

Spike-Timing-Dependent Plasticity in Balanced Random Networks

Abigail Morrison

abigail@brain.riken.jp

Computational Neuroscience Group, RIKEN Brain Science Institute, Wako City, Saitama 351-0198, Japan

Ad Aertsen

aertsen@biologie.uni-freiburg.de

Neurobiology and Biophysics, Institute of Biology III, Albert-Ludwigs-University, 79104 Freiburg, Germany

Markus Diesmann

diesmann@brain.riken.jp

Computational Neuroscience Group, RIKEN Brain Science Institute, Wako City, Saitama 351-0198, Japan

The balanced random network model attracts considerable interest because it explains the irregular spiking activity at low rates and large membrane potential fluctuations exhibited by cortical neurons *in vivo*. In this article, we investigate to what extent this model is also compatible with the experimentally observed phenomenon of spike-timing-dependent plasticity (STDP).

Confronted with the plethora of theoretical models for STDP available, we reexamine the experimental data. On this basis, we propose a novel STDP update rule, with a multiplicative dependence on the synaptic weight for depression, and a power law dependence for potentiation. We show that this rule, when implemented in large, balanced networks of realistic connectivity and sparseness, is compatible with the asynchronous irregular activity regime. The resultant equilibrium weight distribution is unimodal with fluctuating individual weight trajectories and does not exhibit development of structure. We investigate the robustness of our results with respect to the relative strength of depression.

We introduce synchronous stimulation to a group of neurons and demonstrate that the decoupling of this group from the rest of the network is so severe that it cannot effectively control the spiking of other neurons, even those with the highest convergence from this group.

1 Introduction

In order to understand cognitive processing in the cortex, it is necessary to have a solid understanding of its dynamics. The balanced random network model has been able to provide some insight into this matter, as it can account for the large fluctuations in membrane potential and the irregular firing at low rates observed *in vivo* for cortical neurons. Many aspects of this model have been investigated and are now well understood (van Vreeswijk & Sompolinsky, 1996, 1998; Brunel & Hakim, 1999; Brunel, 2000). Until now, studies have mainly focused on the activity dynamics of such networks under the assumption that the strength of a synapse remains constant. However, recent research indicates that the strength of a synapse varies with respect to the relative timing of pre- and postsynaptic spikes (Markram, Lübke, Frotscher, & Sakmann, 1997; Zhang, Tao, Holt, Harris, & Poo, 1998; Bi & Poo, 1998; Debanne, Gähwiler, & Thompson, 1998; Sjostrom, Turrigiano, & Nelson, 2001; Froemke & Dan, 2002; Wang, Gerkin, Nauen, & Bi, 2005), a phenomenon known as spike-timing-dependent plasticity (STDP). This has prompted intense theoretical interest (see, e.g., Kempter, Gerstner, & van Hemmen, 1999; Song, Miller, & Abbott, 2000; van Rossum, Bi, & Turrigiano, 2000; Kistler & van Hemmen, 2000; Rubin, Lee, & Sompolinsky, 2001; Kempter, Gerstner, & van Hemmen, 2001; Gütig, Aharonov, Rotter, & Sompolinsky, 2003; Izhikevich & Desai, 2003; Burkitt, Meffin, & Grayden, 2004; Guyonneau, VanRullen, & Thorpe, 2005), but these studies have thus far tended to focus on the development of the synaptic weights of a single neuron in the absence of the activity dynamics of a recurrent network. Such network studies as there are (Hertz & Prügel-Bennet, 1996; Levy, Horn, Meilijson, & Ruppin, 2001; Izhikevich, Gally, & Edelman, 2004; Iglesias, Eriksson, Grize, Tomassini, & Villa, 2005) have shown divergent results, such as the development of strongly connected neuronal groups in Izhikevich et al. (2004) and systematic synaptic pruning in Iglesias et al. (2005). Additionally, these studies all assume a level of connectivity orders of magnitude below that seen in cortical networks (Braitenberg & Schüz, 1998).

The issue of connectivity is significant, as STDP is sensitive to correlation in neuron activity. Reducing the number of incoming synapses per neuron and scaling up the synaptic strength to compensate for this will inevitably influence this correlation. To ensure that the development of synaptic weights observed is not an artifact of low connectivity, it is necessary to examine the behavior of high-connectivity networks. In this article, we investigate the activity dynamics and the development of the synaptic weight distribution for recurrent networks with biologically realistic levels of connectivity and sparseness. This necessitates the use of distributed computing techniques for simulation, previously developed in Morrison, Mehring, Geisel, Aertsen, & Diesmann (2005). As a starting position, we use the well-understood balanced random network model (van Vreeswijk

& Sompolinsky, 1996, 1998; Brunel & Hakim, 1999; Brunel, 2000), and extend it to incorporate STDP in its excitatory-excitatory synapses. The model of STDP implemented is based on a reexamination of the experimental data of Bi and Poo (1998), resulting in a novel power law update rule. The network is investigated for both the case that no structured external stimulus is applied and that a group of neurons is subjected to an irregular synchronous stimulus.

In section 2 we reexamine the experimental data for STDP and propose a well-constrained power law model for the synaptic weight update. A consideration of rate perturbation further constrains the model to an all-to-all, rather than nearest-neighbor, spike pairing scheme. The neuron and network models are introduced in section 3, and the activity dynamics for a static balanced random network is reviewed. The network is extended to incorporate STDP in section 4. In section 4.1, we demonstrate the mutual equilibrium of weight and activity dynamics in a plastic network with no structured external stimulus and show that no structure is developed. This finding turns out to be robust across changes to the network connectivity, spike pairing scheme, and formulation of the STDP update rule. We examine the behavior of the network if the STDP model exhibits depression that over- or undercompensates for the correlation in network activity. In section 4.2, we apply a synchronous stimulus to a group of neurons and investigate its potential to induce development of structure. We show that a strong stimulus causes the network to enter a pathological state, whereas the structure-enhancing effects of a weaker stimulus are counteracted by the decoupling of the stimulated group from the rest of the network. These results are discussed in section 5, as are ways in which the investigated models might be adapted to enable stimulus-driven development of structure. An algorithmic template for implementing STDP in network simulations that is compatible with distributed computing is given in the appendix.

2 A Novel Power Law Model of STDP

There are many theoretical models of STDP in circulation. The two main sources of disparity are the weight dependence of the weight changes—for example, additive (Song et al., 2000), multiplicative (Rubin et al., 2001), or somewhere in between (Gütig et al., 2003)—and which pairs of spikes contribute to plasticity—for example, all-to-all (Gerstner, Kempter, van Hemmen, & Wagner, 1996) and nearest neighbor (Izhikevich et al., 2004). Other variants include some degree of activity dependence, for example suppression (Froemke & Dan, 2002) or postsynaptic activity-dependent homeostasis (van Rossum et al., 2000). It has already been demonstrated that variations in the formulation of the STDP update rules can lead to qualitatively different results (Rubin et al., 2001; Gütig et al., 2003; Izhikevich & Desai, 2003; Burkitt et al., 2004), but there is as yet no consensus. Therefore, finding a formulation of STDP describing the experimental data as well as possible




Figure 1: Weight and time dependency of change for the Bi and Poo (1998) protocol. (A) Absolute change in excitatory postsynaptic current (EPSC) amplitude as a function of initial EPSC amplitude in log-log representation for potentiation where the presynaptic spike precedes the postsynaptic spike by 2.3 to 8.3 ms (plus signs) and depression where the postsynaptic spike preceded the presynaptic spike by 3.4 to 23.6 ms (circles). The lower black line is a linear fit to the depression data: slope = -1 , offset = 2 . The upper black line is a linear fit to the potentiation data assuming a slope = 0.4 , resulting in an offset = 1.7 . The pale gray line shows an additive STDP prediction for the data, and the dark gray curve shows a multiplicative STDP prediction, with $w_{\max} = 3000$ pA. (B) Percentage change in EPSC amplitude as a function of spike timing. Curves show power law prediction for the data with $\mu = 0.4$, $\tau = 20$ ms, $\lambda = 0.1$, and $\alpha = 0.11$ for different initial EPSC amplitudes: 17 pA, pale gray curve; 50 pA, medium gray curve; 100 pA, dark gray curve. All data from Bi and Poo (1998).

is still a relevant issue. Here we neglect the activity dependence and investigate how the weight dependence of potentiation and depression can best be characterized and what is an appropriate spike pairing scheme to apply.

