LETTER Communicated by Misha Tsodyks

Stationary Bumps in Networks of Spiking Neurons

Carlo R. Laing
Carson C. Chow
Department of Mathematics, University of Pittsburgh, Pittsburgh PA 15260, U.S.A.

We examine the existence and stability of spatially localized "bumps" of neuronal activity in a network of spiking neurons. Bumps have been proposed in mechanisms of visual orientation tuning, the rat head direction system, and working memory. We show that a bump solution can exist in a spiking network provided the neurons fire asynchronously within the bump. We consider a parameter regime where the bump solution is bistable with an all-off state and can be initiated with a transient excitatory stimulus. We show that the activity profile matches that of a corresponding population rate model. The bump in a spiking network can lose stability through partial synchronization to either a traveling wave or the all-off state. This can occur if the synaptic timescale is too fast through a dynamical effect or if a transient excitatory pulse is applied to the network. A bump can thus be activated and deactivated with excitatory inputs that may have physiological relevance.

1 Introduction

Neuronal activity due to recurrent excitations in the form of a spatially localized pulse or bump has been proposed as a mechanism for feature selectivity in models of the visual system (Somers, Nelson, & Sur, 1995; Hansel & Sompolinsky, 1998), the head direction system (Skaggs, Knieram, Kudrimoti, & McNaughton, 1995; Zhang, 1996; Redish, Elga, & Touretzky, 1996), and working memory (Wilson & Cowan, 1973; Amit & Brunel, 1997; Camperi & Wang, 1998). Many of the previous mathematical formulations of such structures have employed population rate models (Wilson & Cowan, 1972, 1973; Amari, 1977; Kishimoto & Amari, 1979; Hansel & Sompolinsky, 1998). (See Ermentrout, 1998, for a recent review.)

Here, we consider a network of spiking neurons that shows such structures and investigate their properties. In our network we find localized time-stationary states (bumps), which may be analogous to the structures measured in the experiments of Colby, Duhamel, and Goldberg (1995) and Funahashi, Bruce, and Goldman-Rakic (1989). The network is bistable, with the bump and the "all–off" state both being stable. Note that the neurons are not intrinsically bistable, as in Camperi and Wang (1998), and the bump solutions do not arise from a Turing–Hopf instability, like that studied by

Bressloff and Coombes (1998) and Bressloff, Bressloff, and Cowan (1999); there is no continuous path in parameter space connecting a bump and the all-off state. A time-stationary solution is one that corresponds to asynchronous firing of neurons where the firing rate is constant at each spatial point but the rate depends on spatial location. We show that the activity profile of the bumps of our model is the same as that of a corresponding population rate model.

However, bumps predicted by the rate model to be stable may in fact be unstable in a model that includes the spiking dynamics of the neurons. The rate model implicitly assumes asynchronous firing and considers only the dynamics of the firing rate. As the synaptic decay time is increased in the spiking network, the bump can lose stability as a result of temporal correlation or "partial synchronization" of neurons involved in the bump. If the initial conditions are symmetric, then this synchronization causes the input to the neurons to drop below the threshold required to keep it firing, leading to cessation of oscillation of the neurons and consequently the rest of the bump. However, for generic initial conditions or with the inclusion of noise, the bump destabilizes to a traveling wave. For fast enough synapses, the wave cannot exist. If some heterogeneity in the intrinsic properties of the neuron is included, then the bump can be "pinned" to a fixed location; the traveling wave does not form, and the bump loses stability to the all-off state.

This instability provides a mechanism for the termination of a bump, as would be required at the end of a memory task (e.g., the delayed saccade task discussed by Colby et al., 1995): if many of the neurons involved in the bump can be caused to fire approximately simultaneously and the synaptic timescale is short, there will not be enough input after this coincident firing to sustain activity, and the network will switch to the all-off state.

2 Neuron Model

We consider a network of N integrate-and-fire neurons whose voltages, v_i, obey the differential equations

$$\frac{dv_i}{dt} = I_i - v_i + \sum_{j,m} \frac{J_{ij}}{N} \alpha(t - t_j^m) - \sum_l \delta(t - t_i^l), \tag{2.1}$$

where the subscript i indexes the neurons, t_j^m is the mth firing of neuron j, defined by the times that $v_j(t)$ crosses the threshold, which we have set to 1, I_i is the input current applied to neuron i, and $\delta(\cdot)$ is the Dirac delta function, which resets the voltage to zero. The function $\alpha(t)$ is a postsynaptic current and is nonzero only for $t > 0$. The connection weight between neuron i and neuron j is J_{ij}. The sum over m and l extends over the entire firing history of the neurons in the network, and the sum over j extends over the network. Each time the voltage crosses the threshold from below the neuron is said to

"fire." The voltage then immediately resets to $v_i = 0$, and a synaptic pulse $\alpha(t)$ is sent to all connected neurons.

In our examination of bump solutions, we will consider subthreshold input ($I_i < 1$) and a weight matrix that is translationally invariant (i.e., J_{ij} depends on only $|i - j|$). It is of the lateral inhibition form (locally excitatory but distally inhibitory); this type of connectivity matrix can be shown to arise from a multilayer network with both inhibitory and excitatory populations if the inhibition is fast, as shown by Ermentrout (1998).

We can formally integrate equation 2.1 to obtain the spike response form (Gerstner, 1995; Gerstner, van Hemmen, & Cowan, 1996; Chow, 1998). This form will allow us to relate the bump profile for the integrate-and-fire network to the profile of a rate model similar to that studied by Amari (1977). Suppose that neuron i has fired in the past at times t_i^l, where $l = 0, -1, -2, \ldots, -\infty$. The neuron most recently fired at t_i^0. We consider the dynamics for $t > t_i^0$. Integrating equation 2.1 yields

$$v_i(t) = I_i(1 - e^{-(t-t_i^0)}) + \sum_{j,m} \frac{J_{ij}}{N} \int_{t_i^0}^{t} e^{s-t} \alpha(s - t_j^m)\, ds. \tag{2.2}$$

By breaking up the integral in equation 2.2 into two pieces, we obtain

$$\begin{aligned} v_i(t) = \; & I_i(1 - e^{-(t-t_i^0)}) + \sum_{j,m} \frac{J_{ij}}{N} \int_{-\infty}^{t} e^{s-t} \alpha(s - t_j^m)\, ds - e^{-(t-t_i^0)} \\ & \times \sum_{j,m} \frac{J_{ij}}{N} \int_{-\infty}^{t_i^0} e^{s-t_i^0} \alpha(s - t_j^m)\, ds, \end{aligned} \tag{2.3}$$

from which we obtain the spike response form

$$v_i(t, s) = I_i - [I_i + u_i(s)]e^{-(t-s)} + u_i(t), \qquad t > s, \qquad v_i(s, s) = 0, \tag{2.4}$$

where

$$u_i(t) = \sum_{j,m} \frac{J_{ij}}{N} \epsilon(t - t_j^m), \tag{2.5}$$

and

$$\epsilon(t) = \int_0^t e^{s-t} \alpha(s) ds. \tag{2.6}$$

We normalize $\epsilon(t)$ so that $\int_0^\infty \epsilon(t)\, dt = 1$. Note that for low rates of firing, $[I + u_i(s)]e^{-(t-s)} \simeq e^{-(t-s)}$.

