

Signal Propagation and Logic Gating in Networks of Integrate-and-Fire Neurons

Tim P. Vogels and L. F. Abbott

Volen Center for Complex Systems and Department of Biology, Brandeis University, Waltham, Massachusetts 02454-9110

Transmission of signals within the brain is essential for cognitive function, but it is not clear how neural circuits support reliable and accurate signal propagation over a sufficiently large dynamic range. Two modes of propagation have been studied: synfire chains, in which synchronous activity travels through feedforward layers of a neuronal network, and the propagation of fluctuations in firing rate across these layers. In both cases, a sufficient amount of noise, which was added to previous models from an external source, had to be included to support stable propagation. Sparse, randomly connected networks of spiking model neurons can generate chaotic patterns of activity. We investigate whether this activity, which is a more realistic noise source, is sufficient to allow for signal transmission. We find that, for rate-coded signals but not for synfire chains, such networks support robust and accurate signal reproduction through up to six layers if appropriate adjustments are made in synaptic strengths. We investigate the factors affecting transmission and show that multiple signals can propagate simultaneously along different pathways. Using this feature, we show how different types of logic gates can arise within the architecture of the random network through the strengthening of specific synapses.

Key words: rate coding; sensory processing; propagation; integrate-and-fire neurons; network models; synfire chains; logic gates

Introduction

Cognitive processing involves signal propagation through multiple brain regions and the activation of large numbers of specific neurons. Computational approaches are useful for studying the nature and mechanisms of this phenomenon. Two different modes of signal propagation have been proposed. In synfire propagation, the signal is carried by a wave of synchronous neuronal activity within a subset of network neurons (Abeles, 1991; Aertsen et al., 1996; Diesmann et al., 1999). In firing-rate propagation, the wave of activity is an asynchronous elevation in the firing rate of the neurons carrying the signal (van Rossum et al., 2002). In both cases, it has proven surprisingly difficult to construct networks that support stable and robust signal propagation over a sufficiently wide dynamic range to support a significant flow of information. Specifically, adjustments must be made to keep the signal from dying out as it propagates through different subsets of network neurons and also to keep the signal from exploding and spreading to all of the neurons in the network. In addition, large amounts of noise must be included for both forms of propagation, in synfire chains to prevent synchronization from spreading across the network (Diesmann et al., 1999; Aviel et al., 2003; Mehring et al., 2003), and in firing-rate propagation to prevent synchronization of the signal-carrying neurons themselves (van Rossum et al., 2002; Litvak et al., 2003; Reyes, 2003).

Noise was included in early studies of signal propagation (Diesmann et al., 1999; van Rossum et al., 2002) by injecting random current into all of the neurons of the network. This corresponds to a source of noise that is external to the network. Using an external noise source has the advantage that the level of noise can be adjusted easily until a level appropriate for signal propagation is found. The propagating signal and noise source are also independent, so they do not interfere with each other. However, an external source of noise is not biologically realistic. Although neurons in real neural circuits show considerable irregularity and variability in their firing (Burns and Webb, 1976; Dean, 1981; Softky and Koch, 1993; Holt et al., 1996; Anderson et al., 2000), the source of this “noisy” activity is internal, not external.

A number of studies have shown that networks of sparsely connected spiking model neurons can produce highly irregular, chaotic activity without any external source of noise (van Vreeswijk and Sompolinsky, 1996, 1998; Brunel, 2000; Mehring et al., 2003; Lerchner et al., 2004; Vogels et al., 2005). These networks allow for the study of signal propagation without the unrealistic injection of external noise. This has been done in the case of synfire chains (Aviel et al., 2003; Mehring et al., 2003), but not for firing-rate propagation. In studies of synfire propagation in irregularly firing networks, a number of complications arose because of interactions between the propagating signal and the ongoing background activity of the network. Signal propagation can be achieved in such networks, but the signals can produce network-wide shock waves that can subsequently silence the background activity of the network (Aviel et al., 2003; Mehring et al., 2003). In light of these complications, we feel it is also important to consider whether firing-rate coded signals can propagate through spontaneously active networks. We must determine

Received April 11, 2005; revised Oct. 6, 2005; accepted Oct. 7, 2005.

This work was supported by the National Science Foundation (IBN-0235463), a National Institutes of Health Pioneer Award, and the Swartz Foundation. We thank Stefano Fusi, Mark Miller, and Jean-Marc Goillard for helpful comments and suggestions.

Correspondence should be addressed to Tim P. Vogels, Volen Center for Complex Systems, MS 013, Brandeis University, Waltham, MA 02454. E-mail: vogels@brandeis.edu.

DOI:10.1523/JNEUROSCI.3508-05.2005

Copyright © 2005 Society for Neuroscience 0270-6474/05/2510786-10\$15.00/0

whether such networks generate sufficient internal “noise” to support propagation, because, unlike the external noise case, the level of noise cannot be adjusted; it is fixed by the network. In addition, we must examine whether the propagating signal interferes with the ongoing background activity enough to disrupt its role as a noise source. We also study synfire propagation in the same networks.

Signal propagation has been studied primarily in feedforward networks, either standing alone or embedded into larger networks. To study both rate-code and synfire transmission through a network with self-sustained activity, we do not embed a feedforward structure into the random network used to generate background activity. Instead, we feed a signal into a selected set of network neurons, find the most likely path of transmission through the existing network, and then modify the properties of neurons and synapses along this candidate pathway to determine whether propagation can be achieved. Thus, we are studying signal propagation within the architecture of a sparsely and randomly connected network, not an artificial feedforward structure.

Although faithful signal propagation is a necessary condition for cognitive processing, by itself it does not represent any type of computational process. Once we have established signal propagation in the networks we study, we find and reinforce subcircuits within the existing architecture that provide computational functions. Specifically, we show how logic gates, switches, and memory units can be formed from multiple interacting signal propagating pathways.

Materials and Methods

Neuron model. The network we study is composed of 10,000 leaky integrate-and-fire neurons. Each integrate-and-fire neuron is characterized by a time constant, $\tau = 20$ ms, and a resting membrane potential, $V_{\text{rest}} = -60$ mV. Whenever the membrane potential crosses a spiking threshold of -50 mV, an action potential is generated and the membrane potential is reset to the resting potential, where it remains clamped for a 5 ms refractory period. To set the scale for currents and conductances in the model, we use a membrane resistance of 100 M Ω .

