Gamma oscillations mediate stimulus competition and attentional selection in a cortical network model

Christoph Börgers[a,1], Steven Epstein[b], and Nancy J. Kopell[b,1]

[a]Department of Mathematics, Tufts University, Medford, MA 02155; and [b]Department of Mathematics and Center for Biodynamics, Boston University, Boston, MA

Contributed by Nancy J. Kopell, September 29, 2008 (sent for review May 30, 2008)

Simultaneous presentation of multiple stimuli can reduce the firing rates of neurons in extrastriate visual cortex below the rate elicited by a single preferred stimulus. We describe computational results suggesting how this remarkable effect may arise from strong excitatory drive to a substantial local population of fast-spiking inhibitory interneurons, which can lead to a loss of coherence in that population and thereby raise the effectiveness of inhibition. We propose that in attentional states fast-spiking interneurons may be subject to a bath of inhibition resulting from cholinergic activation of a second class of inhibitory interneurons, restoring conditions needed for gamma rhythmicity. Oscillations and coherence are emergent features, not assumptions, in our model. The gamma oscillations in turn support stimulus competition. The mechanism is a form of "oscillatory selection," in which neural interactions change phase relationships that regulate firing rates, and attention shapes those neural interactions.

cholinergic modulation | selective attention | gamma rhythm

Cortical gamma frequency (30–90 Hz) oscillations are known to be associated with many aspects of cognition, including sensory processing, attention, memory, and awareness (1). However, the biophysical underpinnings of gamma oscillations and the means by which they contribute to cognitive processing are still not fully understood. Here we continue a series of computational studies investigating how the biophysics of the gamma rhythm, and in particular its neuromodulation, affect network processing associated with attention (2–4).

We focus on the much-cited experiments of Desimone and colleagues (5–7). These showed that orientation-selective principal neurons in macaque V2 and V4 respond to simultaneous presentation of preferred ("good") and nonpreferred ("poor") stimuli in their receptive fields at spike rates well below those elicited by the preferred stimulus alone. Furthermore, when attention is directed to either of the two stimuli, the neuron responds almost as if the unattended stimulus were absent.

The essence of the mechanism we propose here is a relationship between the effectiveness of the excitatory inputs to the target network and the coherence of the network's inhibitory cells: The same drive that produces a robust response when the inhibition is coherent produces a much less robust response when the inhibition is incoherent. We describe how this mechanism can, qualitatively at least, explain the results in (5–7).

Selective attention is associated with an increase in coherence of spiking and in spectral power of oscillations in the gamma frequency band (8–10). The model we offer here does not assume *a priori* that bottom-up or top-down input is made more coherent or more oscillatory as a result of attention (11, 12). Rather, the increased gamma power and coherence of the network output are consequences of the biophysical effects of cholinergic modulation. A critical feature of the model is that the target entity of competing streams of signals is not an individual pyramidal cell (as in refs. 11 and 13), but a network in which there is local inhibitory feedback to the pyramidal cells (4). The gamma oscillation generated by the local network of E cells (excitatory principal neurons) and I cells (inhibitory interneurons) acts to select the most strongly driven E cells and provides a "biased competition" (14) mechanism to suppress cells with weaker drive.

A crucial question is why input from a poor stimulus for a given cell, presented in isolation, increases activity in that cell, whereas adding a poor stimulus to a good stimulus makes the activity decrease. Reynolds *et al.* (7) showed that this is possible in an abstract rate model in which the poor stimulus gives strong inhibition and weak excitation to the target neuron. However, for a variation on their model in which the target is a spiking neuron, we show in supporting information (SI) *Appendix A* that this idea fails: A stimulus that is excitatory when presented by itself cannot be suppressive when added to another stimulus.

This result suggests that there is some ingredient essential for the suppression effect, but absent in the abstract model of Reynolds *et al.*, and in the spiking variation on that model considered in *SI Appendix A*. One natural possibility is that oscillations may play a crucial role. For instance, a stimulus that, presented by itself, gives strong but oscillating inhibition and occasional stochastic bursts of excitation to the target neuron may elicit spiking on the infrequent occasions when the excitation happens to arrive while the inhibition is weak; yet such a stimulus may well have a net inhibitory effect when added to a good stimulus.

A second ingredient that may be essential for the suppression effect is the nonlinearity of stimulus summation. A simplifying assumption made in the model of Reynolds *et al.* (7) and in *SI Appendix A* is that the synaptic input added to the target when the poor stimulus is added to the good stimulus equals the synaptic input affecting the target when only the poor stimulus is presented. When the target is a neuron embedded in a network, this assumption certainly does not hold. It is possible that the nonlinearity of stimulus summation, by itself, suffices to explain the suppression effect (15).

In our model, both oscillations and the nonlinearity of stimulus summation play important roles. A stimulus that is poor for a given neuron gives weak and stochastically fluctuating excitatory input to that neuron. When the poor stimulus is presented in isolation, it also gives rise to inhibition of the target neuron that is strong, but oscillates at gamma frequency. The oscillations create windows of opportunity for the excitatory input stream to elicit some spiking. When the poor stimulus is added to a good stimulus, the overall drive to the I cells becomes so strong that the gamma oscillation disintegrates. The resulting asynchronous inhibition leaves much less room for the noisy excitatory input stream to elicit postsynaptic spiking.

Author contributions: C.B., S.E., and N.J.K. designed research; C.B. performed research; and C.B., S.E., and N.J.K. wrote the paper.

The authors declare no conflict of interest.

