

Rapid Temporal Modulation of Synchrony by Competition in Cortical Interneuron Networks

P.H.E. Tiesinga

tiesinga@physics.unc.edu

Department of Physics and Astronomy, University of North Carolina, Chapel Hill, NC 27599, U.S.A.

T.J. Sejnowski

terry@salk.edu

Sloan-Swartz Center for Theoretical Neurobiology, Salk Institute, La Jolla, CA 92037, Computational Neurobiology Lab, Salk Institute, La Jolla, CA 92037, Howard Hughes Medical Institute, Salk Institute, La Jolla, CA 92037, and Department of Biology, University of California–San Diego, La Jolla, CA 92093, U.S.A.

The synchrony of neurons in extrastriate visual cortex is modulated by selective attention even when there are only small changes in firing rate (Fries, Reynolds, Rorie, & Desimone, 2001). We used Hodgkin-Huxley type models of cortical neurons to investigate the mechanism by which the degree of synchrony can be modulated independently of changes in firing rates.

The synchrony of local networks of model cortical interneurons interacting through GABA_A synapses was modulated on a fast timescale by selectively activating a fraction of the interneurons. The activated interneurons became rapidly synchronized and suppressed the activity of the other neurons in the network but only if the network was in a restricted range of balanced synaptic background activity. During stronger background activity, the network did not synchronize, and for weaker background activity, the network synchronized but did not return to an asynchronous state after synchronizing. The inhibitory output of the network blocked the activity of pyramidal neurons during asynchronous network activity, and during synchronous network activity, it enhanced the impact of the stimulus-related activity of pyramidal cells on receiving cortical areas (Salinas & Sejnowski, 2001). Synchrony by competition provides a mechanism for controlling synchrony with minor alterations in rate, which could be useful for information processing.

Because traditional methods such as cross-correlation and the spike field coherence require several hundred milliseconds of recordings and cannot measure rapid changes in the degree of synchrony, we introduced a new method to detect rapid changes in the degree of coincidence and precision of spike timing.

1 Introduction

Selective attention greatly enhances the ability of the visual system to detect visual stimuli and to store and recall these stimuli. For instance, in change blindness, differences between two images shown after a brief delay are not visible unless the subject pays attention to the particular spatial location of the change (Simons, 2000; Rensink, 2000). The neural correlate of selective attention has recently been studied in macaque monkeys (Connor, Gallant, Preddie, & Van Essen, 1996; Connor, Preddie, Gallant, & Van Essen, 1997; Luck, Chelazzi, Hillyard, & Desimone, 1997; McAdams & Maunsell, 1999; Treue & Martinez Trujillo, 1999; Reynolds, Pasternak, & Desimone, 2000; Fries, Reynolds, Rorie, & Desimone, 2001; Moore & Armstrong, 2003). A key finding is that attention modulates both the mean firing rate of a neuron in response to a stimulus (McAdams & Maunsell, 1999; Reynolds et al., 2000) and the coherence of spiking with other neurons responsive to the stimulus (Fries et al., 2001). Increased synchrony can boost the impact of spikes on subsequent cortical areas to which these neurons project (Salinas & Sejnowski, 2000, 2001, 2002).

These results suggest that mean activity and the degree of coherence may represent independent signals that can be independently controlled. It is unclear biophysically what mechanisms are responsible. For example, when excitatory networks are depolarized, there can be an increase in synchrony but it is always associated with a large increase in the mean output rate (Aoyagi, Takekawa, & Fukai, 2003). Here we explore the hypothesis that interneuron networks are involved in modulating synchrony. The issue of how the interneuron network affects the pyramidal cells is a complex one that is not the subject of this article (but see section 4). We found that interneuron networks can dynamically modulate their synchrony on timescales of 100 ms with only moderate changes in the mean firing rate. Synchrony arose by competition between two groups of interneurons and required a specific range of synaptic background activity. To analyze the temporal dynamics of synchrony modulation, we introduced time-resolved measures for the degree of coincidence and precision. The sensitivity of these measures is compared to the spike field coherence (Fries et al., 2001).

2 Methods

2.1 Network Model. The model consisted of a network of N interneurons connected all-to-all. Each neuron had a sodium I_{Na} , a potassium I_K , and a leak current I_L and was driven by a white noise current η with mean I_i and variance $2D$. The equation for the membrane potential V reads:

$$C_m \frac{dV}{dt} = -I_{Na} - I_K - I_L + \eta(t) + I_{syn}. \quad (2.1)$$

Here, $C_m = 1\mu\text{F}/\text{cm}^2$ is the membrane capacitance (normalized by area), and I_{syn} is the synaptic input from other neurons in the network. For the first N_A neurons the current was $I_i = I_A + \delta I_i$ ($i = 1, \dots, N_A$), and for the remainder, it was $I_i = I_B + \delta I_i$ ($i = N_A + 1, \dots, N$). Heterogeneity in neural properties was represented by δI_i , which was drawn from a gaussian distribution with a standard deviation σ_I . I_A was increased from I to $I + \Delta I$ during a 500 ms interval between $t = 1000$ and 1500 ms; I_B was equal to I for the duration of the simulation. The neurons were connected by GABA_A synapses with unitary strength g_{syn}/N and time constant $\tau_{syn} = 10$ ms. The model equations and implementation are exactly as described by Tiesinga and José (2000a) and are not repeated here. The model used here was adapted from that introduced by Wang and Buzsáki (1996). The standard set of parameters was $N = 1000$, $N_A = 250$, $N_B = N - N_A = 750$, $I = 0.7 \mu\text{A}/\text{cm}^2$, $\Delta I = 0.3 \mu\text{A}/\text{cm}^2$, $g_{syn} = 0.4 \text{mS}/\text{cm}^2$, $\sigma_I = 0.02 \mu\text{A}/\text{cm}^2$, and $D = 0.02 \text{mV}^2/\text{ms}$.

2.2 Statistical Analysis. Spike times were recorded when the membrane potential (V in equation 2.1) crossed 0 mV from below. t_{ni} is the i th spike time by the n th neuron; likewise, t_{mj} is the j th spike time by the m th neuron. N is the total number of neurons in the network, and N_s is the number of spikes (by all neurons) generated during a given simulation run. For some calculations, it is convenient to pool the spike times of all neurons together into one set, $\{t_1, \dots, t_\nu, \dots, t_{N_s}\}$, ordered from low to high values. The ordered set is indexed by ν , where $\nu = 1$ is the earliest spike and $\nu = N_s$ is the latest. Most of the statistical quantities determined from the simulations were spike based. That is, for each measurement y_ν there is an associated spike time t_ν . For instance, y_ν could be an interspike interval, and t_ν could be the first spike time of the interval. The specific choices for y are given below. The simulation runs were divided into three intervals: (I) the baseline state, $0 < t < 1000$ ms; (II) the activated state, $1000 \text{ ms} < t < 1500$ ms, during which N_A of the neurons received a current step, and (III) return to baseline for $1500 \text{ ms} < t < T_{\text{max}}$. Here T_{max} is the duration of the simulated trial. Each interval is divided in subintervals A and B, A is the transient during which the network adapts from the old state to the new state, and B is the steady state reached after the transient. For instance, IIA is the onset of the synchronous state, and IIB is the synchronous state itself. Note that the subdivision in A and B depends on network parameters such as I and D . Averages of y were determined for each of these intervals, using all the y_ν associated with a t_ν in the interval under consideration. Since the t_ν are ordered, this is a contiguous set, $\nu_b \leq \nu \leq \nu_e$ (b stands for begin and e stands for end):

$$\langle y \rangle = \frac{1}{\nu_e + 1 - \nu_b} \sum_{\nu=\nu_b}^{\nu_e} y_\nu, \quad \langle t \rangle = \frac{1}{\nu_e + 1 - \nu_b} \sum_{\nu=\nu_b}^{\nu_e} t_\nu. \quad (2.2)$$

We also determined time-resolved averages, $\langle y \rangle_k$ and $\langle t \rangle_k$, using a sliding window of length T_{av} that was translated along the time axis with increments equal to T_{incr} . The position of the sliding window was indicated by the index k , and the sum is over all v values with $(k-1)T_{incr} \leq t_v < (k-1)T_{incr} + T_{av}$. For convenience, the set of points $(\langle t \rangle_k, \langle y \rangle_k)$ connected by lines is denoted by $y(t)$. T_{av} and T_{incr} are expressed in terms of n_{av} and n_{incr} using the relations $T_{av} = T_{max}/n_{av}$ and $T_{incr} = (T_{max} - T_{av})/n_{incr}$. We used $n_{av} = 30$, $n_{incr} = 90$ and $T_{max} = 3000$ ms.

