

Decoupling through Synchrony in Neuronal Circuits with Propagation Delays

Evgueniy V. Lubenov¹ and Athanassios G. Siapas^{1,*}

¹Division of Biology, Division of Engineering and Applied Science, California Institute of Technology, Pasadena, CA 91125, USA

*Correspondence: thanos@caltech.edu

DOI 10.1016/j.neuron.2008.01.036

SUMMARY

The level of synchronization in distributed systems is often controlled by the strength of the interactions between individual elements. In brain circuits the connection strengths between neurons are modified under the influence of spike-timing-dependent plasticity (STDP) rules. Here we show that when recurrent networks with conduction delays exhibit population bursts, STDP rules exert a strong decoupling force that desynchronizes activity. Conversely, when activity in the network is random, the same rules can have a coupling and synchronizing influence. The presence of these opposing forces promotes the self-organization of spontaneously active neuronal networks to a state at the border between randomness and synchrony. The decoupling force of STDP may be engaged by the synchronous bursts occurring in the hippocampus during slow-wave sleep, leading to the selective erasure of information from hippocampal circuits as memories are established in neocortical areas.

INTRODUCTION

The study of the rules governing the efficacy of synaptic interactions is central to neuroscience, since activity-dependent modifications of synaptic strengths are believed to be critical for processing and storing information in the brain (Hebb, 1949). Recent experiments have shown that the direction and magnitude of synaptic changes are highly dependent on the relative timing of presynaptic inputs and postsynaptic spikes (Markram et al., 1997). In most systems studied to date, presynaptic inputs arriving before postsynaptic spikes lead to synaptic strengthening, while inputs arriving after postsynaptic spikes lead to weakening (Hebbian STDP) (Markram et al., 1997; Bi and Poo, 1998; Feldman, 2000). Cases where presynaptic inputs preceding postsynaptic spikes actually lead to weakening, while the opposite timing leads to strengthening, have also been observed (anti-Hebbian STDP; Bell et al., 1999). Hence, variations of only a few milliseconds in the relative timing of presynaptic and postsynaptic activity can have drastically different consequences for the functional connectivity of neuronal networks. Axonal conduction delays directly affect this timing, and recent studies have

shown that delays can greatly enhance the dynamical repertoire of recurrent networks operating under STDP rules (Izhikevich et al., 2004; Izhikevich, 2006). The functional consequences of STDP have been intensively explored through computational models (Abbott and Blum, 1996; Gerstner et al., 1996; Song et al., 2000; Shon et al., 2004; Morrison et al., 2007; Abbott and Nelson, 2000; Roberts and Bell, 2002), with recent studies showing that Hebbian STDP promotes causal links and network synchronization (Karbowski and Ermentrout, 2002; Nowotny et al., 2003; Zhigulin et al., 2003).

Here we analyze and simulate the behavior of model recurrent networks to show that in the presence of axonal conduction delays, Hebbian STDP rules generate a powerful decoupling force that is engaged whenever neurons fire in transient synchrony, i.e., within population bursts. The resulting weakening of synaptic connections favors the desynchronization of network activity. Conversely, when activity in the network is random, synaptic weights disperse under the influence of the STDP rule, thus producing a coupling and synchronizing effect. The presence of these opposing forces promotes the self-organization of recurrent networks with conduction delays into mixture states at the border between randomness and synchrony.

We further explore several specific consequences of these results. First, we demonstrate that it is possible to desynchronize an oversynchronized (“epileptic”) circuit by stimulating tight population bursts. Thus, figuratively speaking, one can fight synchrony with synchrony, an observation that may be relevant for understanding the therapeutic effects of deep brain stimulation. Second, we study the effects of the endogenous bursts generated within the hippocampus during slow-wave sleep (SWS). In particular, we record the activity of multiple CA1 neurons from freely behaving rats and, consistent with previous experimental studies (Wilson and McNaughton, 1994; Kudrimoti et al., 1999), observe decay in pairwise correlations during SWS. We argue and present experimental evidence suggesting that this decay in correlations results from bursts in CA3 engaging the decoupling force of STDP. We discuss this as a candidate mechanism for the selective erasure of hippocampal memory traces as information is consolidated into long-term cortical memory stores.

RESULTS

Decoupling Force of STDP under Synchronous Activity

The intuition behind the decoupling force of STDP is surprisingly simple and general: neurons that spike within a population burst receive inputs from the rest of the neurons participating in

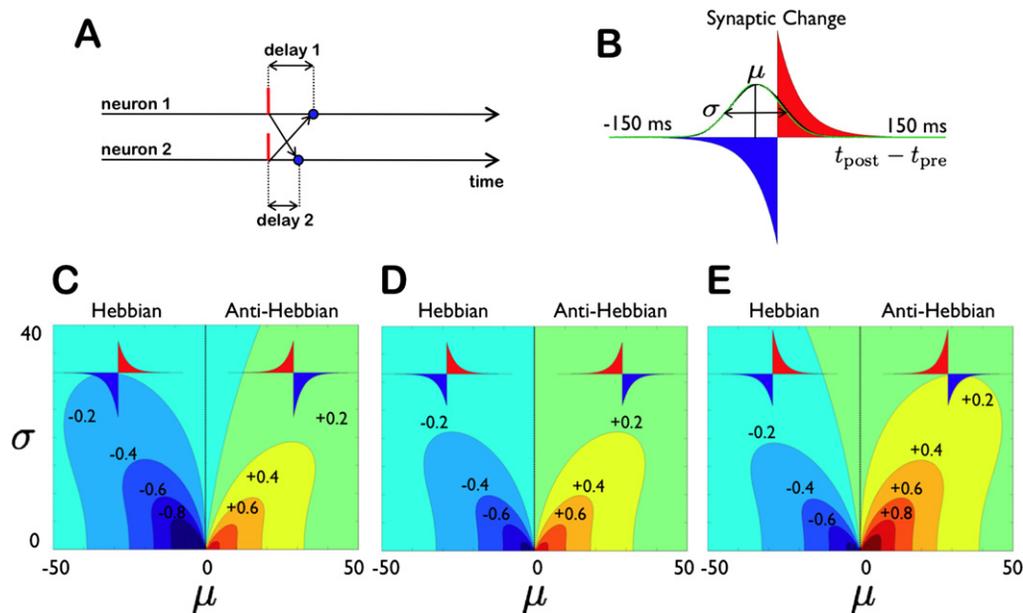


Figure 1. Decoupling Force of STDP

(A) The coincident spiking of two interconnected neurons produces reciprocal inputs that are received with some delays. The red vertical lines mark spike times and the blue dots indicate the arrival times of inputs at the postsynaptic targets.

(B) Hebbian STDP rule shows the prescribed change in synaptic strength as a function of the time difference between the postsynaptic spike and the arrival of the presynaptic spike. These time differences are both negative for the example in (A) and hence lead to weakening of both synapses. The black trace shows the time difference distribution for two neurons from the network in Figure 4, joined through a connection with a 19 ms delay and participating in synchronous bursts. The green curve is a Gaussian fit with $\mu = -18.94$ ms and $\sigma = 21.81$ ms.

(C–E) Convolution of Gaussian kernels of different widths σ with STDP rules with negative integral (C), zero integral (D), and positive integral (E). The pseudocolor panels display the relative change in synaptic strength for connections with delay $|\mu|$ between neurons participating in population bursts of width σ . Since delays are always positive, the $\mu < 0$ and $\mu > 0$ half-planes describe the synaptic changes for a given Hebbian rule and its corresponding anti-Hebbian rule, obtained by reflecting the Hebbian rule about the y axis. All three Hebbian STDP rules have a decoupling effect because the relative synaptic changes for almost all connections are negative (with the exception of short latency connections under wide bursts in [E]). In contrast, anti-Hebbian STDP rules promote the strengthening of almost all connections. This analysis applies when the interburst interval is longer than the time window of the STDP rule.

the burst with delays equal to the axonal conduction times (Figure 1A). Under Hebbian STDP (Figure 1B), this timing relationship between the arrival of presynaptic inputs and postsynaptic firing leads to the selective weakening of the connections between neurons participating in the population burst. In contrast, anti-Hebbian STDP, which is thought of as exerting a decoupling influence (Roberts and Bell, 2002), in fact has the opposite effect, promoting coupling and coincident spiking in recurrent networks with delays.

How does this intuition apply to the case when neurons fire multiple spikes within each burst? How does the decoupling force of STDP depend on the width of the burst, the conduction delay, and the shape of the STDP window function? To address these questions, for a pair of neurons participating in population bursts we approximate the distribution of time differences between postsynaptic spikes and the arrival of presynaptic inputs with a Gaussian. This distribution has mean μ equal to the negative of the axonal conduction delay of the corresponding connection and variance σ proportional to the width of the population bursts (Figure 1B). The cumulative change in synaptic strength is approximated by the integral of the product of the STDP rule and the above distribution of time differences. By varying μ and σ we can thus obtain maps of the magnitude and sign of synaptic changes predicted by a given STDP rule,

as shown in Figures 1C–1E. The maps for Hebbian STDP rules demonstrate the following three important points (Figure S3, available online).

First, population bursts lead to the consistent weakening of synapses and hence to the decoupling of networks over a wide parameter range. Second, the tighter the synchronization (smaller σ), the stronger the weakening of the synapses and hence the more potent the decoupling force. Third, connections with different latencies are differentially affected in any given population burst. In contrast, anti-Hebbian rules produce strengthening of connections ($\mu > 0$ half-planes of Figures 1C–1E). Thus, for circuits in which the level of synchronization grows with the strength of synaptic coupling, Hebbian STDP rules provide a negative feedback mechanism that opposes highly synchronized network activity, while anti-Hebbian STDP rules have the opposite effect.