[1]To whom correspondence may be addressed. E-mail: christoph.borgers@tufts.edu or nk@math.bu.edu.

This article contains supporting information online at www.pnas.org/cgi/content/full/0809511105/DCSupplemental.

The ideas of this article build on previous investigations of properties of E–I networks, and in particular on ideas from ref. 3. There we used the "theta model" (16–18), a caricature of the Hodgkin–Huxley model closely related to the quadratic integrate-and-fire model. We examined mechanisms by which pyramidal-interneuronal network gamma (PING) oscillations (19) are formed and lost as parameters are changed, especially in the presence of noise. In PING, both pyramidal cells and interneurons fire at close to network frequency. The mechanism from ref. 3 most relevant to the current work is one in which the I cells fail to synchronize, and their asynchronous activity suppresses the E cells. Whether the E cells can be suppressed by asynchronous inhibitory activity is determined by the strengths of the inhibitory conductances affecting the E cells and the external drives to the E and I cells. Escape from suppression is aided by lowering external drive to the I cells, or by raising drive to the E cells.

In the present article, we work with conductance-based models for which the analytical results of ref. 3 are not available. Nevertheless, the ideas of suppression of the E cells by incoherence of the I cells can still be made precise; this is done in *SI Appendix B* for the limiting case of pure asynchrony of the I cells. In practice (including our simulations here), there is reduced coherence, but not pure asynchrony. The activity of the E cells is not lost instantaneously as parameters are varied, but over a range in parameter space that can be fairly narrow (see *SI Appendix*, Fig. S1). Motivated by experimental data such as those of Reynolds *et al.* (7), we show that the effect of adding stimuli in the absence of attention can be to move the network parameters toward suppression, and the effect of attention to a target stimulus can be to move them back to the oscillatory regime.

The current work also builds on ideas from ref. 2, which emphasized the role of timing in suppressing a subset of the E cells. If some E cells receive more drive than others, they can govern the frequency of the inhibitory rhythm, making the I cells spike just before the more weakly driven E cells are ready to spike in each gamma cycle. Winner-take-all mechanisms are possible with or without oscillations; however, we showed in ref. 2 that they are particularly effective in the oscillatory regime (see *Discussion*). For other work on fast suppression mechanisms based on oscillatory network dynamics, see refs. 20–24.

Results

We present simulations of a network of excitatory and inhibitory model neurons (see *Methods* and *SI Appendix, SI Methods* for details). Two effects of attention are included in our model. First, we assume that cholinergic modulation subjects all cells of the network to a bath of inhibition, representing activity of cholinergically excited interneurons, perhaps low threshold spiking (LTS) cells (25, 26). We assume that the natural frequency of the I cells, i.e., the frequency with which they would spike if they were isolated from the E cells, is reduced, as a result of the bath of inhibition, approximately to what it would be with a single stimulus present. We associate this nonspecific inhibition with spatial attention directed toward the receptive field containing both the target and the distractor; we assume that it affects nearby, but not distant, receptive fields. Such nonspecific inhibition could be activated via the thalamus (27), via inhibitory systems associated with direct prefrontal pathways (28), or by localized modulation arising from basal forebrain cholinergic projections (29). Our assumptions are consistent with data suggesting that the spatial extent of focal attention in V4 can be as small as a receptive field (30); data from MT suggest that attention can even select subregions of the receptive field (31). The E cells are assumed to be subject to a bath of inhibition as well; this enhances the attentional selection effect, but is not required for it (see Fig. 5). Second, we assume that the attended stimulus is in effect strengthened. This models the effect of a top-down selection signal, and is needed to break the symmetry between the two stimuli; it is implemented here by amplifying the deterministic (constant) component of the drive originating from the attended stimulus, typically by approximately 30%. This enables the oscillatory winner-take-all mechanism of ref. 2.

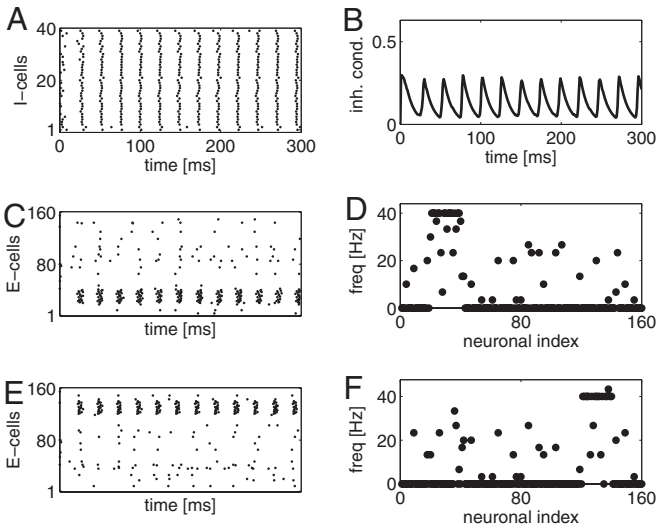

Fig. 1. Network response when a single stimulus is presented. (*A*) Response of I cells. (*B*) Inhibitory conductance affecting each E cell. (*C*) Response of E cells when first stimulus is presented. (*D*) Frequency as a function of neuronal index in *C*. (*E*) Response of E cells when second stimulus is presented. (*F*) Frequency as a function of neuronal index in *E*.

In this article, we assume that inputs are not oscillatory, and show how gamma oscillations arise and contribute to attentional selection locally; in ref. 4, we showed how the oscillations can support attentional selection downstream. See *Discussion* for further comments on the assumption of nonoscillatory input streams.

