

On the Transmission of Rate Code in Long Feedforward Networks with Excitatory–Inhibitory Balance

Vladimir Litvak,¹ Haim Sompolinsky,^{1,2} Idan Segev,^{1,3} and Moshe Abeles^{1,4}

¹The Interdisciplinary Center for Neural Computation, ²The Racah Institute of Physics, ³Department of Neurobiology, Institute of Life Sciences, and ⁴Department of Physiology, Hadassah Medical School, The Hebrew University, Jerusalem 91904, Israel

The capability of feedforward networks composed of multiple layers of integrate-and-fire neurons to transmit rate code was examined. Synaptic connections were made only from one layer to the next, and excitation was balanced by inhibition. When time is discrete and the synaptic potentials rise instantaneously, we show that, for random uncorrelated input to layer one, the mean rate of activity in deep layers is essentially independent of input firing rate. This implies that the input rate cannot be transmitted reliably in such feedforward networks because neurons in a given layer tend to synchronize partially with each other because of shared inputs. As a result of this synchronization, the average firing rate in deep layers will either decay to zero or reach a stable fixed point, depending on model parameters. When time is treated continuously and the synaptic potentials rise instantaneously, these effects develop slowly, and rate transmission over a limited number of layers is possible. However, the correlations among neurons at the same layer hamper reliable assessment of firing rate by averaging over 100 msec (or less). When the synaptic potentials develop gradually, as is the realistic case, transmission of rate code fails. In a network in which inhibition only balances the mean excitation but is not timed precisely with it, neurons in each layer fire together, and this volley successively propagates from layer to layer. We conclude that the transmission of rate code in feedforward networks is highly unlikely.

Key words: rate code; temporal code; synfire chain; network models; excitation–inhibition balance; synchrony; bistability; correlation; synaptic integration; information transmission

Introduction

Our current understanding of the processing of sensory information relies on the notion of multiple stages of feature extraction. This can be implemented as neuronal activity progresses from one cortical area to another, and within each cortical area. Indeed, reaction time for nontrivial perceptual tasks (Thorpe and Fabre-Thorpe, 2001) suggests the existence of several tens of processing stages, assuming ~ 10 msec transmission delay between stages. A very simple model for this type of processing is a feedforward chain of layers of neurons, in which each neuron of a given layer receives multiple synaptic inputs from some of the neurons in the previous layer. Within such a feedforward network, information can be coded in different ways. One possibility is that information is carried in such a system solely through the firing rate of the neurons (Shadlen and Newsome, 1994). In this “rate code” paradigm, neurons in each layer fire at random times (Softky and Koch, 1993) and in an uncorrelated manner with other neurons belonging to the same layer.

In the rate code paradigm, neurons in the next layer compute, within a short time window, the average firing rate of the neurons

in the previous layer and generate an output rate that is related uniquely to the input rate. An alternative for rate coding is the “temporal code” paradigm, in which information is carried by small groups of neurons that fire in synchrony with each other (Abeles, 1982, 1991; Bienenstock, 1995; Softky, 1995; Stevens and Zador, 1998; Diesmann et al., 1999). In a feedforward network, if each layer fires in synchrony, then the next layer will also do so and “synfire” activity will develop (Abeles, 1982). Can feedforward networks of neurons support rate code inherently, or do such networks tend to generate synfire waves of activity spontaneously?

Shadlen and Newsome (1994) developed a model to demonstrate the feasibility of feedforward rate transmission. Their model is based on the notion of balance between excitation and inhibition, whereby each synaptic potential is rather large, but because of this balance, the sum of many random excitatory and inhibitory presynaptic inputs results in a postsynaptic membrane voltage that fluctuates strongly around the resting potential. These random voltage fluctuations occasionally cross the threshold for spike firing and generate random firing in the output neurons. Shadlen and Newsome claim that their model has a rate gain (the ratio between the firing rate of one layer and that of the previous layer) of unity and that the neurons are not sensitive to the timing of their single inputs. They also claim that if each pair of output neurons shares $< 40\%$ of the input neurons, only a small degree of synchrony will be developed, and this ensures an efficient rate code. The feasibility of rate code transmission in unbalanced feedforward layers was studied recently in a model by van Rossum et al. (2002).

The present work aims at a better understanding of the firing

Received June 21, 2002; revised Jan. 10, 2003; accepted Jan. 21, 2003.

This research was supported in part by a grant from the United States–Israel Binational Science Foundation and a grant from the Israeli Science Foundation. I.S. was supported by the Israeli Science Foundation and the Office of Naval Research. We gratefully acknowledge useful correspondence with Michael Shadlen. We thank Carl van Vreeswijk for helpful discussions, particularly for suggesting the calculation that appears in Appendix. H.S. gratefully acknowledges helpful discussions with Daniel D. Lee.

Correspondence should be addressed to Prof. Haim Sompolinsky, Racah Institute of Physics, Hebrew University, Jerusalem 91904, Israel. E-mail: haim@fiz.huji.ac.il.

V. Litvak's present address: Department of Bio-Medical Engineering, Technion–Israel Institute of Technology, Technion City, Haifa 32000, Israel.

Copyright © 2003 Society for Neuroscience 0270-6474/03/233006-10\$15.00/0

dynamics in feedforward networks of neurons. The feasibility of supporting rate code versus temporal code in feedforward networks is discussed.

Materials and Methods

Model neuron. We used three types of model neurons. In two of the models, the counting model described by Shadlen and Newsome (1994) was used. In this model, the membrane potential is not continuous but rather jumps instantaneously in steps of 1 mV whenever a synaptic input arrives. The model was implemented in two different ways. One implementation used continuous time, and the other used discrete time steps. In the third model, synaptic inputs were modeled as current transients having an α shape; in this model, both time and membrane potentials changed continuously.

In the first model, the spike trains were represented as series of specific times when each presynaptic spike occurred. The value of the postsynaptic membrane potential was recalculated analytically at the time of arrival of each presynaptic spike by adding (excitation) or subtracting (inhibition) a step of 1 mV from the membrane voltage. Between synaptic inputs, the membrane potential decays exponentially toward zero, with a time constant of 20 msec. The model has a lower reflecting boundary beyond which the neuron does not hyperpolarize. Whenever the membrane potential hits threshold, the neuron emits an action potential and the membrane voltage is reset immediately to a “reset potential,” after which the dynamics of the membrane potential resumes. We refer to this model as the discrete-PSP continuous-time model. Refractoriness slows down the output at high firing rates. Thus, one cannot hope to obtain linear transmission of firing rates over a wide range of input firing rates. Shadlen and Newsome modeled neurons without refractoriness and claimed that in such a model, linear rate transmission is possible. We found that even under this nonphysiological assumption, faithful transmission of rates is not possible. Adding a refractory period worsens the situation. To gain insight into transmission of rate code, beyond the impairment of refractoriness, we repeated Shadlen and Newsome simulations with their exact model.

In the second model, inputs were the same as in the first model, but the simulation proceeded in time bins of 1 msec, and each spike train was a sequence of zeros (no spike) and ones (spike firing). The rest of the simulation was identical to the previous model. We refer to this model as the discrete-PSP discrete-time model.

The third model neuron had its synaptic inputs modeled by a current with a continuous time course described by an α function. The membrane potential was simulated by the following equation:

$$-\tau_m \frac{dV_m}{dt} = R \sum I_{\text{syn}} + V_m \quad (1)$$

where V_m is the membrane potential, τ_m is 20 msec, and I_{syn} are the synaptic currents for all inputs since the last action potential. Synaptic currents were given by the following equation:

$$I_{\text{syn}} = A \cdot (t - t_0) e^{-(t - t_0)/\tau_{\text{syn}}} \quad (2)$$

where τ_{syn} is 1 msec, t_0 is the firing time of the presynaptic spike, and A was adjusted so that either the peak PSP was 1 ($A = 24,370$) or the total area under the continuous PSP was equal to that of the discrete PSP ($A = 20,140$). A was positive for EPSP and negative for IPSP. When the membrane potential reached the threshold for firing, an action potential was marked, and the membrane potential and all previous synaptic currents were reset to 0. Although this last point is not physiological, it was needed to obtain results with continuous time simulation within reasonable computer time (a few days with 1 GHz processor). When the membrane potential was smaller than -1 , it was clamped at -1 , but the previous synaptic currents were not ignored. There was no explicit refractory period, but because the synaptic currents were reset to 0 after an action potential and new PSPs develop only gradually, it took some time before the neuron fired again. At the highest input rate (200 spikes/sec) and lowest threshold of 8 was examined, the shortest interspike interval was

1.3 msec. We refer to this model as the continuous-voltage continuous-time model.