We model the synaptic connections between these neurons in two different ways, as currents and as conductances, resulting in either current-based (CUBA) or conductance-based (COBA) models. For the CUBA model, the subthreshold membrane potential obeys the following equation:

$$\tau \frac{dV}{dt} = (V_{\text{rest}} - V) + g_{\text{ex}}(E_{\text{ex}} - V) + g_{\text{inh}}(E_{\text{inh}} - V), \quad (1)$$

whereas, in the COBA model, the membrane voltages are calculated as follows:

$$\tau \frac{dV}{dt} = (V_{\text{rest}} - V) + g_{\text{ex}}(E_{\text{ex}} - V) + g_{\text{inh}}(E_{\text{inh}} - V). \quad (2)$$

Reversal potentials are $E_{\text{ex}} = 0$ mV and $E_{\text{inh}} = -80$ mV. The synaptic conductances (or effective synaptic conductance in the case of the CUBA model) g_{ex} and g_{inh} are expressed in units of the resting membrane conductance.

Neurons in the network are either excitatory or inhibitory. When a neuron fires, the appropriate synaptic variable of its postsynaptic targets are increased, $g_{\text{ex}} \rightarrow g_{\text{ex}} + \Delta g_{\text{ex}}$ for an excitatory presynaptic neuron and $g_{\text{inh}} \rightarrow g_{\text{inh}} + \Delta g_{\text{inh}}$ for an inhibitory presynaptic neuron. Otherwise, these parameters obey the following equations:

$$\tau_{\text{ex}} \frac{dg_{\text{ex}}}{dt} = -g_{\text{ex}} \quad (3)$$

and

$$\tau_{\text{inh}} \frac{dg_{\text{inh}}}{dt} = -g_{\text{inh}}, \quad (4)$$

with synaptic time constants $\tau_{\text{ex}} = 5$ ms and $\tau_{\text{inh}} = 10$ ms. In most cases, rather than reporting the values of Δg_{ex} and Δg_{inh} , which are the synaptic strengths, we report the resulting EPSP and IPSP sizes. These are obtained within the active network from spike-triggered average membrane potentials of postsynaptic neurons after spikes evoked within individual network neurons. The integration time step for our simulations is 0.1 ms.

Network architecture. To create self-sustained, asynchronous background activity, we chose a 4:1 ratio of excitatory to inhibitory neurons in a network of 10,000 cells and connected them to each other randomly with a connection probability of 2%. This value was chosen as a compromise between the higher connection probabilities found for neighboring neurons in cortex and the lower values for neurons separated by distance. Our results should apply to models with connection probabilities up to $\sim 10\%$. Except along the signaling pathway (see below), all excitatory synapses took the same strength, as did all of the inhibitory synapses. These two sets of strengths were adjusted to allow asynchronous activity within the network.

Input signals. To test signal propagation through the network, we generated a set of Poisson input spike trains with a firing rate $r_0(t)$. These form a 0th layer, L_0 , that provides input to the network. Input spikes generated by the Poisson process in layer 0 were fed into a layer 1 subpopulation of the network neurons, labeled L_1 , by increasing their excitatory synaptic conductances by $g_{\text{ex}} \rightarrow g_{\text{ex}} + \Delta g_0$ whenever they received an input spike from the layer 0 source. The synaptic strength Δg_0 was tuned so that the firing rates of the layer 1 neurons reproduce the input signal, that is, they track the input firing rate $r_0(t)$. To analyze propagation, we fed various signals $r_0(t)$ into the network. At first, square-wave pulses at 180 Hz lasting for 30 ms were used to assess propagation. Then, constant input firing rates were used to study the elevation of firing rates across different propagation layers L_i for $i = 1, 2, \dots, 6$. Finally, the accuracy of signal propagation was examined by constructing $r_0(t)$ from Gaussian-distributed white noise low-pass filtered at 50 ms and half-wave rectified (van Rossum et al., 2002). In addition, the temporal properties of signal propagation were analyzed using input rates that varied sinusoidally at different frequencies, by measuring the onset delay in each layer for a propagated constant stimulus, and by studying propagation for synchronous stimuli. The results we obtain by injecting layer 0 spikes into the network could also be obtained by injecting current into the layer 1 neurons.

Signaling pathways. Signal propagation is investigated along specific pathways found within the full network. We do not change the network architecture by, for example, adding a feedforward pathway to the network, and we do not require all-to-all connectivity between pathway layers. Instead, we look for existing pathways already in the network. To do this, we look for the most likely candidates for neurons that will be affected by layer 1 activity. A potential signal-propagation pathway is a series of layers of neurons that are connected to each other in a feedforward manner more highly than average because of random fluctuations in network connectivity. By tracing these most likely candidates through the existing network, we uncover a potential signal-propagation pathway.

Specifically, we identify candidate propagation pathways in the following manner. First, we choose 33 neurons randomly as layer 1 neurons that received input from layer 0. Then, by searching the networks, we find 33 neurons, each of which receives three or more synapses from the neurons in layer 1 (although the condition is three or more, the number is three $\sim 90\%$ of the time). These define layer 2. A third layer of neurons is constructed in a similar manner by finding 33 network neurons that receive three or more synaptic connections from the neurons of layer 2. An additional requirement on layer 3 neurons is that they must not receive any direct connections from layer 1. This prevents “short-circuiting” of the multilayered propagation pathway. This procedure is continued with layer i , for $i = 4, 5, 6$, defined as a set of 33 neurons, each of which receives at least three synapses from the neurons of layer $i - 1$, and receives no synapses from layers $i - 2, i - 3, \dots, 1$. In the networks we study, probabilities of connections between neurons are given by binomial distributions. This allows us to compute the expected numbers of neurons in each of the layers of a propagation pathway. The expected number of neurons satisfying the conditions for layer 2 is 250, but the “no

short-circuit" condition reduces the number in subsequent layers significantly. For layer 6, we predict only ~ 10 candidate neurons. These are, of course, mean values. In the actual simulations, we take advantage of fluctuations around the mean (there are 200 ± 14 connections per neuron) that allow for more cells in higher layers. The number 33 was chosen as a compromise between maximizing the number of neurons in each layer and the number of layers that satisfy the prescribed conditions (sometimes for layer 6, or when multiple pathways were identified, < 33 neurons were found that satisfied all of the conditions and a smaller layer size was used). Usually there are not enough cells available to support a seventh layer. Increasing the connectivity of the network decreases the number of cells that fulfill the no short-circuit rule, and hence decreases the number of layers. We call the neurons within layer 1–6 pathway neurons and the synapses between them pathway synapses. We control signal propagation along the pathway by modifying the excitability of pathway neurons or the strength of the pathway synapses.

Logic gates. To study signal processing within the network, we search for candidate subcircuits that fulfill conditions for the computation we want to achieve. The circuits we discuss are not constructed by hand, but instead are found within the random architecture of the network. Once a candidate circuit is found, adjustments are made to the synapses within the circuit, but no new connections are made nor are any existing connections eliminated.