2.2.1 Standard Spike Train Statistics. The i th interspike interval (ISI) of the n th neuron is $\tau_{ni} = t_{n,i+1} - t_{n,i}$. The mean ISI for neuron n is denoted by τ_n , and its average across all neurons is τ . The coefficient of variation (CV) of the ISIs is the standard deviation of the τ_{ni} across all neurons divided by τ . The interval τ_{ni} can be associated with three times: the starting spike time of the interval, t_{ni} , the end spike time of the interval, $t_{n,i+1}$, and the mean, $(t_{ni} + t_{n,i+1})/2$. For the sliding window average, we used $t_v = (t_{ni} + t_{n,i+1})/2$. Each individual ISI was plotted in a scatter plot as a point with coordinates (t_{ni}, τ_{ni}) . The mean firing rate f is defined as the number of spikes divided by the duration of the measurement interval. The mean firing rate is determined for each time interval (I–III) for all the N network neurons together, as well as separately for the N_A activated neurons and the N_B nonactivated neurons.

2.2.2 Binned Spike Train Statistics. A binned representation $X_n(t)$ of the spike train of the n th neuron is obtained by setting $X_n(t) = 1$ when there is a spike time between $t - \Delta t/2$ and $t + \Delta t/2$ and $X_n(t) = 0$ otherwise. We took $\Delta t = 1$ ms. The spike time histogram, $STH = (1000\Delta t/N)X(t)$, is proportional to the sum of X_n over all neurons, $X = \sum_{n=1}^N X_n$ (since Δt is expressed in ms, the factor 1000 is necessary so that STH is in Hz). The local field potential (LFP) is estimated by filtering X by an exponential filter, $\exp(-t/\tau)/\tau$, with timescale $\tau = 5$ ms. The sliding window average, $\sigma_{LFP}(t)$, of the standard deviation of the LFP was used as a synchrony measure. We also determined the normalized autocorrelation, $AC(t) = \langle (X(s) - \bar{X})(X(s-t) - \bar{X}) \rangle_s / \langle (X(s) - \bar{X})(X(s) - \bar{X}) \rangle_s$, where $\langle \cdot \rangle_s$ denotes the time average and $\bar{X} = \langle X \rangle_s$.

2.2.3 Spectral Measures for Neural Synchrony. The coherence of spike trains with the underlying network oscillations was determined using the spike-triggered average (STA) (see Rieke, Warland, de Ruyter van Steveninck, & Bialek, 1997; Dayan & Abbott, 2001, for details). The STA was calculated by taking the LFP from the full network centered around the combined spike times of four neurons, two of which were part of the activated group and the other two were part of the nonactivated group. The STA power spectrum was calculated using the MATLAB routine `psd` with standard windowing using a 2,048 point, fast Fourier transform. The spike field co-

herence (SFC) is the STA power spectrum divided by the power spectrum of the LFP (Fries et al., 2001). The SFC for the standard simulations was calculated based on 300 ms long LFP segments (IB: $700 \text{ ms} \leq t \leq 1000 \text{ ms}$ and IIB: $1200 \text{ ms} \leq t \leq 1500 \text{ ms}$). This was necessary since synchrony was modulated on a fast timescale. Simulations were also performed for the situation where the baseline state and the activated state both lasted for 2000 ms. The SFC in that case was calculated based on 1500 ms long segments.

2.2.4 Event-Based Measures for Neural Synchrony. During synchronized network activity, the histogram $X(t)$ contains peaks of high activity—events—during which many neurons spike at approximately the same time. Spike times t_v were clustered in events using MATLAB's fuzzy k -means clustering algorithm, fcm (Ripley, 1996). For each cluster k , the algorithm returns the indices $v_k, v_k + 1, \dots, v_{k+1} - 1$ of the spike times that are part of the cluster. The number of spike times $N_k = v_{k+1} - v_k$ in event k and their standard deviation σ_k was determined. For low synchrony, peaks in the STH are broad and overlap; as a result, the clustering procedure could not reliably cluster the spikes into events.

2.2.5 Interspike Distance Synchrony Measures. The CV_P measure is based on the idea that during synchronous states, the minimum distance between spikes of different neurons is reduced compared with asynchronous states. The interspike interval of the combined set of network spikes is $\tau_v = t_{v+1} - t_v$. Note that these interspike intervals are between different neurons. The CV is

$$CV_P = \frac{\sqrt{\langle \tau_v^2 \rangle_v - \langle \tau_v \rangle_v^2}}{\langle \tau_v \rangle_v},$$

where P stands for population and $\langle \rangle_v$ denotes the average over all intervals. As with regular ISIs, the interval τ_v can be identified with three times: t_v , t_{v+1} , and the mean $(t_v + t_{v+1})/2$. For the sliding window average, $CV_P(t)$, the last was used.

In an asynchronous network, each neuron fires independently at a constant rate. The sum of the spike trains is to a good approximation a Poisson spike train, with a coefficient of variation, CV_P , equal to one. For a fully synchronous network with N active neurons oscillating with period T , the set τ_v consists of $N_s/N - 1$ intervals equal to T and $N_s - N_s/N$ intervals equal to $\Delta \rightarrow 0$. Hence,

$$CV_P = \sqrt{\alpha} [1 - (\alpha/T + T/\alpha)\beta\Delta],$$

with $\alpha = (N_s - 1)/(N_s/N - 1)$ and $\beta = (N_s/N)(N - 1)/(N_s - 1)$. For large N_s , this reduces to $CV_P \approx \sqrt{N}$. Thus, $(CV_P - 1)/\sqrt{N}$ is a measure for synchrony normalized between 0 and 1. CV_P is sensitive to the precision of the network discharge as well as to the degree of coincidence.

The nearest-neighbor distance of spike t_{ni} from the spikes of neuron m is

$$\Delta_{mni} = \min_j |t_{ni} - t_{mj}|.$$

The minimum is taken over all spike times t_{mj} of neuron m . There are $N_p = (N - 1)N_s$ different pairs. Note that each pair is counted twice, $\Delta_{mni} = \Delta_{nmj}$, because it is associated with t_{ni} as well as with t_{mj} .

The coincidence factor κ is defined as the number of $\Delta_{mni} < P$, divided by the maximum number of coincidences N_c . Here, P is a preset precision, typically equal to 2 ms. The number of $\Delta_{mni} < P$ cannot exceed N_p ; hence, taking $N_c = N_p = (N - 1)N_s$ would yield a number between 0 and 1 (note that N_s is now the total number of spikes in the time interval under consideration). However, when the minimum single-neuron interspike interval is larger than P , the number of coincidences in a pair of neurons (n, m) cannot be larger than the minimum of the number of spikes produced by neurons n and m , formally equal to $\min(\sum_t X_n(t), \sum_t X_m(t))$. Hence, the normalization should be

$$N_c = \sum_{n \neq m} \min \left(\sum_t X_n(t), \sum_t X_m(t) \right). \quad (2.3)$$

During intervals I and III, when the network dynamics is homogeneous, $N_p \approx N_c$ (the difference is 10% or less). However, when there is a bimodal distribution of firing rates, this approximation is not appropriate. The sliding window average $\kappa(t)$ was calculated based on Δ_{mni} , as described above. The density of coincidences was also determined at a finer temporal resolution. $K(t)$ was the number of $\Delta_{mni} < P$ with $t - \Delta t/2 \leq t_{ni} \leq t + \Delta t/2$ normalized by N_p (N_c could not be estimated reliably for these small intervals).

The κ used here is related to the quantity described in Wang and Buzsáki (1996), denoted here by κ_{WB} :

$$\kappa_{WB} = \frac{1}{N(N-1)} \sum_{n \neq m} \frac{\sum_t X_n(t) X_m(t)}{\sqrt{\sum_t X_n(t) \sum_t X_m(t)}}, \quad (2.4)$$

where \sum_t denotes the sum over the discrete bins (width Δt). The measure κ used here differs from κ_{WB} in three aspects:

1. κ_{WB} is based on discrete bins. Hence, it is possible that two spike times with a distance less than Δt are in different bins and considered noncoincident (White, Chow, Ritt, Soto-Treviño, & Kopell, 1998).
2. The normalization of κ_{WB} does not work if the mean firing rates of neurons are significantly different. In that case, κ_{WB} would be smaller than 1 even if all the spikes were coincident. This is because the maximum number of coincident spikes is $\min(\sum_t X_n(t), \sum_t X_m(t))$ rather than $\sqrt{\sum_t X_n(t) \sum_t X_m(t)}$, used in equation 2.4. The normalization by N_c of $\kappa(t)$ accounts for these effects.