Results

We first concentrate on results obtained with the discrete-PSP continuous-time neuron model. This model is identical to the one used by Shadlen and Newsome, but we extended their simulations by analyzing what happens beyond one or two layers.

Input–output relations for the discrete-PSP continuous-time neuron

We simulated a single-neuron model receiving 600 inputs, 300 excitatory and 300 inhibitory. The inputs were long, uncorrelated Poisson spike trains; the average input rate varied between 10 and 100 spikes/sec. In such a precisely balanced situation, the net synaptic current is zero, and the response is driven entirely by the variance in the membrane potential. Shadlen and Newsome found that for this input and for an appropriate choice of parameters, the model neuron exhibits a linear relationship between the mean input rate and the mean firing rate of the neuron. Furthermore, the slope of this linear input–output curve is 1. According to Shadlen and Newsome, this occurs, for example, when the reset potential is zero, the lower barrier is slightly below zero, and the threshold is 15. Although we were able to find parameters for which the input–output relationship of the model neuron, with the above balanced input, is approximately linear (for rates up to 100 spikes/sec), we found that it is very hard to obtain an input–output gain of ~ 1 . After testing a large number of combinations of lower-barrier, reset, and threshold values, the best we could achieve was when the resting potential was set to 0, the reset potential was -0.5 , and the lower barrier was -17 mV. A threshold value of 12 yielded a gain that was closest to 1. These parameters are substantially different from those suggested by Shadlen and Newsome. Our simulations of the Shadlen and Newsome model with the parameters quoted by them (Shadlen and Newsome, 1998, their Fig. 1) failed to replicate their results. In fact, we found that, with the parameters used by Shadlen and Newsome, the input–output curve deviates from a linear curve with a slope of unity. For our model, these deviations are smaller (Fig. 1). At very low input rates, the output rate is below the input rate (because the membrane potential may never reach threshold). At high input rates, the curves are almost linear; however, it is impossible to have the output follow the input in exactly a one-to-one ratio.

Transmission of firing rates in feedforward networks

To test whether the discrete-PSP continuous-time model can transmit the input firing rate in a multilayered network, the following simulations were performed. A feedforward network containing 20 layers was constructed; all neurons in the network were identical, as in Figure 1. Each neuron received 600 inputs, of which exactly 300 were excitatory and 300 were inhibitory. Note that because of the exact balance, activity is driven solely by the variance of the membrane potential. According to Shadlen and Newsome, synchronization among neurons develops only when the neurons share $\geq 40\%$ of their inputs. To avoid this region, we concentrated our simulations on networks in which any pair of neurons in a single layer shared only 10% of their inputs, but we ran sample simulations for other degrees of shared inputs also. To obtain 10% of shared inputs with 300 excitatory and 300 inhibitory neurons, one needs to have 3000 excitatory and 3000 inhibitory neurons in the input layer. Thus, 6000 simulated spike trains from the one layer were used as inputs to the next layer.

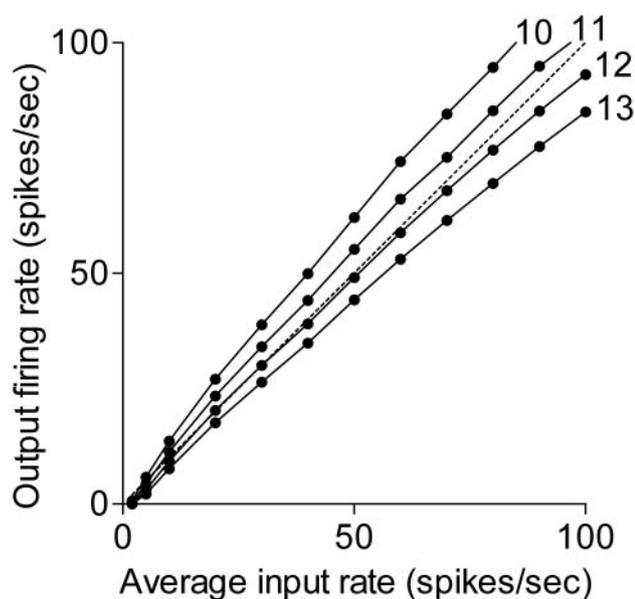


Figure 1. Input–output relations for a continuous-time discrete-PSP model. Simulation of a model neuron used by Shadlen and Newsome receiving 300 excitatory and 300 inhibitory inputs firing independent-Poissonian spike trains. This model is identical to the Shadlen and Newsome model. In this simulation, the membrane time constant was 20 msec, the lower reflecting boundary was -17 , and the reset voltage after reaching threshold was -0.5 mv.

Exactly half of these spike trains were chosen to be excitatory and the other half to be inhibitory. The spike trains in the input layer (layer 1) were random and uncorrelated. Each neuron in the second layer received, as an input, 600 of these spike trains. Simulations differed with respect to the average firing rate of the random spike trains that were used as the input to the first layer. The same connectivity matrix between adjacent layers was maintained throughout the network. To explore the evolution of the dynamics along multiple layers, we analyzed a system of 20 layers, although from a physiological point of view, 20 layers seems unrealistic. As will be shown below, severe transmission problems became apparent after only three layers.

The results of one such simulation are presented in Figure 2. In this simulation, the threshold was 11, the percentage of shared inputs between two neurons at a given layer was 10%, and the initial firing rate was 50 spikes/sec. The activity of 20 randomly chosen spike trains, out of the 6000 produced in each layer, is shown. It can be seen that the uncorrelated firing at a constant rate seen at the input layer is not preserved at subsequent layers. Instead, in the deep layers, the mean rate decreases and the neurons exhibit periods of synchronized activity.

The behavior of the network for various input rates is summarized quantitatively in Figure 3. For input rates of ≥ 30 Hz, the firing rates converge after 20 layers to a common mean rate of ~ 40 Hz. It is instructive to view the mean rate of a layer as an iterative dynamic variable in which different layers correspond to different time units. The results of Figure 3 indicate that for all initial rates of ~ 30 Hz and above, the mean layer rate converges to a common fixed point of ~ 40 Hz independent of the initial value. Conversely, the 10 Hz curve indicates that for low initial rates, the layer rate does not converge to the 40 Hz fixed point. We discuss the low-rate behavior later (see Fig. 9). Changing the thresholds to 12 causes the firing rates to decline rapidly toward zero for all input rates (data not shown); in particular, for this threshold, the only stable fixed point of the system is zero. Changing the threshold to 10 causes the firing rates to build up rapidly to large values, meaning that the nonzero fixed point is the infi-

