

Gating of Signal Propagation in Spiking Neural Networks by Balanced and Correlated Excitation and Inhibition

Jens Kremkow, Ad Aertsen, and Arvind Kumar

Bernstein Center Freiburg, Neurobiology and Biophysics, Faculty of Biology, University of Freiburg, 79104 Freiburg, Germany

Both ongoing and natural stimulus driven neuronal activity are dominated by transients. Selective gating of these transients is mandatory for proper brain function and may, in fact, form the basis of millisecond-fast decision making and action selection. Here we propose that neuronal networks may exploit timing differences between correlated excitation and inhibition (temporal gating) to control the propagation of spiking activity transients. When combined with excitation–inhibition balance, temporal gating constitutes a powerful mechanism to control the propagation of mixtures of transient and tonic neural activity components.

Introduction

Appropriate action selection is one of the fundamental problems that the brain needs to solve in its interaction with the environment. This problem is even enhanced by the intrinsically parallel architecture of the brain. A mechanism based on winner-take-all dynamics is often suggested for this purpose (Ermentrout, 1992). However, winner-take-all dynamics require strong recurrent inhibition, which may not be available in every brain region (Tepper et al., 2004). Moreover, winner-take-all dynamics are typically slow if the competing signals are of comparable strength, and may even result in unpredictable responses in the presence of transient synchrony within and among the competing inputs (Ermentrout, 1992). Another approach to action selection could be based on the gating of signals, such that only one (or few) effectively propagate to the next processing stage. Thus, recently Vogels and Abbott (2009) demonstrated that detailed balance of excitation and inhibition (amplitude gating) can efficiently gate the propagation of firing rates. However, while this mechanism allows for a rapid selection among competing rate inputs, it fails to control activity transients on a timescale of up to a few tens of milliseconds [Vogels and Abbott (2009), their Fig. 3f]. In a natural stimulus environment, such activity transients frequently occur as a consequence of stimulus-driven changes in firing rates or by the emergence of correlations in the network (Buracas et al., 1998; Bar-Yosef and Nelken, 2007; Butts et al., 2007; Hromádka et al., 2008; Haider et al., 2010). Therefore, reliable gating of activity transients is a prerequisite of proper brain function.

Here we propose an alternative gating mechanism that exploits temporal differences between correlated excitation and in-

hibition (Okun and Lampl, 2008; Gentet et al., 2010). We show that small differences between the timing of excitation and inhibition (referred to temporal gating from hereon) can act as a modulator of neuronal integration time and, thereby, serve as a highly selective gate for signal transients. The proposed mechanism allows for both gating at a millisecond timescale and reliable control of onset transients and correlated rate inputs. When combined with amplitude gating, temporal gating constitutes a powerful mechanism to control the propagation of mixtures of transient and tonic neural activity components.

Materials and Methods

Models

Neurons

The neurons in the network were modeled as leaky integrate-and-fire neurons, with the subthreshold dynamics of the membrane potential $V^i(t)$ in neuron i described by the following equation:

$$C \frac{d}{dt} V^i(t) + G_{\text{rest}}[V^i(t) - V_{\text{rest}}] = I_{\text{syn}}^i,$$

where I_{syn}^i is the total synaptic input current into neuron i , and C and G_{rest} denote the passive electrical properties of its membrane at rest (V_{rest}). When the membrane potential reached a fixed spike threshold V_{th} above rest, a spike was emitted, the membrane potential was reset to its resting value, and synaptic integration was halted for 2 ms, mimicking the refractory period in real neurons. The parameters used in the simulations were as follows: $C = 290$ pF, $G_{\text{rest}} = 29$ nS, $V_{\text{rest}} = -70$ mV, and $V_{\text{th}} = -57$ mV (Destexhe et al., 1998; Muller et al., 2007).

Network

We simulated a network of 28,125 neurons, with 22,500 excitatory and 5625 inhibitory neurons (ratio 4:1). The excitatory neurons were arranged on a Cartesian grid of 150×150 , the inhibitory neurons on a grid of 75×75 . Both grids represented the same cortical space of 1×1 mm². To avoid boundary effects, the network was folded as a torus. The recurrent network was sparsely connected with a connection probability of ~5% ($K_{\text{exc} \rightarrow \text{exc}} = 1120$, $K_{\text{exc} \rightarrow \text{inh}} = 280$, $K_{\text{inh} \rightarrow \text{exc}} = 1120$, $K_{\text{inh} \rightarrow \text{inh}} = 280$). As an exception to this rule, excitatory neurons that participated in the signal path received less than 5% connection from the background neurons, to maintain the same in-degree. The local connection probabil-

Received July 26, 2010; revised Sept. 22, 2010; accepted Sept. 23, 2010.

This work was supported by the German Federal Ministry of Education and Research (BMBF Grant 01GQ0420 to Bernstein Center for Computational Neuroscience Freiburg) and the 6th Research Framework Programme of the European Union (Grant 15879-FACETS). We thank Tim Vogels, Benjamin Staude, Ioannis Vlachos, Guillaume S. Masson, and Laurent U. Perrinet for helpful discussions.

Correspondence should be addressed to Arvind Kumar at the above addresses. E-mail: arvind.kumar@biologie.uni-freiburg.de.

DOI:10.1523/JNEUROSCI.3874-10.2010

Copyright © 2010 the authors 0270-6474/10/3015760-09\$15.00/0

ity was described as a function of their cell body distance using a Gaussian profile (Hellwig, 2000). The width of the Gaussian describing the connection probability σ_c was set to 0.6 mm for the excitatory and $\sigma_i = 0.1$ mm for the inhibitory connections (Stepanyants et al., 2008) (Fig. 1*a*). In addition to the local connections from within the network, each neuron received excitatory connections from outside, mimicking both long-range connections within the same cortical area and input from other areas in cortex (Braitenberg and Schüz, 1991; Kumar et al., 2008*a*). The external inputs were modeled as uncorrelated Poisson-type spike trains. The rate and strength of external inputs (1500 excitatory synapses per neuron, each at ~ 2 Hz) were adjusted to obtain a low-firing-rate (~ 3 Hz), asynchronous (pairwise correlation ≈ 0.01), and irregular activity state in the network (Brunel, 2000; Kumar et al., 2008*b*).

Synapses

Synaptic inputs were modeled as transient conductance changes, using exponential functions with $\tau_{exc} = 1.5$ ms and $\tau_{inh} = 10$ ms (Kuhn et al., 2004; Muller et al., 2007). Excitatory and inhibitory synaptic delays in the embedding network were set to 2 ms. The weights (peak conductance) of the excitatory synapses targeting excitatory neurons (J_{EE}) were set to 0.5 nS, such that a postsynaptic potential of ~ 0.15 mV was elicited at rest (Matsumura et al., 1996; Bruno and Sakmann, 2006; Kumar et al., 2008*a*). The weight of excitatory synapses targeting inhibitory neurons (J_{EI}) was set to 1 nS (Cruikshank et al., 2007; Kremkow et al., 2010). The synaptic weights of all inhibitory synapses (J_{IP} , J_{IE}) were set to a 0.5 nS conductance value. This set of synaptic strengths and external input (see above) generated both a realistic effective membrane time constant of ~ 5 ms (Anderson et al., 2000; Destexhe et al., 2003; Léger et al., 2005) and a realistic low-firing-rate, asynchronous-irregular network activity regime (Brunel, 2000; Kumar et al., 2008*b*; Ecker et al., 2010).