3. The calculation of $\kappa(t)$ is better suited to track fast changes in synchrony.

An estimate for the dispersion of spike times in an event can be obtained from Δ_{mni} . For a given spike time t_{ni} , we take the mean of the N_Δ lowest values of Δ_{mni} . This value is averaged across all t_{ni} in an averaging time interval of length T_{av} to obtain $\Delta(t)$. When N_Δ is somewhat smaller than the number of neurons firing on each cycle, $\Delta(t)$ is proportional to the dispersion in spike times. Typically it is slightly less than the dispersion obtained directly from clustering. The relationship between the value of Δ and the dispersion of the spike times on a cycle was studied using a set of $N = 300$ deviates of a gaussian probability distribution with a variance of 1. Δ was equal to 0.53 ± 0.02 for $N_\Delta = 150$, asymptoted to 1.13 ± 0.05 for $N_\Delta = 300$, and reached 1 for $N_\Delta = 278$ (errors were the standard deviation across 100 different realizations).

3 Results

3.1 Competition Leads to Dynamical Modulation of Synchrony. We study the dynamics of an all-to-all connected interneuron network. Key parameters are mean level of depolarizing current I , variance $2D$ of the white noise current, and neuronal heterogeneity modeled as a dispersion σ_I in the driving current across neurons. These three parameters can account for changes occurring *in vivo* in the level of background synaptic activity and the effect of neurotransmitters and neuromodulators. Previous work shows that synchrony is stable only up to $\sigma_I \approx 0.10 \mu\text{A}/\text{cm}^2$ (Wang & Buzsáki, 1996) and $D \approx 0.10 \text{mV}^2/\text{ms}$ (Tiesinga & José, 2000a); for higher values, only the asynchronous state is stable. For moderate values of σ_I and D (below these critical values), synchrony is stable only for a large enough driving current I .

Based on these results, it is expected that applying a current pulse to the network in a low-synchrony state will lead to increased synchrony. Representative simulation results are shown in Figure 1. There were $N = 1000$ neurons in the network; the other parameters were $D = 0.02 \text{mV}^2/\text{ms}$ and $\sigma_I = 0.02 \mu\text{A}/\text{cm}^2$. Initially, the driving current was $I = 0.7 \mu\text{A}/\text{cm}^2$. In this baseline state, the mean firing rate per neuron was $f = 8.18 \text{Hz}$ with a $\text{CV} = 0.38$. During the time interval between 1000 ms and 1500 ms, the drive to all neurons was increased to $I = 1.0 \mu\text{A}/\text{cm}^2$. The firing rate increased by approximately 50% to $f = 12.25$; the CV was 0.25 (see Figure 1A). There was little change in synchrony. In contrast, when the depolarizing current was increased to only 250 of the 1000 interneurons, a robust increase in synchrony was observed (see Figure 1B). During that period, the mean firing rate increased only weakly from $f = 8.18 \text{Hz}$ to $f = 9.48 \text{Hz}$ ($\text{CV} = 0.40 \rightarrow 1.55$).

It is also possible to induce synchrony by increasing the driving current to all neurons in the network. However, the necessary increase in current is

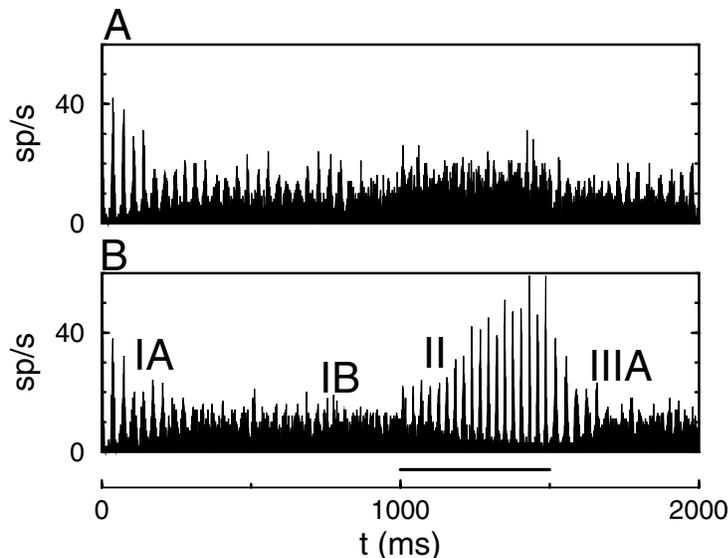


Figure 1: Competition dynamically modulates synchrony. (A) Driving current $I = 0.7 \mu\text{A}/\text{cm}^2$ was increased to $I = 1.0 \mu\text{A}/\text{cm}^2$ during the time interval between 1000 ms and 1500 ms (indicated by the bar in B). (B) Driving current to $N_A = 250$ out of $N = 1000$ neurons was increased from $I = 0.7 \mu\text{A}/\text{cm}^2$ to $I = 1.0 \mu\text{A}/\text{cm}^2$. The other neurons remained at baseline. Roman numerals in B are explained in the text. Parameters for these simulations were $\sigma_I = 0.02 \mu\text{A}/\text{cm}^2$, $D = 0.02 \text{ mV}^2/\text{ms}$, $g_{\text{syn}} = 0.4 \text{ mS}/\text{cm}^2$, and $\tau_{\text{syn}} = 10 \text{ ms}$.

much higher than used in Figure 1A, and the resulting firing rate increase is also much higher compared with the competitive mechanism. In the remainder, we investigate the properties of the synchrony-by-competition mechanism.

3.2 Synchrony Modulation by Competition. We have labeled four distinct time intervals in the histogram shown in Figure 1B. The rastergram for all neurons and histograms for the activated and non-activated neurons are shown separately in Figure 2 for the time interval between 500 ms and 2000 ms. (IA) $0 < t < 500 \text{ ms}$. The membrane potential of the neurons at $t = 0$ was drawn for each neuron independently from a uniform distribution between -70 mV and -50 mV . Initially, the network was partially synchronized due to the particular realization of the initial condition. The amount of synchrony decreased over the course of 500 ms. (IB) $500 < t < 1000 \text{ ms}$: During the baseline state there was only weak synchrony. (II) $1000 < t < 1500 \text{ ms}$: A subset of $N_A = 250$ neurons gets activated. Their firing rate increased, and they became synchronized within

200 ms. Initially, the $N_B = 750$ nonactivated neurons were completely suppressed by the low-synchronous activated neurons, but as the synchrony of the activated neurons increased, the nonactivated neurons resumed firing, albeit at a much lower rate. (IIIA) $1500 < t < 1700$ ms: The additional drive to the activated neurons was reduced back to baseline. Their firing rate decreased, which released the nonactivated neurons from inhibition. Because their membrane potentials were similar, the network was transiently synchronized. Over the course of 200 ms, the network returned to the asynchronous state. During the baseline state, there were brief episodes of spontaneously enhanced synchrony (see below).

The changes in network dynamics are mimicked in the distribution of interspike intervals (see Figure 2B). The interspike intervals (ISI) are shown as dots in a scatter plot (see Figure 2Ba). The sliding window averages of the mean ISI and its coefficient of variation are shown in Figures 2Bb and 2Bc, respectively. In the baseline state, the ISI distribution was homogeneous, but it had a large dispersion (see Figure 2Ba). However, during the synchronized period, the distribution became bimodal: a set of short intervals with a small dispersion corresponding to the activated neurons and a set of long intervals with a large dispersion due to cycle skipping associated with the nonactivated neurons (see Figure 2Ba). The sliding window average of the mean ISI was sensitive to the onset and offset of bimodality associated with synchrony (see Figure 2Bb). Similarly, the CV of the intervals increased during the synchronized period because of the onset of bimodality in the ISI distribution (see Figure 2Bc). However, these time-resolved statistics are not a good measure of synchrony, since synchrony may increase without introducing bimodality in the ISIs.