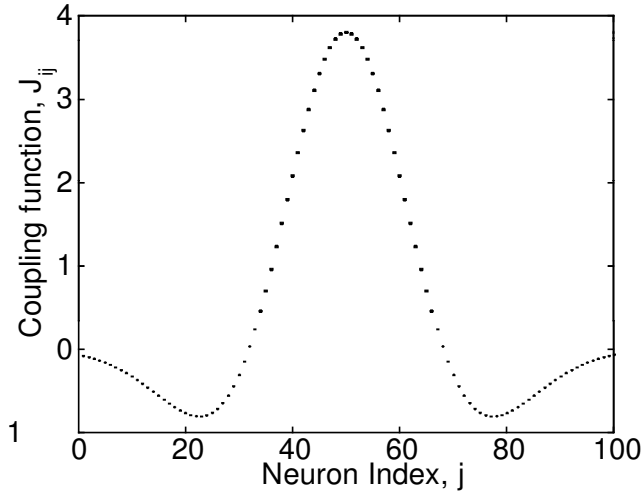

Figure 1: The weight function $J_{ij} = NJ(|i-j|/N)$ for $i = 50$ and $N = 100$, where $J(z) = 5[1.1w(1/28, z) - w(1/20, z)]$, and $w(a, z) = (a\pi)^{-1/2} \exp(-z^2/a)$.

2.1 Numerical Methods. We performed numerical simulations of the integrate-and-fire network using the spike response form equation 2.2. We step equation 2.2 forward using a fixed time step, Δt, until one or more of the voltages is above threshold (set to be 1). Assume that $v_i([n+1]\Delta t) > 1$ while $v_i(n\Delta t) < 1$, for some n. At this point, a backward linear interpolation in voltage is made to determine the approximate firing time of neuron i. To determine the approximate correct value of $v_i([n+1]\Delta t)$, we make an Euler step from the approximate last firing time of neuron i to time $[n+1]\Delta t$, using equation 2.1 for the slope. The equations are then stepped forward again.

The domain is the unit interval with periodic boundary conditions, and the weight function involves a difference of gaussians (see Figure 1 for an example). For the synaptic pulse, $\alpha(t)$, we take

$$\alpha(t) = \beta \exp(-\beta t), \tag{2.7}$$

so that

$$\epsilon(t) = \frac{\beta[e^{-t} - e^{-\beta t}]}{\beta - 1}. \tag{2.8}$$

The parameter β affects the rate at which the postsynaptic current decays. Noise is added to the network as current pulses to each neuron of the form

$$I_{\text{rand}}(t) = 6(e^{-10t} - e^{-15t}),$$

where $t \geq 0$. The arrival times of these pulses have a Poisson distribution with mean frequency 0.05, and there is no correlation between pulse arrival times for different neurons.

3 Existence of the Bump State

We examine the existence of bump solutions of the spike response system described by equation 2.4. A bump solution is spatially localized with spatially dependent average firing rate of the participating neurons. The firing rate is zero outside the bump and rises from zero at the edges to a maximum in the center. The firing times of the neurons are uncorrelated, so the bump is a localized patch of incoherent or asynchronous firing. The state coexists with the homogeneous nonfiring (all-off) state.

It is convenient to define the activity of neuron i as

$$A_i(t) = \sum_l \delta(t - t_i^l), \tag{3.1}$$

where the sum over l is over all past firing times. Our activity differs from the population activity of Gerstner (1995), which considers the activity of an infinite pool of neurons at a given spatial location. We can then rewrite the synaptic input 2.5 in terms of the activity as

$$u_i(t) = \sum_j \frac{J_{ij}}{N} \int_0^\infty \epsilon(s) A_j(t - s)\, ds. \tag{3.2}$$

Consider stationary asynchronous solutions to the spike response equations. Many authors have studied the spatially homogeneous asynchronous state with various coupling schemes (Abbott & Van Vreeswijk, 1993; Treves, 1993; Gerstner, 1995, 1998, 2000). Our approach is similar to that of Gerstner (1995, 1998, 2000). We first rewrite the activity as

$$A_i(t) = A_i^0 + \Delta A_i(t), \tag{3.3}$$

where

$$A_i^0 = \lim_{\tau \to \infty} \frac{1}{\tau} \int_0^\tau A_i(r) dr. \tag{3.4}$$

Substituting equation 3.1 into 3.4 then yields $A_i^0 = \lim_{\tau \to \infty} n(\tau)/\tau$, where $n(\tau)$ is the number of times neuron i fired in the time interval τ. Thus, A_i^0 is the mean firing rate of neuron i.

We now insert equation 3.3 into 3.2 to obtain $u_i(t) = u_i^0 + \Delta u_i(t)$, where

$$u_i^0 = \sum_j \frac{J_{ij}}{N} A_j^0, \tag{3.5}$$

and

$$\Delta u_i(t) = \sum_j \frac{J_{ij}}{N} \int_0^\infty \epsilon(s) \Delta A_j(t-s) ds \tag{3.6}$$

(recall that $\int_0^\infty \epsilon(s)\, ds = 1$). We define the asynchronous state to be one where $\Delta u_i(t)$ is zero in the limit of infinite network size N. In the asynchronous state, the input to neuron i is a constant and given by u_i^0. This implies that the firing times of the neurons are uncorrelated. For a finite system, $\Delta u(t)$ will contribute fluctuations that scale as $N^{-1/2}$.

We now derive the self-consistent equations for the asynchronous state. Substitute $u_i(t) = u_i^0$ into equation 2.4; the local firing period $(A_i^0)^{-1}$ will be given by

$$v_i((A_i^0)^{-1} + s, s) = 1 = I_i - [I_i + u_i^0]e^{-(A_i^0)^{-1}} + u_i^0. \tag{3.7}$$

Solving equation 3.7 yields

$$A_i^0 = G[u_i^0], \tag{3.8}$$

where

$$G[z] = \begin{cases} 0, & z \le 1 - I \\ -1/\ln\left[\frac{I+z-1}{I+z}\right], & z > 1 - I. \end{cases} \tag{3.9}$$

(A plot of $G[z]$ can be seen in Figure 2.) This form is similar to the usual neural network rate equation (Amari, 1977; Kishimoto & Amari, 1979; Hansel & Sompolinsky, 1998; Ermentrout, 1998) except that the gain function we have derived is a result of the intrinsic neuronal dynamics of our model (Gerstner, 1995). Combining equations 3.5 and 3.8, we obtain the condition for a stationary asynchronous solution:

$$u_i^0 = \sum_j \frac{J_{ij}}{N} G[u_j^0]. \tag{3.10}$$

For a finite sized system, the time-averaged firing rate of the neurons follows a profile given by $A_j^0 = G[u_j^0]$.