Embedding of a feedforward network as signal path

To construct a signal path, we embedded a feedforward network, containing three successive groups of neurons (sender, gate, and receiver groups), into the recurrent network. The sender group was composed solely of 100 excitatory neurons, whereas the gate and receiver group contained 100 excitatory and 25 inhibitory neurons each (Fig. 1*b*). The inhibitory neurons provided the balancing inhibition in the signal path and controlled the effective integration time of the excitatory neurons. The three groups were spatially separated to avoid cross talk between groups (Kumar et al., 2008*a*) (Fig. 1*b*). The neurons constituting a group were randomly chosen from a local pool of 300 excitatory and 75 inhibitory neurons. The center of the neuron pool defined the physical location of the group in the network. Each neuron in the signal path received convergent inputs from 60 randomly selected excitatory neurons of the preceding group. The neurons in the gate and receiver groups received excitatory synapses from 1060 neurons ($<5\%$) of the background network to maintain uniform in-degree in the network. The delay of excitatory synapses in the signal path was set to 5 ms (Aertsen et al., 1996; Diesmann et al., 1999). This choice of the excitatory delays of 5 ms is somewhat arbitrary. These delays are supposed to mimic the latency between different neighboring brain regions, and such delays can be in the order of up to 10 ms. Further, we point out that these delays do not influence the dynamics of the embedding network and the embedded feedforward network, and can be reduced to any arbitrary value without influencing the gating. In the default state of the signal path, the delay of inhibitory synapses onto the excitatory neurons within a group was set to 2 ms (Cruikshank et al., 2007; Kremkow et al., 2010).

Embedding multiple signal paths

To construct two signal paths that merged in the gate group, we embedded the first signal path as described above. The sender group of the second signal path was then added. It was spatially separated from the first sender group to avoid cross talk. The excitatory neurons in the gate and receiver groups were common in both signal paths. The inhibitory neurons of the gate group in both signal paths were randomly chosen from a local pool of 75 inhibitory neurons. Thus, both signal paths might share a small fraction of inhibitory neurons in the gate group (see Fig. 5*a*).

Stimuli

We used three different stimulus classes to study signal propagation and gating in the signal path.

Pulse packet input. The input to the sender group was a volley of spikes with α spike times drawn from a Gaussian distribution of width σ (Aertsen et al., 1996; Diesmann et al., 1999). The number of synapses from the presynaptic neurons to the sender group was set to 60, the same as within the signal path.

Uncorrelated rate input. We used a Poisson process to stimulate the sender group with uncorrelated firing rate input (Vogels and Abbott, 2009). The onset time of the stimulus was set to 500 ms to mimic a sudden activation of the signal path (Butts et al., 2007). The rate of the Poisson process was set to 200 Hz to induce a strong onset transient.

Correlated rate input. To introduce correlation in the firing rate input, we used the multiple-interaction process (MIP) (Kuhn et al., 2003). This process introduces correlations by copying spikes from a Poisson “mother process” into a desired number of “children” processes. The copy probability (c) defines the correlation among the children processes. As in the uncorrelated rate input, the onset time constant of the correlated rate stimulus was set to 500 ms. The rate of the MIP population was set to 20 Hz to illustrate that even at low firing rates, spike correlations [$c = 0.5$, e.g. (see Fig. 3*d*)] could induce strong transients.

Data analysis

Event characterization

To characterize the spiking activity of each group and, thus, the α and σ of the pulse packet sent to the next group, we measured “events,” i.e., distinct, short bursts of spikes in the peristimulus time histogram (PSTH) of all excitatory neurons in a group. The spike count of all excitatory neurons in an event described the response strength α . Likewise, σ described the temporal spread of the spike times, that is, the time difference between the start and end of the “event.” Thus, the propagation of activity in the signal path was determined by the response of the three groups of the signal path, quantified by α and σ .

Signal propagation

To determine whether a signal, composed of multiple pulse packets, could successfully propagate from the sender to the receiver, we measured the similarity between the spiking activities in both groups. The similarity was calculated by cross-correlating the PSTHs of the excitatory neurons in both groups. In case of successful propagation, the signal in the receiver group should have a similar temporal sequence of peaks in the PSTH. However, due to the synaptic delay in the signal path, this sequence would be shifted in time. Therefore, when calculating the cross-correlation between sender and receiver, we accounted for this time shift.

Simulation tools

All network simulations were written in Python (<http://www.python.org>) using PyNN (Davison et al., 2008) (<http://neuralensemble.org/trac/PyNN>) as interface to the simulation environment NEST (Morrison et al., 2005; Gewaltig and Diesmann, 2007; Eppler et al., 2008) (<http://www.nest-initiative.org>). The dynamical equations were integrated at a fixed temporal resolution of 0.1 ms. Simulation management was performed using the Python package NeuroTools (<http://neuralensemble.org/trac/NeuroTools>). Simulation data were analyzed in Python using the scientific libraries SciPy (<http://www.scipy.org>) and NumPy (<http://numpy.scipy.org>) and visualized using the plotting library Matplotlib (<http://matplotlib.sourceforge.net>).

Results

To systematically study the concept of temporal gating, we embedded a signal path in a large-scale recurrent cortical network model (Fig. 1*a,b*). The signal path was a feedforward network consisting of three successive groups of neurons. The first group, composed of excitatory neurons activated by the stimulus, is referred to as the sender. The second and third groups contained both excitatory and inhibitory neurons, with the inhibitory neurons providing disinaptic inhibition to the excitatory neurons

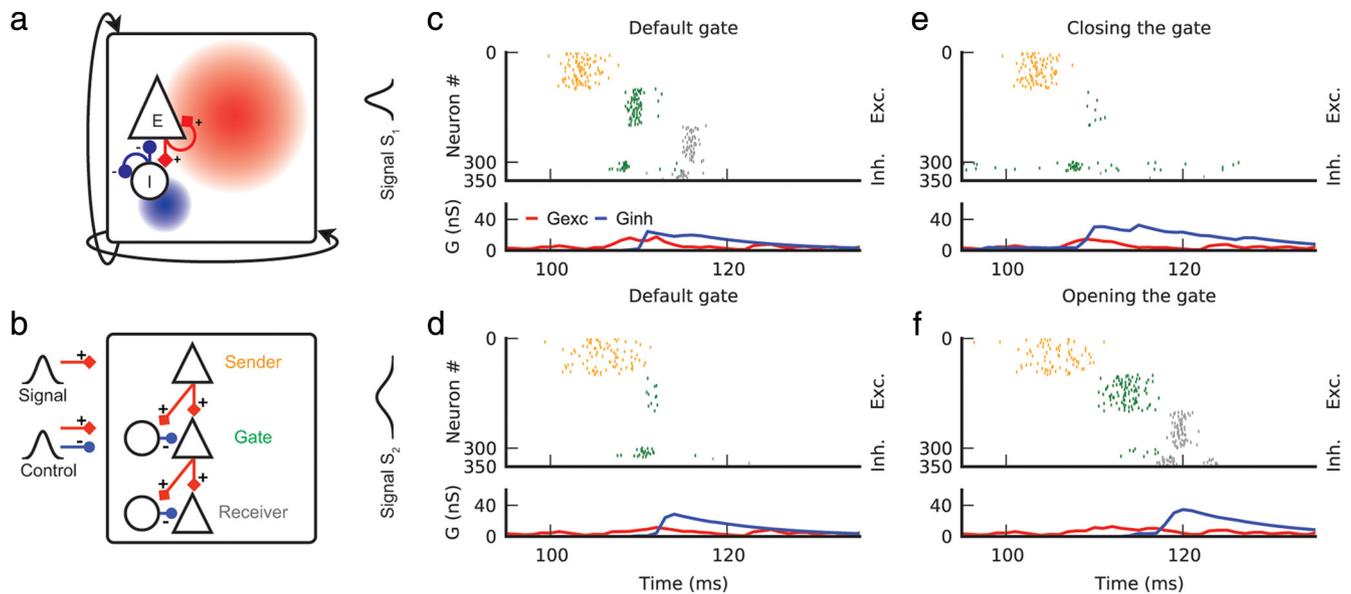


Figure 1. Temporal gating of spiking activity by delayed correlation between excitation and inhibition. **a**, Scheme of cortical network model. Twenty-eight thousand, one hundred twenty-five neurons [80% excitatory (Δ) and 20% inhibitory (\circ)] were arranged as a two-dimensional sheet. Space constants of excitatory and inhibitory neurons are shown as red and blue shaded regions, respectively. **b**, A signal path (sender, gate, receiver) was embedded in the network. The sender neuron group transmitted an input signal via the gate group toward the receiver group. A depolarizing or hyperpolarizing input to the gate was used as control signal. **c, d**, Top, Spiking activity of the sender, gate, and receiver groups. Bottom, Excitatory and inhibitory synaptic conductances (Gexc, Ginh) for a representative excitatory neuron in the gate group. **c**, A strong ($\alpha = 60$), synchronous ($\sigma = 3.5$ ms) signal S_1 propagated through the gating group and successfully activated the receiver group. **d**, By contrast, a less synchronous ($\sigma = 7$ ms) signal S_2 failed to propagate. **e**, Propagation of S_1 could be blocked by advancing inhibitory spiking in the gate group, thereby reducing Δt of the excitatory neurons. **f**, Propagation of S_2 could be rescued by delaying (or abolishing) inhibitory spiking in the gate group, thereby increasing Δt of the excitatory neurons.

