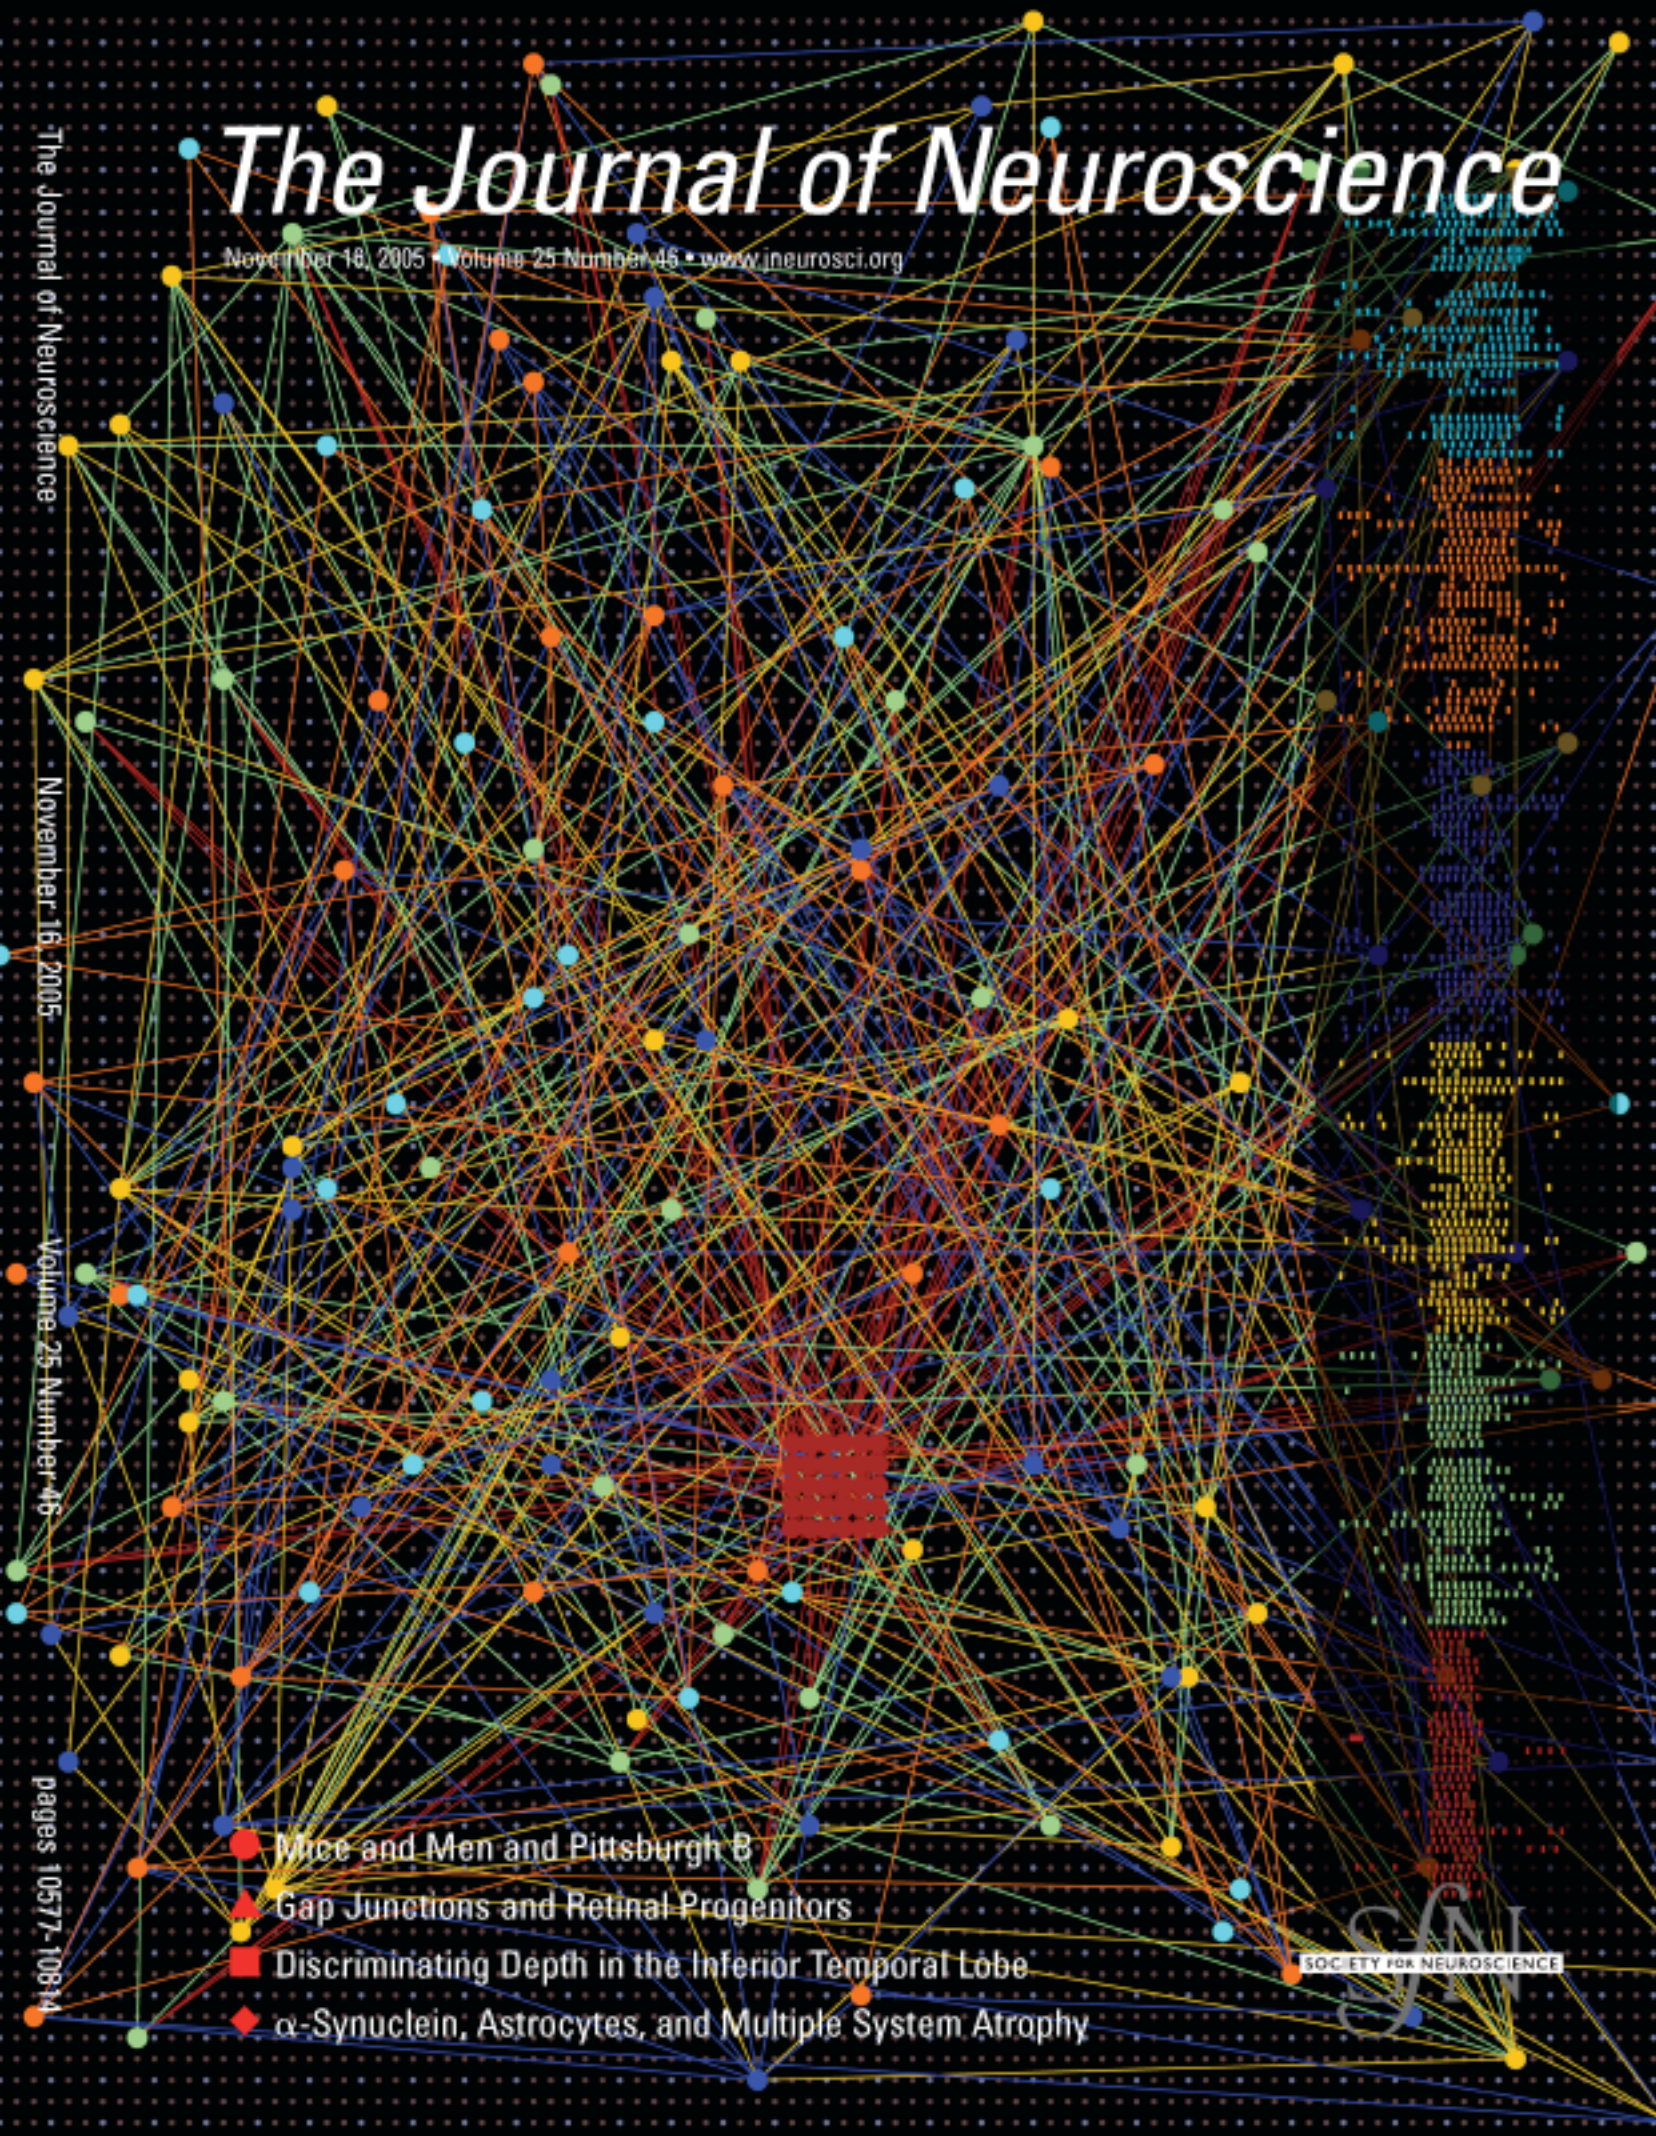
The Journal of Neuroscience
November 16, 2005 • Volume 25 Number 46 • www.jneurosci.org
The Journal of Neuroscience
November 16, 2005 • Volume 25 Number 46
pages 10577–10814
Mice and Men and Pittsburgh B
Gap Junctions and Retinal Progenitors
Discriminating Depth in the Inferior Temporal Lobe
α-Synuclein, Astrocytes, and Multiple System Atrophy
SFN
SOCIETY FOR NEUROSCIENCE

Behavioral/Systems/Cognitive

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 Goaillard 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

 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 MΩ.

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:

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

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 ~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 $i = 1, 2, \ldots, 6$. 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 ~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 $i - 2, i - 3, \ldots, 1$. 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

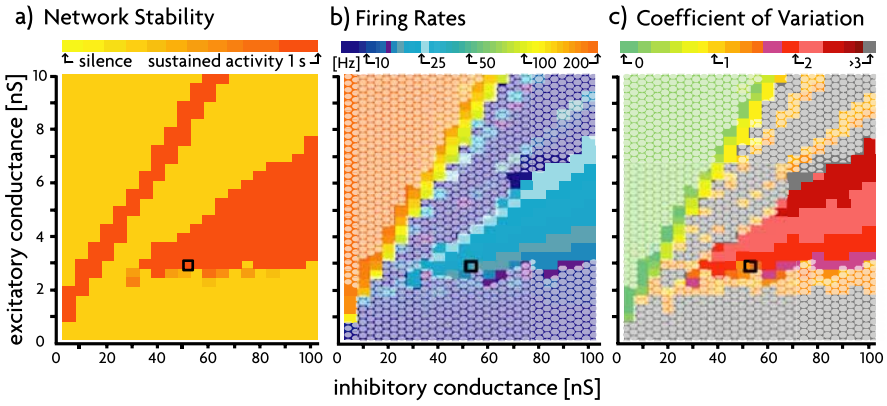

Figure 1. Parameter search. Excitatory and inhibitory conductances refer to the parameters Δg_{ex} and Δg_{inh} converted to nanosiemens assuming a resting neuronal membrane conductance of 100 MΩ. The black square shows the parameter values used in all subsequent COBA simulations. ***a***, Duration of 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.

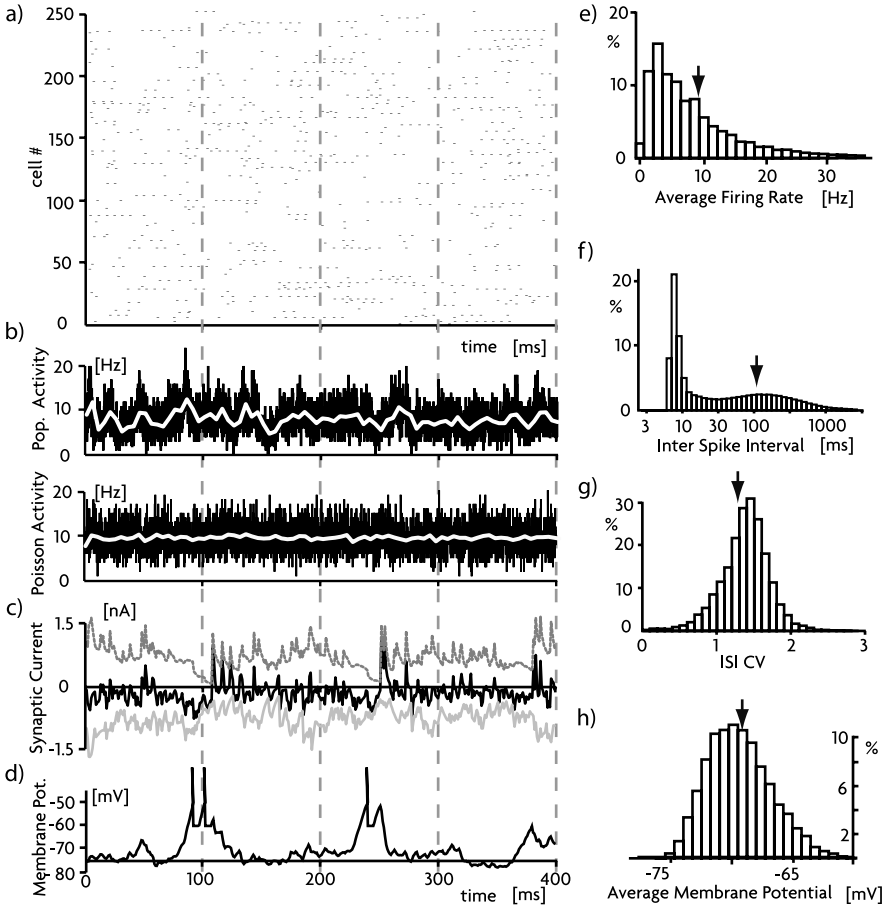

Figure 2. Background activity in a COBA model. ***a***, Spike raster for a sample set of 250 neurons over a simulated time of 400 ms. ***b***, Average firing rate of the entire population and of a Poisson train. The black trace shows the rate computed from 0.1 ms bins, and the white trace shows the same 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 1*a* indicates how long activity was sustained in the network after a brief period of initial stimulation, Figure 1*b* 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 1*c* 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. 2*e*,*g*). Network activity for these parameters is asynchronous. Neither the raster plot of the spike times of 250 sample cells (Fig. 2*a*) nor the average population activity (Fig. 2*b*, 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. 2*b*, top) with equivalent fire rates derived from a Poisson process generating spikes at the same rate (Fig. 2*b*, 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. 2*c*), keeping the average membrane potential of the network neurons at −70 mV (Fig. 2*h*). The sample single-neuron membrane potential trace in Figure 2*d* shows that spiking is irregular but tends to come in bursts. This is further revealed by the distribution of ISIs (Fig. 2*f*), which has a peak at small intervals reflecting this bursting and a second local maximum at a larger interval indicating the typical interburst 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. 3*a*) with 0.25 mV EPSPs and