We first consider mean-field solutions to equation 3.10. We assume that $u_i^0 = u^0$, $I_i = I$ and $\sum J_{ij}/N = J$, yielding

$$u^0 = JG[u^0]. \tag{3.11}$$

If $I > 1$ (oscillatory neurons), then there are no solutions if J is too large

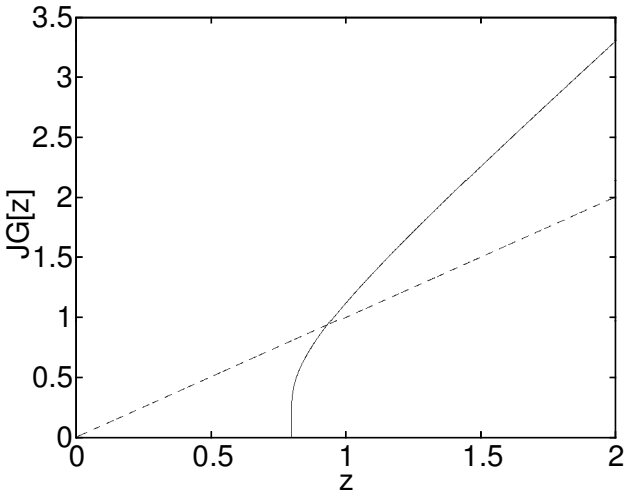

Figure 2: A multiple (J) of the gain function $G[z]$ (see equation 3.9) (solid line) for $I = 0.2$ and $J = 2$, and the diagonal (dashed). The threshold occurs at $z = 1 - I$, when 1 is the voltage threshold for the integrate-and-fire neuron. The intersection of the dashed line and the gain function gives the mean-field solution to equation 3.11.

and one solution if J is small enough. For $I < 1$ (excitatory neurons), equation 3.11 has one solution at $u_0 = 0$ if J is too small and two solutions if J is large enough (See Figure 2, which shows the case $J = 2$.) These two states correspond to an "all-off" state and an "all-on" state, respectively.

3.1 Bump State. In order for a bump to exist, a solution to equation 3.10 for u_i^0 must be found such that $u_i^0 + I_i$ is above threshold ($u_i^0 + I_i > 1$) in a localized region of space. We show example figures of such solutions in Figures 3 and 4. Amari (1977) and Kishimoto & Amari (1979) proved that such a solution can exist for a class of gain functions $G[z]$. Similar to the mean-field solution, we find that for subthreshold input ($I_i < 1$), the all-off state always exists, and the bump state can exist if the weight function has enough excitation. We will discuss stability of the bump in section 4. Stability will be affected by the synaptic timescale, the weight function, the amount of applied current, and the size of the network. For a finite-sized system, we show in the appendix that the individual neurons in a bump do not fire with a fixed spatially dependent period. These finite-sized fluctuations act as a source of noise.

As noted by Gerstner (1995, 1998), the spike response model can be connected to classical neural network or population rate models (Wilson & Cowan, 1972; Amari, 1977; Hopfield, 1984). If we choose $\epsilon(s) = e^{-s}$,

which is true for $\alpha(t) = \delta(t)$, and assume near-synchronous firing so that $A_i(t) \simeq G[u_i(t)]$, then by differentiating equation 3.2 with respect to time we obtain:

$$\frac{d}{dt}u_i(t) = -u_i(t) + \sum_j J_{ij}G[u_j(t)]. \tag{3.12}$$

This is the classical neural network or population rate model. Amit and Tsodyks (1991), Gerstner (1995), and Shriki, Sompolinsky, and Hansel (1999) have previously shown that networks of spiking neurons can be represented with rate models provided there does not exist a large degree of synchrony. The condition for a stationary solution for equation 3.12 is identical to equation 3.10. If the synapse has a more complicated time course such as a difference of exponentials, then a higher-order rate equation could be derived (Ermentrout, 1998). The activity is a functional of the input as well as a function of time. The assumption made in deriving equation 3.12 is that the explicit time dependence of the activity is weak compared to the dependence on the input $u_i(t)$. This is valid only for weakly synchronous or correlated firing. For strongly synchronous firing, the explicit time dependence of the activity would dominate the functional dependence on the inputs.

3.2 Numerical Simulations. We compare stationary bump profiles obtained from both equation 3.12 and a network of integrate-and-fire neurons. A space-time raster plot of the firing times of a stationary bump from a simulation of a network of 100 integrate-and-fire neurons is shown in Figure 3. The network was switched into the bump state by applying a transient spatially localized current for sufficient time to excite the bump. This state has a large basin of attraction; as long as neurons are excited in a localized region, the network relaxes into a bump of the form shown in Figure 3. The raster plot shows that the bump is localized in space and persists for many firing times. The neurons in the center of the bump fire much faster than those near the edges.

Figure 4 shows the profile of the average firing rate (activity) for the integrate-and-fire system, equation 2.2, with three different values of β and no noise. The solid line corresponds to the theoretically predicted profile from equation 3.10. This corresponds to the stable stationary solution of the corresponding rate model, equation 3.12, with the gain function, equation 3.9. The comparison for small β (i.e., slow synapses) is excellent; as β is increased, the agreement lessens but is still very good. We also have found bumps in networks of conductance-based neurons (results not shown; see Gutkin, Laing, Chow, Colby, & Ermentrout, 2000, for an example). We find that the existence of a bump solution for a network of spiking neurons to be robust.