We will first show how the above ideas work in a simple example; later we will discuss more generally the parameter regime in which they work.

Background Network and Assumptions About the Effects of a Single Stimulus. For simplicity, we start with a network that has no activity in the absence of input; this can be relaxed (data not shown). We first consider a single stimulus that strongly drives a subset of E cells; for these cells, it is a "good stimulus," in the terminology of ref. 7. It also noisily and less strongly drives some other E cells, in our model chosen randomly from the rest; for these, it is a "poor stimulus." The stimulus drives all I cells equally; thus, in our model, the I cells are not stimulus-selective. This assumption, too, can be relaxed. The essential requirement here is that for any pair of stimuli, there is a substantial population of I cells driven by both; we comment further on this assumption in *Discussion*.

Presentation of such a stimulus creates a PING oscillation, with the strongly driven E cells and all I cells firing at gamma frequency (2) (Fig. 1 *A–D*). We also consider a second stimulus which has the same effect, but selects a different (nonoverlapping) set of strongly driven cells, and another subset that is driven less strongly and noisily. There is (in the absence of attention) no essential difference between the two stimuli, except from the point of view of a particular cell that may be driven strongly by one and not by the other; see Fig. 1 *E* and *F*.

Presenting Both Stimuli Together Lowers the Rate of Firing and Eliminates the Oscillation. In Fig. 2 we show the effect of presenting both stimuli. The input is now the sum of the inputs from Fig. 1, so the I cells now get twice as much drive as before. No changes have been made to the drive to any of the strongly driven E cells (there

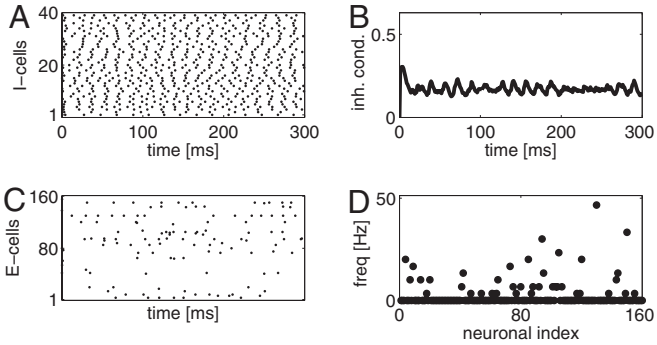

Fig. 2. Network response when both stimuli are presented. (*A*) Response of I cells. (*B*) Inhibitory conductance affecting each E cell. (*C*) Response of E cells. (*D*) Frequency as a function of neuronal index in *C*.

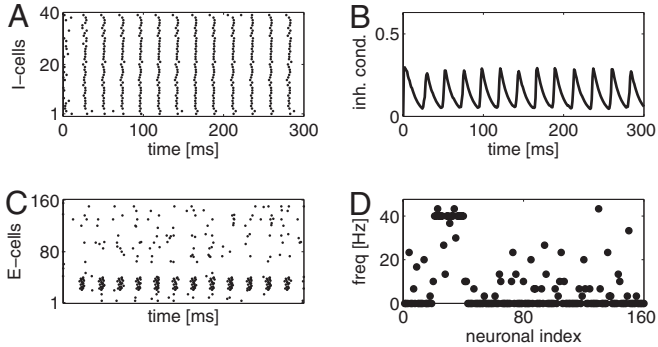

Fig. 4. Network response when both stimuli are presented, with a bath of inhibition, and with the deterministic drive to E cells 21–40 raised by 30%. Panels as in Fig. 2.

are now just more of them), except possibly if some of the strongly driven cells also get some noisy drive from the other stimulus. The stronger drive to the I cells now gives them a higher intrinsic firing rate, which leads to a higher average level of inhibition of the E cells. If this is a strong enough effect, one expects the I cells to become more asynchronous and the E cells to become less active (3). This is exactly what happens, as shown in Fig. 2.

One possible concern is that the I cell asynchrony and loss of E cell activity might be independent effects, contradicting the claim that they are mechanistically related. To address this concern, we present the first stimulus together with a fraction w of the second, with $0 \leq w \leq 1$, to see the change in both the oscillations and the frequency of the E cells as w increases. As shown in Fig. 3, the loss of the oscillation and the suppression of the E cells in fact occur together.

Selective Attention Restores the Oscillation and the Response to the Attended Stimulus. Fig. 4 shows the network response when both stimuli are presented as in Fig. 2, with all cells subject to a bath of inhibition (see *SI Appendix, SI Methods*), and the deterministic component of the drive resulting from the first stimulus raised by 30%, modeling a firing rate increase elicited by a top-down attentional signal. The oscillation is restored, the response to the attended stimulus is strong, and the response to the unattended stimulus is largely suppressed. All E cells respond approximately as they would if only one stimulus had been presented; compare Fig. 1 *C* and *D* with Fig. 4 *C* and *D*.

The bath of inhibition to the E cells is not needed for our mechanism to work, but enhances the effect. Fig. 5 illustrates this; it is the same as Fig. 4, but the bath of inhibition is only applied to the I cells. Neurons 121–140 are now only partially suppressed.

In ref. 2, we argued that attention increases the excitability of pyramidal cells by reducing the M current. Here we argue that attention subjects the pyramidal cell population to a bath of inhibition. The first occurs in a vigilant state, when the drive to the pyramidal cells is weak, the second during attentional selection, when the total excitatory drive is strong. Thus, we suggest that attention may have opposite effects on the pyramidal cells in different behavioral situations.

