A. Aertsen V. Braitenberg (Eds.)

Experiments and Theory



Springer-Verlag



Ad Aertsen Valentino Braitenberg (Eds.)

Information Processing in the Cortex

Experiments and Theory

With 102 Figures

Springer-Verlag Berlin Heidelberg New York London Paris Tokyo Hong Kong Barcelona Budapest Dr. Ad Aertsen Universität Bochum Institut für Neuroinformatik Postfach 102148 Gebäude ND 04 W-4630 Bochum, FRG

Professor Dr. Valentino Braitenberg Max-Planck-Institut für biologische Kybernetik Spemannstraße 38 W-7400 Tübingen, FRG

ISBN 978-3-642-49969-2 ISBN 978-3-642-49967-8 (eBook) DOI 10.1007/978-3-642-49967-8

Library of Congress Cataloging-in-Publication Data. Information processing in the cortex: experiments and theory / edited by Ad Aertsen and Valentino Braitenberg. p. cm. Includes bibliographical references and index. ISBN 978-3-642-49969-2 (U.S.)

1. Cerebral cortex. 2. Human information processing. 3. Cognitive neuroscience. I. Aertsen, Ad (Adrianus), 1948-. II. Braitenberg, Valentino. QP383.I54 1992 612.8'25-dc20 92-19742 CIP

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer-Verlag. Violations are liable for prosecution under the German Copyright Law.

© Springer-Verlag Berlin Heidelberg 1992 Softcover reprint of the hardcover 1st edition 1992

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Typesetting: Camera ready by Margarete Ghasroldashti, MPI, Tübingen 31/3145-5 4 3 2 1 0 - Printed on acid-free paper

Dynamics of Activity in Biology-Oriented Neural Network Models: Stability at Low Firing Rates

Michael Erb¹ and Ad Aertsen²

Max-Planck-Institut für biologische Kybernetik, Spemannstrasse 38, 7400 Tübingen, Germany

We investigated the dynamics of activity in feedback neural network models at low firing rates. The networks were designed to capture the typical features of real cortical networks. Stability analysis of the linearized model and simulations of different degrees of complexity show that stability is only obtained for very fast and sufficiently strong inhibition; otherwise the network activity develops into synchronous oscillations with frequency and amplitude dynamics governed predominantly by the inhibition parameters, but largely independent of (1) the network architecture (uniform, random or structured), (2) the spiking or analog nature of the neural activity, and, albeit to a lesser extent, (3) the linear or nonlinear nature of the neural threshold function. Provided the network connectivity is sufficiently rich and structured, the spike activity exhibits features which resemble those observed in physiological recordings from various cortical areas: cell assembly behaviour with different, simultaneous correlation dynamics (event coherence and rate coherence).

Introduction

In recent years, many different kinds of neural network models have been developed. Addressing a variety of different levels of biological reality, these models occupy a correspondingly wide spectrum of

¹ present address: Fak. für Physik, Abt. Biophysik, Philipps-Universität, Renthof 7, 3550 Marburg, Germany

² present address: Institut für Neuroinformatik, Ruhr-Universität, P.O. Box 102 184, 4630 Bochum, Germany

mathematical formalizations. They range from sets of coupled differential equations, describing the time course of membrane potentials of interconnected individual neurons, to spin-glass models from statistical physics, dealing with the interactions in populations of binary elements. In the context of a more comprehensive investigation aimed at connecting these different levels of modeling (Erb 1991), we investigated the activity dynamics in physiology-oriented networks of model neurons, both of the spiking and of the analog type. In the present report we focus on the issue of stability of the neuronal activity in a highly interconnected model network, in particular on its dependence on various neuron and network parameters. The model was explicitly designed to capture the typical features of real cortical networks, both in terms of anatomy and physiology. An important constraint in this respect is that of 'sparse' firing, i.e. the activity in the network is required to be of low firing rate, such as typically observed in the neocortex (e.g. Abeles et al. 1990). As will be demonstrated, obtaining stability in such a feedback system of sparsely firing neurons is by no means a trivial problem. In fact, stable solutions could only be obtained for a very confined 'island' in parameter space. Outside this restricted range, the system develops instabilities, certain features of which are reminiscent of phenomena observed in recordings of neural activity in the real cortex.

The network model: formulation and simulation

In order to investigate stability neural networks with feedback, we studied the behaviour of a network of 100 spiking model neurons with fixed synaptic connections (Erb et al. 1990). Figure 1a shows the scheme of the network, the corresponding Equations are given in Figure 1b. In this model the neurons are connected only by excitatory synapses, inspired by neuroanatomical findings that about 90% of the cortical synapses is of this type (Braitenberg and Schüz 1991; Schüz, this Volume). The synapses are modeled as lowpass filters with delayed response, transforming the incoming spike activity into EPSP's. These are summated linearly over the dendritic tree to yield the instantaneous value of the membrane potential at the cell body. The probability of spike generation is modeled by a sigmoid function of this membrane potential; this probability, in turn, is modulated by a refractory mechanism, driven by the recent spike history of the neuron. The final output, a sequence of spikes, is obtained by using the firing probability as the instantaneous rate of a stochastic event generator.