3.3 Spectral Analysis of Synchrony. In the experimental analysis of attentional modulation of synchrony, there was an increase in the coherence of multiunit spike trains with the local field potential (LFP) in response to the attended stimuli (Fries et al., 2001). Here we used a filtered version of the spike time histogram (STH) as an estimate for the LFP, and we used the spikes obtained by combining spike trains of four neurons (two activated, two nonactivated) as the multiunit spike train (see Figure 3A). The autocorrelation function $AC(t)$ of the spike time histogram during a 300 ms long synchronized period showed strong periodic peaks, whereas during an equally long asynchronous period, the AC showed periodic modulations but at a much lower amplitude (see Figure 3B). Hence, the AC was sensitive to synchrony. The spike-triggered average (STA) of the LFP triggered on the multiunit spike train was also determined for the same period. The spike field coherence (SFC) is the power spectrum of the STA divided by that of the LFP. For the 300 ms segments used before, there was no clear difference in the SFC between the synchronized and asynchronous segments across a single trial (results not shown). The analysis was also performed on 1500 ms long segments from a longer simulation of the same network (see

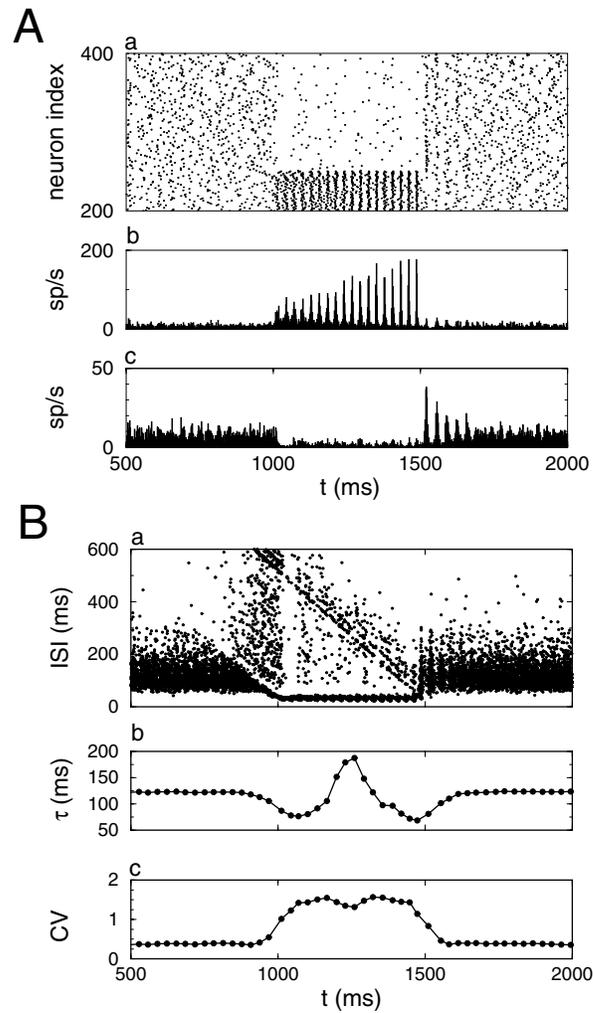


Figure 2: Mechanism for synchrony modulation by competition. (A) Firing rate dynamics. (a) Rastergram showing the spike trains of 200 out of the 1000 neurons. The bottom 50 neurons were activated by the current pulse, and the top 150 neurons received baseline current throughout the simulation. Spike time histogram of the (b) 250 activated neurons and (c) 750 nonactivated neurons. (B) Dynamics of the interspike interval (ISI) statistics. (a) Scatter plot: The y -coordinate is the length of the interval, and the x -coordinate is the starting spike time of the interval. Sliding window average of the (b) mean τ and (c) coefficient of variation CV of the interspike intervals. Data are from the simulation results shown in Figure 1B.

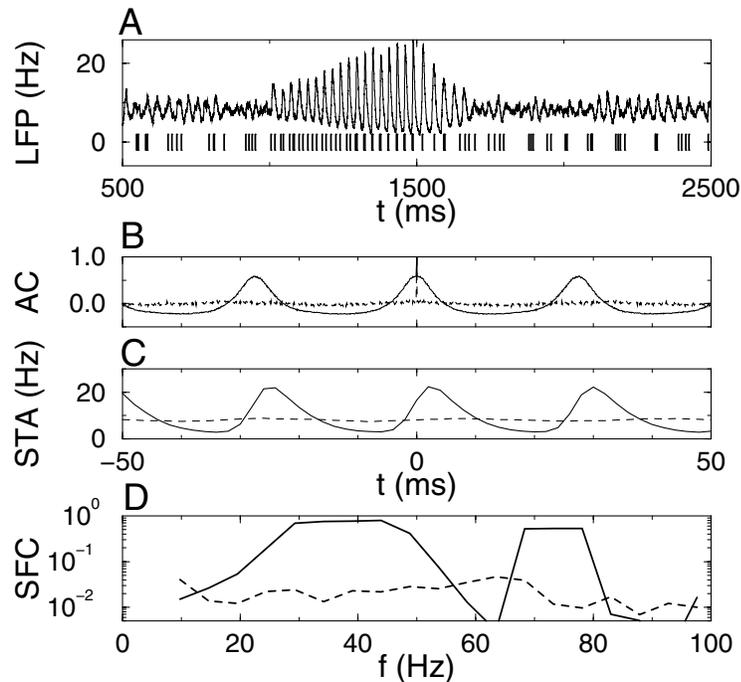


Figure 3: Spike field coherence (SFC) is increased during the synchronous episode. (A, top) The simulated local field potential (LFP) was calculated by convolving the spike time histogram with an exponential filter (time constant 5 ms). (A, bottom) Simulated multiunit recording consisting of the combined spike trains of two activated and two nonactivated neurons. (B) Autocorrelation of the spike time histogram during (solid line) synchronous period (IIB, 1200–1500 ms) and (dotted line) asynchronous period (IB, 700–1000 ms). (C, D) Spike-triggered average (STA) of the LFP triggered on the simulated multiunit recording and the resulting SFC during (solid line) a 1500 ms long synchronous period and (dashed line) an asynchronous period of equal length. Data for A and B are from the simulation results shown in Figure 1B. Data in C and D are from a 6000 ms long simulation with the same parameters, except that the current pulse was applied during the time interval between 2000 and 4000 ms. The synchronous period used for analysis was between 2500 and 4000 ms, and the asynchronous period was 4500 to 6000 ms.

Figures 3C and 3D). The SFC in the gamma frequency band (30–60 Hz) was two orders of magnitude larger for the synchronized segment compared with the asynchronous segment. For our simulation data, the SFC is not a robust measure for changes in synchrony on a 100 ms timescale, but it does work well when longer time intervals are available for averaging.

3.4 Spike-Distance Based Synchrony Measures. During the synchronous state, many neurons fire at similar times; hence, the distance between spikes of different neurons is reduced compared with the asynchronous state. We calculate two synchrony measures using this fact (see section 2).

The coefficient of variation CV_P is obtained by pooling the spike times of all neurons, ordering them from low to high values and then calculating the coefficient of variation of the resulting interspike intervals (see section 2). For an asynchronous population, the combined spike train will form a Poisson process with $CV_P = 1$. For a synchronized network, there will be many small intervals between spikes in a peak and a few long intervals between the first spike of the next peak and the last spike of the preceding peak, yielding $CV_P \sim \sqrt{N}$ (here, N is the number of neurons that fire on a given cycle). Note that CV_P is different from the CV shown in Figure 2Bc: the latter is based on the interspike intervals between spikes of the same neurons. The sliding window average of CV_P could resolve fast changes in neural synchrony and was sensitive to the transient synchrony at the start of the simulation and an episode of spontaneous synchrony at $t \approx 2800$ ms (the circle in Figure 4A). The oscillations in the LFP also increased during synchrony. Hence, the sliding window average σ_{LFP} of the standard deviation of the LFP was also sensitive to changes in synchrony (see Figure 4B). There are two different aspects of synchrony: the number of coincidences (the number of neurons firing at approximately the same time) and the temporal dispersion of the spike times. CV_P and σ_{LFP} confound these aspects.

For CV_P , only the distance between a given spike and the closest other spike time is taken into account. A new time-resolved coincidence measure that is an extension of CV_P takes into account the nearest spike time for all other neurons separately. Hence, for each spike t_{ni} (i th spike of n th neuron), there are $N - 1$ nearest distances Δ_{mni} ($m = 1, \dots, n - 1, n + 1, \dots, N$; see section 2). We determined the set of Δ_{mni} smaller than $P = 2$ ms. The density $K(t)$ of the coincident pairs and the sliding window average $\kappa(t)$ are shown in Figures 4C and D, respectively. $\kappa(t)$ was normalized by the number of expected coincidences, as explained in section 2. Note that when $\kappa(t)$ was normalized by the total number of pairs (N_p ; see section 2), it significantly underestimated the number of coincidences during interval II (filled circles versus solid line in Figure 4D). An estimate of the jitter $\Delta(t)$ (the inverse of the precision) was obtained by calculating the mean distance of the $N_\Delta = 100$ nearest spike times to a given spike time and averaging over all spike times (see Figure 4E).