Requirements on Parameter Values for the Mechanism to Work. Our mechanism breaks down when a substantial population of E cells is strongly driven by both stimuli. (There is no reason to believe that in the experiments discussed, for instance, in ref. 7, there was a substantial number of pyramidal cells driven strongly by both stimuli presented.) In such a situation, simulations (data not shown here) predict that those excitatory cells driven strongly by both stimuli will be active, with most others suppressed.

Our mechanism also requires inhibitory synapses strong enough to enable PING rhythms, but weak enough not to suppress all noise-driven spiking in the presence of the PING rhythm. Each stimulus must add enough drive to the I cells so that two simultaneously presented stimuli desynchronize the I cells. The bath of inhibition to the I cells and drive to the E cells must be enough to counter the effects of the added drive to the I cells. The attended input stream must be strengthened by a sufficiently large amount. For instance, Fig. 4 would look very similar if the attended input stream were strengthened by 20% (in Fig. 4, it was strengthened by 30%), but, if it were strengthened by only 10%, the selection mechanism would not work.

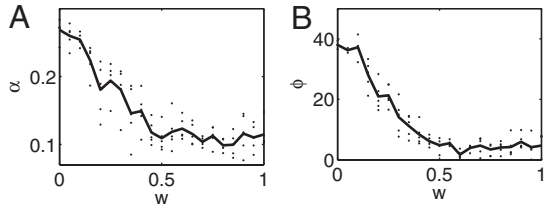

Fig. 3. The first stimulus is presented together with w times the second stimulus. ϕ = mean frequency of E cells 21–40 (*A*) and α = amplitude of gamma oscillation (*B*) as functions of w. The amplitude α is defined to be max $g_I(t)$ − min $g_I(t)$, where g_I denotes the total inhibitory conductance affecting each E cell, and the maximum and minimum are taken over the time interval [300,600] to avoid initialization artifacts. For each value of w, the figure shows 5 different values of α and ϕ (dots) obtained with 5 different seeds of the random number generator, and their average (bold lines).

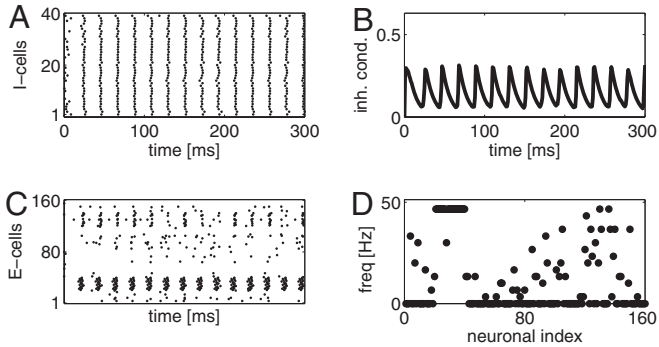

Fig. 5. Network response when both stimuli are presented, with a bath of inhibition affecting the I cells only, and with the deterministic drive to E cells 21–40 raised by 30%. Panels as in Fig. 2.

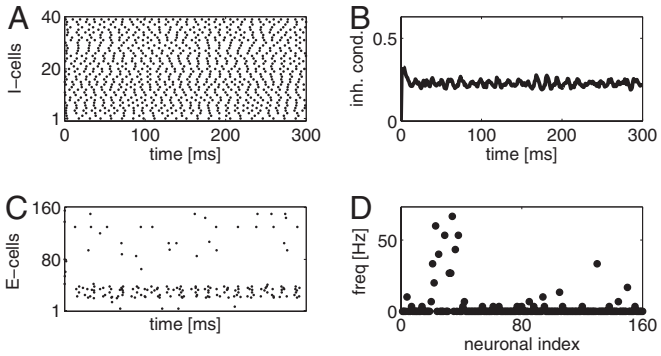

Fig. 6. Network response when both stimuli are presented, without any bath of inhibition, and with the deterministic component of the first stimulus doubled. Panels as in Fig. 2.

Within these requirements, the mechanism is robust. Parameters such as the synaptic decay time constants and the strengths of inhibitory synapses, external drives, and the bath of inhibition can be varied considerably without any qualitative change in behavior; see *SI Appendix*, Figs. S2–S6.

Contrast Enhancement and Selective Attention are Not the Same. An increase in contrast can be thought of as a bottom up increase in the drive originating from a specific stimulus; in this model, it plays the same role as specific top down drive. Our model shows a crucial difference between an increase in contrast and attention, however: In the presence of several stimuli, an increase in the strength of one of them, without diffuse inhibition of the I cells, does not easily restore the oscillations; see Fig. 6.

Buia and Tiesinga (11) have previously hypothesized that attention may correspond to a reduction in driving current to the inhibitory neurons, and have used this idea to explain the attentional effect on contrast-response curves.

Selective Attention Without Gamma Oscillations. It is possible to obtain the attentional selection effect without a rhythm. This is shown in Fig. 6, in which there is no bath of inhibition to the I or E cells. However, to obtain the selection effect, one must raise the strength of the attended stimulus by a factor that is likely to be unrealistically large; in Fig. 6, it was doubled. The I cells now fire asynchronously at a high frequency, suppressing the response to the second (unattended) stimulus. However, the response to the first (attended) stimulus is partially suppressed as well—only approximately half the cells in the strongly driven group (E cells 21–40) fire. Among those cells that do fire vigorously, there is a considerable spread in frequencies, approximately between 20 and 60 Hz. This, at least in our hands, is the typical difference between oscillatory and nonoscillatory competition: Without an oscillation, suppression of the response to unattended stimuli requires a large amount of inhibition, which will affect even the winning (attended) ensemble, as suggested by the analysis surrounding figure 4 of ref. 2. The oscillation allows inhibition to be timed in such a way that it minimally affects the winning ensemble, and maximally affects losing ensembles. In summary, Fig. 6 demonstrates that stimulus competition is possible without a rhythm; it is, however, more effective in the presence of a gamma rhythm.