within the same group. The third group is referred to as receiver, whereas the second group, which is in a strategic position to influence propagation from sender to receiver, is referred to as the gate (Fig. 1*b*). However, the gating mechanism is flexible and in a feed-forward network with more than three groups, any group can be considered as a gate group between its preceding and succeeding partner groups. The embedding network was tuned into an asynchronous-irregular state to provide realistic background activity (Kumar et al., 2008*a*) (cf. Materials and Methods).

Temporal gating of signal propagation

To illustrate the mechanism of temporal gating, we stimulated the sender group with a volley of spikes (pulse packet; cf. Materials and Methods), with α spike times drawn from a Gaussian distribution of width σ .

We first considered the propagation of a synchronous stimulus S_1 (small σ). In the default state of the signal path (cf. Materials and Methods), S_1 generated a transient response in the sender, which successfully propagated through the gate to the receiver (Fig. 1*c*). This propagation can be described in terms of the synfire attractor, a defining feature of feedforward network dynamics (Diesmann et al., 1999; Kumar et al., 2008*a*, 2010; Kremkow et al., 2010), with the fate of signal propagation being determined by the properties of the pulse packet stimulus. The (α, σ) state space is subdivided by a separatrix (schematically shown in Fig. 2*a*, blue line). If the pulse packet starts above the separatrix (class *P* stimulus), the signal converges into a fixed point (FP) and propagates in a stable manner (Figs. 1*c*, 2*a*). If, by contrast, the pulse packet starts below the separatrix (class *F* stimulus), the signal progressively dies out, indicating failure to propagate (Figs. 1*d*, 2*a*). Thus, the location of the separatrix is the critical parameter that controls the propagation of a stimulus (Diesmann et al., 1999; Kumar et al., 2008*a*, 2010; Kremkow et al.,

2010). Consequently, altering the location of the separatrix could change a class *P* stimulus into a class *F* stimulus (or vice versa), thereby providing a simple mechanism for gating the propagation of more or less transient spiking activity (Fig. 2*a*). Various properties of the feedforward network, such as group size, connection probability, and weights determine the location of the separatrix in the (α, σ) space, which, however, can also be modulated by the dynamical properties of the feedforward network, such as noise level of the embedding background activity (Diesmann et al., 1999; Gewaltig et al., 2001). Furthermore, we recently showed that the delay between excitation and inhibition within one group changes the slope of the separatrix by effectively altering the integration time (Δt) of the neurons (Kremkow et al., 2010). Interestingly, recent data from *in vivo* measurements of the delay between excitation and inhibition (Okun and Lampl, 2008) are in the right range to alter the slope of the separatrix in an appreciable way. Several biologically plausible mechanisms could lead to dynamic variations in the excitation/inhibition delay and, thereby, to associated modulations of Δt . For instance, neuromodulators could change Δt by differentially altering neuronal excitability or synaptic strengths (Kruglikov and Rudy, 2008; Antal et al., 2010). Similarly, stimulus properties could introduce variations in Δt (Zhang et al., 2003; Wilent and Contreras, 2005).

We exploited the control over the delay between excitation/inhibition to close the gate, i.e., to block the propagation of the S_1 -type stimulus as shown in Figure 1*c*. Thus, we reduced the temporal delay between excitation and inhibition in the gating group by injecting a small depolarizing control pulse into the inhibitory neurons. As a result, the gate closed: the response in the gate group fell below the separatrix, causing activity in the receiver group to fade (Figs. 1*e*, 2*a*). To further explore the temporal gating mechanism, we systematically varied Δt by altering the inhibitory synaptic delay, while keeping the excitatory delay fixed. For different values of Δt , Figure 2*b* shows α and σ for

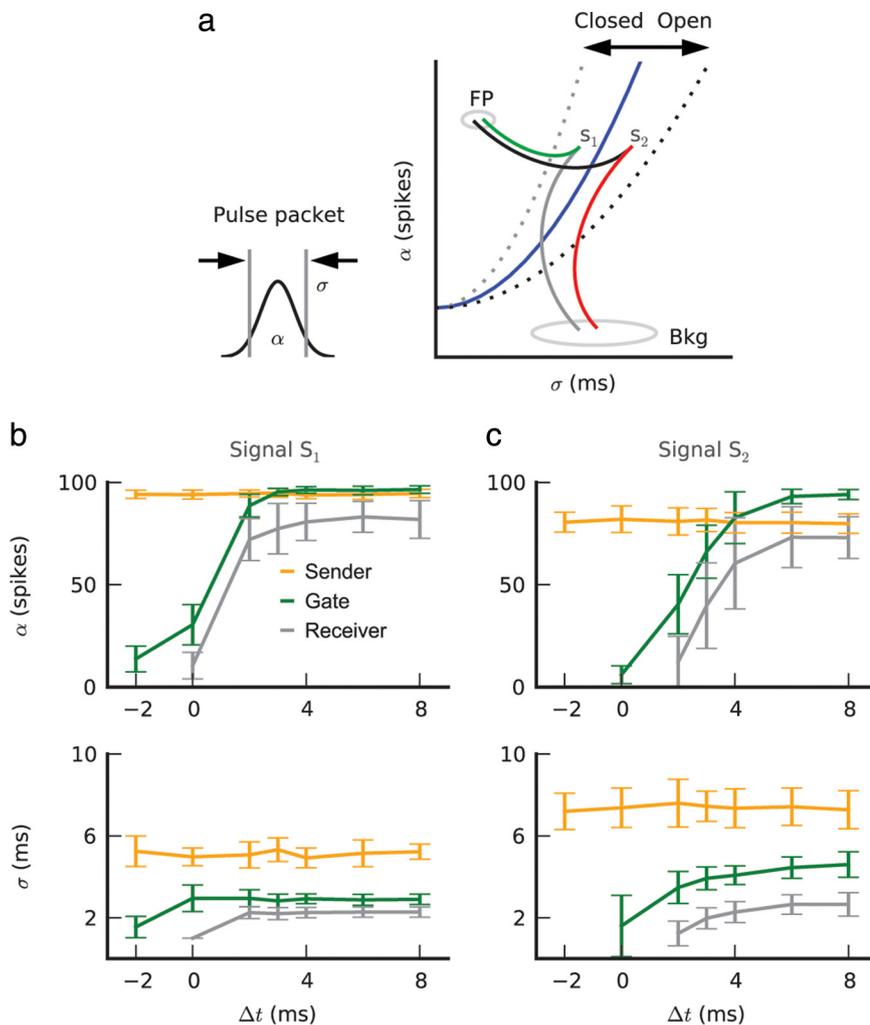


Figure 2. Temporal delay between excitation and inhibition controls signal gating. **a**, Scheme of the pulse packet propagation state space, which is subdivided by a separatrix (Diesmann et al., 1999) (blue line). If the pulse packet starts above the separatrix (S_1), the signal converges into a fixed point (FP) and propagates in a stable manner (green trajectory). If, by contrast, the pulse packet starts below the separatrix (S_2), the signal fails to propagate and fades into the background (Bkg.) (red trajectory). When the effective integration time Δt is reduced, the separatrix moves up (gray dotted line), thus blocking the propagation of S_1 (gray trajectory). By contrast, when the integration time is increased, the separatrix moves down (black dotted line), thus allowing the propagation of S_2 (black trajectory). **b**, α and σ in sender, gate, and receiver groups for the S_1 stimulus (class P) as a function of effective integration time (Δt). Error bars represent SD ($n = 20$). In this example, S_1 successfully propagates in the default state ($\Delta t = 2$ ms). Further increase in Δt does not influence signal propagation. However, a small decrease in Δt (~ 1 ms) suffices to completely block signal propagation. **c**, Same as **b** for stimulus S_2 (class F). In this example, S_2 does not propagate in the default state ($\Delta t = 2$ ms). Propagation can be rescued by increasing Δt . Note that, in contrast to **b**, the prolonged opening of the gate causes activity in the gate to be higher than in the receiver. Also, for increasing integration time (larger Δt), signal propagation becomes more reliable, which is expressed as a reduction of trial-by-trial variability in the receiver group.

