

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

Tim P Vogels^{1,2} & L F Abbott¹

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 receiver¹⁷. 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 interneurons¹⁸. 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

¹Center for Neurobiology and Behavior, Department of Physiology and Cellular Biophysics, Columbia University College of Physicians and Surgeons, New York, New York, USA. ²Volen 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

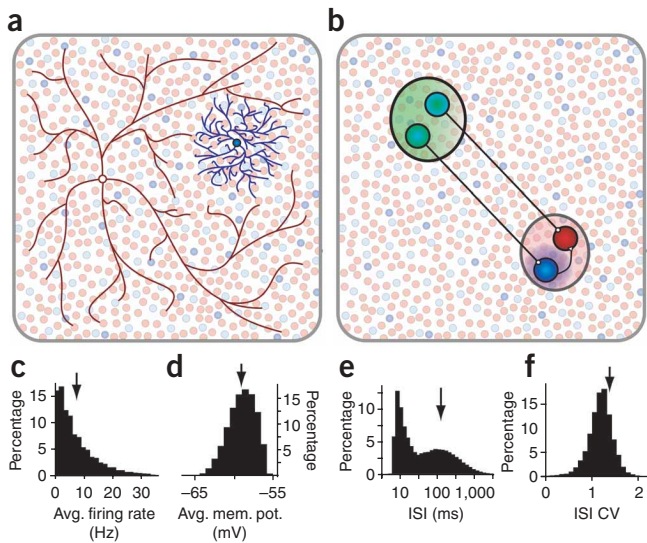


Figure 1 Network connectivity and properties. (a) All excitatory and 65% of the inhibitory neurons are connected randomly with a connection probability of 2% (red). The other 35% of the inhibitory neurons (blue) have local connectivity, targeting their nearest neighbors. (b) An embedded signal pathway 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.

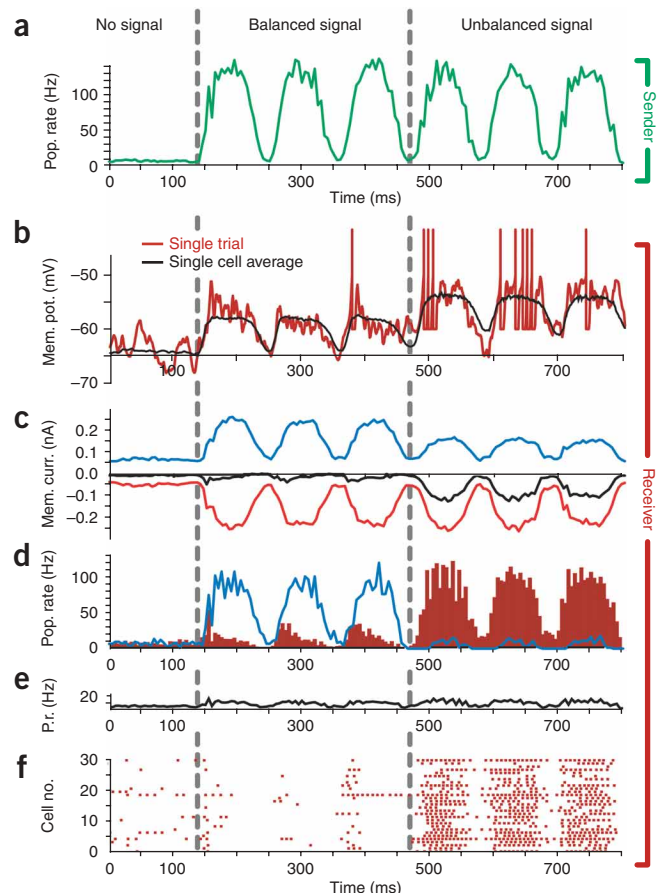
in the sender area to fire in a manner that mimics the input signal (Fig. 2a). In the balanced state, excitation from the sender network, a simple oscillatory signal in the example of Figure 2, activates both excitatory and inhibitory subpopulations in the receiver region. The resulting inhibitory activity (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 region¹⁹. 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 neurons²⁵, corresponds to changing the input–output transfer function so that the same synaptic current generates a smaller response. Although the examples we show focus

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.