The number of coincidences and the dispersion of spike times were calculated using a clustering algorithm for comparison purposes. During the synchronous state, peaks of elevated firing rate occur in the STH (see Figure 5A). The spikes that are part of a peak were determined using a fuzzy k -means clustering algorithm. The number of peaks during a particular time interval was determined by visual inspection and supplied as a parameter to the algorithm. The number N_k of neurons firing on a particular cycle and

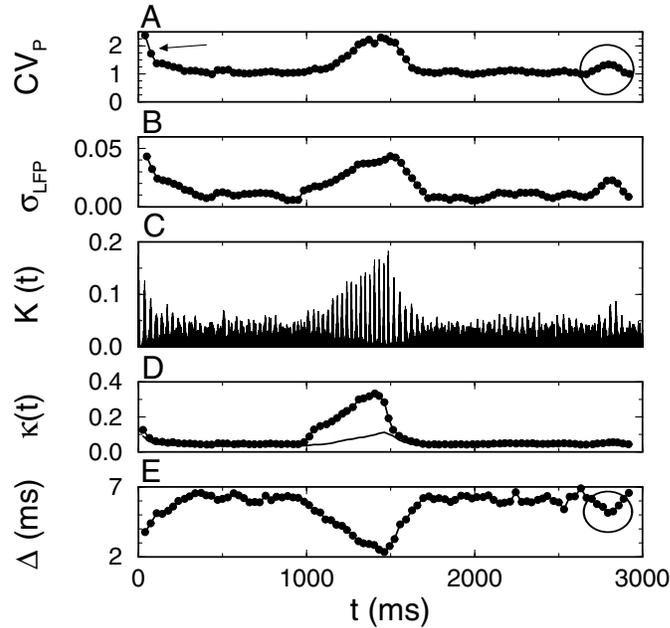


Figure 4: Spike distance measures can track synchrony modulation on fast timescales. (A) Sliding window average of the coefficient of variation (CV_P) of the interspike intervals between all the network neurons. The arrow points to the decay of transient synchrony due to the initial condition, and the circle highlights an episode of spontaneously elevated synchrony. (B) Sliding window average σ_{LFP} of the standard deviation in the local field potential (LFP). The number of spike pairs with a distance Δ_{mi} smaller than 2 ms plotted as (C) a histogram $K(t)$ (bin width was 1 ms) and (D) a sliding window average $\kappa(t)$. The small circles in *D* are κ normalized by N_c , and the solid line is κ normalized by N_p . (E) The sliding window average $\Delta(t)$ of the mean distance of the $N_\Delta = 100$ closest spike times to the given spike time. The analysis in *A* and *B* used all the spike data, whereas the analysis in *C* and *E* was based on the spike trains of 125 activated neurons and 375 nonactivated neurons. Data are from the simulation results shown in Figure 1B.

their temporal dispersion σ_k were calculated and plotted as a function of the centroid of the peak (see Figure 5B). After the onset of the current pulse, the jitter σ_k decreased over time, whereas the number of neurons N_k firing per cycle remained approximately constant. Likewise, when the current pulse was turned off, the jitter increased and did so at a much faster rate of change. Changes in synchrony can be resolved on a cycle-by-cycle basis using the jitter. The temporal modulation of $\Delta(t)$ (see Figure 4E) closely mimics that of σ_k , shown in Figure 5B.

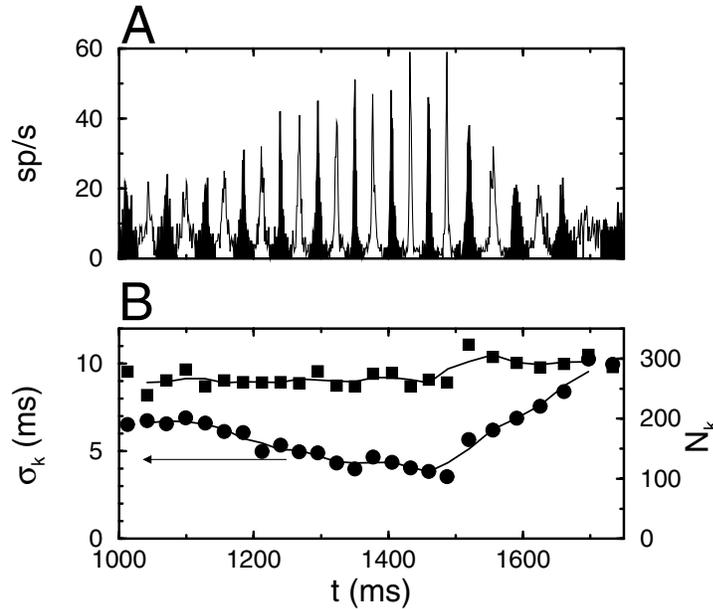


Figure 5: Clustering of spike times during the onset of the synchronized period. Each event corresponds to a cycle of the network oscillation. (A) Spike time histogram with the odd-numbered clusters shown in black and the even-numbered in white. (B, left-hand side scale) The standard deviation σ_k of the spike times and (B, right-hand side scale) the number N_k of neurons that spiked during a cycle plotted versus the mean spike time during a cycle. The oscillation period between 1200 ms and 1500 ms was 27.5 ms. Data are from the simulation results shown in Figure 1B.

3.5 Synaptic Background Activity is Necessary for Dynamic Changes in Synchrony. High synchrony occurs when the network discharge has a jitter (σ_k or Δ) of less than 5 ms; for low synchrony, the network discharges with strong temporal modulation of the firing rate but with a jitter of more than 5 ms. Asynchronous network discharge does not have significant firing rate modulations.

The neurons in the network received incoherent background activity, which was studied by systematically varying the variance D of a balanced white noise current. Increasing the variability in the input current without increasing the mean input current corresponded to a larger value of D (Tiesinga, José, & Sejnowski, 2000; Chance, Abbott, & Reyes, 2002).

For each value of D , a current pulse of $\Delta I = 0.3 \mu\text{A}/\text{cm}^2$ was applied to 250 out of 1000 neurons during the time interval between 1000 ms and 1500 ms. As before, the simulation was run for 3000 ms and divided into

three intervals, I–III. During the first (I) and third (III) intervals, the network was in its baseline state, and during the second (II) interval, it was in its activated state. The spike time histograms are shown in Figure 6A.

For $D = 0 \text{ mV}^2/\text{ms}$, the network synchronized during the first interval and remained synchronized throughout the second and third intervals. During the second interval, the network splits into two populations with different firing rates ($f = 38 \text{ Hz}$ for activated neurons, whereas the nonactivated population spiked on only 1 in 13 cycles of the network oscillation). Because there was neural heterogeneity, $\sigma_I = 0.02 \mu\text{A}/\text{cm}^2$, the jitter σ_k was nonzero, but CV_p was close to its fully synchronized value, \sqrt{N} .

For $D = 0.0004 \text{ mV}^2/\text{ms}$, network synchrony was stable for the baseline network state. Starting from a synchronized state, the network remained synchronized, and starting from an asynchronous state, it became synchronized over time. During the first interval, the network synchronized; synchrony increased even more during the second interval, and in the third interval the degree of synchrony continued to increase. At the start of the third interval, the nonactivated neurons all had a similar voltage. Once these neurons were released from the inhibition projected by the activated neurons, they immediately synchronized into a two-cluster state (see below). The mean firing rate was doubled compared with that during the first interval.

For $D = 0.0084 \text{ mV}^2/\text{ms}$, the network converged to a low-synchrony state during the first interval. The network switched to a high-synchrony state during the second interval. In the third interval, the network returned slowly to the low-synchrony state. The dynamics for $D = 0.0164 \text{ mV}^2/\text{ms}$ was similar, except that the return to the low-synchrony state during the third interval was faster. The same holds for $D = 0.0284 \text{ mV}^2/\text{ms}$, except that the degree of synchrony reached during the second interval was reduced.

For $D = 0.0364 \text{ mV}^2/\text{ms}$, the network was in an asynchronous state during the first and third intervals. The current pulse during the second interval was unable to drive the network to a synchronous state.

The network dynamics for different values of D is also reflected in the time-resolved coincidence measure $\kappa(t)$ and jitter $\Delta(t)$. The state with dynamic modulation of synchrony is characterized by low values of $\kappa(t)$ and high values of $\Delta(t)$ during the first and third intervals and high values of $\kappa(t)$ and low values of $\Delta(t)$ during the second interval. This is the case only for the curves with $D = 0.084 \text{ mV}^2/\text{ms}$ and $D = 0.0204 \text{ mV}^2/\text{ms}$ (see Figure 6B).