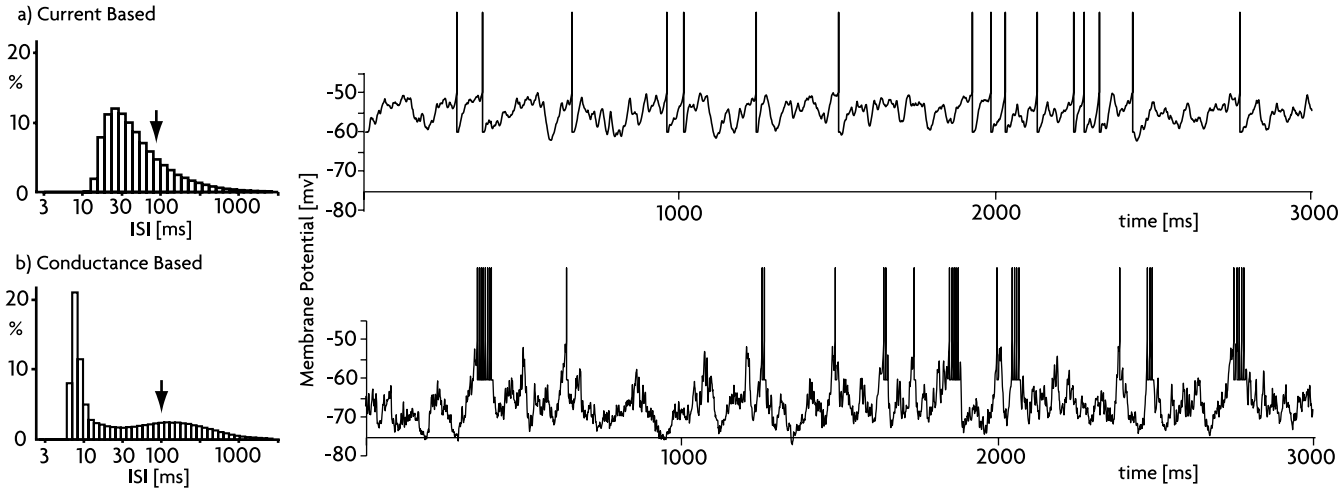

Figure 3. Differences in the background activity of the CUBA (***a***) and COBA (***b***) models. Distribution of ISIs in each model plotted on a semilog scale is shown (left). The arrow marks the mean of each distribution. Membrane potentials of a randomly chosen neuron in each model are shown (right).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

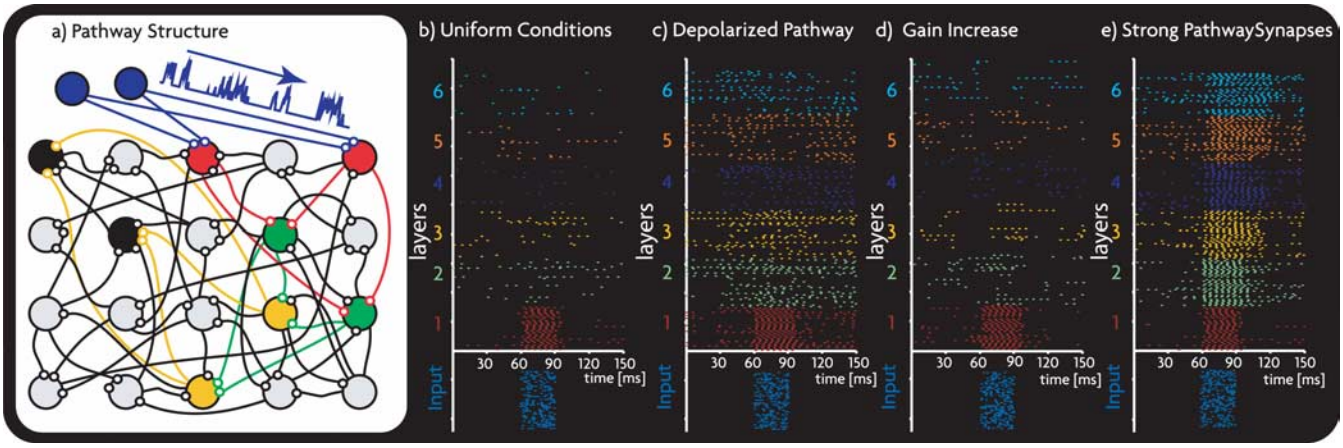

Figure 4. Signal propagation. ***a***, Network diagram showing the layers of a candidate pathway. Input (blue) is fed into the network through strong synapses onto layer 1 neurons (red). In this and the following diagrams, layers 1–6 are indicated by the colors green, yellow, dark blue, orange, and light 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

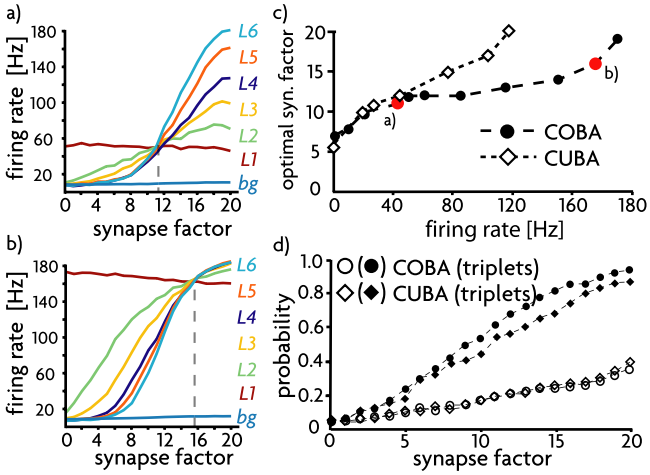

Figure 5. Optimal synaptic enhancement. ***a***, ***b***, Average firing rates of layer 1 (red), 2 (green), 3 (yellow), 4 (dark blue), 5 (orange), and 6 (light blue) in response to a constant layer 0 rate of 50 Hz (***a***) or 170 Hz (***b***). The background rate of nonpathway neurons is indicated by 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

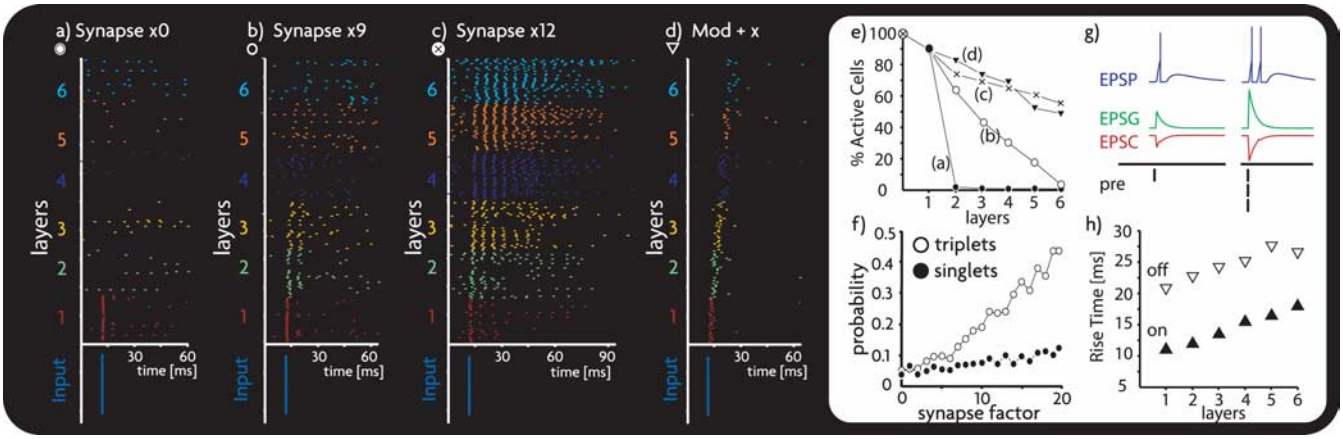

Figure 6. A study of synfire waves in response to a synchronous signal fed into layer 1. ***a***, No propagation occurs in the unaltered network. ***b***, For synapse factor 9, the signal propagates up to layer 3 but then gets washed out by spontaneous activity. ***c***, For synapse factor 12, the signal 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 pathways 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-

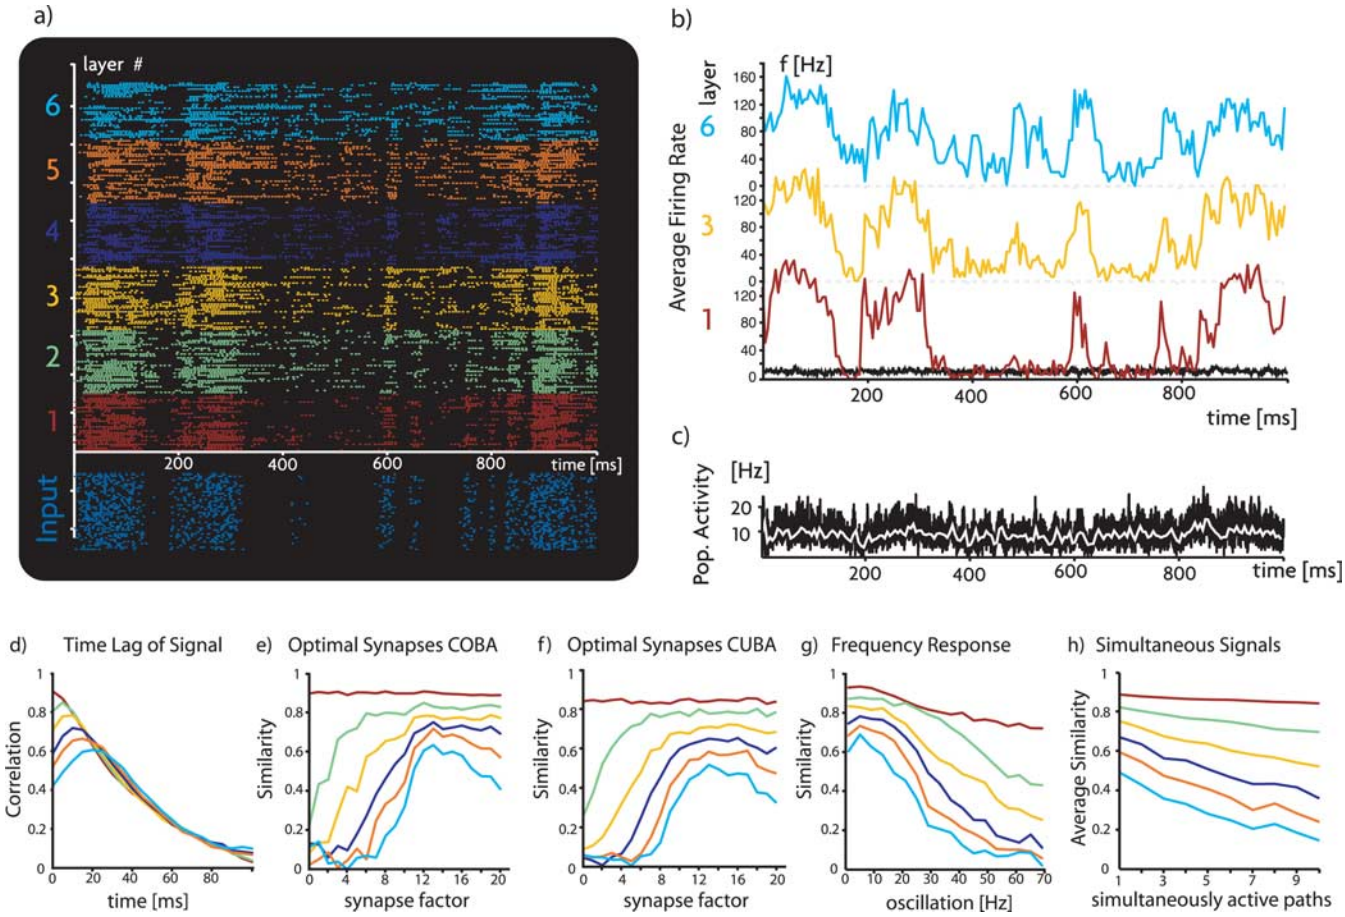

Figure 7. Transmission of time-varying signals in a COBA network. ***a***, Raster of the propagation of a randomly varying layer 0 firing rate through all six layers. ***b***, Average firing rates, calculated in 5 ms bins, for layers 1, 3, and 6 responding to a randomly varying layer 0 rate. ***c***, 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 than in the COBA 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-

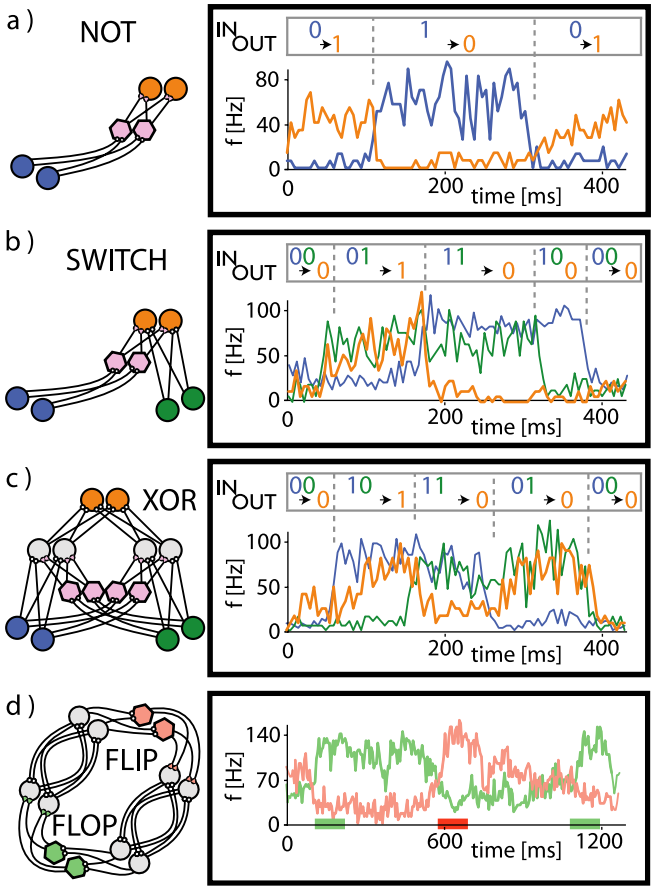

Figure 8. Processing units constructed by synaptically tuning existing network subcircuits. ***a***, NOT gate. ***b***, Switch. ***c***, XOR gate. ***d***, Flip-flop. The left side of each subplot shows the layout of each circuit, with inhibitory neurons drawn as hexagons. On the right, the average firing 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 entwined 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 ~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 ~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 ~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*f*,*g*). 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, Grosof 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 nonrandom 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.

nature neuroscience

Gating multiple signals through detailed balance of excitation and inhibition in spiking networks

Tim P Vogels[1,2] & L F Abbott[1]

Recent theoretical work has provided a basic understanding of signal propagation in networks of spiking neurons, but mechanisms for gating and controlling these signals have not been investigated previously. Here we introduce an idea for the gating of multiple signals in cortical networks that combines principles of signal propagation with aspects of balanced networks. Specifically, we studied networks in which incoming excitatory signals are normally cancelled by locally evoked inhibition, leaving the targeted layer unresponsive. Transmission can be gated 'on' by modulating excitatory and inhibitory gains to upset this detailed balance. We illustrate gating through detailed balance in large networks of integrate-and-fire neurons. We show successful gating of multiple signals and study failure modes that produce effects reminiscent of clinically observed pathologies. Provided that the individual signals are detectable, detailed balance has a large capacity for gating multiple signals.

Experimental observations[1,2] as well as theoretical arguments[3,4] suggest that excitation and inhibition are globally balanced in cortical circuits. In a globally balanced network, each neuron receives large but approximately equal amounts of excitation and inhibition that, on average, cancel each other. Spontaneous activity is driven by fluctuations in the total synaptic input, leading to asynchronous and irregular patterns of spiking[5–8]. Such networks have been used to study signal propagation and to determine conditions that support various signaling schemes[9–16]. However, neurons in these networks are typically only part of a single signaling pathway, and the transmitted signals cannot be gated or rerouted. Cognitive processing requires signal paths to change dynamically according to the information content of the signal and the processing demands of the receiver[17]. This requires precise control and gating of signal-carrying pathways.