In addition to this basic network of spiking neurons, the 'pyramidal cells', the model comprises an inhibitory mechanism for stabilization. This inhibition consists of two parallel, linear branches: a fast one, with a time constant in the range of that for the excitatory synapses (5 ms), and a slow one with a considerably larger time constant (200 ms). The fast branch is intended to mimic the inhibitory action of the 'stellate cells', the slow branch serves to regulate the overall activity in the network towards a preset value (threshold control; Braitenberg 1978; Palm 1982).

In the present study the overall activity in the net was set to an average firing rate in the order of 20 spikes per second per neuron. Unless otherwise indicated, all numerical calculations as well as the network simulations shown here were performed with the following standard (sub)set of parameter values:

$\tau_0=4$ ms	$\tau_1=5 \text{ ms}$	$\tau_2 = 200 \text{ ms}$	
$\Delta_0=1$ ms		$\Delta_2=1$ ms	
$k_0 = 1$		k ₂ =3	
α=40	μ ₀ =0	$\mu_1 = 0.2$	c=0.3

The values of the remaining parameters, in particular D_1 and k_1 , were subject to variation, and are given separately in the appropriate Figure captions. The standard values were chosen primarily for reasons of physiological plausiblity. Moreover, in combination with appropriately chosen values for the remaining parameters, the standard set leads to a stable solution under sparse firing conditions, both for spiking and for analog neurons (cf. Figs. 2c-5c; Erb 1991).

One of the striking findings from simulations of this model is that stability at low firing rates appears to be reached only within an extremely confined region of the parameter space. Outside this stable 'island' the network activity either dies, explodes or develops into strong, coherent oscillations. This is illustrated in Figure 2, which shows the lowpass filtered sum activity of the network in response to a stepwise presentation of a random input pattern for three different sets of model parameters. In each of these cases, the connectivity of the network was described by the same random connectivity matrix, with





Figure 1b. Equations of the Model

Lowpass filters:

$$\dot{s}_{k}(t) = -\frac{1}{\tau_{in}} s_{k}'(t) + \frac{1}{\tau_{in}} s_{k}(t - \Delta_{in})$$
$$\dot{a}_{j}'(t) = -\frac{1}{\tau_{0}} a_{j}'(t) + \frac{1}{\tau_{0}} a_{j}(t - \Delta_{0})$$
$$\dot{\Theta}_{1}'(t) = -\frac{1}{\tau_{1}} \Theta_{1}'(t) + \frac{1}{\tau_{1}} \Theta_{1}(t - \Delta_{1})$$
$$\dot{\Theta}_{2}'(t) = -\frac{1}{\tau_{2}} \Theta_{2}'(t) + \frac{1}{\tau_{2}} \Theta_{2}(t - \Delta_{2})$$



Potentials:

$$\Theta_{1}(t) = k_{1} \cdot \sum_{j=1}^{n} a_{j}'(t)$$

$$\Theta_{2}(t) = k_{2} \cdot \left(\sum_{j=1}^{n} a_{j}'(t) - A_{0}\right)$$

$$d_{i}(t) = \sum_{j=1}^{n} c_{ij} \cdot a_{j}'(t) + \sum_{k=1}^{m} b_{ik} \cdot s_{k}'(t) - \Theta_{1}'(t) - \Theta_{2}'(t)$$

Nonlinear characteristic function:

Г

Event generator:

after spike at t₀

$$\mathbf{r}_{i}(t) = \begin{cases} 0 & \text{for } t - t_{0} < \Delta_{r} \\ 1 - e^{-\frac{t - t_{0}}{\tau_{r}}} & \text{for } t - t_{0} \ge \Delta_{r} \end{cases}$$

$$a_{i}(t) = \begin{cases} 1 & \text{with probability } P(1) = p_{i}(t) \cdot r_{i}(t) \\ 0 & \text{with probability } P(0) = 1 - P(1) \end{cases}$$

Fig. 1. Scheme (a) and Equations (b) of the model neural network

constant mean value (0.3) throughout the matrix and uniformly distributed stochastic fluctuations (range ± 0.1). Clearly the network is stable only in the last example (Fig. 2c), whereas the other two (Figs. 2a,b) exhibit various forms of instabilities, differing both in the rate at which the instability develops and decays and in the oscillation frequency.

