

Chapter 2

Neural Synchrony, Behavioral Priming and Processing Efficiency

Firing rates of cellular inputs are not the only thing that can determine a neuron's responses. The relative timing of the spiking inputs can also matter. The capacitance of a cell's membrane allows individual synaptic input currents occurring at near points in time to sum; if enough depolarizing inputs occur over a short period of time, a spike will be elicited. Estimates in cortical cells suggest the peak magnitude of an average EPSP is between 0.1 and 0.5 millivolts, implying that at least 40 EPSPs must occur simultaneously in order to reach spike threshold from the resting potential (e.g. Komatsu, Nakajima, Toyama, & Fetz, 1988; Mason, Nicoll, & Stratford, 1991; Matsumura, Chen, Sawaguchi, Kubota, & Fetz, 1996; Sayer, Friedlander, & Redman, 1990; Thomson, Girdlestone, & West, 1988). This means that if inputs are coordinated or synchronized, a receiving cell may be more likely to spike than if the inputs arrive at random times (e.g. Konig, Engel, & Singer, 1996; Salinas & Sejnowski, 2000, 2001; Singer, 1999). Accordingly, some researchers have proposed that individual cortical cells function as

coincidence detectors, where the relative timing of spiking inputs matters (e.g. Abeles, 1982, 1991; Softky & Koch, 1993). In contrast, other researchers have argued that cells function more as slow spatio-temporal integrators, where it is primarily the firing rates of neural inputs rather than their precise spike times that matter (e.g. Shadlen & Newsome, 1994, 1998). This latter view is motivated by empirical observations that neural responses *in vivo* appear to be "noisy": The timing of individual spikes in cortical cells is highly variable (Softky & Koch, 1993; Tomko & Crapper, 1974), and neural responses to multiple presentations of the same stimulus also vary across presentations (e.g. Dean, 1981; Schiller, Finlay, & Volman, 1976; Tolhurst, Movshon, & Dean, 1983; Vogels, Spileers, & Orban, 1989). The debate between these two basic views is ongoing.

A similar distinction occurs in artificial neural network models between spiking neural networks and firing-rate networks. In spiking neural networks, the spiking property of neurons is included in individual cells, such that both average firing rates of inputs and their relative timing can potentially have an impact on output responses. Firing-rate models, on the other hand, assume that the relative timing of spiking is unimportant and that only the average firing rates of neural inputs matter. Indeed, the dynamics of spiking networks can be accurately approximated by average firing rate when spiking is relatively random and asynchronous across cells (e.g. Amit & Tsodyks, 1991; Gerstner, 1995; Wilson & Cowan, 1972; see Ermentrout, 1998b, for a review). It is worth noting that the distributed connectionist models discussed in **Section 1.4.3** are a particular form of firing-rate neural network. However, to the extent that neurons spike at the same time, average firing rate approximations become poor for describing the detailed activity dynamics of a population of cells.




Figure 2.1: Spike synchrony across a population of randomly spiking cells leads to fluctuations in the population-averaged firing rate. In contrast, asynchronous or independent random firing across the same population leads to a much more stable estimate of firing rate. The results shown are for 1000 spiking cells, each firing at 50 Hz with correlated random (partially synchronized: every cell shares identical spike times with 5% of the other cells) or independent random spike times generated from a homogenous Poisson distribution. Firing rate across the population was calculated in 10-ms bins at increments of 10 ms. The asynchronous case better meets the stationary requirement of input spiking for the derivation of firing-rate from spiking networks (e.g. Amit & Tsodyks, 1991; Gerstner, 1995; Wilson & Cowan, 1972). Firing rate networks will therefore be better approximations of spiking networks in this case.

This is depicted graphically in Figure 2.1 for 1000 spiking cells that are firing at 50 Hz with either random or partially synchronized spike times (each cell synchronized with 5% of the other cells). The average firing rate of the population was calculated every 10 milliseconds using 10-millisecond time bins. The figure shows that the average firing rate is more steady/stable in the case of random spiking relative to the case of partially

synchronized spiking. When cells fire or cease firing together, the average firing rate necessarily exhibits a transient increase or decrease, respectively. Note that each of the cells is firing at the same average rate (50 Hz) in both the random and synchronized cases. The fluctuations in the synchronized case will therefore not be captured well by firing-rate neural network models that throw away information about individual spike times.

The central idea explored in this thesis is that enhanced synchrony in spiking neural networks permits improved information processing efficiency. If neurons throughout a processing pathway are synchronizing their spike times as they fire at lower rates, fewer pre-synaptic spikes could actually elicit larger and more reliable post-synaptic depolarizations throughout that pathway. This could lead to a situation where lower firing rates give rise to faster reaction times and more efficient neural processing. The impact of changes in spike synchrony is difficult to represent in firing-rate neural networks - particularly when synchrony arises as a result of the dynamical interactions of spiking cells - because the point-like activity of spikes and the resultant transient synaptic currents have been smoothly averaged. This suggests that the spiking property of neural cells may indeed be critical for understanding the relationship between repetition suppression and behavioral priming.