We propose a mechanism for gating based on an extension of the concept of globally balanced networks to local cortical circuits in a form that we call 'detailed balance'. Detailed balance implies that, in addition to an overall or global balance, neurons receive equal amounts of excitation and inhibition on subsets of their synaptic inputs that correspond to specific signaling pathways. Activation of a balanced pathway produces little response in the excitatory neurons of the signal-receiving region, but responses can be gated 'on' by a command signal that disrupts the detailed balance. We analyze properties of the resulting gating mechanism and examine some of its failure modes. We show that the mechanism can gate the propagation of signals from multiple different sources to a single group of neurons, and we determine its capacity for gating large numbers of signals.

RESULTS

We explored the idea of detailed balance in a large network of roughly 20,000 integrate-and-fire neurons with both short- and long-range connectivity (**Fig. 1a,b**; Methods). With appropriately adjusted parameters, this network operates in a globally balanced manner, producing irregular, asynchronous activity in the absence of any time-varying or random external input[5–8]. The distribution of firing rates for the network is approximately exponential with an average firing rate per neuron of 8 Hz (**Fig. 1c**), the distribution of average membrane potentials is approximately gaussian with a mean of −60 mV (**Fig. 1d**), the distribution of interspike intervals (ISIs) is broad with peaks reflecting normal firing and bursting (**Fig. 1e**), and the distribution of coefficients of variation for the ISIs is centered at a value slightly greater than 1 (**Fig. 1f**). Average excitatory and inhibitory membrane currents are of approximately equal magnitude and the net current is near zero, indicating the globally balanced state of the network. This network model is intended to provide a sparse representation of the neurons over a fairly large area, not a full description of a single local circuit such as a cortical column.

Signal gating

To investigate signal gating within the network, we embed a two-layered pathway with 'sender' and 'receiver' subnetworks (**Fig. 1b**). These should be viewed as parts of distinct cortical regions. The connections from the sender region are directed to both excitatory and inhibitory neurons in the receiver region. Such targeting to inhibitory interneurons is consistent with data on the specific targeting of long-range excitatory projections to inhibitory interneurons[18]. To generate a signal, we drive the neurons in the sender subnetwork with a set of external Poisson spike trains at various rates. This causes neurons

[1]Center for Neurobiology and Behavior, Department of Physiology and Cellular Biophysics, Columbia University College of Physicians and Surgeons, New York, New York, USA. [2]Volen Center for Complex Systems, Department of Biology, Brandeis University, Waltham, Massachusetts, USA. Correspondence should be addressed to T.P.V. (timvogels@columbia.edu).

Received 9 December 2008; accepted 15 January 2009; published online 22 March 2009; doi:10.1038/nn.2276

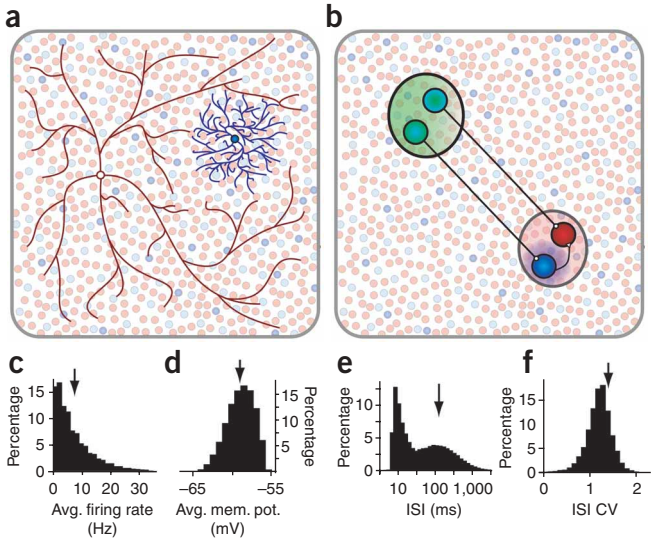

Figure 1 Network connectivity and properties. (**a**) All excitatory and 65% of the inhibitory neurons are connected randomly with a connection probability of 2% (red). The other 35% of the inhibitory neurons (blue) have local connectivity, targeting their nearest neighbors. (**b**) An embedded signal pathway is created by selecting a group of sender neurons (green) that target either excitatory or locally inhibitory neurons (red and blue, respectively, throughout the figures) in a signal-receiving region (red shading) of the network. (**c**–**f**) Asynchronous background activity in the network model. Distributions for network neurons of average (avg.) firing rates (**c**), average membrane potentials (mem. pot.) (**d**), ISIs plotted on a semilog scale (**e**) and coefficients of variation (CV) for those ISIs (**f**). Arrows indicate the means of the distributions.

in the sender area to fire in a manner that mimics the input signal (**Fig. 2a**). In the balanced state, excitation from the sender network, a simple oscillatory signal in the example of **Figure 2**, activates both excitatory and inhibitory subpopulations in the receiver region. The resulting inhibitory activity (blue, **Fig. 2d**) produces a local counter-signal that cancels the excitatory membrane currents (**Fig. 2c**) and generates only modest firing-rate fluctuations in the excitatory receiver neurons (red, **Fig. 2d**,**f**). The signal path is hence gated 'off' in the default (balanced) state of the signal-carrying pathway.

Signal propagation within this network can be gated 'on' in several ways, all of which involve unbalancing the excitatory and inhibitory pathways between the sender and receiver regions. The main requirement is a mechanism that differentially modulates the net excitatory and net inhibitory pathways from the sender to the receiver region[19]. A possible candidate is cholinergic modulation, which satisfies the basic requirements of cell-specific targeting[20,21] as well as relatively rapid response times[22–24]. Rather than modeling such modulation in detail, in the following examples detailed balance is disrupted by decreasing the gain or responsiveness of local inhibitory interneurons in the receiver region. This modulation, in keeping with the strong effects of attention seen for inhibitory neurons[25], corresponds to changing the input–output transfer function so that the same synaptic current generates a smaller response. Although the examples we show focus on modulation of inhibition, any combination of modulations that increases the ratio of excitatory to inhibitory transmission along the signaling pathway will perform similarly. To unbalance the signal in **Figure 2**, we reduced the response gain of the local inhibitory neurons in the receiver region to 15% of its control value. (We discuss more modest gain modulations below.)

Gain modulation that decreases the amplitude of the firing-rate modulations of inhibitory neurons in the receiver region (**Fig. 2d**, blue trace) reduces the inhibition of excitatory neurons in this region, leaving the bulk of the excitatory synaptic current uncanceled (**Fig. 2c**). This produces robust firing in the excitatory receiver neurons that is locked to the temporal pattern of the input signal (**Fig. 2b**,**d**,**f**). Overall network activity is relatively unaffected by these changes (**Fig. 2e**) because the modulated interneurons provide only a small fraction of the total inhibition to the network. Average subthreshold membrane potentials (excluding action potentials and their subsequent refractory periods) of the excitatory receiver neurons in the balanced

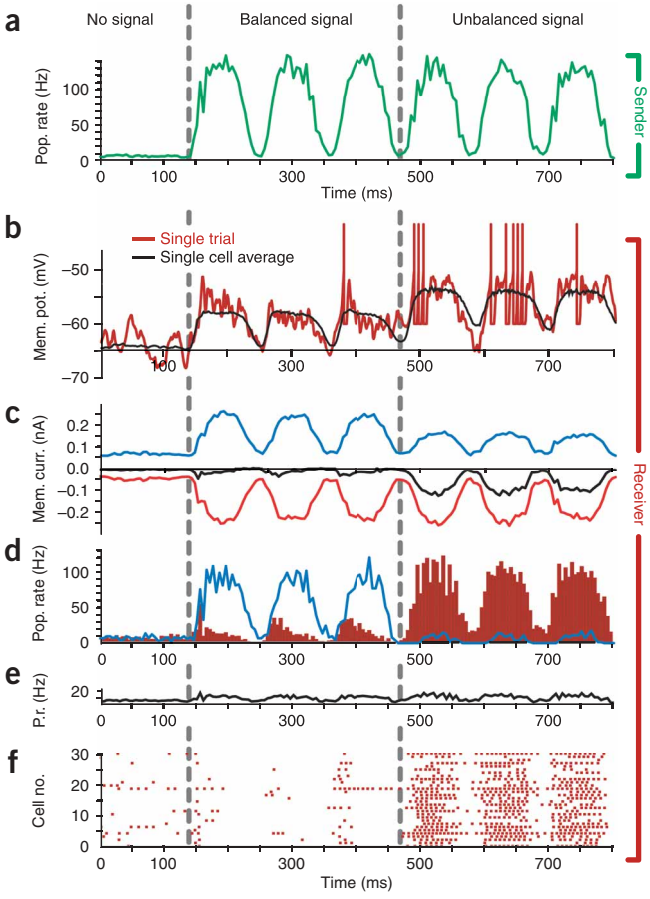

Figure 2 Detailed balance in a network. No signal (left column): all neurons fire at background rates. Balanced signal (middle): sender neurons fire in a correlated manner in response to oscillatory input and project the input firing pattern to both excitatory and inhibitory receiver neurons. Inhibitory receiver neurons reproduce the input pattern, preventing their excitatory neighbors from doing the same. Unbalanced signal (right): by decreasing the responsiveness of the inhibitory receiver neurons, the signal balance in the excitatory receiver neurons shifts in favor of excitation, and the signal is revealed in their firing pattern. All firing rates and averages are calculated in 5-ms bins. (**a**) Population (pop.) average firing rate of the sender neurons responding to a sinusoidally varying input. (**b**) Voltage trace of membrane potential (mem. pot.) in a randomly selected excitatory receiver neuron. Red trace, single trial. Black trace, average subthreshold membrane potential over 100 trials. (**c**) Average membrane currents (curr.) of the excitatory receiver neurons. Red and blue, excitatory and inhibitory currents, respectively; black, net current (including voltage-dependent leak and constant background currents). (**d**) Blue trace, average firing rate of the inhibitory receiver neurons. Red histogram, average firing rate of the excitatory receiver neurons. (**e**) Average population firing rate (p.r.) of the entire network. (**f**) Spike raster for 30 randomly chosen excitatory receiver neurons.

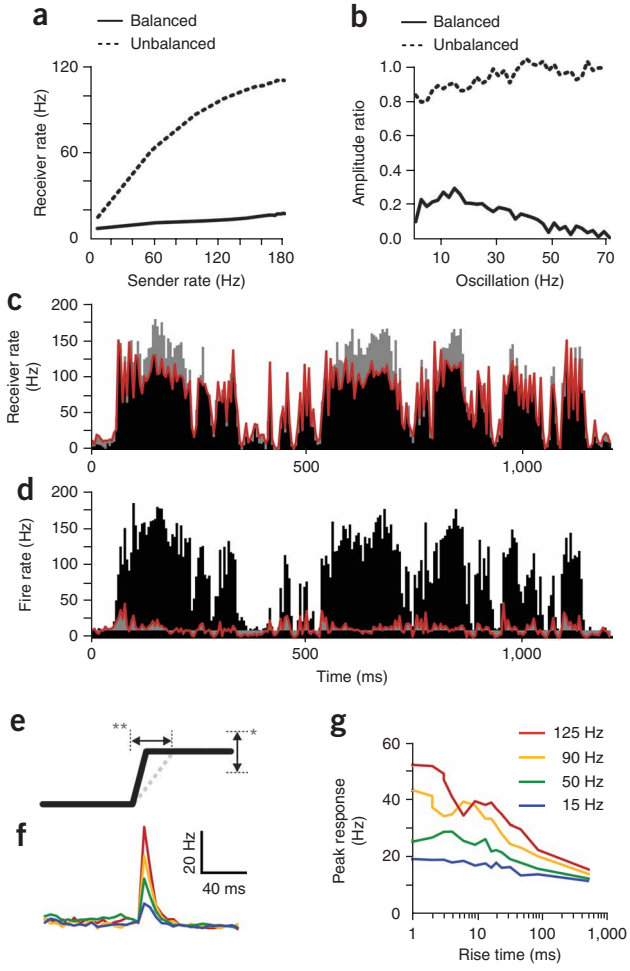

Figure 3 Response analysis. (**a**) Firing rates of the excitatory receiver neurons as a function of different constant sender firing rates, in the balanced (solid trace) and unbalanced (dashed trace) states. (**b**) Ratio of receiver to sender excitatory firing-rate oscillation amplitudes at different oscillation frequencies, in the balanced (solid trace) and unbalanced (dashed trace) states. (**c**,**d**) Response to a random, time-filtered signal in the unbalanced (**c**) and balanced (**d**) states. Red trace, average firing rate of the excitatory receiver neurons; black histogram, rates of the sender neurons. Deviations from the signal in **c** and from the average background rate in **d** are gray. (**e**) Schematic of an input step. Step size (*) and step duration (**) vary independently. (**f**) Average responses of the excitatory receiver neurons in the balanced state to instantaneous steps of different sizes. (**g**) Peak amplitude of the responses in these neurons to steps of different sizes (legend) and rise times (horizontal axis).

and unbalanced states (**Fig. 2b**, black trace) differ by only 4 mV, but this is sufficient to produce markedly different firing patterns.

Response properties

To further quantify the gating mechanism, we studied responses to different types of input (**Fig. 3a**,**b**). The firing rates of excitatory receiver neurons are relatively unaffected by constant input rates in the balanced gated-off state (**Fig. 3a**, solid trace) but rise sharply as a function of input rate when the pathway is gated 'on' (**Fig. 3a**, dashed trace). The rise begins to saturate at high rates because of the residual inhibition produced locally, even at low gain. Gating is also evident in the amplitudes of firing-rate fluctuations for excitatory neurons in the receiver region when the input signal is oscillatory (**Fig. 3b**). In addition, gating occurs when filtered white noise (with a 50-ms time constant[12]) is used as the input signal (**Fig. 3c**,**d**). In the gated-on state, this complex, irregular signal is transmitted with similarity values[12,14] (defined in the Methods) of ~90%, sufficient to propagate the signal across several layers[14]. In the balanced, gated-off state (**Fig. 3d**), the output of the excitatory receiver group is greatly decreased in amplitude, and the similarity between input and output is reduced to ~25%.

A close inspection of the responses in the balanced state (**Fig. 3d**) shows that detailed balance does not completely cancel signals when input firing rates change rapidly. Rapid changes in the signal can evoke a response in the excitatory receiver neurons before inhibition can balance excitation because of the time lag between the monosynaptic excitation and the canceling disynaptic inhibition acting on the receiver neurons. This effect would be even more marked if excitatory and inhibitory synapses were subject to different amounts or types of short-term plasticity. Such partial gating of transients allows large signals with rapid onsets to be propagated, which may induce upstream control circuits to activate gain modulation and open the gate. To further investigate this effect, we activated the sender neurons with step-like input rates of various step sizes and rise times (**Fig. 3e**). Short rise times produce fairly strong responses in excitatory receiver neurons that

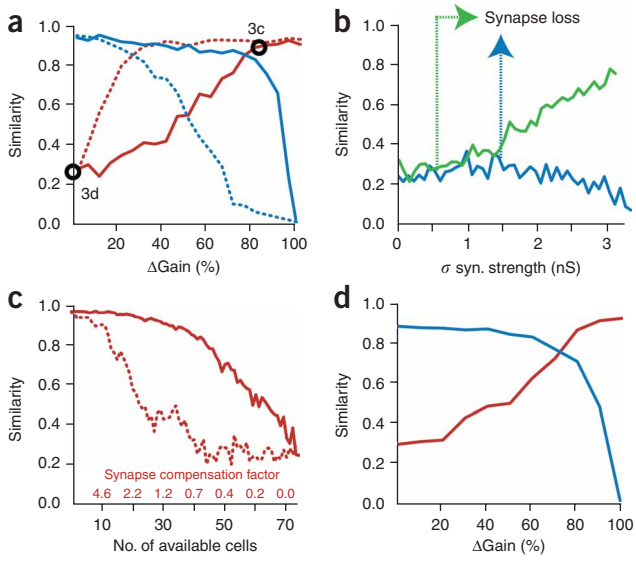

Figure 4 Gain properties. (**a**) Maximum values of the cross-correlations (termed 'similarity'; see Methods) between the sender region and the excitatory (red) and inhibitory (blue) receiver cells for different gains. Solid lines, symmetric gain reduction; dashed lines, asymmetric gain reduction, wherein only the gain of the excitatory synapses onto the inhibitory receiver cells is changed. Circles denote the parameters used in **Figure 3**. (**b**) Similarity values between excitatory receiver activity and the signal in the balanced (gated-off) state as a function of increasing the variability (s.d. σ) of the synaptic (syn.) strengths of the excitatory (green) and inhibitory (blue) pathways. The arrows mark the variability limits beyond which the tails of the strength distributions become rectified to zero. (**c**) Effect of reducing the number of inhibitory receiver neurons on the ability to gate signals off. Similarity values in the balanced state for decreasing numbers of inhibitory receiver cells, without and with synapse strength compensation (solid and dotted lines, respectively). (**d**) Operation of the gating mechanism with only 20 inhibitory receiver neurons by compensating synapse strength and shortened refractory times to allow for more rapid inhibitory firing. Similarity between the signal and the excitatory (red trace) and inhibitory (blue trace) receiver activity is plotted as a function of change in inhibitory gain.