In Figure 7, the mean firing rate and synchrony measures CV_p , κ , and Δ are plotted versus D for the intervals IB, II, and IIIB. For $D < 0.005 \text{ mV}^2/\text{ms}$, there was multistability. Depending on the distribution of initial voltages across the neurons, the network converged to states with two, three, or more clusters present (these clusters should not be confused with those obtained

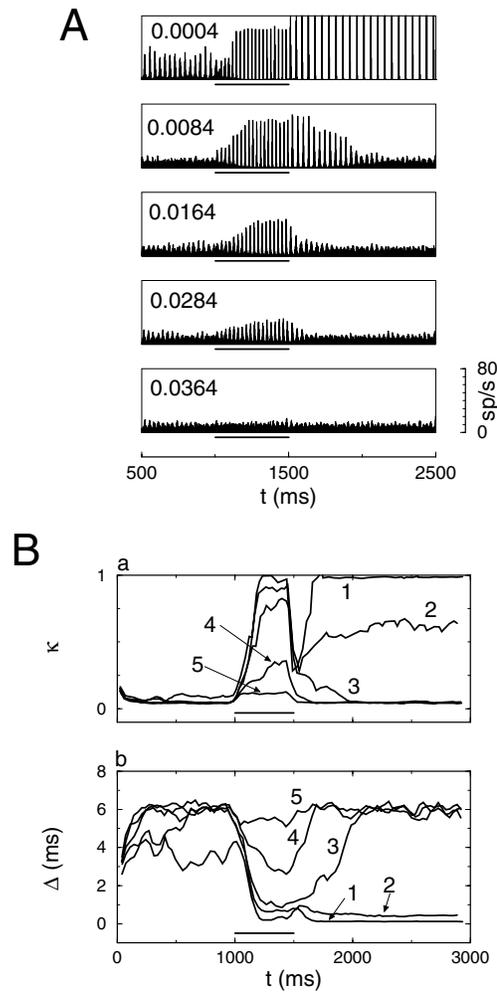


Figure 6: Synaptic background activity is required for fast synchrony modulation. (A) Five spike time histograms with D increasing from top to bottom. The value of D is indicated in the top left corner of each graph. (B) Sliding window average of (a) the number $\kappa(t)$ of coincidences and (b) the estimate $\Delta(t)$ of spike time dispersion for $D = (1) 0.0004, (2) 0.0044, (3) 0.0084, (4) 0.0204$ and $(5) 0.0364$ (D values are in mV^2/ms). A current pulse $\Delta I = 0.3 \mu A/cm^2$ was applied to 250 of the 1000 neurons during the time interval between $t = 1000$ ms and 1500 ms (indicated by the bar in each graph). Network parameters are the same as in Figure 1B except that D is varied.

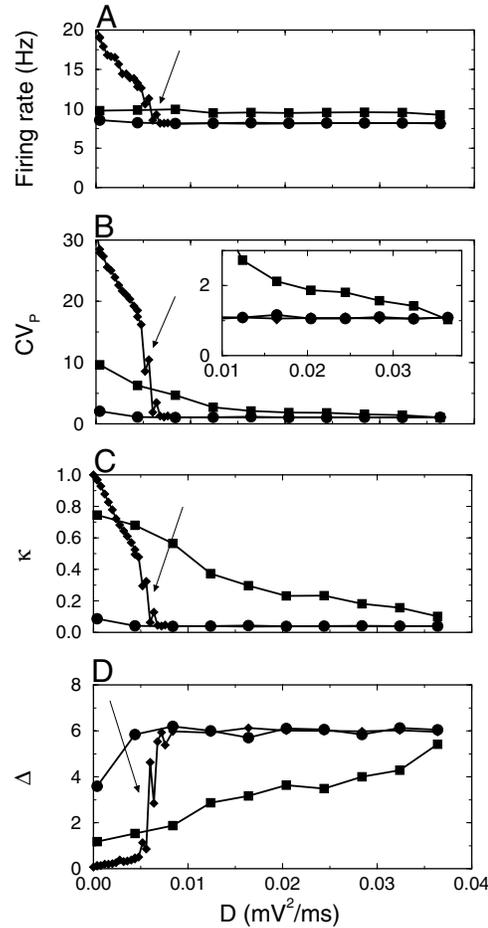


Figure 7: Optimal level of synaptic background activity required for synchrony modulation. (A) The mean firing rate f and (B) CV_p (C) κ and (D) Δ versus noise strength D for three time intervals: (IB, circles) 500–1000 ms, (II, squares) 1000–1500 ms and (IIIB, diamonds) 2000–3000 ms. The inset in B is a close-up. Network parameters are the same as in Figure 1B except that D is varied.

in the preceding section using a clustering procedure). For instance, in a two-cluster state, neurons fired once every two cycles, yielding a firing rate of approximately 20 Hz. Each neuron was in the cluster that was active on either even cycles or odd cycles. The dynamics of states with a higher number of clusters is similar (Kopell & LeMasson, 1994). The number of clusters in IB depended on the initial condition. f , CV_p , κ , and Δ had different val-

ues depending on the number of clusters. The graphs of these quantities versus D would have a ragged appearance due to the sensitivity to initial conditions. Hence, we did not plot them in Figure 7 for small D values in intervals I and II. However, after the application of the current pulse during the second interval, the network settled into the most synchronous state with the highest firing rate. As a result, the network state reached in IIIB did not depend on initial conditions. For higher D values, these synchronized states became unstable, and the firing rate as well as the values of CV_p and κ in the third interval dropped precipitously (see Figure 7, arrow).

In summary, synchrony can be modulated in time only for a finite region of background activity consistent with D values around $D = 0.02 \text{ mV}^2/\text{ms}$. For weaker background activity, it is possible to increase synchrony, but it cannot be shut off. Furthermore, the firing rate during synchrony is much higher compared with that for lower synchrony. For stronger background activity, synchrony cannot be established.

4 Discussion

4.1 Measuring Synchrony. Physiological correlates of attention have been observed in the response of cortical neurons in macaque monkeys to either orientated gratings (area V4) or random dot patterns moving in different directions (area MT) (McAdams & Maunsell, 1999; Treue & Martinez Trujillo, 1999). When attention is directed into the receptive field of the recorded neuron, the gain of the firing rate response may increase (McAdams & Maunsell, 1999; Treue & Martinez Trujillo, 1999), as well as the coherence of the neuron's discharge with the local field potential (Fries et al., 2001). The coherence was quantified using the SFC, the power spectrum of the spike triggered average divided by that of the local field potential. The experimental recordings were 500 ms to 5000 ms in duration—long enough to get statistical differences with attention focused either into the receptive field or outside the receptive field. Here we study how synchrony can be modulated dynamically as it would be during normal behavior. The speed with which attention can be covertly shifted is quite rapid—in the range of a few hundred ms (Shi & Sperling, 2002). Attentional shifts caused by sudden changes in stimulus properties, such as brightness and motion, can be as low as 100 ms. For our simulations, the SFC across a single trial could not resolve rapid changes in synchrony that were obvious from visual inspection of the spike time histogram.

For oscillatory synchrony in the gamma frequency range, the fastest timescale on which synchrony can change is the cycle length. Event-based methods could resolve changes in synchrony on these timescales (see Figure 5). We introduced two nonevent-based measures that could track the temporal dynamics of synchrony on timescales of about 100 ms: the sliding window average of the standard deviation of the LFP and CV_p . These measures were able to resolve the changes in synchrony that were visible in

the spike time histogram. These measures confound two different aspects of synchrony: the number of coincident spikes per event and their temporal dispersion corresponding to N_k and σ_k , respectively, in the event-based analysis. Either increasing the number of spikes or decreasing their temporal dispersion will lead to increases in the value of σ_{LFP} or CV_p . However, the distribution of the pairwise distances between spikes used to define κ and Δ distinguishes between changes in the number of coincident neurons and the temporal dispersion. There were two parameter values— P , the precision of the coincidences, and N_Δ , the number of pairwise distances—in the average for Δ . For κ , all $N_p = (N - 1)N_s \sim N^2 f T_{\max}$ pairwise distances need to be determined. The computational load increases quadratically with network size N and linearly with firing rate f and the length T_{\max} of the measurement interval; hence, it is less efficient than CV_p . $\kappa(t)$ is a time-resolved version of the quantity used by Wang and Buzsáki (1996). We adapted κ so that it was correctly normalized for networks with a bimodal distribution of firing rates.

In summary, CV_p is a fast and easy statistic to quantify synchrony modulations, but $\kappa(t)$ and $\Delta(t)$ delineate changes in temporal dispersion from changes in the number of coincident spikes.

4.2 Modulating Synchrony. The synchronization properties of model networks of interneurons has recently been intensively investigated (Traub, Whittington, Colling, Buzsáki, & Jeffreys, 1996; Wang & Buzsáki, 1996; White, Chow, Ritt, Soto-Treviño, & Kopell, 1998; Chow, 1998; Ermentrout & Kopell, 1998; Brunel & Hakim, 1999; Kopell, Ermentrout, Whittington, & Traub, 2000; White, Banks, Pearce, & Kopell, 2000; Tiesinga & José, 2000a; Hansel & Golomb, 2000; Bressloff & Coombes, 2000; Gerstner, 2000; Tiesinga, Fellous, José, & Sejnowski, 2001; Aradi, Santhakumar, Chen, & Soltesz, 2002; Bartos et al., 2002; Borgers & Kopell, 2003; Olufsen, Whittington, Camperi, & Kopell, 2003; Hansel & Mato, 2003; Fransen, 2003). The degree of synchrony, and its robustness against noise (D) and heterogeneity (σ_I), depends primarily on the synaptic coupling strength g_{syn} and degree of depolarization I (Wang & Buzsáki, 1996; White et al., 1998; Tiesinga & José, 2000a). For $I = 0.7 \mu\text{A}/\text{cm}^2$, $g_{syn} = 0.4 \text{ mS}/\text{cm}^2$, the network studied here is asynchronous (with $D = 0.02 \text{ mV}^2/\text{ms}$ and $\sigma_I = 0.02 \mu\text{A}/\text{cm}^2$). During the application of a current pulse to the N_A network neurons, the other N_B neurons were almost completely suppressed. Hence, the full network behaves approximately as a network of N_A neurons with effective synaptic coupling $g_{eff} = (N_A/N)g_{syn} = 0.1 \text{ mS}/\text{cm}^2$ and driving current $I = 1.0 \mu\text{A}/\text{cm}^2$. For these parameters, the network is synchronous (Wang & Buzsáki, 1996; Tiesinga & José, 2000a).