the three excitatory neuron groups in the signal path. For small integration time ($\Delta t \leq 1$ ms in this example), the gate response was weak and fell below the separatrix. Consequently, the response in the receiver group was indistinguishable from the background network activity. However, increasing the integration time ($\Delta t \approx 2$ ms in this example) shifted the gate activity above the separatrix, inducing a strong response (high α , small σ) in the receiver group. The same mechanism can be used to open the gate, i.e., to facilitate propagation of class F stimuli, which would otherwise fail to propagate (Fig. 1*f*). To test this, we applied a hyperpolarizing (rather than depolarizing) control pulse to the inhibitory neurons, thereby delaying or even suppressing their spiking activity. As a result, the integration time was increased, such that more excitatory neurons in the gate group responded. Its activity thereby rose

above the separatrix and, hence, ensured a strong response (high α , small σ) in the receiver group (Fig. 1*f*). A systematic increase in the integration time Δt of the gate neurons resulted in a monotonic increase in both α and σ of the gate activity (Fig. 2*c*). Due to the synfire attractor, the activity of the receiver group was further synchronized. Note that, in contrast to Figure 2*b*, the prolonged opening of the gate caused activity in the gate to be higher than in the receiver. Also, it is noteworthy that for increasing integration time (higher Δt), signal propagation became more reliable, as expressed in a reduction in trial-by-trial variability in the receiver group (Fig. 2*c*, top panel).

In summary, the sigmoidal relationship between Δt and the activity in the gate and receiver groups demonstrates the effectiveness of the temporal gating (Fig. 2*b,c*).

As there is an ongoing debate (Kuhn et al., 2004; Vogels and Abbott, 2005; Kumar et al., 2008*a,b*) on whether synapses in a point neuron model are best described by conductance transients (as done here) or current transients [the more common choice, e.g., Brunel (2000)], we resimulated the network model with current-based synapses. The sigmoidal relationship between Δt and the activity in the gate and receiver groups remained unchanged when synapses were modeled as current transients (compare supplemental Fig. S1, available at www.jneurosci.org as supplemental material, and Fig. 2*b,c*), demonstrating that the “temporal gating” does not depend on the choice of synapse or input activity model.

Gating of signals with strong transients

Next we tested whether temporal gating is also effective in controlling the propagation of signal onset transients. Such signal onset transients arise when the input signal is abruptly changed (Deger et al., 2009). To mimic sharp changes in the input, the sender group, we stimulated the sender group with Poisson-type spiking activity, starting at time $t = 500$ ms (Fig.

3*a*). In the default state of the gate group, with balanced excitation and inhibition, the onset transient propagated to the receiver group, whereas the tonic part of the input was blocked (Fig. 3*b*). The control of such transients, however, is critical because in a dynamic natural environment, such transients are likely to occur frequently (Buracas et al., 1998; Haider et al., 2010). In the default state of the gate, the effective integration time is large enough for the transients to elicit a response in the downstream neurons (gate group). Thus, an effective way to control the flow of transients is to reduce the integration time of the downstream neurons (e.g., by temporal gating) (see Figs. 1, 2). Indeed, when the delay between excitation and inhibition in the gate group was reduced, while still keeping the overall excitation and inhibition in balance, it was possible to control the propagation of both the

stimulus transient and its tonic component (Fig. 3*c*). Note that in this particular example, when the gate is closed (Fig. 3*c*), only the inhibitory population in the gate group is active, while the activity of the excitatory neurons is completely suppressed by the balancing inhibitory input. Reducing the gain of the inhibitory neurons would result in a biologically more realistic residual spiking activity in the excitatory neurons in the gate group. Transients in the spiking activity in neural systems may also arise due to spike time correlations among neurons (Alonso et al., 1996; Dan et al., 1998; Chen et al., 2006). The emergence of such correlations could be due to large-scale fluctuations in ongoing network activity (Arieli et al., 1996; Kenet et al., 2003) or stimulus-induced network activity (Bar-Yosef and Nelken, 2007; Alonso et al., 2008; Hromádka et al., 2008). To generate a signal composed of tonic and transient components, we used the MIP model (Kuhn et al., 2003) (Fig. 3*d*; cf. Materials and Methods). Once again, when the sender group was stimulated by a MIP-type input signal, the control of the activity was imprecise and signal transients reflecting the spike correlations propagated uninhibited (Fig. 3*e*). However, reducing the delay between excitation and inhibition at the gate group provided an effective control over the propagation of signal transients (Fig. 3*f*). In summary, Figure 3, *b* and *e*, illustrates the limitations of the Vogels and Abbott mechanism in gating of both onset transients (Fig. 3*b*) and transients induced by correlations (Fig. 3*e*). Both types of transients were fast enough to pass before the delayed feedforward inhibition could quench them. By contrast, Figure 3, *c* and *f*, shows the improved effectiveness of the temporal gating mechanism in controlling the propagation of these activity transients.

Combining temporal and amplitude gating

The effectiveness of temporal gating in controlling the propagation of activity transients has its price: it precludes the propagation of sustained (tonic) activity devoid of transients (Fig. 3*c,f*). While the importance of such sustained average neural activity can be argued, in view of the ubiquitous fluctuations in the ongoing activity and dynamic natural environment, it is of general interest to know whether it is possible to do both, to control the flow of activity transients, while allowing tonic activity to propagate. Because temporal gating is most effective in controlling the propagation of activity transients, whereas detailed balance (or amplitude gating) can be used to control the propagation of tonic activity, a combination of the two might provide a general mechanism to control the propagation of arbitrary neural activity patterns. Therefore, we examined whether a combination of temporal gating and amplitude gating would be possible within the same network.