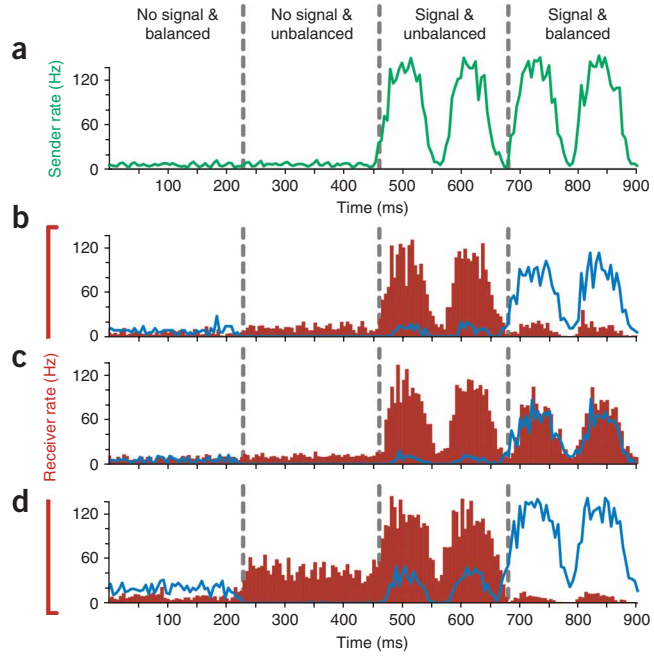

Figure 5 Network pathologies. (**a**) Average firing rate of the sender neurons without and with an oscillatory input. (**b**–**d**) Responses of excitatory (red histogram) and inhibitory (blue trace) receiver neurons with correct tuning (**b**); weakened local inhibition, leading to a gating deficit (**c**); or a hyperactive receiver region causing a response to the gating modulation (**d**). Conditions shown in the different columns are no signal and no modulation, no signal but gated on, signal on and gated on, signal on but gated off. Firing rates are calculated in 5-ms bins.

depend on the step size (**Fig. 3f,g**), even in the balanced gated-off state, but these diminish as the rise time of the step increases, illustrating the transient nature of this transmission.

The gain changes used to gate signals on have been fairly large, so we next examined different degrees and types of gain modulation in the inhibitory receiver neurons. Beginning with no gain change (ΔGain = 0), we decreased the responsiveness and thus the firing rate of the inhibitory receiver population. This causes the firing rate of the excitatory receiver neurons and its similarity to the sender signal in the gated-on state (**Fig. 4a**, solid red trace) to increase. At ΔGain $\sim$80%, the signal similarity of the activity of the inhibitory neurons goes rapidly to zero (**Fig. 4a**, solid blue trace), and the similarity of the excitatory receiver activity plateaus at $\sim$90%. Alternatively, it is possible to reach this same plateau level with a gain shift of only 30% (**Fig. 4a**, dotted traces) by modulating the inhibitory population asymmetrically, which means modifying only the responsiveness to excitatory inputs.

The gating mechanism is robust to many (but not all; see below) different perturbations of the network. To study the effect of synaptic variability, we computed the similarity of responses in the receiver region to the signal, when it is gated off, as a function of the variability in the strengths of the inhibitory (**Fig. 4b**, blue trace) or excitatory (**Fig. 4b**, green trace) sender synapses onto the excitatory receiver cells. Synaptic strength variance does not have a large effect in either case until the variance becomes high enough to force substantial numbers of synapses to zero strength (which occurs at a different point for excitatory and inhibitory synapses because of their different initial strengths), changing the mean synaptic strength. After this point, the high degree of variability in the excitatory synaptic strengths makes it difficult to shut the signal off (**Fig. 4b**).

We also tested the robustness of gating a signal off to the loss of its most critical components, the inhibitory neurons in the receiver region (**Fig. 4c**). The effect of decreasing the number of available interneurons (originally 73) is roughly linear (solid trace), until gating off fails completely when less than 40 cells are available. It is possible to partially rescue gating by upregulating the strengths of all remaining inhibitory synapses proportionately to the number of deleted neurons and thus deleted synapses (dotted trace). However, gating still fails when less than 25 inhibitory cells are available because such a small population of inhibitory neurons cannot fire a sufficient number of action potentials to provide balancing membrane currents, even if their synapses are strengthened to unrealistically high values. If the inhibitory receiver neurons are allowed to spike at rates as high as 600 Hz, it is possible to further rescue the balance mechanism and to successfully gate signals in as many as 600 excitatory cells with as few as 20 inhibitory neurons (**Fig. 4d**).

Pathologies

The basic requirement to achieve a state of detailed balance is local inhibition strong enough to cancel signals in the gated-off state. In addition, the gain modulation used to unbalance and gate 'on' a pathway must not have an excessively destabilizing effect on the global excitatory–inhibitory balance of the network. With this in mind, we examined more ways in which network gating can fail when tuning is relaxed by studying gated off and gated on states with no signal and in the presence of an oscillatory signal (**Fig. 5**). With proper tuning (**Fig. 5b**), the excitatory neurons of the receiver subnetwork respond robustly only when the signal is present and gating is on, although there is a weak transient response when the signal is present but gating is off.

We considered two different detuning conditions. First, we reduced the strength of all synapses from the local inhibitory neurons by 60% (**Fig. 5c**). This causes baseline firing rates in the absence of a signal to rise slightly, but the effect is not large because the bulk of the inhibition is not affected. Little change is seen in the response to the gain modulation alone, but activating the input shows that the gating mechanism has been compromised. Because of the weakened inhibition, excitatory inputs to the excitatory receiver neurons cannot be fully cancelled by local inhibition, and the signal can never be fully gated off.

We also detuned the detailed balance by increasing the strength of excitatory synapses within the receiver area by 60% (**Fig. 5d**). Excitatory synapses onto excitatory and inhibitory neurons were both modulated in this way, so a rough balance is still maintained within the receiver region. As in the case of reduced inhibition, enhanced excitation slightly elevates firing rates in the gated-off, no-signal condition. Activating gain modulation to open the gate causes a substantial elevation in the firing rate of excitatory receiver neurons, even when no signal is present. Thus, with altered excitation, the receiver neurons respond to the gating signal as if it were an input. This means that the network falsely transmits internally generated activity (the gating signal) as if it were an external signal. By contrast, in this condition signal responses in both the gated-on and gated-off states appear normal. We address the implications of these findings in the Discussion.

Multiple signals

One of the advantages of the gating mechanism we propose is that a group of receiver neurons can remain responsive to one set of incoming signals even while other sets are being cancelled by balancing inhibition. This gives the mechanism the capacity to gate multiple signals. As a first example, the network we have been considering is expanded so that it

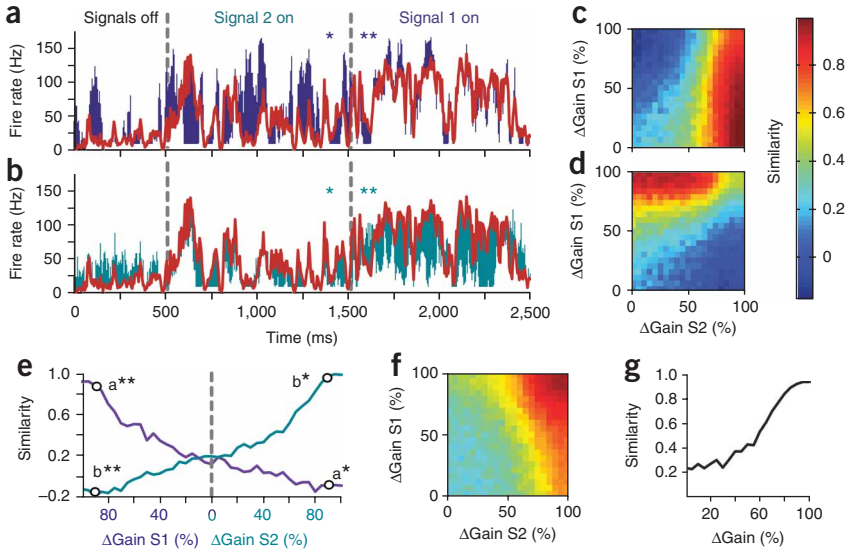

Figure 6 Gating two signals in a network. (**a**,**b**) Average response of the excitatory receiver subnetwork to two simultaneously delivered signals (S1 and S2). Colored areas indicate difference between the average firing rate in the receiver region (red) and either S1 (purple, **a**) or S2 (green, **b**). First column: both signal pathways are balanced, the signals are off. Second and third columns: signal pathways 2 and 1, respectively, are unbalanced by shifting the gain of their respective inhibitory receiver populations to 15% of their control values. (**c**,**d**) Similarity values between S1 and the excitatory receiver activity (**c**) and S2 and the same excitatory receiver activity (**d**) for all possible combinations of the two gain modulations. Both signals reach similarity values of >85%. (**e**) Similarity values for S1 and S2 for independent gain changes. Left of the gray line only the gain for S1 is manipulated while the gain for S2 remains 100%, and vice versa on the right side. Black circles indicate the gain values used for **a**,**b** for the regions denoted by the asterisks. (**f**) Similarity values as in **c**,**d** but measured for the combined signal S1 + S2. (**g**) Similarity values between S1 + S2 and the excitatory receiver subnetwork activity as a function of combined (equal) inhibitory gains, taken from the results along the diagonal of **f**.

can gate two signals, rather than one. We then discuss the capacity of detailed balance for switching a large number of signals using a model that simulates a single neuron in the networks we have been using.

We introduced a second signal to the receiver group. (To accommodate two signals, S1 and S2, we modified the network architecture slightly to allow two separate groups of ~70 inhibitory neurons; see Methods.) We then compared the average firing rate of the excitatory receiver subnetwork compared to each of these signals, S1 and S2 (**Fig. 6a**,**b**). The colored bars indicate the difference between the average firing rate in the receiver region (red) and each of these signals. When both pathways are balanced ('signals off'), firing in the receiver region stays roughly at the background level, except for transient responses as described above. When one of the pathways is unbalanced to allow propagation of its signal, the response of the receiver group follows that signal accurately, as indicated by the small divergence between the appropriate signal and response pair. Activity of the receiver subnetwork does not follow the signal that is in the off state, as indicated by the larger difference regions for the gated-off signals. This finding can also be confirmed using the similarity to S1 and S2, respectively, for all combinations of the two levels of inhibitory gain reduction (**Fig. 6c**,**d**). The regions where similarity is high for either signal are well separated from each other, and both signals reach similarity values above 85% in the regions where they are gated 'on'. Furthermore, when only one of the two signals is gated on, the other signal tends to weaken, and similarity between receiver activity and the gated-off signal can even become negative because of inhibitory overshoot (**Fig. 6e**). To compare the gating of two signals with the gating of one, it is useful to examine the similarity value between the combined signal S1 + S2 and the firing rate of the receiver cells (**Fig. 6f**,**g**). This is much the same as in the single-signal case (**Fig. 4a**, solid red trace).

How many signals can be canceled and then gated 'on' by a population of inhibitory neurons? To address this question, we studied a single excitatory receiver neuron, rather than the full network that we have been considering up to this point. As discussed above for the two-signal case, configuring a network for multiple signals takes a fair amount of modification and readjustment, and this made it unpractical to consider a wide range of different numbers of signals within a full network. Instead, we set up the mechanism of detailed balance in a single integrate-and-fire neuron that receives 800 excitatory and 200 inhibitory inputs. The critical component in determining the capacity of detailed balance for switching multiple signals is the population of inhibitory afferents, because they are less numerous than their excitatory counterparts and must cancel the excitatory effects of the signals while at the same time allowing a particular signal to get through when modulated. We chose 200 inhibitory inputs to match the number that seem to influence a single pyramidal neuron in cortical circuits[26].

The single-neuron model acts much like any of the excitatory receiver neurons in the full network because we adjusted its input to match what a typical neuron receives when the network is intact. The excitatory and inhibitory neurons that provide input to this model neuron are represented by Poisson spike trains generated from firing rates that encode various numbers of signals directly rather than through other model neurons. For this reason, we refer to them as inputs or afferents rather than as neurons—although, of course, they correspond to neurons in the full network. Detailed balance is achieved by distributing the signals across the excitatory and inhibitory afferents and adjusting synaptic strengths so that all signals are cancelled in the default state. To gate a particular signal on, we set the gains of all the inhibitory afferents that carry that signal to zero, essentially shutting them off. We used this extreme form of modulating because we wished to compute the maximum capacity of the system, not a capacity limited by restricting the amount of modulation.

Each signal consisted of a mean firing rate plus independent filtered white noise fluctuations (as used in **Fig. 3c**,**d**). To begin (**Fig. 7**), we distributed M signals across the afferents so that each excitatory and inhibitory afferent carries only one signal. The maximum number of signals that can be distributed in this way is $M = 200$, the number of inhibitory afferents. Performance, measured in terms of a similarity index (see Methods), falls rapidly as a function of the number of signals being gated, and this way of distributing and gating multiple signals does not allow switching of more than ~10 signals (**Fig. 7a**,**d**). The problem is not in canceling out the signals that are not being gated 'on' but in being able to fully gate the chosen signal on.

One problem with the scheme of having one signal per inhibitory afferent is that the number of afferents being gain-modulated to zero to gate a given signal on is small. For example, only 10 inhibitory afferents

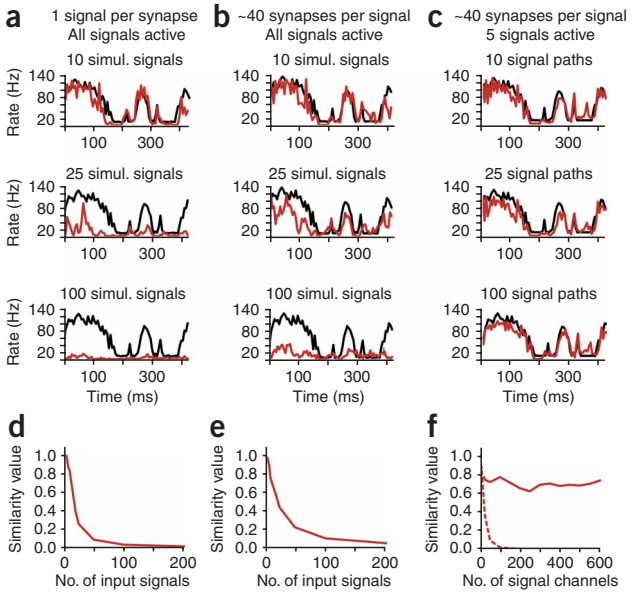

Figure 7 Multiple signals into a single cell. (**a**–**c**) Firing rates of the output signal (red), averaged over 200 runs, compared to the input signal (black). First row: 10 simultaneous (simul.) signal paths. Second row: 25 simultaneous signals paths. Third row: 100 simultaneous signal paths. (**a**) Each afferent to the model neuron carries only one signal. (**b**) Each afferent carries 40 signals. (**c**) Each afferent can carry 40 signals but only 5 signals are present at a given time. (**d**–**f**) Similarity values as a function of the number of signals being gated in **a**–**c**. (**d**) With one signal per afferent, gating is limited to less than about 20 signals. (**e**) Overlapping several signals onto each afferent improves performance slightly. (**f**) When only 5 signals are present at a given time, large numbers of signals can be gated when 40 signals are carried on each afferent (solid trace), but performance is still limited if each afferent carries a single signal (dashed trace).

carry any particular signal when $M = 20$. This problem can be addressed by distributing the M signals across the excitatory and inhibitory afferents so that each afferent carries more than one signal (**Fig. 7b**). Gating works best if each signal is carried on ~40 inhibitory and ~160 excitatory afferents (chosen randomly for each signal from all available afferents of each type), which means that each inhibitory afferent carries, on average, $40M/200$ signals, rather than 1 as before. In this case, because of the overlap in the signals, when a particular signal is gated 'on' by setting the gains of the inhibitory afferents carrying that signal to zero, this upsets the detailed balance for the other, ungated signals. To compensate for this, the gains of the remaining inhibitory afferents are adjusted so that the ungated signals are cancelled as nearly as possible by the remaining active inhibitory afferents. In other words, through a procedure discussed in the Methods, the gains on the remaining active afferents are increased to compensate for those missing owing to gating 'on' of the chosen signal. Performance with this distribution of signals is better than in the one-signal-per-afferent case, but detailed balance still cannot handle more than 30 signals (**Figs. 7b**,**e**).