Stability of linear networks with analog neurons and uniform connectivity matrix

For a quantitative investigation of the stability issue we introduced a number of simplifications. First, instead of actually generating discrete spikes by means of a stochastic event generator, we interpreted the firing probability of each neuron directly as its output activity. Thus, we obtain a deterministic network of so-called 'analog' neurons, i.e. neurons with continuously varying 'firing rates'. Secondly, in order to allow for a formal stability analysis, we studied a linear approximation of the model. Under these conditions it is possible to explicitly formulate the transfer function of the network; this transfer function obviously depends on the various neuron and network parameters. Finally, in order to obtain a more amenable expression for the transfer function, we drastically reduced the complexity of the network connectivity. In the simplified case of a uniform connectivity matrix, i.e. the limiting case (for fluctuations down to zero) of the random connectivity in Fig. 2, the transfer function is given by

$$G_g(s) = \frac{\frac{f'}{\tau_0 s + 1} e^{-\Delta_0 s}}{1 - \frac{f'}{\tau_0 s + 1} e^{-\Delta_0 s} \left(c - \frac{k_1}{\tau_1 s + 1} e^{-\Delta_1 s} - \frac{k_2}{\tau_2 s + 1} e^{-\Delta_2 s}\right)}$$

where f' is the slope of the nonlinear characteristic function f at the fixpoint.

Analysis of the pole positions of this transfer function yields a measure for the asymptotic stability of the system. In particular, the pole with the largest real part determines the amplitude dynamics of the network activity: it decides whether the effects of a disturbance will continue to grow or will eventually die out. Clearly, a positive real part signifies an increasing and, hence, unstable solution, whereas a negative real part corresponds to a decreasing, stable solution. Moreover, the imaginary part of this pole specifies the fine grain time course of the network response, in particular its oscillation frequency. Figure 3 shows three examples of such transfer functions, each one for a different set of model parameters. The logarithm of the absolute value of the transfer function, the 'gain' in dB, is coded according to a





Fig. 2. Lowpass filtered sum activity of the spiking network in response to an input step at t=0 of a random input pattern for three different sets of model parameters (a: $\Delta_1=0.1$ ms, $k_1=0.1$, $k_2=3$; b: $\Delta_1=1$ ms, $k_1=10$, $k_2=30$; c: $\Delta_1=0.1$ ms, $k_1=10$, $k_2=3$). Observe that the network is stable at low rate only in c, whereas both in a and b the network exhibits distinct instabilities, differing both in the time course with which the instability develops and decays and in the frequency of oscillation

grey scale; black spots mark the positions of the poles. For the pole with the largest real part, i.e. the one which lies most to the right, we calculated in each case the growth exponent from the real part, and the oscillation frequency f and period T from the imaginary part. The results are given in Figure 3.

In order to test the predictive value of this stability analysis, we also calculated the step responses of the linearized model for the same three sets of parameter values. The results are shown in Figure 4. Observe that in each case the behavior of the network activity in response to a step input - stable vs. unstable, temporal development of instability, oscillation frequency - agrees very well with the predictions from the pole analysis. In the first two examples, the transfer function has a pole with a positive real part (rightmost black spots in Figs. 3a,b). Accordingly, in both cases the network shows a clearly unstable step response, characterized by a vigorous oscillation (Figs. 4a,b). Observe also that, in accordance with the results of the stability analysis, the first step response (Fig. 4a) exhibits a slower amplitude rise and a higher frequency of oscillation than the second one (Fig. 4b). This reflects the fact that the corresponding pole in the first example (rightmost black spot in Fig. 3a) has a smaller positive real part and a larger imaginary part than its counterpart in the second case (Fig. 3b). In contrast to this, all pole positions in the third example have a negative real part (Fig. 3c). This results in a stable step response with rapidly decreasing amplitude of oscillation (Fig. 4c).

Using this approach of pole analysis, we investigated the influence of the various model parameters on network stability. Unfortunately, the dimensionality of the parameter space is quite high (cf. Fig. 1b), causing an exhaustive exploration to be out of the question. On the other hand, however, the number of truly free parameters, as well as their relevant ranges can be restricted considerably, by using anatomical and physiological plausibility as a constraint. Guided by this criterion, we made a number of exploratory tests in selected subdomains of the parameter space. As it turned out, stability could only be obtained for a very limited range of parameter values. The most crucial of these appeared to be the delay D_1 and the gain k_1 of the fast inhibitory pathway. These were investigated systematically at high resolution, while keeping the others fixed at physiologically plausible and potentially stable values (the standard set defined above). We found that under these circumstances, stability at low firing rate could only be achieved within a very confined 'stability island' (cf. Figs. 3c and 4c as

Figure 3. Logarithmic Plot (dB) of the Absolute Value of the Transfer Function for three Different Parametersets