2.1 The Importance of Metabolic Efficiency in Neural Processing

The metabolic cost of neural information processing has often been neglected by theorists attempting to understand the basic mechanisms of brain function. However, it has undoubtedly been an important constraint in the evolution of the nervous system (e.g. Aiello & Wheeler, 1995; Allman, 1990). The brain has been estimated to account for 20% of the body's resting metabolic energy use while it accounts for only 2% of total body mass (Kety, 1957; Sokoloff, 1960; Rolfe & Brown, 1997; see Ames, 2000; Laughlin, 2001, for recent reviews). Approximately 50-80% of the mammalian brain's energy use appears to be due to neural signaling (i.e. 10-15% of total resting energy consumption), with an estimated 40-50% consumed by action potentials and 30-40% by synaptic potentials (Attwell & Laughlin, 2001; Rothman, Sibson, Hyder, Shen, Behar, & Shulman, 1999; Sibson, Dhankar, Mason, Rothman, Behar, & Shulman, 1998). The large cost of neural processing suggests that efficiency, not just efficacy of signaling, will be critical for an organism's survival. Accordingly, a few theorists have begun to consider the implications of metabolic efficiency for neural coding (e.g. Balasubramanian, Kimber, & Berry, 2001; Levy & Baxter, 1996; Sarpeshkar, 1998; Schreiber, Machens, Herz, & Laughlin, 2002). These researchers have argued that sparse, distributed neural representations strike an effective balance between the rate of information transfer and metabolic cost. The sparseness of the theoretically optimal energy-efficient code depends fundamentally on the ratio of signaling costs to the fixed cost of maintaining cells (see Laughlin, 2001; Levy & Baxter, 1996, for discussion).

2.2 Mechanisms Underlying Spike Synchrony in Artificial Neural Networks

If more synchronous firing can permit a reduction in the number of spikes necessary to process a stimulus effectively, thus reducing the incurred metabolic cost, what are the mechanisms that might underlie enhanced synchronization? Spike synchronization in large networks of neurons has been investigated by a wide range of theorists using a variety of models and formalisms (e.g. Abbott & van Vreeswijk, 1993; Bressloff & Coombes, 1999; Chow, 1998; Ermentrout & Kopell, 1998; Gerstner, van Hemmen, & Cowan, 1996; Golomb & Hansel, 2000; Hansel, Mato, & Meunier, 1995; Hopfield & Herz, 1995; Hopfield & Brody, 2001; Terman & Wang, 1995; Terman, Kopell, & Bose, 1998; Tsodyks, Uziel, & Markram, 2000; van Vreeswijk, Abbott, & Ermentrout, 1994). One method that has proven particularly useful in understanding the mechanisms of synchronization is to reduce networks of biologically realistic spiking cells (i.e. based on Hodgkin-Huxley style equations to represent current-voltage kinetics in cell membrane: Hodgkin & Huxley, 1952) to a corresponding "phase" model, where each repetitively firing neuron is characterized as an oscillator with a certain period, phase, and form of interaction with the other neural oscillators (e.g. Ermentrout & Kopell, 1984, 1991; Hansel, Mato, & Meunier, 1993; Kopell, 1988; Kuramoto, 1975, 1984). For example, consider two reciprocally coupled neurons with dynamics governed by the phase variables ϕ_1 and ϕ_2 :

$$d\phi_1/dt = \omega_1 + H_1(\phi_2 - \phi_1)$$

$$d\phi_2/dt = \omega_2 + H_2(\phi_1 - \phi_2)$$

where ϕ represents the spike time or phase of a neuron in relation to its unperturbed limit cycle when uncoupled, ranging from 0 to T (its period length); ω represents the firing frequency of each neuron when uncoupled ($1/T$); and $H(\phi)$ is the interaction function that depends on the response function of the neurons (the phase-response curve, or PRC, that is often determined empirically by measuring how very small inputs advance or lag a neuron's phase), the time-varying synaptic current (such as AMPA or GABA), and the phase difference between the neurons. This reduction is reasonable as long as all neurons fire periodically when they are uncoupled, their firing rates fall within a narrow range, and their coupling is relatively weak (meaning that interactions are not so strong that the shape of the response function is changed) (Ermentrout & Kopell, 1991; Kopell, 1988; Kuramoto, 1984). If we let ψ represent the phase difference between the two neurons ($\psi = \phi_2 - \phi_1$) and assume that both neurons have identical firing rates ($\omega_1 = \omega_2$) and H functions:


$$\begin{aligned} d\psi/dt &= d\phi_2/dt - d\phi_1/dt \\ &= H(-\psi) - H(\psi) \\ &= -2 H_{odd}(\psi) \end{aligned}$$

where $H_{odd}(\psi)$ is the odd part of $H(\psi)$ ⁸. There is a phase-locked solution for this system if $d\psi/dt = 0$ for some ψ (e.g. when the neurons are phase-locked at synchrony, $\psi = 0$). Since odd periodic functions all have 0 and $T/2$ as solutions, synchrony and anti-synchrony are solutions, and depending on the particular response function and synaptic current, there can be additional solutions. The stability of these phase-locked solutions is

⁸ $H(\psi)$ can be expanded into even and odd parts where $H(\psi) = H_{even}(\psi) + H_{odd}(\psi)$. By definition of even and odd functions, $H_{even}(\psi) = H_{even}(-\psi)$ and $H_{odd}(\psi) = -H_{odd}(-\psi)$. Therefore, the even parts of $d\psi/dt = H(-\psi) - H(\psi)$ cancel, leaving $d\psi/dt = -2H_{odd}(\psi)$.

determined by the slope of $H_{odd}(\psi)$ at the solution; if $H_{odd}'(\psi) > 0$, the solution is stable, else it is unstable. To see this, consider what happens if $H_{odd}'(\psi)$ is positive at a solution and a small value is added to or subtracted from ψ . Adding a small positive value to ψ leads to a negative $d\psi/dt$, and ψ decreases slowly back to the solution where it stops (since $d\psi/dt = 0$ there). Subtracting a small value from ψ leads to a positive $d\psi/dt$, and ψ increases back to the solution. On the other hand, if $H_{odd}'(\psi) < 0$, ψ moves further away from the solution when perturbed. Figure 2.2 shows an example of H_{odd} for two neurons coupled with excitation versus inhibition (model and parameters taken from type I cells in Gutkin et al., 2001, without I_{AHP}). The synchronous solution is unstable with excitation (slope negative at $\psi=0$), yet stable with inhibition (slope positive at $\psi=0$).