A specific prediction of the above analysis is that the decoupling force of STDP must be engaged by the synchronous population bursts, known as sharp-wave (SPW) bursts, that are generated within the hippocampus during SWS. The axonal delays, μ , within the densely recurrent CA3 subfield range up to 10 ms (Section S1 in the Supplemental Data, available online), while the width of SPW bursts corresponds to σ of approximately 25 ms (Figure 6C). The maps for Hebbian STDP indicate that

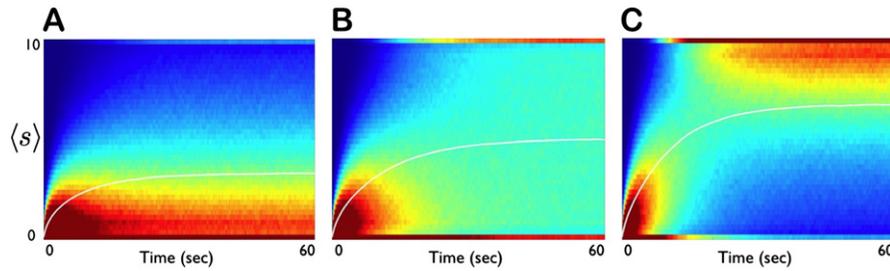


Figure 2. Diffusion-like Dispersion of Synaptic Weights under Random Network Activity

The time evolution of synaptic weights s between 50,000 uncorrelated spike train pairs was studied over a minute for (A) negatively biased rule with $A_+ = 1, A_- = -1.1$; (B) unbiased rule with $A_+ = 1, A_- = -1$; and (C) positively biased rule with $A_+ = 1.1, A_- = -1$. All rules were all-to-all with $\tau_+ = \tau_- = 20$ ms. Connection weights were initialized to 0 and allowed to vary between 0 and 10. Each column in the pseudocolor plots displays the distribution of synaptic weights at a point in time. These distributions were constructed every 10 ms and stacked to form the displayed matrices. Notice that although individual weights continued to change throughout the shown period, their distributions converged approximately halfway through to stationary distributions $P_\infty(s)$. The white curves show the evolution of the mean synaptic weights $\langle s \rangle$, which converged to values $\langle s \rangle_\infty$. Notice that each rule produces a different steady-state synaptic weight distribution with a different mean $\langle s \rangle_\infty$. These distributions are controlled by the shape of the STDP rule and are independent of the initial conditions (data not shown). Notice that $\langle s \rangle_\infty > 0$ even for the negatively biased rule in (A). Thus, an uncoupled system will evolve toward a more coupled state under random firing.

bursts with these parameter values should lead to decoupling of the recurrent connections within CA3. We present experimental evidence in support of this prediction further below.

Coupling Force of STDP under Random Activity

The decoupling force of STDP is engaged in the presence of synchronous firing. In contrast, when activity in the network is uncorrelated, synaptic weights undergo the equivalent of a random walk, biased according to the sign of the STDP rule integral and bounded by the range of allowed weights. If firing in the network were to remain random indefinitely, the weights would settle into an equilibrium distribution $P_\infty(s)$ with mean $\langle s \rangle_\infty$ (Figure 2). This distribution can be computed by studying the evolution of synaptic weights when the STDP rule is driven by uncorrelated spike trains. Even in the presence of bias in the random walk, not all weights end up at the extremes, but disperse ac-

ording to $P_\infty(s)$ (Figure 2). Therefore, for as long as firing in the network is random, the mean weight $\langle s \rangle$ will evolve toward $\langle s \rangle_\infty$. When $\langle s \rangle_\infty$ is sufficiently large to fall within the coupling range associated with correlated network activity, the diffusion-like dispersion of synaptic weights exerts a coupling and synchronizing influence (Figure 3 and Figure S7).

Opposing Forces of STDP in a Model Network

The analysis above considers isolated pairs of neurons. To illustrate how the opposing forces of Hebbian STDP interact within a circuit, we consider a randomly driven recurrent network of excitatory regular spiking (RS) Izhikevich neurons with axonal delays (Figure 4) (Izhikevich, 2003, 2006). In this simple case the network dynamics are principally controlled by the recurrent connections (see Experimental Procedures). To quantify the level of network synchronization, we introduce a simple order

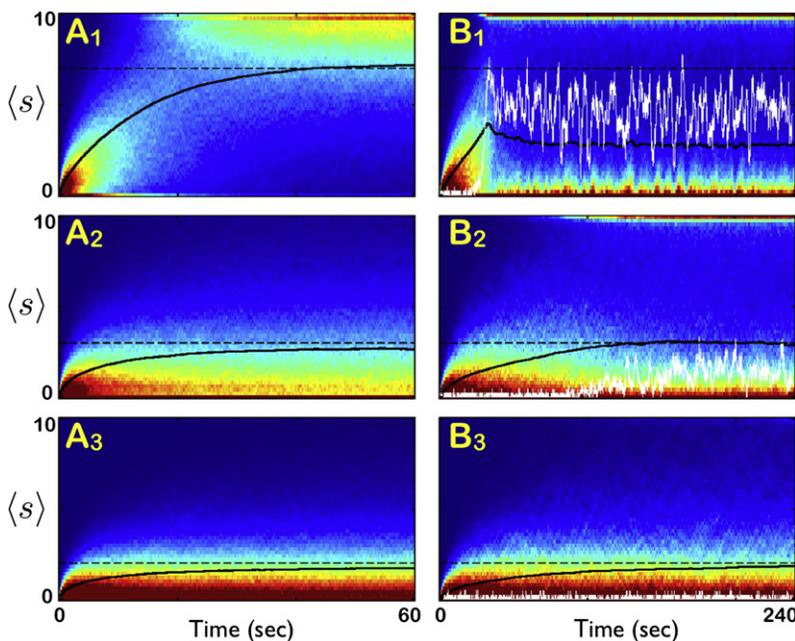


Figure 3. Coupling Force of Synaptic Weight Diffusion

The evolution of synaptic weight distribution under weight diffusion, (A1–A3), was compared with the actual evolution of weights in the network from Figure 4, (B1–B3), under three different STDP rules (Equations 1–3). The rules were chosen so that $\langle s \rangle_\infty$ was above (Equation 1), below (Equation 3), or right at (Equation 2) the critical mean weight $\langle s \rangle_c$ separating random firing from synchronous activity for this network. The evolution of the mean weight $\langle s \rangle$ is shown in black, and the order parameter ψ , scaled by a factor of 10 in white. Notice that as long as the activity in the network remained uncorrelated ($\psi \approx 0$), the evolution of the synaptic weight distribution matched closely the one predicted under weight diffusion. Notice that mixture states are reached only in (B1) and (B2), i.e., when $\langle s \rangle_\infty \geq \langle s \rangle_c$. Notice that as soon as activity in the network becomes correlated ($\psi > 0$), the evolution of the synaptic weight distribution diverges from the one predicted under weight diffusion (B1 and B2).

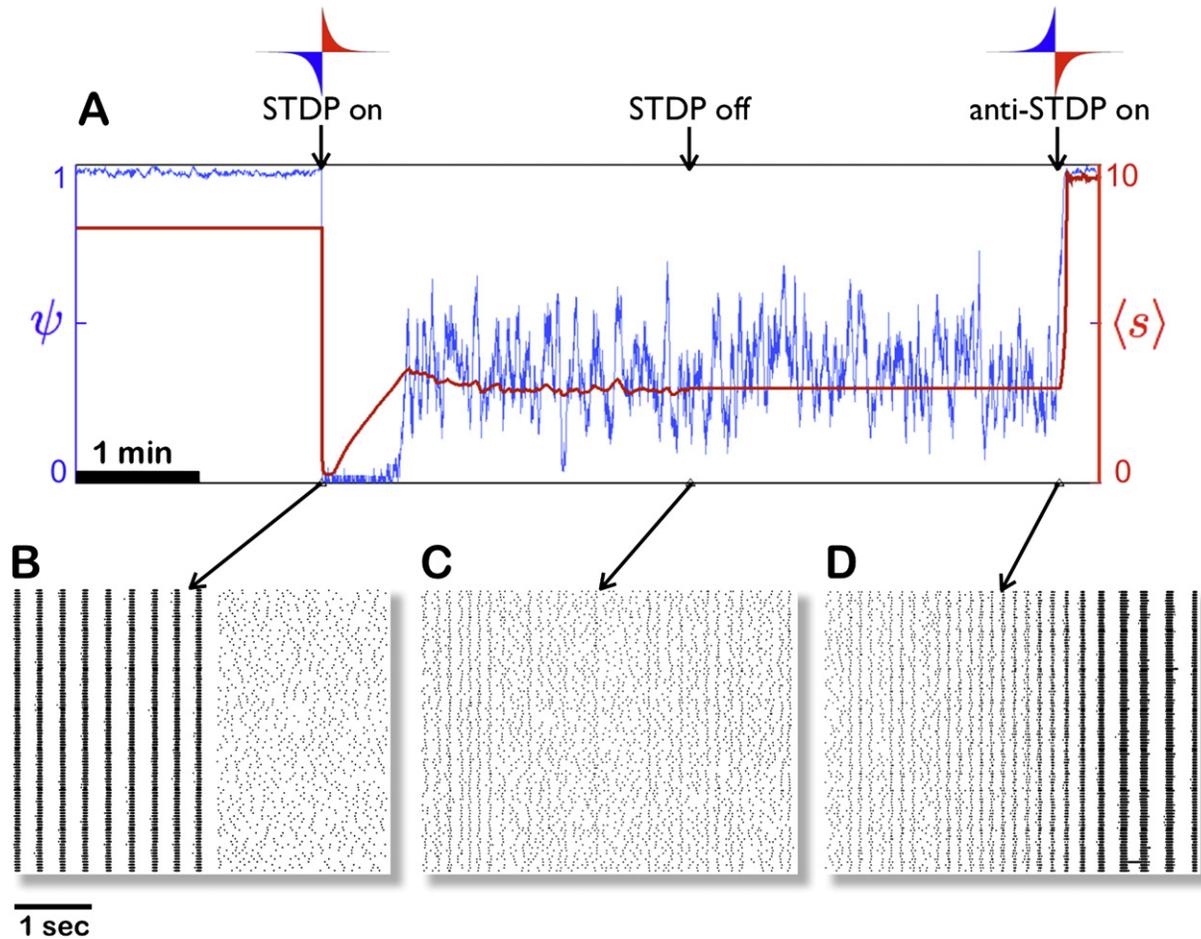


Figure 4. Illustration of the Desynchronizing and Synchronizing Effects of STDP

(A) Evolution of the mean synaptic weight $\langle s \rangle$ (red trace) and order parameter ψ (blue trace) in a recurrent network of 100 neurons with connection delays between 1 and 20 ms, driven by uncorrelated Poisson inputs. The order parameter quantifies the relative fraction of time spent in synchronized activity (left half of [B], $\psi \approx 1$) versus random activity (right half of [B], $\psi \approx 0$).