To identify a candidate NOT gate, we search for a group of inhibitory cells that is highly connected (at least three synapses) to an output layer of cells spontaneously firing at rates between 30 and 40 Hz. We also seek out a third layer consisting of excitatory cells that have three or more synapses to the inhibitory cell population. The entire candidate gate consists of ~ 60 neurons. To make this candidate unit functional, excitatory synapses are strengthened by sixfold and inhibitory synapses by twofold.

To identify a neuronal "switch," we search for a regular pathway consisting of three excitatory layers, as described above. Additionally, we found an inhibitory cell population with three or more synapses per cell to the output layer of the excitatory pathway. Driving these inhibitory cells with a population of highly connected excitatory neurons makes it possible to control propagation along the excitatory pathway. To make this candidate unit functional, excitatory synapses are strengthened by 10-fold and inhibitory synapses by 2-fold.

The candidate pathway for an XOR gate uses two reciprocally connected switches that funnel their outputs into a common layer of neurons. In such XOR gates, two input layers drive layers of both excitatory and inhibitory interneurons, with the inhibitory neurons synapsing onto the excitatory interneurons of the other switch. Neurons are chosen to be maximally connected, as discussed in the previous paragraph. Input into one path, thus, drives the output layer while inhibiting propagation along the other path of the gate. A functional XOR gate involves ~ 220 neurons. To make such a candidate unit functional, excitatory synapses are strengthened by 10-fold and inhibitory synapses by 2-fold.

To find a candidate "flip-flop," two recurrently connected loops of excitatory neurons are identified, with the additional requirement that each loop synapses onto a layer of inhibitory interneurons that in turn contacts the other loop. With sufficiently strengthened synapses, it is possible to maintain an elevated firing rate in one loop while suppressing it in the other, even in the absence of external input. A functional flip-flop involves ~ 200 neurons. To make this candidate unit functional, excitatory synapses are strengthened by 10-fold and inhibitory synapses by 8-fold. For this network structure, fine-tuning that depends on the actual connectivity of the given network is necessary.

Analysis. To characterize sustained asynchronous network activity, we monitor individual neuronal membrane potentials, the population firing rate (the average of the firing rates across the network), average membrane potentials, and interspike intervals (ISIs). The asynchronous irregular network activity we seek has an approximately constant population firing rate and coefficients of variation (CVs) near 1 for the ISI distributions of the neurons. The ISI CV for a neuron is the ratio of the SD of the ISI distribution to its mean. CV values close to 0 indicate regular spiking patterns, values near 1 indicate irregular spiking, and values > 1 indicate, in our simulations, burstiness in the firing pattern, in which a neuron is

likely to fire several spikes in a short interval followed by a longer period of silence.

Signal propagation is analyzed by determining how similar the firing rates of the different layers along the propagation pathway are to the layer 0 input rate. To do this, we calculate the firing rates of pathway neurons in 5 ms bins by counting their spikes. We use the notation $r_i(t)$ for the average of the firing rates of all of the neurons in layer i for the time bin centered at time t , and we denote the time average of this firing rate by \bar{r}_i . As a measure of similarity of the firing of neurons in layer i to those in layer 0, we compute the correlation coefficient as follows:

$$C_i(\tau) = \frac{\langle (r_0(t) - \bar{r}_0)(r_i(t + \tau) - \bar{r}_i) \rangle_t}{\sqrt{\langle (r_0(t) - \bar{r}_0)^2 \rangle_t \langle (r_i(t + \tau) - \bar{r}_i)^2 \rangle_t}}, \quad (5)$$

where the brackets denote an average over t values. We use the activity of layer 0 as a reference, rather than layer 1, to distinguish signal transmission from propagation of fluctuations arising spontaneously in layer 1. Signal propagation up to layer i is characterized by reporting the maximum value of $C_i(\tau)$, which we call the similarity. The value of τ , where this maximum occurs, is the propagation delay for that layer. To test for the propagation of synfire chains, synchronous spikes are fed into layer 1. Spike number conservation across layers is used as an estimate for propagation success.

Results

We are interested in studying signal propagation within a network that sustains, by itself, ongoing activity characterized by irregular firing that is asynchronous across the network. This pattern of firing generates the noise needed for propagation. We begin by discussing the nature and properties of this irregularly and asynchronously firing network, and then address how signals propagate through it.

Sustained activity

Previous studies have shown that large networks of sparsely connected integrate-and-fire model neurons can sustain irregular asynchronous activity (van Vreeswijk and Sompolinsky, 1996, 1998; Amit and Brunel, 1997; Brunel, 2000; Mehring et al., 2003; Lerchner et al., 2004; Vogels et al., 2005). In most of these studies, synaptic connections between neurons were simulated by injecting an exponentially decaying pulse of current into the postsynaptic neuron whenever the presynaptic neuron fired an action potential. We call such models CUBA for "current based." An alternative is to transiently change the conductance of the postsynaptic neuron after a presynaptic action potential. We call these models COBA for "conductance based." Because biological synapses involve conductance changes and not current injection, it might appear that COBA models would inevitably be more realistic than CUBA models, but this is not necessarily the case. The individual neurons in large network simulations are typically represented by single compartments that must approximate the full dendritic structure of the real neurons being modeled. Depending on the placement of synapses on the dendritic tree, a CUBA or a COBA single-compartment model, or a mixture of the two, may be a more accurate representation. For this reason, we consider both CUBA and COBA models in our studies. Most of the results we show are for COBA models, but every effect reported was also reproduced in an analogous CUBA model.

Large, sparsely coupled networks display self-sustained irregular asynchronous firing of the type needed to support signal propagation if their parameters are adjusted appropriately. For our purposes, three conditions need to be met: sustained activity (for at least as long as it takes to examine signal propagation), relatively low firing rates, and ISI CVs near 1. To find such activity, we performed a parameter search of a 10,000 neuron COBA




Figure 1. Parameter search. Excitatory and inhibitory conductances refer to the parameters Δg_{ex} and Δg_{inh} converted to nanosiemens assuming a resting neuronal membrane conductance of 100 M Ω . The black square shows the parameter values used in all subsequent COBA simulations. **a**, Duration of network activity. Parameter pairs in which network activity was sustained over the length of the simulation (1000 ms) are colored in orange. Pairs leading to silent networks are shown in yellow and the same regions are denoted by a white mesh in **b** and **c**. **b**, Average network firing rates. Firing rates in configurations with sustained activity range from 8 to 200 Hz. **c**, Average CV of ISIs. CVs range from 0 (very regular) to 3 (very bursty) over the range in which activity was sustained.