2.1 Weight Dependence of STDP. Figure 1A shows the dependence of synaptic weight change on initial synapse strength for the experimental data from Bi and Poo (1998). The difference between this plot and the classical plot (Figure 5 in Bi and Poo) is that here, the absolute, rather than relative, change in the synaptic strength is plotted and a double logarithmic representation is used. The exponent of the dependence can be obtained from the slope of a linear fit to the data. In the case of depression, neglecting the largest outlier results in an exponent of -1 . We will therefore assume a purely multiplicative update rule: $\Delta w_- \propto w$. In the case of potentiation, depending on the treatment of the largest outlier, an exponent in the range 0.3 to 0.5 results. This suggests a power law rule of the form $\Delta w_+ \propto w^\mu$, for which we shall assume $\mu = 0.4$. This novel rule fits the data better than either additive ($\Delta w_+ = c$) or multiplicative ($\Delta w_+ \propto w_{\max} - w$) rules, shown

in Figure 1A for comparison. Although all the rules can be fitted to give similar results for synapses in the 30 to 100 pA range, the greatest difference in prediction is for the behavior of very weak synapses. A multiplicative rule predicts a greater absolute potentiation for synapses close to 0 pA than that obtained for synapses in the 30 to 100 pA range, and an additive rule predicts the same absolute potentiation regardless of synaptic strength. In contrast to these rules, the power law rule predicts very small absolute increases when very weak synapses are potentiated.

2.2 Constraining the Model Parameters. Having established the weight dependence for the change in synaptic strength over the entire 60-pair protocol, we now constrain the model parameters. The update rule underlying the synaptic modifications over the course of the protocol has a weight and a time component; specifically, we assume it takes the form

$$\begin{aligned}\Delta w_+ &= \lambda_{60} w_0^{1-\mu} w^\mu e^{-\frac{|\Delta t|}{\tau}} & \text{if } \Delta t > 0 \\ \Delta w_- &= -\lambda_{60} \alpha w e^{-\frac{|\Delta t|}{\tau}} & \text{if } \Delta t < 0,\end{aligned}\tag{2.1}$$

where τ is the time constant, λ_{60} is the learning rate over the whole protocol, w_0 is a reference weight, and α scales the strength of depressing increments with respect to the strength of potentiating increments. We take Δt to be the difference between the postsynaptic and presynaptic spikes arriving at the synapse, that is, after the delays due to axonal- and backpropagation, as suggested by Debanne et al. (1998). If the spikes arrive exactly synchronously, $\Delta t = 0$, no alteration is made to the weight.

We start by fitting the potentiation data from Figure 1A to the theoretical form of equation 2.1. We have:

$$\log(\Delta w_+) = \mu \log(w) + (1 - \mu) \log(w_0) + \log\left(\lambda_{60} e^{-\frac{|\Delta t|}{\tau}}\right).\tag{2.2}$$

Without loss of generality, we set $w_0 = 1$ pA, and substitute for Δt the mean time interval $\overline{\Delta t} = 6.3$ ms of the potentiation data. The value of $\log(\lambda_{60} e^{-\frac{\overline{\Delta t}}{\tau}})$ is thus the offset of the linear fit to the potentiation data. Assuming $\tau = 20$ ms in line with other theoretical work on STDP (e.g., Song et al., 2000; van Rossum et al., 2000), this results in the value $\lambda_{60} = 7.5$. However, to implement STDP in a simulation, we need to know the value of the learning rate for just one pair, λ instead of the learning rate over 60 pairs, λ_{60} . This can be done numerically: by substituting $\Delta w_+ = w$ in equation 2.2 and solving for w , the value of the weight that would experience 100% potentiation in this protocol can be calculated as $w_{(0)} = w_0 e^{\frac{\log(\lambda_{60}) - \overline{\Delta t}/\tau}{1-\mu}} = 17$ pA. We start at

this initial weight, and apply the single pair update rule 60 times:

$$\begin{aligned}
 w_{(1)} &= w_{(0)} + \lambda w_0^{1-\mu} w_{(0)}^\mu e^{-\frac{|\Delta t|}{\tau}} \\
 w_{(2)} &= w_{(1)} + \lambda w_0^{1-\mu} w_{(1)}^\mu e^{-\frac{|\Delta t|}{\tau}} \\
 &\vdots \\
 w_{(60)} &= w_{(59)} + \lambda w_0^{1-\mu} w_{(59)}^\mu e^{-\frac{|\Delta t|}{\tau}}.
 \end{aligned}$$

A value of $\lambda = 0.1$ results in $w_{(60)} = 2w_{(0)} = 34$ pA, as required.

As the range of time intervals used to produce the depression data is too large to have much confidence in the offset of the linear fit, we determined α in an analog fashion to give 40% depression for $\overline{\Delta t} = 6.3$ ms. For the rest of this article, unless otherwise stated, we will use the following update rule for a pair of spikes:

$$\begin{aligned}
 \Delta w_+ &= \lambda w_0^{1-\mu} w^\mu e^{-\frac{|\Delta t|}{\tau}} & \text{if } \Delta t > 0 \\
 \Delta w_- &= -\lambda \alpha w e^{-\frac{|\Delta t|}{\tau}} & \text{if } \Delta t < 0,
 \end{aligned} \tag{2.3}$$

with $\mu = 0.4$, $\tau = 20$ ms, $w_0 = 1$ pA, $\lambda = 0.1$, and $\alpha = 0.11$. The fit of this update rule to the time-dependence data of Bi and Poo (1998) is depicted in Figure 1B for three different initial EPSC amplitudes. This produces a similar window function to that reported in studies on cortical rather than hippocampal plasticity (e.g., Froemke & Dan, 2002), where a maximum potentiation of 103% and a maximum depression of 51% was reported. Although the above calculation assumes that the time constants for potentiation and depression are equal, $\tau_+ = \tau_- = \tau$, it is trivial to adjust the parameters to accommodate $\tau_+ < \tau_-$, as reported in several experimental studies (e.g., Feldman, 2000; Froemke & Dan, 2002). An appropriate efficient algorithmic implementation of STDP suitable for distributed computing is given in the appendix.

2.3 Spike Pairing Scheme. A self-consistent rate is a necessary condition for stability in the balanced random network model. It is known that different spike pairing schemes result in qualitatively different weight dynamics as a function of postsynaptic rate (Kempster et al., 2001; Izhikevich & Desai, 2003; Burkitt et al., 2004). Therefore, to establish which pairs of spikes should be considered for the synaptic weight update, we consider the effects of different spike pairing schemes when a self-consistent rate is perturbed. We implemented a nearest-neighbor scheme, whereby a presynaptic spike is paired with only the last postsynaptic spike to effect depression, and a postsynaptic spike is paired with only the last presynaptic spike to effect

potentiation. Conversely, in an all-to-all scheme, a presynaptic spike is paired with all previous postsynaptic spikes to effect depression, and a postsynaptic spike is paired with all previous presynaptic spikes to effect potentiation.

For this investigation, 1000 neurons were provided with input designed to match the network model described in section 3, namely, 9000 independent excitatory and 2250 independent inhibitory Poisson processes at 7.7 Hz to model input from the local network, and a further 9000 independent excitatory Poisson processes at 2.32 Hz to model external excitation. (For all synaptic and neuronal parameters, refer to section 3). Of the excitatory inputs modeling local network input, 1000 of them were mediated by synapses implementing the power law STDP update rules described above; the rest were static. The 1×10^6 plastic synapses were initialized from a gaussian distribution with a mean of 45.61 pA and a standard deviation of 4.0 pA. In the case of the all-to-all pairing scheme, we set $\alpha = 0.1021$ and $\lambda = 0.0973$, and in the case of the nearest-neighbor scheme, $\alpha = 0.0976$ and $\lambda = 0.116$. This choice of parameters results in a firing rate of 7.7 Hz and a stable and unimodal synaptic weight distribution measured across all the plastic synapses with a mean of 45.5 pA and a standard deviation of 4.0 pA. A unimodal distribution was expected for the power law formulation of STDP, as bimodal distributions have been reported only for additive or near-additive rules (see Gütig et al., 2003), with very weak or no weight dependence, whereas the power law formulation has a strong dependence on the initial synaptic strength.

To perturb the self-consistent rate, a current was injected into the neurons from 50 s onward. A current of 30 pA increased the firing rate by approximately 2 Hz; a current of -35 pA decreased it by the same amount. For all configurations of spike pairing scheme and injected current, the unimodal nature of the distribution was conserved. As can be seen in Figure 2, the choice of pairing scheme has a significant effect on the consequent weight development. If the postsynaptic rate is increased, an all-to-all pairing scheme results in a net decrease of synaptic weights (see Figure 2B), whereas a nearest-neighbor scheme results in a net increase (see Figure 2A). Conversely, if the postsynaptic rate is decreased with respect to the input rate, an all-to-all scheme increases the mean synaptic weight (see Figure 2B), whereas a nearest-neighbor scheme decreases it (see Figure 2A). In other words, the all-to-all scheme acts like a restoring force, counteracting the rate disparity, whereas the nearest-neighbor scheme acts in such a direction as to magnify any rate disparity. These results do not depend on the time constants for depression and potentiation being equal; results obtained for $\tau_- = 34$ ms and $\tau_+ = 14$ ms as reported by Froemke and Dan (2002), were qualitatively the same (all-to-all scheme: $\alpha = 0.042$, $\lambda = 0.0973$; nearest-neighbor scheme: $\alpha = 0.0458$, $\lambda = 0.116$; data not shown).