(B–D) Five-second rasters of spikes illustrating the patterns of network activity corresponding to different values of the order parameter. Before plasticity is enabled, activity in the network is completely synchronized in a 3–4 Hz oscillation. Within a few seconds of enabling Hebbian STDP, $\langle s \rangle$ plummets and activity in the network become completely asynchronous, reflecting the random pattern of external inputs (B). Over 1 min $\langle s \rangle$ stabilizes and the activity in the network becomes a mixture of short synchronous bursts interleaved with asynchronous firing (C). This pattern of activity persists even when plasticity is disabled, suggesting that plasticity is not necessary for the maintenance of this state, even though it is critical for its stability in the presence of strong external perturbations. (D) Soon after anti-Hebbian STDP is enabled, the strength of synaptic coupling saturates and activity in the network becomes completely synchronized.

parameter ψ which is 0 when network activity is random and 1 when activity is highly synchronized (Experimental Procedures). In the simulation shown in Figure 4, plasticity is disabled during the first 2 min, during which the network exhibits regular 3–4 Hz oscillations (left half of Figure 4B, $\psi \approx 1$). Within seconds after plasticity is enabled, there is a rapid decoupling, as shown by the sharp decrease of the mean synaptic weight $\langle s \rangle$ toward zero, which results in the desynchronization of network activity ($\psi \approx 0$). Subsequently, synaptic weights begin redistributing in accordance with the limiting weight distribution under random firing $P_\infty(s)$, which results in an increase of $\langle s \rangle$ toward $\langle s \rangle_\infty$ (Figure 3). This coupling process persists until the network settles into a mixture state characterized by random activity patterns interspersed with synchronized population bursts and intermittent oscillatory episodes (Figure 4C). Finally, as predicted by

our analysis, anti-Hebbian STDP promotes coincident spiking and rapidly drives the system to excessive coupling and synchrony (Figure 4D). This suggests that anti-Hebbian STDP is likely to be present only in feedforward pathways. The same qualitative behavior can be observed in more realistic networks (Figure S4) and under different STDP implementations (Experimental Procedures and Figure S5).

Self-Organization in Mixture States

The simultaneous existence of the synchronizing and desynchronizing forces described above suggests that Hebbian STDP promotes self-organization into mixture states at the border between randomness and synchrony. In the case of randomly driven networks, in which the level of synchrony is principally controlled by the strength of the recurrent connections, the

reasoning behind this prediction is very intuitive. Activity in an uncoupled network must be uncorrelated under random input, as there are no interactions that can give rise to correlated firing. At the other extreme, when the coupling between neurons is very strong, all elements behave as one, therefore producing highly synchronous firing under any input. Such dependence of network activity patterns on the level of coupling has been experimentally demonstrated in area CA3 of the hippocampus (Bains et al., 1999). So as long as STDP decouples neurons when network activity is synchronous and couples them when activity is random, the only allowed equilibrium state must be at the border between randomness and synchrony.

To examine this hypothesis directly, we generated a family of networks with varying connectivities and synaptic coupling strengths and studied their activity with and without plasticity (Figure 5). When plasticity is disabled, we find that for any given level of connectivity the mean synaptic weight determines the level of synchronization, and that there is a narrow range of coupling strengths that are associated with mixture states (Figure S6B). When Hebbian STDP is enabled, the mean synaptic weight for the majority of the networks converges to the range associated with mixed activity patterns. The only exceptions are very sparse networks for which the upper limit for synaptic coupling is not high enough to enable synchronization. Furthermore, the overall proportion of time spent in random activity versus transient synchronization as reflected by the mean value of the order parameter converges as well. The mean order parameter to which networks converge is principally controlled by the shape of the STDP rule. It is insensitive to the initial synaptic weight distribution and varies only slightly as a function of connectivity (Figure S6A).

Thus, Hebbian STDP promotes the self-organization of spontaneously active neuronal networks into mixture states whose order parameter depends on the shape of the STDP rule (Figure S7). This is clearly the case for the family of networks we consider, in which activity is principally patterned by the recurrent connections. As we discuss in more detail below, this is also likely to be true for the hippocampus during SWS. In other networks and under different input conditions the same principles operate, but they interact with all other circuit elements that shape network activity, i.e., sources of structured input, inhibitory and pacemaker populations, etc. Since STDP need not have control over these additional influences, convergence to mixture states under these conditions need not occur. For example, during active exploratory behavior and REM sleep, the hippocampus exhibits rhythmic theta activity driven by entorhinal and septal inputs. This is in contrast to the mixture-like state seen in SWS when the hippocampus is spontaneously active.

Mixture States and Hippocampal Activity in SWS

During SWS, activity in CA1 and CA3 comprises SPW bursts occurring against a background of low-rate irregular firing (Figure 6A). This regime is qualitatively similar to the mixture states observed in the model networks. Indeed, since the recurrent circuits in CA3 are spontaneously active during SWS, the opposing forces of STDP described above should promote self-organization into a state at the edge of synchrony. But are mixture states

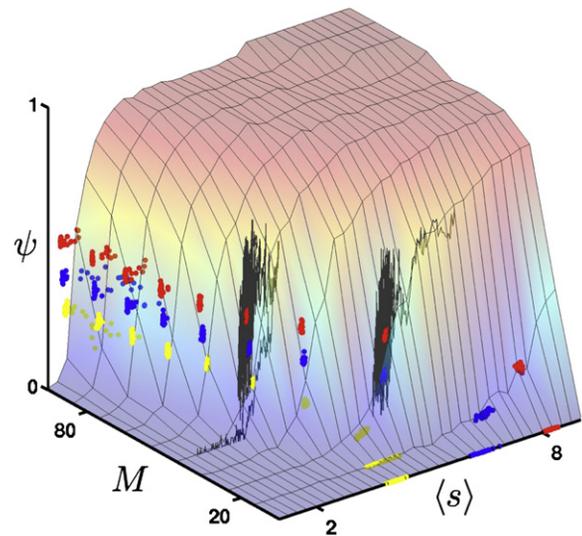


Figure 5. Recurrent Networks Self-Organize in Mixture States under the Influence of Local STDP Rules

A fully connected recurrent network with 100 neurons was progressively pruned to generate a set of ten networks with connection probabilities M between 5% and 95%. Each of the ten network topologies was initialized with 33 bimodal synaptic weight distributions such that the mean synaptic weight ($\langle s \rangle$) ranged between 10% and 90% of the maximum value. This resulted in a matrix of 330 networks with constant connection probabilities along the rows and constant mean synaptic weights along the columns. Activity was simulated in each network with plasticity disabled, and the corresponding order parameter is displayed as the colored surface. Activity was also simulated in each network under three different Hebbian STDP rules, and the colored dots display the corresponding mean weight and average order parameter after steady state is reached. Different dots of the same color correspond to networks with different initial $\langle s \rangle$ (33 for each M), while different colors represent the three different STDP rules (yellow, negative integral; blue, zero integral; red, positive integral). Notice that for $M > 25$ there is a sharp transition between random firing ($\psi = 0$) and highly synchronized activity ($\psi = 1$) as a function of $\langle s \rangle$. Thus, for a given network topology, the mean synaptic weight can be viewed as a bifurcation parameter that controls the degree of synchronization in the network. When Hebbian STDP is enabled, the mean synaptic weight ($\langle s \rangle$) and order parameter ψ leave their initial position on the colored surface and evolve along trajectories, such as the ones shown in black for two different networks ($M = 45, M = 25$). For all networks capable of synchrony ($M > 25$), $\langle s \rangle$ settles within the interval associated with intermediate ψ corresponding to mixture network activity patterns. The time averaged value of the order parameter associated with the steady state does not depend on the initial mean weight (dots of the same color cluster for each M), and depends only slightly on M (clusters of the same color have similar ψ), but is instead controlled by the integral of the STDP rule (clusters of different colors have different ψ).

quantitatively compatible with the statistics of neuronal firing observed during SWS?

To address this question we allowed the 1000 neuron network (Figure S4) to converge to a mixture state when randomly driven at 1 Hz or 0.5 Hz, and then compared the statistics of RS cells in the model with CA1 pyramidal neurons recorded during SWS (Table 1). Notice that the shape, width, and magnitude of cross-correlation functions between pairs of neurons are in good agreement with experimental observations (Figures 6C–6H).

There are two quantitative differences. First, background firing rates are somewhat higher in the model. Second, a larger fraction

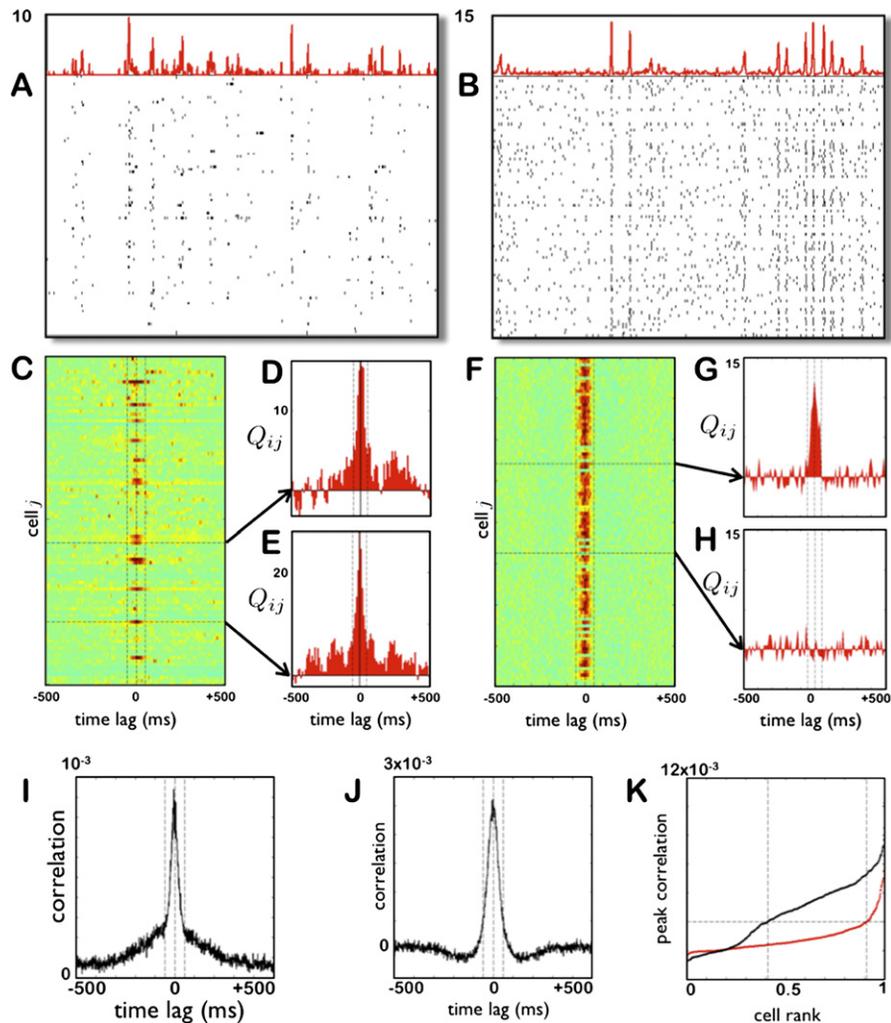


Figure 6. Synchronous Bursts and Correlations in the Hippocampus and in the Model