What is limiting the ability of the detailed balance scheme to switch large numbers of signals? The limitation is, in fact, not a deficit of the detailed balance scheme but a fundamental problem with encoding multiple signals using firing rates, which no switching scheme can avoid. This is the problem of keeping firing rates positive. As discussed above, each signal corresponds to a mean firing rate plus positive and negative fluctuations about this mean. Although the mean rate carries no information about the signal, it cannot be set to zero or half of the signal would be lost owing to firing-rate rectification. Adding together M signals results in a total input that has a mean proportional to M and a fluctuating signal that is proportional to only $M^{1/2}$ because the M signals are independent of one another. Thus, there is a strong tendency for the mean input to drown out the signals to which we want the postsynaptic neuron to respond[27]. Of course, because of the balancing inhibition, the mean excitatory input is canceled, but this cancellation is subject to fluctuations due to the spiking nature of the inputs. As a result, there is a fundamental limitation in the number of signals from which a single signal can be extracted, independent of the method by which this is done.

To show that the limited capacity (**Fig. 7a**,**b**) is a result of this fundamental restriction on firing-rate coding and not a limitation of the detailed-balance approach, we restricted the number of signals present on the afferents to the neuron at any given time. In other words, we set up the postsynaptic neuron so that it could extract any one out of M input signals, but at any given time we restricted the number of signals present to 5 out of these M possibilities (**Fig. 7c**). This seems reasonable for an *in vivo* switching situation: out of the myriad of possible stimuli that can activate a neuron, only a few are likely to be present at any given time. The number 5 is arbitrary; the key is to restrict the number of signals at any given time to a value that does not make the signals undetectable owing to the $M^{1/2}$ scaling problem discussed above. When the input signals are restricted in this way, the capacity is still limited when each afferent is only allowed to carry one signal, but when each signal is distributed across 40 afferents, the switching capacity is much larger (**Fig. 7c**,**f**). There is no decrease in similarity for the gated on signal over the entire range from 1 to 600 possible signals. This shows that the limited performance for the one-signal-per-afferent case (**Fig. 7a**) is a result of having too few afferents per signal. By contrast, the poor performance when 40 afferents are used per signal (**Fig. 7b**) does not represent a limitation of detailed-balance gating but rather a basic limitation of rate-based coding. When this latter limitation is avoided, detailed balance can switch very large numbers of signals.

Table 1 Synapse modifications

		Synapse groups	Healthy	Inhibition deficit	Hyperexcitable receiver	Two signals
Δg_{ex}^{global}	0.8 nS	$\Delta g_{ex\rightarrow ex}^{sender\rightarrow receiver}$	1.1	1.1	1.1	0.9
		$\Delta g_{ex\rightarrow ex}^{net\rightarrow receiver}$	1.0	1.0	1.6	1.0
		$\Delta g_{ex\rightarrow loc\ inh}^{sender\rightarrow receiver\ 1}$	1.0	1.0	1.0	1.1
		$\Delta g_{ex\rightarrow loc\ inh}^{net\rightarrow receiver\ 1}$	1.0	1.0	1.6	1.3
Δg_{inh}^{local}	1.5 nS	All	1.0	0.4	1.0	1.0
		$\Delta g_{loc\ inh\rightarrow ex}^{receiver\ 1\rightarrow receiver}$	3.1	3.1*	3.1	3.3
		$\Delta g_{loc\ inh\rightarrow ex}^{receiver\ 2\rightarrow receiver}$				4.0
Δg_{inh}^{global}	7.5 nS	$\Delta g_{inh\rightarrow loc\ inh}^{net\rightarrow receiver\ 1}$	1.25	1.25	1.25	1.32
		$\Delta g_{inh\rightarrow loc\ inh}^{net\rightarrow receiver\ 2}$				1.2
		$\Delta g_{inh\rightarrow ex}^{net\rightarrow receiver}$				0.75

Overview of synaptic alterations, sorted by synapse type (rows) and purpose of modification (columns). The synaptic strengths appearing in column 2 were multiplied by the factors appearing in columns 4, 5 and 6. The asterisk indicates that this value is 3.1 times the scaled value (scaled by 0.4).

DISCUSSION

We have proposed an extension of the concept of global balance that offers an alternative to the more traditional model of gating by using inhibitory neurons to deactivate a signaling pathway[14,28,29]. This inverts the usual scheme of allowing signal propagation by default and disrupting the signal flow to gate a pathway off. Instead, in our model, the balanced gated off state is the default. In complex networks, gating presumably occurs in parallel along the many pathways responsible for transmitting different aspects of a stimulus. It may be easier for a system to keep track of what is 'interesting' in a broadband signal stream than to keep track of all the 'uninteresting' stimulus features that should be suppressed. If one feature of a stimulus warrants further processing, a control mechanism can select it by unbalancing its respective module, allowing the signal to propagate further downstream. We used gain modulation to unbalance signal-carrying pathways and gate signals on, but ordinary subtractive inhibition of inhibitory interneurons could also be used in our scheme.

Other discussions of signal switching in balanced neural circuits, either by shifting inhibition or through gain modulation, can be found in refs. 2,19,30. Continuous temporally and strengthwise correlated (balanced) excitatory and inhibitory input activity has been reported *in vivo* through recordings from pairs of pyramidal cells in the rat somatosensory cortex during spontaneous and sensory-evoked activities[31]. It has also been demonstrated that inhibition can be used to adjust the gain of a downstream circuit that receives balanced input and can thus control behavioral responses in a context-dependent manner[32]. These findings fit well into the framework of our hypothesis.

Although we found that gating is robust to several perturbations, it would be interesting to study how a homeostatic mechanism might impose and maintain a detailed balance. Spike timing–dependent plasticity has been shown to generate a global balance between excitation and inhibition[33,34], but detailed balance is likely to require some further competitive as well as homeostatic mechanisms. In developing systems, this might involve tuning AMPA and immature, excitatory GABA synapses to the same degree.

Failure to maintain a precise balance between excitation and inhibition due to various abnormalities of synaptic transmission is commonly hypothesized as a basis for mental disorders such as schizophrenia[35–39] and autism[40,41]. Although it is natural to think that such an imbalance might lead to basic instabilities, such as those associated with epilepsy, it is more difficult to understand how they would lead to cognitive and behavioral disorders. Our results provide a suggestion. If we associate the local inhibitory neurons in our model network with parvalbumin-positive inhibitory neurons in cortex, the failure of gating in this model with reduced inhibition (**Fig. 5c**) could provide a functional basis for the hypothesis that reduced GABA production in parvalbumin-positive interneurons may contribute to gating problems in schizophrenia[35]. Similarly, the inability to discriminate between external and internal activity could be related to the hallucinations and delusions that have been hypothesized to arise from defective dopaminergic regulation[36,37] or NMDA current anomalies[38]. Although these latter anomalies correspond to hypofunction of NMDA conductances, this is associated with a hyperexcitability of the affected circuits[39], so we modeled the overall effect by increasing excitation (**Fig. 5d**).

The mechanism we have proposed makes a distinctive prediction concerning inhibitory activity in a signal-receiving region. Although excitatory neurons receiving a signal should respond more vigorously in an attentive (gating on) than in a nonattentive (gating off) state, at least some local inhibitory interneurons should follow a signal more reliably in the inattentive state and should decorrelate their activity from the stimulus with attention.

In studying the capacity of detailed balance to switch multiple signals, we encountered a fundamental limitation of multisignal encoding using firing rates that affects any switching scheme. However, once this limitation was avoided by keeping too many of the possible signals from being present at any given time, we found a large capacity to gate multiple signals. Thus, detailed balance offers a powerful and dynamic way of controlling signal flow in complex and multiply interconnected circuitry.

METHODS

Neuron model. The model used for all our simulations is a leaky integrate-and-fire neuron, 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 fires 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 used a membrane resistance of 100 MΩ.

We modeled synapses onto each neuron as conductances, so the subthreshold membrane voltage obeys

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

Reversal potentials were $E_{\text{ex}} = 0$ mV and $E_{\text{inh}} = -80$ mV. The synaptic conductances g_{ex} and g_{inh} are expressed in units of the resting membrane conductance. When the neuron receives a presynaptic action potential, the appropriate postsynaptic variable increases, $g_{\text{ex}} \rightarrow g_{\text{ex}} + \Delta g_{\text{ex}}$ for an excitatory spike and $g_{\text{inh}} \rightarrow g_{\text{inh}} + \Delta g_{\text{inh}}$ for an inhibitory spike. Otherwise, these parameters obey the equations

$$\tau_{\text{ex}} \frac{dg_{\text{ex}}}{dt} = -g_{\text{ex}} \quad \text{and} \quad \tau_{\text{inh}} \frac{dg_{\text{inh}}}{dt} = -g_{\text{inh}}$$

with synaptic time constants $\tau_{\text{ex}} = 5$ ms and $\tau_{\text{inh}} = 10$ ms. I_{b} is a constant background current used to maintain network activity (see below). The integration time step for the simulations was 0.1 ms. We implemented all simulations in C.

Network architecture. We studied a network of 20,164 leaky integrate-and-fire neurons, laid out on a 142 × 142 grid. Neurons were either excitatory or inhibitory. The ratio of inhibitory neurons was roughly 1 in 4, but the geometric organization of neurons on the grid constrained the final numbers to 15,123 excitatory cells and 5,041 inhibitory cells. Inhibitory neurons were divided into two groups of 3,361 and 1,680 neurons differing in their connectivity pattern. All excitatory neurons and 65% of the inhibitory neurons had a random connectivity of 2% to the rest of the network (red cells in **Fig. 1a**). The 1,680 inhibitory neurons of the second group each targeted 40% of their 500 closest neighbors, thus acting locally[42] (blue cells in **Fig. 1a**). To avoid boundary effects, we implemented the network with the topology of a torus. We used 20,000 cells because this was the largest network that we could study within reasonable computation times. It has been shown previously that the activity in such networks becomes independent of size at about 10,000 neurons[8,11,16]. Other network parameters were chosen in keeping with both general properties of cortical circuits and previous work[11,12,14,16,42].

Signal path. In additional to the general architecture, we introduced a specific pathway from one region of the network to another, which we call sender and receiver subnetworks (**Fig. 1b**). The two subnetworks were chosen to be sufficiently distant from each other to exclude possible interactions through local inhibitory neurons. Synapses from a given excitatory sender neuron were allocated to contact either the excitatory or the locally inhibitory neurons of the receiver region, but not both. This division is made for the sake of tuning simplicity. The numbers of neurons and projecting synapses for the sender subnetwork were chosen to supply each of the ~500 excitatory receiver and ~70 inhibitory receiver neurons with 50 synapses from the sender subnetwork, a number necessary for critical spike propagation with appropriate tuning of

synaptic strengths (see below). In addition, the number of neurons in both subnetworks was chosen to be as large as possible without interfering with overall network activity during signal propagation[14]. The final numbers were 494 excitatory–excitatory sender neurons, 234 excitatory–inhibitory sender neurons, 463 excitatory receiver neurons and 73 inhibitory receiver neurons.

Tuning conditions. Except for synapses along the signaling pathway and those mentioned further below, all synapses of the same cell type had the same strength. A background current (I_b) of 0.03 nA was delivered to every neuron and the three sets of strengths were adjusted to allow asynchronous background activity within the network. The postsynaptic conductances of $\Delta g_{ex}^{global} = 0.8$ nS, $\Delta g_{inh}^{local} = 1.5$ nS and $\Delta g_{inh}^{global} = 7.5$ nS correspond to 0.5 mV excitatory postsynaptic potentials and −0.4 mV and −1.1 mV inhibitory postsynaptic potentials, respectively, as obtained from spike-triggered averages in the active network. To propagate signals from the sender subnetwork to the excitatory receiver subnetwork, we set the synapses between those two groups to $\Delta g_{ex\to ex}^{sender\to receiver} = 0.9$ nS. The synapses between sender neurons and inhibitory receiver neurons were left unchanged, and the inhibitory synapses between the inhibitory and the excitatory neurons in the receiver subnetwork were strengthened to $\Delta g_{loc\ inh\to ex}^{receiver\to receiver} = 4.65$ nS. The extra tonic excitation that the inhibitory receiver neurons receive from background activity through the projections from the sender subnetwork was compensated by strengthening global inhibition to these cells to $\Delta g_{inh\to loc\ inh}^{net\to receiver} = 9.4$ nS. Under these conditions the system is sufficiently balanced to prevent correlated inputs in the sender subnetwork from modulating the firing pattern of the excitatory receiver subnetwork.

To allow propagation, the balance between the excitatory and the inhibitory signal was modified by decreasing the gain of the inhibitory receiver neurons. In integrate-and-fire neurons such a gain change is equivalent to reducing the strength of all synapses by 85% in the case of symmetric gain changes and by 30% in the asymmetric case in which only the response amplitude to excitatory inputs is altered. These values were chosen to minimize similarity values in the gated-off state (**Fig. 4a**).

To compensate for cell loss (**Fig. 4c**), we calculated the decrease of overall synaptic strength and redistributed the difference equally among the remaining synapses in the pathway. The gating mechanism can function with only 20 inhibitory receiver neurons when we strengthen their synapses threefold, to 15 nS (approximately the same strength as two globally inhibitory neurons), and allow them to fire at rates of up to 600 Hz (**Fig. 4d**).

Two input signals. To propagate and control an additional signal to the receiver region, a second inhibitory receiver subnetwork is necessary. Some of the globally acting inhibitory neurons in the receiver region can be recruited as locally inhibitory for that purpose by generating a new local architecture for them. An additional set of Poisson input spike trains is generated and connected synaptically to both the shared excitatory receiver subnetwork and the new inhibitory receiver subnetwork. As before, the synapses of the Poisson population are tuned to drive their respective target subnetworks in the absence of additional correlated signal input. To balance two active signals at the same time, some of the synapses within the network must also be retuned. See **Table 1** for a complete listing of all modified synapses.

Pathologies. We chose two ways to disrupt the detailed balance mechanism (**Table 1**). First, we introduced a deficiency in the locally inhibitory neurons, including those of the inhibitory receiver subnetwork, by decreasing their synaptic strengths to $\Delta g_{inh}^{local} = 0.6$ nS. Second, we induced a hyperexcitability of the receiver region by increasing all the excitatory synapses from the rest of the network onto the receiver neurons to $\Delta g_{ex\to ex}^{net\to receiver} = 1.28$ nS and $\Delta g_{ex\to loc\ inh}^{net\to receiver} = 1.28$ nS. These two manipulations are independent and can be combined without retuning.

Multiple input signals to a single cell. In the later part of the paper, we modeled the gating of multiple signals in a single integrate-and-fire cell that receives 800 excitatory and 200 inhibitory synapses modeled as Poisson processes with temporally changing spiking probabilities. To avoid unrealistically high membrane currents when many inputs arrive at the cell, synaptic strengths were tuned down to $g_{ex} = 0.014$ nS and $\Delta g_{inh} = 0.044$ nS with postsynaptic potentials of 0.13 mV and −0.22 mV at V_{rest}, respectively. When we drove each afferent with more than one signal, the overlap (effectively a summation of spiking probabilities for each synapse) demanded a rescaling of the input rates to a dynamic signal range between 0 and 150 Hz.

The following procedure is used to compute the inhibitory gain factors needed to gate 'on' one signal among many. First, to specify which signal is connected to which inhibitory afferent, we define an $N \times M$ matrix B with $B_{ia} = 1$ if inhibitory afferent i receives signal a and $B_{ia} = 0$ if it does not. The total inhibitory input due to signal a is then $C_a = \sum_i B_{ia}$. We next choose a particular signal, say signal 1, to gate on. We do this by setting the gains for all the inhibitory neurons receiving signal 1—that is, all neurons with $B_{i1} = 1$—to zero. We then adjust the firing rates (gains) of the remaining inhibitory afferents to compensate for these missing afferents for all other signals (missing because their gains are zero). If n afferents receive signal 1, we define an $(N-n) \times (M-1)$ matrix $\tilde{B}$, which is just B with signal 1 and all of the afferents connected to signal 1 removed. Define $\tilde{C}$ to be a vector with $M-1$ components given by $\tilde{C}_a = C_{a+1}$ for $a = 1, 2, \ldots, M-1$. The inhibition missing because the afferents receiving signal 1 have been turned off can be replaced if the firing rates of the afferents not receiving signal 1 are multiplied by a row vector of gain factors α such that $\alpha\tilde{B} = \tilde{C}$. This equation is 'solved', in the sense of minimizing the square of the difference between the two sides summed over a, by setting $\alpha = \mathrm{pinv}(\tilde{B})\tilde{C}$, where $\mathrm{pinv}(\tilde{B})$ is the pseudoinverse of $\tilde{B}$. The gains of all inhibitory afferents receiving signal 1, by contrast, are set to zero. This determines the complete set of gains used to gate signal 1 on and leave all other signals off. To gate signal 1 off, all the gains are set to 1. A similar procedure is used for any other signal that we wish to gate.

Response properties of the balance mechanism. To supply a signal to the network, we generated Poisson input spike trains with a firing rate $r_0(t)$ as a source of input to the network. Each input spike generated by that group increased the excitatory synaptic conductances in neurons of the sender region by $g_{ex} \to g_{ex} + \Delta g_0$. The synaptic strength Δg_0 was tuned so that the firing rates of the sender neurons reproduce the input signal, that is, they track the input firing rate $r_0(t)$.