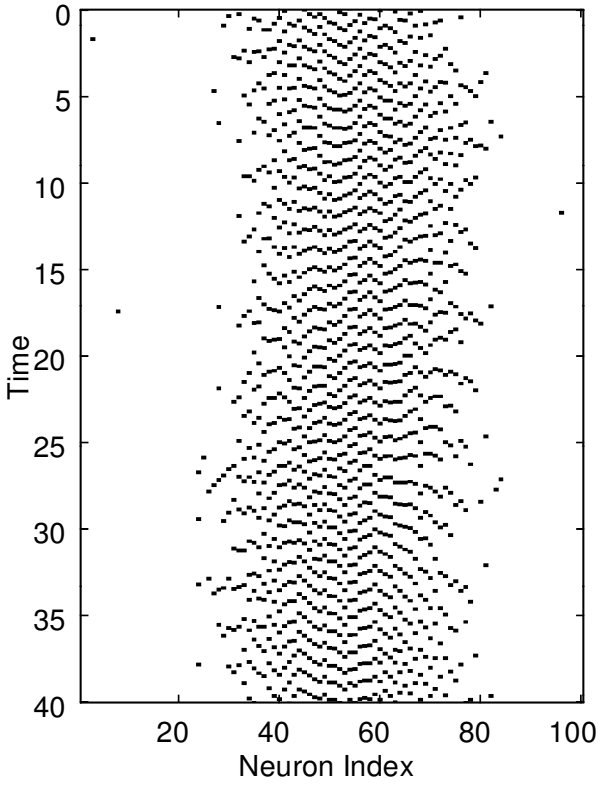

Figure 3: An example of a stationary bump. The dots represent the firing times of a network of 100 neurons. Noise in the form of random current pulses with a mean frequency of 0.05 as discussed in section 2.1 was included in the simulation. Parameter values used were $\beta = 1.5$, $I = 0.9$, and J_{ij}, as in Figure 1.

4 Stability and the Synaptic Timescale

In order for the bump to be observable, it must be stable to perturbations. A stability analysis of the spike response system can be performed by considering the linear behavior of small perturbations around the stationary bump state. However, for the spatially inhomogeneous bump, this computation is quite involved. Instead, we infer the conditions for stability of the bump from a stability analysis of the homogeneous asynchronous state of the spike response model and confirm our conjectures with numerical simulations.

Stability of the bump state has previously been examined in a first-order

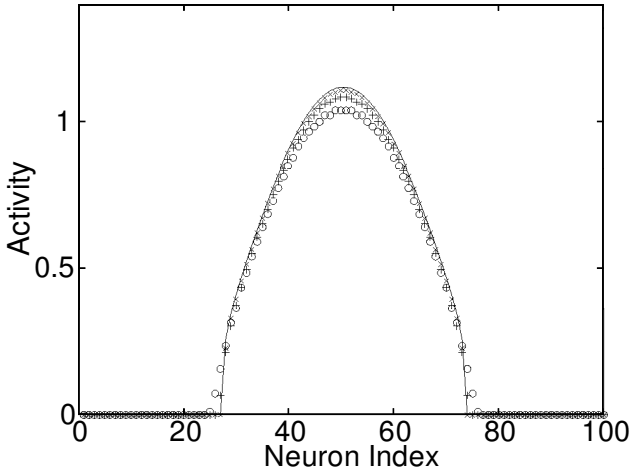

Figure 4: Activity profile (mean firing rate) for $\beta = 2.5$ ($\circ$), $\beta = 1.5$ ($+$), and $\beta = 0.5$ ($\times$) for the integrate-and-fire network and for the rate model, equation 3.12, (solid line). $I = 0.9$, and the weight function is as in Figure 1. Note that for smaller β, the agreement between the firing-rate model and the integrate-and-fire model is better. The errors in the activity values are all less than 0.01.

rate model. Amari (1977) and Kishimoto and Amari (1979) found for saturating gain functions that the stability of the bump in the rate model depended on a relationship between the weight function and the applied current. Hansel and Sompolinsky (1998) found the stability constraints for a model with a simplified weight function on a periodic domain and a piecewise linear gain function. However, as discussed in section 3.1 the rate model is valid only for infinitely fast synapses and asynchronously firing neurons. If correlations develop between firing times of neurons, the rate model is no longer valid.

It has been shown previously for homogeneous networks (Abbott & Van Vreeswijk, 1993; Treves, 1993; Gerstner, 1995, 1998, 2000) that the asynchronous state of a network of integrate-and-fire neurons is unstable to oscillations with fast excitatory coupling. Gerstner (1998, 2000) has shown that oscillations will develop at harmonics of the average firing rate (activity). The addition of noise helps to stabilize the asynchronous state. However, for low to moderate levels of noise, a fast enough synapse can still cause an instability. These calculations are based on perturbing the stationary asynchronous state for an infinite number of neurons using a Fokker-Planck or related integral formalism. We note that there are two sources of noise in the simulations: the randomly arriving currents, I_{rand}, and the fluctuations due to the finite size of the network. The latter is a manifes-

tation of unpredictability in a high-dimensional, deterministic, dynamical system.

The integrate-and-fire network has a parameter, β, that controls the timescale of the synaptic input. We find that varying β has a profound effect on the stability of the bump solution. If β is slowly increased (i.e., the synaptic timescale is slowly shortened) while all other parameters are held constant, the bump will eventually lose stability. The initial location of a bump is determined by the spatial structure of the initial current stimulation. If the initial condition is symmetric and no noise is added, the bump will remain in place and then lose stability directly to the all-off state at a critical value of β. However, as a result of the invariance of the network under spatial translations, a bump is marginally stable with respect to spatial translation; there is a continuum of attractors parameterized by their spatial location rather than a finite set of isolated attractors. A consequence of this marginal stability is that a bump may "wander" under the influence of noise or finite size fluctuations. With asymmetric initial conditions or noise, the bump will begin to wander as β is increased (see Figure 5). (Figure 10 shows the bump wandering for a fixed value of β.)

At a larger value of β, the wandering bump loses stability to a traveling wave. Note that the wave speed increases as β increases. The same type of behavior was observed for larger networks (results not shown). If the wave hits an obstruction, it will switch to the all-off state. The bump can be pinned in place if a small amount of disorder or heterogeneity is included in the input current. In Figure 6 we have added disorder by randomly choosing the fixed currents, I_i, keeping the average of these values across the network equal to the value used in Figure 5. The bump migrates to a local maximum in the current and remains there. For sufficiently large β, it destabilizes into the all-off state, as in Figure 6. The pinning due to disorder is sufficiently strong that traveling waves cannot persist. Note that small patches of neurons can be activated by noise, but they cannot persist if β is too large.

We conjecture that the loss of stability in the bump for large β is due to a loss of stability of the asynchronous bump state due to the synchronizing tendency of the neurons with fast excitatory coupling, as is seen in the homogeneous network. Integrate-and-fire neurons belong to what is known as type I or class I neurons (Hansel, Mato, & Meunier, 1995; Ermentrout, 1996). It is known that for type I neurons, fast excitation has a synchronizing tendency, whereas slow excitation has a desynchronizing tendency (Van Vreeswijk, Abbott, & Ermentrout,1994; Gerstner, 1995; Hansel et al., 1995). Chow (1998) showed that a network with heterogeneous input is still able to synchronize.