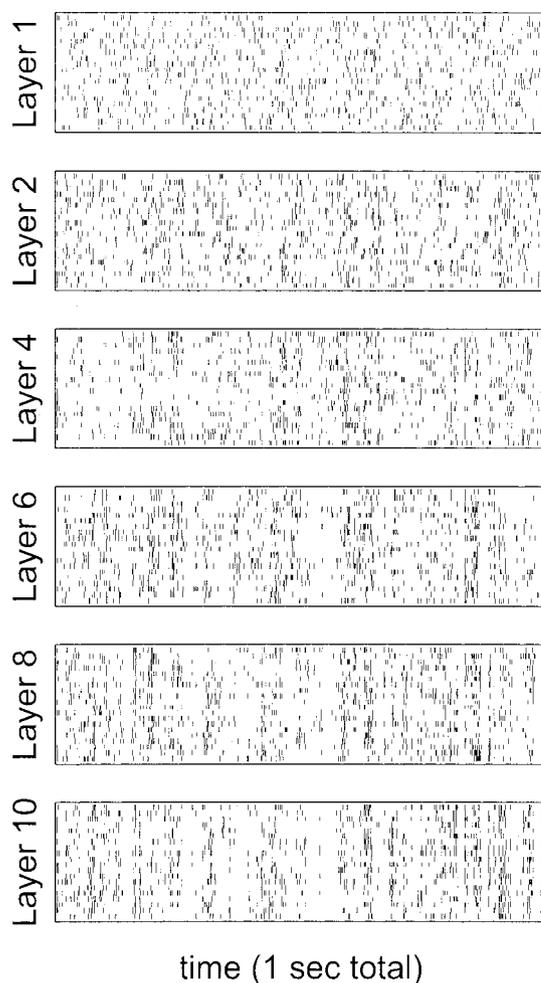


Figure 2. Dot displays for a network of continuous-time discrete-PSP neurons. One second of activity of 20 randomly selected neurons from the input stream (layer 1) and several layers along the network. Each layer consisted of 6000 neurons; the percentage of shared inputs between two neurons at a given layer was 10%. Each neuron received 300 excitatory and 300 inhibitory inputs. The same connection matrix was used for all the layers. Spike trains in the input layer (layer 1) are random and uncorrelated, with an average rate of 50 spikes/sec. Threshold was set to 11, at which the decline in firing rate in deeper layers was very gradual. Buildup of synchrony is obvious.

nite (or, more realistically, saturated) rate. Thus, the optimal threshold for the network behavior is found to be slightly lower than the optimal threshold for a single neuron, which is 12. We return to this point below.

Figure 3 indicates that (except for low initial rates), the mean rates of the layers approach a common fixed point value after 20–30 layers. This by itself would allow for the possibility that rate information can be transmitted across ~ 10 layers, when the mean rates corresponding to different input rates are still apart from each other. Transmitting information by firing rates is meaningful if the firing rates of the layers can be estimated by sampling a limited number of neurons over a limited time period. We thus estimated the errors that would occur in rate estimation by summing the activity of 600 neurons over 100 msec. The number 600 was chosen because it is reasonable to assume that the connectivity of a “read-out” neuron will be similar to that of the neurons in our network. The time window of 100 msec is on the order of magnitude of the minimum time of meaningful sensory processing as measured in psychophysical experiments. The resultant errors are shown by the vertical bars in Figure 3,

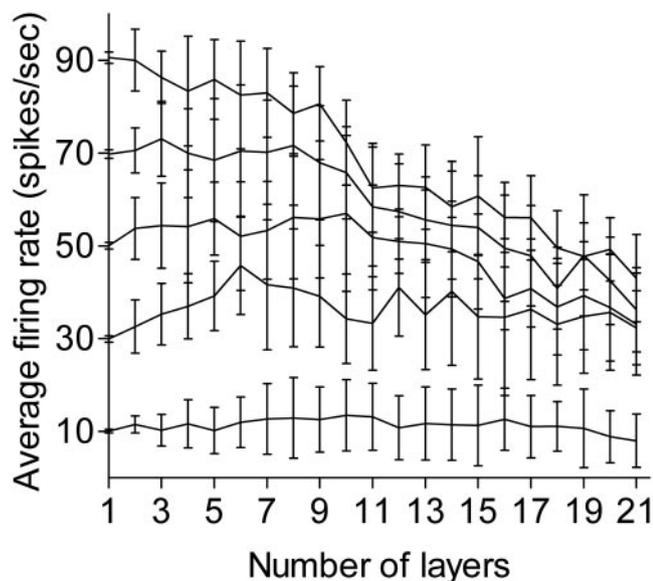


Figure 3. Modulation of firing rates along a feedforward network of continuous-time discrete-PSP neurons. Same network as in Figure 2 but with different initial firing rates at input. The error bars show SDs in estimation of rates based on observing 600 neurons for 100 msec. Note the big difference for the errors for the input layer (layer 1) and the next layer (layer 2). After three layers, the error bars of adjacent curves start to overlap.

which correspond to ± 1 SD. After three layers, the error bars, corresponding to input frequencies that are 20 spikes/sec apart, start to overlap. Even after one layer, these error bars become much larger than the error bars of the input (independent Poissonian) spike trains. The reason for the rapid increase in the estimation error is the correlations that develop between the neurons, as discussed below.

Emergence of synchrony in the network

Examination of Figure 2 shows that two types of synchrony appear. Over a long time scale, periods of high firing rates alternate with periods of low firing rates. Over very short time periods, vertical lines appear, indicating precise synchrony. To quantify this synchrony, we measured the cross-correlograms between neuronal pairs within the same layer. Figure 4A shows the mean cross-correlograms of 200 randomly chosen pairs from the input sources (*Layer 1*) and various subsequent layers. The correlation builds up gradually as we proceed along the layers. This is shown quantitatively in Figure 4B, in which spike trains were converted sequences of zeros and ones, with time steps of 1 msec and then cross-correlated. The ratio of the peak in the cross-correlation divided by the peak of the autocorrelation is plotted. The rate of buildup of synchrony depends on the level of shared inputs among pairs of neurons in the same layer. If there were no shared inputs, neurons would fire independently and the graphs in Figure 1 could be used to evaluate the transmission of rates between layers. If neurons shared all their inputs (100% shared inputs), all the neurons in a given layer would have exactly the same inputs, and they would fire in unison. The results with 10% shared inputs demonstrate that even with a limited degree of shared inputs, substantial synchrony builds up relatively rapidly.

Comparing Figures 3 and 4 indicates that although the estimation error seems to saturate in the sixth to seventh layers (see the size of the error bars in Fig. 3), the cross-correlogram peaks continue to increase roughly linearly up to the 20th layer. This indicates that the main source of estimation error is the covaria-

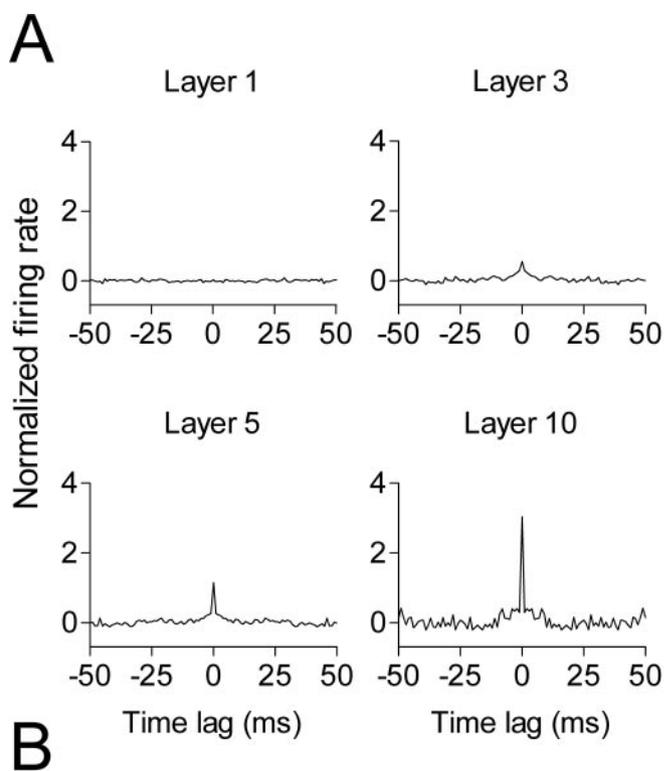


Figure 4. Buildup of synchrony along a feedforward network of continuous-time discrete-PSP neurons. Same network as in Figures 2 and 3. Cross-correlations are based on averaging 200 pairwise correlations. *A*, At input layer 1, the spike trains are uncorrelated. A small correlation appears and grows slowly at deeper layers. All graphs were normalized by the product of average rates of the two neurons. *B*, Development of correlation for different initial rates. Values are the ratio between the peak of the cross-correlation and the peak of the autocorrelations.

tion of rates over long time scales. These rate correlations are not captured by the cross-correlograms in Figure 4, which measure the synchrony over time scales of tens of milliseconds.