Fig. 3. Transfer function of the linearized network for three different sets of model parameters (a: $\Delta_1=0.1$ ms, $k_1=40$; b: $\Delta_1=0.3$ ms, $k_1=10$; c: $\Delta_1=0.1$ ms, $k_1=10$). The logarithm of the absolute value of the transfer function ('gain' in dB) over a restricted portion of the complex plane is coded according to a grey scale: the larger the value, the darker the grey. Hence, black spots mark the positions of the poles. The growth exponent, the oscillation frequency f and period T, calculated from the real and the imaginary parts of the pole positions are indicated in the Figures. The 3-D plot (d) shows the gain for the same parameter combination as in c over an extended range of the complex plane and rotated over 90 degrees. This plot serves to illustrate the multitude and pairwise arrangement of the poles, with the increasing imaginary parts reflecting the various time delays

opposed to 3a,b and 4a,b): the inhibition must be of sufficient but not too great strength and, most critically, should be faster than the excitatory influences (e.g. a 0.1 ms delay for inhibition as compared to 1 ms for excitation). Also, the slope at the operating point of the sigmoid nonlinearity must lie within a certain range; other parameters appear to play a less significant role.



Fig. 4. Lowpass filtered sum activity of linearized network in response to a uniform step input for the same three sets of parameter values as used in Fig. 3. Periods and corresponding oscillation frequencies, directly measured from these step responses are a: T=1.81 ms, f=553 Hz; b: T=2.64 ms, f=378 Hz; c: T=1.67 ms, f=600 Hz. Notice the good correspondence of these values with the corresponding ones in Fig. 3

Figure 4. Simulation of the Linear Model

Stability of nonlinear networks

Clearly, a linear stability analysis can only be expected to reflect the local network behaviour, i.e. the activity dynamics in the vicinity of the fixpoint. For a study of the global dynamics of the network activity, it is necessary to take into account the influence of higher order terms, which come into play because of the sigmoid threshold function of the pyramidal cells (cf. Fig. 1).

After incorporating the non-linear characteristic function into the model network, we again calculated the time course of the network activity in response to an input step. The results are shown in Figure 5 for the same three sets of parameter values used in Figures 3 and 4. Surprisingly, also the behaviour of the nonlinear solutions complies, at least qualitatively, with the predictions from the linear stability analysis. Again, only the third example (Fig. 5c) shows the decay that is indicative of a stable system, whereas the other two (Figs. 5a,b) once more exhibit an unstable, progressively growing oscillation. Observe also, however, that the introduction of the nonlinearity gave rise to a kind of temporal scaling of the step responses: both the rise times of the envelopes and the oscillation frequencies, while approximately retaining their relative magnitude, have clearly grown as compared to their linear counterparts (notice the different time scales in Figs. 4 and 5). In addition, one observes how the amplitude of the oscillations, in contrast to the linear case, is bounded by the saturating nature of the nonlinearity. Moreover, the oscillations exhibit a typical nonlinear distortion, the shape of which is determined by the position of the fixpoint; in our case it is situated in the expansive part of the nonlinearity, due to the imposed constraint of low firing rate (see Erb 1991 for a more elaborate discussion of the nonlinear effects).

Stability of networks with non-uniform connectivity: the emergence of cell assemblies

Thus, an analysis of the pole positions of the transfer function for the simplified model network enables us to make testable predictions regarding the stability of the network, both in the linear and the nonlinear case. The question arises to what extent these predictions remain valid for networks with a more interesting connectivity, in which not all connections are of the same strength. In order to answer this question, we will continue to alleviate, step by step, the simplifying assumptions we made when reducing the full model (Fig. 1) to the linearized transfer function (Equ. 1).



Figure 5. Simulation of the Nonlinear Model

Fig. 5. Lowpass filtered sum activity of the nonlinear network in response to a uniform step input for the same three sets of parameter values as used in Figs. 3 and 4. Observe the qualitative correspondence in the behaviour of the linear and nonlinear solutions: stable vs. unstable, relation between oscillation frequencies and time constants of build-up and decay of oscillation. Notice also the new features introduced by the sigmoid nonlinearity: temporal scaling (compare the different time scales in Figs. 4 and 5), bounded amplitude and distorted waveform of the oscillations.