We also note that gamma oscillations occur with great robustness in E/I networks, as long as the E cells are strongly excited, the I cells are not too strongly excited, the synaptic interactions between the two cell populations are strong, and the decay time constant $\tau_{D,I}$ of inhibition is approximately on the order of 10 ms (32). Shorter (33) and longer (34) time constants have been reported for $GABA_A$-receptor mediated synaptic transmission under some circumstances. The precise value of $\tau_{D,I}$ is not important for our mechanism; see *SI Appendix*, Figs. S2 and S3.

Stimulus Competition With Gamma Oscillations in the Absence of Attention. *SI Appendix*, Fig. S7 shows that the suppression of a good stimulus by a poor one, in the absence of attention, can occur, in a somewhat different parameter regime, without a reduction in gamma power. Selection by means of a bath of inhibition and moderate strengthening of the attended input is possible in this parameter regime as well (see *SI Appendix*, Fig. S8). However, for reasons discussed in detail in *SI Appendix, SI Figures*, we find it less plausible that this parameter regime is physiologically realistic.

Discussion

Model Assumptions. In our model, attention inside the receptive field leads to reduced drive to the I cells associated with that receptive field. This models the presumed effect of cholinergic activation of some network of interneurons, perhaps LTS cell (25, 26), that inhibit fast-firing interneurons. [Although Beierlein *et al.* (35) have shown slow oscillations generated by LTS cells coupled by gap junctions in thalmo-cortical slices, such synchronization would not necessarily preclude their providing the bath of inhibition that we postulate; this bath does not have to be tonic.] Froemke *et al.* (36) have shown that activation of nucleus basalis, the main source of acetylcholine in the cortex, leads to a rapid reduction of synaptic inhibition in rat primary auditory cortex; our assumption is in line with this result.

We also assume that attention amplifies the specific input generated by the attended stimulus. We note that this firing rate increase need not be the same as that seen when attention is directed to a single stimulus in the receptive field, which has been found to be weak or nonexistent in some cases (for example, see ref. 6, but also ref. 37). Our assumptions are in agreement with Roberts *et al.* (38), who proposed that "the effects of spatial attention are mediated by an interaction of cholinergic input and feedback connections. from higher cortical areas." Herrero *et al.* (39) have recently demonstrated that muscarinic acetylcholine receptors play an important role in mediating effects of selective attention.

Our simulations were done assuming that the E cell ensembles selected by the two stimuli do not overlap much. This does not preclude broad tuning curves, as long as there are no cells that are strongly driven by both stimuli. We predict that in experiments in which the overlap is substantial, the results of ref. 7 about firing rates and attention will not hold.

Our mechanism requires that there be a population of fast-spiking interneurons that are strongly excited by both stimuli. This does not imply that all such neurons must be nonselective (see refs. 40–45). For instance, the results of ref. 45 are consistent with our assumption, even though they demonstrate stimulus selectivity of inhibitory interneurons in area TA of the inferior temporal cortex.

In the present article, inputs are nonoscillatory, but the network generates oscillatory outputs. We hypothesize that this is how selective attention may enhance gamma frequency coherence at early stages of processing. In ref. 4, inputs were oscillatory, and the modeling showed how more coherent oscillations can shut out less coherent, distracting ones. We hypothesize that the mechanism of ref. 4 may enhance selection at later stages of processing. Simulations (data not shown) indicate that the mechanism of the present article can in fact also operate at a stage when inputs are already oscillatory.

The idea that two classes of interneurons participate in the attentional effects has independently been suggested in a recent article by Buia and Tiesinga (46). There are substantial differences between their model and ours, however. In ref. 46, gamma

oscillations are created by inhibitory cells responding to top-down signals; in our model, they are created locally in the network of pyramidal cells and fast-spiking interneurons (19). In ref. 46, the local fast-spiking interneurons are stimulus selective; in our model, it is crucial that they be *less* stimulus selective than the pyramidal cells. Also, in contrast with our work, asynchrony does not play a role in ref. 46.

Implications of our Model. Our reasoning in this article is based on the connection between the effectiveness of excitatory inputs to a target network and the coherence of the inhibitory cells in the network (2, 3). We showed how one can use this notion to understand the results of Reynolds *et al.* (7) on the changes in firing rates associated with competing inputs and attention. (Refs. 12, 13, 15, and 47–50 also address those data.) Multiple inputs can decrease coherence of the interneurons by increasing their drive, and attention can increase their coherence by providing a bath of inhibition that reduces their drive. Although the simulations were for two stimuli, the ideas apply to multiple stimuli, provided that there is saturation of the total amount of inhibition.