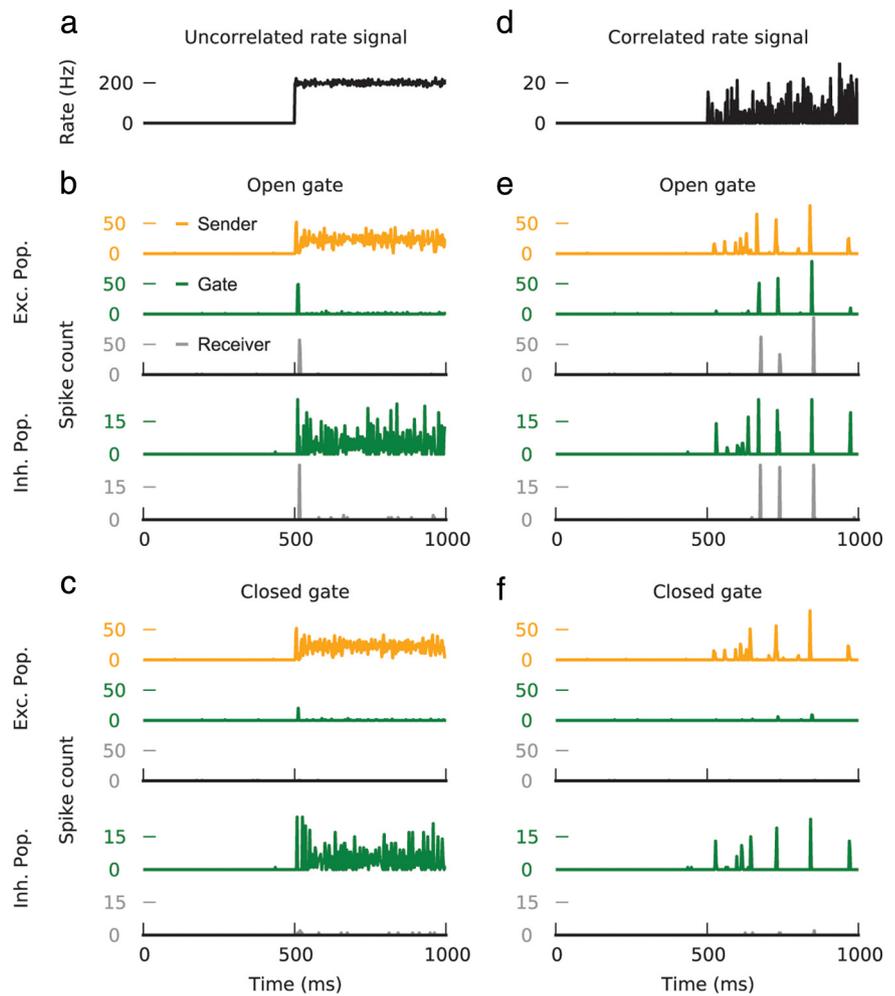


Figure 3. Effectiveness of gating tonic activity and activity transients. *a*, Uncorrelated rate stimulus as input to the signal path. *b*, Population activity in the excitatory (top) and inhibitory (bottom) neurons in the signal path. In the default state ($\Delta t = 2$ ms), gating by balancing excitation and inhibition successfully blocks propagation of the tonic component of the stimulus, but fails to block the stimulus onset transient. *c*, By reducing the effective integration time of the excitatory neurons ($\Delta t \leq 1$ ms), temporal gating effectively blocks both the onset transient and the tonic stimulus component. *d*, Firing rate stimulus with transient input correlations. *e*, As in *b*, gating by balancing excitation and inhibition fails to block activity transients, here induced by transient input correlations. *f*, As in *c*, temporal gating effectively blocks activity transients.

For this we stimulated the signal path with uncorrelated Poisson activity, starting at time $t = 500$ ms. This type of stimulus contains both a transient onset and a tonic phase, which makes it optimal for probing the effectiveness of both gating mechanisms. The amplitude gating mechanism was implemented by varying the strength of the inhibition in the gate group, thereby controlling the balance of excitation and inhibition in the excitatory neurons of the gate group (Vogels and Abbott, 2009). To this end, we varied the synaptic strength from the excitatory neurons in the sender group onto the inhibitory population (J_{EI}), while keeping the synaptic strength from the excitatory neurons onto the excitatory neurons (J_{EE}) fixed. Thus, we define the inhibitory gain factor as J_{EI}/J_{EE} . In all examples shown so far (i.e., Figs. 1–3), the inhibitory gain factor was 2. This gain factor ensured an effective feedforward inhibition in the neurons of the gate group (cf. Materials and Methods) (Cruikshank et al., 2007; Kremkow et al., 2010). Figure 4 shows the results of systematically varying both the degree of temporal gating and amplitude gating. Simultaneous closing of both the temporal ($\Delta t = -2$ ms) and amplitude (inhibitory gain = 2) gates blocked both the transient and the tonic component in the gate group (Fig. 4*a–c*). Opening the tem-

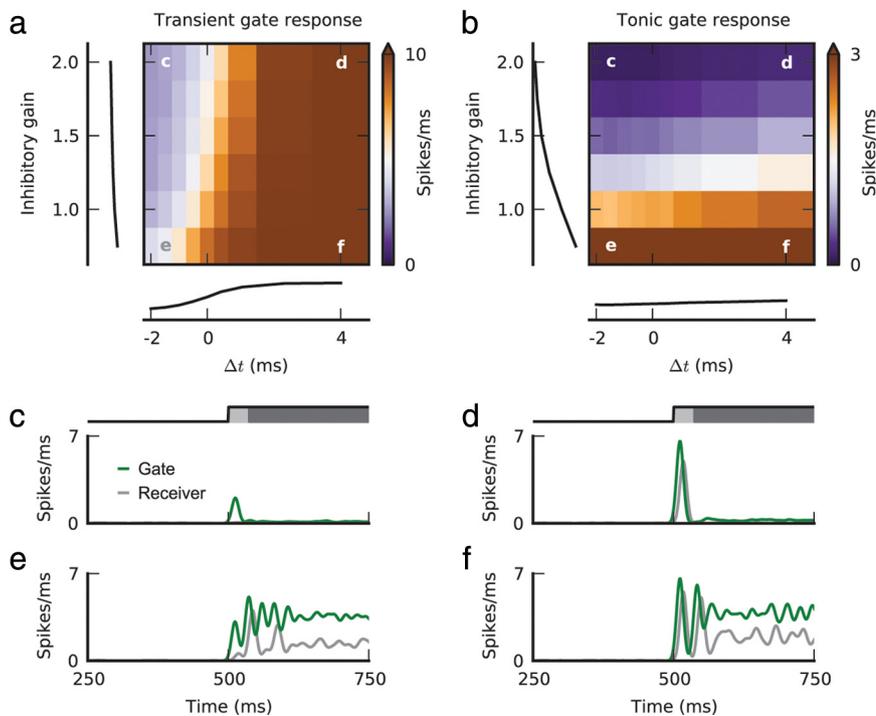


Figure 4. Combining temporal and amplitude gating. **a**, Transient response (first 10 ms, marked light gray in **c** and **d**) in the gate as a function of effective integration time Δt (temporal gating) and inhibitory gain (amplitude gating) with uncorrelated rate stimulus as input to the signal path. Mean transient response in the gate does not change as a function of inhibitory gain for all values of Δt , indicating that inhibitory gain cannot gate the propagation of transient component of the activity (left). By contrast, mean transient response in the gate as a function of Δt for all values of inhibitory gain shows a sigmoidal shape, indicating that Δt can control the propagation of transient component of the activity (bottom). **b**, Tonic response (excluding first 10 ms, marked dark gray in **c** and **d**) in the gate as a function of effective integration time Δt and inhibitory gain. Mean tonic response in the gate as a function of inhibitory gain for all values of Δt , varies in a threshold linear manner, indicating that inhibitory gain can control the propagation of tonic component of the activity (left). By contrast, mean tonic response in the gate does not change as a function of Δt for all values of inhibitory gain, indicating that Δt cannot gate the tonic component of the activity (bottom). **c–f**, Population activity of the excitatory neurons in the gate and receiver group during four distinct gating modes, marked with corresponding letters in **a** and **b**. **c**, Simultaneously closing of the temporal and the amplitude gates blocks both the onset transient and the tonic component, resulting in a very small activity in the receiver group. **d**, Opening the temporal gate, while keeping the amplitude gate closed, allows the onset transient to propagate. **e**, Opening the amplitude gate, while keeping the temporal gate closed, allows the tonic component to propagate. **f**, Both the transient and the tonic component are propagated when both gates are open.

poral gate ($\Delta t = 4$ ms), while keeping the amplitude gate closed (inhibitory gain = 2), allowed only the onset transient to induce activity in the gate group (Fig. 4*a,b,d*). These two examples, thus, show the effectivity of temporal gating at high inhibitory gain (compare Figs. 3 and 4*c,d*). By contrast, opening the amplitude gate (e.g., inhibitory gain = 0.75) resulted in tonic activity in the gate group (Fig. 4*a,b,e,f*). Here, the weaker feedforward inhibition could not balance the direct excitation during the tonic phase (Kremkow et al., 2010). However, during the transient phase the inputs from the sender group arrived synchronously, activating the inhibitory neurons in the gate group simultaneously, thereby providing sufficient balancing inhibition for the temporal gating to function and block the transient [compare closed temporal gate (Fig. 4*e*) and open temporal gate (Fig. 4*f*)]. Note that using inhibitory gain as a gating variable may result in an attenuated response in the receiver group. However, because a high signal-to-noise ratio is maintained, this reduced activity could easily be amplified in a subsequent stage.