Deterministic networks of analog neurons

We addressed the issue of nonuniformity of connections by two different approaches. First, we returned to the original, random connectivity matrix with constant mean value and uniformly distributed stochastic fluctuations around it (cf. Fig. 2), and gradually increased the fluctuation range. Analysis of the step responses (not shown here) demonstrated that, for a nonlinear network with mean connectivity of 0.3 and fluctuations in the range of %0.1 (i.e. for a dynamic range in the connectivity of a factor of 2), the time course of the summated activity could hardly be distinguished from that in the uniformly connected network (Erb 1991). Moreover, the activity of the individual neurons followed very closely the time course of the summated activity for the entire net. Clearly, the averaging over the many different synaptic connections taking place in each neuron separately gives rise to very similar membrane potential trajectories. Thus, the variability among the activity profiles of different neurons is strongly reduced. As a consequence, under these conditions the individual dynamics and, hence, the mean of the dynamics closely follow the dynamics of the mean. This tendency towards homogeneity in large, highly interconnected networks only breaks down when the dynamic range of the connectivity strengths becomes larger and larger and/or when the connectivity matrix exhibits a clear structure.

In a second approach we investigated the activity dynamics for a structured network connectivity. To this end we used a connectivity matrix resulting from a study of associative memory and the performance of different types of learning rules (Erb 1985, 1991; Palm 1986, 1987). In particular, we used a fixed connectivity matrix in which were embedded the memory traces of a set of 10 randomly generated input patterns. The activity in response to a uniform input step for a nonlinear network with this structured connectivity matrix is depicted in Figure 6, again for the same three sets of parameter values as used in the foregoing. Each of the three Figures 6a-c shows the time course of the summated network activity (righthand panel), together with the individual neuron activities (superimposed traces in lefthand panel). Observe that, with the exception of Fig. 6c, the time course of the summated network activity is hardly discernible from the activity in the uniformly connected network (Fig. 5), which once more emphasizes the strong homogenizing effect of averaging.

More interestingly, however, the time courses of the single neuron activities now exhibit two different kinds of nonuniformities. In the first and second example (Figs. 6a,b), the network activity develops into a synchronous oscillation in which *all* neurons are partaking, each one, however, with a somewhat different amplitude as manifested by the broadening of the band of superimposed traces in the lefthand panels. We note that this coherent oscillation is imposed by the global inhibition, which itself is oscillating in counterphase. Interestingly, in

Figure 6. Simulation of the Nonlinear Model with a Learned Matrix



Time Course of Single Neuron Activity and Total Activity

Fig. 6. Dynamic behaviour of the activity in the nonlinear network with analog neurons and structured connectivity matrix. Response to a uniform step input for the same three sets of parameter values as used in Figs. 3 to 5. Time course of the lowpass filtered sum activity (righthand panels) and individual neuron activities (superimposed traces in lefthand panels). Observe the striking similarity of the summated activity with that in the uniformly connected network (Fig. 5). Notice also, however, the different kinds of nonuniformities in the time courses of the single neuron activities. In particular, Fig. 6c reveals the emergence of a cell assembly: the dynamic instability temporarily dissociates a sub-group, defined through it interconnections, from the network it is embedded in

the third example (Fig. 6c) after a transient, global oscillation in which all neurons participate (0-60 ms), the nonlinearity splits the net into two subgroups: a small group of neurons which increase their firing rates and develop a coherent oscillation, whereas the majority of neurons does not feel addressed by the stimulus and rapidly settles at a very low firing rate. Since the mean activity level in the 'ignited' subgroup (Fig. 6c) is higher than the overall level calculated for the entire network with uniform connectivity (Fig. 5c), the active subgroup develops an unstable oscillation, even at parameter values for which the uniform net was stable. Evidently this last case is the more interesting one from the point of view of information processing. The selective instability provides a mechanism which temporarily dissociates a particular subgroup, defined through it interconnections, from the network it is embedded in. This behaviour closely corresponds to Hebb's notion of a 'cell assembly' (Hebb 1949), and, hence, provides a functional mechanism for signalling the occurrence of some interesting constellation in the 'outside' world.

Stochastic networks of spiking neurons

In a final step, in order to return from the simplified description to the full model, we reinstated the stochastic description with spiking neurons. Using the same structured connectivity matrix as in the deterministic case (Fig. 6), we again measured the step responses. The results for the same three sets of parameter values are shown in Figure 7: the lowpass filtered spike trains for the entire network (righthand panels in Figs. 7a-c) and for the individual neurons (superimposed traces in lefthand panels in Figs. 7a-c and raster displays in Figs. 7df). Notice that the activity dynamics in these stochastic, spiking networks resemble those in the deterministic network very strongly. The time course of the summated activity, particularly in the first two examples, is quite comparable to that in Figure 6. In addition, we observe similar departures from homogeneity among the individual neuron activities as in the deterministic case. Either practically all neurons participate in a coherent global oscillation, albeit with varying amplitudes and differing degrees of synchrony (Figs. 7a,b,d,e) or, alternatively, after a transient oscillation (0-100 ms), the net splits up into separate subgroups: a small group of neurons increase their firing rates, while the majority rapidly settles at a very low rate (Figs. 7c,f).