A.



B.




Figure 2.2: (A) Synchrony is unstable in Type I model cells with excitatory coupling, and (B) synchrony is stable, as is anti-synchrony in this case, for Type I cells with inhibitory coupling. The stability of particular phase-locked solutions is determined by the slope of H_{odd} where it crosses $\omega_1 - \omega_2$: Positive slope indicates stability and negative slope indicates instability (see text for explanation).

2.2.1 Synchrony in Networks with Excitatory or Inhibitory Synapses

In order to characterize the impact of various factors on the stability of synchrony and other phase-locked solutions, researchers have found it useful to distinguish between neurons with different types of phase response curves (PRCs) (Ermentrout, 1996; Hansel et al., 1995). A type I PRC is always positive, meaning that these neurons will always fire earlier (are phase-advanced) in response to a depolarizing input. A type II PRC is negative following a spike and crosses over to positive values later in the period, causing neurons to fire either later or earlier in response to a depolarizing input, depending on the timing of the input. Leaky integrate-and-fire (IAF) models and some more detailed conductance-based models of cortical neurons appear to be of type I (e.g. Hansel et al., 1995; Reyes & Fetz, 1993; van Vreeswijk et al., 1994; Wang & Buzsaki, 1996). Type I neurons when coupled with excitatory synaptic interactions generally do not synchronize unless the rise and decay times of the synaptic current are very fast relative to the period of firing and the membrane time constant (Chow, 1998; Ermentrout, 1996; Hansel et al., 1995; van Vreeswijk et al., 1994; see **Appendix D** for a discussion of the membrane time constant). If excitatory synaptic interactions are slow relative to the firing rates and the the membrane time constants, cells generally lock in anti-phase or fire asynchronously with no stable states. This suggests that very low firing rates and long membrane time constants will be required to observe synchrony in type I excitatory networks. On the other hand, inhibitory coupling in networks of type I neurons is generally synchronizing unless the rise and decay times of the synaptic currents are very fast or very slow relative to the firing rates and time constants of the cells (Chow, 1998; Terman, Kopell, & Bose, 1998; van Vreeswijk et al., 1994; White, Chow, Ritt, Soto-Trevino, & Kopell, 1998). In

other words, networks of type I neurons coupled with excitation require sufficiently fast synaptic interactions, low firing rates, and slow membrane time constants in order to synchronize, whereas inhibitory coupling synchronizes with slower synaptic interactions throughout a wider range of firing rates and membrane time constants. In contrast, excitation in networks of type II neurons can be synchronizing throughout a wider range of firing rates and time constants as long as the excitation is not too slow (Hansel et al., 1995). The original Hodgkin-Huxley model of the squid giant axon is an example of a neuron with a type II PRC⁹.

More recently, researchers have been exploring what happens when certain assumptions of the phase-reduction method are violated or when additional cellular mechanisms are included. For example, when each cell in a coupled network receives different amounts of driving current, some degree of heterogeneity will result in the firing rates of the individual neural oscillators. The impact of firing-rate heterogeneity necessarily rules out exact synchrony as a solution, since $d\psi/dt$ (from the example above for 2 neurons) becomes:

$$\begin{aligned} d\psi/dt &= d\phi_2/dt - d\phi_1/dt \\ &= \omega_2 - \omega_1 + H_2(-\psi) - H_1(\psi) \end{aligned}$$

Phase-locked solutions ($d\psi/dt=0$) will therefore only exist at some ψ where $H_2(\psi) - H_1(\psi) = \omega_1 - \omega_2$ (where $0 \leq \psi \leq T$ of the fastest neuron). Because $H_1(0) = H_2(0) = 0$, synchrony can only be a solution if $\omega_1 = \omega_2$. Figure 2.3 depicts this situation graphically.

⁹ Published results are scant on the effect of inhibitory coupling in networks of type II neurons. However, given that reversing the sign of the coupling strength will tend to reverse the sign and slope of H_{odd} , one might speculate that inhibition should be more de-synchronizing with type II neurons.




Figure 2.3: Heterogeneity of firing rate (indicated by non-zero $\omega_1 - \omega_2$) degrades the stability of the synchronous solution in two Type I neurons with inhibitory coupling. Since H_{odd} is always 0 at the beginning and end of the period of firing, it cannot be equal to $\omega_1 - \omega_2$. It should be noted that this depiction is only for illustrative purposes and is not precise because the calculation of H_{odd} here actually assumes equal firing rates.

However, if the heterogeneity in firing rates is small, stable phase-locked solutions near synchrony can still exist as long as the original homogeneous solutions are stable (Chow, 1998; Neltner, Hansel, Mato, & Meunier, 2000; White et al., 1998). Since the synaptic coupling strengths between the neurons can scale the height of H , these values partially determine whether or not a phase locked solution of $H_2(\psi) - H_1(\psi) = \omega_1 - \omega_2$ exists (i.e. scaling the synaptic strengths by a fixed value may alter the existence, but not the stability of phase-locked states; Chow, 1998).