In summary, Figure 4 shows that amplitude gating is effective in gating tonic component of the activity (Fig. 4*b*), while temporal gating is effective in gating the activity transients (Fig. 4*a*). When the two gating mechanisms are implemented within the same network,

together they provide a powerful mechanism to control the propagation of a wide variety of temporal activity profiles.

Gating of multiple signals

So far we considered how isolated transients or those that form part of stimulus-evoked or ongoing activity modulations propagate along a single signal path. However, in a modular system like the brain, action selection can be thought of in terms of directing one of multiple signals to a next processing stage or, alternatively, as directing a particular signal to one of multiple subsequent processing stages (Bienenstock, 1995). In both these scenarios, the brain needs a mechanism to control the propagation of activity over multiple (and interacting) signal pathways. Here, we test the effectiveness of temporal gating in selecting between two neuronal activity signals that are statistically identical but uncorrelated, and dominated by transients.

Figure 5*a* shows the scheme of two signal paths, embedded in a recurrent random network (cf. Materials and Methods). This arrangement of sender, gate, and receiver groups is similar to that used in Figures 1–4, except that here two spatially separated sender groups (Sender_X and Sender_Y) can send signals (S_X , S_Y) to the gate. To control the propagation of either signal S_X or S_Y , we systematically varied the delay between excitation and inhibition from the respective sender to the gate (Δt_X and Δt_Y , respectively). To quantify the propagation of the two signals, we measured the cross-correlations (ρ_X and ρ_Y) between the receiver group output and each of the two input signals (cf. Materials and Methods). Figure 5*b* shows the difference between the two: $\rho = \rho_X - \rho_Y$ as a function of both Δt_X and Δt_Y . Positive values of ρ indicate that the receiver group output is more similar to the input signal S_X , and negative values reflect a higher similarity with input signal S_Y . When ρ is close to zero, the receiver output is equally correlated (including uncorrelated) with both input signals; in this case, the scenario is evidently not suitable for action selection.

When the two signals were presented simultaneously, selective transmission of signal S_X (Fig. 5*b,c*) could be obtained by increasing Δt_X and decreasing Δt_Y (effectively closing signal path S_Y) (compare Fig. 2). Similarly, selective transmission of signal S_Y (Fig. 5*b,f*) could be obtained by increasing Δt_Y and decreasing Δt_X (closing signal path S_X). The exact value of Δt_X to gate signal S_X depended on the value of Δt_Y (Fig. 5*b*) and vice versa. When the gate was closed for both signal paths (Fig. 4*e*), neither of the two input signals propagated to the receiver and, hence, ρ was close to zero. On the other hand, when the gate was opened for both signal paths (Fig. 5*d*), both signals propagated and the response of the receiver group was equally correlated to both input signals, again resulting in ρ being close to zero.

Thus, the temporal gating mechanism appears well suited for selective gating of multiple signals to a next processing stage.

Similarly, temporal gating can also be used to selectively direct a particular signal to one of several next processing stages (not shown here).

Discussion

In a modular system like the brain, appropriate gating or control over activity propagation is critical for information processing and action selection (Bienenstock, 1995; Kumar et al., 2010). In fact, for various biological reasons, gating of neural activity emerges as a superior strategy over winner-take-all dynamics for action selection (Ermentrout, 1992; Tepper et al., 2004). In any case, a key requirement for an appropriate gating mechanism is that it should be able to control neural activity with arbitrary statistical properties.

In particular, it is important to be able to gate neuronal activity that is dominated by strong transients, which are ubiquitous in a dynamic sensory environment, as has been reported under various conditions *in vivo* in all sensory pathways. For example, it has been demonstrated in the rodent whisker-barrel system that cortex is primarily driven by transient synchronous activity of thalamic neurons (Bruno and Sakmann, 2006). Similarly, neurons with overlapping receptive fields of the same polarity (on- or off-center) in the cat visual thalamus show spike correlation in the millisecond range and, consequently, drive postsynaptic neurons in the visual cortex much more effectively (Alonso et al., 1996). Furthermore, transient subthreshold dynamics have been measured intracellularly in the cat visual cortex *in vivo* during presentation of natural stimuli (Haider et al., 2010; Y. Fregnac, personal communication). Likewise, subthreshold membrane potential *in vivo* in the rat auditory cortex showed large transients, induced by synchronous synaptic input, occurring irregularly in an otherwise quiet background (DeWeese and Zador, 2006). Together, these observations in different sensory pathways in several species corroborate the importance of activity transients. Consequently, a proper gating mechanism should be able to cope with these transients.

Recently, Vogels and Abbott (2009) proposed a gating mechanism based on the detailed balance of excitation and inhibition to control the propagation of neuronal activity in feedforward networks. This gating strategy, however, is only effective in controlling the propagation of tonic components of neuronal activity modulation—transients “escape” from such gating and propagate to the next stage in any case. This is because the mechanism fails in gating transients that are fast enough to pass the small time window, induced by the lagging inhibition, i.e., in the range of some tens of milliseconds [Vogels and Abbott (2009), their Fig. 3f]. As a consequence, gating based only on detailed balance of excitation and inhibition is truly limited in its applicability.

The temporal gating mechanism proposed here extends this idea by changing the delay between excitation and inhibition, and

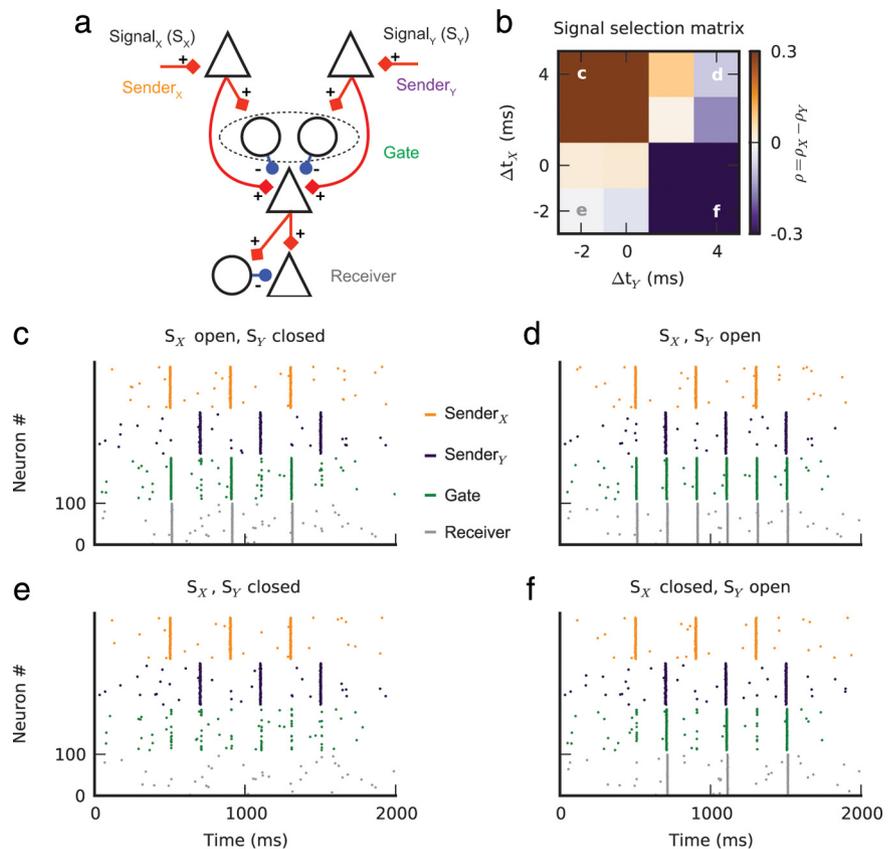


Figure 5. Selective gating of multiple signals. **a**, Scheme of two interacting signal paths. **b**, Selectivity of signal propagation in two interacting signal paths. The color map shows the selectivity measure $\rho = \rho_x - \rho_y$ as a function of effective integration time Δt_x and Δt_y . Positive values (brown color) of ρ indicate that the receiver group output is more similar to signal S_x , and negative values (blue color) reflect a high similarity with signal S_y (i.e., selective gating of either S_x or S_y). At $\rho \approx 0$, the receiver output is equally correlated with both signals (i.e., no selective gating). **c–f**, Spiking activity of the sender, gate, and receiver groups during four distinct gating modes, marked with corresponding letters in **b**. **c**, With the gate open for S_x and closed for S_y , the activity in the receiver group resembles the signal S_x . **d**, With both gates open, the activity in the receiver group is a superposition of S_x and S_y , thereby reducing ρ to zero. **e**, Closing both gates blocks the propagation of both signals. **f**, With the gate open for S_y and closed for S_x , the activity in the receiver group resembles the signal S_y .