Figure 7. Simulation of the Spiking Model with a Learned Matrix

Time Course of Single Neuron Activity and Total Activity

Fig. 7. Dynamic behaviour of the activity in the nonlinear network with stochastic, spiking neurons and structured connectivity matrix. Response to a uniform step input for the same three sets of parameter values as used in Figs. 3 to 6. Time course of the lowpass filtered sum activity (a-c: righthand panels) and individual neuron activities (a-c: superimposed traces in lefthand panels; d-f: raster displays). Observe that the activity dynamics in these stochastic, spiking neuron activities. In contrast to the analog network, however, the assembly activity (c, f) reveals a difference in coherence between spike events and firing rates, demonstrating that this class of networks may exhibit different

Thus, also for stochastically spiking neurons, widespread stimulation of a structured feedback network, after a short epoch of global oscillation, may induce a selective activation and, hence, a temporary dissociation of a particular subgroup from the network.

A more detailed inspection of the activity patterns developed by the analog and the spiking networks reveals a new feature, which is only exhibited by the spiking network. Interestingly, and in contrast to the analog network, the results in Figs. 7c (left) and 7f show that the spike trains of neurons in the activated cell assembly rapidly loose their coherence upon reaching their increased activity level. This fast modulation of spike correlation is not reflected in the firing rates of the neurons involved. These gradually build up and stay at a relatively high level, independently of the reorganization taking place in the precise timing of the individual events. As a consequence, the sum activity of the assembly, after a short epoch of synchronous oscillation following 'ignition', quickly decays to a noisy constant high level. This suggests that the spiking network may be inherently more stable than the analog network with otherwise unchanged parameters, and that when, for some reason, a subgroup instability does develop, it lasts for a rather short time. Indeed, results from additional simulations (not shown here) indicate that in spiking networks with sparse firing, the stability problem is less severe than it appears in the stability analysis of the linearized model with 'analog' neurons.

In addition, this difference in coherence between spike events and firing rates demonstrates that this class of networks may exhibit different correlation dynamics simultaneously, each of them characterized by its own time constant. The correlation among the activities of different neurons may show rapid synchronization of spike trains (*event coherence*) and, simultaneously yet more or less independently, slower co-variations of firing rates (*rate coherence*). Such different correlation dynamics have also been observed in physiological recordings from various cortex preparations (e.g. Krüger and Mayer 1990; Aertsen and Gerstein 1991; Vaadia et al. 1991; Aertsen et al. 1992). A more extensive discussion of their functional significance will be given elsewhere (Neven and Aertsen, in preparation).

correlation dynamics simultaneously, each of them characterized by its own time constant: rapid synchronization of spike trains (event coherence) and, simultaneously yet more or less independently, slower co-variations of firing rates (rate coherence)

Discussion

On the basis of a stability analysis of the linearized network model, we found stability at low firing rates only for a very restricted range of parameter values. The most crucial parameters appear to be the delay and the gain of the fast inhibitory pathway. In particular, one needs shorter delays in this pathway than in the excitatory connections (e.g. 0.1 ms delay for inhibition vs. 1 ms for excitation). Such delays are incorporated implicitly in most neural network models, since updating of the membrane potential takes place only one simulation time step after the generation of spikes. Also the slope at the operating point of the sigmoid nonlinearity must be in an adequate range. This can be achieved from outside the network by adjusting either the mean input strength or the set value for the slow inhibition. Other parameters appear to play a less significant role.

The dynamic behaviour of networks consisting of neurons with a nonlinear threshold function was qualitatively comparable to that of the linearized model, indicating that the nonlinearity does not play a crucial role as far as stability is concerned. The same holds for the influence of the connectivity matrix, as long as its dynamic range remains withing certain bounds. Even when the variation among the connection strengths becomes large, the sum activity of the network still basically follows the general pattern described for the simplified model network. Summarizing, we observe very similar dynamics of the summated network activity in all cases studied: a stable solution for a very confined 'island' in parameter space, characterized primarily by very fast and sufficiently strong inhibition, and in the majority of cases a synchronous oscillation with frequency and amplitude dynamics governed predominantly by the inhibition parameters, but largely independent of (1) the network architecture (uniform, random or structured), (2) the spiking or analog nature of the neural activity, and, albeit to a lesser extent, (3) the linear or nonlinear nature of the neural threshold function. Finally, simulations of networks with stochastic, spiking neurons suggest that in such networks the problem of stability at low firing rates may be less severe than in networks of deterministic, 'analog' neurons.