Input–output relations for a neuron with continuous membrane voltage and continuous time

The discrete-PSP continuous-time model is peculiar in that it makes a big difference whether an EPSP arrives immediately before or after an IPSP. A burst of excitatory spikes can trigger a spike in the output neuron even if these EPSPs are followed immediately by IPSPs. To study how different the discrete-PSP

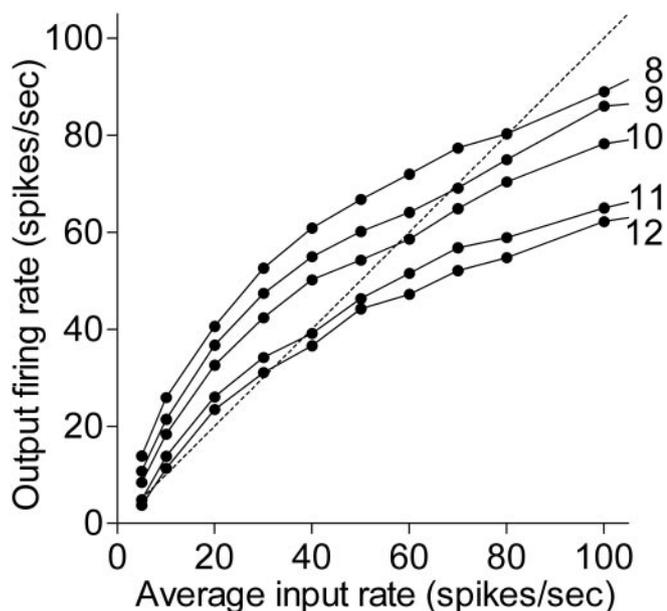


Figure 5. Input–output relations for a continuous-time continuous-voltage model. Same conventions as in Figure 1. The model neuron here had PSPs modeled by a current having an α shape (see Materials and Methods). The lower-voltage barrier was set to -1 , and the reset voltage after hitting threshold was 0 mV. At very low input rates, the curves must be convex, but for most of the range, they are concave. Rates were estimated by measuring the time required to generate 400 spikes. Error bars are $\pm 5\%$.

model is from a more plausible continuous-voltage model, we investigated a model in which synaptic potentials were generated by current pulses having the shape of an α -function (Rall, 1967).

Figure 5 shows the input–output relations for such a neuron with various thresholds. The deviation from linearity of the curves is substantially higher than those of the discrete-PSP model (compare with Fig. 1). At very low rates, they are similar to those of the discrete-PSP model, whereas at higher rates, they flatten considerably. The gradual rise of the PSPs allows for integration of EPSCs and IPSCs before they exert their full effect on the membrane potential, thereby reducing the gain of the neuron. The relations in Figure 5 were obtained from simulations with a lower barrier of -1 . Lowering the barrier to -8 or -17 causes even larger flattening at high rates.

The considerable difference between the continuous-voltage and the discrete-PSP models may be appreciated by comparing the corresponding membrane voltage fluctuations in the two models during a 16 msec simulation (Fig. 6). Input firing rates were low (10 spikes/sec), and the output neuron did not fire. The discrete model shows multiple upswings and downswings, whereas the continuous model (*thick line*) tends to average them out. Had the input firing rate been increased 10-fold (to 100 spikes/sec), the discrete model would look almost the same, but with time squashed to 1.6 msec. The continuous model cannot vary much within 1.6 msec and would look much smoother.

Simulation of 20 layers with 6000 neurons, each composed of continuous-time continuous-voltage model neurons, is not practical even with fast computers. However, one can approximate this model by using discrete time steps with discrete voltage jumps. This model first computes the difference between EPSPs and IPSPs in a single time step and only then updates the membrane potential. Although this model still shows discrete membrane potential jumps, it allows for partial averaging out of EPSPs and IPSPs at high input rates. This model is similar to the one used by Salinas and Sejnowski (2000).

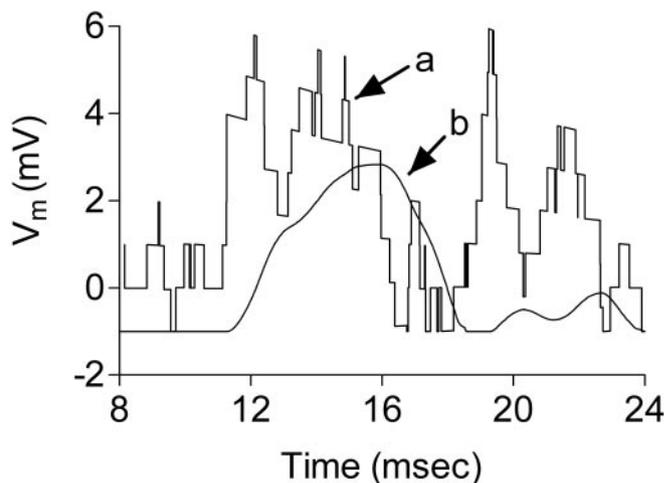


Figure 6. Membrane potentials for neuronal models. The membrane potential during the 16 msec span is shown for a simulation with 300 excitatory and 300 inhibitory neurons, each firing at 10 spikes/sec. The resting potential of the model neurons is 0 mV, and the lower barrier is -1 mV. Curve *a* shows the behavior of the continuous-time discrete-PSP model. The membrane potential shows frequent up and down jumps of 1 mV in size. Curve *b* shows the behavior of the continuous-time continuous-voltage model. The EPSPs and IPSPs tend to average out. Note that, for the continuous-voltage model (graph *b*; time, 8 – 12 msec), hyperpolarizing currents may continue to pull the membrane potential down even when the potential is clamped by the lower barrier. In contrast, in the discrete model (graph *a*), once the membrane potential hits the lower barrier, all the previous IPSPs are “forgotten.”

Input–output relations for the discrete-PSP discrete-time neuron

We repeated the simulations leading to Figures 1 and 5 with the discrete-time discrete-PSP neuron with time steps of 1 msec. As before, we used 600 inputs, 300 excitatory and 300 inhibitory. The inputs were long, uncorrelated Poisson spike trains; the average input rate varied between 10 and 100 spikes/sec. The resting potential was set to 0 , and a lower reflecting barrier was set to -1 . The threshold for spike firing was varied between 12 and 17 to monitor the sensitivity of the model to changes in threshold value. The results are presented in Figure 7. The curves tend to flatten at high input rates, as for the continuous-time and continuous-voltage model (Fig. 5). The gain clearly depends strongly on the threshold value. Therefore, fine-tuning of parameters is necessary to obtain a gain close to unity even for a restricted range of input rates. On the basis of these results, we conclude that a threshold value of 15 is the optimal value for generating a gain close to unity in an appreciable range of input rates.

The curves in Figure 7 look similar to those of an integrate-and-fire neuron driven by a net depolarizing current. However, the mechanism is very different. Here, the firing is highly irregular because it is driven by the variance of the membrane potential. In an integrate and fire neuron with constant depolarizing current, the firing is very regular.

The results of one such simulation are presented in Figure 8. In this simulation, the percentage of shared inputs was 10% and the initial firing rate was 50 spikes/sec. Twenty randomly chosen spike trains, out of the 6000 produced in each layer, are shown. The uncorrelated firing with a rate of 50 spikes/sec seen at the input layer (layer 1) builds up rapidly toward 90 spikes/sec, and correlations appear.

Transmission of firing rates in feedforward networks

To test whether the discrete-PSP discrete-time model can transmit the input firing rate in a multilayered network, we repeated

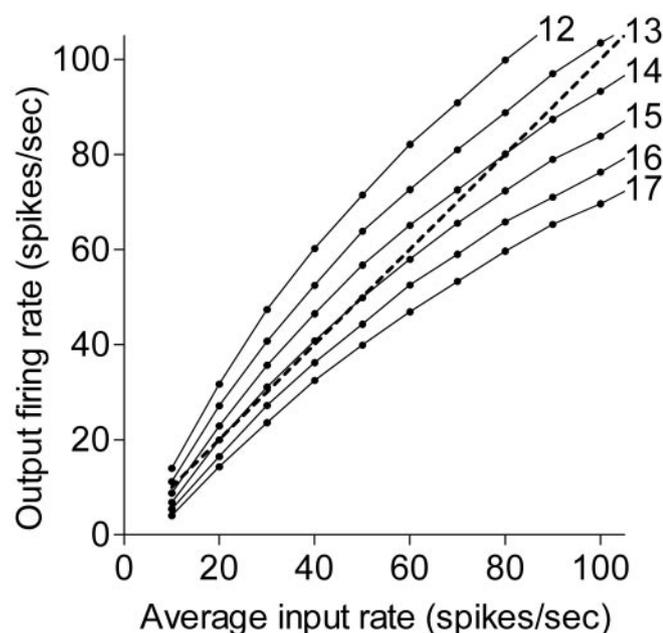


Figure 7. The gain of a discrete-time discrete-PSP model neuron is a nonlinear function of the input rate and depends on firing threshold. The model neuron received 300 excitatory and 300 inhibitory inputs and was simulated for 300 sec. The membrane time constant was 20 msec, and the lower reflection barrier was -1 . Firing threshold was varied between 12 and 17 (numbers at right). The curves are more similar to those of Figure 5 than to Figure 1.

simulations for a network as in Figures 2–4 but with the discrete-time neuron.