Figure 2. Background activity in a COBA model. **a**, Spike raster for a sample set of 250 neurons over a simulated time of 400 ms. **b**, Average firing rate of the entire population and of a Poisson train. The black trace shows the rate computed from 0.1 ms bins, and the white trace shows the same activity computed using 5 ms bins. The top panel is computed from the network, and the bottom panel, for comparison, from an equivalent number of Poisson processes with a 5 ms refractory period producing spikes at the same rate as the network. **c**, Membrane currents of a randomly chosen neuron. Inhibitory currents are in dark gray, excitatory ones are in light gray, and the total synaptic current is shown in solid black. **d**, Membrane potential of a randomly chosen neuron. **e**, Distribution of firing rates of the network neurons. **f**, Distribution of ISIs of the network neurons. **g**, Distribution of CVs of ISIs of the network neurons. **h**, Distribution of average membrane potentials of the network neurons. **e–h**, The arrow marks the mean of the distribution. Pop., Population; Pot., potential.

model with 2% connectivity by systematically varying the strengths of its excitatory and inhibitory synapses (Fig. 1). Figure 1a indicates how long activity was sustained in the network after a brief period of initial stimulation, Figure 1b gives the average firing rate of the network neurons (for the period while the activity lasted in the case of transiently active networks), and Figure 1c shows the ISI CV averaged over all network neurons. The black square in these figures denotes a state that satisfies the requirements listed above (actually, the “sustained” activity in the COBA network lasts for as long as our simulations run, but it will not last indefinitely) with conductances that correspond, on average, to 1 mV EPSPs and -2.6 mV IPSPs within the active network. This state is used for the remaining figures in which a COBA model appears.

With the parameters described by the black square in Figure 1, the network stays active with an average firing rate of 9 Hz and an average ISI CV of 1.2 (Fig. 2e,g). Network activity for these parameters is asynchronous. Neither the raster plot of the spike times of 250 sample cells (Fig. 2a) nor the average population activity (Fig. 2b, top) shows obvious temporal structure, and autocorrelations also do not reveal any obvious firing patterns (data not shown). To quantify the asynchronous nature of the firing, we compared the population firing rate obtained from the network (Fig. 2b, top) with equivalent fire rates derived from a Poisson process generating spikes at the same rate (Fig. 2b, bottom). The Poisson spikes produce a smoother firing rate than the network, indicating that there is some temporal structure in the network activity, but the overall levels of the fluctuations are similar. For example, the variance of the firing rate measured in 0.1 ms bins for the network is 1.5 times that of the Poisson train.

Excitatory and inhibitory currents are balanced and tend to cancel each other (Fig. 2c), keeping the average membrane potential of the network neurons at -70 mV (Fig. 2h). The sample single-neuron membrane potential trace in Figure 2d shows that spiking is irregular but tends to come in bursts. This is further revealed by the distribution of ISIs (Fig. 2f), which has a peak at small intervals reflecting this bursting and a second local maximum at a larger interval indicating the typical inter-burst interval.

It is possible to obtain activity similar, but not identical, to that shown for the COBA model in Figure 2 using a CUBA model (Fig. 3a) with 0.25 mV EPSPs and




Figure 3. Differences in the background activity of the CUBA (*a*) and COBA (*b*) models. Distribution of ISIs in each model plotted on a semilog scale is shown (left). The arrow marks the mean of each distribution. Membrane potentials of a randomly chosen neuron in each model are shown (right). Note that the two models are operating in different parameter regimes (see Results).

–2.25 mV IPSPs. Although both the CUBA (Fig. 3*a*) and COBA (Fig. 3*b*) models shown have the same average firing rate and both display irregular asynchronous activity, they differ in two important respects. First, spiking in the CUBA model shows much less bursting than in the COBA model. This is revealed by the absence of both a short-interval peak and a long-interval interburst peak in the ISI distribution of the CUBA model (Fig. 3*a,b*, compare left panels) as well as being apparent in the sample membrane potential traces (Fig. 3*a,b*, right panels). Interspike intervals in the CUBA model approximately follow a Poisson distribution, except for small ISIs that are suppressed by the imposed refractory period. Another important distinction is that the CUBA model includes a constant current injected into all of the neurons of the network sufficient to depolarize them by 11 mV. This causes the average membrane potential in the active network to be –55 mV, whereas the average membrane potential for the COBA models is –70 mV. Without this additional input, the CUBA network does not sustain activity. In contrast, the COBA model sustains activity without any injected current. Note that the two network configurations shown differ significantly in the sizes of their postsynaptic potentials. It is possible to build a COBA model very similar to the CUBA model shown in Figure 3*a*, but only a COBA model can produce reasonable results in the high synaptic-strength regimen shown in Figure 3*b*.

The ability of the COBA model to generate and sustain asynchronous activity without current injection is attributable to the voltage dependence of the EPSP and IPSP amplitudes for conductance-based synapses (Kuhn et al., 2004; Kumar et al., 2005; Moreno-Bote and Parga, 2005). In the CUBA model, EPSP and IPSP amplitudes are voltage independent, but in the COBA model, driving-force effects cause EPSPs to shrink and IPSPs to grow with increased depolarization, or EPSPs to grow and IPSPs to shrink with increased hyperpolarization. This provides a stabilization mechanism that holds the membrane potential much steadier than it is in the CUBA model. The resulting buffering protects the network against excessively low or high firing rates and creates enough stability to sustain activity.

Signal propagation

We study signal propagation by randomly choosing a set of 33 neurons, called layer 1 or L_1 , and injecting an input signal into them (see Materials and Methods). For Figure 4, the input was a

square pulse of presynaptic activity from layer 0 (shown as the blue raster under the axis in Fig. 4*b–e*), large enough to significantly elevate firing in the L_1 neurons (Figs. 4*b–e*, red raster). We then examined the effect that this layer 1 firing had on other network neurons. We chose to monitor the neurons most likely to be affected by defining a set of layer 2 neurons that were connected to the layer 1 neurons by at least three synapses. We defined layer 3 neurons as a set connected to the layer 2 neurons by at least three synapses and proceeded in this manner to define six layers of neurons along the potential signal propagation pathway (see Materials and Methods). Figure 4*a* shows the structure of such a pathway, and the activities of the neurons in its different layers are indicated by the different colored rasters in Figure 4*b–e*.

If we make no modifications to the synapses or neurons along the candidate pathway in the network, firing-rate signals fail to propagate beyond layer 1 (Fig. 4*b*). The way the network is set up, only three synapses typically extend from layer 1 to a particular neuron in layer 2. In addition to the input from layer 1, each cell of layer 2 receives input from ~140 excitatory and 40 inhibitory neurons. Even accounting for the elevation in the firing rates of layer 1 neurons above the average activity of 10 Hz within the network, the signal from layer 1 represents a small minority of the EPSPs that a layer 2 neuron receives. With such a signal-to-noise ratio, it is not surprising that signals fail to propagate.

