the activity patterns discussed in this work.

E -cell excitatory input to J -cells affects their w -nullclines as shown in Fig. 2. Hence, sufficiently strong excitation to a J -cell at rest causes this J -cell to generate an action potential. Cortical excitation can analogously enable the E - and J -cells that it impacts to fire, if the target cells are sufficiently close to their rest states.

Remarks:

1. Note that we omit the sodium current in (1), (2). Hence, the active phase of each cell corresponds to a calcium-induced burst, with sodium spikes omitted; sample voltage traces from an E -cell and a corresponding J -cell are shown in Fig. 3. Frequencies of firing discussed below thus represent interburst frequencies.

2. We will use the term *synchrony* among coupled cells to denote simultaneous or near-simultaneous firing (see Bose et al., 2000, for a precise notion of near-simultaneity). When we refer to *asynchrony* among a group of active cells, we mean that all cells in the group are firing within some time interval, but they do so at different times in the interval.

3. Simulations

3.1. Bumps of Activity in a Variety of Coupling Architectures

We generated sustained, localized bumps of activity in numerical simulations of (1), (2) with each of the architectures discussed below and in Appendix A. In these simulations, we used chains of 50 E -cells and 50 J -cells; for clarity, we can think of the E - and J -cell populations as linear chains of ordered cells. As mentioned in the Discussion (Section 5), these orderings are based on connectivity but do not necessarily correspond to spatial positions.

We set up networks in which the GABA_A inhibition to cell E_i comes from the three closest J -cells—namely, J_{i-1} , J_i , J_{i+1} ; we say that the GABA_A inhibition from the J -cells to the E -cells has a footprint width of three. In (2) in our simulations, each J -cell receives excitation from only one E -cell, as in the classic Amari model (Amari, 1977); see Fig. 1. That is, the excitation in the network has a narrower footprint than the fast inhibition across the layers, as observed in thalamic networks (Sohal and Huguenard, 1998). Also in (2), each J -cell sends GABA_A inhibition to its two nearest neighbors (but not to itself). For all fast inhibitory and excitatory connections in the network, we scale the conductances of inputs to cells at the ends of the chains for uniformity. While the exact footprint widths used are not important for our results, the relative footprint sizes shape the activity patterns we observe.

In our simulations, we observed bumps in several different architectures of GABA_B tickler connections. These appear in Figs. 4, 5, and 6. In an *off-center* architecture (not shown), each J -cell sends ticklers to a small group of E -cells a certain distance (along the chain) away, with no connections to closer E -cells. In a *uniform* architecture, each J -cell sends ticklers to a uniformly distributed subset of E -cells within a footprint width of 29 (i.e., a subset of the 14 E -cells to either side of the nearest E -cell for each J -cell). We use this exaggerated tickler footprint width in all of our networks

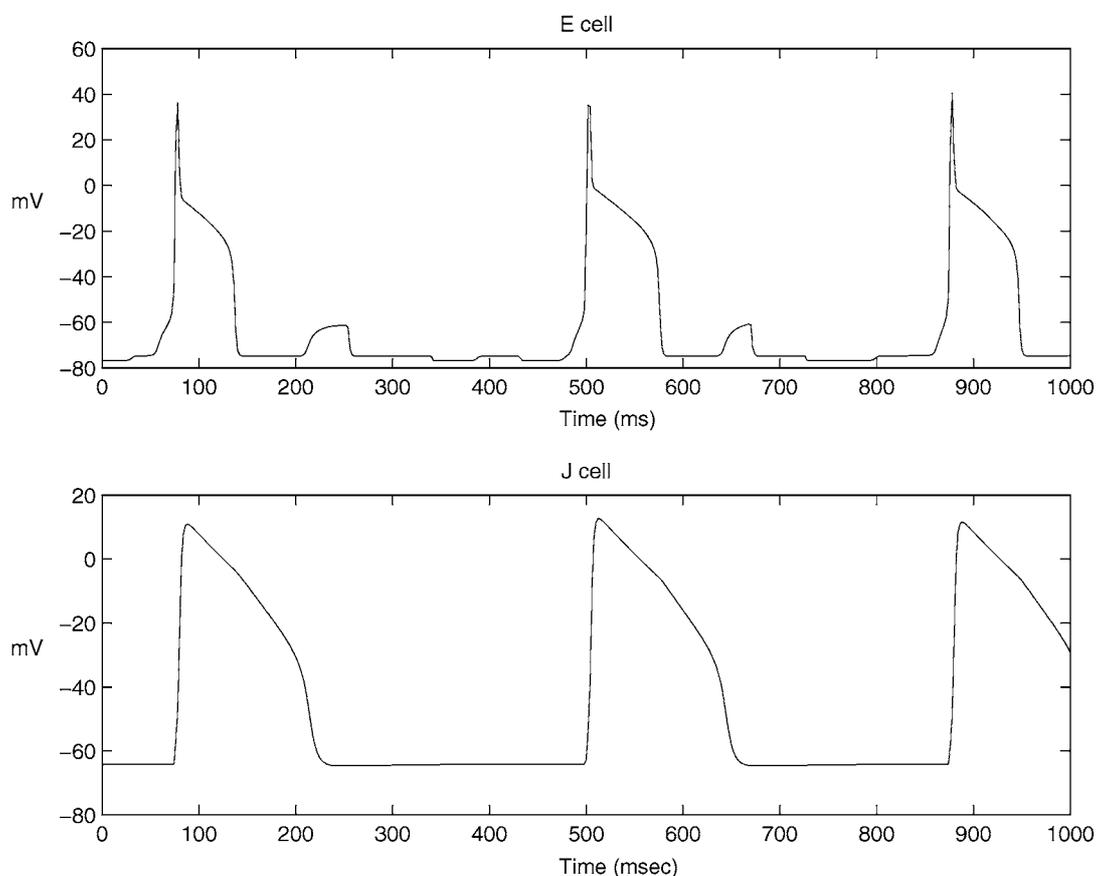


Figure 3. Voltage versus time for an *E*-cell and the corresponding *J*-cell in a simulation of the uniform architecture. Note the relatively fast transitions between prolonged silent and active phases. Note also that after the *E*-cell fires, it receives GABA_A input from the responding *J*-cell. This causes the *E*-cell to make a fast jump to a lower right branch, decreasing its voltage towards 0 mV.

to produce wider bumps for illustrative purposes; however, this is not necessary for bump generation. A bump of activity in a uniform architecture is shown in Fig. 4.

In a *balanced random* architecture, each *E*-cell receives five tickler inputs coming from *J*-cells randomly selected from within ± 14 places of it. As in Sohal and Huguenard (1998), cells with identical numbers of tickler inputs have identical maximal GABA_B conductances. We scale the conductances of inputs to *E*-cells near the ends of the *E*-chain so that the maximal tickler input strength to all *E*-cells is identical. A bump of activity in a balanced random architecture is shown in Fig. 5.

In an *unbalanced random* architecture, each *J*-cell sends out five tickler inputs to a random set of *E*-cells within ± 14 places of it. A *J*-cell may send more than one of its ticklers to the same *E*-cell. A bump of activity in an unbalanced random architecture is shown in

Fig. 6. We did not scale the inputs to *E*-cells near the ends of the *E*-chain for this architecture because by design, different *E*-cells receive input from different numbers of ticklers. Note in Fig. 6, however, that we simulated sufficiently large networks that the bumps did not interact with the ends of the chains of cells (also see Appendix A).

3.2. Spread of Activity

Initial conditions were selected with all *E*-cells at rest. All *J*-cells also began at rest except for a small subset of cells in the middle of the chain; some of these were started in the active phase and others were started in the silent phase but not at rest. Alternating excitation and PIR spread activity from the central *J*-cells. At the start of a simulation, the small group of active *J*-cells

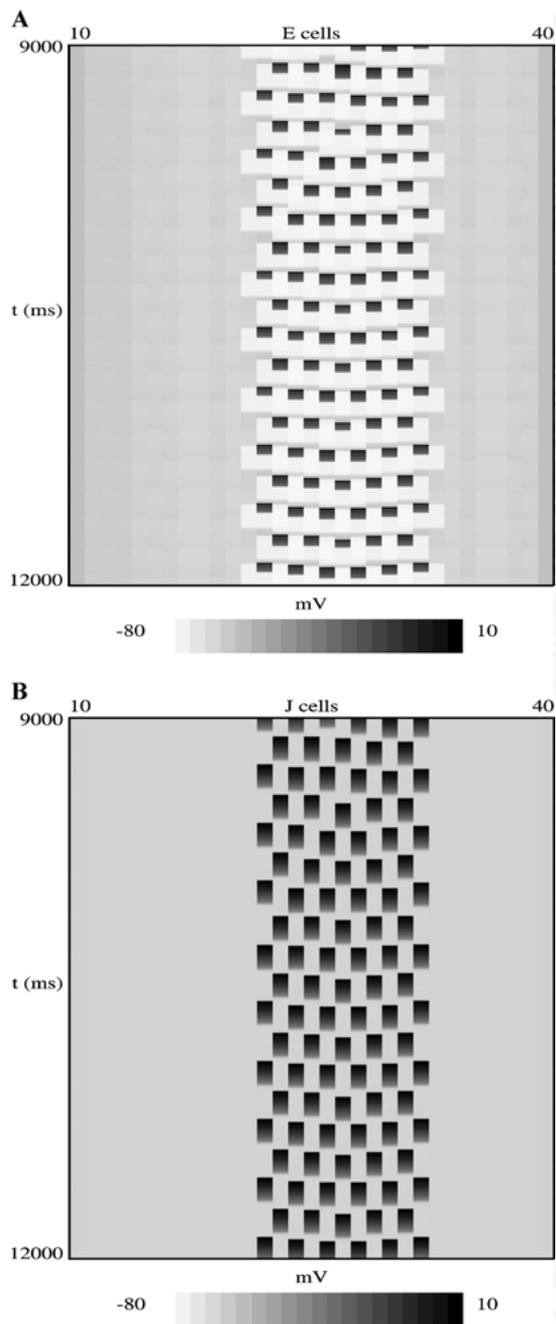


Figure 4. Activity bumps in **A** *E*-cells and **B** *J*-cells in the uniform architecture. Each plot shows voltage levels for cells 10 to 40 in the center of a chain of 50 cells, with the grey scale shown at the bottoms of the plots. We have cut off the scale at 10 mV, although voltage may transiently reach higher levels (see Fig. 3). Initially, cells J_{23} , J_{24} were excited. These plots show activity after the resulting bump of width 11 was established, with time (in milliseconds) evolving downwards. Note that *E*-cells have a shorter firing duration than *J*-cells and that *E*-cells adjacent to the bump become hyperpolarized but do not rebound. Parameter values are given in Appendix A.