In the first group of simulations, we used a set of precisely balanced connection matrices in which each neuron at each layer received exactly 300 excitatory and 300 inhibitory inputs. This guaranteed a precise balance of excitation and inhibition.

When the threshold was set at 15, which is the optimal value for a gain of unity in a single neuron, and the neurons shared 10% of their inputs, activity declined quickly to zero for all initial rates. We next examined whether we could preserve firing rate by fine-tuning the threshold. The results (with 10% shared connections and an input firing rate of 50 spikes/sec) are summarized in Figure 9A. The threshold of the model neurons was varied between 10 and 15. Decay of firing rates was observed for threshold values >12 . The decay started after an initial increase in the firing rate in the first three to four layers. For threshold values of <12 , the rate stabilized at a constant value after the initial increase. The smaller the threshold for spike firing, the higher the final firing rate was.

How does the stable firing rate at deep layers correspond to the values of the input rate? Figure 9B shows that, for a firing threshold of 12, the firing rates at deep layers converged to the same stable fixed values (fixed point of the dynamics) of ~ 90 spikes/sec for all the initial conditions except for the lowest (10 spikes/sec). This situation is similar to the trend observed in the continuous-time discrete PSP model (Fig. 3). As in Figure 3, averaging over many (600) neurons for a short time span (100 msec) already does not produce accurate rate estimation, even after three to four layers, well before the mean layer rates converge to the asymptotic fixed-point value. This is partly because, at high initial rates, the mean layer rates approach each other well before they converge to the fixed point and partly because of the growth of the estimation error bars as a result of the buildup of correlations.

Figures 3 and 9A indicate that at low input rates, the layer rates do not converge to the fixed point obtained with high input rates. Does this mean that rate transmission is possible in this system at

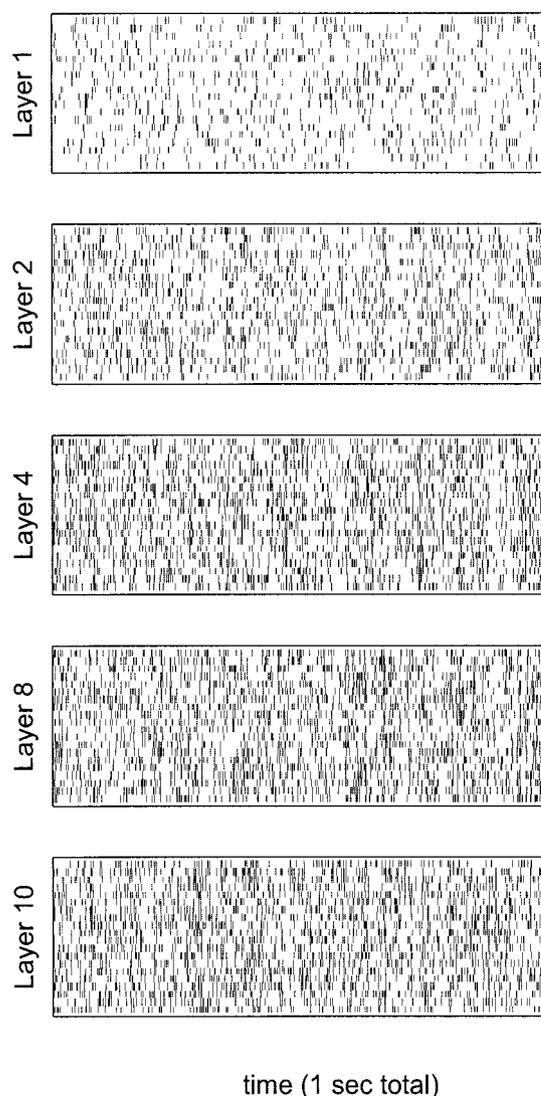


Figure 8. Dot displays for a network of discrete-time discrete-PSP neurons. One second of activity of 20 randomly selected neurons from the input stream (layer 1) and several layers along the network is shown. Each layer consisted of 6000 neurons; the percentage of shared inputs between two neurons at a given layer was 10%. Each neuron received 300 excitatory and 300 inhibitory inputs. The same connection matrix was used for all the layers. This is the same architecture as in Figure 2. Trains in the input layer (layer 1) are random, uncorrelated spike trains with an average rate of 50 spikes/sec. Threshold was set to 12, at which the firing rates changed only moderately. Firing rate buildup and synchrony are obvious. Synchrony is expressed both by bands of dense and sparse firing and by short vertical lines of precise synchrony.

low rates? To answer this question, we simulated networks with input rates of 5–25 spikes/sec, in steps of 5 spikes/sec. For the initial firing rate of 5 spikes/sec, the firing rate decayed to zero after a few layers. For the other input rates, the network switched between periods of silence and periods of high firing rates. Detailed analysis showed that the states of the network for different input rates differed in number and length of silent periods and not in the firing rates at periods when neurons do fire (Fig. 10, top).

We explain this behavior by suggesting that the network shifted constantly between two stable fixed points. The lower fixed point has a firing rate equal to zero, and the higher fixed point has a high firing rate whose value does not depend on the input to the network. The percentage of time spent at each fixed point determines the average firing rate in each case. Evaluating the firing rate of the network, given such behavior, requires aver-

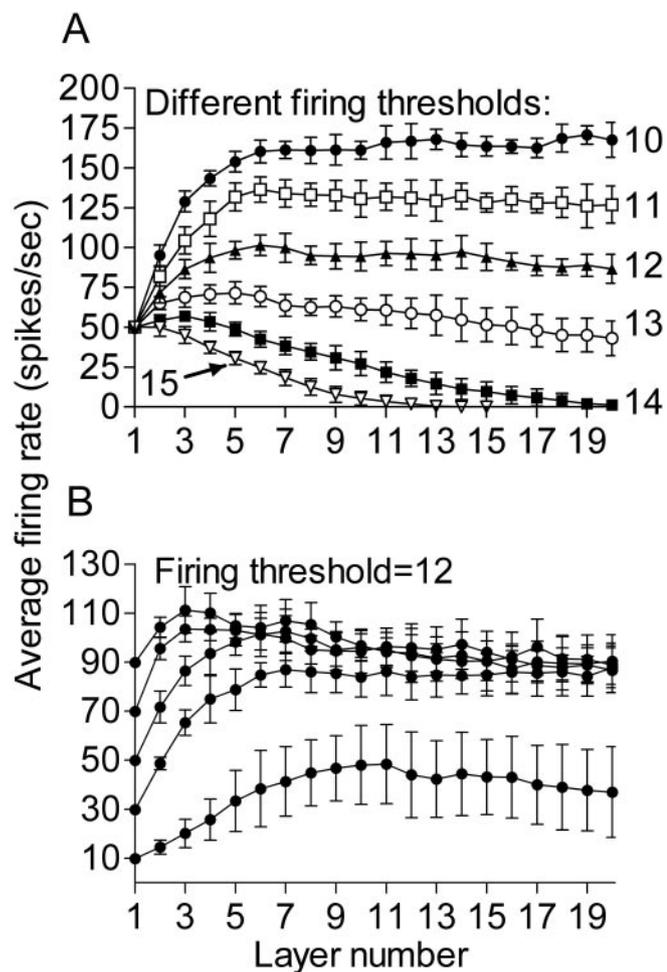


Figure 9. Dependence of model behavior on firing threshold. Simulations had the same network architecture as in Figure 8. *A*, Firing threshold was varied between 10 and 15, with initial firing rate of 50 spikes/sec in all cases. For threshold values <15 , the rate increased in the first two to three layers and then either stabilized (for thresholds 10, 11, and 12) or decayed to 0 (for thresholds 13, 14, and 15). Note that, for a threshold equal to 15, the rate does not change significantly between the first and second layers, which is consistent with the single-neuron simulation shown in Figure 7. Other parameters are as in Figure 8. *B*, Initial firing rate varied between 10 and 90 spikes/sec, with a firing threshold of 12 in all cases. Error bars were computed as described in Figure 3.

aging over a very long time. Because all the neurons switch states together, averaging over large number of neurons rather than over a long time does not help reduce errors.