Our finding that stability in sparsely firing feedback networks requires a considerably shorter delay of inhibition than of excitation is in agreement with observations by Abeles (1991), as well as with

results of a number of theoretical studies (e.g. Marcus and Westervelt 1989; Amit and Treves 1989; Rubin and Sompolinsky 1989; Golomb et al. 1990). Moreover, it fits well with the notion that inhibition acts locally, i.e. directly on the cell body or on proximal dendrites, whereas excitation is projected onto the dendritic trees, and is more prominent in (long-range) cortico-cortical connections, thus involving longer delays (Braitenberg and Schüz 1991; Schüz, this Volume). Our study further demonstrates that, provided the network connectivity is sufficiently rich and structured, subgroups of neurons may be briefly dissociated from the net by the temporal coherence of the activity patterns that develop upon stimulation. This behaviour is precisely what was predicted by the 'cell assembly' hypothesis (Hebb 1949; Gerstein et al. 1989) and the concept of the 'synfire chain' (Abeles 1982). Moreover, it conforms to a number of related physiological observations (e.g. Eckhorn et al. 1988 and this Volume; Gray and Singer 1989a,b; Aertsen and Gerstein 1991; Vaadia et al. 1991). Finally, the different simultaneous correlation dynamics exhibited by the spiking network model, event coherence and rate coherence, also find their physiological counterpart in observations on cortical recordings (Krüger and Mayer 1990; Vaadia et al. 1991; Aertsen et al. 1992).

One of the salient conclusions of our study was that the task of 'tuning' a feedback network to a stable, low firing rate proved to be quite difficult. This was reflected in the unsolicited experience that, before availing ourselves of the guidance of the stability analysis, we spent impressive amounts of time searching for appropriate parameter settings that would keep the network from oscillating. At this point it should be stressed that none of our network components has any intrinsic tendency to oscillate; the interactions in the network fully suffice to make the network go into oscillation, as anyone who ever tried to build an amplifier will appreciate. In view of this observation, the question how the real cortex, assuming that it works along similar principles, solves this inherent stability problem becomes only more pertinent. It is conceivable that not all parts of the cortex manage this problem equally well, or that the solutions that evolved in different species are not all equally efficient in coping with the unavoidable variations of 'network parameters'. In this respect it is tempting to speculate that experimentally observed oscillations in some physiological cortical networks, e.g. in cat visual cortex (Gray and Singer 1989a,b; Eckhorn et al. 1988 and this Volume) might possibly be instances of systems which failed to solve this stability problem adequately, especially when confronted with effective stimuli and, hence, correspondingly higher levels of input activity. If this were indeed the case, the question might be not so much how the brain functions by virtue of oscillations, as most researchers working on cortical oscillations seem to assume, but rather how it manages to do so in spite of them.

When having to work with a system, which lives on the brink of instability and every once in a while develops partial, transitory 'explosions', another possibility comes to mind. If short, activityrelated epochs of instability of portions of the net are apparently unavoidable and bound to occur from time to time, one might as well use them for a convenient purpose. An obvious candidate for such an application, in fact, the one usually put forward in connection with cortical oscillations, is that of 'dynamical linking', in which temporal coherence among active neurons defines their short-lived functional association (e.g. von der Malsburg 1981, 1986). Certainly our finding that networks with strongly nonuniform connectivity exhibit short epochs of synchronized instability, with subgroups of neurons temporarily dissociated from the network (Figs. 6,7) would be consistent with this idea (see also Aertsen and Gerstein 1991; Aertsen and Preissl 1991). Also, the manifestation of two distinct correlation dynamics, event coherence and rate coherence, exhibited both by the spiking network model and in physiological recordings, appears to be relevant in this respect (Neven and Aertsen, in preparation).

Expanding along these lines, one could, in fact, rethink the original goal of striving towards a stable network with fixpoint dynamics and, instead, consider the information processing capabilities of networks with more complex, dynamic attractor states. Clearly the limit cycle attractor is the first one that comes to mind when faced with the oscillating nature of our network dynamics, but obviously there are other, more elaborate possibilities, such as dynamics governed by quasi-periodic 'strange' attractors (e.g. Skarda and Freeman 1987; Preissl and Aertsen, this Volume). Clearly, an essential requirement in this context would be that not the entire net partakes in the instability, since such a trivial global state would not leave any room for a selective coding or computational mechanism to establish itself. Rather, the membership of the 'ignited' subgroups should be selective and transient, preferably determined by the dynamic interplay of connectivity within the net and neural activity feeding into it. In this respect, both the assembly-like activity patterns described above and the rapid,

stimulus-dependent modifications of functional connectivity observed earlier (Erb et al. 1986, 1989; Aertsen and Preissl 1991) seem to support this as a viable alternative. A study, specifically aimed to elucidate this dynamic concept of neuronal computation and memory, is currently in progress.

Acknowledgements

We thank Günther Palm, Hubert Preißl and Stefan Rotter for many stimulating discussions. Partial funding for this project was received from the Bundesministerium für Forschung und Technologie (BMFT) and the Stiftung Volkswagenwerk (I/64605).