(A) Spike rasters of 101 simultaneously recorded CA1 pyramidal cells over a period of 15 s in SWS. (B) Spike rasters of 101 regular spiking neurons from the model of Figure S4 over a period of 15 s after the network has converged to a mixture state. The model network is randomly driven at 0.5 Hz. The red curves above the rasters in (A) and (B) are the instantaneous mean firing rates (Hz) of the average cell activity. Notice the prominent vertical striations, corresponding to population bursts, and the low-rate firing during the interburst intervals. (C) Standardized cross-covariances Q_{ij} (Siapas et al., 2005) between a given pyramidal cell i and each of the simultaneously recorded pyramidal cells j are displayed as a pseudocolor plot. Notice that all significant correlations, indicated by hot colors, occur within a narrow window around zero (dotted lines are at ± 50 ms). The two examples displayed in (D) and (E) are representative of the shape of the time difference distributions between the spiking of pairs of correlated neurons. The median μ and σ of the Gaussian fits over all interacting pairs from all eight data sets were 0 and 26 ms, respectively. (F) Same as in (C) but for cells from the model. Notice that the widths of the cross-covariances are very similar in the data and model. (G and H) Examples of correlated (G) and uncorrelated (H) pairs in the model. (I and J) The mean correlation coefficient as a function of time lag for CA1 cells (I) and for cells in the model (J). (K) The peak correlation coefficients for each pair of CA1 cells (red) and cells in the model (black) in ascending order. The values of peak correlations for CA1 and model cells are very similar. A larger fraction of cells in the model is highly correlated, which accounts for the differences in the peak values in (I) and (J).

of neurons participate in each burst in the model (Figure 6K). The first discrepancy is due to the fact that the model is 300 times smaller than CA3. Our analysis indicates that networks approaching the size of CA3 are expected to support mixture states with synchronous bursts occurring against background rates as low as the experimentally observed values (Section S3). The second discrepancy may be due to a variety of factors, including differences in network size, connectivity structure, inhibitory interactions, etc. The understanding of the precise mechanisms controlling the size of synchronous bursts in CA3 remains incomplete; therefore, further studies are required to resolve this issue.

Decoupling through Stimulation

One counterintuitive prediction of the above results is that it should be possible to desynchronize recurrent networks exhibiting persistent synchrony by oversynchronizing them. The reason is that the strength of the decoupling force is inversely proportional to the width of the population bursts, and therefore, tighter bursts should engage the decoupling force more effectively. To test this prediction we simulate the activity in the network of Figure 4 under a Hebbian STDP rule with a large positive integral, which corresponds to a high-order parameter equilibrium mixture state. One minute into the simulation and after the equilibrium

Table 1. Comparison of the Activity of CA1 in SWS and Network Models in the Mixture State

	Firing Rate (Hz)	CV_{isi}	Burst Rate (Hz)	CV_{ibi}	Cells per Burst (%)	CC(0)
CA1	0.20	2.00	0.54	1.71	7.81	0.0008
Model at 0.5 Hz	0.97	0.94	0.78	0.88	46.82	0.0025
Model at 1.0 Hz	3.01	0.65	3.48	0.36	55.80	0.0045

The following values are reported in the columns: (1) number of spikes per second; (2) coefficient of variation (CV) of interspike intervals; (3) number of synchronous bursts per second; (4) CV of interburst intervals; (5) percent of neurons active in each burst; and (6) cross-correlation coefficient (CC) between cell pairs at zero lag (1 ms bins).

mixture state is reached, a quarter of the neurons are driven by external stimulation once a second. This stimulus is capable of generating a very tight population burst that spreads to all neurons in the network (Figure 7C) and thus effectively engages the decoupling force of STDP. Soon after stimulation is turned on, the mean synaptic weight and order parameter converge to lower values (Figures 7A, 7C, and 7F) and network activity between stimulation pulses becomes random, thus confirming our prediction. Notice that after stimulation is turned back off, activity in the network settles back to the original highly synchronous mixture state (Figures 7A, 7D, and 7G). It is possible that the stimulation itself reduces the amount of network synchrony by directly affecting membrane dynamics. However, this is not the case, since the same pattern of stimulation with plasticity disabled is incapable of affecting the order parameter in the network (Figure 7). Thus, the stimulated bursts must engage the decoupling force of STDP in order to desynchronize activity. This observation suggests that plasticity mechanisms may contribute to some of the therapeutic effects of deep brain stimulation.

Decoupling in Hippocampal Networks

Given that SPW bursts produce cross-correlations with a narrow peak around zero (Figure 6C), our analysis predicts that they should lead to decoupling in the recurrent circuits of CA3. How can such decoupling be experimentally detected, without the ability to directly measure the strength of multiple synapses in freely behaving and naturally sleeping animals? We examine the behavior of two different proxies of synaptic coupling: the mean pairwise firing-rate correlation and the slope of the field excitatory postsynaptic potential (fEPSP) in stratum radiatum of CA3 produced by electrical stimulation of the recurrent connections.

In the mixture state the mean synaptic weight, $\langle s \rangle$, is approximately proportional to the pairwise firing-rate correlations averaged over all cell pairs, $\langle r \rangle$ (Figure 8A). Of course, correlations do not in general reflect the level of synaptic coupling. For example, in a network driven by structured external input, the correlations will likely reflect the structure of the input more than anything else. However, in the case of recurrent networks driven by random activity, correlations do provide an experimentally accessible, albeit indirect, measure of coupling (Figure 8A). Furthermore,

this special case is the relevant one, as previous studies show that SPW bursts are generated within the CA3 recurrent network when external input to the network through the entorhinal cortex is highly attenuated and CA3 is driven by random internal fluctuations in activity (Buzsaki, 1986; Chrobak and Buzsaki, 1996; Csicsvari et al., 2000). Furthermore, under these conditions the firing of CA3 cells is the dominant cause of spiking in CA1 (Csicsvari et al., 2000); hence, CA1 recordings can be used as a “read-out” of CA3 activity during SPW bursts. Consequently, a testable prediction of the model is that population bursts in the hippocampus should produce a decrease in the mean pairwise correlation $\langle r \rangle$ between CA3, as well as CA1, cells.

Two previous studies examining the reactivation of CA1 activity patterns in sleep have reported such a decay in correlations (Wilson and McNaughton, 1994; Kudrimoti et al., 1999), but the reason behind the phenomenon has remained a mystery. The present work not only offers a mechanistic hypothesis of how this phenomenon occurs, but also extends the experimental findings in light of decoupling. In particular, since the goal is to make inferences about the level of coupling in CA3 from recordings obtained in CA1, pairwise correlations were computed based on the firing of cells during ripples, when CA1 is almost exclusively driven by inputs intrinsically generated within the recurrent circuits of CA3 (Csicsvari et al., 2000). We find that in all eight sleep sessions from three animals, the mean correlation computed from ripples occurring in the first halves of SWS epochs is always higher than the corresponding mean computed from ripples occurring in the second halves of SWS epochs (Figure 8B and Experimental Procedures). Furthermore, this decay in correlations is statistically significant not only overall ($p < 0.01$, sign test over eight data sets; $p < 0.0005$, paired t test, one comparison per SWS epoch), but also can be demonstrated within five of the eight individual sleep sessions ($p < 0.05$, nonparametric bootstrap; Figure 8B). In contrast, pairwise correlations computed from inter-ripple intervals need not reflect the level of coupling in CA3, because the identity of the inputs driving activity in CA1 during these periods is less clear. Consistently we find no decay in correlations based on interripple intervals (Figure S10C). To control for imperfect ripple identification, we also compared correlations computed in the first versus second halves of SWS epochs without ripple segmentation and confirmed the presence of decay both overall ($p < 0.01$, sign test) and in all eight individual sleep sessions ($p < 0.05$, nonparametric bootstrap; Figure S10A).