We tried several different strategies to promote propagation: depolarizing neurons along the pathway, increasing their response gain, or increasing the strengths of the excitatory synapses between pathway neurons. Depolarization of the pathway neurons leads to a general increase in their firing rates, but the signal fails to propagate beyond layer 1 (Fig. 4*c*). Increasing the gain of neurons, which is usually thought of as changing the slope of the input/output function to modulate intrinsic excitability, is, in the case of the integrate-and-fire neurons we use, equivalent to increasing the strengths of all of the synapses (excitatory and inhibitory) that they receive. This sensitizes the pathway neurons to all of their inputs, not just the signal-carrying ones. In this case, the average firing rate does not increase significantly, because the depolarizing and hyperpolarizing membrane currents remain balanced. Even with gain modulation of pathway neurons, the signal fails to propagate beyond layer 1 (Fig. 4*d*).

Propagation of the firing-rate signal through all six pathway




Figure 4. Signal propagation. *a*, Network diagram showing the layers of a candidate pathway. Input (blue) is fed into the network through strong synapses onto layer 1 neurons (red). In this and the following diagrams, layers 1–6 are indicated by the colors green, yellow, dark blue, orange, and light blue, respectively. The gray-filled circles denote nonpathway neurons of the network. For this figure, layer 0 activity consists of a 30 ms pulse of activity at ~ 180 Hz. *b*, In a network with uniform excitatory and inhibitory synaptic strengths and neuronal parameters, no propagation occurs. *c*, Depolarization of pathway neurons by 15 mV fails to induce propagation, although firing rates in all affected cells increase. *d*, Gain increase of pathway neurons. Because gain modulation maintains the excitatory/inhibitory balance, firing rates do not increase significantly, but activity fails to propagate. *e*, Strengthening of pathway synapses by ~ 10 -fold results in signal propagation.

layers without synchronization can be achieved by increasing the strengths of the synapses that carry the signal from one layer of pathway neurons to the next by ~ 10 -fold (Fig. 4*e*). Other network synapses not along the propagation pathway are not altered. We define the synapse enhancement factor, or synapse factor, for short, as the ratio of the strength of excitatory pathway synapses to excitatory synapses not along the pathway minus 1. In other words, for a synapse enhancement factor x , the strength of an excitatory pathway synapse is $1 + x$ times that of an excitatory nonpathway synapse. The result in Figure 4*e* indicates that the spontaneous activity within the network is sufficient to support signal propagation.

To further investigate the effects of strengthening pathway synapses, we fed a constant input signal from layer 0 into the network. Figure 5, *a* and *b*, shows the firing rates induced in layers 1–6 (different colored lines) of a COBA network by this constant input for two different layer 0 rates (*a*, 50 Hz; *b*, 170 Hz). The background firing rate of neurons not along the propagation pathway is also shown (lowest blue line). The induced layer 1 firing rate (red curve) is close to the layer 0 firing rate and is approximately independent of the synapse enhancement factor, because the synapses from layer 0 to layer 1 are held fixed. For low values of pathway synaptic strength (synapse enhancement factors not much greater than 0), the firing rates in layers 2–6 are unaffected by the layer 0 input, and they remain near background levels. As the strength of pathway synapses is increased, layers 2–6 increasingly respond to the input. Eventually, at the point in Figure 5, *a* and *b*, at which all of the different colored curves cross (other than the curve for the background rate), all of the layers respond to the input by firing at approximately the same rate. Such agreement in the firing rates indicates faithful rate propagation through the pathway layers of the network. We call the synaptic factor at which this crossing occurs the optimal synapse enhancement factor. If the synapse enhancement is further increased beyond the optimal value, spontaneous as well as evoked activity is transmitted along the pathway, which drowns out the signal.

We extracted optimal synapse enhancement factors for different layer 0 firing rates from sets of curves similar to those in Figure 5, *a* and *b*, for both COBA and CUBA models. For the COBA model, the curve of optimal synapse factor versus firing




Figure 5. Optimal synaptic enhancement. *a*, *b*, Average firing rates of layer 1 (red), 2 (green), 3 (yellow), 4 (dark blue), 5 (orange), and 6 (light blue) in response to a constant layer 0 rate of 50 Hz (*a*) or 170 Hz (*b*). The background rate of nonpathway neurons is indicated by the straight black line at the bottom. The ratio of the strength of pathway synapses to nonpathway synapses is 1 plus the synapse factor. The optimal synapse enhancement factor is indicated by the vertical dashed line. *c*, The optimal synapse enhancement factors in COBA and CUBA models for different layer 0 firing rates. The examples shown in *a* and *b* are filled in red. *d*, The probability of a postsynaptic spike within a 5 ms window of a presynaptic single spike (open symbol) or a synchronous triplet of presynaptic spikes (solid symbols), plotted as a function of the synapse factor for both COBA and CUBA models.

rate has a long plateau at a value of 12, stretching from a layer 0 firing rate of 30 Hz to 160 Hz (Fig. 5*c*, solid symbols). This means that a wide range of firing rates can be transmitted across a pathway using a fixed factor for the enhancement of pathway synapses. When the layer 0 rate is too low, the signal gets lost in the 9 Hz background firing of the network, and network neurons cannot fire faster than 200 Hz because of the imposed refractory period. Thus, transmission occurs pretty much over the entire range in which signal propagation could have been expected, using a single value of synapse enhancement. The plateau value of the optimal enhancement was ~ 13 -fold stronger than the nonpathway network synapses. In the active network, these synapses create average EPSPs of 8 mV. In the CUBA model (Fig. 5*c*, open symbols), no plateau exists, indicating that the range of firing




Figure 6. A study of synfire waves in response to a synchronous signal fed into layer 1. *a*, No propagation occurs in the unaltered network. *b*, For synapse factor 9, the signal propagates up to layer 3 but then gets washed out by spontaneous activity. *c*, For synapse factor 12, the signal evokes a response in all layers, but packet length increases with every layer because of secondary spiking. *d*, When the pathway neurons are gain modulated to decrease their responsiveness by 10-fold and the synapse factor is 30, the difference between pathway and nonpathway synapses can be increased enough to propagate a synfire wave through six layers. Mod + x, Gain modulated and synapses strengthened. *e*, Average number of active cells in the first propagated wave front of a synfire event, plotted for the different cases described in *a–d*. *f*, Probability of evoking a secondary spike within 5 ms of the end of the refractory period of the first spike, plotted for presynaptic singlets (solid) and triplets (open). Each point is calculated from 5000 stimuli delivered to randomly chosen neurons. *g*, Secondary spikes are evoked when the postsynaptic conductance is large. *h*, Rise-time delays for on (solid) and off (open) signals consisting of pulses between 0 and 100 Hz at synapse factor 12. Higher layers have slower rise times, because they are affected by the rise times of their precursors as well.

rates that can be transmitted in the CUBA model is much more restricted than in the COBA model for a fixed enhancement of pathway synapses.