2.2.2 Synchrony in Networks with Both Excitation and Inhibition

In some select cases, researchers have been able to analyze the impact on synchronization of allowing excitatory neurons to interact with inhibitory interneurons (e.g. Borgers & Kopell, in press; Brunel, 2000; Ermentrout & Kopell, 1998; Ermentrout, Pascal, & Gutkin, 2001; Karbowski & Kopell, 2000; Terman & Wang, 1995). For example, Ermentrout, Pascal, and Gutkin (2001) have recently shown using the reduced-phase method that when an interacting pair of neurons, one excitatory and one inhibitory, are coupled to other excitatory-inhibitory pairs through the excitatory neurons (shown in Figure 2.4), synchronous or near-synchronous solutions are stable throughout a wider range of firing rates compared to networks involving only excitation. In other words, adding local inhibitory interactions to networks dominated by excitation can enhance synchronization. Ermentrout et al. (2001) have shown that this phenomenon is a result of the way inhibition alters the PRC's - and thus $H_{odd}(\psi)$ - of the excitatory neurons. A similar example provided in **Appendix B** shows that synchronization in networks of coupled excitatory-inhibitory pairs (type I neurons) depends on firing rate in the same way that synchronization in all-excitatory networks does: The synchronous or near-synchronous solution is stable at lower firing rates and breaks down gradually as firing rate increases. The inclusion of inhibition leads synchrony or near-synchrony to break down at higher and higher rates, extending the range of firing rates over which interactions promote synchrony.




Figure 2.4: Network architecture assumed in analyzing the interaction of coupled excitatory (e) and inhibitory (i) neuron pairs (e.g. Ermentrout et al., 2001).

Using different analytic methods, Ermentrout and Kopell (1998; also Karbowski & Kopell, 2000; Kopell, Ermentrout, Whittington, & Traub, 2000) have shown that excitatory-inhibitory pairs can be good at promoting synchrony even when there are long conduction delays between pairs. In networks with all-to-all or sparse, random connections between two populations of neurons, one excitatory and one inhibitory, Borgers and Kopell (in press) have recently shown that excitatory-inhibitory interactions can lead to near-synchrony even when individual cells receive substantially different numbers of inputs (i.e. input heterogeneity). In less formal, numerical simulations, they also found that further incorporating within-population connections (e.g. excitatory-excitatory and inhibitory-inhibitory) did not alter their results substantively. However, this is not to say that synchrony is the only possible outcome of excitatory and inhibitory interactions. Brunel (2000) has shown analytically that networks of simple excitatory and inhibitory integrate-and-fire (IAF) neurons with sparse, random connectivity and

randomly firing inputs can exhibit a rich range of dynamical states, including regular (periodic) and irregular (a-periodic) asynchronous spiking, as well as regular and irregular synchronous spiking. The synchronous regular spiking regime generally occurs when excitation is stronger than inhibition, whereas the synchronous irregular regime occurs when inhibition is stronger than excitation and when the random (Poisson) inputs fire at either high or low rates (asynchronous irregular occurring for intermediate input rates). Both of these regimes largely give way to asynchronous firing (either regular or irregular) as the transmission delays between cells become too heterogeneous, although both regimes are stable over a range of uniform delays.

A number of researchers, both theorists and experimentalists, have focused more specifically on the role that special properties of interneurons, such as "gap junctions", might play in synchronizing populations of excitatory and inhibitory cells (e.g. Beierlein, Gibson, & Connors, 2000; Benardo, 1997; Chow & Kopell, 2000; Galarreta & Hestrin, 1999, 2001; Gibson, Beierlein, & Connors, 1999; Lewis & Rinzel, 2000, in press). Gap junctions are fast electrical synapses that have been observed between interneurons, particularly between those with similar physiological characteristics (e.g. fast-spiking vs. low-threshold-spiking interneurons) (e.g. Gibson et al., 1999; Amitai, Gibson, Beierlein, Patrick, Ho, Connors, & Golomb, 2002). The direct electrical coupling between interneurons enforces a certain degree of synchronous spiking with strong coupling; at weaker coupling, gap junctions can still be synchronizing, provided that the spike amplitude and width are not too large and the firing rates of the interneurons are not too high (Chow & Kopell, 2000; Lewis & Rinzel, in press). Networks of interneurons might then help to synchronize populations of excitatory cortical cells by providing a common

synchronized input, regardless of the dynamical effects of slower chemical neurotransmission (Beierlein et al., 2000; Galarreta & Hestrin, 2001). Neuromodulation due to slow effects of glutamate and/or acetylcholine has been shown to suppress activity in fast-spiking (FS) interneurons while enhancing activity in low-threshold-spiking (LTS) interneurons. Since LTS cells are coupled nearly exclusively with gap junctions, these and other neuromodulators may help to enhance synchronization in neocortical cells (Beierlein et al., 2000).