thereby obtaining the means to control the temporal scale of transients that can pass or, alternatively, are blocked. The delay between excitation and inhibition affects the neurons integration time (Pouille and Scanziani, 2001; Cruikshank et al., 2007). Thereby, it can be used to block or facilitate activity propagation. Thus, this gating mechanism presents an example of exploiting a single neuron-level property to modulate network behavior—in this case, the location of the state space separatrix—such as to control the propagation of spiking activity within the network. As our study showed, the temporal gating mechanism is a simple yet powerful mechanism to control the propagation of activity transients. Thus it is complementary to the amplitude gating mechanism proposed by Vogels and Abbott (2009). When the two mechanisms are combined, they can control neural activity with different temporal profiles.

The key parameter of temporal gating is the delay between excitation and inhibition. The delayed correlated inhibition in our model was induced by disynaptic inhibition, a connectivity pattern observed ubiquitously in the CNS [hippocampus (Buzsáki, 1984), thalamocortical connection (Swadlow, 2003; Cruikshank et al., 2007), local cortical network (Kapfer et al., 2007; Silberberg and Markram, 2007), and long-range horizontal cortical network (Hirsch and Gilbert, 1991; Tucker and Katz, 2003a,b)]. In addition, neuromodulators (Kruglikov and Rudy,

2008; Antal et al., 2010) and stimulus properties (Zhang et al., 2003; Wilentz and Contreras, 2005) could change the latency between excitation and inhibition in a neuron specific manner, by differentially altering neuronal excitability or synaptic strengths. Thus, the required neuronal hardware for the proposed temporal gating mechanism is quite general and not restricted to any specific brain area. Furthermore, the gating mechanism does not depend on the choice of the synapse model (current based vs conductance based) (Fig. 2 and supplemental Fig. S1, available at www.jneurosci.org as supplemental material). Moreover, the delay values used here were well within the range measured experimentally *in vivo* (Okun and Lampl, 2008). The effectivity of temporal gating in controlling the propagation of activity transients, however, precludes propagation of sustained (tonic) neural activity devoid of transients (Fig. 3*c,f*). Because detailed balance (or amplitude gating) is suitable in controlling the propagation of tonic activity, a combination of the two mechanisms provides a general solution to control the propagation of arbitrary neural activity patterns.

Here we presented a possibility of combining the two mechanisms and showed that, indeed, by comodulating excitation–inhibition balance and their latencies, it is possible to control the propagation of neural activity for combinations of transient and tonic components (Fig. 4). In addition, more sophisticated mechanisms might be invoked. For instance, dynamic modulation of excitation–inhibition balance can be achieved with depressing or facilitating synapses. Thus, introducing synapses with short-term plasticity (Klyachko and Stevens, 2006) into our temporal gating mechanism might be an alternative way of combining the strengths of the two gating mechanisms.

Gating of transients is most effective when the activity transients are close to the separatrix (Fig. 2*a*). If, by contrast, the activity transients are close to the stable fixed point (FP) (Fig. 2*a*) or too similar to the background fluctuations (Bkg) (Fig. 2*a*), temporal gating may not be as effective. Thus, if temporal coding is indeed used in neuronal systems, it should be possible to predict the spectrum of neural activity transients, given the properties of synapses and the connection probabilities between processing stages. In a similar vein, knowledge of the neural activity transients may provide important hints concerning the connectivity between subsequent processing stages.

Preliminary experimental evidence already points toward the involvement of a temporal gating mechanism in decision making and action selection. For instance, animals can make decisions on a millisecond timescale (Yang et al., 2008). Mechanisms based on winner-take-all dynamics would not support such rapid decision making, while temporal gating allows for signal gating and, thereby, action selection (decision making) on a millisecond timescale. Moreover, precisely timed perturbations of the cortical network can systematically alter the behavior of the animal (Seidemann et al., 1998). However, more direct tests of this type of gating mechanism have not been performed thus far. As noted above, the temporal gating method exploits circuit level behavior, which makes this mechanism particularly interesting for *in vivo* experiments. Recent developments in optogenetic techniques, which allow selective activation/inactivation of specific groups of neurons (Zhang et al., 2007), could be used to test our proposal of temporal gating *in vivo* in a highly controlled manner.

References

- Aertsen A, Diesmann M, Gewaltig MO (1996) Propagation of synchronous spiking activity in feedforward neural networks. *J Physiol Paris* 90:243–247.
- Alonso JM, Usrey WM, Reid RC (1996) Precisely correlated firing in cells of the lateral geniculate nucleus. *Nature* 383:815–819.
- Alonso JM, Yeh CI, Staelzel CR (2008) Visual stimuli modulate precise synchronous firing within the thalamus. *Thalamus Relat Syst* 4:21–34.
- Anderson JS, Carandini M, Ferster D (2000) Orientation tuning of input conductance, excitation, and inhibition in cat primary visual cortex. *J Neurophysiol* 84:909–926.
- Antal M, Acuna-Goycolea C, Pressler RT, Blitz DM, Regehr WG (2010) Cholinergic activation of m2 receptors leads to context-dependent modulation of feedforward inhibition in the visual thalamus. *PLoS Biol* 8:e1000348.
- Arieli A, Sterkin A, Grinvald A, Aertsen A (1996) Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* 273:1868–1871.
- Bar-Yosef O, Nelken I (2007) The effects of background noise on the neural responses to natural sounds in cat primary auditory cortex. *Front Comput Neurosci* 1:3.
- Bienenstock E (1995) A model of neocortex. *Netw Comput Neural Syst* 6:179–224.
- Braitenberg V, Schüz A (1991) *Anatomy of the cortex: statistics and geometry*. Berlin: Springer.
- Brunel N (2000) Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *J Comput Neurosci* 8:183–208.
- Bruno RM, Sakmann B (2006) Cortex is driven by weak but synchronously active thalamocortical synapses. *Science* 312:1622–1627.
- Buracas GT, Zador AM, DeWeese MR, Albright TD (1998) Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron* 20:959–969.
- Butts DA, Weng C, Jin J, Yeh CI, Lesica NA, Alonso JM, Stanley GB (2007) Temporal precision in the neural code and the timescales of natural vision. *Nature* 449:92–95.
- Buzsáki G (1984) Feed-forward inhibition in the hippocampal formation. *Prog Neurobiol* 22:131–153.
- Chen Y, Geisler WS, Seidemann E (2006) Optimal decoding of correlated neural population responses in the primate visual cortex. *Nat Neurosci* 9:1412–1420.
- Cruikshank SJ, Lewis TJ, Connors BW (2007) Synaptic basis for intense thalamocortical activation of feedforward inhibitory cells in neocortex. *Nat Neurosci* 10:462–468.
- Dan Y, Alonso JM, Usrey WM, Reid RC (1998) Coding of visual information by precisely correlated spikes in the lateral geniculate nucleus. *Nat Neurosci* 1:501–507.
- Davison AP, Brüderle D, Eppler J, Kremkow J, Müller E, Pecevski D, Perrinet L, Yger P (2008) PyNN: A common interface for neuronal network simulators. *Front Neuroinformatics* 2:11.
- Deger M, Cardanobile S, Helias M, Rotter S (2009) The Poisson process with dead time captures important statistical features of neural activity. *BMC Neuroscience* 10:P110.
- Destexhe A, Contreras D, Steriade M (1998) Mechanisms underlying the synchronizing action of corticothalamic feedback through inhibition of thalamic relay cells. *J Neurophysiol* 79:999–1016.
- Destexhe A, Rudolph M, Paré D (2003) The high-conductance state of neocortical neurons *in vivo*. *Nat Rev Neurosci* 4:739–751.
- DeWeese MR, Zador AM (2006) Non-Gaussian membrane potential dynamics imply sparse, synchronous activity in auditory cortex. *J Neurosci* 26:12206–12218.
- Diesmann M, Gewaltig MO, Aertsen A (1999) Stable propagation of synchronous spiking in cortical neural networks. *Nature* 402:529–533.
- Ecker AS, Berens P, Keliris GA, Bethge M, Logothetis NK, Tolias AS (2010) Decorrelated neuronal firing in cortical microcircuits. *Science* 327:584–587.
- Eppler JM, Helias M, Müller E, Diesmann M, Gewaltig MO (2008) PyNEST: a convenient interface to the NEST simulator. *Front Neuroinformatics* 2:12.
- Ermentrout B (1992) Complex dynamics in winner-take-all neural nets with slow inhibition. *Neural Netw* 5:415–431.
- Gentet LJ, Avermann M, Matyas F, Staiger JF, Petersen CCH (2010) Membrane potential dynamics of GABAergic neurons in the barrel cortex of behaving mice. *Neuron* 65:422–435.
- Gewaltig MO, Diesmann M (2007) Nest (neural simulation tool). *Scholarpedia* 2:1430.
- Gewaltig MO, Diesmann M, Aertsen A (2001) Propagation of cortical synfire activity: survival probability in single trials and stability in the mean. *Neural Netw* 14:657–673.