Can systematic changes in firing rates account for the observed decay in correlations? To address this issue we used a combination of two analytical approaches. First, we coded the ripples of each SWS epoch in binary format, giving a matrix whose ij entry is 1 if cell i fired one or more spikes in ripple j . Second, before comparing the mean correlation computed from the first half of this binary matrix to the corresponding mean from the second half, for each row (cell) we equalized the number of 1's between the two halves by random insertion. This way neurons are made equally likely to participate in ripples early and late in SWS, and the only difference is in their tendency to do so in a coordinated manner. Furthermore this analysis is insensitive to firing rates within ripples and changes in firing rates across ripples. We find that even after this manipulation, a highly significant correlation decay is observed ($p < 0.005$, paired t test, one

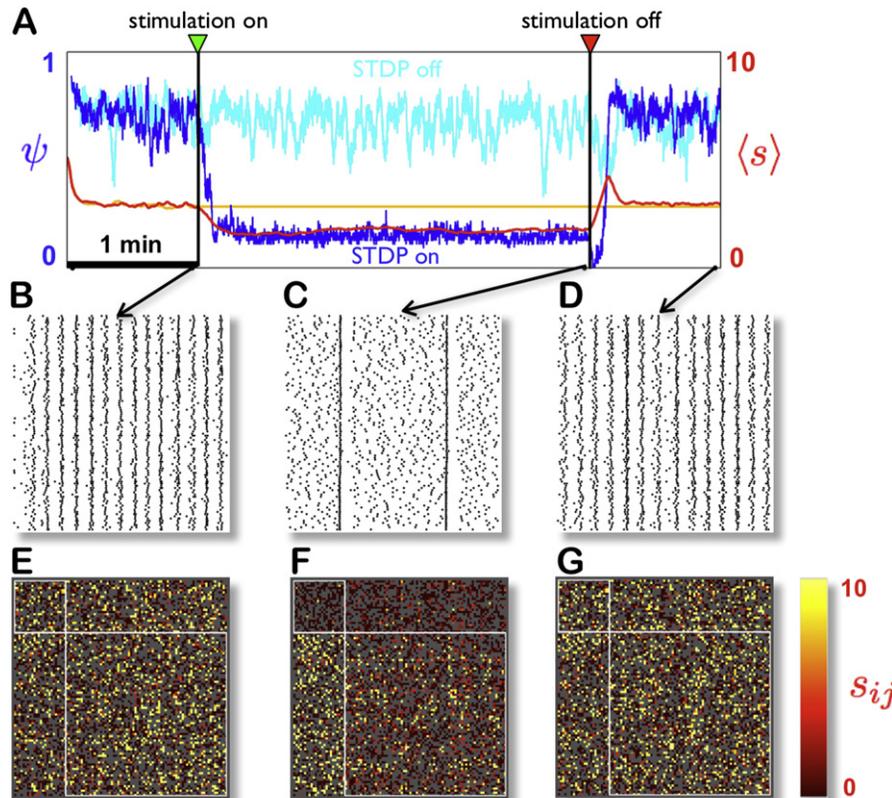


Figure 7. Decoupling through Stimulation in Recurrent Networks

The same network as in Figure 4 was simulated under the STDP rule shown in Figure 1E, which has a positive integral and hence a bias toward synaptic strengthening. Such an STDP rule can be considered abnormal in the sense that it poises network activity in a state that is highly synchronized, as can be seen from the high value of ψ in the first minute of simulation (A). Stimulation consisting of 5 ms current injections every second given to neurons 1 through 25 (upper left square marked by the white lines in [E]–[G]) is turned on 1 min into the simulation. Soon after, activity in the network becomes markedly desynchronized, as can be seen by the drop in ψ (blue line). This is accompanied by a corresponding drop in the mean synaptic weight (red line). The stimulation is given for 3 min, and as long as it is present, the activity in the network is desynchronized. Soon after the stimulation is turned off, the network returns to the synchronized state. If STDP is disabled when the stimulation is turned on, the order parameter (light blue line) and mean synaptic weight (orange line) are not affected by the stimulation. Also shown are two-second spike rasters (B–D) and snapshots of the synaptic weight matrix (E–G) right before stimulation is turned on (B and E), right before stimulation is turned off (C and F), and at the end of the run (D and G). Notice that stimulation leads to strengthening of the connections from the stimulated subset to the rest of the neurons (bottom left block in [F]), while interconnections within the stimulated subset are weakened (top left block in [F]).

comparison per SWS epoch; see [Experimental Procedures](#)). This strongly argues that firing rate changes are not the source of the correlation decay.

In order to further investigate whether the observed decay in correlations is indeed due to weakening of the CA3 recurrent synapses, we studied the evolution of the fEPSP slope in stratum radiatum of CA3 evoked by electrical stimulation of the fimbria in a freely behaving, chronically implanted rat. At the beginning of each experiment, the stimulus intensity was adjusted to evoke a clear, low-latency unimodal fEPSP in CA3 (Figure 8D). Next, probe pulses of fixed intensity were delivered every 30 s and the corresponding fEPSP slope was computed while the rat slept or ran on a linear track for reward (Figure 8C). We found that in the sleep period following experience, there was a gradual and highly significant decrease of fEPSP slope ($r = -0.43$, $p < 10^{-10}$), consistent with the prediction that SPW bursts in CA3 engage the decoupling force of STDP and thus weaken the recurrent synaptic connections.

DISCUSSION

Generality of the Decoupling Force

The main conclusion of our results is that synchronous bursts, occurring within recurrent circuits with delays, should lead to the selective decoupling of the coactivated neurons. We predict that such a decoupling force is present in circuits that meet the following common criteria: (1) presence of recurrent connections, (2) nonzero axonal conduction delays, (3) Hebbian STDP in the recurrent connections, and (4) conditions, such as internal dynamics or external input, producing synchronous population bursts. Since many networks in the adult and developing brain meet the above criteria, we expect the decoupling force of STDP to be a very general and important phenomenon.

The convergence of networks to mixture states is a consequence of the coupling and decoupling forces of STDP that applies to randomly driven networks for which the level of synchronous activity is principally controlled by the strength of the

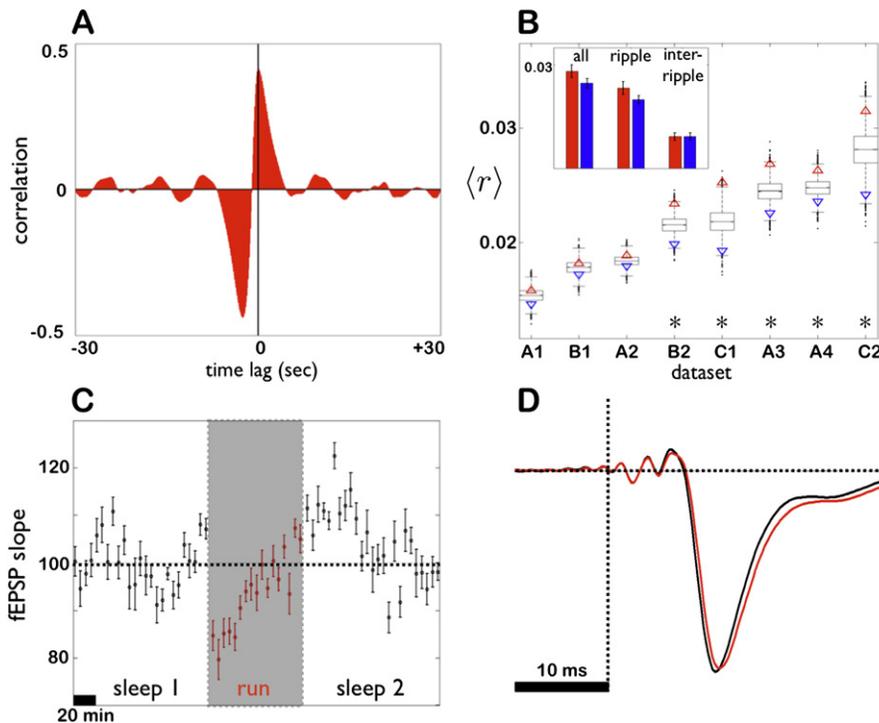


Figure 8. Evidence for Decoupling through Synchrony in the Hippocampus

(A) Correlation coefficient between the mean pairwise correlation (r) and the mean synaptic weight (s) as a function of time lag, computed from the activity in a model network poised in a mixture state. Notice the significant positive correlation at zero lag, justifying the use of $\langle r \rangle$ as a measure of coupling. The significant negative correlation peaking at a negative lag is due to the decoupling force: intervals of high synchrony, marked by high correlations $\langle r \rangle$, lead to decoupling and are thus followed by intervals of low $\langle s \rangle$.

(B) Mean pairwise correlations based on activity during ripples from the first halves of SWS epochs, $\langle r \rangle_{R1}$ (red), and the second halves, $\langle r \rangle_{R2}$ (blue). Results from eight sessions (A1–A4, B1 and B2, and C1 and C2) in three animals (A, B, and C) are shown. Notice that for all sessions $\langle r \rangle_{R1} > \langle r \rangle_{R2}$, indicating a decay in pairwise correlations. Boxplots show the sampling distributions of the mean pairwise correlations estimated for each session with a nonparametric bootstrap procedure (Experimental Procedures). Significant decay in correlations within a session is indicated by an asterisk (five out of eight sessions, $p < 0.05$). In the sessions for which the decay in correlations could not be demonstrated as significant, the actual mean correlations were low, suggesting an issue with detectability due to the random sampling of a larger fraction of noninteracting cells. Across sessions the decay in correlations was significant both when all activity was considered and when activity during ripples only was considered ($p < 0.01$, sign test; inset, left and middle bar pairs). There was no decay in correlations when activity during interripple intervals only was considered ($p > 0.25$, sign test; inset, right bar pair).

(C) Time evolution of the slope of CA3 fEPSPs evoked by electrical stimulation of the ipsilateral fimbria (Experimental Procedures). Stimulation pulses are delivered once every 30 s and each point represents the average slope of ten fEPSP measurements normalized by the average fEPSP slope (1.12 mV/ms). During the period marked by the gray rectangle, the rat runs back and forth on a linear track for liquid reward. Consistent with the hypothesis that experience leads to strengthening of CA3 recurrent connections, fEPSP slopes increase throughout the run period. Consistent with the conjecture that SPW bursts in SWS lead to decoupling of the CA3 network, fEPSP slopes progressively decay in the sleep period following experience.

(D) Mean fEPSPs in the first and second half of sleep 2 (black and red traces, respectively).

recurrent excitatory connections. Under these circumstances STDP acts as a local rule that tunes the strength of synaptic interactions to poise the global network dynamics at the border between randomness and synchrony. The coupling and decoupling forces of STDP operate in more complex networks as well, but convergence to mixture states need not occur in general, as there may be a number of additional influences that pattern network activity and affect the level of synchrony besides the strength of synaptic coupling, such as structured input, pacemaker neurons, excitatory-inhibitory interactions, etc.