2.2.3 The Impact of Short-term Plasticity on Synchronization

For the purposes of the current work described here, it is important to understand how short-term plasticity processes such as firing-rate adaptation and synaptic depression alter the synchronization properties within networks of spiking neurons. The impact of both of these processes has been analyzed previously using phase-reduction and other methods (Bressloff, 1999; Crook, Ermentrout, & Bower, 1998; Ermentrout, Pascal, & Gutkin, 2001; Loebel & Tsodyks, 2002; van Vreeswijk & Hansel, 2001). Crook et al. (1997) and Ermentrout et al. (2001) have shown that spike frequency adaptation - both due to the voltage-dependent M-current (I_M) and the calcium-activated potassium current ($I_{K(Ca)}$) - helps to synchronize networks of excitatory cells (type I). It does this partially by reducing firing rates and partially by altering the shape of the PRC so that synchrony or near-synchrony becomes a stable solution. Indeed, the alteration of the PRC can be strong enough in the case of the M-current that an early negative component can emerge in the PRC, changing the otherwise type I neuron into a type II neuron (recall from **Section 2.2.1** above that type II neurons can synchronize more readily when coupled with

excitation) (Ermentrout et al., 2001). This implies that pyramidal cells with strong M-currents may be closer to type II than to type I, whereas inhibitory interneurons appear to be of type I. The synchronizing effect of adaptation is perhaps not so surprising in the context of the discussion in **Section 2.2.2** above. An excitatory neuron with firing-rate adaptation is mechanistically analogous to a single excitatory-inhibitory neuron pair, where the time-course of the inhibition is dictated by the rise and decay times of the adaptation current.

Similarly to firing-rate adaptation, synaptic depression can synchronize networks of excitatory cells, as well as networks with both excitatory and inhibitory cells, by reducing firing rates (e.g. Bressloff, 1999). An example of this is shown in **Appendix C** at varying levels of synaptic depression for two excitatory-inhibitory neuron pairs coupled through the excitatory neurons. In contrast to adaptation, there is very little evidence from this example that synaptic depression can produce additional dynamical effects due to an alteration of the PRC shape (although see Bressloff, 1999). It appears then that synaptic depression may only play a rate-reducing role in synchronizing cells. Indeed, the simulations below will show that its impact on firing rate reductions is considerably stronger than that due to firing-rate adaptation when each is parameterized to match values recorded in neurophysiological experiments.

Both firing-rate adaptation and synaptic depression may also play an additional role in reducing firing-rate heterogeneity. As discussed above, exact synchrony is not possible when different cells receive very different degrees of depolarization or when other cellular properties are not uniform (e.g. Chow, 1998; Neltner et al., 2000; White et al., 1998). As adaptation and synaptic depression build up across each spike and reduce

firing rates, they do so in a way that is proportional to the initial firing rate (or alternatively, proportional to the driving current; see **Appendix A** for mathematical details and discussion). This means that the distribution of firing rates across the population will be compressed as the firing rates decrease. It is possible that the reduction in firing rate variability will contribute to greater synchronization, above and beyond the effect due to overall firing rate reduction.

Recently, Tsodyks, Uziel, and Markram (2000; Loebel & Tsodyks, 2002) and van Vreeswijk and Hansel (2001) have analyzed the impact of synaptic depression and firing-rate adaptation, respectively, on synchronized bursts within networks of excitatory and inhibitory neurons. Synchronized bursts reflect activity that is asynchronous on a spike-to-spike basis, yet where firing-rate activity changes simultaneously across a population of cells, alternating between bursts of high firing rate and periods of quiescence. When recurrent excitation is strong, firing rates can start out low and build up rapidly due to positive feedback, yielding a burst of high-rate spiking and spike-to-spike asynchrony. Firing-rate adaptation and synaptic depression can then counteract the high firing rates through strong negative feedback, terminating the burst. The onset of the next burst will then depend on the recovery rates of adaptation and synaptic depression. If these mechanisms are included within a network representing two perceptual alternatives that compete through lateral inhibition (each alternative is a subgroup of neurons coupled with both excitation and inhibition), it is possible to account for a wide range of psychophysical phenomena associated with perceptual rivalry (e.g. binocular rivalry, Laing & Chow, 2002).

2.3 Empirical Evidence for Neural Synchrony

Much of the work on spike synchrony using *in vivo* recording techniques in the neocortex has focused on the role that synchrony may play in the temporal binding of visual features during visual segmentation and figure-ground segregation (e.g. Engel, Konig, Gray, & Singer, 1990; Engel, Konig, Kreiter, & Singer, 1991; Gray, Konig, Engel, & Singer, 1989). When two oriented bars are present in a visual display in different colors, the visual system is somehow capable of attributing the correct color to the correct shape, even though it is equally capable of representing either shape in either color. Visual perception researchers have referred to this as the "binding problem". One potential (although not the only) solution to this problem is that neurons that participate in representing different visual features of the same object (e.g. color, shape, etc.) fire synchronously, allowing neurons representing another object to fire simultaneously yet at a different phase without creating any perceptual ambiguities. A number of studies employing multi-neuron recording techniques in the visual cortex of anesthetized and awake animals have garnered a range of evidence consistent with this view. Synchronization in the gamma band (35-70 Hz) has been observed between neurons located in homologous areas of visual cortex in the left and right hemispheres, as well as between neurons in different visual cortical regions within the same hemisphere (Engel et al., 1991; Gray, Engel, Konig, & Singer, 1990; Konig, Engel, & Singer, 1995; Roelfsema, Engel, Konig, & Singer, 1997). Synchronization also appears to occur between visual neurons in different cortical columns with similar tuning properties for orientation and direction of motion (Gray, Konig, Engel, & Singer, 1989; Engel, Konig, Gray, & Singer, 1990). In one well-known study, Gray et al. (1989) recorded visual