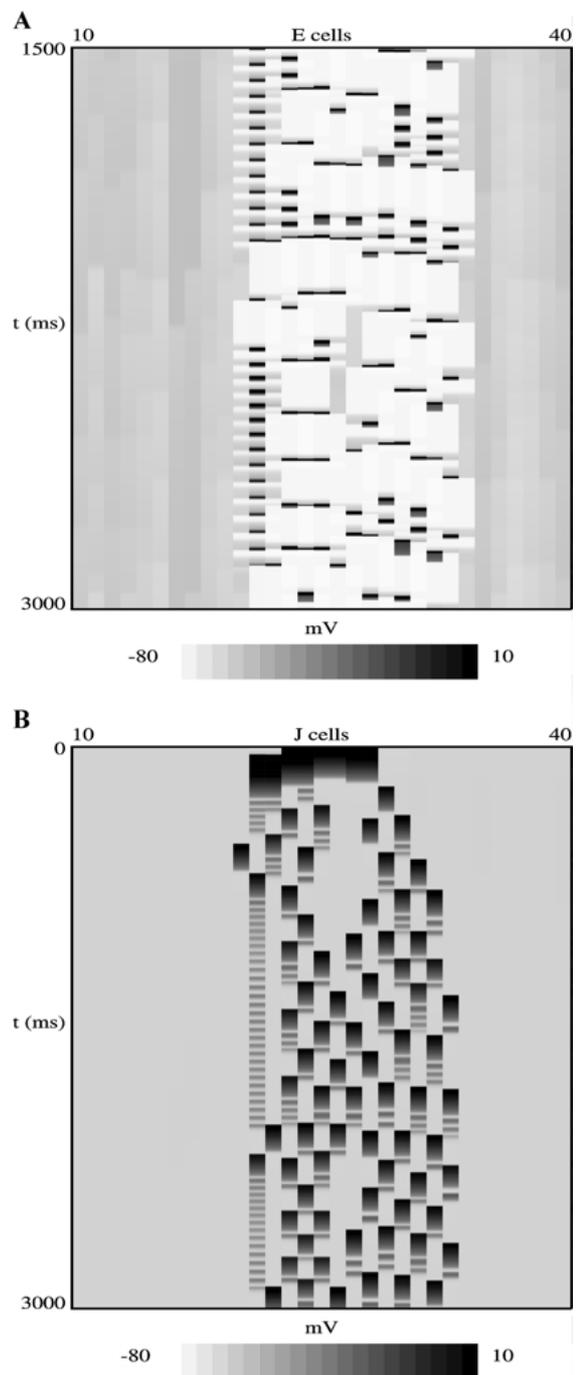


Figure 5. Activity bumps in **A** *E*-cells and **B** *J*-cells in the balanced random architecture. Each plot shows voltage levels for cells 10 to 40 in the center of a chain of 50 cells. B shows how an activity bump of width 13 developed directly from the start of simulation with a central group of excited *J*-cells; A shows the second half of the simulation shown in B. Note the relative disorganization in firing times (especially of *E*-cells, which have a shorter firing duration than *J*-cells). Parameter values are given in Appendix A.

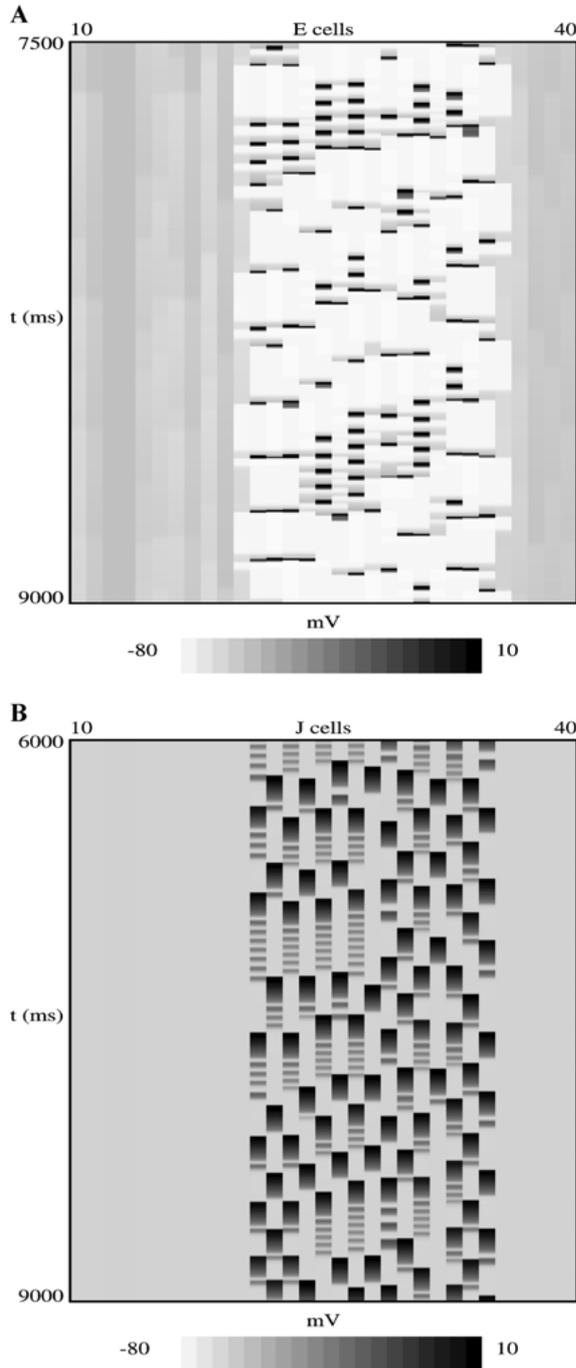


Figure 6. Activity bumps in **A** *E*-cells and **B** *J*-cells in the unbalanced random architecture. Each plot shows voltage levels for cells 10 to 40 in the center of a chain of 50 cells. These plots show activity after a bump of width 15 was established from initial conditions with cells J_{23} , J_{24} active. The location of activity showed a slight rightward drift before stabilizing. Note the difference in time scales, selected because *E*-cells have a shorter firing duration than *J*-cells. Parameter values are given in Appendix A.

send GABA_A inhibition to the *E*-cells to which they are coupled. Since the GABA_A wears off quickly after the *J*-cells become inactive, the *E*-cells are able to fire by PIR. Each *E*-cell sends excitation back to its partner *J*-cell, which fires in response if parameters are chosen appropriately. With the model given in the Section 2, GABA_A starts to decay soon after the *J*-cells become inactive, such that *J*-cells have little recovery time before being excited. Thus, for the *J*-cells to respond, parameters must be chosen such that the excitation from the *E*-cells is sufficiently strong; the turn-on of excitation is not too fast; and the *J*-cells need very little recovery time before they can fire again (Rubin and Terman, 2000). This is not difficult to achieve in simulations.

Since each *J*-cell inhibits three *E*-cells, the number of *E*-cells that rebound exceeds the number of *J*-cells that are initially active. Since each *E*-cell excites a distinct *J*-cell, after one such cycle the number of active *J*-cells has grown, and the bump starts spreading. Some previously active *J*-cells may not recover in time to fire on the next cycle, and clusters may develop as discussed in Section 3.4, but this does not interfere with the spread of activity, and such cells fire again on later cycles. This mechanism for the spread of activity is similar to the propagation of spindle waves in thalamic networks (Destexhe et al., 1996; Golomb et al., 1996; Rinzel et al., 1998; Terman et al., submitted). However, certain differences arise when our network's tickler connections come into play, which we discuss in subsequent subsections.

3.3. Block of Propagation

The long-range, slowly decaying tickler inputs localize the bumps. This occurs as long as more GABA_B builds up in outlying *E*-cells than in *E*-cells involved in a bump. Note that due to the wide footprint of the ticklers, outlying *E*-cells receive GABA_B input before any GABA_A input arrives. Although individual tickler inputs are weak, once enough *J*-cells are firing, certain outlying *E*-cells may receive sufficient GABA_B such that when GABA_A inhibition does reach them and then wears off, they cannot rebound. As long as the frequency of oscillations of active cells in the bump is faster than the time scale on which GABA_B decays, the GABA_B level of these outlying *E*-cells is maintained at a sufficiently high level to suppress their activity. This suppression acts as a firewall to prevent the spread of activity beyond these *E*-cells in the chain. We give a

mathematical analysis of this block of propagation and the size of the resulting bumps in Section 4.1.

3.4. Activity Within Bumps

As seen in Figs. 4, 5, and 6, cells within bumps can be loosely grouped into clusters that alternate firing. Even within clusters, however, activity is somewhat disorganized, and cluster membership varies significantly over time in an architecture-dependent manner; indeed, as we show below, bump existence requires a certain degree of asynchrony in cell activity, as also seen by Laing and Chow (forthcoming). A GABA_A footprint width of three generally limits the network to the coexistence of two alternating clusters at any one time.