A recent study has reported that synchronous bursts lead to the selective strengthening of recurrent connections between coactivated neurons, instead of exerting a decoupling force as we observe (Morrison et al., 2007). The explanation for the discrepancy is that for technical reasons there were only dendritic,

and not axonal, delays implemented in Morrison et al. (2007). As a result, when two reciprocally connected neurons fire simultaneously, spikes arrive at the presynaptic sites instantaneously and before the action potentials have backpropagated to the postsynaptic sites, thus producing input-output time differences that lead to strengthening under Hebbian STDP. This is physiologically unrealistic because one generally expects axonal delays to be larger than dendritic delays. Furthermore, consistent with our prediction and contrary to Morrison et al. (2007), experiments that were used to demonstrate STDP in CA3 synapses directly showed that making a pair of reciprocally connected cells fire synchronously leads to weakening of the synapses between them (Debanne et al., 1998). Using only dendritic delays corresponds to using an anti-Hebbian STDP rule within our analysis framework. Under these conditions synchronous bursts lead to

synaptic coupling, thus engaging a positive feedback mechanism that is the source of the reported “synfire explosions” (Morison et al., 2007). This is contradicted by the observation that synchronous bursts in the hippocampus are associated with a decorrelation of network activity, rather than lead to highly synchronized states.

Role of Causality and Temporal Order

The intuition behind the decoupling force of STDP is very simple when the synchronous burst in the network is driven by external input (Figure 1A and Figure 7). When the synchronous burst builds up internally within the network, however, some recurrent synapses must causally contribute to the firing of each neuron, and these synapses will be strengthened by Hebbian STDP. But how can STDP then lead to desynchronization if it is not weakening the synapses causally responsible for the generation of the synchronous burst?

This apparent paradox is resolved by noticing that when activity builds up randomly within the network, different synapses are responsible for firing the postsynaptic neuron at different points in the population burst and across bursts (Figure S11). Hence one cannot speak of the set of synapses causally responsible for synchrony because this set is not fixed—its membership changes with time. Under such random buildup conditions, any given connection contributes only occasionally, and therefore weakly, to the firing of the postsynaptic neuron. As a result, pairwise firing-rate cross-correlation functions are peaked around zero and their width reflects the duration of population bursts (Figure 1).

Indeed, the decoupling force of STDP would be defunct if all bursts progressed in a stereotyped, causal, chain-like fashion. Under these circumstances the cross-correlation function between cell pairs would be narrowly peaked and centered at a nonzero offset, reflecting axonal conduction delays and the consistent ordering of cell firing. This is not the case for either the model networks (Figure 1A) or the hippocampus (Figure 6C), where cross-correlation functions are broader and centered around zero. Furthermore, notice that the presence of temporal order alone—cell A fires before cell B—need not interfere with the decoupling force. A stronger condition is necessary: the temporal order must be causal, i.e., cell B fires after it has received the input from cell A.

Sequential Firing in the Hippocampus in SWS

Several studies have reported evidence for the sequential reactivation of experience-specific activity patterns in the hippocampus within population bursts during SWS (Skaggs and McNaughton, 1996; Nadasdy et al., 1999; Lee and Wilson, 2002; O’Neil et al., 2006). Do the experimental observations in these studies violate the temporal ordering condition necessary for decoupling?

Each of the above studies takes one of the following two analytical approaches. In the first, pairwise cross-correlation asymmetries from the awake state are compared with the corresponding asymmetries from a subsequent sleep period (Skaggs and McNaughton, 1996; O’Neil et al., 2006). The presence of a significant positive correlation between the awake and sleep asymmetries constitutes the evidence for sequence replay. However, as discussed in O’Neil et al. (2006), the actual value of the correlation

coefficient is small ($r < 0.1$) and most asymmetries fall in a cloud around the origin (Skaggs and McNaughton, 1996), indicating that many pairs do not have a sequence bias. Furthermore, cell pairs with no reactivated sequence bias and nearly symmetric cross-correlation histograms still replay cofiring patterns during sleep (O’Neil et al., 2006). This strongly argues that cells that cofire in the awake state cofire within SWS bursts, but rarely (yet above chance) in an order consistent with the awake state.

In the second approach, activity in the awake state is used to define an order template, and subsequent sleep firing patterns are searched for sequences matching the order template (Nadasdy et al., 1999; Lee and Wilson, 2002). The fact that sleep sequences contain more matches than expected by chance constitutes the evidence for sequential reactivation. This is not contradicted by the observation that the vast majority of sequences are in an order that does not match the awake template. In particular, in Lee and Wilson (2002) there are 655 cell pair matches and 600 nonmatches out of 1255 trials, a deviation which is not significantly different from the 628 matches expected by chance. For triplets and longer sequences, the deviations from chance are significant, but small. The numbers are 57 observed matches instead of 43 expected by chance out of 259 trials for triplets, and 35 observed matches instead of 11 expected by chance out of 270 trials for longer sequences. Finally, because these studies encode bursts of spikes as single events occurring at the time of the first spike in the burst, the order bias does not directly translate into a cross-correlation asymmetry bias, which is the relevant quantity as far as decoupling is concerned.

In summary, previous studies have demonstrated the presence of a significant bias toward sequential reactivation of activity patterns during population bursts in SWS. Yet the ordered patterns constitute a small fraction of the total number of bursts, and therefore do not violate the ordering condition for decoupling.

Relevance for Hippocampal Function

It is believed that mnemonic information initially established in hippocampal circuits is gradually transferred into neocortical regions, and SPW bursts have been implicated in this process (Buzsaki, 1989; Squire, 1992; Siapas and Wilson, 1998). It remains unclear, however, how old information is selectively and gradually erased from hippocampal circuits. Our findings suggest that SPW bursts may not only establish strong connections to and between target cortical areas, but simultaneously act to gradually erase intrahippocampal associational connections within the CA3 subfield as mnemonic information is transferred outside the hippocampal formation.

In particular, because of the spatial selectivity of hippocampal neurons and their precise firing with respect to theta oscillations, awake experience results in consistent temporal patterning of firing within the window of plasticity (Skaggs et al., 1996). This is believed to lead to the formation of strongly connected subnetworks in CA3, linking the subset of neurons active in the awake state. The addition of these strong connections increases the level of coupling above the equilibrium point. As a consequence, during subsequent SWS, the CA3 network is more synchronized than usual and its relaxation to the equilibrium mixture state is associated with the preferential expression of patterns biased in favor of the neurons from the strongly connected subnetwork. Our results

argue that this burst re-expression of patterns decouples the strongly connected CA3 subnetwork, thus effectively erasing the experience-specific hippocampal memory trace and returning the CA3 network to equilibrium. Notice that the decoupling force of STDP will not completely erase the given memory trace, but will weaken it until it matches the typical strength of older traces.

What regulates the identity of neuronal patterns expressed by the hippocampus during sleep? As discussed above, awake behavior increases the probability of expression of certain experience-specific neuronal patterns (memories) during subsequent SWS. In contrast, because of the decoupling force of STDP, the expression of any given pattern decreases the probability of its future re-expression due to the weakened connections between the participating neurons. Together these two mechanisms ensure that the number of times a pattern is expressed in sleep is proportional to how strongly it was embedded in the hippocampal network. These mechanisms also play a homeostatic role, preventing any given memory from permanently taking over circuit resources.

EXPERIMENTAL PROCEDURES

Computational Model

We use a modified and extended version of the network simulator introduced by Izhikevich (Izhikevich, 2003, 2006). All but one (Figure S4) of the networks included in the text consisted of 100 RS neurons with dynamics given by

$$\frac{dv}{dt} = 0.04v^2 + 5v + 140 - u + I \quad (1)$$

$$\frac{du}{dt} = a(bv - u) \quad (2)$$

with spikes generated when $v \geq 30$ and associated afterspike resetting rules $v \leftarrow c$, and $u \leftarrow u + d$, where $a = 0.02$, $b = 0.2$, $c = -65$ and $d = 8$.

Synaptic interactions and external input were introduced via the I term. In particular, consider neuron n in a network of size N , let the conduction delay from neuron m to n be δ_{nm} and the synaptic strength at time t be $s_{nm}(t)$. Let neuron m generate spikes at times τ_{mi} where individual spikes are indexed by i , and let neuron n receive external input at times t_{nj} via a synapse of strength r_n . Then

$$I_n(t) = \sum_{m=1}^N \sum_{\tau_{mi} \leq t} s_{nm}(t) \delta(t - \tau_{mi} - \delta_{nm}) + \sum_{t_{nj} \leq t} r_n \delta(t - t_{nj}). \quad (3)$$

Thus, a spike fired by neuron m at time τ_{mi} causes the membrane voltage of neuron n to instantaneously jump by $s_{nm}(t)$ (mV) at time $\tau_{mi} + \delta_{nm}$. In this sense $s_{nm}(t)$ represents the peak amplitude of the EPSP seen in neuron n in response to input from m . The EPSP peak is attained instantaneously upon input arrival at n , and the EPSP decays according to the membrane dynamics (Equation 1).

No inhibitory interneurons were included in these simulations because it was determined that their presence was not essential for the establishment of network oscillations and synchronization. In particular, a model network of 1000 neurons that incorporated inhibition, diverse neuronal populations, connection divergence, and distribution of delays mimicking the hippocampal CA3 field, as well as “nearest-neighbor” STDP implementation, behaved in qualitatively the same way (Figure S4) as the simpler systems we focused on. We also confirmed that the decoupling force was present in a larger network consisting of 10,000 neurons. Since the decoupling force of Hebbian STDP depends principally on the presence and structure of population bursts, irrespective of the mechanisms that produce them, the main result is insensitive to the details of the model. In the 100 neuron networks the number of postsynaptic targets per neuron ranged between 5 and 95. Each connection was given a random delay, δ_{nm} , between 1 and 20 ms. This range and distribution of delays is consistent with experimental measurements of axonal conduction delays in

cortico-cortical connections (Swadlow, 1994). The corresponding range for the hippocampal CA3 field is 1 to 10 ms (Miles et al., 1988; Ishizuka et al., 1990). Unless otherwise stated, neurons were driven by 10 Hz uncorrelated Poisson trains of $r_n = 20$ mV EPSPs that were capable of firing neurons roughly 50% of the time. The time and identity of all spikes were stored and used in the calculation of the order parameter evolution. Similarly, the synaptic weight matrix was stored every 1 ms for subsequent analysis.