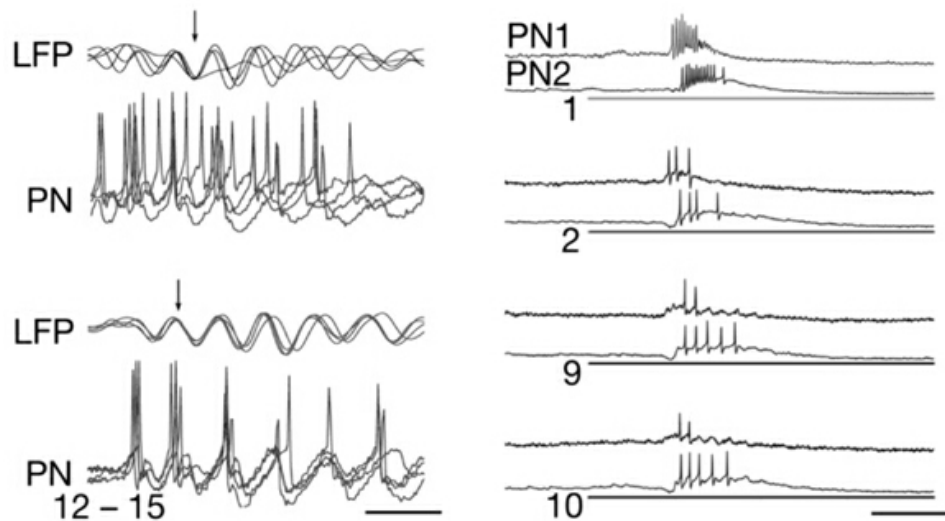
cortical neurons with non-overlapping receptive fields and moved either one long bar simultaneously through both receptive fields or two smaller bars, each swept in the same direction through the receptive fields or in opposite directions. The neurons with different receptive fields synchronized strongly in the one-bar case, whereas they did not synchronize with two bars moving in different directions, and they synchronized only weakly with two bars moving in the same direction. While this type of study demonstrates sensitivity of synchronization to stimulus properties and the global stimulus configuration, studies have yet to establish simultaneously within-object neural synchronization and between-object de-synchronization for visual displays with multiple objects. Similar evidence of neural synchronization in auditory perceptual processing has been observed by deCharms and Merzenich (1996). They found that neurons in primary auditory cortex synchronize during the presentation of a sustained tone. While firing rates showed transient increases at the onset of the tone, only neural synchronization persisted throughout the tone presentation.

More recently, a number of studies have documented enhanced spike synchronization during states of directed attention. A study by Steinmetz, Roy, Fitzgerald, Hsiao, Johnson, and Niebur (2000) involved training monkeys to perform two different psychophysical tasks, one visual and one tactile. In the visual task, the monkey had to detect the dimming of a target spot in the visual display while ignoring tactile stimuli that were irrelevant to task performance. The tactile task involved the same tactile and visual stimuli as in the visual task, but in this case, correct performance required detecting if the raised pattern of bumps presented to the fingertips matched the pattern in the visual display. For a given block of trials, the animals were cued as to

which of the two tasks to perform, and recordings were made from pairs of neurons in somatosensory cortex (S2) during the two different task conditions. While firing rates increased in somatosensory cortex during the tactile task, greater spike synchrony was observed during the tactile task than during the visual task, and the degree of synchrony was greater than that expected due to the increased rates. A different study by Fries, Reynolds, Rorie, and Desimone (2001) examined changes in synchrony in visual cortex (V4) when monkeys were trained to attend to one of two visual stimuli presented simultaneously at positions of equal eccentricity. One of the stimuli was presented inside the V4 receptive field of the neuron being recorded, while the nearby local field potential (LFP) was recorded on a separate electrode. The neuron being recorded became more synchronized in the gamma band (35-70 Hz) with the LFP when the monkey was instructed to attend to the object inside the neuron's receptive field than when attending to the other object. Critically, the visual contrast of the stimuli was manipulated so as to minimize the differences in firing rate across the two conditions, preventing the changes in synchrony from being attributed simply to changes in rate. For these results and for those discussed above, it is interesting to note that synchrony is enhanced in the gamma band in the hippocampus and neocortex by cholinergic (muscarinic) agonists and disrupted by antagonists (Rodriguez, Kallenbach, Singer, & Munk, 2001; Fellous & Sejnowski, 2000; Tiesinga, Fellous, Jose, & Sejnowski, 2001). Taken together with other evidence for the role of acetylcholine in attention and memory processing (e.g., Coull, 1998; Hasselmo, 1995; Marrocco & Davidson, 1998; Robbins, 1997), these observations may suggest an important role for neuromodulators in the enhanced synchrony observed under these task conditions.

For the current work, perhaps the most central issue is to understand how recent experience and practice with a particular stimulus alters neural synchronization. While few studies to date have evaluated how short-term stimulus repetition influences spike synchrony in neocortex, it has been examined in the insect olfactory system. Stopfer and Laurent (1999) repeatedly presented 1-second air puffs of particular odors at a rate of one per 10 seconds (0.1 Hz) to the antennae of locusts. They simultaneously recorded activity of individual neurons in the locust antennal lobe and nearby LFPs. In locusts, as in monkeys (e.g. Miller et al., 1991, 1993), the initial presentation of an odor elicited strong spiking responses, whereas each subsequent presentation (tested up to 10 repetitions) yielded fewer and fewer spikes. However, as firing rates decreased, spikes in individual neurons became progressively more synchronized with the LFP (shown in Figure 2.5).

A.



B.