3.4.1. Cluster Formation. The formation of clusters within bumps of activity, through the fast decay of GABA_A, is analogous to cluster formation among RE and TC cells in the spindle sleep rhythm (Steriade et al., 1993; Wang et al., 1995; Terman et al., 1996; Rubin and Terman, in press). Suppose that one group of *J*-cells fire in the first cycle of oscillations after activity is initiated in our model and a larger group of *E*-cells then fires by PIR, exciting a corresponding group of *J*-cells. Once the *J*-cells fall down from the active phase, the input to those *E*-cells that they have inhibited quickly wears off. Some of the inhibited *E*-cells had been at rest and fire as a result (assuming GABA_B input does not stop them). The *E*-cells that have already fired once, however, will not have fully recovered yet, assuming the *J* active phase is not too long. Hence, they will not fire on this cycle but will continue towards a rest state in the silent phase until firing at the next PIR opportunity.

3.4.2. Role of Ticklers. To a large extent, the frequency of oscillations within the bumps is determined by the duration of the *J*-cell active phase, since some *E*-cells fire as soon as each set of active *J*-cells returns to the silent phase. Since GABA_A turns off as soon as *J*-cells become inactive, it does not contribute to setting oscillatory frequency. Slow decay of tickler inputs, however, may delay the firing of *E*-cells and slightly prolong the period of some cycles.

Tickler connections are also crucial for maintaining asynchrony even within clusters of active cells and for controlling the extent to which cluster membership varies over time. In the uniform architecture, after activity is initiated within the *J*-population, all active *E*-cells receive similar numbers of tickler inputs, such that

fairly regular clusters develop within a localized bump of activity through the mechanism described above.

In the balanced and unbalanced random architectures, the randomness in tickler connections leads to networks with variability in the numbers of tickler inputs that *E*-cells receive from active *J*-cells. In both cases, this produces two groups of cells within localized bumps of activity, with different firing frequencies in each group. We refer here to “groups” rather than “clusters” because in these architectures, cells with similar frequencies do not necessarily fire together. *E*-cells in the lower frequency group received more of their tickler inputs from active *J*-cells than did *E*-cells in the higher frequency group, giving them a disadvantage in rebounding when released from GABA_A.

In our simulations with the balanced random architecture, *E*-cells in one group fired at about 6 Hz while *E*-cells in the other group fired at about 3 Hz. Neighboring *E*-cells could fire together repeatedly before eventually diverging. The firing patterns of cells in the balanced random architecture were generally more disorganized than those in the unbalanced random architecture. A simple combinatorial calculation shows that the balanced random architecture leads to more variability in the number of tickler connections from active *J*-cells to each *E*-cell, which may cause this extra disorganization. In our simulations with the unbalanced random architecture, *E*-cells in the two groups fired at about 4 and 10 Hz, respectively. Moreover, group membership tended to alternate between neighboring *E*-cells, due to GABA_A-induced PIR.

3.4.3. Effects of Inhibition Between *J*-Cells. Since thalamic reticular cells inhibit each other, we included inhibition between *J*-cells in simulations through two different experiments. We performed these experiments with the balanced random architecture, since it was observed to be the least conducive to cluster formation. First, after a bump was established, we introduced nearest-neighbor *J*-to-*J* inhibition by switching g_A^J from 0 to 0.5—namely, one half of the strength of the GABA_A inhibitory conductance from the *J*-cells to the *E*-cells. Second, we set $g_A^J = 0.5$ from the start of simulation. In both cases, the cells formed clusters which took turns firing, with neighboring cells of the same type (*E* or *J*) generally ending up belonging to different clusters; however, the *E*-clusters were less regular than the *J*-clusters.

It is clear why inhibition between *J*-cells promotes clustering (Rubin and Terman, 2000). Suppose several neighboring *E*-cells fire. If their firing times vary

slightly, then some J -cells receive excitation earlier than others and fire, inhibiting their neighbors in the J -chain. When the firing J -cells fall down to the silent phase, they release these neighboring J -cells, along with the E -cells they inhibited. These previously inhibited J -cells are ready to fire in response to excitation, whereas the previously active J -cells still need to recover, such that an alternation of firing between neighboring J -cells results.

Some alternation between neighboring E -cells is encouraged by the alternation between J -cells together with the GABA_A footprint width of three. That is, the E -cells corresponding to active J -cells only receive GABA_A from one J -cell each, whereas the neighboring E -cells, corresponding to silent J -cells, actually each receive GABA_A from two J -cells. This allows the neighboring E -cells to recover slightly faster than those E -cells corresponding to active J -cells, such that they have a slight advantage in rebounding when the J -cells' activity finishes.

3.5. Control of Bumps Through External Inputs

3.5.1. Initiation. We performed simulations of the balanced random architecture with applied excitatory inputs. Starting with the network in the silent rest state, application of cortical excitatory input to five central E -cells initiated a sustained, localized bump of activity. This bump spread to 12 cells in each population and then stabilized. As in the other simulations, this spread occurs through PIR induced by the inhibitory connections from the J -cells to the E -cells.

Application of excitatory input to central J -cells also leads to bump initiation, as shown in Fig. 7. However, when the initial excitation was applied to a larger group of J -cells, certain central cells in the resulting bump showed less activity in early activity cycles than did other cells in the bump. This occurs because outlying E -cells adjacent to the excited region are the first to fire by PIR after the initial excitation, and after they rebound, the activity must propagate back inward to the center of the initially excited group.

3.5.2. Termination. Once a bump was established, it could be terminated by application of a brief cortical excitatory input to all E - or J -cells, as shown in Fig. 8. This input caused all the cells in the excited population to fire together. When all E -cells were excited, all J -cells fired in response. In both cases, network activity ceased after the J -population fired.

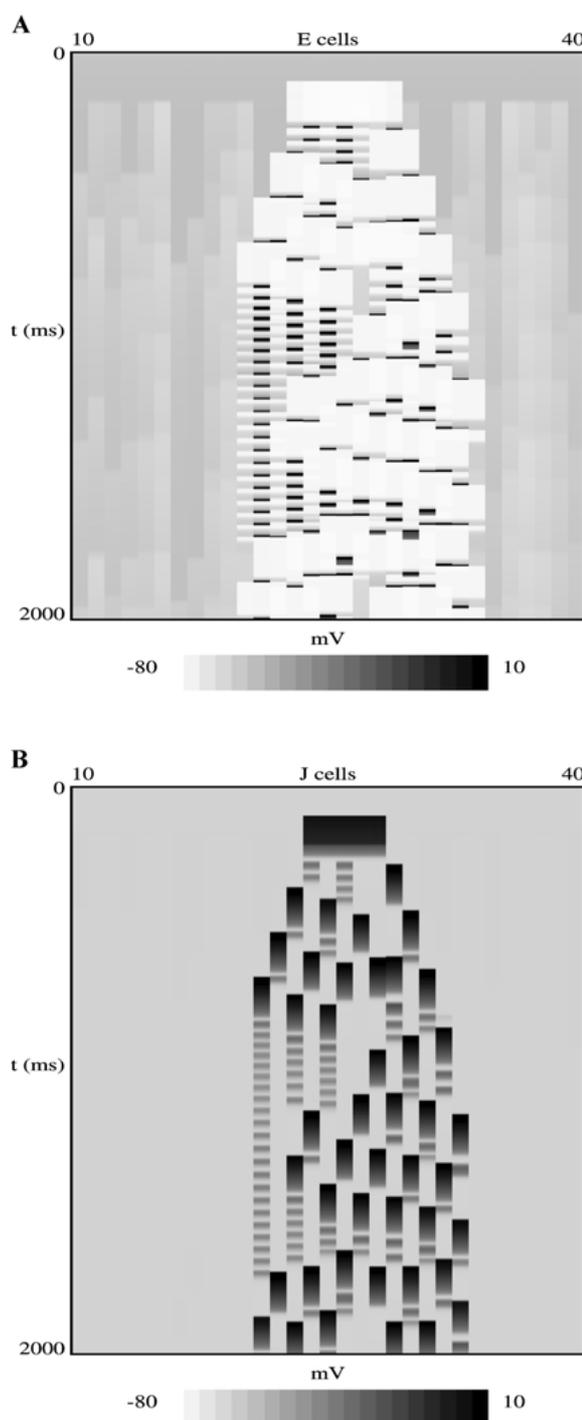


Figure 7. Initiation of activity via applied excitation. A network with the balanced random architecture was simulated with initial conditions at rest. After 100 ms, g_{ctv} was set to 1.0 for cells J_{24} , J_{25} , J_{26} , J_{27} , J_{28} to simulate cortical excitation to those cells. The cortical input was removed by returning g_{ctv} to zero after 100 ms, and an activity bump developed.

If a sufficiently large subset of J -cells was excited together (either directly through I_{ctx} or indirectly through I_{ctx} to the corresponding E -cells), then under appropriate parameter choices, activity still terminated. Cortical input to only a small subset of the active cells in the appropriate population did not prove sufficient to terminate activity, however.

3.5.3. Translation. If a network represents thalamic head-direction cells, then bumps of activity within that network should move in concert with the head. That is, there should exist a mechanism to induce translation of established bumps.

We attempted to induce translation of established bumps in our various network architectures, with results depending on the architecture simulated. In the uniform architecture, we set $g_{ctx} = 1$ ($g_{ctx'} = 1$) for 100 msec for the E -cell (J -cell) adjacent to one side of the bump. This caused the absorption of that cell and the corresponding cell in the other layer into the bump along with the removal from the bump of the pair of cells at the opposite edge of the bump (see Fig. 9). By experimenting with the nature of the current applied, we found that application of excitatory input to J -cells produced more robust bump translation than input to E -cells. This is consistent with the IPSP dominance in TC cells in thalamocortical networks (Destexhe et al., 1998, 1999), which implies that cortical excitation to RE cells has a stronger effect than that to TC cells, such that cortical inputs to RE cells would likely be used for bump control. Induction of bump translation in random architectures was less successful. This coincides with past observations that heterogeneous input interferes with translation invariance of bumps (Tsodyks and Sejnowski, 1995).