Our model predicts that when two stimuli are in the same receptive field, the spike-spike coherence (SSC) should decrease significantly, provided that there is no attentional selection, the stimuli are strong enough, and there is no group of E cells driven strongly by both. To our knowledge, there are no experimental data that address this prediction directly; such data would provide a critical test for our model. An SSC analysis was not done in ref. 7. In ref. 51, this analysis was done, but the stimuli were in different receptive fields. In some of the early literature on binding, experiments were described in which two stimuli in the same receptive field failed to abolish a local gamma oscillation; see for instance figure 5 D–F of ref. 52. However, there was no control for attention in ref. 52. There may also be cases in which two stimuli in the same receptive field, even in the absence of attention, do not abolish the gamma oscillation because they are not strong enough. In ref. 53, a gamma oscillation was found to be evoked by a single moving grating, and abolished by adding a second grating (plaids), in keeping with our prediction, even though some E cells were driven strongly by both stimuli. However, in contrast with our model, the data of ref. 53 indicate that the gamma oscillation is lost without a reduction in the activity of the E cells; this does not occur in the parameter regime of our modeling study, although it is possible when the drives to the E and I cells are quite strong. (In figure 7C of ref. 3, the region in which asynchronous activity of both the E and the I cells occurs is the region above the solid and below the dashed line, in the right upper corner, corresponding to strong drive to both E and I cells.)

Our model also provides a framework in which to understand recent results of Mitchell *et al.* (54), showing that attention lowers the Fano factor (ratio of spike count variance to mean spike rate) more for the fast-spiking interneurons than for the pyramidal cells. This may be a consequence of the PING mechanism: In PING, the excitation synchronizes the interneurons more tightly than the inhibition does the pyramidal cells (55).

In the current article, we described a situation in which both stimuli lie in the same receptive field. In this case, there is a large shared pool of relevant interneurons. The condition is related to, but not identical with, that of Fries *et al.* (51), in which the target and distractor lie in different, although possibly nearby receptive fields. In the latter case, there may still be overlapping pools of interneurons, but also other interneurons driven by only one of the two stimuli.

The transition from states of little coherence to gamma rhythms always goes together with an increase in firing rates of individual E cells in our model. However, the coherence of the gamma oscillation can be tightened without any change in the population frequency (see for instance ref.9). In our model, cells participating strongly in the gamma oscillation always fire at the population frequency. A more complex model, perhaps involving more stochasticity in drives (2), could allow cells to skip cycles. In such a model, it should be possible to tighten gamma coherence without altering the firing frequency of participating cells much.

Sustained and Selective Attention. In ref. 2, we modeled sustained attention (vigilance) with a low amplitude noisy gamma rhythm, which provides a windowing of the inhibition, allowing low level excitation to break through. We showed that disturbing the coherence of the inhibitory cells leads to a much lower response to excitatory drive. Weaker stimuli are suppressed by an oscillatory winner-take-all mechanism: The E cells coding for the salient input are given more drive and cause the I cells to fire; they themselves are ready to fire again when the inhibition is lowest. By contrast, the less excited E cells corresponding to the distractor input are not ready to fire before the more excited cells induce another barrage of inhibition and prevent firing. The windowing of the inhibition is thus a barrier to a "poor" stimulus. It is crucial to this mechanism that the poor stimulus not be strong enough to disturb the rhythm formed by the interaction of the I cells and the set of E cells coding for the more salient input.

In the present work, the two stimuli are equivalent in strength in the absence of attention. The interaction of the two stimuli leads to a reduction or suppression of the rhythm, with an associated reduction of firing rate. The effect of attention is to make one of the two inputs stronger than the other (the result of a selective top-down signal) and to reduce the effect of the distractor input on the I cells with the bath of inhibition. This effectively recreates the situation of ref. 2, in which the locked oscillation of the cells coding for the stronger input can suppress the effects of the weaker input.

Attention and Oscillatory Selection. The gamma rhythm plays three separate roles in our models of attention. First, it increases sensitivity to weak excitatory stimuli by creating windows of opportunity (2). Second, it allows a more strongly driven ensemble to efficiently suppress a less strongly driven one, as described in ref. 2 and the present article; this is a timing effect (23, 24, 56). Third, the gamma rhythm allows some signals, the tightly coherent ones, to be processed downstream while others are suppressed (4).

All three of the above mechanisms are examples of what one might call "oscillatory selection." We see this as closely related to the idea of "communication through neuronal coherence" proposed by Pascal Fries (57). Womelsdorf *et al.* (58) wrote, "We provided evidence suggesting that neuronal interactions mechanistically depend on the phase relationship between rhythmic activities." We agree but emphasize here that phase relationships, their functional consequences, and even the capacity of the networks to phase lock or oscillate at all, mechanistically depend on neuronal interactions.

Methods

We only sketch our model here (refer to *SI Appendix, SI Methods* for details). Our E cell model is from Ermentrout and Kopell (59), and our I cell model is from Wang and Buzsáki (60). Our networks contain 160 E cells and 40 I cells. The external drive to each cell has a constant component, chosen from a Gaussian distribution with coefficient of variation 0.1. In addition, there is a Poisson train of excitatory input pulses to each neuron. Synaptic interactions are modeled as in ref. 2. Network connectivity is all-to-all; our mechanism would work equally well with sparse, random network connectivity, as long as each E cell receives synaptic input from a sufficiently large number of I cells and vice versa (55). We assume strong E $\rightarrow$ I, I $\rightarrow$ E, and I $\rightarrow$ I synapses. Although our mechanism works when there is recurrent excitation within each of the stimulus-driven ensembles of E cells (data not shown), E $\rightarrow$ E synapses are not

essential, and are omitted here. In a realistic network, they might effectively add tonic excitation to the E cells (61). The numerical methods are as in ref. 2.