As an alternative measure of synaptic strength, we determined the average probability that a spike from a neuron in the propagation pathway evoked a spike within 5 ms in one of its postsynaptic targets in the next layer of the pathway (Fig. 5*d*). Over the range of synapse enhancement factors considered, this increased from near 0 to 0.4. Synchronous triplets of presynaptic inputs increase the probability of postsynaptic spiking by approximately threefold, similar to a linear summation of three separate events.

In addition to firing-rate propagation, we tested for synfire propagation within the network by evoking synchronous events in layer 0, with submillisecond precision. In the unaltered network, a synfire wave failed to propagate beyond layer 1 (Fig. 6*a*), as expected from earlier results, in which many more than three “spike units” were needed for successful synfire propagation (Diesmann et al., 1999). Strengthening pathway synapses by a factor of 9 produces pulses of activity in the first few layers, but the synfire wave ultimately gets dispersed by the background activity of the network (Fig. 6*b*). At a synaptic enhancement of 12, propagation occurs but a different problem arises (Fig. 6*c*). The initial tight synfire packet expands as it propagates through the layers. This phenomenon is somewhat reminiscent of what has been described as synfire explosions (Mehring et al., 2003), a surge of high activity that holds the entire network hostage and subsequently silences it completely. Explosions like these do not occur in our network, probably because of its sparse and random connectivity. The activity stays confined to the pathway, but the duration of the synfire packet increases across the layers.

It is possible to achieve synfire propagation in the network by combining a negative gain modulation along the pathway (i.e., weakening all synapses onto pathway neurons) with a strengthening of pathway synapses (Fig. 6*d*). This decreases the effect of the background activity along the pathway but keeps the rest of the network from becoming synchronized with the synfire pulse by maintaining high noise levels for nonpathway neurons. Figure 6*e* summarizes the results of Figure 6*a–d* in terms of the percentage of spiking neurons within each layer during the propagation (or nonpropagation) of the leading edge of the synfire wave.

The problem of lengthening synfire packets seen in Figure 6*c* is a general property of network propagation. The same kind of lengthening can be seen for a firing-rate pulse in Figure 4*e*. The probability to evoke a postsynaptic spike with a synchronous triplet increases linearly with the number of synchronous presynaptic spikes (Fig. 5*d*), but the total number of evoked spikes grows faster because second and third postsynaptic spikes are evoked (Fig. 6*f,g*). With sufficiently strong synapses, a single synchronous volley evokes a train of two or three synchronous spikes in the next layer, which in turn evoke an even longer train in subsequent layers (Fig. 6*c*). Packet lengthening limits the frequency response for firing-rate propagation (but to realistic levels) (Fig. 7*g*), but it is fatal to synfire propagation, because this requires the faithful reproduction of a signal containing substantial high-frequency components. The same synaptic enhancement factor that allows the leading edge of the synfire packet to propagate causes the trailing edge to lag more and more from one layer to the next. Figure 6*h* shows the rise time delays for on and off signals from 0 to 100 Hz and vice versa as solid and open symbols, respectively. The longer rise time delay for off-signals is attributable to this packet-lengthening effect.

We further tested firing-rate propagation by generating a time-varying firing rate in layer 0 from filtered white noise (see Materials and Methods) (Fig. 7*a*). Figure 7*b* shows average firing rates in three of the six pathway layers. The firing rates in different layers clearly track each other, but occasionally the network transmits “ghost signals” [as for example at $t \approx 500$ ms (Fig. 7*b*)] that arise spontaneously. Figure 7*c* shows that the induced fluctuations along the pathway do not significantly modify the overall firing rate of the entire network.

To quantify the quality of signal transmission, we calculated the correlation coefficients between the firing rate of layer 0 and the average firing rates of the neurons in the six network pathway layers (see Materials and Methods). The resulting six correlation coefficients depend on the time delay at which the firing rate correlation is computed, peaking at a value that gives the time delay for propagation up to each layer (Fig. 7*d*, different colored lines). These delays vary from 0 in layer 1 to 20 ms in layer 6. In Figure 7*e–h*, we plot the value of the correlation coefficient at its maximum, which we call the similarity. Figure 7, *e* and *f*, reiter-




Figure 7. Transmission of time-varying signals in a COBA network. *a*, Raster of the propagation of a randomly varying layer 0 firing rate through all six layers. *b*, Average firing rates, calculated in 5 ms bins, for layers 1, 3, and 6 responding to a randomly varying layer 0 rate. *c*, Average network firing rates (black trace, 0.1 ms bins; white trace, 5 ms bins) are relatively unaffected by the propagating fluctuations. *d*, Correlations of layer 1 (red), 2 (green), 3 (yellow), 4 (dark blue), 5 (orange), and 6 (light blue) firing rates with the layer 0 rate computed at various time delays. *e*, Similarity values of layer 1–6 firing rates with the layer 0 rate as a function of the synaptic enhancement factor in the COBA model. Optimal transmission occurs at a synapse factor of 12, corresponding to a 13-fold increase in synaptic strength. *f*, Same as *e*, but for the CUBA model. *g*, Similarity values for layer 0 rates oscillating sinusoidally at different frequencies. *h*, Similarity values when different numbers of signals propagate through the network along 10 different pathways.

ates the point made in Figure 5 about optimal synapse enhancements and show that the tuning of this factor is most critical for propagation through the later layers of the pathway in both COBA (Fig. 7*e*) and CUBA (Fig. 7*f*) models. In the COBA model with optimally strengthened synapses, the firing rate in layer 6 matches the layer 0 rate with a similarity value of 60%. Similarity values are somewhat smaller in the CUBA model because of the plateau in the optimal synaptic enhancement curve for the COBA model in Figure 5*c*.

We next tested transmission using sinusoidally varying layer 0 firing rates and found that transmission fidelity depends on signal frequency (Fig. 7*g*). Transmission is most accurate at ~5 Hz, and accuracy falls off rapidly beyond ~20 Hz. The reason for this can be seen in the raster of Figure 4*e* (see also Fig. 6*c,h*). The network is a poor temporal edge detector. As discussed above, the duration of the response to a square wave input grows as the signal propagates through the layers, and it is significantly lengthened by layer 6. Figure 6*h* shows that it takes each layer ~15 ms to adjust the layer 6 firing rate from background activity to a steady state of 100 Hz, and even longer to return to the background level. This translates into a maximum change of ~6 Hz/ms. Once the signal changes faster than this, responses start to wash into each other and the signal gets lost. The peak at 5 Hz occurs because oscillations at this frequency best mask spontaneously arising

ghost signals. The frequency maximum and fall-off seen in Figure 7*g* match quite well with frequency responses measured in cortical neurons (Movshon et al., 1978; Orban et al., 1985; Hawken et al., 1996).