We also achieved bump translation via less narrowly focused inputs, with input timing set independently of E - and J -cell activity. After a bump was established, we set $g_{ctx'} = 1$ for sets of 5 J -cells for 10 msec intervals out of each 50 msec, with $g_{ctx'} = 0$ for all J -cells during the other 40 msec. As we slid the target set of J -cells receiving these inputs through the J -population, the center of the bump of activity of E - and J -cells slid to follow the input location, while the bump itself remained wider than 5 cells. This translation worked at both sliding speeds—namely, 1 E -cell per 100 msec and 1 E -cell per 150 msec—that we tried (data not shown). This issue will be explored further in future work.

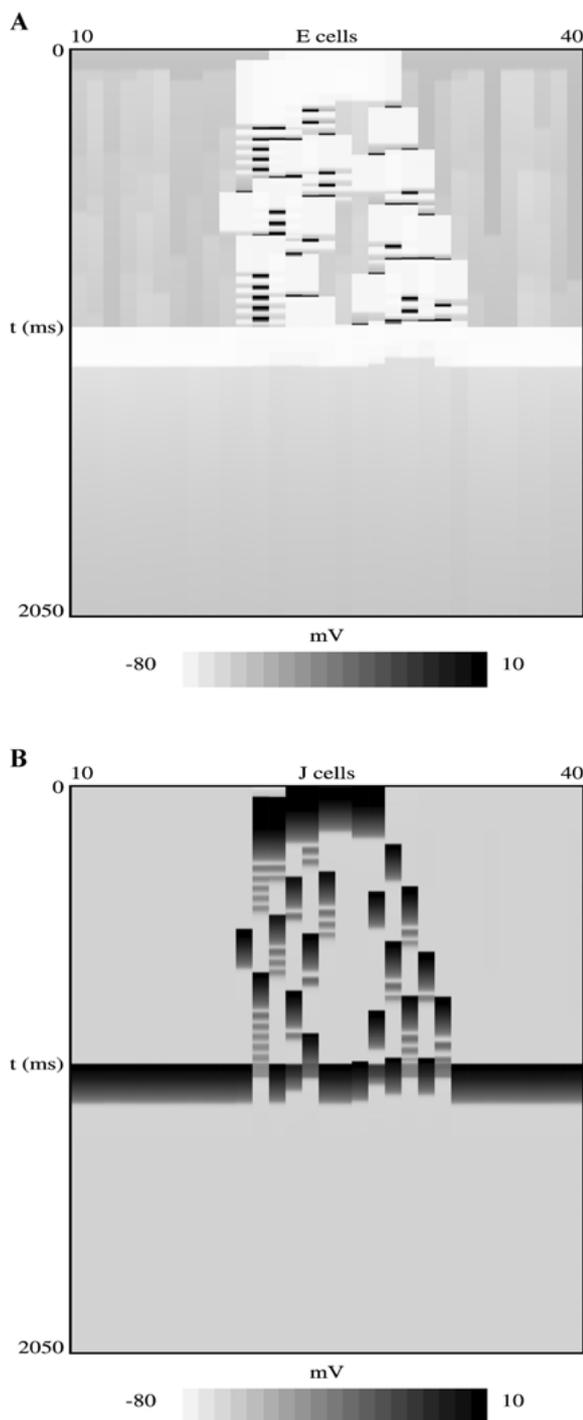


Figure 8. Termination of activity via applied excitation. In simulation of the balanced random architecture with a bump of activity, $g_{ctx'}$ was set to 0.1 for 50 ms to simulate cortical excitation to J -cells. E -cells were hyperpolarized by resulting J -cell firing, and then the network became silent. Plots show cells 10 to 40 in the center of each chain.

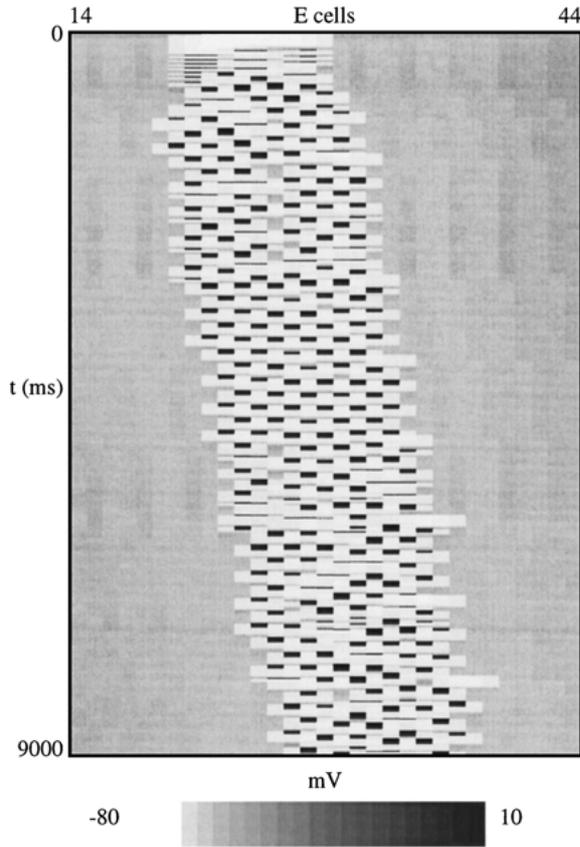


Figure 9. Bump translation by applied excitation. In this simulation, a bump of width 11 (cells 20 to 30) was established in a network with the uniform architecture with no cortical input. From time 1500 to 1600, g_{ctx} was set to 1 for J -cell J_{31} . This caused J_{31} and E_{31} to join the bump and eliminated J_{20} , E_{20} from the bump. Similarly, application of $g_{ctx} = 1$ to J_{32} from time 3000 to 3100 and to J_{33} from time 4000 to 4100 allowed J_{32} , E_{32} to replace J_{21} , E_{21} in the bump and then J_{33} , E_{33} to replace J_{22} , E_{22} in the bump, and so on (however, at some later stages, excitation was applied to more than one J -cell to translate the bump effectively). Note that at each stage, excitation did not spread beyond the new cell added.

4. Analysis and Mechanisms

4.1. Block of Propagation and Bump Size

Here we derive an analytic expression for the size of the bump. This will not be a perfect estimate quantitatively, but it performs reasonably well. Most important, the associated analysis elucidates the mechanism for propagation block, including the roles of biophysical parameters in Eqs. (1) and (2) and associated synaptic currents, in this mechanism. We can attain this insight because we use the conductance-based formulation of

Eqs. (1) and (2) and incorporate synaptic dynamics (see also Appendix B). Recall from Section 3.3 that propagation of activity is blocked if outlying E-cells receive sufficient GABA_B inhibition to prevent cells from rebounding when GABA_A inhibition wears off. Hence, to estimate the bump size, we need to compute two things: (1) the amount of GABA_B inhibition each cell receives if the bump is of a given size and (2) the amount of GABA_B inhibition that is required to prevent a cell from rebounding.

In what follows, we consider the continuum limit of Eqs. (1) and (2) as the number of cells in each population becomes unbounded. Thus the variables v , y , w , z depend on space and time. In this limit, the GABA_A input to the E -cell at position x and time t is given by $I_A(x, t) = g_A(v(x, t) - v_{inh}) \int_{-\infty}^{\infty} W_A(x - y) s_A(y, t) dy$ for a weight function $W_A(x)$, and inputs $I_B(x, t)$, $I_A'(x, t)$, $I_E(x, t)$ have similar forms. We assume that the bump is of size L , to be determined, and cells within the bump correspond to $0 < x < L$.

Assume that a bump solution exists and that GABA_B inhibition wears off slowly. This implies that to a first approximation, the GABA_B conductance is time-independent—call it $S(x, L)$. Hence, we can set $I_B(x, t) = S(x, L)(v(x, t) - v_{inh})$ for

$$S(x, L) = g_B \int_0^L W_B(x - y) dy.$$

Let $\sigma_L = S(0, L) = S(L, L)$. One can evaluate the integral explicitly if W_B is simple enough. This is done later for the case of a square, off-center architecture.

We next determine an expression for how much GABA_B inhibition is required to prevent a cell from rebounding. We first recall the mechanism for PIR (e.g., Rubin and Terman, forthcoming). While a silent cell receives GABA_A inhibition, its (v, h) values lie on the left branch of the v -nullcline corresponding to the total inhibition level it receives. At some time, suppose that its GABA_A inhibition is removed. Then, under the dynamics of (1), its (v, h) values quickly jump toward the v -nullcline corresponding to setting $I_A = 0$ and keeping I_B constant. If the trajectory lies above the left knee of this new nullcline, then the cell will jump up; if it lies below the left knee, then the cell will be unable to rebound (see Fig. 10). This observation separates cells into those that fire in response to the removal of GABA_A inhibition and those that do not, based on the level of GABA_B input they have received. It is this dichotomy that we use to estimate bump size below.

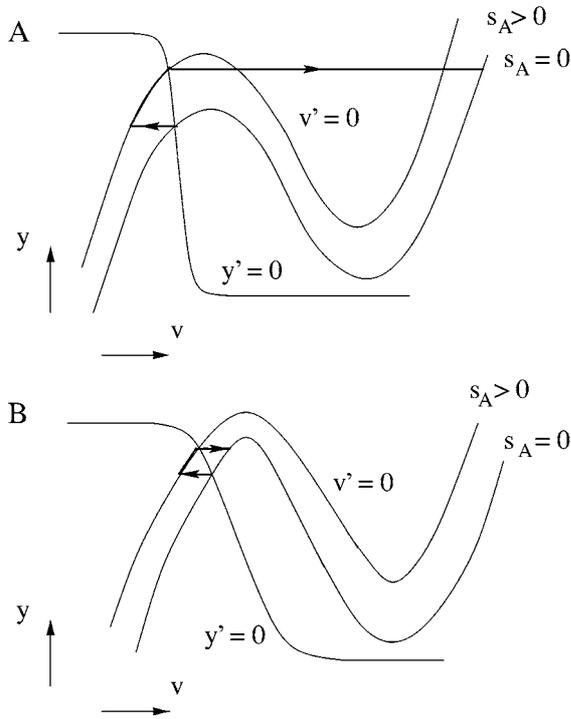


Figure 10. Diagram illustrating firing by post-inhibitory rebound (PIR) and block of rebound. In **A**, when inhibition is applied (s_A raised from 0) to an E -cell at rest, the cell approaches a new fixed point. Then, when inhibition is removed, the dynamics of the v -equation in (1) carry it to the right branch of the v -nullcline corresponding to $s_A = 0$: it rebounds. In **B**, when inhibition is applied, the E -cell approaches a new fixed point. But, since this fixed point lies below the left knee of the $s_A = 0$ v -nullcline, the cell cannot rebound when inhibition is removed.