STDP Implementation

Let presynaptic neuron m generate a total of M spikes at times τ_j , $1 \leq j \leq M$, and postsynaptic neuron n generate a total of N spikes at times t_i , $1 \leq i \leq N$. Let the conduction delay from m to n be δ_{nm} and the synaptic strength at time t be $s_{nm}(t)$. Let $\Delta t_k = t_i - (\tau_j + \delta_{nm})$, $1 \leq k \leq NM$ be the time difference between postsynaptic spike i and the arrival of presynaptic spike j at neuron n . The total change in synaptic strength produced by the pairing Δt_k is $F(\Delta t_k)$, where

$$F(\Delta t) = \begin{cases} A_+ e^{-\frac{\Delta t}{\tau_+}} & \text{if } \Delta t \geq Z \\ 0 & \text{if } \Delta t \in (-Z, +Z) \\ A_- e^{\frac{\Delta t}{\tau_-}} & \text{if } \Delta t \leq -Z \end{cases} \quad (4)$$

The parameter Z can be used to nullify the effect of pairings in the interval $(-Z, Z)$ and the familiar form of the rule is recovered by setting $Z = 0$. All STDP variants we consider are additive in the sense that the total change in synaptic strength results from adding the contributions of individual pairings.

Depending on which pairings Δt_k contributes to the synaptic change, we distinguish two classes of implementations. In “all-to-all” implementations all NM possible pairings contribute. In nearest-neighbor implementations, at most, $N + M$ pairings contribute; i.e., each presynaptic and postsynaptic spike is paired only with its immediately preceding postsynaptic and presynaptic spike, respectively.

Within each class two additional manipulations can be independently enabled. First, presynaptic and postsynaptic spikes can be assigned efficacies when suppression is enabled, following the model in Froemke and Dan (2002). Second, pairings falling in the interval $(-Z, Z)$ can be ignored by setting $Z > 0$, thus nullifying $F(\Delta t)$ in the interval $(-Z, Z)$.

The following general equations can be specialized for the implementations discussed above. For every neuron k define a potentiation function $S_+^k(t)$ and a depression function $S_-^k(t)$ as follows:

$$S_+^k(t) = \sum_{t_i \leq t} \epsilon_{pre}^k(t_i) A_+ e^{-\frac{t-t_i}{\tau_+}} I(t, t_i, t_{i+1}) \quad (5)$$

$$S_-^k(t) = \sum_{t_i \leq t} \epsilon_{post}^k(t_i) A_- e^{-\frac{t-t_i}{\tau_-}} I(t, t_i, t_{i+1}) \quad (6)$$

where

$$\epsilon_{pre}^k(t_i) = 1 - e^{-\frac{t_i - t_{i-1}}{\tau_{pre}}}$$

and

$$\epsilon_{post}^k(t_i) = 1 - e^{-\frac{t_i - t_{i-1}}{\tau_{post}}}$$

are spike efficacies, and I takes on the values of 0 and 1 only and specifies the time interval over which the preceding term is present. In particular, in all-to-all implementations $I(t, t_i, t_{i+1}) = H(t - t_i - Z)$, where $H(t)$ is the Heaviside step function, thus specifying the interval $[t_i + Z, \infty)$. In nearest-neighbor implementations $I(t, t_i, t_{i+1}) = H(t - t_i - Z)H(t_{i+1} + Z - t)$, thus specifying the interval $[t_i + Z, t_{i+1} + Z)$. The instantaneously updating synaptic strength $s'_{nm}(t)$, where m is the presynaptic and n the postsynaptic neuron, is given by:

$$s'_{nm}(t) = \sum_{t_i \leq t} \epsilon_{post}^n(t_i) S_+^m(t_i - \delta_{nm}) H(t - t_i) + \sum_{\tau_j \leq t - \delta_{nm}} \epsilon_{pre}^m(\tau_j) S_-^n(\tau_j + \delta_{nm}) H(t - \tau_j - \delta_{nm}) \quad (7)$$

where the first term is the contribution of postsynaptic spike t_i , and the second term is the contribution of presynaptic spike τ_j .

Finally, when no saturations occur, the actual synaptic strength $s_{nm}(t)$ is given by $s'_{nm}(t)$ smoothed with a kernel with exponential impulse response $h(t)$, i.e.,

$$s_{nm}(t) = h(t) * s'_{nm}(t) \quad (8)$$

$$h(t) = \frac{1}{\tau_{STDP}} e^{-\frac{t}{\tau_{STDP}}} \quad (9)$$

If $s_{nm}(t)$ reaches either boundary of the interval $[0, s_{max}]$, only updates moving $s_{nm}(t)$ away from the extremes and toward the range of admissible strengths are applied.

Unless stated otherwise, all STDP rules were additive with $\tau_+ = \tau_- = 20$ ms, and connection weights, $s_{nm}(t)$, were allowed to vary between 0 and 10 mV. For all simulations included in the main text, biases in the STDP rules were introduced via the A_+ and A_- parameters and were as follows: all unbiased rules $A_+ = 1$, $A_- = -1$; Figure 1C, $A_+ = 1$, $A_- = -1.4$; Figure 1E, $A_+ = 1.4$, $A_- = -1$; red dots in Figure 5, $A_+ = 1.1$, $A_- = -1$; blue dots in Figure 5, $A_+ = 1$, $A_- = -1$; yellow dots in Figure 5, $A_+ = 1$, $A_- = -1.1$. We used an STDP time constant $\tau_{STDP} = 1$ s and suppression time constants $\tau_{pre} = 28$ ms and $\tau_{post} = 88$ ms. Suppression was disabled by default by setting all efficacy terms to 1. Zeroing of the central region of the STDP rule was disabled by default by setting $Z = 0$.

Mechanism of Network Oscillations in the Model

Buildup of excitation in recurrent networks occurs when the simultaneously active fraction of neurons at a given point in time is capable of firing a larger fraction at a subsequent point in time. This condition is satisfied when the product of connection divergence and interaction strength exceeds some threshold. Once a critical fraction of neurons, dependent on the above product, becomes simultaneously active (for example, due to coincident inputs), activity in the network avalanches into a population burst. In the networks we study this population burst is terminated because of the dynamics of individual neurons. In particular, the threshold for firing of RS neurons increases with each spike fired and recovers slowly. Therefore, as neurons get recruited in the avalanche of activity, their firing thresholds keep increasing with every spike they fire, until ultimately they become completely refractory. Thus, the avalanche of excitation leads to a population burst, which is followed by a network refractory period in which no neuron fires. Under conditions when the critical fraction of neurons needed to trigger a population burst is small, random input drives the network into its next avalanche as soon as it has recovered from the refractory period associated with the previous one. This leads to the regular sequence of population bursts followed by silent refractory periods, which together constitute the apparent network oscillation.

Order Parameter

The order parameter ψ quantifies the fraction of small time bins for which the instantaneous firing rate in the network is outside of the range expected under random firing (Figure S1). By definition ψ can only detect significant deviations in the level of simultaneous firing and is therefore insensitive to cross-correlation structure at nonzero time lags. Despite its simplicity, ψ captures well the level of synchronization in the network, and for a wide range of dynamic regimes, including mixture states, it appears to provide nearly identical quantification to the more complex and computationally expensive r^2 measure of synchrony introduced by Pinsky and Rinzel (1995) (Figure S2). An order parameter can be similarly computed based on the activity in a subset of the network (Figure S4).

Experimental Analysis

Using chronic multitetrode arrays we recorded the simultaneous activity of CA1 neurons from three animals (A, B, and C) during eight 2-hr sleep sessions (except for A3, which lasted 53 min). In particular, rats were implanted with custom 24 tetrode microdrive arrays and data were collected using a 24-bit data acquisition system developed in our lab. Only putative pyramidal cells that met isolation criteria and had mean firing rates smaller than 1 Hz were included in the analysis. The number of cells analyzed and the total number of cells recorded in each session were as follows: A1, 59/104; A2, 97/156;

A3, 69/108; A4, 101/159; B1, 83/132; B2, 76/113; C1, 33/56; C2, 27/42. For each session the spiking of N neurons was binned in $T/100$ ms bins, yielding an $N \times T$ matrix Q , with instantaneous firing rates in the rows, and “population vectors” in the columns (Kudrimoti et al., 1999). SWS epochs were identified based on the power content of local field potentials (LFPs) and video recordings of the sleep sessions. Columns falling outside of SWS epochs were deleted from Q (on average 28% of all columns). All remaining columns were labeled as occurring in the first “1” or second “2” halves of their corresponding SWS epochs, and as overlapping a ripple “R” or falling within an interripple interval “I.” By definition 50% of the bins were labeled “1”, and on average, 10% of the bins were labeled “R”.

The objective of the analysis was to compare the mean pairwise firing-rate correlation in the first halves of SWS epochs to the corresponding mean in the second halves. Pairwise correlations were computed in three different ways: based on all activity (Figure S10A), based on activity during ripples only (Figure S10B), and based on activity during interripple intervals only (Figure S10C). This was done by using different subsets of the columns of Q to compute the mean correlation between all unique pairs of rows and then averaging. In particular, when all activity was used, Q was split into Q_1 and Q_2 , according to the “1” or “2” label assigned to each column, and the mean pairwise correlation between the rows of each matrix was computed, giving $\langle r \rangle_1$ and $\langle r \rangle_2$ (Figure S10A). When activity during ripples only (interripple intervals only) was used, the submatrix $Q_R(Q_i)$ was formed first, based on the “R” (“I”) column labels. The rest of the analysis proceeded as above to give the matrices Q_{R1} and $Q_{R2}(Q_{i1}$ and $Q_{i2})$ and the corresponding mean correlations $\langle r \rangle_{R1}$ and $\langle r \rangle_{R2}$ ($\langle r \rangle_{i1}$ and $\langle r \rangle_{i2}$) (Figures S10B and S10C). Finally, we repeated the entire analysis starting with a binary version of the matrix Q , in which every nonzero entry was set to 1.

The advantage of using all data when computing correlations is that the analysis does not depend on successful ripple identification and therefore provides a control against improper segmentation. A major disadvantage, however, is that correlations are sensitive to the firing rate nonstationarity introduced by lumping together high-firing-rate ripples with low-firing-rate interripple intervals. In this case factors that need not relate to coupling, such as the rate of ripple occurrence, can influence the magnitude of pairwise correlations. In contrast, analysis based on activity during ripples only depends on successful ripple identification, yet correlations are more sensitive to the identity of neurons participating in SPW bursts than to nonspecific factors. Furthermore, since activity in CA1 during ripples is driven by inputs intrinsically generated within the recurrent circuit of CA3, the mean correlation in CA1 during ripples is more likely to reflect the level of coupling in CA3. Finally, correlations based on interripple intervals need not relate to coupling in CA3, because the identity of the inputs driving CA1 activity in these periods is less clear.