We used a correlation measure[12,14] to determine how similar the firing rates in the receiver region were to the input. To do this, we calculated the population firing rate $r(t)$ in 5-ms bins by counting spikes and also determined its time-averaged value $\bar{r}$. The correlation is then

$$C(\tau) = \frac{\langle (r_0(t) - \bar{r}_0)(r(t+\tau) - \bar{r})\sigma_{r_0} \rangle_t}{\langle (r_0(t) - \bar{r}_0)(r_0(t+\tau) - \bar{r}_0)\sigma_r \rangle_t}$$

where the brackets denote an average over time, $r_0(t)$ and $\bar{r}_0$ are the firing rate and average for the input, and σ_r and σ_{r_0} are the s.d. values of the corresponding firing rates. We used the activity of the input as a reference rather than the sender subnetwork to distinguish signal transmission from propagation of fluctuations arising in the sender. Signal propagation between subnetworks is then characterized by reporting the maximum value (over τ) of $C(\tau)$, which we call the similarity[14]. For the analysis of multisignal gating in a single integrate-and-fire cell, we used a similarity measure that was not normalized by σ_r to avoid overestimating the quality of the output signal in cases when the output firing rate was greatly diminished.

ACKNOWLEDGMENTS

The idea of detailed balance was originally suggested to us by G. Turrigiano. Research supported by the US National Science Foundation (IBN-0235463), the Swartz Foundation, the Patterson Trust Fellowship Program in Brain Circuitry and a US National Institutes of Health (NIH) Director's Pioneer Award, part of the NIH Roadmap for Medical Research, through grant number 5-DP1-OD114-02. Thanks to J. Peelle, M. Schiff, P. Jercog and R. Yuste for suggestions.

Published online at http://www.nature.com/natureneuroscience/
Reprints and permissions information is available online at http://npg.nature.com/reprintsandpermissions/

1. Shu, Y., Hasenstaub, A. & McCormick, D.A. Turning on and off recurrent balanced cortical activity. *Nature* **423**, 288–293 (2003).
2. Haider, B., Duque, A., Hasenstaub, A.R. & McCormick, D.A. Neocortical network activity in vivo is generated through a dynamic balance of excitation and inhibition. *J. Neurosci.* **26**, 4535–4545 (2006).

3. Shadlen, M.N. & Newsome, W.T. Noise, neural codes and cortical organization. *Curr. Opin. Neurobiol.* **4**, 569–579 (1994).
4. Troyer, T.W. & Miller, K.D. Physiological gain leads to high ISI variability in a simple model of a cortical regular spiking cell. *Neural Comput.* **9**, 971–983 (1997).
5. Amit, D.J. & Brunel, N. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb. Cortex* **7**, 237–252 (1997).
6. van Vreeswijk, C. & Sompolinsky, H. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* **274**, 1724–1726 (1996).
7. Brunel, N. Dynamics of networks of randomly connected excitatory and inhibitory spiking neurons. *J. Physiol. (Paris)* **94**, 445–463 (2000).
8. Kumar, A., Schrader, S., Aertsen, A. & Rotter, S. The high-conductance state of cortical networks. *Neural Comput.* **20**, 1–43 (2008).
9. Abeles, M. *Corticonics: Neural Circuits of the Cerebral Cortex* (Cambridge University Press, Cambridge, UK, 1991).
10. Aertsen, A., Diesmann, M. & Gewaltig, M.O. Propagation of synchronous spiking activity in feedforward neural networks. *J. Physiol. (Paris)* **90**, 243–247 (1996).
11. Diesmann, M., Gewaltig, M.O. & Aertsen, A. Stable propagation of synchronous spiking in cortical neural networks. *Nature* **402**, 529–533 (1999).
12. van Rossum, M.C., Turrigiano, G.G. & Nelson, S.B. Fast propagation of firing rates through layered networks of noisy neurons. *J. Neurosci.* **22**, 1956–1966 (2002).
13. Vogels, T.P., Rajan, K. & Abbott, L.F. Neural networks dynamics. *Annu. Rev. Neurosci.* **28**, 357–376 (2005).
14. Vogels, T.P. & Abbott, L.F. Signal propagation and logic gating in networks of integrate-and-fire neurons. *J. Neurosci.* **25**, 10786–10795 (2005).
15. Destexhe, A. & Contreras, D. Neuronal computations with stochastic network states. *Science* **314**, 85–90 (2006).
16. Kumar, A., Rotter, S. & Aertsen, A. Conditions for propagating synchronous spiking and asynchronous firing rates in a cortical network model. *J. Neurosci.* **28**, 5268–5280 (2008).
17. Posner, M.I. ed. *Cognitive Neuroscience of Attention* (Guilford Press, New York, 2004).
18. Germuska, M., Saha, S., Fiala, J. & Barbas, H. Synaptic distinction of laminar-specific prefrontal-temporal pathways in primates. *Cereb. Cortex* **16**, 865–875 (2006).
19. Salinas, E. Context-dependent selection of visuomotor maps. *BMC Neurosci.* **5**, 47–68 (2004).
20. Disney, A.A., Aoki, C. & Hawken, M.J. Gain modulation by nicotine in macaque v1. *Neuron* **56**, 701–713 (2007).
21. Disney, A.A. & Aoki, C. Muscarinic acetylcholine receptors in macaque V1 are most frequently expressed by parvalbumin-immunoreactive neurons. *J. Comp. Neurol.* **507**, 1748–1762 (2008).
22. Disney, A.A., Domakonda, K.V. & Aoki, C. Differential expression of muscarinic acetylcholine receptors across excitatory and inhibitory cells in visual cortical areas V1 and V2 of the macaque monkey. *J. Comp. Neurol.* **499**, 49–63 (2006).
23. Xiang, Z., Huguenard, J.R. & Prince, D.A. Cholinergic switching within neocortical inhibitory networks. *Science* **281**, 985–988 (1998).
24. Gil, Z., Connors, B.W. & Amitai, Y. Differential regulation of neocortical synapses by neuromodulators and activity. *Neuron* **19**, 679–686 (1997).
25. Mitchell, J.F., Sundberg, K.A. & Reynolds, J.H. Differential attention-dependent response modulation across cell classes in macaque visual area V4. *Neuron* **55**, 131–141 (2007).
26. Binzegger, T., Douglas, R.J. & Martin, K.A. A quantitative map of the circuit of cat primary visual cortex. *J. Neurosci.* **24**, 8441–8453 (2004).
27. Abbott, L.F. Theoretical neuroscience rising. *Neuron* **60**, 489–495 (2008).
28. Anderson, C.H. & Van Essen, D.C. Shifter circuits: a computational strategy for dynamic aspects of visual processing. *Proc. Natl. Acad. Sci. USA* **84**, 6297–6301 (1987).
29. Olshausen, B.A., Anderson, C.H. & Van Essen, D.C. A neurobiological model of visual attention and invariant pattern recognition based on dynamical routing of information. *J. Neurosci.* **13**, 4700–4719 (1993).
30. Pouille, F. & Scanziani, M. Routing of spike series by dynamic circuits in the hippocampus. *Nature* **429**, 717–723 (2004).
31. Okun, M. & Lampl, I. Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nat. Neurosci.* **11**, 535–537 (2008).
32. Baca, S.M., Marin-Burgin, A., Wagenaar, D.A. & Kristan, W.B. Jr. Widespread inhibition proportional to excitation controls the gain of a leech behavioral circuit. *Neuron* **57**, 276–289 (2008).
33. Song, S., Miller, K.D. & Abbott, L.F. Competitive Hebbian learning through spike-timing dependent synaptic plasticity. *Nat. Neurosci.* **3**, 919–926 (2000).
34. Morrison, A., Aertsen, A. & Diesmann, M. Spike-timing-dependent plasticity in balanced random networks. *Neural Comput.* **19**, 1437–1467 (2007).
35. Lewis, D.A., Hashimoto, T. & Volk, D.W. Cortical inhibitory neurons and schizophrenia. *Nat. Rev. Neurosci.* **6**, 312–324 (2005).
36. Seeman, P. Dopamine receptors and the dopamine hypothesis of schizophrenia. *Synapse* **1**, 133–152 (1987).
37. Moore, H., West, A.R. & Grace, A.A. The regulation of forebrain dopamine transmission: relevance to the pathophysiology and psychopathology of schizophrenia. *Biol. Psychiatry* **46**, 40–55 (1999).
38. Tamminga, C.A. Schizophrenia and glutamatergic transmission. *Crit. Rev. Neurobiol.* **12**, 21–36 (1998).
39. Jackson, M.E., Homayoun, H. & Moghaddam, B. NMDA receptor hypofunction produces concomitant firing rate potentiation and burst activity reduction in the prefrontal cortex. *Proc. Natl. Acad. Sci. USA* **101**, 8467–8472 (2004).
40. Rubenstein, J.L. & Merzenich, M.M. Model of autism: increased ratio of excitation/inhibition in key neural systems. *Genes Brain Behav.* **2**, 255–267 (2003).
41. Tabuchi, K. *et al.* A neuroligin-3 mutation implicated in autism increases inhibitory synaptic transmission in mice. *Science* **318**, 71–76 (2007).
42. Aviel, Y., Mehring, C., Abeles, M. & Horn, D. On embedding synfire chains in a balanced network. *Neural Comput.* **15**, 1321–1340 (2003).

pronounced than in *drbp* nulls (*5*, *14*). Functionally, *drbp* and *bruchpilot* phenotypes appear similar: Both demonstrate decreased and desynchronized evoked SV release with atypical short-term facilitation. However, the deficits in evoked SV release are much more severe in *drbp* nulls than in *bruchpilot* nulls [i.e., release occurs at 5% versus 30% (*5*) of the respective wild-type level]. DRBP levels were clearly reduced in *bruchpilot* mutants (fig. S7), whereas gross Bruchpilot levels were not altered in *drbp* mutants (Fig. 2B). Given that even a partial loss of DRBP causes marked reduction in SV release (Fig. 3A), deficits in *bruchpilot* mutants might be explained, at least in part, by a concomitant loss of DRBP, and DRBP probably serves functions beyond the structural and Ca^{2+} channel–clustering roles of Bruchpilot.

Taken together, we identified DRBP as a central part of the AZ cytomatrix. How, in detail, DRBP functionally integrates into this protein network is subject to future analyses. Notably, the short-term plasticity phenotype of *drbp* mutants is reminiscent of mammalian *munc13-1* KO and *caps-1* and *caps-2* DKO mutants (*25*, *26*), which implicates functional links between priming factors and DRBP. Consistent with the functional importance of the DRBP protein family suggested by our study, human genetics recently identified two *rbp* loci associated with autism with high confidence (*27*, *28*).

References and Notes

1. Y. Jin, C. C. Garner, *Annu. Rev. Cell Dev. Biol.* **24**, 237 (2008).
2. S. J. Sigrist, D. Schmitz, *Curr. Opin. Neurobiol.* **21**, 144 (2011).
3. S. Schoch, E. D. Gundelfinger, *Cell Tissue Res.* **326**, 379 (2006).
4. L. Siksou, A. Triller, S. Marty, *Curr. Opin. Neurobiol.* **21**, 261 (2011).
5. R. J. Kittel *et al.*, *Science* **312**, 1051 (2006).
6. Y. Wang, S. Sugita, T. C. Sudhof, *J. Biol. Chem.* **275**, 20033 (2000).
7. H. Hibino *et al.*, *Neuron* **34**, 411 (2002).
8. S. A. Spangler, C. C. Hoogenraad, *Biochem. Soc. Trans.* **35**, 1278 (2007).
9. T. Mittelstaedt, S. Schoch, *Gene* **403**, 70 (2007).
10. S. W. Hell, *Science* **316**, 1153 (2007).
11. J. Bückers, D. Wildanger, G. Vicidomini, L. Kastrup, S. W. Hell, *Opt. Express* **19**, 3130 (2011).
12. J. Hou, T. Tamura, Y. Kidokoro, *J. Neurophysiol.* **100**, 2833 (2008).
13. F. Kawasaki, R. Felling, R. W. Ordway, *J. Neurosci.* **20**, 4885 (2000).
14. W. Fouquet *et al.*, *J. Cell Biol.* **186**, 129 (2009).
15. F. Kawasaki, S. C. Collins, R. W. Ordway, *J. Neurosci.* **22**, 5856 (2002).
16. K. J. Venken, Y. He, R. A. Hoskins, H. J. Bellen, *Science* **314**, 1747 (2006).
17. L. Siksou *et al.*, *J. Neurosci.* **27**, 6868 (2007).
18. P. Rostaing, R. M. Weimer, E. M. Jorgensen, A. Triller, J. L. Bessereau, *J. Histochem. Cytochem.* **52**, 1 (2004).
19. S. Hallermann *et al.*, *J. Neurosci.* **30**, 14340 (2010).
20. E. O. Gracheva, E. B. Maryon, M. Berthelot-Grosjean, J. E. Richmond, *Front. Synaptic Neurosci.* **2**, 141 (2010).
21. E. Neher, T. Sakaba, *Neuron* **59**, 861 (2008).
22. F. Kawasaki, B. Zou, X. Xu, R. W. Ordway, *J. Neurosci.* **24**, 282 (2004).
23. P. S. Kaeser *et al.*, *Cell* **144**, 282 (2011).
24. Y. Han, P. S. Kaeser, T. C. Südhof, R. Schneggenburger, *Neuron* **69**, 304 (2011).
25. C. Rosenmund *et al.*, *Neuron* **33**, 411 (2002).
26. W. J. Jockusch *et al.*, *Cell* **131**, 796 (2007).
27. M. Bucan *et al.*, *PLoS Genet.* **5**, e1000536 (2009).
28. D. Pinto *et al.*, *Nature* **466**, 368 (2010).

Acknowledgments: This work was supported by Deutsche Forschungsgemeinschaft (DFG) grants (SFB 665, SFB 958, and EXC 257) to S.J.S and D.S., as well as Bundesministerium für Bildung und Forschung (The German Federal Agency of Education and Research) funding for Deutsche Zentrum für Neurodegenerative Erkrankungen (DZNE) to D.S. Also, M.S. was supported by a Ph.D. fellowship from the Max Delbrück Center for Molecular Medicine and a Boehringer Ingelheim Fonds Ph.D. fellowship. E.K. and S.W. were supported by Ph.D. fellowships from the graduate school GRK 1123 funded by the DFG. M.M. was supported by a fellowship of the Swiss National Science Foundation (PBSKP3-123456/1).

Supporting Online Material

www.sciencemag.org/cgi/content/full/334/6062/1565/DC1
Materials and Methods
Figs. S1 to S7
References (*29*–*39*)

22 August 2011; accepted 2 November 2011
10.1126/science.1212991

Inhibitory Plasticity Balances Excitation and Inhibition in Sensory Pathways and Memory Networks

T. P. Vogels,[1]*† H. Sprekeler,[1]* F. Zenke,[1] C. Clopath,[1,2] W. Gerstner[1]

Cortical neurons receive balanced excitatory and inhibitory synaptic currents. Such a balance could be established and maintained in an experience-dependent manner by synaptic plasticity at inhibitory synapses. We show that this mechanism provides an explanation for the sparse firing patterns observed in response to natural stimuli and fits well with a recently observed interaction of excitatory and inhibitory receptive field plasticity. The introduction of inhibitory plasticity in suitable recurrent networks provides a homeostatic mechanism that leads to asynchronous irregular network states. Further, it can accommodate synaptic memories with activity patterns that become indiscernible from the background state but can be reactivated by external stimuli. Our results suggest an essential role of inhibitory plasticity in the formation and maintenance of functional cortical circuitry.

The balance of excitatory and inhibitory membrane currents that a neuron experiences during stimulated and ongoing activity has been the topic of many studies (*1*–*11*). This balance, first defined as equal average amounts of de- and hyperpolarizing membrane currents (from here on referred to as "global balance"), is essential for maintaining stability of cortical networks (*1*, *2*). Balanced networks display asynchronous irregular (AI) dynamics that mimic activity patterns observed in cortical neurons. Such asynchronous network states facilitate rapid responses to small changes in the input (*2*, *3*, *12*), providing an ideal substrate for cortical signal processing (*4*, *13*, *14*).

Moreover, input currents to cortical neurons are not merely globally balanced but also coupled in time (*5*, *6*, *15*) and cotuned for different stimulus features (*7*, *8*). The tight coupling of excitation and inhibition suggests a more precise, detailed balance, in which each excitatory input arrives at the cell together with an inhibitory counterpart (Fig. 1A), permitting sensory inputs to be transiently (*9*) or persistently turned on by targeted disruptions of the balance (*10*, *11*).