We introduce some notation. Let $Y_{FP}(S)$ and $Y_{LK}(S)$ denote the y -coordinates of the fixed point and left knee, respectively, along the left branch of the cubic with total inhibitory input S . We denote by σ_A the level of GABA_A inhibition that each cell receives during a cycle in which it is silent. Note that σ_A should, in general, depend on both space and time; however, we will ignore this dependence. This is discussed in more detail later, when we derive an explicit expression for σ_A .

It follows that a bump exists with size L if

$$Y_{FP}(\sigma_A + \sigma_L) = Y_{LK}(\sigma_L) \quad (3)$$

as illustrated in Fig. 11, and

$$Y_{FP}(\sigma_A + S(x, L)) > Y_{LK}(S(x, L)) \quad \text{for } 0 < x < L. \quad (4)$$

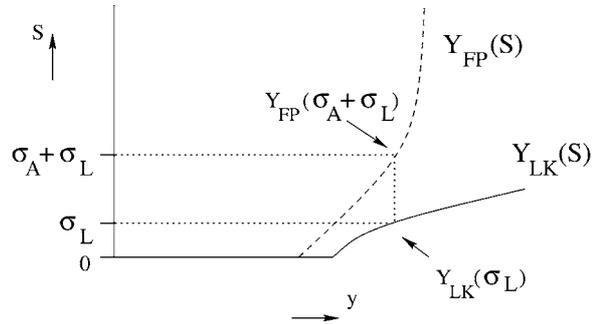


Figure 11. Diagram illustrating the amount of GABA_B inhibition at the edge of a bump. For fixed σ_A , this amount, σ_L , is exactly the level of inhibition such that $Y_{FP}(\sigma_A + \sigma_L) = Y_{LK}(\sigma_L)$.

Note that (4) implies that cells within the bump are able to rebound, while (3) implies that the boundary of the set of such cells is at $x = 0$ and $x = L$.

Consider (3) only for the moment. We find sufficient conditions so that if σ_A is fixed, then there must exist a solution of (2) for some σ_L . Note that $Y_{FP}(S) < 1$ for all S ; moreover, we show in Appendix B that $Y_{LK} > 1$ for S sufficiently large. It then follows that $Y_{FP}(\sigma_A + \sigma_L) < Y_{LK}(\sigma_L)$ if σ_L is sufficiently large. We assumed in Section 2 that if E -cells receive sufficient GABA_A inhibition, then in the absence of GABA_B, they can rebound once their GABA_A input is removed. Assume that σ_A represents sufficient GABA_A inhibition for this rebound. Then, $Y_{FP}(\sigma_A) > Y_{LK}(0)$. It now follows that there must exist a solution of (3) for some σ_L .

To obtain an explicit expression for the solution σ_L , we need to have expressions for the curves $Y_{LK}(S)$ and $Y_{FP}(S)$. Suppose that $Y_{LK}(S)$ is linear—say,

$$Y_{LK}(S) = y_{LK} + \alpha S. \quad (5)$$

This is usually a good approximation, and it is not difficult to estimate y_{LK} and α for a given model (see Appendix B). Recall that $Y_{FP}(S) < 1$ for all S . Hence, $Y_{LK}(\sigma_L)$ is strictly less than one, which gives the a priori bound

$$\sigma_L < \frac{1 - y_{LK}}{\alpha}.$$

We next assume that $Y_{FP}(S)$ is linear as long as $Y_{FP}(S)$ is not too close to 1—say,

$$Y_{FP}(S) = y_{FP} + \beta S. \quad (6)$$

Our rebound assumption implies that $y_{FP} < y_{LK}$, and we assume now that $\beta < \alpha$. This latter assumption

is reasonable since $Y'_{FP}(S) \rightarrow 0$ and $Y'_{LK}(S)$ is quite large for large S , as discussed in Appendix B. Then (3) together with (5) and (6) imply that

$$\sigma_L = \min \left\{ \frac{y_{FP} + \beta\sigma_A - y_{LK}}{\alpha - \beta}, \frac{1 - y_{LK}}{\alpha} \right\} \quad (7)$$

as illustrated in Fig. 11. We note that $y_{FP} + \beta\sigma_A > y_{LK}$, since $Y_{FP}(\sigma_A) > Y_{LK}(0)$.

We now consider σ_A . Recall that this is the level of GABA_A inhibition that the cells receive while they are silent. We suppose that the architecture is such that GABA_A has a square wave footprint—say, $W_A(x) = 1$ if $|x| < \lambda_A$ and $W_A(x) = 0$, otherwise. We also assume that at any given time, a certain fraction of the cells are active. We denote this fraction by ρ . For example, if the network breaks up into two clusters that take turns firing, then $\rho = 1/2$. It follows that the GABA_A inhibition that the cells at $x = 0$ and $x = L$ receive is $\sigma_A = \rho g_A \lambda_A / 2$.

Finally, we need an explicit expression for the amount of GABA_B input a cell receives from a bump of size L . We then substitute this expression for σ_L into (7) to obtain a formula for the bump size. We suppose the network has a square, off-center GABA_B architecture. That is, there exist positive γ and ω such that $W_B(x) = 1$ if $\gamma < |x| < \gamma + \omega \equiv \lambda_B$ and $W_B(x) = 0$, otherwise. We will assume that $\lambda_B > \lambda_A$. Then a straightforward calculation demonstrates that

- (a) If $L > \lambda_B$, then $S(x, L) = \omega g_B$ for $0 < x < L - \lambda_B$ and $S(x, L) < \omega g_B$ for $L - \lambda_B < x < L$.
- (b) If $\gamma < L < \lambda_B$, then $\sigma_L = g_B(L - \gamma)$ and $S(x, L) < \sigma_L$ for $0 < x < L$.
- (c) If $L < \gamma$, then $S(x, L) = 0$ for $0 < x < L$.

We note that the level of GABA_B inhibition within the bump must be largest at the bump's boundaries—that is, we must have $S(x, L) < \sigma_L$ for $0 < x < L$. Therefore, the size of the bump must satisfy (b) above; in this case $\sigma_L = g_B(L - \gamma)$. If we substitute this into (7), then we obtain the desired formula for the bump size—namely,

$$L = \gamma + \frac{1}{g_B} \min \left\{ \frac{y_{FP} + \beta \rho g_A \lambda_A / 2 - y_{LK}}{\alpha - \beta}, \frac{1 - y_{LK}}{\alpha} \right\}.$$

Remarks:

1. Since $L < \lambda_B$, this leads to a lower bound on the maximal GABA_B conductance g_B .

2. Similar analysis can be done for nonrandom discrete tickler architectures, such as the uniform architecture.
3. A numerical example of these curves, together with trajectories of a cell that belongs to a bump and a cell outside of the bump, appears in Fig. 12. Note that the assumptions of linearity on $Y_{LK}(S)$ and $Y_{FP}(S)$ are accurate locally: Y_{FP} is approximately linear at least up to about $y = 0.6$ and, after an initial nonlinearity, Y_{LK} is quite close to linear. Note too that if the cells' positions are not close to $Y_{FP}(S)$ when GABA_A turns off, as in Fig. 12, then y_{FP} should be replaced by a smaller parameter found by translating the critical point curve to the cells' turn-off positions. Numerical estimate of parameters in (7), together with this translation, yields $\sigma_L \approx .21$, which only slightly overestimates the value of s for the silent cell after GABA_A turns off in Fig. 12. Division by the small parameter g_B in the estimate for L magnifies errors in σ_L . In typical simulations, the formula nonetheless overestimates bump size only by about 2 cells.

4.2. Bump Termination via External Inputs

Consider a bump of active E - and J -cells. If I_{ctx} is applied to all E -cells, causing all E -cells to fire, then all J -cells fire in response. The E -cells then return to the silent phase before the J -cells but cannot rebound when the J -cells later inactivate and the GABA_A input to the E -cells turns off. Hence, activity terminates. Unlike the synchronization instability of the bumps in (Laing and Chow, forthcoming), this synchronization-induced silence is not solely controlled by synaptic timescales. Indeed, it can work through two different mechanisms, one depending on intrinsic E - and J -cell properties instead of synaptic timescales.