The stationary bump corresponds to an asynchronous network state where the neurons receive heterogeneous input. For fast enough excitatory coupling, the asynchronous state might lose stability, and oscillations will develop. The synchronous oscillations need not be exceptionally strong.

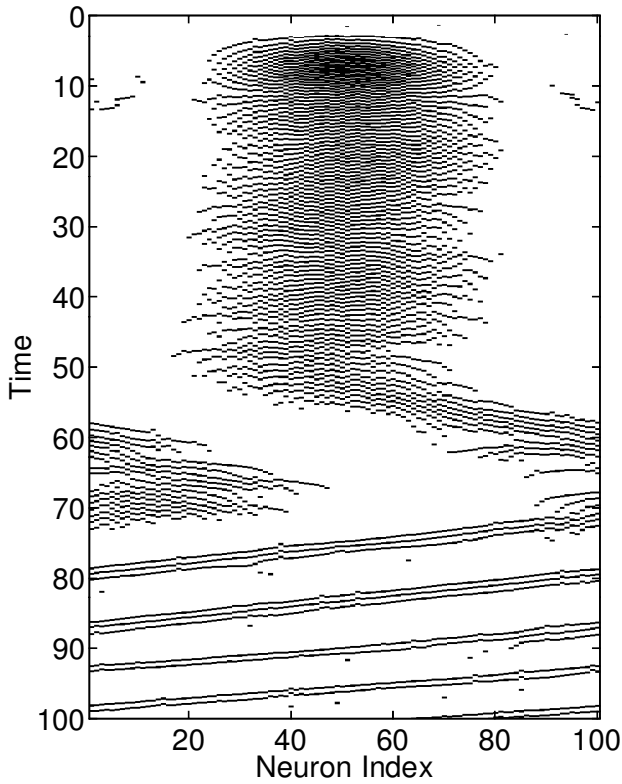

Figure 5: Bump destabilization due to quasi-static increase of β. Here, β is linearly increased in time from $\beta = 3.1$ at the top of the figure to $\beta = 8.3$ at the end. The weight function is $J(z) = 5[1.16w(1/28, z) - w(1/20, z)]$ (see the caption of Figure 1 for definiton of w), $I = 0.8$, and the noise is as in Figure 3. The bump loses stability first to wandering and then to a traveling wave.

All that is required is that large enough oscillations develop and induce a traveling wave. For symmetric initial conditions, the symmetry prevents the formation of a traveling wave. The neurons oscillate in place, and for enough synchronization, the synaptic input is not at a high enough level to push some of the neurons in the bump above threshold, when it is time for them to fire. The bump solution would then collapse and switch the network to the all-off state. This also occurs for the case with heterogeneity. The pinned bump develops stationary oscillations and switches to the all-off state.

For a fixed weight function, we find that when I is at a value for which the rate model predicts the bump to be stable, the bump in the integrate-

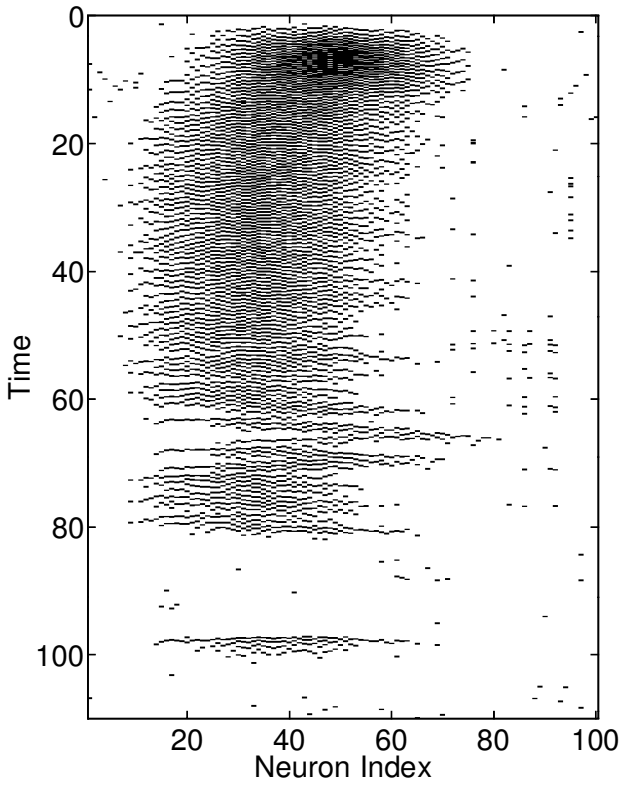

Figure 6: Bump destabilization due to quasi-static increase of β with disorder added to the static background current. The range of β, the weight function, the noise, and the average value of I are as in Figure 5. Disorder pins the bump and prevents the formation of a traveling wave.

and-fire network is stable for small β, but becomes unstable as β increases. This is shown in Figure 7, where we show the region of the $I - \beta$ plane in which there is a stable bump solution in the integrate-and-fire network with no noise for $N = 100$. The stability region, as well as the size and shape of the bump, depends on the applied current I, the weight function $J(x)$, and the size of the network, but is unchanged by adding noise of the intensity used in other simulations in this article.

The finite size fluctuations are important for the dynamics. Figure 8 shows a plot of $u_i(t)$ at the center of a bump containing 100 neurons for two different values of β. The traces show noisy oscillations at a given frequency around an average value. The dominant frequency of the oscillations is at the neuron's average firing rate. As β increases, the amplitude of the

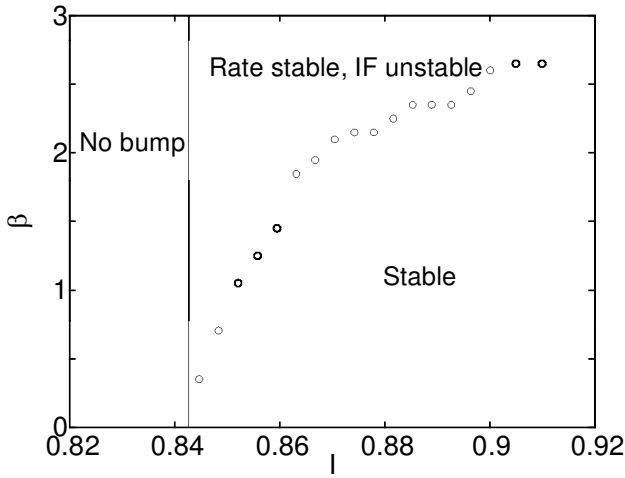

Figure 7: Numerically obtained stability boundary (circles) for a bump in the noise-free integrate-and-fire for J_{ij} as in Figure 1 and $N = 100$, as a function of I and β. The bump is stable below the curve marked by circles. The solid line shows the existence region for the bump in the rate model: it exists to the right but not to the left of the line. Above the curve marked by circles, the bump exists and is stable in the rate model but unstable in the integrate-and-fire model. As discussed in the text, the curve marked by circles will move to larger β as N is increased. This curve is not significantly changed when noise of the level used in the other simulations shown is added.