In the absence of some sliding threshold mechanism (e.g., Bienenstock, Cooper, & Munro, 1982) or postsynaptic homeostasis (see Turrigiano, Leslie,




Figure 2: Development of mean synaptic weight for (A) nearest-neighbor and (B) all-to-all spike pairing schemes, averaged over 1×10^6 synapses. From 50 s onward, a current is injected into the postsynaptic neurons to increase their rate (upward triangles) or decrease it (downward triangles) by 2 Hz. Dashed curves show mean weight development when no current is injected. See the text for further details of the simulation.

Desai, Rutherford, & Nelson, 1998), an all-to-all spike pairing scheme is clearly better suited to situations where a self-consistent rate is required and shall be used for the remainder of this letter.

3 Network Model

We consider networks of current-based integrate-and-fire neurons. For each neuron, the dynamics of the membrane potential V is

$$\dot{V} = -\frac{V}{\tau_m} + \frac{1}{C}I,$$

where τ_m is the membrane time constant, C is the capacitance of the membrane, and I is the input current to the neuron. The current arises as a superposition of the synaptic currents and any external current. The synaptic current I_s due to one incoming spike is represented as an α -function:

$$I_s(t) = w \frac{e}{\tau_\alpha} t e^{-t/\tau_\alpha},$$

where w is the peak value of the current and τ_α is the rise time. When the membrane potential reaches a given threshold value Θ , the membrane potential is clamped to zero for an absolute refractory period τ_r . The values for these parameters used in this article, unless otherwise stated, are:

τ_m : 10 ms
 C : 250 pF

Θ : 20 mV
 τ_r : 0.5 ms
 w : 45.61 pA
 τ_w : 0.33 ms

The parameters are chosen such that the generated postsynaptic potentials (PSPs) have a rise time of 1.7 mV and a half-width of 8.5 ms (e.g., Fetz, Toyama, & Smith, 1991). To avoid transients in the dynamics at the beginning of the simulation, the initial membrane potentials are drawn from a gaussian distribution with a mean of 5.7 mV and a standard deviation of 7.2 mV.

The network model was derived from Brunel (2000), where the larger number of excitatory neurons is balanced by increasing the strength of the inhibitory connections. In order to realize biologically realistic values for both the number of connections per neuron ($\simeq 10^4$) and the connection probability ($\simeq 0.1$) in the same network, it is necessary to consider networks with of the order of 10^5 neurons, approximately the number of neurons in a cubic millimeter of cortex (Braitenberg & Schüz, 1998). Specifically, we simulated networks with 90,000 excitatory and 22,500 inhibitory neurons. In the static network, each neuron receives 9000 connections from excitatory and 2250 connections from inhibitory neurons, whereby the peak current of an inhibitory synapse is larger than that of an excitatory synapse by a factor of -5 . A neuron is not permitted to have a connection to itself. The total number of synapses in the network is 1.27×10^9 . Additionally, each neuron receives an independent Poissonian stimulus corresponding to 9000 excitatory spike trains, each at 2.32 Hz. The strength of these external excitatory connections is the same as the local excitatory connections. The propagation delay between neurons is 1.5 ms, and the networks are simulated with a computational resolution of 0.1 ms.

The activity of the network described above, if all synapses are static, is in the asynchronous irregular (AI) regime (Brunel, 2000), whereby each neuron fires irregularly at a low rate (coefficient of variation for the interspike interval $CV_{ISI} = 0.91$, rate $\nu = 7.8$ Hz), and the network activity is not characterized by synchronous events. However, even in the AI regime, the network is not devoid of synchronous activity. In general, both slow and fast oscillations can be excited (Brunel & Hakim, 1999; Brunel, 2000; Tetzlaff, Morrison, Timme, & Diesmann, 2005). The slow oscillations are eradicated by the use of a noisy external stimulus and drawing the initial membrane potentials from a distribution. The fast oscillations are somewhat more tenacious and are clearly visible in the cross-correlogram in Figure 3. They cause a high variability in the spike count, which can be quantified by the Fano factor, calculated by binning the spike times, and dividing the variance of the spike count by the mean. In this article, all Fano factors are calculated by binning the spike trains from 1000 neurons (unless otherwise stated), using a bin size of 3 ms, which resulted in a mean spike count of at least 10 spikes per bin for all the data sets investigated.




Figure 3: Cross-correlogram of the spiking activity in the static balanced random network, averaged over 500 pairs of neurons and 50 s, using a bin size of 0.1 ms. This is calculated as the cross-covariance normalized by the square root of the product of the variances. As can be seen in the inset, the cross-correlogram is shifted to the right by the synaptic propagation delay of 1.5 ms, as it is recorded from the perspective of the synapse, and the delay is entirely dendritic (see the appendix).

The above network has a Fano factor for the spike count of $F_{SC} = 7.6$, significantly higher than the value of 1, which is a characteristic of Poissonian statistics.

In the case of an idealized network with pure Poissonian firing statistics, where the excitatory synapses exhibit STDP as described in section 2, the fix point of the resulting synaptic weight distribution can be easily determined: $w^* = w_0 \alpha_p^{\frac{1}{\mu-1}}$. Setting $w^* = 45.61$ pA in line with the static network described above, this determines $\alpha_p = 0.1$. However, as will be demonstrated in section 4, to retain the activity dynamics of the static network, a slightly larger $\alpha > \alpha_p$ is required to compensate for the structured raw cross-correlation induced by the fast oscillations.

4 Dynamics of Network Activity and Weight Distribution

4.1 Unstimulated Network. Replacing the excitatory-excitatory connections in the static network with synapses following the power law




Figure 4: Synaptic weight distribution and correlation in the plastic balanced recurrent network with $\alpha = 1.057\alpha_p$. (A) The percentage area difference of the weight distribution from the final distribution (at 2000 s) as a function of time. The data were recorded from the outgoing excitatory-excitatory synapses of 900 neurons: 8.1×10^6 synapses. (B) Histogram of the equilibrium synaptic weight distribution, averaged over 10 samples in the final 500 s. The black curve indicates the gaussian distribution, with the mean and standard deviation obtained from the histogram ($\mu_w = 45.65$, $\sigma_w = 3.99$). (C) Cross-correlogram of the spiking activity, averaged over 500 pairs of neurons and 50 s, using a bin size of 0.1 ms. The recording was made after the weight distribution had stabilized (400 to 450 s). (D) Difference between the cross-correlograms of the plastic network and a static network using synaptic strengths drawn from the equilibrium distribution shown in (B). Cross-correlograms calculated as in Figure 3.

STDP update rules given by equation 2.3, a stable network configuration is achieved for $\alpha = 1.057\alpha_p$. The resultant weight distribution settles within 200 s to approximately gaussian ($\mu_w = 45.65$, $\sigma_w = 3.99$; see Figures 4A and 4B).

The activity of the plastic network is close to the activity of the static network described in section 3: a slightly higher rate of 8.8 Hz is observed. The other spike statistics are also comparable ($CV_{ISI} = 0.88$, $F_{SC} = 8.5$). Some

difference is to be expected, as in the static network all excitatory-excitatory synapses have the same weight, whereas in the plastic network, the distribution of weights contributes to the fluctuations of the membrane potential and so increases the rate. Indeed, if the weights in the static network are chosen from a gaussian distribution with the same mean and standard deviation as the equilibrium distribution, the activity statistics are very similar ($\nu = 8.9$ Hz, $CV_{ISI} = 0.9$, and $F_{SC} = 8.5$). For the rest of this letter, the static network using this distribution of synaptic weights will be considered a control case.

The similarity of the dynamics is also demonstrated by the cross-correlogram in Figure 4C. The cross-correlogram has qualitatively the same shape as that of the static network, and the difference of the two (see Figure 4D) shows no structure whatsoever. Therefore, we can conclude that the two networks exhibit the same activity dynamics and that the power law STDP update rule described in section 2 is compatible with balanced recurrent networks in the AI regime. However, we do not claim that this property is unique to this STDP model. It would seem likely that any STDP formulation that results in a unimodal equilibrium distribution of synaptic weights would also be compatible, at least for some parameter range. This is confirmed for a formulation of the Gütig et al. (2003) update rules in the next section.