- Haider B, Krause MR, Duque A, Yu Y, Touryan J, Mazer JA, McCormick DA (2010) Synaptic and network mechanisms of sparse and reliable visual cortical activity during nonclassical receptive field stimulation. *Neuron* 65:107–121.
- Hellwig B (2000) A quantitative analysis of the local connectivity between pyramidal neurons in layers 2/3 of the rat visual cortex. *Biol Cybern* 82:111–121.
- Hirsch JA, Gilbert CD (1991) Synaptic physiology of horizontal connections in the cat's visual cortex. *J Neurosci* 11:1800–1809.
- Hromádka T, DeWeese MR, Zador AM (2008) Sparse representation of sounds in the unanesthetized auditory cortex. *PLoS Biol* 6:e16.
- Kapfer C, Glickfeld LL, Atallah BV, Scanziani M (2007) Supralinear increase of recurrent inhibition during sparse activity in the somatosensory cortex. *Nat Neurosci* 10:743–753.
- Kenet T, Bibitchkov D, Tsodyks M, Grinvald A, Arieli A (2003) Spontaneously emerging cortical representations of visual attributes. *Nature* 425:954–956.
- Klyachko VA, Stevens CF (2006) Excitatory and feed-forward inhibitory hippocampal synapses work synergistically as an adaptive filter of natural spike trains. *PLoS Biol* 4:e207.
- Kremkow J, Perrinet LU, Masson GS, Aertsen A (2010) Functional consequences of correlated excitatory and inhibitory conductances in cortical networks. *J Comput Neurosci* 28:579–594.
- Kruglikov I, Rudy B (2008) Perisomatic GABA release and thalamocortical integration onto neocortical excitatory cells are regulated by neuromodulators. *Neuron* 58:911–924.
- Kuhn A, Aertsen A, Rotter S (2003) Higher-order statistics of input ensembles and the response of simple model neurons. *Neural Comput* 15:67–101.
- Kuhn A, Aertsen A, Rotter S (2004) Neuronal integration of synaptic input in the fluctuation-driven regime. *J Neurosci* 24:2345–2356.
- Kumar A, Rotter S, Aertsen A (2008a) Conditions for propagating synchronous spiking and asynchronous firing rates in a cortical network model. *J Neurosci* 28:5268–5280.
- Kumar A, Schrader S, Aertsen A, Rotter S (2008b) The high-conductance state of cortical networks. *Neural Comput* 20:1–43.
- Kumar A, Rotter S, Aertsen A (2010) Spiking activity propagation in neuronal networks: reconciling different perspectives on neural coding. *Nat Rev Neurosci* 11:615–627.
- Léger JF, Stern EA, Aertsen A, Heck D (2005) Synaptic integration in rat frontal cortex shaped by network activity. *J Neurophysiol* 93:281–293.
- Matsumura M, Chen D, Sawaguchi T, Kubota K, Fetz EE (1996) Synaptic interactions between primate precentral cortex neurons revealed by spike-triggered averaging of intracellular membrane potentials *in vivo*. *J Neurosci* 16:7757–7767.
- Morrison A, Mehring C, Geisel T, Aertsen AD, Diesmann M (2005) Advancing the boundaries of high-connectivity network simulation with distributed computing. *Neural Comput* 17:1776–1801.
- Muller E, Buesing L, Schemmel J, Meier K (2007) Spike-frequency adapting neural ensembles: beyond mean adaptation and renewal theories. *Neural Comput* 19:2958–3010.
- Okun M, Lampl I (2008) Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nat Neurosci* 11:535–537.
- Pouille F, Scanziani M (2001) Enforcement of temporal fidelity in pyramidal cells by somatic feed-forward inhibition. *Science* 293:1159–1163.
- Seidemann E, Zohary E, Newsome WT (1998) Temporal gating of neural signals during performance of a visual discrimination task. *Nature* 394:72–75.
- Silberberg G, Markram H (2007) Disynaptic inhibition between neocortical pyramidal cells mediated by Martinotti cells. *Neuron* 53:735–746.
- Stepanyants A, Hirsch JA, Martinez LM, Kisvárdy ZF, Ferecskó AS, Chklovskii DB (2008) Local potential connectivity in cat primary visual cortex. *Cereb Cortex* 18:13–28.
- Swadlow HA (2003) Fast-spike interneurons and feedforward inhibition in awake sensory neocortex. *Cereb Cortex* 13:25–32.
- Tepper JM, Koós T, Wilson CJ (2004) Gabaergic microcircuits in the neostriatum. *Trends Neurosci* 27:662–669.
- Tucker TR, Katz LC (2003a) Recruitment of local inhibitory networks by horizontal connections in layer 2/3 of ferret visual cortex. *J Neurophysiol* 89:501–512.
- Tucker TR, Katz LC (2003b) Spatiotemporal patterns of excitation and inhibition evoked by the horizontal network in layer 2/3 of ferret visual cortex. *J Neurophysiol* 89:488–500.
- Vogels TP, Abbott LF (2005) Signal propagation and logic gating in networks of integrate-and-fire neurons. *J Neurosci* 25:10786–10795.
- Vogels TP, Abbott LF (2009) Gating multiple signals through detailed balance of excitation and inhibition in spiking networks. *Nat Neurosci* 12:483–491.
- Wilent WB, Contreras D (2005) Dynamics of excitation and inhibition underlying stimulus selectivity in rat somatosensory cortex. *Nat Neurosci* 8:1364–1370.
- Yang Y, DeWeese MR, Otazu GH, Zador AM (2008) Millisecond-scale differences in neural activity in auditory cortex can drive decisions. *Nat Neurosci* 11:1262–1263.
- Zhang F, Wang LP, Brauner M, Liewald JF, Kay K, Watzke N, Wood PG, Bamberg E, Nagel G, Gottschalk A, Deisseroth K (2007) Multimodal fast optical interrogation of neural circuitry. *Nature* 446:633–639.
- Zhang LI, Tan AYY, Schreiner CE, Merzenich MM (2003) Topography and synaptic shaping of direction selectivity in primary auditory cortex. *Nature* 424:201–205.