Networks with imprecise balance

The scenario with exactly the same number of inhibitory and excitatory inputs to every neuron is not easy to achieve in biology. We tested this assumption by running two types of simulations. In the first, simulation of the connectivity between layers was probabilistic. In the second, exactly half of the inputs were inhibitory, but they were not synchronized with the excitatory inputs.

In the first group of simulations, we tried to achieve a more realistic condition in which the connections matrix was created randomly with a given probability for a contact between neurons in two subsequent layers; excitatory and inhibitory neurons were chosen randomly, with a probability of 0.5. The results of these simulations show that the changes in the mean firing rates between subsequent layers were not monotonous. Fluctuations with an amplitude of tens of spikes/sec were observed between

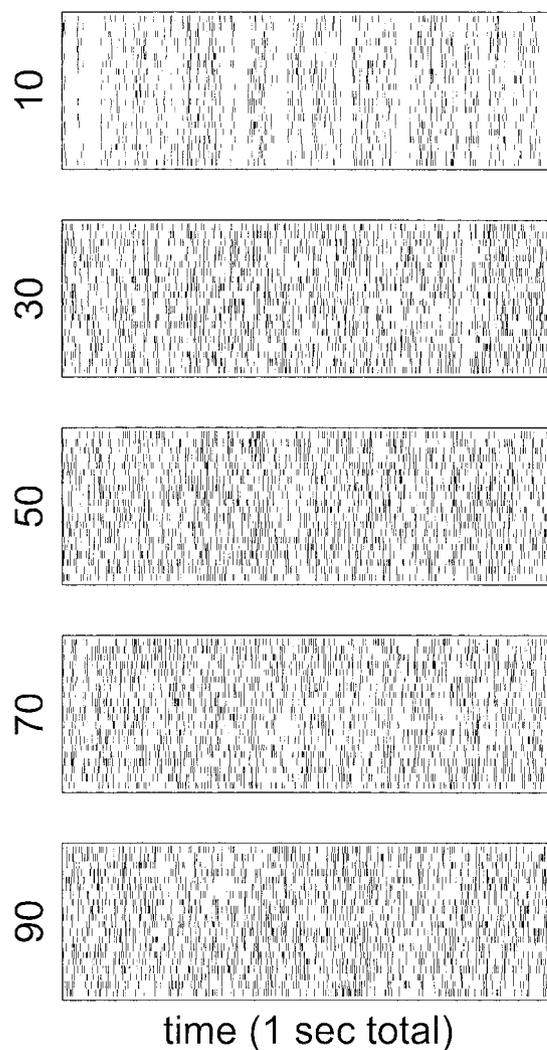


Figure 10. Firing rate in deep layers is essentially independent of initial input rate. Dot displays of 20 randomly chosen neurons from layer 10 for the network shown in Figure 9*B*. Numbers to the left of displays denote the initial firing rate in spikes/sec.

subsequent layers. Eventually, at deep layers, activity was either eliminated completely in the case of high thresholds (Fig. 11*A*) or it underwent broad fluctuations with a firing rate that is independent of the initial rate for low thresholds (Fig. 11*B*). These huge rate fluctuations were induced by the violations of the exact balance between excitation and inhibition caused by random choice of inhibitory and excitatory neurons.

Development of synfire waves in networks with excitation and inhibition that are not timed precisely with each other

Previous studies (Abeles, 1991; Herrmann et al., 1995; Diesmann et al., 1999) have shown that in long-feedforward networks, waves of synchronous activity appear and propagate in a stable manner through many layers. Networks in which these synfire waves were observed differ from the balanced network used above in several respects. Our purpose in this section is to pinpoint the primary factor that differentiates networks in which synfire activity can develop under a wide range of conditions and networks in which synfire activity is unstable.

Shadlen and Newsome (1998) argued that synfire waves develop only in networks with sparse connectivity, a high percentage of shared inputs, and low firing rates. They claim that in