4.1.1 Survival of Strong Synapses. One of the main sources of interest in STDP is its apparent ability to create neuronal assemblies (e.g., Izhikevich et al., 2004). However, searching for a stable structure in a network with 8.1×10^8 excitatory-excitatory synapses presents a significant data analysis problem in itself. To discover whether such assemblies are self-organizing in the balanced recurrent network, we consider the persistence of the strongest synapses in the network. If structure were developing, strong synapses should stay strong or become even stronger. During the course of the simulation, the outgoing plastic synapses of 900 neurons were monitored every 5 s, and those that were above a threshold of $\mu + 1.3\sigma$ (50.8 pA) of the equilibrium distribution shown in Figure 4, were recorded (approximately the strongest 10%). The persistence of strong synapses for various different initial times is shown in Figure 5A. For a group of synapses that are strong at a given time, the number that remain strong decays exponentially with a time constant of approximately 60 s, and almost none of an initial group remains strong after 800 s. No synapses remained strong for longer than 1000 s, although the simulation ran for 2000 s. This behavior was not systematically dependent on the time at which the initial population was defined, as is demonstrated by the similar courses of the curves in Figure 5A. These results suggest that structure is not developing in the balanced recurrent network.

The survival statistics are qualitatively different for a low-connectivity network, as can be seen in Figure 5B. Here, the number of neurons in the




Figure 5: Persistence of strong synapses. (A) High-connectivity network. From a total population of 8.1×10^6 synapses (see Figure 4) and for a given initial time, approximately the 10% strongest synapses are recorded. The number of these synapses that remain strong is plotted as a function of survival time. The shading of the lines indicates the initial times: 100 s (palest gray curve) in steps of 200 s until 1500 s (black curve). (B) Low-connectivity network. As in A, but with a total population of 8.1×10^4 synapses and an earliest initial time of 500 s (palest gray curve).

network was reduced to 900 excitatory and 225 inhibitory neurons while retaining a connection probability of 0.1. To compensate for the lower number of synaptic inputs, the rate of the external input was increased to 34.66 Hz, the peak current of the static excitatory synapses was increased by a factor of 4 to 182.44 pA, and the peak current of the inhibitory synapses was a factor of -18 larger than the excitatory synapses. In the case of the plastic excitatory synapses, the scaling factor of 4 was applied postsynaptically, so for the purposes of the STDP update, the synaptic weights stayed in the range of approximately 30 to 70 pA. This permitted values of $\alpha = 1.109\alpha_p$ and $\lambda = 0.1$ to be used, resulting in an STDP window function as close as possible to that used for the full-sized network, and much the same activity statistics ($\nu = 7.9$ Hz, $CV_{ISI} = 0.91$, $F_{SC} = 8.6$ (900 neurons)). The equilibrium weight distribution was unimodal, with a mean of 45.52 pA and a standard deviation of 6.52 pA.

In this case, strong synapses are much more stable, with some remaining strong for thousands of seconds. The decay of the number of strong synapses does not vary systematically with the time at which the initial population was defined, as was the case for the full-sized network. However, the decay is no longer well described by an exponential function, but by a power law with an exponent of -1.02 . Despite the longer endurance of strong synapses, it still could not be said that structure is spontaneously

developing in the network, as although the number of strong synapses did not quite decay to zero over the 5000 s measuring period, the remainder comprises less than 0.2% of the total number of plastic synapses. Similar results are obtained when using the nearest-neighbor spike pairing scheme described in section 2.3, resulting in a unimodal distribution with a mean of 45.6 pA, a standard deviation of 6.62 pA, and power law survival statistics with an exponent of -0.93 (data not shown). The results are also robust with respect to the formulation of the STDP update rules: power law STDP with asymmetric time constants ($\tau_- = 34$ ms, $\tau_+ = 14$ ms, $\alpha = 0.048$) and Güting et al. (2003) STDP ($w_{\max} = 100$ pA, $\lambda = 0.005$, $\mu = 0.4$, and $\alpha = 1.188$) produce unimodal distributions, with means of 45.7 pA and 45.57 pA, standard deviations of 7.13 pA and 5.17 pA, respectively, and power law survival statistics with exponents of -0.62 and -1.3 , respectively (data not shown).

These results show that the connectivity of a network has a significant effect on the stability of individual weights, as strong synapses in the low-connectivity networks persist for much longer periods than in the network with biologically realistic connectivity. Further, they demonstrate that the nondevelopment of structure is quite a robust observation. It is not restricted to high-connectivity networks and does not depend critically on either our specific formulation of the STDP update rules or the spike pairing scheme.

4.1.2 Sensitivity to Scaling of Depression. The value of $\alpha = 1.057\alpha_p$ was determined through a tuning process to give a mean synaptic weight close to the weight of excitatory synapses in the static network, and thus a similar self-consistent firing rate. This raises the question of what the effect on the network activity dynamics would be if α is chosen too high, leading to a net depression of the synapses, or too low, leading to a net potentiation. Clearly, any net change of the synaptic weights will lead to a change in the correlation structure, which will affect the weights, and so on. A network can be stable only if its weight distribution and correlation structure are compatible, as in the network described above. How easy is this to accomplish?

We found that if α is chosen 2% higher— $\alpha = 1.078\alpha_p$ —the network does indeed settle to a new equilibrium, with a near-gaussian weight distribution at a slightly lower mean ($\mu_w = 45.33$ pA) and a slightly increased standard deviation ($\sigma_w = 4.1$ pA). The network firing rate is, of course, lower, $\nu = 3.4$ Hz, and the activity is still in the AI regime ($CV_{ISI} = 0.92$, $F_{SC} = 3.9$). Interestingly, the network is in a more asynchronous state than either the heterogeneous static network or the plastic network with $\alpha = 1.057\alpha_p$ ($F_{SC} = 8.5$), as reflected in its cross-correlogram (see Figure 6). The central and side peaks are smaller in amplitude and less well defined, and the difference in the cross-correlograms of this network and the static network (see Figure 6, inset) shows clear structure around $\tau = 0$. This lower degree of synchrony is presumably due to the increased dominance of inhibition in the local recurrent network as the excitatory weights decrease.




Figure 6: Cross-correlogram of the spiking activity in the plastic balanced random network with $\alpha = 1.078 \alpha_p$, averaged over 500 pairs of neurons and 50 s, using a bin size of 0.1 ms. The recording was made after the weight distribution had stabilized (950–1000 s). The inset shows the difference in the cross-correlograms for the plastic and the heterogeneous static networks. Cross-correlograms calculated as in Figure 3.

If α is chosen 2% lower— $\alpha = 1.035 \alpha_p$ —a very different picture emerges. Up to 24 s, the mean of the synaptic weight distribution slowly increases to 45.83 pA, and its standard deviation decreases to 3.85 pA, retaining its near-gaussian form. Over this period, the network firing rate increases smoothly from 8 Hz to about 27 Hz. Between 24 s and 26 s, there is a sudden transition. The network firing rate increases rapidly to 200 Hz (see Figure 7, inset), and the synaptic weight distribution splinters into clusters (see Figure 7, main panel).

Simultaneously, the network activity dynamics undergoes a transition. The global activity changes from asynchronous irregular firing (see Figure 8A) to strong peaks of activity interspersed with periods of silence (see Figure 8B). Within each peak, a highly regular pattern of activity builds up within a few milliseconds and then decays abruptly, to be replaced with a different pattern of activity (see Figure 8, lower panels).

These results show that α is a critical parameter for the configuration of network and plasticity model investigated here and that the value of $\alpha = 1.057 \alpha_p$ is quite close to the bifurcation point. For values of α greater than this, the network is stable, albeit with weaker synaptic weights and consequently lower firing rates. Note that even extremely weak weights will




Figure 7: Weight distribution and rate development in the plastic balanced recurrent network with $\alpha = 1.035 \alpha_p$. The gray histogram shows the distribution of weights for the outgoing excitatory-excitatory synapses of 1000 neurons (i.e., 9×10^6 synapses) after 24 s. The black histogram shows the weight distribution for the same group of synapses after 30 s. The inset shows the average rate of these neurons as a function of time.

not extinguish the network activity due to the external input. For values of α lower than the bifurcation point, the network is unstable and leaves the asynchronous irregular regime.

4.2 Stimulated Network. In section 4.1.1 it was demonstrated that if an appropriate scaling of the plasticity window is chosen such that a self-consistent rate ensues, no structure develops spontaneously. In this section, we investigate whether structure can be induced in the network by generating systematic positive correlation in the network. The most obvious way to accomplish this is to induce a group of neurons to fire synchronously. The greater the number of inputs a neuron has from this group, the more likely it is to fire shortly after the synchronous event.