noisy oscillations increases. We conjecture that the increase in the size of the oscillations is partially due to the dynamical synchronizing effect described above. For small N, the noisy oscillations are dominated by finite size fluctuations, but for large N, it will be dominated by the synchronizing effect of the fast excitation. We believe that the destabilizing effect of the oscillations will be present even for infinite N. While increasing N does decrease the size of the fluctuations, this is more than compensated by the increase in oscillation size due to increasing β. In Figure 9 we plot the maximum value of β for which a bump is stable as a function of N for various noise values. The plot shows that the stability boundary asymptotes for a fixed value of β for large values of N. This plot suggests that for any N, there is a finite β above which the bump is no longer stable even in the presence of noise. As expected, the bump can tolerate higher levels of β as the noise level increases, although the maximum β seems to saturate as noise increases.

From Figures 4 and 8, we see that as β is increased, the activity and mean value of u at each neuron decrease slightly despite the fact that we have compensated the synaptic strength so the integrated synaptic input is

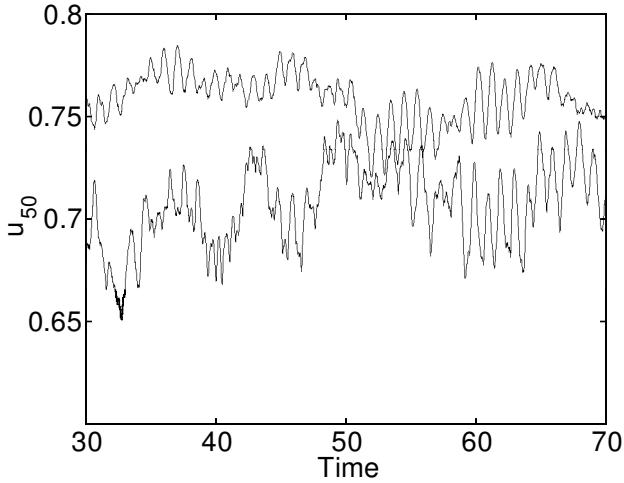

Figure 8: $u_i(t)$ at the center of a bump for $\beta = 0.9$ (upper line) and $\beta = 1.9$ (lower line) for $I = 0.9$, J_{ij} as in Figure 1, and noise as in Figure 3. As β increases, the size of the noisy oscillations increases.

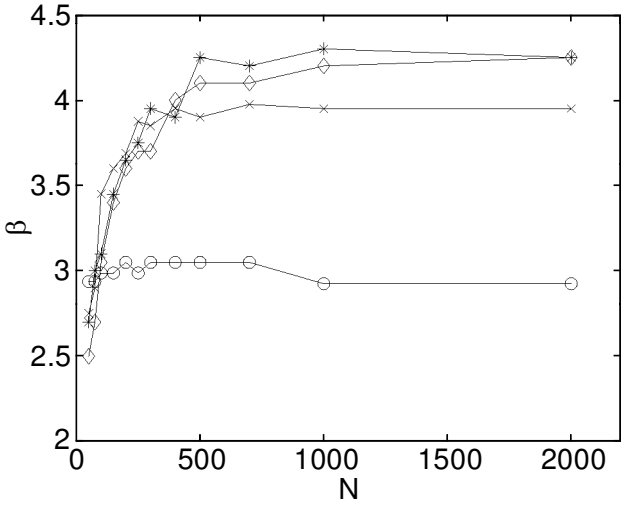

Figure 9: Values of β above which the bump is not stable as a function of N in the network for $I = 0.91$, J_{ij} as in Figure 1, and the mean frequencies of the random current pulses have values: no noise ($\circ$), 0.05 ($\times$), 0.1 ($*$), and 0.2 ($\diamond$).

constant. We believe that this small decrease is a result of the fluctuations in the input due to the finite number of neurons and the synchronizing dynamical effect. The termination of the bump as β is increased is not due to this overall decrease in synaptic input. If the coupling weight is increased to keep the mean value of u constant (by multiplying $J(z)$ by a factor slightly greater than 1), the bump is still seen to terminate as β is increased. Indeed, numerical results suggest that the mean value of u decreases linearly (albeit weakly) with β, while its standard deviation increases at a faster than linear rate (results not shown).

5 Initiating and Terminating Bumps

Since bumps are thought to be involved in such areas as working memory (Colby et al., 1995; Funahashi et al., 1989) and head direction systems (Redish et al., 1996; Zhang, 1996), understanding their dynamics on a timescale much longer than the period of oscillation of the individual neurons is important. In particular, being able to "turn on" and "turn off" a bump is of interest. Since the input current I is subthreshold, the network is bistable, the two stable states being the bump and the all-off state, where none of the neurons fires. Turning the bump on is simply a matter of shifting the system from the all-off state into the basin of attraction of the bump. This is done by applying a spatially localized input current to the already existing uniform current for a short period of time—on the order of five oscillation periods of the neuron in the center of the bump (whose activity is highest), as is shown in Figure 10. Note that the bump persists after the stimulation has been removed.

As we have seen, partially synchronizing the neurons can cause the bump to terminate. One mechanism for causing some of the neurons to synchronize is to apply a brief, strong, excitatory current to most or all of the neurons involved in the bump. This will cause a number of the neurons to fire together (i.e., be temporarily synchronized), leading to termination of the bump if the synapse is fast enough. An example is shown in Figure 10, where a short stimulus is applied to all of the neurons involved in the bump, although not to all of the neurons in the network. The stimuli to turn on and turn off the bump are both excitatory; the stimulus to turn off a bump does not have to be inhibitory, although that is also successful.

The reason the bump turns off is a dynamical effect. An alternate means of turning off the bump with excitation is to use the lateral inhibition in the network. If the parameters are tuned correctly, then a strong excitatory input to all of the neurons in the network will induce enough inhibition to reduce the synaptic input below threshold and extinguish the bump. As discussed in section 4, the bump is seen to wander in Figure 10. A small amount of disorder could prevent this occurrence by pinning the bump to a given location. We explore the initiation and termination of the bumps more carefully in a companion publication. (Gutkin et al., 2000).