Synchrony by competition is effective when the baseline state is asynchronous and reached quickly from a synchronized network state (corresponding to an initial condition with all neurons having a similar mem-

brane potential). The activated state should be synchronous and should be established quickly from an asynchronous state (corresponding to an initial condition with a wide dispersion of membrane potential across different neurons). In a previous study (Tiesinga & José, 2000a), the focus was on the degree of synchronization in the asymptotic state of the network. Here, we study a network with competition between otherwise identical neurons, and the focus is on the speed with which the network can switch between synchronous and asynchronous states.

The strength of the white noise current D , corresponding to the level of balanced synaptic background activity, was critical. In the noiseless network ($D = 0$), there are multiple stable cluster states. For the baseline state, we found states with two, three, or four approximately equal-sized clusters that were reached from the set of random initial conditions used in the simulations (Golomb & Rinzel, 1993, 1994; Kopell & LeMasson, 1994; van Vreeswijk, 1996; Tiesinga & José, 2000b). The cluster states were stable against weak noise. The transition to the activated state could switch the baseline state from one cluster state to another, but it would not affect the degree of synchrony. The firing rate did, however, vary strongly. For strong noise, neither the baseline nor the activated state was synchronized. Only for intermediate noise strengths could the degree of synchrony be modulated. In this regime, the degree of synchrony reached during the activated state decreased with increasing D , but the transition from the synchronized state to the asynchronous state during interval IIIA was speeded up.

For what parameters values can synchrony by competition be obtained? We varied the value of the amplitude of the current pulse, ΔI ; the mean of the common driving current, I ; the number of activated neurons, N_A ; the total number of neurons in the network, N ; and the strength of synaptic coupling, g_{syn} (data not shown). For large networks, $N > 100$, with all-to-all coupling, the network state did not depend on N (Tiesinga & José, 2000a). Let us denote the firing rate of the baseline network by f_1 and that of the N_A neurons in the activated network by f_2 . f_2 should be below but within 10% of the oscillation frequency (40 Hz for gamma frequency range). Mean activity was constant when $f_1/f_2 \approx N_A/N$. Networks with $N_A > 100$ robustly synchronize (Wang & Buzsáki, 1996), and we found that $N/N_A > 2$ works best. The values of I , g_{syn} , and ΔI are less critical: I and g_{syn} should lead to a firing rate f_1 in the appropriate range, whereas $I + \Delta I$ and $(N_A/N)g_{syn}$ should lead to a synchronous state with a firing rate equal to f_2 . This is possible since the higher firing rate state is usually more synchronous (Tiesinga & José, 2000a). The specific value of ΔI determines how fast synchrony is attained in interval II.

These parameter values are specific to the model used in our investigation. This raises the issue whether and to what extent the results reported here generalize to other models. When the maximum conductance of the sodium and potassium currents in the model were varied by less than 10%,

we obtained similar results to those reported here. Hence, there is robustness against small changes in model neuron parameters. The key requirements are that the model network has a synchronous state in the gamma frequency range and an asynchronous state and that the degree of synchrony depends on the coupling strength and driving current. These properties hold across different Hodgkin-Huxley-type models (Wang & Buzsáki, 1996; White et al., 1998) and for leaky integrate-and-fire model neurons (Brunel & Hakim, 1999; Hansel & Golomb, 2000). This suggests that synchrony modulation by competition will also be present in these models.

The firing rate of neurons recorded *in vitro* adapts in response to a sustained tonic depolarizing current: at the onset of current injection, neurons fire at a high rate but the rate decreases and saturates at a lower value (McCormick, Connors, Lighthall, & Prince, 1985; Shepherd, 1998). The role of adaptation in synchronous oscillations has been studied in models (van Vreeswijk & Hansel, 2001; Fuhrman, Markram, & Tsodyks, 2002). We expect that synchrony modulation by competition also works in the presence of adaptation, but that the temporal dynamics of the transition between the activated and nonactivated states may be different. Further investigation of this issue is needed.

4.3 Functional Consequences of Temporal Modulation of Synchrony.

The driving hypothesis of this investigation is that attention modulates the synchrony of cortical interneuron networks. We have shown that inhibitory networks can modulate their synchrony on timescales of the order 100 ms without increasing their mean activity. The question is, How do these synchrony modulations relate to the effects of attention on putative pyramidal neurons recorded *in vivo* (Connor et al., 1996, 1997; Luck et al., 1997; McAdams & Maunsell, 1999; Treue & Martinez Trujillo, 1999; Reynolds et al., 2000; Fries et al., 2001)?

Attentional modulation can be studied in models using the following conceptual framework. The neuron is thought to receive two sets of inputs: predominantly excitatory, stimulus-related inputs from upstream cortical areas and modulatory synaptic inputs (Chance et al., 2002). Some of the modulatory inputs are from local interneuron networks that project to the principal output neurons in cortical layer 2/3 (Galarreta & Hestrin, 2001). The response to stimulus-related inputs can be characterized using the firing rate versus input ($f - I$) curve. The $f - I$ curve is determined by measuring the firing rate while the amplitude I of the tonic depolarizing current is systematically varied (sometimes the firing rate of presynaptic excitatory afferents is varied instead). Modulatory inputs alter the $f - I$ curve, leading to a different firing rate in response to the same input. The modulatory changes fall in two categories: leftward shifts of the $f - I$ curve, leading to higher sensitivity as the neuron can respond to weaker inputs, and gain changes, where the firing rate response to any input is multiplied by a constant factor. The attentional modulation of firing rate tuning curve observed

in McAdams and Maunsell (1999) and Treue and Martinez Trujillo (1999) can be explained as a multiplicative gain change of the $f - I$ curve.

Recent investigations reveal that the gain of the $f - I$ curve is modulated by the amount of tonic inhibition that a neuron receives (Prescott & DeKoninck, 2003; Mitchell & Silver, 2003). However, measurements of the LFP (Fries et al., 2001) indicate that neurons receive temporally patterned inputs in the gamma frequency range. The spike time coherence in the gamma-frequency range increased with attention (Fries et al., 2001). Our hypothesis is that the modulations in gamma-frequency coherence are due to synchrony modulation of local interneuron networks. We studied the $f - I$ curves and the coherence with the LFP of model neurons receiving synchronous inhibition in the gamma-frequency range (José, Tiesinga, Fellous, Salinas, & Sejnowski, 2002). The results are summarized here, full details will be reported elsewhere. We found that changing the degree of synchrony, quantified as the jitter (σ_k in Figure 5), led to gain changes as well as shifts. The mean number of inputs per oscillation cycle (N_k in Figure 5) determined whether the change in gain dominated the shift or vice versa. For $N_k > 50$, the modulation of the $f - I$ curve was shift dominated. In that case, the firing rate of output neurons could saturate for strong inputs, and an increase in synchrony led to enhanced coherence with the LFP but not to an increase in firing rate. For weak inputs, the input synchrony acted as a gate. For low-synchrony inhibitory inputs, the receiving cortical neuron was shut down, whereas for high synchrony, it could transmit stimulus-related information in its output. For $N_k \approx 10$, there were large changes in the gain of the $f - I$ curve that were consistent with the gain modulation of orientation tuning curves reported in McAdams and Maunsell (1999). The modulation of firing rate with input synchrony depended on the oscillation frequency of the interneuron network. We found that the strongest modulation was obtained with oscillations in the gamma-frequency range.

To summarize, synchrony by competition in the gamma-frequency range can selectively gate cortical information flow at approximately constant metabolic expense in the interneuron networks. Further investigation is necessary to determine what type of regulation can be performed using this mechanism.