4.2.1 Stimulus Protocol. A stimulus is applied to a group of neurons ($N = 500$) at random intervals. At each stimulation event, a rectangular pulse current is injected into each neuron, whereby the injection time for each neuron is drawn from a gaussian distribution ($\sigma = 0.5$ ms). In this way, there is no systematic relationship between the firing times of the stimulated group during a stimulus event beyond that given by the gaussian. Due to the random connectivity of the network, the number of connections K_{synch}

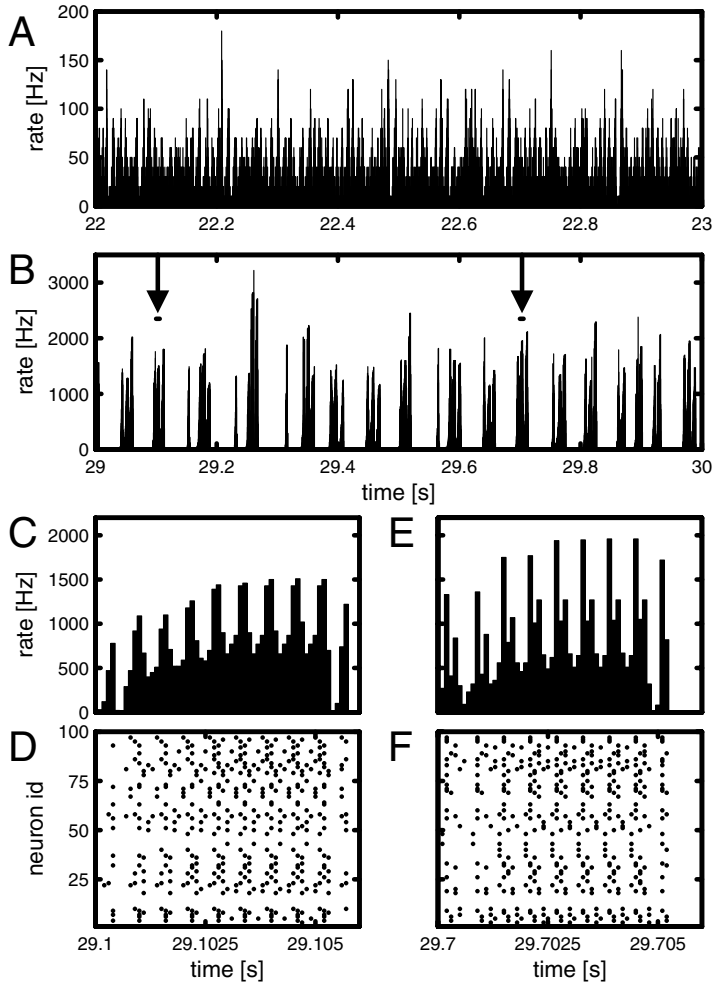


Figure 8: Activity dynamics in the plastic balanced recurrent network with $\alpha = 1.035\alpha_p$. (A) Instantaneous rate as a function of time between 22 s and 23 s for a sample of 1000 neurons with a bin size of 0.1 ms. (B) As in A, but between 29 s and 30 s. Arrows indicate time periods treated in lower panels. (C) Instantaneous rate as a function of time between 29.1 s and 29.106 s. (D) Corresponding raster plot for 100 neurons. (E) As in C, but for the time period 29.7 s to 29.706 s. (F) Corresponding raster plot; same neurons as in D.

each neuron receives from the stimulated group is binomially distributed with a mean of 50. The higher the convergence a neuron has from this group, the more likely it is to be affected by the stimulation. We therefore define a high-convergence group with respect to the stimulated group, such




Figure 9: Effect of stimulation on the network. Raster plot for 100 arbitrarily indexed neurons from the following populations: (A) stimulated group ($N = 500$), (B) high-convergence group ($N = 568$, $K_{\text{synch}} \geq 69$), (C) random-convergence group ($N = 1000$). The stimulus consists of a rectangular pulse current of amplitude 6840 pA and duration 0.1 ms at irregular intervals with a frequency of 3 Hz. For each stimulation event, the injection times for each neuron are drawn from a gaussian distribution with a standard deviation of 0.5 ms.

that $K_{\text{synch}} \geq 69$. This value was chosen as it results in a high-convergence group of approximately the same size as the stimulated group ($N = 568$). The spiking activity for the stimulated group (see Figure 9A), the high-convergence group (see Figure 9B), and a group of neurons ($N = 1000$) with random convergences (see Figure 9C) shows that although the stimulus has a strong effect within the stimulated group, its initial effect on the high-convergence group is only moderate and its effect on the random-convergence group weaker still.

4.2.2 Activity Development of Stimulus-Driven Network. Initially, it seems as if the stimulus has a catastrophic effect on the stability of the network. In Figure 10A, a sudden rate transition can be seen at 224 s. All three recorded neuron populations reach rates of over 200 Hz. A closer inspection of the spiking activity at the transition point (see Figures 10C and 10D) yields some insight. A strong burst in the stimulated group triggers an oscillation

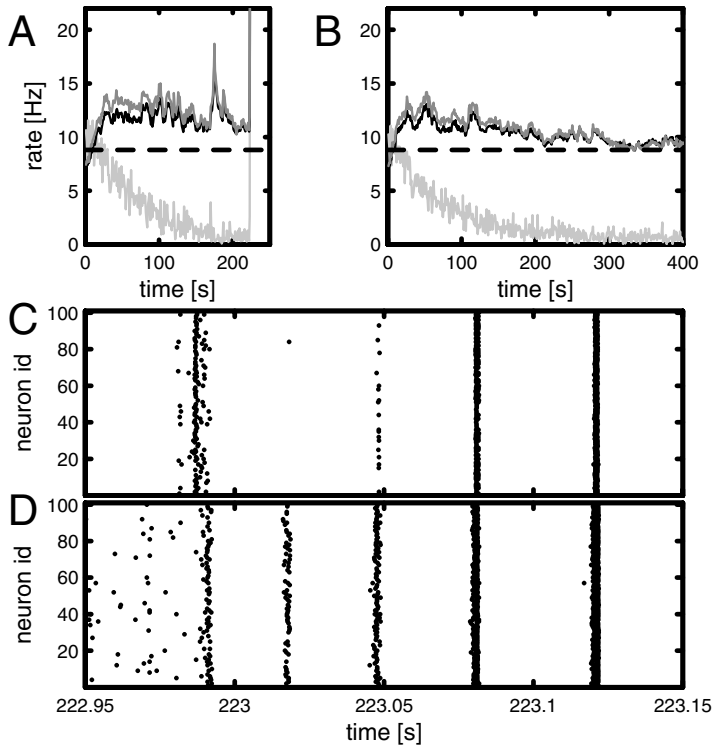


Figure 10: Activity development in the stimulated plastic network: (A) Average rate as a function of time (bin size 1 s) for stimulated group, pale gray curve; high-convergence group, dark gray curve; random-convergence group, solid black curve. The dashed black line indicates the equilibrium rate for the unstimulated plastic network. (B) As in A, but with no connections permitted within the stimulated group. (C) Raster plot for stimulated group during activity state transition, arbitrary neuron indices. (D) Raster plot for random-convergence group during activity state transition, arbitrary neuron indices.

in the random-convergence group (which is representative for the activity in the whole network). This synfire explosion was first reported by Mehring, Hehl, Kubo, Diesmann, & Aertsen (2003). Here, however, the effect is magnified rather than dying away, as the network correlation and the synaptic strengths have a positive feedback relationship through the STDP rule. The emergent activity is characterized by bursts of strongly patterned activity separated by periods of no activity, as was also observed for the unstimulated network with $\alpha = 1.035\alpha_p$ (see Figure 8).

Quite apart from its consequences for the network activity, the triggering of the synfire explosion is in itself an interesting effect, given that the

stimulated group initially had only a weak effect on the network (compare Figures 9A and 9C). This is partly due to the implementation of synaptic delays. As neurons were connected randomly, the neurons in the stimulated group also receive input from other neurons within that group. Since the synaptic delay was implemented as dendritic rather than axonal, when two neurons fire simultaneously, the presynaptic spike arrives at the synapse before the postsynaptic spike. This causes a systematic increase in the weights of the intragroup synapses when a synchronous stimulus is repeatedly applied. Although in general the response to the stimulus within the group is weakened, as will be discussed below, this increase of weights can occasionally cause an amplification of the response. In this case, more of its neurons fire within a few milliseconds, which naturally has a stronger effect on the rest of the network. If the network is connected such that there are no connections between neurons belonging to the stimulated group, the amplification does not take place, and no synfire explosion is triggered for a stimulated group of this size (see Figure 10B). However, increasing the group size triggers a synfire explosion without the amplifying effects of the intragroup synapses: for a group size of 600 stimulated neurons, this occurs after 249 s; for a group size of 700 after 127 s; and for a group size of 800 after only 7 s (data not shown).