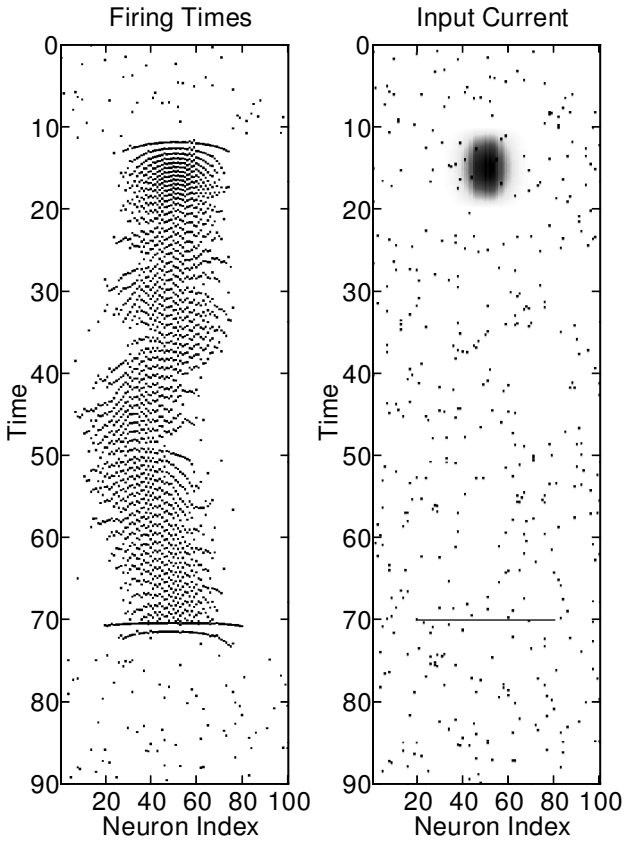

Figure 10: Turning on and off a bump with excitation. The plot on the left shows the firing times and that on the right shows the input current (maximum amplitude above background 0.4). The coupling weight is as in Figure 1, $\beta = 2.4$, $I = 0.9$, and noise is as in Figure 3.

6 Discussion and Conclusion

We have shown that a one-dimensional network of spiking neurons can sustain spatially localized bumps of activity and that the profiles of the bumps agree with those predicted by a corresponding population rate model. However, when the synapses occur on a fast timescale, bumps can no longer be sustained in the network. They either lose stability to traveling waves or completely switch off. We also find that heterogeneity or disorder can pin the bumps to a single location and keep them from wandering. We conjec-

ture that the loss of stability of the bump is due to partial synchronization between the neurons. It is known for homogeneous networks of type 1 neurons that fast excitatory synapses have a synchronizing tendency. We use this instability to turn off bumps with a brief excitatory stimulus to synchronize the neurons partially.

For the network sizes that we have probed, we have found that bumps can be sustained by synapses with decay rates as fast as three to four times the firing rate of the fastest neurons in the bump. If we consider neurons in the cortex to be firing at approximately 40 Hz, this would correspond to synaptic decay times of the order of 5 to 10 ms, which is not unreasonable. Results with conductance-based neurons have found that the synaptic timescale can be speeded up to well within the AMPA range and still sustain a bump state (Gutkin et al., 2000). We also find that as the network size increases, the bump may tolerate faster synapses. While the stability of the bump depends crucially on the synaptic timescale, the activity profile of the bump depends on only the connection weights and the gain function. Thus, it may be possible to make predictions on the connectivity patterns of experimental cortical systems from the firing rates of the neurons within the bump and the firing rate (F-I) curve of individual neurons.

If these recurrent bumps are involved in working memory tasks, then our results lead to some experimental predictions. For example, if it is possible to speed up the excitatory excitations in the cortex pharmacologically, bump formation and hence working memory may be perturbed. A brief applied stimulus applied to the cortical area where the working memory is thought to be held may also disrupt a working memory task.

Among other authors who have produced similar work are Hansel and Sompolinsky (1998), Bressloff et al. (1999), and Compte, Brunel, and Wang (1999). Hansel and Sompolinsky (1998) consider a rate model similar to that studied by Amari (1977) and Kishimoto and Amari (1979), using a piecewise linear gain function (our $G[z]$) and retaining only the first two Fourier components of the weight function J, which allows them to make analytic predictions about the transitions between different types of behavior. They also show the existence of a bump in a network of conductance-based model neurons and show that bumps can follow moving spatially localized current stimulations, a feature that may be relevant for head-direction systems such as those studied by Redish et al. (1996) and Zhang (1996).

Bressloff and Coombes (1998) and Bresslof et al. (1999) study pattern formation in a network of coupled integrate-and-fire neurons, but their systems consider suprathreshold input ($I_i > 1$) so that the all-off state is not a solution. They find that by increasing the coupling weight between neurons, the spatially uniform synchronized state (all neurons behave identically) becomes unstable through a Turing-Hopf bifurcation, leading to spatial patterns similar to those shown in Figure 4. They find bistability between a bump and a spatially uniform synchronized state, whereas we find bistability between a bump and the all-off state. This difference is crucial if the

system is to be thought of as modeling working memory as investigated by, among others, Colby et al. (1995) and Funahashi et al. (1989).

Compte et al. (1999) have demonstrated the existence of a bump attractor in a two-layer network of excitatory and inhibitory integrate-and-fire neurons. Their network involves strong excitation and inhibition in a balanced state. It is possible that a corresponding rate model could be found for this network to obtain the shape of the profile. They were also able to switch the bump off and on with an excitatory stimulus. However, it is believed that their switching-off mechanism is due to the inhibitory input induced from the excitation provided.

We have also observed bumps in two-dimensional networks. A number of other authors, including Fohlmeister, Gerstner, Ritz, and van Hemmen (1995) and Horn and Opher (1997), have investigated two-dimensional networks of integrate-and-fire neurons, but not in the context of bumps. Preliminary investigations (results not shown) indicate that bumps can be initiated and terminated in two-dimensional networks in exactly the same fashion as for one-dimensional networks. We emphasize that while the bumps are spatially localized, they need not be contiguous. It may be such that a certain amount of disorder in the synaptic connectivity would lead to a bump that is distributed over a given region. This disorder may ever confer some beneficial effects by breaking the translational symmetry and keep the bump from wandering when under the influence of stochastic firing. We hope to elucidate these effects in the future.

Appendix: Nonperiodicity of Bump Solutions

Here we show that locally periodic solutions are not possible for a bump in a finite-sized integrate-and-fire network. We consider the ansatz of periodic firing given by $t_i^{m_i} = (\phi_i - m_i)T_i$, where ϕ_i is a phase, m_i is an integer, and T_i is the local firing period. We suppose that the neuron previously fired at t_i^{-1} and will next fire at t_i^0. We now substitute this ansatz into the firing equation, equation 2.4, to obtain

$$1 = I_i - [I_i - u_i(t_i^{-1})]e^{-T_i} + u_i(t_i^0), \tag{A.1}$$

where

$$u_i(t) = \int J_{ij} \sum_{m_j} \epsilon(t - (\phi_j - m_j)T_j). \tag{A.2}$$

In order to find a solution for T_i that is constant in time, we require $u_i(t)$ to be constant or T_i periodic in t. However, $\sum_{m_j} \epsilon(t - (\phi_j - m_j)T_j) \equiv f(t)$ is T_j periodic and T_j is not constant in j for a bump solution. This implies that $u_i(t)$ cannot be T_i periodic for all i (except for infinite N or for highly singular cases where all the periods are rationally related). Hence, a locally

periodic bump solution is not possible for the finite-sized spike response model and hence for the integrate-and-fire network.