With 10,000 network neurons and ~200 pathway cells, it is possible to find several different signaling pathways within a network and feed different signals into them simultaneously. Figure 7*h* shows similarity values for the six layers along from 1 to 10 pathways. Although transmission quality drops with multiple pathways because of interference between them, it is still possible to transmit signals at levels greater than the noise along multiple pathways. With this many pathways, the probabilistic nature of network connectivity does not allow us to find 33 neurons that satisfy all of the conditions for layer 6. If we use the average similarity value in layer 5, rather than layer 6, as a criterion for transmission quality to compensate for this, we see that the accuracy of the reproduced spike patterns in layer 5 falls below 50% for four simultaneously active pathways. With ~200 neurons involved in each pathway, this corresponds to using 8% of the network neurons for signal propagation.

Signal processing

Signal propagation is a prerequisite for signal processing, and we now show that with signal propagation accomplished, the net-




Figure 8. Processing units constructed by synaptically tuning existing network subcircuits. *a*, NOT gate. *b*, Switch. *c*, XOR gate. *d*, Flip-flop. The left side of each subplot shows the layout of each circuit, with inhibitory neurons drawn as hexagons. On the right, the average firing rates of input and output layers are plotted along with the Boolean interpretation above the traces. In *d*, external input to the different loops is indicated by color-coded bars along the *x*-axis.

work we are studying can also perform computations. To do this, we identify candidate processing circuits within the existing architecture of the network (by procedures discussed in Materials and Methods), just as we identified candidate pathways previously. Also, as before, we then strengthen synapses within the identified circuit to turn a candidate processing unit into a functioning one. A common feature of all four circuits shown in Figure 8, already discussed for the excitatory pathway in Figure 7*h*, is their relatively sluggish responses, amplified further by the longer inhibitory synaptic time constant.

Figure 8*a* illustrates the structure for a candidate NOT gate, in which a set of inhibitory interneurons controls an upstream layer, so that driving the inhibitory layer silences the otherwise active output layer. In Boolean terms, this represents an output of 1 for input 0 and an output of 0 for input 1. Figure 8*b* shows an extension of the strategy used to trace out the NOT gate, to create a switch that controls propagation along an excitatory pathway. Driving the inhibitory cell population of the circuit impedes any signal propagation along the excitatory pathway. Although the example shown is an on/off switch, the same circuit can be used to modulate a propagating signal in an analog manner by varying the firing rate along the inhibitory pathway (data not shown).

To create an XOR gate (an exclusive OR gate that propagates signals when either one of two pathways is active but not both), two intertwined switches synapsing onto the same output layer are identified. When all synapses are strengthened sufficiently,

the output layer fires above 60 Hz for single inputs and well below 40 Hz for simultaneous inputs along both paths. Translated into Boolean terms, this represents the output of an XOR gate (Fig. 8*c*).

Figure 8*d* shows the result of strengthening synapses within a candidate flip-flop circuit as a basic memory unit. A flip-flop must maintain high firing rates in one of two recurrent loops, even in the absence of an external signal, while suppressing activity in the other through a set of inhibitory interneurons. When one of the two loops is driven by an external source, the loop sustains its firing rate even after the input is shut off. The flip-flop can reverse its state when the other loop is stimulated. Such activity is seen in Figure 8*d*. Although the circuit can function as a flip-flop, it is not perfect. Because of fluctuations in the background activity, the inhibitory activity from one loop is sometimes insufficient to silence the other. Such a failure is seen \sim 800 ms into the simulation of Figure 8*d*, at which the flip-flop spontaneously changes its state, thus failing to maintain a memory. Elevated and asynchronous sustained firing rates in such units are possible only because of sufficient background activity, but this activity can cause failures too.

Discussion

We studied signal propagation in two different types of network models, COBA and CUBA. For the CUBA network, we used irregular asynchronous activity as observed in previous work, and for the COBA network, we chose parameters that also created irregular asynchronous but, more importantly, self-sustained activity. In such a network configuration, the size of postsynaptic events is approximately one magnitude larger than in the CUBA network, the average membrane potentials are lower, and the resulting activity has a more burst-like character. In addition, we tested some of our results in the low conductance regimen of the COBA network, with similar outcomes. Both models provide an internal, nonadjustable source of noise sufficiently large to prevent synchronization but not large enough to destroy signals and, thus, support rate-mode signal propagation. Furthermore, interference between the propagating signal and the background activity did not prove problematic. Signal transmission in the COBA model is more accurate, in the sense that a wider range of firing rates can be transmitted across the layers of the signaling pathway without parameter adjustment. This is because of the presence of conductance-based synapses, rather than the fact that the CUBA and COBA models operate in different parameter regimens.

With sufficiently strengthened pathway synapses, a rate-coded signal can travel through at least six synaptic stages with a transmission delay of \sim 20 ms. A 13-fold increase in synapse strength was needed to optimally transmit signals through a network. This corresponds to evoked EPSPs, within the active network, of \sim 8 mV, which is relatively large but not unheard of (Song et al., 2005). In the 10,000 neuron, randomly connected networks we studied, propagation involved only three pathway synapses. One way to get more synapses involved in the signal propagation and, thus, to lessen their strength, is to use a structured architecture, more like the real cortex, with columns and targeted branching axons. This is a topic of ongoing research.

The optimal synapse enhancement depends on the nature of the transmitted signal, especially in the CUBA model. The critical factors seem to be the ratio of silent to active periods and the distribution of firing rates within the active signal. A signal with long silences or low firing rates requires a smaller optimal synapse enhancement than one with large amounts of high-frequency

firing. The poorer performance of the CUBA model may be attributable to the fact that the membrane potential for CUBA neurons hovers close to the threshold, which increases the propagation of ghost signals, as well as the decreased membrane time constant in the COBA model.

Signal propagation in the networks studied falls off as a function of frequency much as it does in the real cortex (Movshon et al., 1978; Orban et al., 1985; Hawken et al., 1996). The failure of the network to transmit high-frequency signals is the root of its problems in supporting synfire waves. Although the probability to evoke a single postsynaptic response to an input rises linearly with the number of presynaptic spikes, the probability of evoking a second spike increases as well. Thus, synapses sufficiently strong to propagate the leading edge of a synfire wave create a wake of secondary spikes. Achieving synfire propagation requires a fine balance between noise and synaptic strength. For synfire chain propagation to succeed in the network we studied, gain modulation had to be used to decrease the effect of noise along the pathway. The parameter regimen over which synfire propagation is possible is rather small.