This intrinsic mechanism for synchronization-induced silence occurs for a choice of parameters such that the J active phase is short. If the J active phase lasts a sufficiently short time, then E -cells cannot recover sufficiently for PIR. A second, synaptic mechanism for this activity termination arises through tickler inputs when the J active phase is not so short. When all J -cells fire and stay active for a sufficiently long time, all E -cells receive tickler inputs. The resulting GABA_B inhibition to the E -cells persists beyond the fast turn-off of GABA_A. It is known that a slow decay of inhibition in a synchronized E -cell population can lead to absence of subsequent activity (Rubin and Terman,

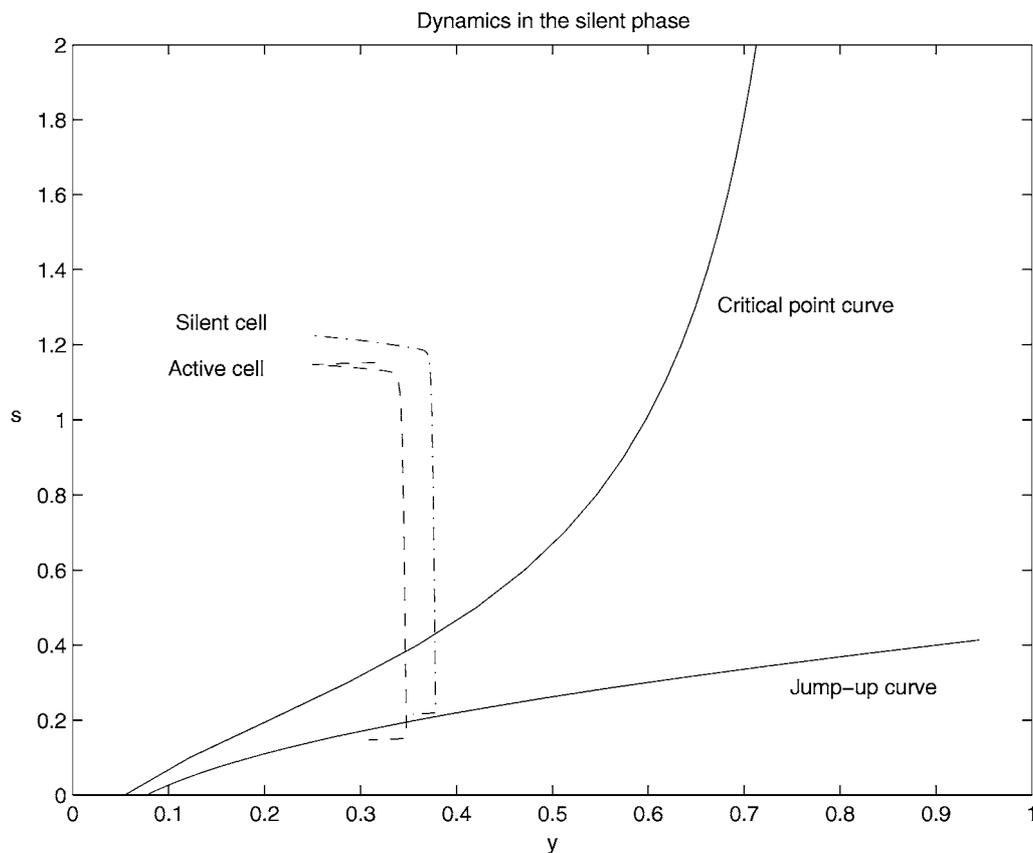


Figure 12. Dynamics of E -cells in the silent phase. The solid lines denote the critical point curve $\{y = Y_{FP}(s)\}$ and the jump-up curve $\{y = Y_{LK}(s)\}$, respectively. The other curves are trajectories of E -cells generated by (1) with fixed $GABA_A$ input but different levels of $GABA_B$ inputs. When $GABA_A$ turns off, both E -cells experience rapid drops in their total synaptic inputs s , as indicated by the dashed lines. (Note that s here denotes $g_A s_A + g_B s_B$.) The active cell (dashed curve) hits the jump-up curve at some $(y, s) \approx (0.35, 0.18)$ and then jumps to the active phase where its y value begins to decrease (also shown as the part of the dashed curve below the jump-up curve). The tickler input to the silent cell keeps its s value around 0.2 (corresponding to $s_B \approx 5$ since $g_B = 0.04$), and it stays just above the jump-up curve. Thus, the silent cell (dash-dotted curve) remains in the silent phase, where its y value decreases as it approaches the critical point curve from above.

2000) (although it can produce synchronized activity under other parameter choices; see Wang and Rinzel, 1993 and Terman et al., 1998). Similar mechanisms underlie the termination results, presented in subsection 3.6.2, when J -cells receive cortical excitation. For example, in Fig. 8, some E -cells that are hyperpolarized due to cortical excitation of the J -cells are prevented from rebounding by ticklers.

4.3. Bump Translation via External Inputs

The mechanism for bump translation via localized application of cortical input involves tickler connections in two ways. First, the activity of the externally excited cell that joins the bump must cause activity to

cease fairly quickly in an E -cell on the other side of the bump. This occurs through the added tickler input from the newly active J -cell. Of course, excitation of multiple cells may be required for bump translation when changes in activity and corresponding tickler inputs due to excitation of single E - or J -cells are insufficient to achieve this suppression. Second, the activity must not spread to previously silent E -cells adjacent to the newly active E -cell. This spread is blocked by ticklers from the various active J -cells in the translated bump. Both of these effects are displayed in Fig. 13 and can be analyzed with similar calculations to those in Section 4.1.

An analogous mechanism leads to translation in our experiments with less narrowly focused sliding inputs.

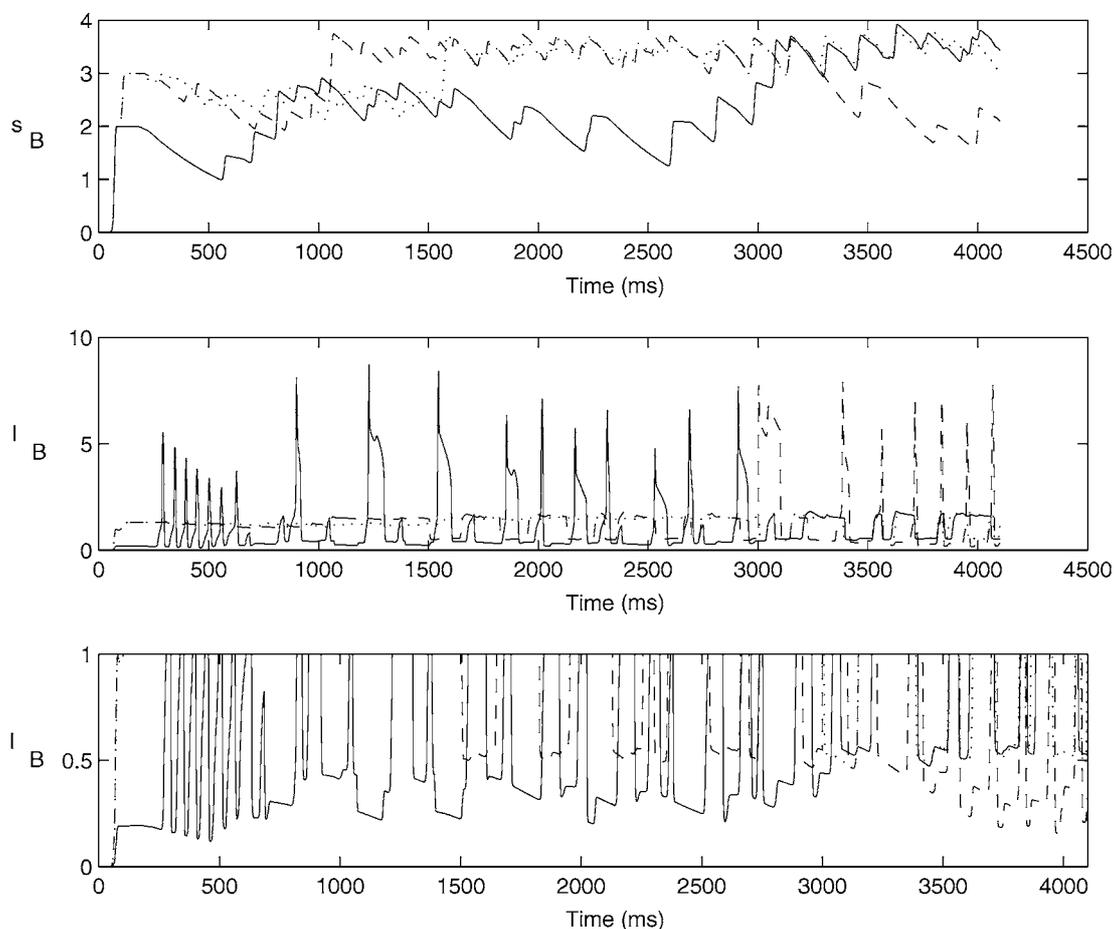


Figure 13. Tickler inputs in bump translation. In this simulation, current was applied to E -cell 32 in a chain during time 3000 to 3100, bringing it into an established bump of activity and pushing out E -cell 21. Top plot shows s_B for E_{21} (solid), E_{32} (dashed) and E_{33} (dotted), while the middle plot shows I_B for these cells. The bottom plot contains the same data as the middle plot, but we have zoomed in on the lower range of I_B values. Note that the active cell among the three has lowest s_B . I_B to the cells is more variable; however, as the bottom plot shows, only the active cell ever has its I_B drop below a certain level.

There, periodic driving to a target zone that slides through the J -population knocks out activity in one or more (but not all) previously active J -cells and easily prevents extra silent cells, adjacent to newly recruited ones, from becoming active.

Translation via highly localized inputs often failed in random architectures. This can occur because an inactive cell that is excited outside one side of the bump may not affect cells on the other side of the bump, given the irregularity of the tickler connections in the network. In both types of random architecture, the probability that a central E -cell (i.e., an E -cell not within 14 units of the edge of the E -chain) receives no ticklers from a given J -cell within ± 14 units of it is $(28/29)^5 = .8391$.

We did not systematically explore changes in activity in unsuccessful translation attempts or stimulation of individual cells away from the edges of bumps.

5. Discussion

We have considered a model network of reciprocally connected excitatory and inhibitory cells, representing a minimal RE-TC thalamic network. Our model incorporates biophysical representations of thalamic cells and the synaptic coupling between them, with a variety of connection architectures. In this setting, we show that thalamic networks can support stable localized bumps of sustained oscillatory activity. These bumps

spread and persist through recurrent, alternate volleys of excitation and localized fast inhibition, which induces post-inhibitory rebound. They are localized by long-range slow inhibition. Activity within the bumps is clustered but disorganized; synchronization eliminates activity. Computational HD cell models based on classic voltage-based (Amari, 1977) or activity-based (Wilson and Cowan, 1973) population models (Redish et al., 1996; Zhang, 1996) require recurrent excitatory connections (Ermentrout, 1998; Blair, 1996), which are not known to exist in the thalamus, to support such solutions. Our model yields bump attractors without including such connections.