Irrespective of whether a synfire explosion was triggered, a notable feature of the network behavior is the significant reduction in the rate of the stimulated group. This change in rate can be explained with reference to the changes in the synaptic weight distributions (see Figure 11A). The positive correlation between the stimulated group and the high-convergence and random-convergence groups corresponds to a systematic negative correlation from the perspective of the incoming synapses of the stimulated group. As a consequence, the means of the distributions of synaptic weights from the random-convergence and high-convergence groups to the stimulated group decrease from the mean weight determined in section 4.1 by 13.3% and 21.1%, respectively. This represents a significant reduction in the input received by the neurons of the stimulated group and causes them to fire much less frequently. Note that this effect can be generalized to any stimulus that creates a systematic positive correlation between a stimulated group and another group of neurons receiving input from it.

This decoupling of the stimulated group from the rest of the network counteracts the development of structure observed in the development of the weights of the outgoing synapses of the simulated group. The means of the distributions from the stimulated group to the random-convergence and high-convergence groups increase from the mean weight value determined in section 4.1 by 13.8% and 25.5%, respectively. Over the same period, the mean of the synaptic weights unconnected with the stimulated group increased by just 0.2%. Moreover, a clear relationship can be seen between the degree of convergence a given neuron has from the stimulated group and the mean weight of those synapses (see Figure 11B). Up to about the




Figure 11: Development of synaptic weights in the stimulated plastic network, with no connections within the stimulated group. (A) Distribution of synaptic weights between different neuron populations after 400 s: stimulated group to high-convergence group, solid black curve; stimulated group to random-convergence group, dashed black curve; high-convergence group to stimulated group, solid dark gray curve; random-convergence group to stimulated group, dashed dark gray curve. The pale gray curve indicates the distribution of synaptic weights for synapses unconnected to the stimulated group. (B) Average synaptic weight after 400 s as a function of convergence from the stimulated group (dots) and the high-convergence group (crosses). The vertical black line indicates the cut-off point for membership in the high-convergence group ($K_{\text{synch}} \geq 69$), and the gray line is a linear fit to the data below this point.

convergence chosen as the cut-off point for membership in the high-convergence group, the dependence of the average synaptic weight on the degree of convergence is well described by a linear relationship. Interestingly, beyond this point, the average synaptic weight increases much faster than the linear prediction. Although this suggests that a development of functional structure is occurring, such that the volley of near-synchronous spikes is reliably transferred from the stimulated group to the high-convergence group, the development of the weights of the outgoing synapses of the high-convergence group does not support this. Even after 400 s of repeated stimulation, no increase of the average synaptic weight as a function of the convergence from the high-convergence group can be observed. Clearly, the high-convergence group does not echo the stimulus sufficiently reliably as to have a corresponding effect on its own downstream high-convergence group, even though the synaptic weights between the stimulated group and the high-convergence group have increased so significantly.

The failure of the stimulated group to drive the high-convergence group effectively, despite the increased synaptic weights, can also be seen in the cross-correlogram of their spiking activity in Figure 12. Ideally, we would




Figure 12: Expected and actual correlation development. (A) Initial correlation: cross-correlation in the static network between the stimulated group and an external group with the same convergence with respect to the stimulated group and the rest of the network as the high-convergence group; synaptic weights drawn from distribution shown in Figure 4B. Averaged over 500 pairs of neurons and 50 s, using a bin size of 0.1 ms. (B) Expected final correlation: as in A, but synaptic weights for the external group are drawn from the appropriate distributions shown in Figure 11A. (C) Actual final correlation: cross-correlation between the stimulated group and the high-convergence group in the plastic network after stabilization of rates. Averaged over 500 pairs of neurons between 300 to 400 s with a bin size of 0.1 ms. Cross-correlograms calculated as in Figure 3.

like to compare the cross-correlation between the stimulated group and the high-convergence group at the beginning and end of the simulation. However, the rates are changing too quickly at the beginning for a cross-correlation to be meaningful. In order to gain some insight into the initial cross-correlation between the stimulated and the high-convergence groups, the static network with heterogeneous weights was used to provide input to an external group of 500 neurons. This group was connected such that it had the same convergence with respect to the stimulated group and the

rest of the network as the high-convergence group, with synaptic weights drawn from the equilibrium distribution shown in Figure 4B. The static network was then stimulated as described in section 4.2.1, to give a reasonable idea of the cross-correlation structure between the stimulated and high-convergence groups in the plastic network at the beginning of the simulation. This is depicted in Figure 12A. To gain insight into how the cross-correlation should have changed as a result of the increased synaptic weights between the stimulated group and the high-convergence group, the experiment was repeated with synaptic weights for the external group drawn from the appropriate distributions shown in Figure 11A. This gives a reasonable idea of the expected cross-correlation at the end of the 400 s simulation, under the assumption that the rate of the stimulated group retains its initial value. The expected cross-correlation is shown in Figure 12B, and a significant increase from the initial cross-correlation is observable. However, the actual cross-correlation between the stimulated and high-convergence groups at the end of the plastic network simulation, shown in Figure 12C, instead indicates a significant reduction in correlation from the initial condition. This suggests that although the initial strong correlation between the stimulated and high-convergence groups causes a systematic increase in the synaptic weights, which should increase the reliability of transferral of the near-synchronous volley of spikes and thus induce development in the outgoing synapses of the high-convergence group, this effect is entirely counteracted by the decoupling of the stimulated group from the rest of the network, which lowers the mean membrane potential of the stimulated group and thereby its responsiveness to the stimulus.

5 Discussion

5.1 Compatibility of STDP with the Balanced Random Network Model. We have shown that in a network with biologically realistic connectivity and sparseness, the weight dynamics of excitatory STDP synapses and the network dynamics can reach a mutual equilibrium. The equilibrium weight distribution is unimodal, and the network dynamics is that of a balanced random network in the asynchronous irregular regime. Specifically, the dynamics is almost identical to that of a static network with weights drawn from the equilibrium distribution of the plastic network. No spontaneous development of structure can be observed, as although the weight distribution is stable, the weights of individual synapses are not. A strong synapse does not stay strong indefinitely, but decays with a characteristic time constant and is replaced by other, previously weak, synapses. In other words, the number of strong synapses remains essentially constant, but the membership to this group is permanently in flux. Smaller networks with lower connectivity did not exhibit spontaneous development of structure either, although strong synapses persisted much longer, with a power law rather than exponential decay. The qualitatively different survival

statistics demonstrate the importance of investigating STDP in networks with biologically realistic connectivity.

These results are in contrast to those of Izhikevich et al. (2004), where the development of neuronal groups was observed, but also in contrast to those of Iglesias et al. (2005), where a principal feature was the systematic pruning of the majority of synapses. It is difficult to determine exactly why the results are so divergent, as there are so many differences in the network, neuron, and plasticity models used. However, we have demonstrated the robustness of our findings. Neither a change in the connectivity in the network, nor a change in the spike pairing scheme, nor even in the specific formulation of STDP altered the essential finding that structure does not spontaneously develop in the balanced random networks studied.

In order to achieve a stable weight distribution and network activity with the desired mean weight and activity statistics, it was necessary to adjust the strength of depressing increments, while keeping the strength of potentiating increments constant, to compensate for the nonvanishing cross-correlation brought about by network oscillations. This tuning amounts to a 6% increase in the strength of depressing increments, compared to that required to maintain the desired mean weight with purely Poissonian firing statistics. Interestingly, this adjustment brings the strength of a depressing increment closer to the value of $\alpha = 0.11$ determined from the experimental data in section 2.2: the analytical value for Poissonian statistics is 8% lower, the adjusted value only 3% lower than the experimental value. If the strength of depressing increments is increased further in the direction of the experimental value, we show that a new stable state is found at a lower mean weight and with a lower degree of synchrony. In contrast, if the strength is decreased, no new stable state is reached within the asynchronous irregular regime. After an initial smooth rate increase, a sudden transition can be observed in the activity dynamics and the weight distribution. The activity is characterized by bursts of strongly patterned activity at high rates, interspersed with network silence. The weight distribution, which remains unimodal until the transition, splits into several distinct clusters.