Although the excitatory-inhibitory balance plays an important role for stability and information processing in cortical networks, it is not understood by which mechanisms this balance is established and maintained during ongoing sensory experiences. Inspired by recent experimental results (*7*), we investigated the hypothesis that synaptic plasticity at inhibitory synapses plays a central role in balancing the excitatory and inhibitory inputs a cell receives.

We simulated a single postsynaptic integrate-and-fire neuron receiving correlated excitatory and inhibitory input signals. The cell received input through 1000 synapses (Fig. 1B), which were divided into eight independent groups of 100 excitatory and 25 inhibitory synapses. All excitatory and inhibitory synapses within each group followed the same temporally modulated rate signal (time constant $\tau \sim 50$ ms) to mimic ongoing sensory activity (*13*, *16*). Spikes were generated from independent Poisson processes, leading to 125 different spike trains per signal. This architecture allowed each signal to reach the cell simultaneously through both excitatory and inhibitory synapses (Fig. 1B). To mimic glutamatergic and γ-aminobutyric acid (GABAergic) transmission, the synapses were conductance-based

[1]School of Computer and Communication Sciences and Brain-Mind Institute, École Polytechnique Fédérale de Lausanne, 1015 Lausanne EPFL, Switzerland. [2]CNRS, UMR 8119, Université Paris Descartes, 45 Rue des Saints Pères, 75270 Paris Cedex 06, France.

*These authors contributed equally to this work.
†To whom correspondence should be addressed. E-mail: tim.vogels@epfl.ch

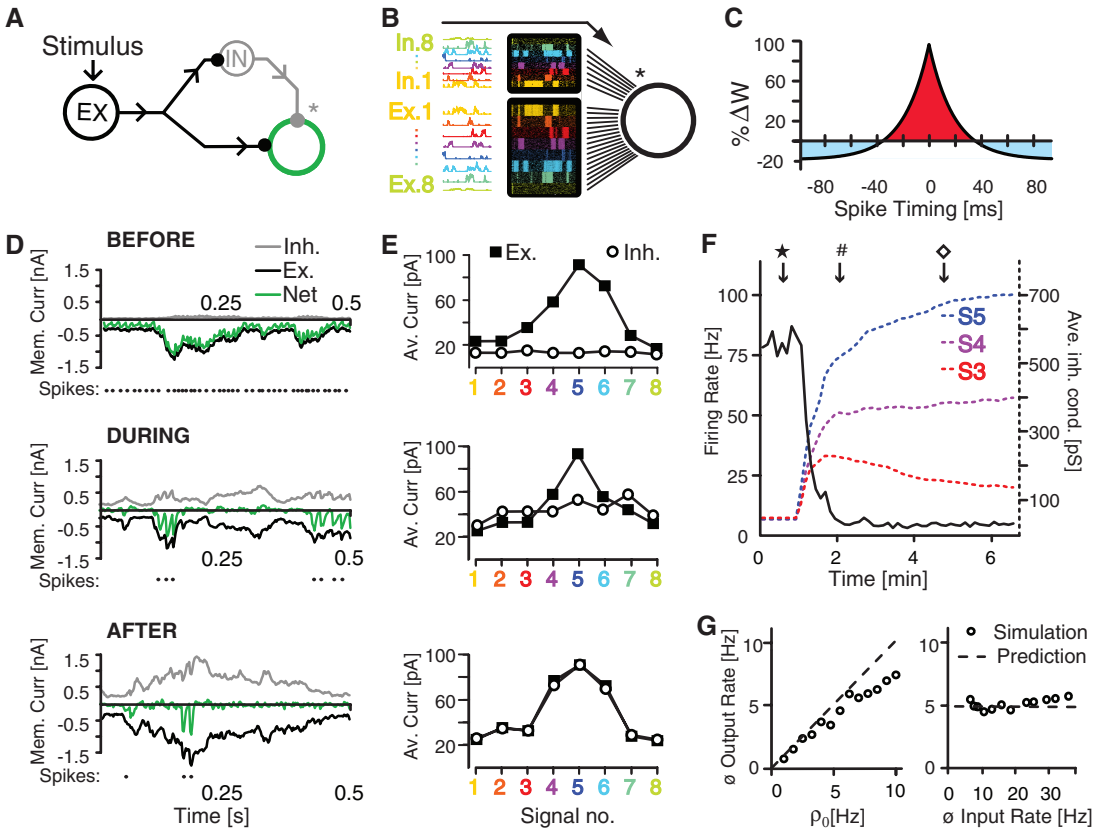

Fig. 1. Inhibitory synaptic plasticity balances excitation and inhibition. (**A**) Feedforward inhibition: Excitatory input reaches a target region through both direct excitation and indirect disynaptic inhibition. (**B**) Feedforward inhibition for a single postsynaptic cell: Eight groups of 100 excitatory and 25 inhibitory synapses each deliver spikes to a single postsynaptic cell. Spiking probabilities are homogeneous within the groups but vary in time, simulating eight separate (color-coded) signal channels that reach the cell simultaneously through excitatory and inhibitory synapses. (**C**) Spike-timing–dependent learning rule: Near-coincident pre- and postsynaptic spikes potentiate inhibitory synapses [marked with * in (A) and (B)], whereas every presynaptic spike causes synaptic depression. (**D**) Total excitatory (black), inhibitory (gray), and net (green) membrane currents before, during, and after inhibitory synaptic plasticity. The resulting spikes are indicated as dots underneath each current plot. (**E**) Excitatory and inhibitory membrane currents (black and white symbols, respectively) evoked by each signal channel, averaged over 4 s, before, during, and after inhibitory synaptic plasticity (top, middle, and bottom, respectively). (**F**) Temporal evolution of the postsynaptic firing rate (solid line) and the average synaptic weights of the inhibitory synapses associated with three representative signals (dotted lines). ★, #, and ◊ indicate the times at which the top, middle, and bottom graphs of (D) and (E) were recorded. (**G**) Average firing rate of the postsynaptic neuron after learning, plotted for different values of target firing rate ρ_0 (left) and different input rates (right). The dashed lines in both graphs show the analytical predictions.

with reversal potentials $V^E = 0$ mV and $V^I = -80$ mV and time constants $\tau^E = 5$ ms, and $\tau^I = 10$ ms for excitation and inhibition, respectively [see supporting online material (SOM)]. The strength of the inhibitory synapses was initially weak but could change according to a spike-timing–dependent plasticity rule, in which near-coincident pre- and postsynaptic spikes induce potentiation of the synapse (*17*–*19*). Additionally, every presynaptic spike leads to synaptic depression (*17*, *18*) (Fig. 1C). This learning rule can be summarized as

$$\Delta w = \eta(pre \times post - \rho_0 \times pre) \qquad (1)$$

where Δw denotes the change in synaptic efficacy, *pre* and *post* are the pre- and postsynaptic activity, η is the learning rate, and ρ_0 is a constant that acts as a target rate for the postsynaptic neuron (see SOM Sec. 2 for a mathematical analysis).

Whereas inhibitory synapses were plastic, the efficacies of the excitatory model synapses were fixed at the beginning of a simulation and left unchanged unless otherwise noted. Analogous to frequency- or orientation-tuned sensory neurons, excitatory synapses were tuned to have a preferred signal (Fig. 1E). Because all excitatory synapses were set to nonzero strengths, the postsynaptic neuron fired at high rates when the inhibitory synapses were weak at the beginning of a simulation (Fig. 1, D and E, top, and F). The resulting high number of pairs of pre- and postsynaptic spikes led to relatively indiscriminate strengthening of all inhibitory synapses (Fig. 1, D and E, middle) until excitatory and inhibitory membrane currents became approximately balanced and the postsynaptic firing rate was dramatically reduced (Fig. 1F). In this globally balanced state, only unbalanced excitatory signals led to coincident pairs of pre- and postsynaptic spikes, consequently strengthening underpowered inhibitory synapses. Those inhibitory synapses that were stronger than their excitatory counterparts kept the postsynaptic side unresponsive and were thus weakened (because of sole presynaptic firing) until they allowed postsynaptic spiking again. Over time, this led to a precise, detailed balance of excitatory and inhibitory synaptic weights for each channel (Fig. 1, D and E, bottom). In agreement with the mathematical analysis, the postsynaptic firing rate was determined mainly by the depression factor, ρ_0, but not by the average input firing rate to the postsynaptic neuron (Fig. 1G). The mechanism was robust to plausible delays of several milliseconds. However, because detailed balance requires a correlation between excitatory and inhibitory synaptic inputs, the balance deteriorated when the delay between excitation and inhibition increased to values larger than the autocorrelation time of the input signals and the coincidence time of the Hebbian learning rule, but global balance still persisted (fig. S2).

To investigate how the state of the balance affects the neuron's response properties, we presented a fixed stimulus sequence to the neuron (Fig. 2A) and compared the spiking response over 50 trials to the input rates of each signal. In the globally balanced state (Fig. 2B, top) in which inhibitory synapses were distributed so that excitation and inhibition were balanced only on average across all channels, the peristimulus time histogram (PSTH) faithfully reproduced the firing rates of the preferred signals. The other, non-preferred input signals evoked more inhibition than excitation and thus had no impact on the cell's firing behavior. An additional steplike input rate protocol, in which 100-ms-long pulses of various step sizes (Fig. 2C) were presented to one channel at a time, revealed that spiking responses are largely insensitive to stimulus intensity and indeed narrowly tuned to the preferred stimulus, giving rise to an all-or-none response (Fig. 2, D and E).

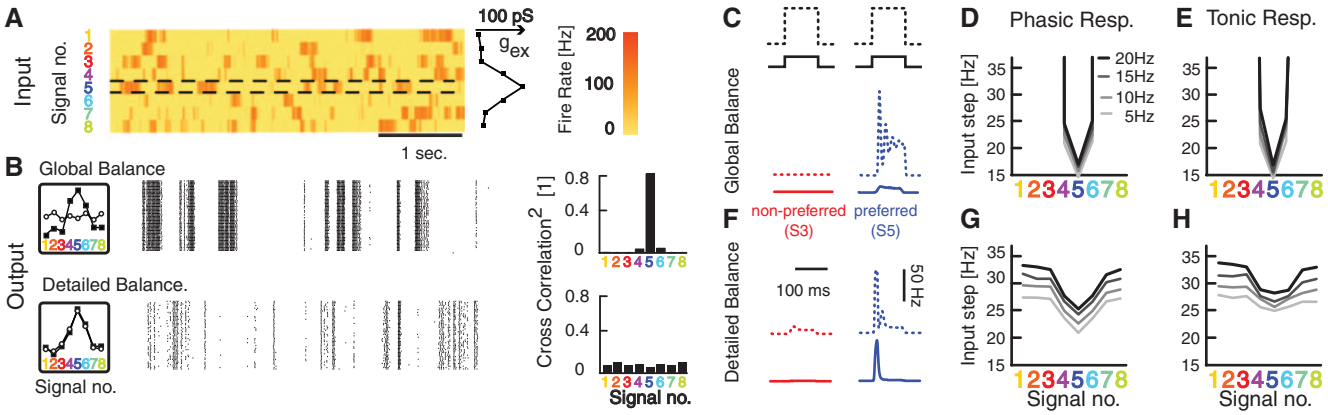

Fig. 2. Inhibitory synaptic plasticity sparsifies and democratizes receptive fields. (**A**) A fixed sequence of eight stimuli of varying firing rates is fed repetitively into a postsynaptic cell. Excitatory synapses are strength-tuned by signal group (see conductance graph on the right) so that signal five (marked also by dashed lines) is the preferred signal. (**B**) Postsynaptic spikes over 50 trials with globally or detailed balanced inhibitory synapses (top and bottom graphs, respectively) as indicated by the schematics on the left (compare with Fig. 1E). The normalized and squared cross-correlation coefficients between each input signal and the PSTH are also shown (right). (**C**) Schematic of a step stimulus delivered with large and small step sizes (solid and dotted black lines respectively); Sample PSTHs for nonpreferred (red) and preferred (blue) stimuli to both step sizes are shown for a globally balanced cell. (**D** and **E**) Iso-response contour lines of the postsynaptic cell in the globally balanced regime during the onset (phasic) (0 to 50 ms) (D) and tonic (50 to 100 ms) (E) parts of the response. (**F**) Sample responses for nonpreferred (red) and preferred (blue) stimuli to both step stimuli [as in (C)]. (**G** and **H**) Iso-response contour lines [as in (D) and (E)] for a detailed balanced cell.

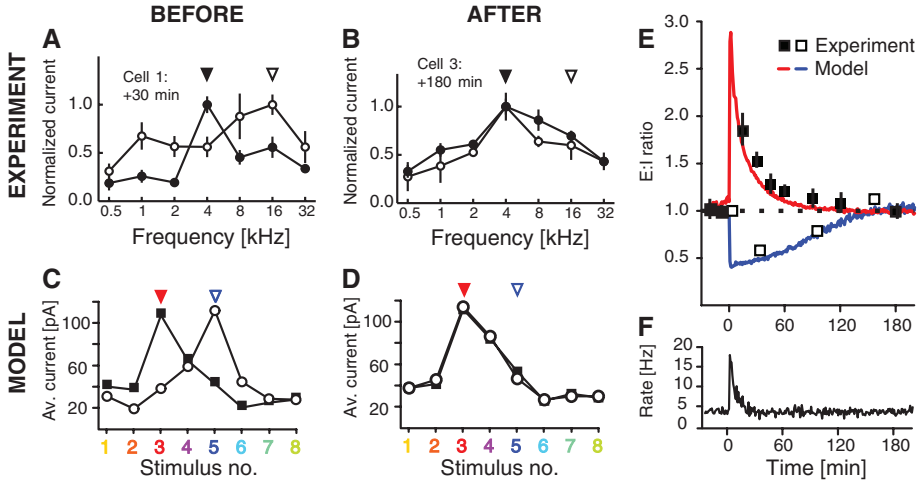

Fig. 3. Temporal dynamics of inhibitory plasticity, experiment, and model. Frequency-tuned excitatory and inhibitory membrane currents (black and white symbols, respectively) as recorded from pyramidal cells in the primary auditory cortex of adult rat (*7*) (**A**) 30 min and (**B**) 180 min after a stimulus protocol shifted the preferred frequency of the excitatory membrane currents from 16 to 4 kHz. Similarly stimulus-tuned input currents in a simulation (**C**) 30 min and (**D**) 180 min after (manually) changing the excitatory tuning curve. Solid and open arrowheads indicate the previous and the new preferred stimuli in all panels. (**E**) Summary plot of the ratios of excitatory and inhibitory current amplitudes of previously preferred stimuli and new preferred stimuli, as indicated in (A) to (D), in the experiment (open and solid symbols, respectively) and simulations (blue and red lines, respectively). (**F**) Firing rate of the simulated neuron over the time of the simulation in (E). Error bars indicate SEM. [(A), (B), and (E) adapted from (*7*) with permission]

In the detailed balanced state, the response of the cell was sparse (Fig. 2B, bottom) and reminiscent of experimental observations (*16*, *20*–*22*) across many sensory systems. Spikes were caused primarily by transients in the input signals, during which the faster dynamics of the excitatory synapses momentarily overcame inhibition. Sustained episodes of presynaptic firing, on the other hand, caused steady membrane currents that canceled each other and thus failed to evoke a reliable postsynaptic response. Seemingly indifferent to the tuning of the excitatory synapses, each signal contributed an equal part to the PSTH of the output signal, but the effect of the excitatory synaptic weights was uncovered by the steplike input protocol (Fig. 2F). The broad, graded responses (as opposed to all-or-none) to preferred and nonpreferred stimuli (Fig. 2, G and H) were in accord with experimental results (*5*, *7*, *8*, *23*, *24*) and confirm earlier theoretical studies arguing that sharp tuning is not a necessary feature for a sparse sensory representation (*25*, *26*). The sparsity of the response to each signal was a direct consequence of the detailed balance of correlated excitatory and inhibitory synapses as described above, not of the specificity of the tuning curve.

The self-organizing dynamics of inhibitory plasticity imply that the excitatory-inhibitory balance is maintained, even in the presence of ongoing excitatory plasticity (Fig. 3). Experiments (*7*) in which a stimulus alters the frequency tuning of excitatory input currents to pyramidal neurons in rat primary auditory cortex point in a similar direction: The disrupted cotuning of excitatory and inhibitory input currents (Fig. 3A) prompts a compensatory response that subsequently changes the amplitude of the inhibitory input currents. After 180 min, the cell returns to a cotuned state, albeit with a different preferred frequency (Fig. 3B). When we disturbed the cotuning of a simulated neuron in a similar way (Fig. 3C), inhibitory plasticity rebalanced the excitatory input currents (Fig. 3, D and E) and stabilized the output firing rates of the postsynaptic neurons (Fig. 3F). Quantitative agreement with the rebalancing dynamics observed in the experiment (for both synaptic depression and potentiation) was achieved by adjusting η, ρ_0, and the average firing rate of the inhibitory input neurons.