Occurrence of bump attractors in past models also may require a sufficiently fast membrane time constant for inhibitory cells (Redish et al., 1996) or a sufficiently fast time constant for inputs to these cells (Ermentrout, 1998). Moreover, these models do not take into account synaptic delays or refractoriness of cells. By considering conductance-based membrane potential equations and dynamic synaptic conductances, we incorporate more temporal details into our network dynamics. In particular, our model takes into account the states of cells (e.g., active or silent) when evaluating how they are affected by the synaptic inputs they receive and allows for synaptic inputs to persist after the corresponding presynaptic cells become inactive, when appropriate. Thus, we can clarify the roles of various biophysical features (e.g., the parameters that control nullcline shapes and the durations of various phases of cellular activity) in the mechanisms for bump generation, maintenance, and spread. We also obtain more detailed information about the activity of cells belonging to bump solutions, beyond their firing rates or averaged voltages. These results offer an important example of how inhibition is useful in generating complex activity patterns in bursting cells and highlight the importance of the timescales of inhibitory synaptic dynamics in pattern formation.

Since HD cells in ATN cease discharging without volitional input (Taube, 1995), model networks should include mechanisms for bump generation and bump elimination. Since all cells in our model are intrinsically silent, we have bistability of bump solutions with a co-existent stable silent state of no activity in the network. This contrasts with previous HD models, in which activity bumps develop from random or arbitrary initial states (Redish et al., 1996; Zhang, 1996). We demonstrate bump elimination, through synchronization initiated by a widespread transient excitatory input, and

bump generation, through a transient excitatory input to a small group of cells; both can be achieved via input to either population in the network.

While nearby cells in the chains in our model are active at the same time, this does not contradict the finding that some nearby HD cells in ATN have different preferred directions (Taube, 1995). The cells in our model should be interpreted as a group of cells that share synaptic connections, representing only a subset of a larger overall cell population. Correspondingly, the cells in our chains are ordered by connectivity, not by actual spatial location. The cells involved in the activity patterns we have described can thus be scattered throughout a spatial region.

Our finding that the firing rates of cells are not constant across a bump of activity is consistent with experimental measurements of HD firing in ATN (Taube, 1995). Our model does not, however, reproduce certain other features associated with HD firing. For example, while we present a mechanism for changing membership in bumps, we do not consider how connections decrease a cell's firing rate as the head moves away from the center of its preferred firing range. Indeed, unlike bumps in previous computational studies (e.g., Amari, 1977; Ermentrout, 1998; Laing and Chow, forthcoming), our bumps do not show a smooth, pulse-like profile of activity across the cells within the bump, with activity levels decreasing near the edges; instead, each cell in our simulations of bump translation essentially fires at a fixed nonzero rate or a zero rate. Further, our model does not produce a nonzero background firing rate for cells outside of their preferred ranges. Both of these features might be attainable with the inclusion of additional randomness and noise in the model. For example, fluctuations in tickler conductances could allow certain cells to fire less frequently as the head neared the outskirts of their preferred ranges. Even without noise, weak excitation to an appropriate part of the network could possibly decrease a certain cell's firing rate, through increased tickler inputs, while stronger excitation would eliminate that cell's activity, as seen in our bump translation results. Careful investigation of these issues remains for future consideration.

We did not try to reproduce quantitative details of HD firing rates, to estimate the widths of preferred firing ranges, or to compute speeds of responses to change in inputs. Our aim was to develop and explore a minimal thalamic network model that supports bump activity. Correspondingly, details of how external (visual) cues or vestibular inputs affect the ATN head direction

system and consideration of its interactions with the PoS are beyond the scope of this work.

Appendix A: Numerics

In Eqs. (1) and (2), the voltage-dependent functions take the form $X_\infty(v) = \frac{1}{1+\exp[(v-\theta_X)/\sigma_X]}$, where X can be m , y , or z , and $\tau_X(v) = \tau_X^0 + \frac{\tau_X^1}{1+\exp[(v-\theta_X^1)/\sigma_X^1]}$ where X can be y or z .

We model the GABA_A inhibitory current to cell E_i by $I_A = g_A s_A = g_A (\sum_j s_j)(v_i - v_{inh})$ and, similarly, the GABA_B inhibitory current to E_i by $I_B = g_B s_B = g_B (\sum_k s_k)(v_i - v_{inh})$. The sums in these expressions are taken over the J -cells that are coupled to cell E_i via GABA_A and GABA_B, respectively. Each synaptic variable s_j (s_k) depends on the voltage w_j (w_k) of the j th (k th) J -cell.

Since GABA_A turns on and off relatively quite quickly in response to changes in presynaptic membrane potential, we choose the specific form $I_A = g_A (s_\infty(w_{i-1}) + s_\infty(w_i) + s_\infty(w_{i+1}))(v_i - v_{inh})$, where $s_\infty(v) = 1/(1 + \exp(-(v - \theta_A)/\sigma_A))$ for constants θ_A, σ_A , in our simulations. Here cell E_i receives inputs from cells J_{i-1}, J_i, J_{i+1} . The parameter v_{inh} is such that $(v - v_{inh}) > 0$, and hence GABA_A input affects E -cell v -nullclines as shown in Fig. 2. Alternately, our synaptic variables s_k evolve according to $s'_k = \alpha F(w_k)(1 - s_k) - \beta s_k$ for a monotone increasing nonlinear function F . Experiments have shown a delay between the firing of an RE cell and the GABA_B inhibition of coupled TC cells (Golomb et al., 1994; Huguenard and Prince, 1994; Destexhe and Sejnowski, 1995). Hence, in our simulations, we use the slow, indirect scheme of Golomb et al. (1994), developed to replicate these effects, for F in I_B .

The GABA_A current $I_{A'}$ to the i th J -cell has the form $I_{A'} = g_{A'} s_{A'} = g_{A'} (\sum_j s_j)(w_i - v_{inh})$, where the sum is

over the J -cells inhibiting cell J_i . The excitatory input from E -cells to cell J_i also turns on and off quickly and hence we model the resulting excitatory current as $I_E = g_E s_{E_i}(w_i - v_{exc})$ where $s'_{E_i} = \alpha_E k_\infty(v_i)(1 - s_{E_i}) - \beta_k s_{E_i}$ for relatively large constants α_k, β_k . We take $k_\infty(v) = s_\infty(v)$ for simplicity. The parameter v_{exc} is such that $(w - v_{exc}) < 0$, and hence E -cell input to J -cells affects their w -nullclines as shown in Fig. 2.

Finally, cortical inputs to E -cells generate a current that we model as $I_{ctx} = g_{ctx}(v - v_{ctx})$ for constants g_{ctx}, v_{ctx} . We allow the strength of g_{ctx} to vary within the E -cell population. Cortical inputs to J -cells take an analogous form.

Parameter values used in simulations of our model are given in Tables 1 and 2. These do not include units because for our qualitative results, only relative parameter sizes are relevant. Note that for simulations with the random networks, we used $g_A = 1.0, g_B = 0.04$ instead of $g_A = 0.5, g_B = 0.03$.

Appendix B: Formulas for Curves of Fixed Points and Knees

The analysis of bump size in Section 4 relies on comparison of the rest states that E -cells approach in the silent phase while inhibited to the knees of the appropriate nullclines as inhibition wears off. Thus, it is useful to examine how one can compute locations of curves of fixed points and knees, as well as the slopes of these curves.

First, we need to define mathematical representations of certain features in (v, y) phase space. Write the first equation in (1) as

$$v' = f(v, y) - g_{syn}s(v - v_{syn}) \equiv F(v, y, s) \quad (8)$$

where the g_{syn} term subsumes all synaptic currents to the E -cell. We can do this by defining $s = s_A +$

Table 1. Basic set of parameter values for the conductance-based E -cell model.

Parameter	Value	Parameter	Value	Parameter	Value	Parameter	Value
g_{Ca}	1.5	v_{Ca}	90.0	θ_m	-45.0	σ_m	-9.0
g_L	0.2	v_L	-60.0	θ_y	-72.0	σ_y	5.0
				τ_y^0	100.0	τ_y^1	500.0
				θ_y^τ	-78.0	σ_y^τ	3.0
g_A	0.5	v_{inh}	-80.0	θ_A	-40.0	σ_A	2.0
g_B	0.03			α	2.0	β	0.002
g_{ctx}	1.0	v_{ctx}	0.0	ϕ	0.75		

Table 2. Basic set of parameter values for the conductance-based J -cell model.

Parameter	Value	Parameter	Value	Parameter	Value	Parameter	Value
g_{Ca}	1.5	v_{Ca}	90.0	θ_m	-45.0	σ_m	-7.4
g_L	0.2	v_L	-65.0	θ_z	-72.0	σ_z	5.0
				τ_z^0	100.0	τ_z^1	500.0
				θ_z^τ	-78.0	σ_z^τ	3.0
$g_{A'}$	0.5	v_{inh}	-80.0	θ_A	-40.0	σ_A	2.0
g_E	0.1	v_{exc}	0.0	α_k	2.0	β_k	0.1
$g_{ctx'}$	1.0	$v_{ctx'}$	0.0	ψ	1.0		

g_{BSB}/g_A and $g_{syn} = g_A$. The left branch of the corresponding v -nullclines can be denoted by $v = \Phi(y, s)$, where $F(\Phi(y, s), y, s) = 0$. Note that the location of this curve depends on s .

We have assumed that there is a critical point of (1) on the left branch for each s . This is given by taking $v = \Phi(y, s)$, such that $v' = 0$, and solving

$$y_\infty(\Phi(y, s)) - y = 0 \quad (9)$$

such that $y' = 0$ in (1); this is easily done numerically (see Fig. 12). Label the solution curve $\{y = Y_{FP}(s)\}$.