5.2 Effects of Induced Synchronous Activity. We investigated the effects of irregular but synchronous stimuli on a group of neurons within the stable plastic network. If the stimulus is too strong due to amplification within the group or simply because the group size is too large, a synfire explosion is triggered. Unlike the behavior observed for a static network (Mehring et al., 2003), this explosion does not die out but drives the network into a pathological state. If the stimulus is not strong enough to trigger a synfire explosion, the weights of the outgoing synapses of the stimulated group develop as expected: a net increase is seen, particularly for synapses where the postsynaptic neuron has a high degree of convergence from the stimulated group. However, the causal relationship between the activity

of the stimulated group and the rest of the network is an acausal relationship when viewed from the perspective of the incoming synapses of the stimulated group. These synapses are therefore systematically weakened, which leads to a dramatic fall in the firing rate of the stimulated group. This loss of background input reduces their response to the stimulation, with the end result that the correlation between the stimulated group and the high-convergence group is, in fact, less after 400 s than at the beginning of the simulation, despite the increase in weights. No functional structure is developed under these conditions: the high-convergence group does not echo the stimulus strongly enough to strengthen the synapses of its own downstream high-convergence group.

5.3 Power Law Formulation of STDP. The general framework for normalized STDP update rules established by Gütig et al. (2003) has four free parameters: the maximum weight w_{\max} , the asymmetry parameter α , the step size λ , and the exponent of the weight dependence μ . The time intensity of the simulation study is such that a systematic investigation of all these parameters was unfeasible; therefore, a reexamination of the experimental data of Bi and Poo (1998) was undertaken in order to constrain this framework as tightly as possible. However, when the absolute rather than relative changes in EPSC are considered as a function of the initial EPSC amplitude, a power law relationship emerges that cannot be expressed within the Gütig et al. framework. This novel update rule is a better fit to the data than either additive (Song et al., 2000; van Rossum et al., 2000; Izhikevich et al., 2004) or multiplicative rules (Rubin et al., 2001), or anything in between (Gütig et al., 2003). A property of our formulation of STDP is that the absolute potentiation of very weak synapses is small rather than maximal (for multiplicative rules) or equal to the absolute potentiation of a stronger synapse (for additive rules). This is a clear prediction that can be tested experimentally to distinguish between the available models more conclusively. We combined this formulation of the STDP update rules with an all-to-all spike pairing scheme, which we showed has a tendency to correct rate disparities, in contrast with a nearest-neighbor scheme, which, in the absence of additional stabilization mechanisms, tends to amplify any such disparities.

5.4 Limitations and Perspectives. Inevitably, this study was limited by the long duration of the simulations. The weight dynamics involve timescales up to three orders of magnitude longer than the activity dynamics in a static network: the transient of the weight distribution investigated in section 4.1 lasted about 200 s, whereas the transient in activity dynamics before a stable asynchronous irregular state is established in a static network is more like 200 ms. This is the main problem with such simulations, rather than the increased computational complexity of implementing plastic synapses; this slows our network simulations by a factor of less than

10. Currently, a 1000 s simulation of the unstimulated network at 8.8 Hz requires about 60 hours of computational time on our PC cluster (20×2 processors, AMD Opteron 2.4 GHz, Dolphin/Scali interconnect). Simulating networks with reduced connectivity and scaled synapses ameliorated the problem, but at the cost of changing the survival statistics. It has yet to be established to what extent the network can be reduced from its biologically realistic scale while maintaining qualitatively the same behavior. Regardless of whether this can be achieved, there is a clear need for analytical approaches that simultaneously account for plasticity and activity dynamics.

Given these limitations, there are nonetheless clear indications of how this work can be extended to investigate STDP in a balanced recurrent network more systematically. The discovery that the implementation of delays as purely dendritic can trigger a synfire explosion by strengthening the synapses within a stimulated group suggests that the distribution of propagation delay between the axon and the dendrite is potentially a crucial parameter for the network dynamics. Future work will need to determine to what extent the pathological state exhibited by the network as a result of a synfire explosion or insufficiently strong depression is affected by an increased proportion of axonal delay. Moreover, the possibility that the patterns observed in these states are at least partially artifacts of discrete time step simulation cannot be ruled out. We are currently investigating efficient methods of incorporating continuous spike times within a discrete time simulation scheme (Morrison, Straube, Plesser, & Diesmann, 2007), which should allow clarification of this issue. Another major area for development is the relative simplicity of the plasticity model used. It is, of course, necessary to start with simple models; otherwise, when confronted with complex behavior, it is impossible to say which aspect of the model is causing what. However, recent research reveals a rich variety of plasticity expression such as suppression (Froemke & Dan, 2002), activity-dependent postsynaptic homeostasis (Turrigiano et al., 1998), and nonlinear interactions between potentiation and depression (Wang et al., 2005). These effects may turn out to have fundamental consequences for network behavior. The model used also incorporates no upper bound on synaptic strength, which seems unbiological. However, as no upper bound could be extracted from the experimental data of Bi and Poo (1998), further investigation is required to establish a reasonable saturation behavior for the power law update rule proposed here.

If structure is to develop in a balanced random network as a response to synchronous stimuli, the stimulated group must continue to receive enough background input so that its response to the stimulus does not diminish over time. To achieve this, one or both of the model components in this study must be adjusted. A natural adjustment to the plasticity model would be to incorporate postsynaptic activity dependent homeostasis as used in van Rossum et al. (2000). This would increase the weights of the incoming synapses of the stimulated group as a response to its falling rate.

Alternatively, the connectivity in the network model could be adjusted. First, in the model used here, all propagation delays are identical, and each neuron has the same number of incoming synapses. Introducing heterogeneities in these quantities has been shown to reduce network oscillations (Brunel, 2000; Tetzlaff et al., 2005), which might prevent the network from entering a pathological state under stronger stimulation than was possible to use here. Second, a connection pattern resulting in a long-tailed rather than binomial convergence distribution for the stimulated group would lessen the acausal correlation between the embedding network and the stimulated group, thus reducing the drop in the mean weight of its incoming synapses.

Appendix: Algorithmic Implementation of STDP

The following implementation assumes that a synapse that connects the axon of neuron i to the dendrite of neuron j is stored in a list belonging to i . In the case of a distributed application, this list is assumed to be physically on the machine of the postsynaptic neuron, j (see Morrison et al., 2005). Each list of postsynaptic targets maintains two variables: t_{old} , the time of the last presynaptic spike, and K_+ , which describes the current time-dependent weighting of the STDP update rule for potentiation. Each synapse maintains its current weight w_{ji} and its synaptic delay $d_i = d_i^A + d_i^D$, which is composed of its axonal delay d_i^A and its dendritic delay d_i^D , such that $d_i^D \geq d_i^A$. Note that these propagation delays are taken into account when calculating the difference in time between presynaptic and postsynaptic spikes: spike timing is calculated from the perspective of the synapse and not the soma (see Debanne et al., 1998). Additionally, the synapse is equipped with two functions $F_+(w)$ and $F_-(w)$, which perform the weight-dependent update for potentiation and depression, respectively. For example, in the case of power law STDP, $F_-(w) = \lambda \alpha w$. The target list performs its weight updates and sends spikes to its postsynaptic neurons when its spike method is invoked for a presynaptic spike at time t :

```

spike(t):
  for each postsynaptic neuron j:
    history ← j.get_history(told + djA - djD, t + djA - djD)
    for each spike tj in history:
      dt ← (tj + djD) - (told + djA)
      if dt ≠ 0:
        wj ← wj + F+(wj) · K+ · exp(-dt/tau)
      K- ← j.get_K_value(t + djA - djD)
      wj ← wj - F-(wj) · K-
      send spike (wj, dj) to neuron j
  K+ ← K+ · exp(-(t - told)/tau) + 1
  told ← t

```

The postsynaptic neuron j maintains the variables t_{old} , containing its last spike, K_- , which describes the current time-dependent weighting of the STDP update rule for depression, and N_{syn} , the number of STDP incoming synapses. In addition, a dynamic data structure `spike_register` is required, which stores spike information in the form of tuples: $(t_{\text{sp}}, K_{\text{sp}}, \text{counter}_{\text{sp}})$, where t_{sp} is the spike time, K_{sp} is the value of K_- at time t_{sp} , and $\text{counter}_{\text{sp}}$ counts how many times the spike information has been accessed by synapses. When neuron j spikes at time t , the `spike_register` has to be updated:

```
update_register(t):
    K_- ← K_- · exp(-(t - t_old)/tau) + 1
    while length of spike_register ≥ 1:
        if counter_sp ≥ N_syn:
            remove first element
    store (t, K_-, 0) as last element
    t_old ← t
```

To enable the synapses to access this information, two other methods have to be available: `get_history(t1, t2)`, which returns the spike times recorded in the `spike_register` in the range $(t_1, t_2]$, and `get_K_value(t)`, which returns the value of K_- for the time t . They may be implemented as follows:

```
get_history(t1, t2):
    starting at beginning of spike_register:
    while t_sp ≤ t1:
        move to next element
    history ← t_sp
    counter_sp ← counter_sp + 1
    while t_sp ≤ t2:
        append t_sp to history
        counter_sp ← counter_sp + 1
        move to next element
    return history

get_K_value(t):
    starting at end of spike_register:
    while t_sp ≥ t:
        move to previous element
    return K_sp · exp(-(t - t_sp)/tau)
```

This implementation permits a wide class of STDP models to be simulated in a distributed application. Maintaining a `spike_register` means that a neuron does not require knowledge of the locations of its incoming

synapses in order to update their weights when it spikes; synaptic weights are updated only when a presynaptic spike occurs. Using the K_{\pm} variables to implement the sums of exponential functions and tracking `countersp`, the number of times synapses have accessed a particular spike, allows spike information that is no longer relevant to be discarded, thus keeping the number of retained spikes to a minimum. This means that the memory required does not increase for longer simulations, and the list operations in `gethistory` remain fast.