A useful way to think about propagation is as an avalanche in which each presynaptic neuron triggers activity in a postsynaptic neuron with a certain probability (Harris, 1963; Zapperi et al., 1995; de Carvalho and Prado, 2000; Beggs and Plenz, 2003). A critical cascade occurs when, on average, each presynaptic neuron fires one of its postsynaptic targets. In the pathways we consider, the average number of synaptic targets of pathway neurons is three, so we would expect the probability that a presynaptic spike generates a postsynaptic spike in any one of these targets to be 0.33 at criticality. Figure 5*d* shows that the spiking probability is ~ 0.25 at the optimal synaptic strength. This is somewhat lower than the critical value; even synchronous stimulation with three spikes evokes a postsynaptic spike only 80% of the times. The reason for this discrepancy is the probability of evoking a secondary spike (Fig. 6*fg*). The added probabilities for primary and secondary spikes are ~ 0.3 , close to the critical value. Although this does not take into consideration the temporal slurring that must occur as a result of delivering 20% of the stimulus 5 ms late, it is reassuring that these numbers are fairly close to criticality, guaranteeing spike number conservation.

The existence of multiple pathways in a network introduces the possibility of gating interactions between them. Finding candidate pathways makes it possible to create logic gates, switches, and memory units by strengthening selective synapses. Noise was required for the logic gates to function properly, but it also caused them to fail sometimes. Remarkably, all of these circuits could be found within the random network, so only synaptic enhancement was required to make them functional. We did this strengthening by hand, but it would be interesting to investigate whether various synaptic plasticity mechanisms could do this autonomously.

The problem we address, how to get signals over a wide dynamic range to propagate reliably across neural circuits, is an important element for understanding any sort of brain function. We showed that networks based on our understanding of background activity provide a sufficient source of noise to support signal propagation and that useful computations can be performed by interacting pathways found within such networks.

References

Abeles M (1991) *Corticonics: neural circuits of the cerebral cortex*. Cambridge, UK: Cambridge UP.

Aertsen A, Diesmann M, Gewaltig MO (1996) Propagation of synchronous

spiking activity in feedforward neural networks. *J Physiol (Paris)* 90:243–247.

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.

Anderson JS, Lampl I, Gillespie DC, Ferster D (2000) The contribution of noise to contrast invariance of orientation tuning in cat visual cortex. *Science* 290:1968–1972.

Aviel Y, Mehring C, Abeles M, Horn D (2003) On embedding synfire chains in a balanced network. *Neural Comput* 15:1321–1340.

Beggs JM, Plenz D (2003) Neuronal avalanches in neocortical circuits. *J Neurosci* 23:11167–11177.

Brunel N (2000) Dynamics of networks of randomly connected excitatory and inhibitory spiking neurons. *J Physiol (Paris)* 94:445–463.

Burns BD, Webb AC (1976) The spontaneous activity of neurones in the cat's visual cortex. *Proc R Soc Lond B Biol Sci* 194:211–223.

Dean AF (1981) The variability of discharge of simple cells in the cat striat cortex. *Exp Brain Res* 44:437–440.

de Carvalho JX, Prado CP (2000) Self-organized criticality in the Olami-Feder-Christensen model. *Phys Rev Lett* 84:4006–4009.

Diesmann M, Gewaltig MO, Aertsen A (1999) Stable propagation of synchronous spiking in cortical neural networks. *Nature* 402:529–533.

Harris TE (1963) *The theory of branching processes*. Berlin: Springer.

Hawken MJ, Shapley RM, Groszof DH (1996) Temporal-frequency selectivity in monkey visual cortex. *Vis Neurosci* 13:477–492.

Holt GR, Softky WR, Koch C, Douglas RJ (1996) Comparison of discharge variability in vitro and in vivo in cat visual cortex neurons. *J Neurophysiol* 75:1806–1814.

Kuhn A, Aertsen A, Rotter S (2004) Neuronal integration of synaptic input in the fluctuation-driven regime. *J Neurosci* 24:2345–2356.

Kumar A, Schrader S, Rotter S, Aertsen A (2005) Dynamics of random networks of spiking neurons with conductance-based synapses. Abstract for Computational and Systems Neuroscience (Cosyne), 153. Salt Lake City, UT.

Lerchner A, Ahmadi M, Hertz J (2004) High-conductance states in a mean-field cortical network model. *Neurocomput* 58–60:935–940.

Litvak V, Sompolinsky H, Segev I, Abeles M (2003) On the transmission of rate code in long feedforward networks with excitatory-inhibitory balance. *J Neurosci* 23:3006–3015.

Mehring C, Hehl U, Kubo M, Diesmann M, Aertsen A (2003) Activity dynamics and propagation of synchronous spiking in locally connected random networks. *Biol Cybern* 88:395–408.

Moreno-Bote R, Parga N. 2005. Membrane potential and response properties of populations of cortical neurons in the high conductance state. *Phys Rev Lett* 94:088103.

Movshon JA, Thompson ID, Tolhurst DJ (1978) Spatial and temporal contrast sensitivity of neurones in areas 17 and 18 of the cat's visual cortex. *J Physiol (Lond)* 283:101–120.

Orban GA, Hoffmann KP, Duysens J (1985) Velocity selectivity in the cat visual system. I. Responses of LGN cells to moving bar stimuli: a comparison with cortical areas 17 and 18. *J Neurophysiol* 54:1026–1049.

Reyes AD (2003) Synchrony-dependent propagation of firing rate in iteratively constructed networks in vitro. *Nat Neurosci* 6:593–599.

Softky WR, Koch C (1993) The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci* 13:334–350.

Song S, Sjöström PJ, Reigl M, Nelson SB, Chklovskii DB (2005) Highly non-random features of synaptic connectivity in local cortical circuits. *PLoS Biol* 3:e68.

van Rossum MC, Turrigiano GG, Nelson SB (2002) Fast propagation of firing rates through layered networks of noisy neurons. *J Neurosci* 22:1956–1966.

van Vreeswijk C, Sompolinsky H (1996) Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* 274:1724–1726.

van Vreeswijk C, Sompolinsky H (1998) Chaotic balanced state in a model of cortical circuits. *Neural Comput* 10:1321–1371.

Vogels TP, Rajan K, Abbott LF (2005) Neural networks dynamics. *Annu Rev Neurosci* 28:357–376.

Zapperi S, Baekgaard Lauritsen K, Stanley HE (1995) Self-organized branching processes: mean-field theory for avalanches. *Phys Rev Lett* 75:4071–4074.