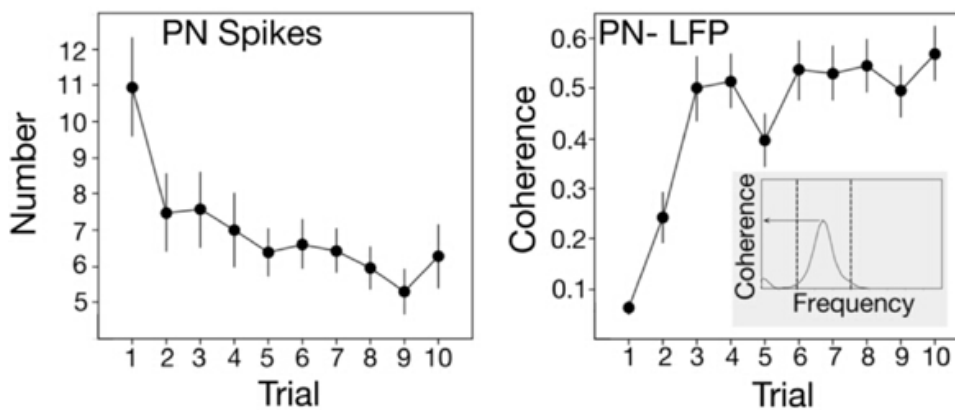


Figure 2.5: Short-term repetition of odor puffs to the antennae of a locust leads to decreased spike rates in antennal lobe neurons while at the same time leading to enhanced synchronization, as reported by Stopfer and Laurent (1999): (A) local field potentials (LFPs) and voltage traces of projection neurons (PNs) initially exhibit relatively random timing relationships, yet after 12-15 odor puffs, spikes in the PNs are locked to the peaks of the LFP; (B) Firing rates in the PNs decrease across repeated odor puffs, while simultaneously showing greater coherence with the LFP - indicating better spike synchrony across the population. From "Short-term memory in olfactory network dynamics," by M. Stopfer and G. Laurent, 1999, *Nature*, 402, p. 664. Copyright 1999 by Macmillan Magazines Ltd. Permission pending.

Decreased neural firing rates and increased spike synchronization varied together across a range of stimulus durations (0.25-2 seconds) and interstimulus intervals (2.5-20 seconds), and responses for a particular odor returned to initial levels following a delay of 12-15 minutes. They also established that the short-term changes in firing rate and synchrony were odor-specific in the sense that they did not carry over to the presentation of a chemically distinct odor, whereas some degree of generalization did occur across chemically related odors. When they submitted the spike trains recorded in several antennal lobe neurons during the presentation of two different odors to a classification algorithm (discussed in MacLeod, Backer, & Laurent, 1998), odor classification was actually better using the responses to later presentations than to the initial presentation even though the firing rates were odor-specific and were initially more intense. Previous work in the same lab had already demonstrated that artificially de-synchronizing antennal lobe neurons using a block of inhibitory neurotransmitter reduced the specificity of downstream neural responses and impaired odor discrimination in behaving insects (MacLeod et al., 1998; Stopfer, Bhagavan, Smith, & Laurent, 1997). While it is currently unknown if short-term stimulus repetition has a similar effect on the synchronization of neural responses in mammalian cortex, these results indicate that short-term repetition, decreased firing rates, and enhanced synchronization do occur in some neural systems - and in those systems, synchronization appears to be important for perceptual processing.

One recent study has examined the relationship between longer-term stimulus repetition and neural synchronization in rat sensory neocortex. Vazquez, Pandya, Engineer, Moucha, Rathbun, and Kilgard (2001) repeatedly paired the presentation of auditory tones with electrical stimulation in the cholinergic basal forebrain and recorded

the spiking activity of multiple neurons simultaneously in rat primary auditory cortex (A1). The stimulation of the basal forebrain enhanced release of cortical acetylcholine - and hence, likely increased the degree of long-lasting activity-dependent plasticity (LTP/LTD) in the auditory cortex during tone presentation (see also Brocher, Artola, & Singer, 1992; Huerta & Lisman, 1993; Kilgard & Merzenich, 1998). Consistent with this interpretation, one month of pairing tones with basal forebrain stimulation led to cortical map and receptive field expansion for practiced tones and a large enhancement of spike synchronization between individual neurons. These results suggest that longer-term changes in neural synchronization may occur with practice in mammalian neocortex, and hence, may be relevant to understanding longer-lasting neural and behavioral priming effects in humans (e.g. Cave, 1997; van Turennout et al., 2000; Wagner et al., 2000).

2.4 Basic Hypothesis and Overview of Simulations: Short-term Plasticity Enhances Neural Synchrony and Processing Efficiency

As discussed in the sections above, networks of excitatory and inhibitory neurons tend to synchronize better at lower rather than higher firing rates. Short-term plasticity processes of firing-rate adaptation and synaptic depression should help neurons to synchronize because they reduce firing rates, and in the case of adaptation, synchrony should be further supported by the alteration of the response functions (PRCs) of the excitatory neurons. Adaptation and synaptic depression may also help to promote synchrony by reducing the variability in firing rates present across a network of interacting cells. As neurons synchronize, processing is metabolically more efficient because fewer spikes are

required to evoke responses in subsequent cortical regions. If the changes in synchrony are large enough relative to decreases in firing rate, processing can actually be facilitated, giving rise to behavioral priming effects. These ideas are similar in some respects to a recent proposal by Salinas and Sejnowski (2000, 2001) that synchrony can help to gate the "flow" of information processing in the cortex without appreciably modifying the "content" of processing. In other words, synchrony can help to scale the efficacy of certain neural signals over others, but it doesn't necessarily alter the nature of the information being sent. A cell helping to represent certain stimuli by virtue of its weighted connectivity with other cells will tend to represent those same stimuli whether synchronized with similarly tuned cells or not. The current proposal differs from the Salinas and Sejnowski proposal mainly in its emphasis on the way that synchronization improves with dynamic, practice-dependent decreases in rate (they focused on cases in which synchrony changes without large changes in rate), as well as in its focus on providing a mechanistic basis for how synchrony can occur in populations of cells that interact. Enhanced synchronization will tend to increase metabolic efficiency in either case (performance relative to energy spent), but the mechanisms involved in the two cases are likely to be somewhat different (discussed in detail in *Simulation 8*). The role of synchrony in promoting efficiency is also somewhat more transparent in the case of dynamic firing rate decreases because less and less absolute energy is spent while performance is nevertheless improving.