This template algorithm can be trivially extended to models that require clipping of the weights at upper and lower boundaries or where the presynaptic and postsynaptic time constants differ. Different spike pairing schemes can be investigated by changing the way the K_{\pm} variables are updated. With more modifications, it can also be extended to more complex models such as the suppression model put forward by Froemke and Dan (2002).

The biggest limitation of the algorithm is the requirement that $d_i^D \geq d_i^A$. This is necessary to ensure causality is not violated—the postsynaptic neuron must not be required to give information on its future state. In this article, for simplicity, the assumption was made that the delay is purely dendritic: $d_i^D = d_i$, $d_i^A = 0$. Using the above algorithm, other distributions of the delay between the dendrite and the axon can be investigated up to the limit of $d_i^D = d_i^A = d_i/2$. This limitation can be circumvented, but at the cost of more complicated techniques, essentially queueing the event until the postsynaptic neuron has reached a state where the required information is available.

Acknowledgments

This work was carried out when A.M. and M.D. were based at the Bernstein Center for Computational Neuroscience, Albert-Ludwigs-University, Freiburg. We are particularly grateful to Guo-Qiang Bi and Mu-ming Poo for providing us with their original data. We acknowledge constructive discussions with the members of the NEST initiative (www.nest-initiative.org). This work would not have been possible without a dedicated grant for a high-performance computer cluster, HBFG Chapter 1423 Title 812 59, and the exertions of Martin Walter and Ulrich Gehring of the Freiburg University Computer Center. This work was partially funded by DAAD 313-PPP-N4-1k, DIP F1.2, and BMBF Grant 01GQ0420 to the Bernstein Center for Computational Neuroscience Freiburg.

References

- Bi, G.-q., & Poo, M.-m. (1998). Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.*, *18*, 10464–10472.

- Bienenstock, E. L., Cooper, L. N., & Munro, P. W. (1982). Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex. *J. Neurosci.*, *2*(1), 32–48.
- Braitenberg, V., & Schüz, A. (1998). *Cortex: Statistics and geometry of neuronal connectivity* (2nd ed.). Berlin: Springer-Verlag.
- Brunel, N. (2000). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *J. Comput. Neurosci.*, *8*(3), 183–208.
- Brunel, N., & Hakim, V. (1999). Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural Comput.*, *11*(7), 1621–1671.
- Burkitt, A. N., Meffin, H., & Grayden, D. B. (2004). Spike-timing-dependent plasticity: The relationship to rate-based learning for models with weight dynamics determined by a stable fixed point. *Neural Comput.*, *16*, 885–940.
- Debanne, D., Gähwiler, B. H., & Thompson, S. M. (1998). Long-term synaptic plasticity between pairs of individual CA3 pyramidal cells in rat hippocampal slice cultures. *J. Physiol. (Lond.)*, *507*, 237–247.
- Feldman, D. E. (2000). Timing-based LTP and LTD at vertical inputs to layer II/III pyramidal cells in rat barrel cortex. *Neuron*, *27*, 45–56.
- Fetz, E., Toyama, K., & Smith, W. (1991). Synaptic interactions between cortical neurons. In A. Peters (Ed.), *Cerebral cortex* (Vol. 9, pp. 1–47). New York: Plenum.
- Froemke, R. C., & Dan, Y. (2002). Spike-timing-dependent synaptic modification induced by natural spike trains. *Nature*, *416*(6879), 433–438.
- Gerstner, W., Kempter, R., van Hemmen, J. L., & Wagner, H. (1996). A neuronal learning rule for sub-millisecond temporal coding. *Nature*, *383*, 76–78.
- Gütig, R., Aharonov, R., Rotter, S., & Sompolinsky, H. (2003). Learning input correlations through nonlinear temporally asymmetric Hebbian plasticity. *J. Neurosci.*, *23*(9), 3697–3714.
- Guyonneau, R., VanRullen, R., & Thorpe, S. J. (2005). Neurons tune to the earliest spikes through STDP. *Neural Comput.*, *17*, 859–879.
- Hertz, J., & Prügel-Bennet, A. (1996). Learning synfire-chains by self-organization. *Network: Comput. Neural Systems*, *7*(2), 357–363.
- Iglesias, J., Eriksson, J., Grize, F., Tomassini, M., & Villa, A. (2005). Dynamics of pruning in simulated large-scale spiking neural networks. *Biosystems*, *79*, 11–20.
- Izhikevich, E. M., & Desai, N. S. (2003). Relating STDP to BCM. *Neural Comput.*, *15*, 1511–1523.
- Izhikevich, E. M., Gally, J. A., & Edelman, G. M. (2004). Spike-timing dynamics of neuronal groups. *Cereb. Cortex*, *14*, 933–944.
- Kempter, R., Gerstner, W., & van Hemmen, J. L. (1999). Hebbian learning and spiking neurons. *Phys. Rev. E*, *59*, 4498–4514.
- Kempter, R., Gerstner, W., & van Hemmen, J. L. (2001). Intrinsic stabilization of output rates by spike-based Hebbian learning. *Neural Comput.*, *12*, 2709–2742.
- Kistler, W. M., & van Hemmen, J. L. (2000). Modeling synaptic plasticity in conjunction with the timing of pre- and postsynaptic action potentials. *Neural Comput.*, *12*, 385–405.
- Levy, N., Horn, D., Meilijson, I., & Ruppin, E. (2001). Distributed synchrony in a cell assembly of spiking neurons. *Neural Networks*, *14*, 815–824.
- Markram, H., Lübke, J., Frotscher, M., & Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*, *275*, 213–215.

- Mehring, C., Hehl, U., Kubo, M., Diesmann, M., & Aertsen, A. (2003). Activity dynamics and propagation of synchronous spiking in locally connected random networks. *Biol. Cybern.*, *88*(5), 395–408.
- Morrison, A., Mehring, C., Geisel, T., Aertsen, A., & Diesmann, M. (2005). Advancing the boundaries of high connectivity network simulation with distributed computing. *Neural Comput.*, *17*(8), 1776–1801.
- Morrison, A., Straube, S., Plesser, H. E., & Diesmann, M. (2007). Exact subthreshold integration with continuous spike times in discrete time neural network simulations. *Neural Comput.*, *19*(1), 47–79.
- Rubin, J., Lee, D., & Sompolinsky, H. (2001). Equilibrium properties of temporally asymmetric Hebbian plasticity. *Phys. Rev. Lett.*, *86*, 364–367.
- Sjostrom, P., Turrigiano, G., & Nelson, S. (2001). Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron*, *32*, 1149–1164.
- Song, S., Miller, K. D., & Abbott, L. F. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nat. Neurosci.*, *3*(9), 919–926.
- Tetzlaff, T., Morrison, A., Timme, M., & Diesmann, M. (2005). Heterogeneity breaks global synchrony in large networks. In *Proceedings of the 30th Göttingen Neurobiology Conference. Neuroforum* (Supp. 1), 206b.
- Turrigiano, G. G., Leslie, K. R., Desai, N. S., Rutherford, L. C., & Nelson, S. B. (1998). Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature*, *391*, 892–896.
- van Rossum, M. C. W., Bi, G.-Q., & Turrigiano, G. G. (2000). Stable Hebbian learning from spike timing-dependent plasticity. *J. Neurosci.*, *20*(23), 8812–8821.
- van Vreeswijk, C., & Sompolinsky, H. (1996). Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, *274*, 1724–1726.
- van Vreeswijk, C., & Sompolinsky, H. (1998). Chaotic balanced state in a model of cortical circuits. *Neural Comput.*, *10*, 1321–1371.
- Wang, H.-X., Gerkin, R. C., Nauen, D. W., & Bi, G.-Q. (2005). Coactivation and timing-dependent integration of synaptic potentiation and depression. *Nat. Neurosci.*, *8*(2), 187–193.
- Zhang, L. I., Tao, H. W., Holt, C. E., Harris, W. A., & Poo, M. M. (1998). A critical window for cooperation and competition among developing retinotectal synapses. *Nature*, *395*, 37–44.