The learning rule for inhibitory synapses does not rely on a feedforward structure to achieve low firing rates. It simply matches excitatory and inhibitory synapses that show correlated activity. We therefore tested whether inhibitory plasticity was able to stabilize the dynamics of recurrent networks. In simulations of such networks (*13*)

with plastic inhibitory synapses that were initially weak (Fig. 4A), the resulting high firing rates and subsequent increase in inhibitory synaptic strengths caused by the plasticity rule indeed produced globally balanced input currents that led to a self-organized AI network state (Fig. 4B) with firing rates between 3 and 15 Hz.

We wondered whether it was possible to introduce associative memories to the stabilized network by strengthening specific excitatory connections within dedicated groups of neurons. First proposed by Hebb (*27*), such "cell assemblies" aim to provide a physiologically plausible explanation of how groups of neurons form a memory. Groups of highly connected neurons have since been successfully embedded into large spiking networks (*28*) and shown to self-sustain their activity without disrupting the global dynamics of the host network (*13*, *29*, *30*), but the parameter space that guarantees stable performance is narrow and tuning is arduous. The question has been raised how useful such memory attractors can be for long-term memory systems if only one of all stored memories can be active at a time, and potentially remains active for long periods of time, broadcasting the stored information into the network (*29*).

Inhibitory plasticity can solve some of these problems. After two arbitrarily chosen groups of excitatory neurons were turned into Hebbian assemblies by strengthening the excitatory connections within the groups fivefold, the assemblies temporarily fired at high rates and raised the background firing rate across the network (Fig. 4C). The resulting increase of coincident spike pairs caused inhibitory plasticity to increase the inhibitory synapses onto neurons in both assemblies until the global AI state was reestablished (Fig. 4D). After the excitatory and inhibitory inputs onto these neurons had been rebalanced, the firing rates of neurons in the cell assemblies became indistinguishable from the rest of the network, despite the imprinted memory traces in the excitatory synapses. Electrophysiological recordings of neuronal activity would thus not reveal the presence of a synaptic memory trace in this state.

Retrieval of previously quiescent memory items could be achieved by momentarily disrupting the balance within a cell assembly, for example, through additional excitatory input. It was sufficient to drive a small fraction of the cells of one assembly to reactivate all cells of that assembly. Notably, the recall was asynchronous and irregular, as indicated by low correlations between neurons and large variability of the interspike intervals (Fig. 4E). Although we embedded two overlapping assemblies into the network, only one was activated. The rest of the network remained nearly unperturbed in the AI state. Unlike traditional attractor networks, both assemblies could also be activated in unison by driving cells of both memories simultaneously (figs. S4 and S5), and their activity decayed to the background state after the stimulus was turned off.

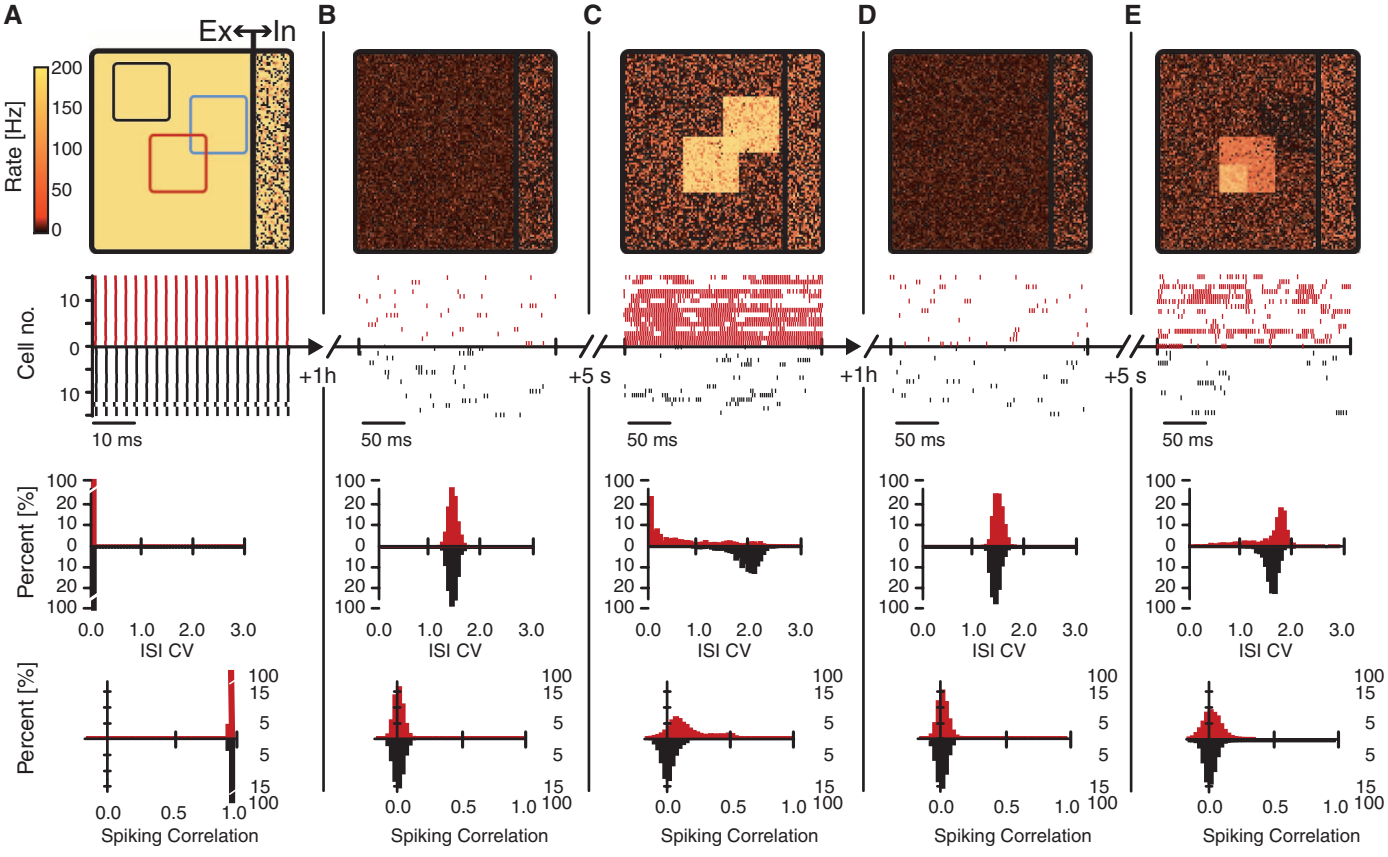

Fig. 4. Inhibitory plasticity in recurrent networks. Five consecutive snapshots of the momentary activity of a network of 10,000 integrate-and-fire cells with inhibitory plasticity. (**A**) Synchronous regular network dynamics with high firing rates at the beginning of the simulation with weak inhibitory synapses. (**B**) Establishment of the AI (steady) state with low firing rates through up-regulation of inhibitory synaptic weights by the synaptic plasticity rule. (**C**) The introduction of two synaptic memory patterns (cell assemblies) by fivefold increased excitatory synaptic weights between neurons outlined in red and blue in (A) leads to high firing rates. (**D**) Recovery of the AI state at low firing rates. (**E**) Memory retrieval through externally driving the lower left quarter of the red cell assembly with an additional excitatory stimulus. Each snapshot (A) to (E) shows (from top to bottom) the following: (i) The momentary (1-s) average firing rate of all neurons on a grid of 100^2 cells and separated into excitatory and inhibitory cells [left and right of the vertical line in (A), respectively]. Three groups of neurons play the role of either a cell assembly (red and blue outlines) or a control group (black outline). (ii) A raster plot of 30 randomly drawn neurons from one (red) cell assembly and the control group, indicated by a red and a black square in the plot above. (iii) The distributions of coefficients of variation of interspike intervals (ISI CVs) recorded from the neurons in the red and black groups. (iv) The distributions of spiking correlations between spike trains from neurons in the same designated groups. For methods and additional statistics, please see SOM.

Our results offer an explanation for how long-term memories can be stably embedded into networks as quiescent and overlapping Hebbian assemblies. Unlike previous studies, our network does not exhibit the behavior of an attractor network, in which activated cell assemblies will compete with each other and the winning pattern often exhibits persistent elevated activity. Instead, the network remains quiet unless the balance of one or more assemblies is modulated in favor of the excitation and returns to the background state when the modulation is turned off. We have shown this effect here by driving a subset of cells with an external stimulus, but there are several conceivable methods to modulate the balance of excitation and inhibition (SOM). The possibility to activate several patterns simultaneously allows the analog combination of patterns into larger composite memories. The capacity of storable and retrievable patterns is likely to depend on complex interactions between dynamics, size, and connectivity of the assemblies and the host network, as well as several other parameters.

We show that a simple, Hebbian plasticity rule on inhibitory synapses leads to robust and self-organized balance of excitation and inhibition that requires virtually no fine-tuning (figs. S6 to S9) and captures an unexpected number of recent experimental findings. The precision of the learned balance depends on the degree of correlation between the excitatory and the inhibitory inputs to the cell, ranging from a global balance in the absence of correlated inputs to a detailed balance for strong correlations. The phenomenon is robust to the shape of the learning rule, as long as it obeys two fundamental requirements: Postsynaptic activity must potentiate activated inhibitory synapses, whereas in the absence of postsynaptic firing inhibitory synapses must decay. Because the balance is self-organized, inhibitory plasticity will most likely maintain balance also in the presence of excitatory plasticity, as long as excitation changes more slowly than inhibition or when excitatory plasticity events are rare.

The mammalian brain hosts a wide variety of inhibitory cell types with different synaptic time scales, response patterns, and morphological target regions. Presumably, these cell types serve different functions, and consequently their synapses may obey several different plasticity rules (*31*). In our simplified model, the dynamics of inhibitory plasticity powerfully contributes to the functional state of cortical architectures and may have a strong impact on cortical coding schemes.

References and Notes

1. N. Brunel, *J. Comput. Neurosci.* **8**, 183 (2000).
2. C. van Vreeswijk, H. Sompolinsky, *Science* **274**, 1724 (1996).
3. M. Tsodyks, T. Sejnowski, *Network Comput. Neural Syst.* **6**, 111 (1995).
4. A. Renart *et al.*, *Science* **327**, 587 (2010).
5. M. Wehr, A. M. Zador, *Nature* **426**, 442 (2003).
6. M. Okun, I. Lampl, *Nat. Neurosci.* **11**, 535 (2008).
7. R. C. Froemke, M. M. Merzenich, C. E. Schreiner, *Nature* **450**, 425 (2007).
8. J. de la Rocha, C. Marchetti, M. Schiff, A. D. Reyes, *J. Neurosci.* **28**, 9151 (2008).
9. B. K. Murphy, K. D. Miller, *Neuron* **61**, 635 (2009).
10. Y. Shu, A. Hasenstaub, D. A. McCormick, *Nature* **423**, 288 (2003).
11. T. P. Vogels, L. F. Abbott, *Nat. Neurosci.* **12**, 483 (2009).
12. W. Gerstner, *Neural Comput.* **12**, 43 (2000).
13. T. P. Vogels, L. F. Abbott, *J. Neurosci.* **25**, 10786 (2005).
14. A. Kumar, S. Rotter, A. Aertsen, *Nat. Rev. Neurosci.* **11**, 615 (2010).
15. J. Cafaro, F. Rieke, *Nature* **468**, 964 (2010).
16. T. Hromádka, M. R. Deweese, A. M. Zador, *PLoS Biol.* **6**, e16 (2008).
17. M. A. Woodin, K. Ganguly, M. M. Poo, *Neuron* **39**, 807 (2003).
18. V. Kilman, M. C. W. van Rossum, G. G. Turrigiano, *J. Neurosci.* **22**, 1328 (2002).
19. K. Hartmann, C. Bruehl, T. Golovko, A. Draguhn, *PLoS One* **3**, e2979 (2008).
20. M. R. DeWeese, M. Wehr, A. M. Zador, *J. Neurosci.* **23**, 7940 (2003).
21. H. Yao, L. Shi, F. Han, H. Gao, Y. Dan, *Nat. Neurosci.* **10**, 772 (2007).
22. S. Crochet, J. F. Poulet, Y. Kremer, C. C. Petersen, *Neuron* **69**, 1160 (2011).
23. L. M. Aitkin, D. J. Anderson, J. F. Brugge, *J. Neurophysiol.* **33**, 421 (1970).
24. I. O. Volkov, A. V. Galazjuk, *Neuroscience* **43**, 307 (1991).
25. P. Seriès, P. E. Latham, A. Pouget, *Nat. Neurosci.* **7**, 1129 (2004).
26. J. Beck, V. R. Bejjanki, A. Pouget, *Neural Comput.* **23**, 1484 (2011).
27. D. Hebb, *The Organization of Behavior; a Neuropsychological Theory* (Wiley-Interscience, New York, 1949).
28. W. Gerstner, R. Ritz, J. L. van Hemmen, *Biol. Cybern.* **69**, 503 (1993).
29. D. J. Amit, N. Brunel, *Cereb. Cortex* **7**, 237 (1997).
30. A. Renart, R. Moreno-Bote, X.-J. Wang, N. Parga, *Neural Comput.* **19**, 1 (2007).
31. M. A. Woodin, A. Maffei, *Inhibitory Synaptic Plasticity* (Springer, New York, 2010).

Acknowledgments: Research was supported by Swiss National Science Foundation grant no. 200020 13287 (Coding Characteristics) and CRSIK0 122697 (Sinergia). Additionally, T.P.V. was supported by the European Community's Seventh Framework Marie Curie International Reintegration grant no. 268436, and H.S. and F.Z. by the European Community's Seventh Framework Program under grant agreement no. 243914 (BRAIN-I-NETS) and 237955 (FACETS-ITN), respectively. C.C. received additional support from a French National Science grant ANR-08-SYSC-005. Thanks to G. Hennequin and A. Woodruff for helpful discussions.

Supporting Online Material

www.sciencemag.org/cgi/content/full/science.1211095/DC1
Materials and Methods
SOM Text
Figs. S1 to S10
Tables S1 and S2
References (*31*–*49*)

13 July 2011; accepted 20 October 2011
Published online 10 November 2011;
10.1126/science.1211095

Autophagy-Dependent Anticancer Immune Responses Induced by Chemotherapeutic Agents in Mice

Mickaël Michaud,[1,2,3]* Isabelle Martins,[1,2,3]* Abdul Qader Sukkurwala,[1,2,3] Sandy Adjemian,[1,2,3] Yuting Ma,[2,3,4,5] Patrizia Pellegatti,[6] Shensi Shen,[1,2,3] Oliver Kepp,[1,2,3] Marie Scoazec,[2,7] Grégoire Mignot,[8,9] Santiago Rello-Varona,[1,2,3] Maximilien Tailler,[1,2,3] Laurie Menger,[1,2,3] Erika Vacchelli,[1,2,3] Lorenzo Galluzzi,[1,2,3] François Ghiringhelli,[8,9] Francesco di Virgilio,[6] Laurence Zitvogel,[2,3,4,5]† Guido Kroemer[1,2,10,11,12]†

Antineoplastic chemotherapies are particularly efficient when they elicit immunogenic cell death, thus provoking an anticancer immune response. Here we demonstrate that autophagy, which is often disabled in cancer, is dispensable for chemotherapy-induced cell death but required for its immunogenicity. In response to chemotherapy, autophagy-competent, but not autophagy-deficient, cancers attracted dendritic cells and T lymphocytes into the tumor bed. Suppression of autophagy inhibited the release of adenosine triphosphate (ATP) from dying tumor cells. Conversely, inhibition of extracellular ATP-degrading enzymes increased pericellular ATP in autophagy-deficient tumors, reestablished the recruitment of immune cells, and restored chemotherapeutic responses but only in immunocompetent hosts. Thus, autophagy is essential for the immunogenic release of ATP from dying cells, and increased extracellular ATP concentrations improve the efficacy of antineoplastic chemotherapies when autophagy is disabled.

Transplantable or primary murine cancers respond to chemotherapy with anthracyclines or oxaliplatin much more efficiently when they grow in syngenic immunocompetent mice than in immunodeficient hosts (*1*, *2*). Similarly, clinical studies indicate that severe lymphopenia negatively affects the chemotherapeutic response of solid cancers (*3*), and immune defects are negative predictors of the response to chemotherapy with anthracyclines or oxaliplatin (*2*, *4*, *5*). Apparently, some successful chemotherapeutics can induce a type of tumor cell stress and death that is immunogenic (*6*–*8*), implying that the patient's dying cancer cells serve as a therapeutic vaccine that stimulates an antitumor immune response, which in turn can control residual cancer cells (*9*, *10*). Immunogenic cell death is characterized by the preapoptotic exposure of calreticulin (CRT) on the cell surface (*11*), postapoptotic release of the chromatin-binding protein high mobility group B1 (HMGB1) (*2*), and secretion of adenosine triphosphate (ATP) (*4*).