In networks of coupled excitatory and inhibitory spiking neurons, I will demonstrate through a series of simulations that the short-term plasticity processes of firing-rate adaptation and synaptic depression can account simultaneously for the basic

neural and behavioral effects associated with short-term repetition priming. The following provides a brief overview of the points addressed by each simulation in the thesis:

2.4.1 Short-term Plasticity, Synchrony, Priming, and Processing Efficiency (Simulations 1-5)¹⁰

- **Simulation 1** - This simulation examines the effects of short-term stimulus repetition and interstimulus interval on the mean and variance of firing rates in a large pool of excitatory and inhibitory integrate-and-fire (IAF) neurons that are parameterized to match the physiological characteristics of cortical cells. Results are compared to repetition suppression data acquired through extracellular single-unit recordings of inferotemporal cortex neurons (Miller et al., 1991, 1993) and hemodynamic responses measured in neuroimaging studies (Grill-Spector & Malach, 2001; Jiang et al., 2000).
- **Simulation 2** - The effects of short-term stimulus repetition on spike synchronization are evaluated in the context of the changes in firing rate demonstrated in *Simulation 1*. The robustness of repetition-related changes in synchronization to input heterogeneity, longer interstimulus intervals, and synaptic delays is examined.
- **Simulation 3** - The contributions of inhibitory interneurons, firing-rate adaptation, and synaptic depression to the changes in spike synchronization observed in *Simulation 2* are evaluated by selectively blocking each. The potential of each mechanism to contribute a dynamical effect independent of changes in rate by

¹⁰ Some of the results presented in *Simulations 1-4* have been presented in preliminary form in Gotts and Chow (2001).

altering the PRCs of the excitatory neurons was assessed by carefully matching the means and variances of firing rates in the coupled pool before and after the blocking of each mechanism.

- **Simulation 4** - This simulation explores the impact of firing rate and synchronization of a population of spiking neurons that represent processing in more posterior cortical regions on the firing rate of a single output neuron that is taken to be representative of neurons in more frontal/motor regions. The latency required to reach a certain threshold number of spikes in this output neuron is used as a simple proxy of reaction time, and a map is created to translate firing rates and levels of synchrony measured in *Simulation 2* into reaction times. The relation of this vastly simplified response system to data on the neurophysiological bases of reaction times (Hanes & Schall, 1996) is discussed.
- **Simulation 5** - Given the claim that decreased rates and increased spike synchrony enhance metabolic efficiency, two simple measures of efficiency are calculated based on the reaction times in *Simulation 4* and the number of spikes emitted in generating a response. This permits an assessment of how stimulus repetition alters efficiency.

2.4.2 Interaction of Short-term Plasticity and Neuromodulation (*Simulations 6-8*)¹¹

- **Simulation 6** - The short-term plasticity process of synaptic depression is incorporated into a firing-rate connectionist model of auditory word recognition (phonology through semantics), along with several basic actions of

¹¹ The results presented in *Simulations 6* and *7* have been published in Gotts and Plaut (2002).

neuromodulation (acetylcholine and norepinephrine). This simulation explores the ability of two different types of neurological damage (removing intracortical connections versus reducing levels of neuromodulation) to account for the basic contrasting pattern of data associated with "access/refractory" and "degraded-store" semantic impairments in neurologically damaged patient populations (Warrington & Shallice, 1979; Shallice, 1988). In particular, the simulation explores the ability of neuromodulatory deficits that result in large synaptic depression effects to explain the habituation-like performance of "access/refractory" patients performing semantic tasks with repeated stimuli.

- ***Simulation 7*** - The ability of different combinations of the two damage types explored in *Simulation 6* (damage to intracortical connections versus neuromodulation) to account for common "mixed" patterns of access/refractory and degraded-store performance characteristics is evaluated (e.g. Rapp & Caramazza, 1993). The degree of constraint present in the model (the entire range of behaviors under all possible damage combinations) is also characterized and several predictions are elaborated.
- ***Simulation 8*** - This simulation attempts to show that the priming effects observed in the simulations involving spiking neurons (*Simulations 1-5*) and the habituation-like effects observed in the firing-rate simulations (*Simulations 6-7*) can both be produced by the same model. It explores the idea that the repetition-related synchronization effects and behavioral facilitation are observed at moderate to high levels of neuromodulators (acetylcholine or norepinephrine) where activity is sufficiently high for firing-rate adaptation and synaptic

depression effects to apply, reducing firing-rate variability and contributing dynamical effects through the alteration of the excitatory and inhibitory neurons' PRCs. At lower levels of neuromodulation, firing may reach such low levels that firing-rate adaptation does not build up sufficiently to help neurons synchronize, and the PRCs of the excitatory and inhibitory neurons are altered in such a way that asynchronous firing becomes more likely. Under these circumstances, rates can decrease across repetitions due to synaptic depression but synchrony does not improve, leading to slowed responses and habituation.