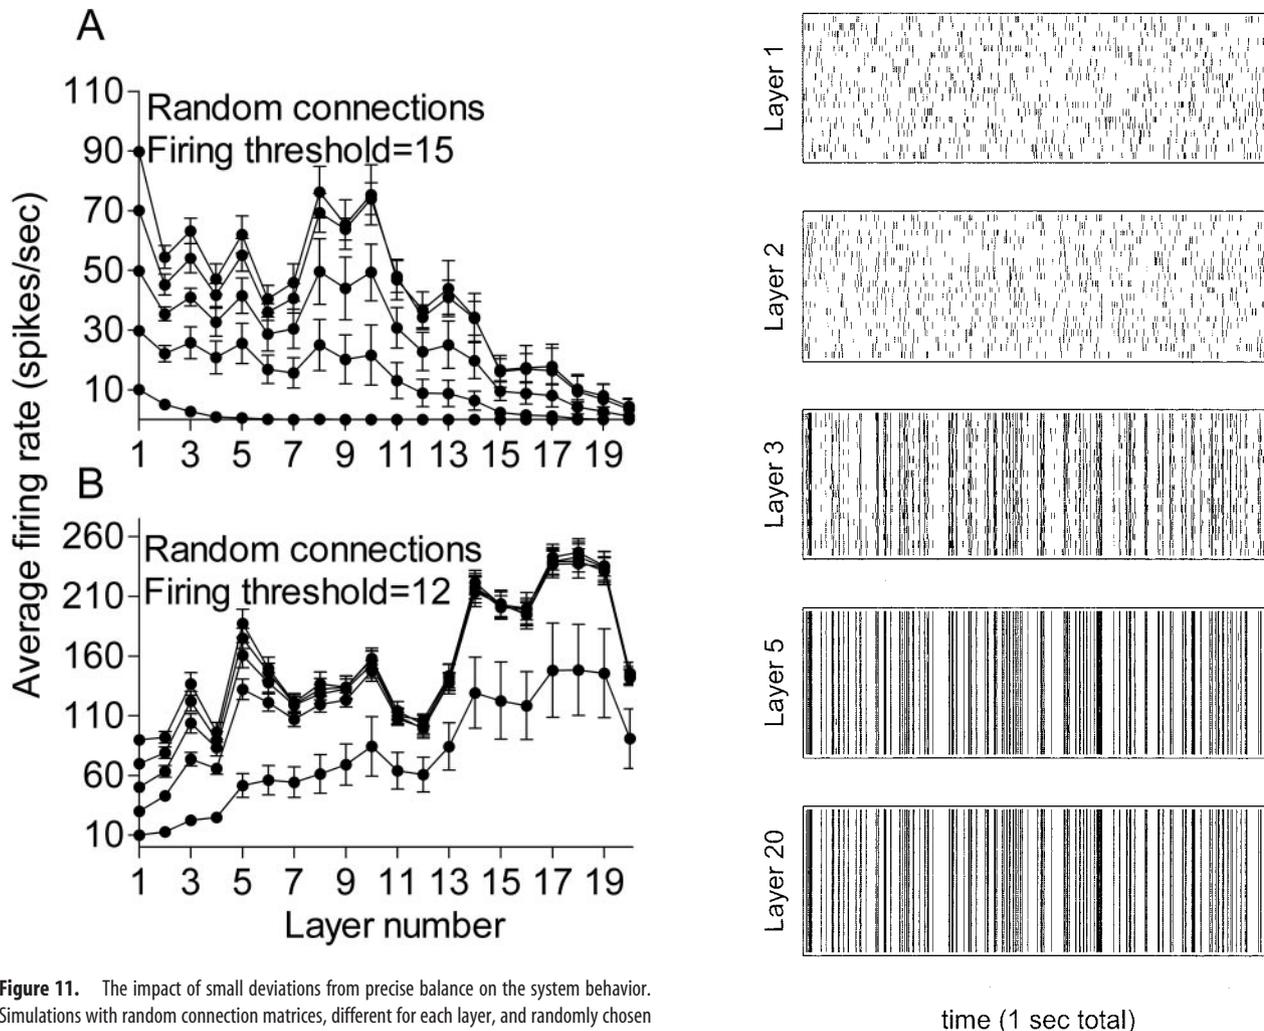


Figure 11. The impact of small deviations from precise balance on the system behavior. Simulations with random connection matrices, different for each layer, and randomly chosen inhibitory neurons (with a probability of 0.5). *A*, Simulation with the same single-neuron parameters as in Figure 9*A* (threshold 15). Decay of the rate is observed here also, but it is not monotonous. *B*, Simulation with threshold 12. The firing rate does not decay.

networks in which each neuron receives multiple inputs at any given time (“high-input regime”) and the percentage of shared inputs is <40%, no substantial synchrony can develop. To test this hypothesis, we performed a simulation in which inhibitory inputs had the same average rate as the excitatory inputs but were not synchronized precisely with each other. To achieve this, the spike trains that were chosen as inhibitory at each layer were replaced with random spike trains, with an average rate equal to the rate of the excitation. The results of this simulation can be seen in Figure 12, in which only the excitatory spike trains are shown. Synfire waves were formed after a few layers, interleaved with periods of no activity. These waves propagated stably for any number of layers tested. Thus, the precise synchronization of the excitatory and inhibitory inputs with each other, rather than the high-input conditions, is the reason for the stability of the partial synchrony in the feedforward network.

Discussion

In this study, we have shown that in feedforward networks with an exact balance between excitation and inhibition, it is difficult to transmit the population firing rate faithfully through many layers. Thus, the idea that a balance between excitation and inhibition in a feedforward network accounts for randomness of firing time and lack of synchrony and is highly problematic for

Figure 12. Emergence of precisely synchronized waves. Spike trains of inhibitory neurons produced by the model at each layer were replaced by random spike trains with average firing rates of the excitatory inputs. This preserved the balance of firing rates between excitation and inhibition but prevented the two types of inputs from synchronizing with each other. In this simulation, precisely synchronized waves of activity formed after three to four layers and propagated stably for any number of layers tested.

transmitting rates. This is in marked contradiction to the conclusions of Shadlen and Newsome (1998), which are based primarily on studies of only one layer. There, the problems of convergence to a fixed point, the buildup of synchrony, and the inability to distinguish among different input firing rates in a short time span are minor. Here, we show that in the full-layered network, severe problems for rate transmissions appear. They are associated with (1) single-neuron input–output properties (2) the dynamics of the mean layer rates, (3) the buildup of rate variances, and (4) the sensitivity to deviations from balance conditions. These issues are discussed below.

Single-neuron input–output properties

The starting point of the Shadlen and Newsome model is a single neuron with a gain that is close to unity. Our results show that at the level of a single neuron, such a gain is extremely hard to achieve for plausible single-neuron model parameters. Even if an approximately linear input–output relationship is obtained (Fig. 1), achieving a gain of unity requires extreme fine-tuning of model parameters.

Dynamics of mean rates

A faithful rate transmission in the feedforward model requires that the mean rate of the layers will remain roughly the same as the input rate; namely, that the input–output gain of rate for the whole system will be close to unity. Our results show that for the neuron model used by Shadlen and Newsome, it is possible to achieve approximately a unity gain for a system with <10 layers (Fig. 3). For longer chains, the layer rates converge to a fixed point, independent of the input rates (except for low-input-rate regimes).

The slow convergence of the layer rate to a common fixed-point value is probably a result of the peculiarity of the continuous-time discrete-PSP model, in which even an infinitesimal time difference between excitatory and inhibitory inputs makes a big difference in the chances of hitting threshold. Comparison of Figure 1 with Figure 5 and examination of Figure 6 show the huge difference between this model and the more realistic model with gradually rising PSPs.

Maintaining reasonable mean rates across the layers required fine-tuning of the single-neuron gain. Changing the single-neuron threshold by <10% induced either a fast decay of the rates or rapid growth to unrealistically high levels. Intuitively, one might think that the optimal value for rate transmission is a gain of unity. Indeed, this was the rationale behind the Shadlen and Newsome model. However, we have shown (Fig. 9) that with this choice, the activity at deeper layers decays to zero because of the partial synchrony that develops in the network. This synchrony between either a pair of excitatory cells or a pair of inhibitory cells in a given layer increases the variance of the input of this layer to the neurons in the next layer, whereas the synchrony between excitatory and inhibitory cells decreases this variance. As we show in the Appendix, the net effect of the synchrony in a given layer is to reduce the variance of the input to the next layer. This will cause a decrease in the mean firing rate in successive layers. Thus, a gain larger than unity is required to maintain a persistent activity across the chain. In this case (Figs. 3, 9), the firing of each layer will settle into a state with a nonzero firing rate and a mild level of synchrony.

Buildup of rate variance

The feasibility of rate transmission depends not only on the propagation of the mean rates across the layers but also, critically, on the variance of these rates. We have shown here that even when the mean rates are transmitted faithfully, rate information is lost because the fluctuations of the population rates build up quickly even after only three to four layers. This results from the emergence of correlations between the rates of different neurons because of their common input. In our simulations, we have used 100 msec as the window of integration time for the rate estimation. To ensure accurate rate estimates, the fluctuations in population rates have to be reduced by at least a factor of 3, and for this, the window of integration time needs to be increased to ~1 sec. Thus, a simple spatial averaging of the activity of each layer will not transmit rate information faithfully, because the spatial averaging will not suppress the random fluctuations in the rates efficiently, as would occur in the uncorrelated case. Thus, we conclude that in a feedforward network, firing rates may be used as codes only for a small number of processing stages. Whether a more sophisticated decoding scheme can overcome this problem has yet to be studied (Yoon and Sompolinsky, 1998).

Requirement of precise balance

The scenario of both excitation and inhibition being fed forward from layer to layer is probably not physiological when dealing

with layers of neurons that reside in separate cortical areas. Transmission between cortical areas is excitatory. Thus, the balanced network cannot emulate rate transmission between cortical areas. Here, the simulations shown in Figure 12 are more appropriate. Even within a cortical column, the physiological adequacy of a scheme with identical inhibitory and excitatory neurons is questionable. The local axonal distribution of excitatory and inhibitory neurons is very different, as are the intrinsic neuronal properties (Thomson and Deuchars, 1994; Markram et al., 1998). In view of the relatively small percentage of inhibitory neurons and inhibitory synapses in most cortical regions, one may prefer to consider a pool of inhibitory neurons that receives excitation from many excitatory neurons in the region and delivers inhibition to both excitatory and inhibitory neurons, ignoring their layer membership. Under these conditions, when the inhibitory neurons are not part of the feedforward chain, the inhibition can balance the excitation without being timed precisely with it. This type of architecture leads to the realization of the third scenario of network behavior, the synfire chain, depicted in Figure 12.

These instances of synchronous firing propagating in a robust way from one layer to the next are, in fact, identical to the synfire waves first suggested by Abeles (1982) to account for experimentally observed precise firing patterns in recordings from monkey frontal cortex. The stability of such waves has also been confirmed by a number of theoretical studies (Abeles, 1991; Bienenstock, 1995; Herrmann et al., 1995; Diesmann et al., 1999), and synfire-like phenomena have also been observed recently *in vitro* by Reyes (2002).