Acknowledgments

We thank D. Pinto and G. B. Ermentrout for stimulating discussions that spurred this research. We also thank A. Compte and N. Brunel for some clarifying conversations. We especially thank Boris Gutkin for many fruitful discussions and for his critical insight. This work was supported in part by the National Institutes of Health (C. C.) and the A. P. Sloan Foundation (C. C.).

References

Abbott, L. F., & Van Vreeswijk, C. (1993). Asynchronous states in a network of pulse-coupled oscillators. *Phys. Rev. E*, 48, 1483–1490.

Amari, S. (1977). Dynamics of pattern formation in lateral–inhibition type neural fields. *Biol. Cybern.*, 27, 77–87.

Amit, D. J., & Brunel, N. (1997). Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb. Cortex*, 7, 237–252.

Amit, D. J., & Tsodyks, M. V. (1991). Quantitative study of attractor neural network retrieving at low spike rates: I. Substrate—spikes, rates and neuronal gain. *Network*, 2, 259–273.

Bresslof, P. C., Bresslof, N. W., & Cowan, J. D. (1999). *Dynamical mechanism for sharp orientation tuning in an integrate-and-fire model of a cortical hypercolumn*. Preprint 99/13, Department of Mathematical Sciences, Loughborough University.

Bresslof, P. C., & Coombes, S. (1998). Spike train dynamics underlying pattern formation in integrate-and-fire oscillator networks. *Phys. Rev. Lett.*, 81(11), 2384–2387.

Camperi, M., & Wang, X.-J. (1998). A model of visuospatial working memory in prefrontal cortex: Recurrent network and cellular bistability. *J. Comput. Neurosci.*, 5, 383–405.

Chow, C. C. (1998). Phase-locking in weakly heterogeneous neuronal networks. *Physica D*, 118, 343–370.

Colby, C. L., Duhamel, J.-R., & Goldberg, M. E. (1995). Oculocentric spatial representation in parietal cortex. *Cerebral Cortex*, 5, 470–481.

Compte, A., Brunel, N., & Wang, X.-J. (1999). *Spontaneous and spatially tuned persistent activity in a cortical working memory model*. Abstract, Computational Neuroscience Meeting.

Ermentrout, G. B. (1996). Type I membranes, phase resetting curves, and synchrony. *Neural Comp.*, 8, 979–1001.

Ermentrout, G. B. (1998). Neural networks as spatio-temporal pattern-forming systems. *Reports on Progress in Physics*, 61(4), 353–430.

Fohlmeister, C., Gerstner, W., Ritz, R., & van Hemmen, J. L. (1995). Spontaneous excitations in the visual cortex: Stripes, spirals, rings and collective bursts. *Neural Comp.*, 7(5), 905–914.

Funahashi, S., Bruce, C. J., & Goldman-Rakic, P. S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J. Neurophysiology*, 61(2), 331–349.

Gerstner, W. (1995). Time structure of the activity in neural networks models. *Phys. Rev. E*, 51, 738–758.

Gerstner, W. (1998). Populations of spiking neurons. In W. Maass & C. M. Bishop (Eds.), Pulsed neural networks (pp. 261–295). Cambridge, MA: MIT Press.

Gerstner, W. (2000). Population dynamics of spiking neurons: Fast transients, asynchronous states, and locking. *Neural Comput.*, 12, 43–89.

Gerstner, W., van Hemmen, J. L., & Cowan, J. (1996). What matters in neuronal locking? *Neural Comput.*, 8, 1653–1676.

Gutkin, B. S., Laing, C. R., Chow, C. C., Colby, C., & Ermentrout, G. B. (2000). *Turning on and off with excitation: The role of spike-timing asynchrony and synchrony in sustained neural activity*. Preprint.

Hansel, D., Mato, G., & Meunier, C. (1995). Synchrony in excitatory neural networks. *Neural Comput.*, 7, 307–335.

Hansel, D., & Sompolinsky, H. (1998). Modeling feature selectivity in local cortical circuits. In C. Koch & I. Segev (Eds.), *Methods in neuronal modeling* (2nd ed.). Cambridge, MA: MIT Press.

Hopfield, J. J. (1984). Neurons with graded response have collective computational properties like those of two-state neurons. *Proc. Natl. Acad. Sci. USA*, 81, 3088–3092.

Horn, D., & Opher, I. (1997). Solitary waves of integrate-and-fire neural fields. *Neural Comp.*, 9(8), 1677–1690.

Kishimoto, K., & Amari, S. (1979). Existence and stability of local excitations in homogeneous neural fields. *J. Math. Biology*, 7, 303–318.

Redish, A. D., Elga, A. N., & Touretzky, D. S. (1996). A coupled attractor model of the rodent head direction system. *Network*, 7, 671–685.

Shriki, O., Sompolinsky, H., & Hansel, D. (1999). *Rate models for conductance based cortical neuronal networks*. Preprint.

Skaggs, W. E., Knierim, J. J., Kudrimoti, H. S., & McNaughton, B. L. (1995). A model of the neural basis of the rat's sense of direction. In G. Tesauro, D. Touretzky, & T. Leen (Eds.), *Advances in neural information processing systems, 7* (pp. 173–180). Cambridge, MA: MIT Press.

Somers, D. C., Nelson, S. B., & Sur, M. (1995). An emergent model of orientation selectivity in cat visual cortical simple cells. *J. Neuroscience*, 15, 5448–5465.

Treves, A. (1993). Mean-field analysis of neuronal spike dynamics. *Network*, 4, 259–284.

Van Vreeswijk, C., Abbott, L. F., & Ermentrout, G. B. (1994). When inhibition not excitation synchronizes neural firing. *J. Comp. Neurosci.*, 1, 313–321.

Wilson, H. R., & Cowan, J. D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical J.*, 12, 1–24.

Wilson, H. R., & Cowan, J. D. (1973). A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Kybernetic*, 13, 55–80.

Zhang, K. (1996). Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensembles: A theory. *J. Neuroscience*, 16, 2112–2126.

Received August 4, 1999; accepted September 25, 2000.