Differentiating equation (9) with respect to s for $y = Y_{FP}(s)$ yields the equation

$$Y'_{FP}(s) = \frac{(\partial y_\infty / \partial v)(\partial \Phi / \partial s)}{1 - (\partial y_\infty / \partial v)(\partial \Phi / \partial y)}.$$

Since y_∞ is a monotone decreasing function of v , consideration of (8) and (1) shows that $Y'_{FP}(s) > 0$. Further, since $y_\infty \rightarrow 1$ as the E -cell is hyperpolarized, $\partial y_\infty / \partial v \rightarrow 0$ and $Y'_{FP}(s) \rightarrow 0$ as well.

The knees of the v -nullcline are defined by the condition $f_v - g_{syn}s = 0$, with the left knee given by solving this equation for $v = \Phi(y, s)$. Denote the resulting solution curve by $\{y = Y_{LK}(s)\}$. We can solve numerically and plot this so-called jump-up curve in (y, s) space; we can then track the evolution of E -cells in the silent phase in this space, as shown in Fig. 12. When a trajectory of a cell hits $\{y = Y_{LK}(s)\}$, the cell fires an action potential and jumps up to the active phase.

Differentiating the equation $F(\Phi(Y_{LK}(s), s), Y_{LK}(s), s) = 0$ with respect to s yields an equation for $Y'_{LK}(s)$ [see also (Rubin and Terman, 2000)]. For the specific currents in (1), this is given by

$$Y'_{LK}(s) = \frac{g_{syn}(\Phi - v_{syn})}{-g_{Ca}m_\infty^2(\Phi)(\Phi - v_{Ca})},$$

which, as the ratio of two positive quantities, is positive. Note that $m_\infty(\Phi)$ approaches zero, such that this derivative becomes quite large, as the cell becomes hyperpolarized.

The curves Y_{FP} , Y_{LK} can be obtained numerically (see Fig. 12), along with the various quantities in the formulas for their derivatives.

Acknowledgments

Research for this article was supported in part by the NSF grants DMS-9802339 (D.T.) and DMS-9804447 (J.R.), by the Sloan Foundation (C.C.), and by NIH grant K01 MH01508 (C.C.). C.C. thanks Bill Skaggs for fruitful discussions and for introducing him to the head-direction system.

References

- Amari S (1977) Dynamics of pattern formation in lateral-inhibition type neural fields. *Biol. Cybern.* 27:77-87.
- Amit DJ, Brunel N (1997) Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb. Cortex* 7:237-252.
- Blair HT (1996) A thalamocortical circuit for computing directional heading in the rat. In: Touretzky D, Moser M, Hasselmo M, eds. *Advances in Neural Information Processing Systems* (Vol. 8). MIT Press, Cambridge, MA.
- Bose A, Kopell N, Terman D (2000) Almost-synchronous solutions for mutually coupled excitatory neurons. *Physica D* 140:69-94.
- Cox CL, Huguenard JR, Prince DA (1996) Heterogeneous axonal arborizations of rat thalamic reticular neurons in the ventrobasal nucleus. *J. Comp. Neurol.* 366:416-430.
- Cox CL, Huguenard JR, Prince DA (1997) Nucleus reticularis neurons mediate diverse inhibitory effects in thalamus. *Proc. Natl. Acad. Sci. USA* 94:8854-8859.
- Destexhe A, Bal T, McCormick D, Sejnowski T (1996) Ionic mechanisms underlying synchronized oscillations and propagating waves in a model of ferret thalamic slices. *J. Neurophysiol.* 76:2049-2070.

- Destexhe A, Contreras D, Steriade M (1998) Mechanisms underlying the synchronizing action of corticothalamic feedback through inhibition of thalamic relay cells. *J. Neurophysiol.* 79:999–1016.
- Destexhe A, Contreras D, Steriade M (1999) Cortically induced coherence of a thalamic-generated oscillation. *Neuroscience* 92:427–443.
- Destexhe A, McCormick DA, Sejnowski TJ (1993) A model for 8–10 Hz spindling in interconnected thalamic relay and reticularis neurons. *Biophys. J.* 65:2474–2478.
- Destexhe A, Sejnowski T (1996) Synchronized oscillations in thalamic networks: insights from modeling studies. In Steriade M, Jones E, McCormick D, eds. *Thalamus*. Elsevier, Amsterdam. Vol. 2, pp. 331–372.
- Destexhe A, Sejnowski TJ (1995) G protein activation kinetics and spillover of γ -aminobutyric acid may account for differences between inhibitory responses in the hippocampus and thalamus. *Proc. Natl. Acad. Sci. USA* 92:9515–9519.
- Ermentrout B (1998) Neural nets as spatio-temporal pattern forming systems. *Rep. Prog. Phys.* 61:353–430.
- Ermentrout B (2000) XPPAUT. Available at <http://www.pitt.edu/~phase/>.
- Golomb D, Wang XJ, Rinzal J (1994) Synchronization properties of spindle oscillations in a thalamic reticular nucleus model. *J. Neurophysiol.* 72:1109–1126.
- Golomb D, Wang XJ, Rinzal J (1996) Propagation of spindle waves in a thalamic slice model. *J. Neurophysiol.* 75:750–769.
- Goodridge JP, Taube JS (1997) Interaction between the postsubiculum and anterior thalamus in the generation of head direction cell activity. *J. Neurosci.* 17:9315–9330.
- Hansel D, Sompolinsky H (1998) Modeling feature selectivity in local cortical circuits. In: Koch C, Segev I, eds. *Methods in Neuronal Modeling: From Ions to Networks*. 2nd ed. MIT Press, Cambridge, MA. pp. 499–567.
- Huguenard J, Prince D (1994) Clonazepam suppresses GABA_B-mediated inhibition in thalamic relay neurons through effects in nucleus reticularis. *J. Neurophysiol.* 71:2576–2581.
- Kultas-Ilinsky K, Yi H, Ilinsky IA (1995) Nucleus reticularis thalami input to the anterior thalamic nuclei in the monkey: A light and electron microscopic study. *Neurosci. Letters* 186:25–28.
- Laing CR, Chow CC (Forthcoming) Stationary bumps in networks of spiking neurons. *Neural. Comp.*
- Rao SG, Williams G, Goldman-Rakic PS (1999) Isodirectional tuning of adjacent interneurons and pyramidal cells during working memory: Evidence for microcolumnar organization in PFC. *J. Neurophysiol.* 81:1903–1916.
- Redish AD, Elga AN, Touretzky DS (1996) A coupled attractor model of the rodent head direction system. *Network* 7:671–685.
- Rinzal J, Terman D, Wang XJ, Ermentrout B (1998) Propagating activity patterns in large-scale inhibitory neuronal networks. *Science* 279:1351–1355.
- Rubin J, Terman D (In press) Analysis of clustered firing patterns in synaptically coupled networks of oscillators. *J. Math. Biol.*
- Rubin J, Terman D (Forthcoming) Geometric singular perturbation analysis of neuronal dynamics. In: Fiedler B, ed. *Handbook of Dynamical Systems. Vol. 3: Towards Applications*. Elsevier, Amsterdam.
- Rubin J, Terman D (2000) Geometric analysis of population rhythms in synaptically coupled neuronal networks. *Neural Comput.* 12:597–645.
- Shibata H (1992) Topographic organization of subcortical projections to the anterior thalamic nuclei in the rat. *J. Comp. Neurol.* 323:117–127.
- Skaggs WE, Knierim JJ, Kudrimoti HS, McNaughton BL (1995) A model of the neural basis of the rat's sense of direction. In: Tesauro G, Touretzky D, Leen T, eds. *Advances in neural information processing systems*. MIT Press, Cambridge, MA. Vol. 7, pp. 173–180.
- Sohal V, Huguenard J (1998) Long-range connections synchronize rather than spread intrathalamic oscillatory activity: computational modeling and *in vitro* electrophysiology. *J. Neurophysiol.* 80:1736–1751.
- Somers DC, Nelson SB, Sur M (1995) An emergent model of orientation selectivity in cat visual cortical simple cells. *J. Neurosci.* 15:5448–5465.
- Steriade M, Jones E, McCormick D, eds. *Thalamus*. Vol. 2. Elsevier, Amsterdam.
- Steriade M, McCormick DA, Sejnowski TJ (1993) Thalamocortical oscillations in the sleep and aroused brain. *Science* 262:679–685.
- Taube JS (1995) Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. *J. Neurosci.* 15:70–86.
- Taube JS, Muller RU, Ranck JB Jr (1990) Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *J. Neurosci.* 10:420–435.
- Terman D, Bose A, Kopell N (1996) Functional reorganization in thalamocortical networks: Transition between spindling and delta sleep rhythms. *Proc. Natl. Acad. Sci. USA* 93:15417–15422.
- Terman D, Ermentrout G, Yew A (Submitted) Propagating activity patterns in thalamic neuronal networks.
- Terman D, Kopell N, Bose A (1998) Dynamics of two mutually coupled inhibitory neurons. *Physica D* 117:241–275.
- Tsodyks M, Sejnowski T (1995) Associative memory and hippocampal place cells. *Int. J. Neur. Syst.* 6(supp.):81–86.
- Van Groen T, Wyss JM (1990) The postsubicular cortex in the rat: characterization of the fourth region of the subicular cortex and its connections. *Brain Res.* 529:165–177.
- Van Groen T, Wyss JM (1995) Projections from the anterodorsal and anteroventral nucleus of the thalamus to the limbic cortex in the rat. *J. Comp Neurol.* 358:584–604.
- Wang XJ, Golomb D, Rinzal J (1995) Emergent spindle oscillations and intermittent burst firing in a thalamic model: specific neuronal mechanisms. *Proc. Natl. Acad. Sci. USA* 92:5577–5581.
- Wang XJ, Rinzal J (1993) Spindle rhythmicity in the reticularis thalamic nucleus: synchronization among mutually inhibitory neurons. *Neuroscience* 53:899–904.
- Wilson HR, Cowan JD (1973) A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Kybernetik* 13:55–80.
- Zhang K (1996) Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensembles: A theory. *J. Neurosci.* 16:2112–2126.