This work has highlighted the difficulty of achieving neuronal variability through balance between excitation and inhibition in purely feedforward architecture. This situation should be contrasted with the balanced state in recurrent networks. As shown by van Vreeswijk and Sompolinsky (1996, 1998), in these networks, the balance between excitation and inhibition is generated by the internal feedback via the dynamic adjustment of the firing rates of the excitatory and inhibitory populations. Consequently, there is no need to fine-tune the connectivity parameters. Furthermore, the firing rates in the balanced state of the recurrent networks vary linearly with the rate of their external input. Thus, transmission of rates in a long chain of neuronal layers may be feasible if the layers possess appropriate lateral feedback.

Our results do not rule out completely the possibility of rate transmission in a strictly feedforward network. The question of whether there is still a theoretical possibility of maintaining rate transmission in feedforward networks requires additional modeling studies. Van Rossum et al. (2002) found that the input-rate to output-rate of a single-neuron model may be linearized while large variability is added to output timing by incorporating a bias and large membrane noise. In a network of such neurons, firing rates may be maintained with low time correlations. The study by van Rossum et al. differs from ours in several key aspects. First, to obtain their results, all neurons were injected with a Gaussian noise with positive mean. Fine-tuning of the noise parameters was necessary to achieve propagation of rate coding. Furthermore, in these tuned parameters, each neuron was firing in almost periodic manner (particularly in the first layers), as shown in their Figure 4A. This resulted in a low variability in the total spike counts, which is required for the decoding of the rate. These regular patterns of individual neurons are very different from the observed cortical activity. Both the Shadlen and Newsome model and ours were aimed at exploring propagation of rate code via firing activity patterns, which are highly irregular, as observed in

cortex. Another unrealistic feature in the van Rossum et al. model is the small numbers of neurons per layer and all-to-all connectivity between layers. Again, our model assumed a degree of connectivity and population size that mimic cortical architecture more realistically. How these differences will affect the van Rossum network compared with ours is an issue that needs additional detailed investigation.

Appendix

In this appendix, we show, using analytical methods, that the emergence of synchrony in a long feedforward network with balanced excitation and inhibition lowers the effective gain of the network layers. This phenomenon is the reason for decay of the rate in a feedforward network with single neurons that have a gain approaching 1 (Fig. 9). The input to a neuron during a single integration time is the sum of its excitatory and inhibitory inputs, as follows:

$$I = \sum_{i=1}^k x_i - \sum_{j=1}^k y_j \quad (3)$$

where x_i and y_j denote excitatory and inhibitory inputs from single presynaptic cells, respectively, and k represents the numbers of excitatory and inhibitory inputs, which are equal in the balanced model.

The mean of the input over time is zero as a result of the balance between excitation and inhibition, as follows:

$$\langle I \rangle = \left\langle \sum_{i=1}^k x_i - \sum_{j=1}^k y_j \right\rangle = \sum_{i=1}^k \langle x_i \rangle - \sum_{j=1}^k \langle y_j \rangle = kr - kr = 0 \quad (4)$$

where r represents the average firing rate.

The variance of the input is given by the equation

$$\begin{aligned} \langle I^2 \rangle &= \left\langle \left(\sum_{i=1}^k x_i - \sum_{j=1}^k y_j \right)^2 \right\rangle = \left\langle \left(\sum_{i=1}^k x_i \right)^2 - 2 \left(\sum_{i=1}^k x_i \right) \left(\sum_{j=1}^k y_j \right) \right. \\ &\quad \left. + \left(\sum_{j=1}^k y_j \right)^2 \right\rangle = \sum_{i=1}^k \langle x_i^2 \rangle + \sum_{i \neq j} \langle x_i x_j \rangle - 2 \sum_{ij} \langle x_i y_j \rangle \\ &\quad + \sum_{j=1}^k \langle y_j^2 \rangle + \sum_{i \neq j} \langle y_i y_j \rangle = k(r^2 + \sigma) \\ &\quad + k(k-1)(r^2 + c) - 2k^2(r^2 + c) + k(r^2 + \sigma) \\ &\quad + k(k-1)(r^2 + c) = 2k\sigma - 2kc \end{aligned} \quad (5)$$

where σ represents the variance of a single input and c represents the average covariance of pairs of inputs. We use the fact that all the presynaptic neurons are identical in their properties, and therefore, all the pairs of different presynaptic neurons covary in their rates in the same way. Note that for $c = \sigma$, the current vanishes. In this case, all neurons fire in full synchrony; hence, there is perfect cancellation of the excitatory and inhibitory currents at any time step.

The effect of synchrony on the transmission of firing rates by the network can be summarized as follows. The synchrony between either a pair of excitatory or inhibitory cells in a given layer increases the variance of the input of this layer to a neuron in the

next layer. Conversely, the synchrony between excitatory and inhibitory cells decreases the variance of their input to the next layer. The total number of pairs of excitatory and inhibitory cells in the input to each cell is $k(k-1)$, and the total number of excitatory-inhibitory pairs is k^2 . Hence, the net effect of the synchrony in a given layer is to reduce the variance of the input to the next layer, which will cause a decrease in the mean rate of this layer. As a result, the gain of the single neuron, when fine-tuned to unity in the absence of synchrony ($c = 0$) is insufficient to preserve the firing rate in the presence of synchrony. In a recent work by Salinas and Sejnowski (2000), it was claimed that when the correlations are uniform across the network, as is assumed here, the current variance is unaffected by the correlations, because the correlations within the two populations cancel exactly the contribution of the correlations between the two populations. Here, we show that this cancellation is not exact but rather is valid only on the order of k^2 contribution to the current variance. Conversely, the term that is linear in the connectivity k is nonzero and contributes negatively to the current variance. For this reason, the layer rates decay quickly to zero unless the single-neuron gain is chosen to be greater than unity to compensate for the effect of correlations, as shown in Figure 9.

References

- Abeles M (1982) Local cortical circuits. New York: Springer.
- Abeles M (1991) Corticonics: neural circuits of the cerebral cortex. Cambridge, UK: Cambridge UP.
- Bienenstock E (1995) A model of neocortex. *Network* 6:179–224.
- Diesmann M, Gewaltig MO, Aertsen A (1999) Stable propagation of synchronous spiking in cortical neural networks. *Nature* 402:529–533.
- Herrmann M, Hertz J, Prugel-Bennett A (1995) Analysis of synfire chains. *Network* 6:403–414.
- Markram H, Gupta A, Uziel A, Wang Y, Tsodyks M (1998) Information processing with frequency-dependent synaptic connections. *Neurobiol Learn Mem* 70:101–112.
- Rall W (1967) Distinguishing theoretical synaptic potentials computed for different soma-dendritic distribution of synaptic inputs. *J Neurophysiol* 30:1138–1168.
- Reyes AD (2002) Synchrony-dependent propagation of signals in feedforward networks *in vitro*. *Soc Neurosci Abstr* 28:344.7.
- Salinas E, Sejnowski TJ (2000) Impact of correlated synaptic input on output firing rate and variability in simple neuronal models. *J Neurosci* 20:6193–6209.
- Shadlen MN, Newsome WT (1994) Noise, neural codes and cortical organization. *Curr Opin Neurobiol* 4:569–579.
- Shadlen MN, Newsome WT (1998) The variable discharge of cortical neurons: implications for connectivity, computation and information coding. *J Neurosci* 18:3870–3896.
- Softky WR (1995) Simple codes versus efficient codes. *Curr Opin Neurobiol* 5:239–247.
- Softky WR, Koch C (1993) The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci* 13:334–350.
- Stevens CF, Zador AM (1998) Input synchrony and the irregular firing of cortical neurons. *Nat Neurosci* 1:210–217.
- Thomson AM, Deuchars J (1994) Temporal and spatial properties of local circuits in neocortex. *Trends Neurosci* 17:119–126.
- Thorpe SJ, Fabre-Thorpe M (2001) Neuroscience: seeking categories in the brain. *Science* 291:312–316.
- van Rossum MCW, Turrigiano GG, Nelson SB (2002) Fast propagation of firing rates through layered networks of noisy neurons. *J Neurosci* 22:1956–1966.
- 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.
- Yoon H, Sompolinsky H (1998) The effect of correlations on the Fisher information of population codes. In: *Advances in neural information processing systems*, Vol 11, pp 167–173. Cambridge, MA: MIT.