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



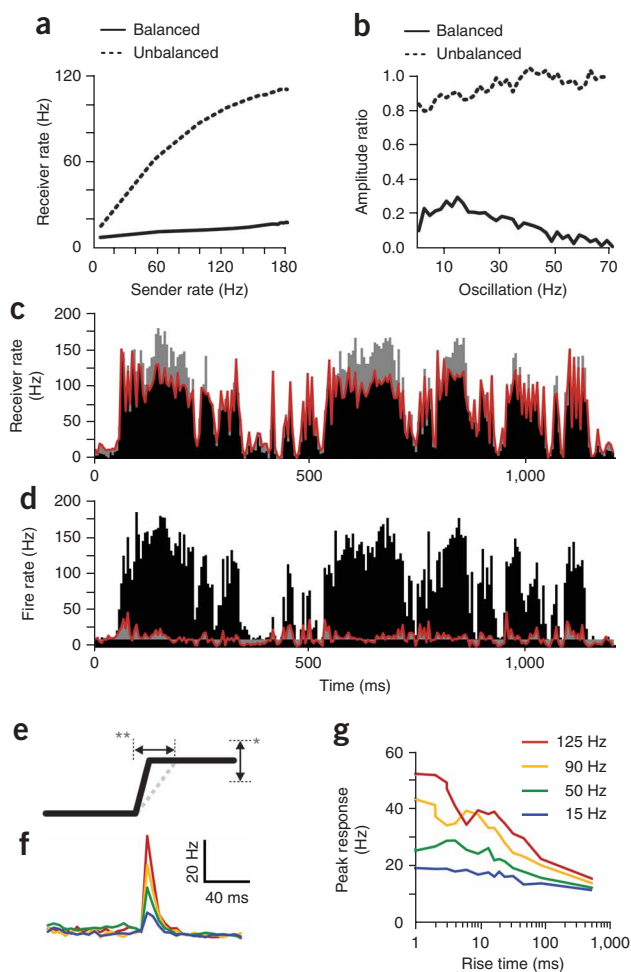


Figure 3 Response analysis. (a) Firing rates of the excitatory receiver neurons as a function of different constant sender firing rates, in the balanced (solid trace) and unbalanced (dashed trace) states. (b) Ratio of receiver to sender excitatory firing-rate oscillation amplitudes at different oscillation 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).

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 constant¹²) 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 layers¹⁴. 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%.

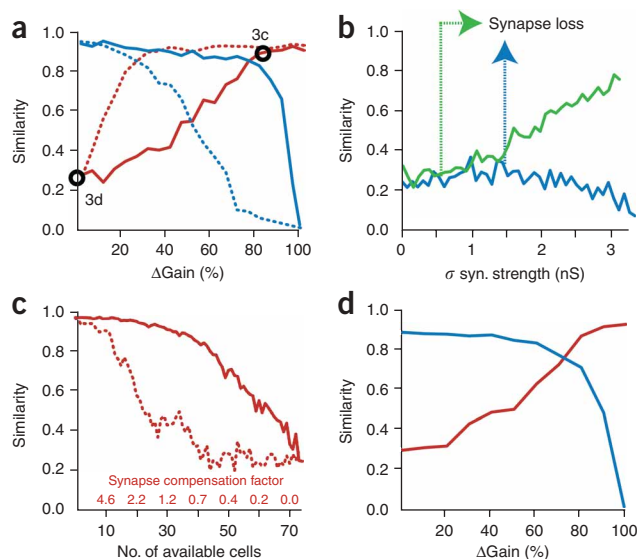
A close inspection of the responses in the balanced state (Fig. 3d) shows that detailed balance does not completely cancel signals when input firing rates change rapidly. Rapid changes in the signal can evoke a response in the excitatory receiver neurons before inhibition can balance excitation because 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

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

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.