Estimating the significance of the difference between the correlations computed in the first versus second halves of SWS epochs is difficult, because a sample of all unique pairwise correlations is not composed of independent measurements. For example, knowing that cell pairs ij and ik are strongly correlated implies that cells j and k are likely to be correlated as well. Thus, a t test over all individual pairwise correlations greatly overestimates the significance of the difference. We used two separate approaches instead. First, instead of individual pairwise correlations (21,105 comparisons), we compared the means of pairwise correlations from the first versus second halves of each SWS epoch (31 comparisons, paired t test, $p < 0.0005$), or each data set (eight comparisons, sign test, $p < 0.01$, insert of Figure 8B). Second, we used a non-parametric bootstrap procedure to estimate the sampling distribution of the mean correlation in each case. In particular, given a matrix Q_X (i.e., Q , Q_R , or Q_I), we generated 10,000 random subsets Q_{Xi} by drawing half of the columns of Q_X at random, and computed the corresponding mean pairwise correlations $\langle r \rangle_{Xi}$. The significance of the effect was determined based on the location of the actual correlations with respect to the distribution of $\langle r \rangle_{Xi}$ (Figure S10).

We found that when activity during ripples only was considered, the mean correlations in the first halves of SWS were significantly higher than the corresponding means in the second halves (Figure S10B). This decay in correlations was present also when all activity was considered (Figure S10A), but was absent when activity during interripple intervals only was analyzed (Figure S10C). The exact same results were obtained when we used the binary version of the matrix Q , indicating that pairwise correlations reflected the patterns of

coactive neurons, rather than global comodulation of firing rates across the cell population. Overall these findings are in excellent qualitative and quantitative agreement with previous work (Wilson and McNaughton, 1994; Kudrimoti et al., 1999). They extend previous work by showing that the decay in correlations is detectable in the mean that is taken over all unique cell pairs, without consideration of activity patterns during previous periods of wakefulness, as well as in the binary correlations. Furthermore, the decay is demonstrated with a conservative statistical procedure within individual sleep sessions. Finally, a statistically significant decay in correlations is demonstrated when activity during ripples only is considered, whereas previous work has reported a nonsignificant trend in the same direction (Kudrimoti et al., 1999).

Stimulation Experiments

For the stimulation experiments a rat was chronically implanted with a recording tetrode ($B = -4.3$, $L = 4$) and a monopolar stimulating electrode ($B = -3.3$, $L = 4$; stainless steel, 140 μm diameter, A-M Systems). After recovery from surgery the recording tetrode was slowly lowered to stratum radiatum of CA3 and the stimulating electrode was positioned in the ipsilateral fimbria (approximately 4 mm below the surface of the brain). During each experimental session the rat was allowed to sleep for 125 min (PRE), then ran on a linear track for water reward for 90 min (RUN), and was allowed to sleep again for 125 min (POST). Throughout the recording sessions a probe stimulation pulse of fixed intensity was delivered to the fimbria every 30 s (200 μs biphasic pulse) and the slope of the resulting fEPSP was measured. The stimulus intensity was adjusted before the start of each session to produce a low-latency, half-maximal unimodal fEPSP (current between 40 and 100 μA). The DC bias current of the stimulus isolator was carefully monitored and maintained in the nA range throughout the recordings.

SUPPLEMENTAL DATA

The Supplemental Data for this article can be found online at <http://www.neuron.org/cgi/content/full/58/1/118/DC1/>.

ACKNOWLEDGMENTS

We thank Ming Gu and Casimir Wierzynski for generously contributing data for the experimental analysis and offering many helpful suggestions. We also thank Gilles Laurent, Erin Schuman, and Andreas Tolia for many useful discussions and comments on the manuscript. This work was supported by the Caltech IST Center for Biological Circuit Design, a 21st Century McDonnell Foundation Award, the Bren Foundation, the McKnight Foundation, and the Whitehall Foundation.

Received: June 13, 2007

Revised: September 29, 2007

Accepted: January 25, 2008

Published: April 9, 2008

REFERENCES

- Abbott, L., and Blum, K. (1996). Functional significance of long-term potentiation for sequence learning and prediction. *Cereb. Cortex* 6, 406–416.
- Abbott, L., and Nelson, S. (2000). Synaptic plasticity: taming the beast. *Nat. Neurosci.* 3, 1178–1183.
- Bains, J., Longacher, J., and Staley, K. (1999). Reciprocal interactions between CA3 network activity and strength of recurrent collateral synapses. *Nat. Neurosci.* 2, 720–726.
- Bell, C., Han, V., Sugavara, Y., and Grant, K. (1999). Synaptic plasticity in the mormyrid electrosensory lobe. *J. Exp. Biol.* 202, 1339–1347.
- Bi, G., and Poo, M. (1998). Synaptic modification in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* 18, 10464–10472.
- Buzsaki, G. (1986). Hippocampal sharp waves: Their origin and significance. *Brain Res.* 398, 242–252.
- Buzsaki, G. (1989). Two-stage model of memory trace formation: a role for "noisy" brain states. *Neuroscience* 31, 551–570.
- Chrobak, J., and Buzsaki, G. (1996). High-frequency oscillations in the output networks of the hippocampal-entorhinal axis of the freely behaving rat. *J. Neurosci.* 16, 3056–3066.
- Csicsvari, J., Hirase, H., Mamiya, A., and Buzsaki, G. (2000). Ensemble patterns of hippocampal CA3–CA1 neurons during sharp wave associated population events. *Neuron* 28, 585–594.
- Debanne, D., Gähwiler, B., and Thompson, S. (1998). Long-term synaptic plasticity between pairs of individual ca3 pyramidal cells in rat hippocampal slice cultures. *J. Physiol.* 507, 237–247.
- Feldman, D. (2000). Timing-based LTP and LTD at vertical inputs to layer II/III pyramidal cells in rat barrel cortex. *Neuron* 27, 45–56.
- Froemke, R., and Dan, Y. (2002). Spike-timing-dependent synaptic modification induced by natural spike trains. *Nature* 416, 433–438.
- Gerstner, W., Kempter, R., van Hemmen, J., and Wagner, H. (1996). A neuronal learning rule for sub-millisecond temporal coding. *Nature* 383, 76–81.
- Hebb, D. (1949). *The Organization of Behavior: A Neuropsychological Theory* (New York: Wiley).
- Ishizuka, N., Weber, J., and Amaral, D. (1990). Organization of intrahippocampal projections originating from CA3 pyramidal cells in the rat. *J. Comp. Neurol.* 295, 580–623.
- Izhikevich, E. (2003). Simple model of spiking neurons. *IEEE Trans. Neural Netw.* 14, 1569–1572.
- Izhikevich, E. (2006). Polychronization: Computation with spikes. *Neural Comput.* 18, 245–282.
- Izhikevich, E., Gally, J., and Edelman, G. (2004). Spike-timing dynamics of neuronal groups. *Cereb. Cortex* 14, 933–944.
- Karbowsky, J., and Ermentrout, G. (2002). Synchrony arising from a balanced synaptic plasticity in a network of heterogeneous neural oscillators. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* 65, 031902.
- Kudrimoti, H., Barnes, C., and McNaughton, B. (1999). Reactivation of hippocampal cell assemblies: effects of behavioral state, experience, and EEG dynamics. *J. Neurosci.* 19, 4090–4101.
- Lee, A., and Wilson, M. (2002). Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron* 36, 1183–1194.
- Markram, H., Lubke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275, 213–215.
- Miles, R., Traub, R., and Wong, R. (1988). Spread of synchronous firing in longitudinal slices from the CA3 region of the hippocampus. *J. Neurophysiol.* 60, 1481–1496.
- Morrison, A., Aertsen, A., and Diesmann, M. (2007). Spike-timing-dependent plasticity in balanced random networks. *Neural Comput.* 19, 1437–1467.
- Nadasdy, Z., Hirase, H., Czurko, A., Csicsvari, J., and Buzsaki, G. (1999). Replay and time compression of recurring spike sequences in the hippocampus. *J. Neurosci.* 19, 9497–9507.
- Nowotny, T., Zhigulin, V., Selverston, A., Abarbanel, H., and Rabinovich, M. (2003). Enhancement of synchronization in a hybrid neural circuit by spike-timing dependent plasticity. *J. Neurosci.* 23, 9776–9785.
- O'Neil, J., Senior, T., and Csicsvari, J. (2006). Place-selective firing of CA1 pyramidal cells during sharp wave/ripple network patterns in exploratory behavior. *Neuron* 49, 143–155.
- Pinsky, P., and Rinzel, J. (1995). Synchrony measures for biological neural networks. *Biol. Cybern.* 73, 129–137.
- Roberts, P., and Bell, C. (2002). Spike timing dependent plasticity in biological systems. *Biol. Cybern.* 87, 392–403.
- Shon, A., Rao, R., and Sejnowski, T. (2004). Motion detection and prediction through spike-timing dependent plasticity. *Network* 15, 179–198.
- Siapas, A., and Wilson, M. (1998). Coordinated interactions between hippocampal ripples and cortical spindles during slow-wave sleep. *Neuron* 21, 1123–1128.

Siapas, A., Lubenov, E., and Wilson, M. (2005). Prefrontal phase-locking to hippocampal theta oscillations. *Neuron* 46, 141–151.

Skaggs, W., and McNaughton, B. (1996). Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science* 274, 1216–1217.

Skaggs, W., McNaughton, B., Wilson, M., and Barnes, C. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* 6, 149–172.

Song, S., Miller, K., and Abbott, L. (2000). Competitive Hebbian learning through spiketiming-dependent synaptic plasticity. *Nat. Neurosci.* 3, 919–926.

Squire, L. (1992). Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychol. Rev.* 99, 195–231.

Swadlow, H. (1994). Efferent neurons and suspected interneurons in motor cortex of the awake rabbit: axonal properties, sensory receptive fields, and subthreshold synaptic inputs. *J. Neurophysiol.* 71, 437–453.

Wilson, M., and McNaughton, B. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science* 265, 676–679.

Zhigulin, V., Rabinovich, M., Huerta, R., and Abarbanel, H. (2003). Robustness and enhancement of neural synchronization by activity-dependent coupling. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* 67, 021901.