ACKNOWLEDGMENTS. We thank Robert Desimone and Pascal Fries for their thorough reading of this article and for constructive comments that led to numerous improvements. We also thank Murray Sherman for helpful conversations. This work was supported by National Science Foundation Grant DMS-0418832 (to C.B.) and National Institutes of Health grant R01 NS-46058 (to N.J.K.) as part of the National Science Foundation/National Institutes of Health Collaborative Research in Computational Neuroscience Program.

1. Jensen O, Kaiser J, Lachaux JP (2007) Human gamma-frequency oscillations associated with attention and memory. *Trends Neurosci* 30:317–324.
2. Börgers C, Epstein S, Kopell N (2005) Background gamma rhythmicity and attention in cortical local circuits: A computational study. *Proc Natl Acad Sci USA* 102:7002–7007.
3. Börgers C, Kopell N (2005) Effects of noisy drive on rhythms in networks of excitatory and inhibitory neurons. *Neural Comp* 17:557–608.
4. Börgers C and Kopell N (2008) Gamma oscillations and stimulus selection. *Neural Comp* 20:383–414.
5. Luck SJ, Chelazzi L, Hillyard SA, Desimone R (1997) Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *J Neurophysiol* 77:24–42.
6. Moran J, Desimone R (1985) Selective attention gates visual processing in the extrastriate cortex. *Science* 229:782–784.
7. Reynolds JH, Chelazzi L, Desimone R (1999) Competitive mechanisms subserve attention in macaque areas V2 and V4. *J Neurosci* 19:1736–1753.
8. Bichot NP, Rossi AF, Desimone R (2005) Parallel and serial neural mechanisms for visual search in macaque area V4. *Science* 308:529–534.
9. Fries P, Reynolds JH, Rorie AE, Desimone R (2001) Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* 291:1560–1563.
10. Womelsdorf T, Fries P (2006) Neuronal coherence during selective attentional processing and sensory-motor integration. *J Physiol (Paris)* 100:182–193.
11. Buia C, Tiesinga P (2006) Attentional modulation of firing rate and synchrony in a model cortical network. *J Comput Neurosci* 20:247–264.
12. Zeitler M, Fries P, Gielen S (2008) Biased competition through variations in amplitude of gamma-oscillations. *J Comput Neurosci* 25:89–107.
13. Mishra J, Fellous JM, Sejnowski TJ (2006) Selective attention through phase relationship of excitatory and inhibitory input synchrony in a model cortical neuron. *Neural Networks* 19:1329–1346.
14. Desimone R, Duncan J (1995) Neural mechanisms of selective visual attention. *Annu Rev Neurosci* 18:193–222.
15. Ardid S, Wang XJ, Compte A (2007) An integrated microcircuit model of attentional processing in the neocortex. *J Neurosci* 27:8486–8495.
16. Ermentrout GB, Kopell N (1986) Parabolic bursting in an excitable system coupled with a slow oscillation. *SIAM J Appl Math* 46:233–253.
17. Gutkin BS, Ermentrout GB (1998) Dynamics of membrane excitability determine interspike interval variability: A link between spike generation mechanisms and cortical spike train statistics. *Neural Comp* 10:1047–1065.
18. Hoppensteadt FC, Izhikevich EM (1997) *Weakly connected neural networks* (Springer, New York).
19. Whittington MA, Traub RD, Kopell N, Ermentrout GB, Buhl EH (2000) Inhibition-based rhythms: Experimental and mathematical observations on network dynamics. *Int. J. of Psychophysiol* 38:315–336.
20. Jin DZ, Seung HS (2002) Fast computation with spikes in a recurrent neural network. *Phys Rev E* 65:051922.
21. Olufsen M, Whittington M, Camperi M, Kopell N (2003) New functions for the gamma rhythm: Population tuning and preprocessing for the beta rhythm. *J Comp Neurosci* 14:33–54.
22. Tiesinga PHE, Sejnowski TJ (2004) Rapid temporal modulation of synchrony by competition in cortical interneuron networks. *Neural Comput* 16:251–275.
23. Wang W, Slotine J-JE (2003) K-winners-take-all computation with neural oscillators. Available at http://arxiv.org/pdf/q-bio/0401001.
24. Coultrip R, Granger R, Lynch G (1992) A cortical model of winner-take-all competition via lateral inhibition. *Neural Networks* 5:47–54.
25. Xiang Z, Huguenard JR, Prince DA (1998) Cholinergic switching within neocortical inhibitory networks. *Science* 281:985–988.
26. Beierlein M, Gibson JR, Connors BW (2003) Two dynamically distinct inhibitory networks in layer 4 of the neocortex. *J Neurophysiol* 90:2987–3000.
27. Zikopoulos B, Barbas H (2006) Prefrontal projections to the thalamic reticular nucleus form a unique circuit for attentional mechanisms. *J Neurosci* 26:7348–7361.
28. Medalla M, Lera P, Feinberg M, Barbas H (2007) Specificity in inhibitory systems associated with prefrontal pathways to temporal cortex in primates. *Cerebral Cortex* 17:136–150.
29. Price J, Stern R (1983) Individual cells in the nucleus basalis-diagonal band complex have restricted axonal projections to th cerebral cortex in the rat. *Brain Res* 269:352–356.
30. Balz GW, Hock HS (1997) The effect of attentional spread on spatial resolution. *Vision Res* 37:1499–1510.