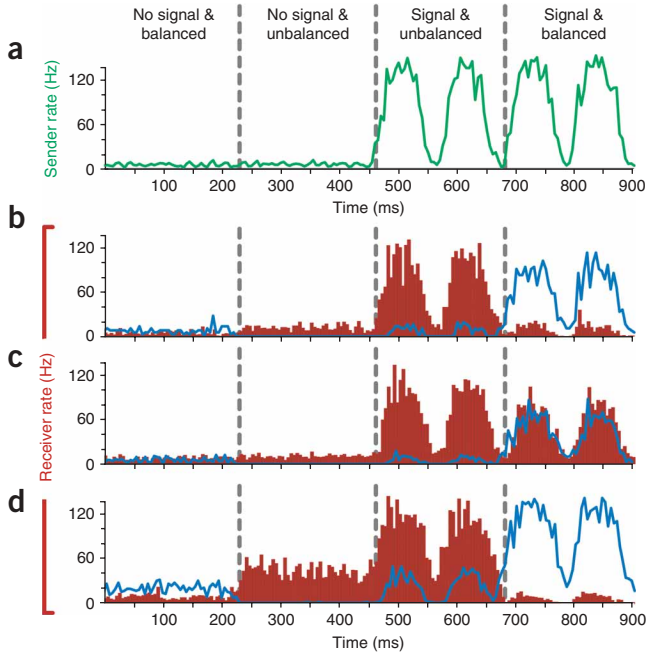


Figure 5 Network pathologies. (a) Average firing rate of the sender neurons without and with an oscillatory input. (b–d) Responses of excitatory (red histogram) and inhibitory (blue trace) receiver neurons with correct tuning (b); weakened local inhibition, leading to a gating deficit (c); or a hyperactive 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 ($\Delta\text{Gain} = 0$), 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 $\Delta\text{Gain} \sim 80\%$, 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

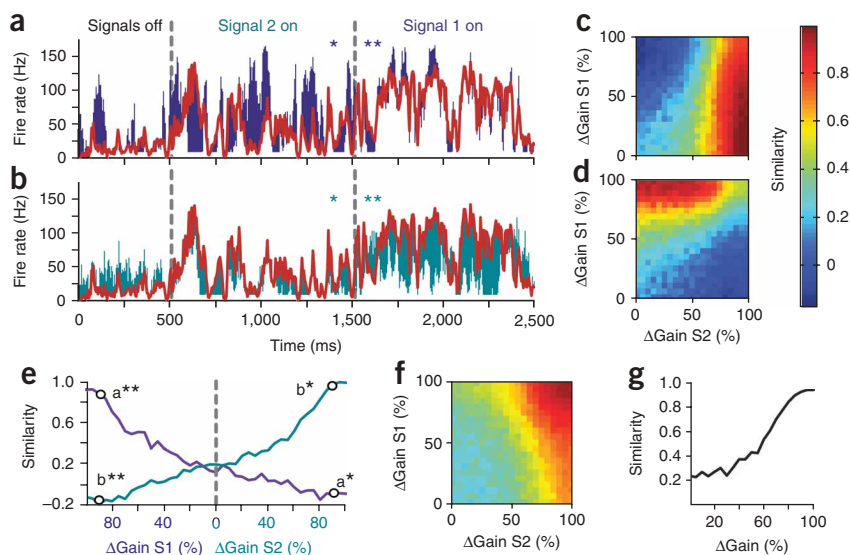


Figure 6 Gating two signals in a network. **(a, b)** Average response of the excitatory receiver subnetwork to two simultaneously delivered signals (S1 and S2). Colored areas indicate difference between the average firing rate in the receiver region (red) and either S1 (purple, **a**) or S2 (green, **b**). 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 circuits²⁶.