References

- Aoyagi, T., Takekawa, T., & Fukai, T. (2003). Gamma rhythmic bursts: Coherence control in networks of cortical pyramidal neurons. *Neural Computation*, *15*, 1035–1062.
- Aradi, I., Santhakumar, V., Chen, K., & Soltesz, I. (2002). Postsynaptic effects of GABAergic synaptic diversity: Regulation of neuronal excitability by changes in IPSC variance. *Neuropharmacology*, *43*, 511–522.
- Bartos, M., Vida, I., Frotscher, M., Meyer, A., Monyer, H., Geiger, J., & Jonas, P. (2002). Fast synaptic inhibition promotes synchronized gamma oscillations.

- tions in hippocampal interneuron networks. *Proc. Natl. Acad. Sci.*, 99, 13222–13227.
- Borgers, C., & Kopell, N. (2003). Synchronization in networks of excitatory and inhibitory neurons with sparse random connectivity. *Neural Computation*, 15, 509–538.
- Bressloff, P., & Coombes, S. (2000). Dynamics of strongly-coupled spiking neurons. *Neural Computation*, 12, 91–129.
- Brunel, N. and Hakim, V. (1999). Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural Computation*, 11, 1621–1671.
- Chance, F., Abbott, L., & Reyes, A. (2002). Gain modulation from background synaptic input. *Neuron*, 35, 773–782.
- Chow, C. (1998). Phase-locking in weakly heterogeneous neuronal networks. *Physica D*, 118, 343–370.
- Connor, C., Gallant, J., Preddie, D., & Van Essen, D.C. (1996). Responses in area V4 depend on the spatial relationship between stimulus and attention. *J. Neurophysiol.*, 75, 1306–1308.
- Connor, C., Preddie, D., Gallant, J., & Van Essen, D.C. (1997). Spatial attention effects in macaque area V4. *J. Neurosci.*, 17, 3201–3214.
- Dayan, P., & Abbott, L. (2001). *Theoretical neuroscience*. Cambridge, MA: MIT Press.
- Ermentrout, G., & Kopell, N. (1998). Fine structure of neural spiking and synchronization in the presence of conduction delays. *Proc. Natl. Acad. Sci.*, 95, 1259–1264.
- Fransen, E. (2003). Coexistence of synchronized oscillatory and desynchronized rate activity in cortical networks. *Neurocomputing*, 52–54, 763–769.
- Fries, P., Reynolds, J., Rorie, A., & Desimone, R. (2001). Modulation of oscillatory neuronal synchronization by selective visual attention. *Science*, 291, 1560–1563.
- Fuhrman, G., Markram, H., & Tsodyks, M. (2002). Spike frequency adaptation and neocortical rhythms. *J. Neurophys.*, 88, 761–770.
- Galarreta, M., & Hestrin, S. (2001). Electrical synapses between GABA-releasing interneurons. *Nat. Rev. Neurosci.*, 2, 425–433.
- Gerstner, W. (2000). Population dynamics of spiking neurons: Fast transients, asynchronous states, and locking. *Neural Computation*, 12, 43–89.
- Golomb, D., & Rinzel, J. (1993). Dynamics of globally coupled inhibitory neurons with heterogeneity. *Phys. Rev. E*, 48, 4810–4814.
- Golomb, D., & Rinzel, J. (1994). Clustering in globally coupled inhibitory neurons. *Physica D*, 72, 259–282.
- Hansel, D., & Golomb, D. (2000). The number of synaptic inputs and the synchrony of large sparse neuronal networks. *Neural Computation*, 12, 1095–1139.
- Hansel, D., & Mato, G. (2003). Asynchronous states and the emergence of synchrony in large networks of interacting excitatory and inhibitory neurons. *Neural Computation*, 15, 1–56.
- José, J., Tiesinga, P., Fellous, J., Salinas, E., & Sejnowski, T. (2002). Is attentional gain modulation optimal at gamma frequencies? *Society of Neuroscience Abstract*, 28, 55.6.

- Kopell, N., Ermentrout, G., Whittington, M., & Traub, R. (2000). Gamma rhythms and beta rhythms have different synchronization properties. *Proc. Natl. Acad. Sci.*, *97*, 1867–1872.
- Kopell, N., & LeMasson, G. (1994). Rhythmogenesis, amplitude modulation, and multiplexing in a cortical architecture. *Proc. Natl. Acad. Sci.*, *91*, 10586–10590.
- Luck, S., Chelazzi, L., Hillyard, S., & Desimone, R. (1997). Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *J. Neurophys.*, *77*, 24–42.
- McAdams, C., & Maunsell, J. (1999). Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J. Neurosci.*, *19*, 431–441.
- McCormick, D., Connors, B., Lighthall, J., & Prince, D. (1985). Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. *J. Neurophys.*, *54*, 782–806.
- Mitchell, S., & Silver, R. (2003). Shunting inhibition modulates neuronal gain during synaptic excitation. *Neuron*, *38*, 433–445.
- Moore, T., & Armstrong, K. (2003). Selective gating of visual signals by microstimulation of frontal cortex. *Nature*, *421*, 370–373.
- Olufsen, M., Whittington, M., Camperi, M., & Kopell, N. (2003). New roles for the gamma rhythm: Population tuning and preprocessing for the beta rhythm. *J. Comput. Neurosci.*, *14*, 33–54.
- Prescott, S., & De Koninck, Y. (2003). Gain control of firing rate by shunting inhibition: Roles of synaptic noise and dendritic saturation. *Proc. Natl. Acad. Sci.*, *100*, 2071–2081.
- Rensink, R. (2000). The dynamics representation of scenes. *Visual Cognition*, *7*, 17–42.
- Reynolds, J., Pasternak, T., & Desimone, R. (2000). Attention increases sensitivity of V4 neurons. *Neuron*, *26*, 703–714.
- Rieke, F., Warland, D., de Ruyter van Steveninck, R. R., & Bialek, W. (1997). *Spikes: Exploring the neural code*. Cambridge, MA: MIT Press.
- Ripley, B. (1996). *Pattern recognition and neural networks*. Cambridge: Cambridge University Press.
- Salinas, E., & Sejnowski, T. (2000). Impact of correlated synaptic input on output variability in simple neuronal models. *J. Neurosci.*, *20*, 6193–6209.
- Salinas, E., & Sejnowski, T. (2001). Correlated neuronal activity and the flow of neural information. *Nat. Rev. Neurosci.*, *2*, 539–550.
- Salinas, E. and Sejnowski, T. (2002). Integrate-and-fire neurons driven by correlated stochastic input. *Neural Computation*, *14*, 2111–2155.
- Shepherd, G. (1998). *Synaptic organization of the brain* (4th ed.). New York: Oxford University Press.
- Shi, S., & Sperling, G. (2002). Measuring and modeling the trajectory of visual spatial attention. *Psychological Review*, *109*, 260–305.
- Simons, D. (2000). Current approaches to change blindness. *Visual Cognition*, *7*, 1–15.
- Tiesinga, P., Fellous, J.-M., José, J., & Sejnowski, T. (2001). Computational model of carbachol-induced delta, theta and gamma oscillations in the hippocampus. *Hippocampus*, *11*, 251–274.

- Tiesinga, P., & José, J. (2000a). Robust gamma oscillations in networks of inhibitory hippocampal interneurons. *Network*, *11*, 1–23.
- Tiesinga, P., & José, J. (2000b). Synchronous clusters in a noisy inhibitory network. *J. Comp. Neurosci.*, *9*, 49–65.
- Tiesinga, P., José, J., & Sejnowski, T. (2000). Comparison of current-driven and conductance-driven neocortical model neurons with Hodgkin-Huxley voltage-gated channels. *Phys. Rev. E*, *62*, 8413–8419.
- Traub, R., Whittington, M., Colling, S., Buzsáki, G., & Jeffreys, J. (1996). Analysis of gamma rhythms in the rat hippocampus in vitro and in vivo. *J. Physiol.*, *493*, 471–484.
- Treue, S., & Martinez Trujillo, J. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature*, *399*, 575–579.
- van Vreeswijk, C. (1996). Partial synchronization in populations of pulse-coupled oscillators. *Phys. Rev. E*, *54*, 5522–5537.
- van Vreeswijk, C., & Hansel, D. (2001). Patterns of synchrony in neural networks with spike adaptation. *Neural Computation*, *13*, 959–992.
- Wang, X., & Buzsáki, G. (1996). Gamma oscillation by synaptic inhibition in a hippocampal interneuronal network model. *J. Neurosci.*, *16*, 6402–6413.
- White, J., Banks, M., Pearce, R., & Kopell, N. (2000). Networks of interneurons with fast and slow gamma -aminobutyric acid type A (GABAA) kinetics provide substrate for mixed gamma-theta rhythm. *Proc. Natl. Acad. Sci.*, *97*, 8128–8133.
- White, J., Chow, C., Ritt, J., Soto-Treviño, C., & Kopell, N. (1998). Synchronization and oscillatory dynamics in heterogeneous, mutually inhibited neurons. *J. Comput. Neurosci.*, *5*, 5–16.