31. Womelsdorf T, Anton-Erxleben K, Pieper F, Treue S (2006) Dynamic shifts of visua receptive fields in cortical area MT by spatial attention. *Nat Neurosci* 9:1156–1160.
32. Xiang Z, Huguenard JR, Prince DA (2002) Synaptic inhibition of pyramidal cells evoked by different interneuronal subtypes in layer V of rat visual cortex. *J Neurophysiol* 88:740–750.
33. Bartos M, *et al.* (2002) Fast synaptic inhibition promotes synchronized gamma oscillations in hippocampal interneuron networks. *Proc Natl Acad Sci USA* 99:13222–13227.
34. Banks MI, Li T-B, Pearce RA (1998) The synaptic basis of $GABA_{A,slow}$. *J Neurosci* 18:1305–1317.
35. Beierlein M, Gibson JR, Connors BW (2000) A network of electrically coupled interneurons drives synchronized inhibition in neocortex. *Nat Neurosci* 3:904–910.
36. Froemke RC, Merzenich MM, Schreiner CE (2007) A synaptic memory trace for cortical receptive field plasticity. *Nature* 450:425–429.
37. McAdams CJ, Maunsell JH (1999) Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J Neurosci* 19:431–441.
38. Roberts M, Delicato LS, Herrero J, Gieselmann MA, Thiele A (2007) Attention alters spatial integration in macaque V1 in an eccentricity-dependent manner. *Nat Neurosci* 10:1483–1491.
39. Herrero JL, *et al.* (2008) Acetylcholine contributes through muscarinic receptors to attentional modulation in V1. *Nature* 454:1110–1114.
40. Azouz R, Gray CM, Nowak IG, McCormick DA (1997) Physiological properties of identified interneurons in cat striate cortex. *Cerebral Cortex* 7:534–545.
41. Desimone R, Albright T, Gross C, Bruce C (1984) Stimulus-selective properties of inferior temporal neurons in the macaque. *J Neurosci* 4:2051–2062.
42. Martin KAC, Somogyi P, Whitteridge D (1983) Physiological and morphological properties of identified basket cells in the cat's visual cortex. *Exp Brain Res* 50:193–200.
43. Roerig B, Chen B (2002) Relationships of local inhibitory and excitatory circuts to orientation preference maps in ferret visual cortex. *Cerebral Cortex* 12:187–198.
44. Swadlow HA (2003) Fast-spike interneurons and feedforward inhibition in awake sensory neocortex. *Cerebral Cortex* 13:25–32.
45. Tamura H, Kaneko H, Kawasaki K, Fujita I (2004) Presumed inhibitory neurons in the macaque inferior temporal cortex: Visual response properties and functional interactions with adjacent neurons. *J Neurophysiol* 91:2782–2796.
46. Buia CI, Tiesinga PH (2008) Role of interneuron diversity in the cortical microcircuit for attention. *J Neurophysiol* 99:2158–2182.
47. Deco G, Rolls ET (2005) Neurodynamics of biased competition and cooperation for attention: A model with spiking neurons. *J Neurophysiol* 94:295–313.
48. Niebur E, Koch C, Rosin C (1993) An oscillation-based model for the neural basis of attention. *Vis Res* 33:2789–2802.
49. Tiesinga PHE (2005) Stimulus competition by inhibitory interference. *Neural Comput* 17:2421–2453.
50. Tiesinga PHE, Fellous J-M, Salinas E, José JV, Sejnowski TJ (2004) Synchronization as a mechanism for attentional gain modulation. *Neurocomputing* 58–60:641–646.
51. Fries P, Womelsdorf T, Oostenveld R, Desimone R (2008) The effects of visual stimulation and selective visual attention on rhythmic neuronal synchronization in macaque area V4. *J Neurosci* 28:4823–4835.
52. Kreiter AK, Singer W (1996) Stimulus-dependent synchronization of neuronal responses in the visual cortex of the awake macaque monkey. *J Neurosci* 16:2381–2396.
53. Lima B, Chen N, Singer W, Neuenschwanter S (2005) Plaids stimuli disrupt gamma activity in V1 of the behaving monkey. *Soc Neurosci 2005 Abstract Viewer/Itinerary Planner* (Program No. 855.8). Avalailable at http://sfn.scholarone.com/itin2005.
54. Mitchell JF, Sundberg KA, Reynolds JH (2007) Differential attention-dependent response modulation across cell classes in macaque visual area v4. *Neuron* 55:131–141.
55. Börgers C, Kopell N (2003) Synchronization in networks of excitatory and inhibitory neurons with sparse, random connectivity. *Neural Comp* 15:509–539.
56. Burchell TR, Faulkner HJ, Whittington MA (1998) Gamma frequency oscillations gate temporally coded afferent inputs in the rat hippocampal slice. *Neurosci Lett* 255:151–154.
57. Fries P (2005) A mechanism for cognitive dynamics: Neuronal communication through neuronal coherence. *Trends Cognit Sci* 9:474–480.
58. Womelsdorf T, *et al.* (2007) Modulation of neuronal interactions through neuronal synchronization. *Science* 316:1609–1612.
59. Ermentrout GB, Kopell N (1998) Fine structure of neural spiking and synchronization in the presence of conduction delay. *Proc Natl Acad Sci USA* 95:1259–1264.
60. Wang XJ, Buzsáki G (1996) Gamma oscillation by synaptic inhibition in a hippocampal interneuronal network model. *J Neurosci* 16:6402–6413.
61. Morita K, Kalra R, Aihara K, Robinson HPC (2008) Recurrent synaptic input and the timing of gamma-frequency-modulated firing of pyramidal cells during neocortical ''up'' states. *J Neurosci* 28:1871–1881.