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

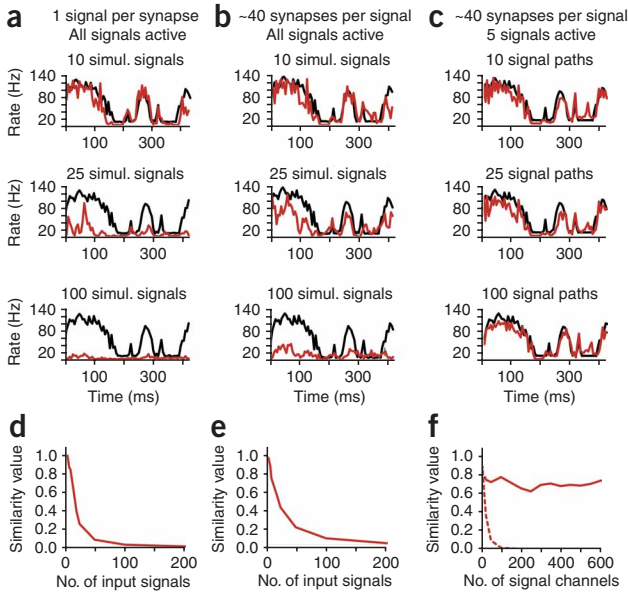


Figure 7 Multiple signals into a single cell. (a–c) Firing rates of the output signal (red), averaged over 200 runs, compared to the input signal (black). First row: 10 simultaneous (simul.) signal paths. Second row: 25 simultaneous signals paths. Third row: 100 simultaneous signal paths. (a) Each afferent 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 canceled 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 respond²⁷. 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 |
| Δg_{inh}^{global} | 7.5 nS | $\Delta g_{inh \rightarrow ex}^{receiver\ 1 \rightarrow receiver}$ | 3.1 | 3.1* | 3.1 | 3.3 |
| | | $\Delta g_{inh \rightarrow ex}^{receiver\ 2 \rightarrow receiver}$ | | | | 4.0 |
| | | $\Delta g_{inh \rightarrow ex}^{net \rightarrow receiver\ 1}$ | 1.25 | 1.25 | 1.25 | 1.32 |
| | | $\Delta g_{inh \rightarrow ex}^{net \rightarrow receiver\ 2}$ | | | | 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 activities³¹. 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 manner³². 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 schizophrenia³⁵. 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 anomalies³⁸. Although these latter anomalies correspond to hypofunction of NMDA conductances, this is associated with a hyperexcitability of the affected circuits³⁹, 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 $100 \text{ M}\Omega$.

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 locally⁴² (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 addition 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 propagation¹⁴. 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 $\Delta g_{\text{ex}}^{\text{global}} = 0.8$ nS, $\Delta g_{\text{inh}}^{\text{local}} = 1.5$ nS and $\Delta g_{\text{inh}}^{\text{global}} = 7.5$ 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 $\Delta g_{\text{ex} \rightarrow \text{ex}}^{\text{sender} \rightarrow \text{receiver}} = 0.9$ 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 $\Delta g_{\text{inh} \rightarrow \text{ex}}^{\text{receiver} \rightarrow \text{receiver}} = 4.65$ 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 $\Delta g_{\text{inh} \rightarrow \text{inh}}^{\text{net} \rightarrow \text{receiver}} = 9.4$ 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 $\Delta g_{\text{inh}}^{\text{local}} = 0.6$ 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 $\Delta g_{\text{ex} \rightarrow \text{ex}}^{\text{net} \rightarrow \text{receiver}} = 1.28$ nS and $\Delta g_{\text{ex} \rightarrow \text{inh}}^{\text{net} \rightarrow \text{receiver}} = 1.28$ 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 $g_{\text{ex}} = 0.014$ nS and $\Delta g_{\text{inh}} = 0.044$ 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 $a = 1, 2, \dots, M - 1$. 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 $\alpha = \text{pinv}(\tilde{B})\tilde{C}$, where $\text{pinv}(\tilde{B})$ 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 $g_{\text{ex}} \rightarrow g_{\text{ex}} + \Delta g_0$. 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 similarity¹⁴. 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. Jercoc 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/>

- Shu, Y., Hasenstaub, A. & McCormick, D.A. Turning on and off recurrent balanced cortical activity. *Nature* **423**, 288–293 (2003).
- 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. *Corticons: 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).