References

- Abeles M. (1982) Local Cortical Circuits. An electrophysiological study. Springer, Berlin Heidelberg New York
- Abeles M. (1991) Corticonics. Neural circuits in the cerebral cortex. Cambridge University Press, Cambridge, U.K.
- Abeles M, Vaadia E, Bergman H (1990) Firing patterns of single units in the prefrontal cortex and neural network models. Network 1:13-35
- Aertsen A, Gerstein G (1991) Dynamic aspects of neuronal cooperativity: Fast stimulus-locked modulations of 'effective connectivity'. In: Krüger J (ed) Neuronal Cooperativity. Springer, Berlin Heidelberg New York, pp 52-67
- Aertsen A, Preissl H (1991) Dynamics of activity and connectivity in physiological neuronal networks.
 In: Schuster H (ed) Nonlinear Dynamics and Neuronal Networks. VHC Verlag, Weinheim, pp 281-301
- Aertsen A, Vaadia E, Abeles M, Ahissar E, Bergman H, Karmon B, Lavner Y, Margalit E, Nelken I, Rotter S (1992) Neural interactions in the frontal cortex of a behaving monkey: Signs of dependence on stimulus context and behavioral state. J f Hirnforsch (in press)
- Amit DJ, Treves A (1989) Associative memory neural networks with low temporal spiking rates. Proc Natl Acad Sci USA 86:7851-7875

- Braitenberg V (1978) Cell assemblies in the cerebral cortex. In: Heim R, Palm G (eds) Theoretical Approaches to Complex Systems. Lecture Notes in Biomathematics, Vol. 21, Springer, Berlin Heidelberg New York, pp 171-188
- Braitenberg V, Schüz A (1991) Anatomy of the Cortex. Statistics and Geometry. Springer, Berlin Heidelberg New York
- Eckhorn R, Bauer R, Jordan W, Brosch M, Kruse W, Munk M, Reitboeck HJ (1988) Coherent oscillations: a mechanism of feature linking in the visual cortex? Multiple electrode and correlation analysis in the cat. Biol Cybern 60:121-130
- Erb M (1985) Computersimulation eines assoziativen Speichers als Modell für die Funktion der Großhirnrinde. MSc Thesis (in German), Eberhard-Karls-Universität Tübingen (FRG)
- Erb M, Palm G, Aertsen A, Bonhoeffer T (1986) Functional versus structural connectivity in neuronal nets. In: Strukturbildung und Musteranalyse, Proc. 9th Cybernetics Congress (DGK). Göttingen FRG, p 23
- Erb M, Aertsen A, Palm G (1989) Functional connectivity in neuronal systems: context-dependence of effective network organization does not require synaptic plasticity. In: Elsner N, Singer W (eds) Dynamics and plasticity in neuronal systems. Thieme, Stuttgart New York, p 445
- Erb M, Aertsen A, Palm G (1990) Stability in physiology-oriented neuronal network models. In: Elsner N, Roth G (eds) Brain - Perception - Cognition. Thieme, Stuttgart New York, p 527
- Erb M (1991) Simulation neuronaler Netze: Stabilität, Plastizität und Konnektivität. PhD Thesis (in German), Eberhard-Karls-Universität Tübingen (FRG)
- Gerstein GL, Bedenbaugh P, Aertsen AMHJ (1989) Neuronal Assemblies. IEEE Trans Biomed Engin 36:4-14
- Golomb D, Rubin N, Sompolinsky H (1990) Willshaw model: Associative memory with sparse coding and low firing rates. Phys Rev A 41:1853-1854
- Gray CM, Singer W (1989a) Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. Proc Natl Acad Sci USA 86:1698-1702
- Gray CM, König P, Engel AK, Singer W (1989b) Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. Nature 338:334-337
- Hebb DO (1949) The organization of behavior. A neuropsychological theory. Wiley, New York
- Krüger J, Mayer M (1990) Two types of neuronal synchrony in monkey striate cortex. Biol Cybern 64:135-140
- Marcus CM, Westerveld RM (1989) Stability of analog neural networks with delay. Phys Rev A 39:347-359
- Malsburg, C. von der (1981) The correlation theory of brain function. Internal report 81-2. Max-Planck-Institute for Biophysical Chemistry, Göttingen (FRG)
- Malsburg, C. von der (1986) Am I thinking assemblies? In: Palm G, Aertsen A (eds) Brain Theory. Springer, Berlin Heidelberg New York, pp 161-176
- Palm G (1982) Neural Assemblies. An alternative approach to artificial intelligence. Springer, Berlin Heidelberg New York

- Palm G (1986) Associative networks and cell assemblies. In: Palm G, Aertsen A (eds) Brain Theory. Springer, Berlin Heidelberg New York, pp 211-228
- Palm G (1987) On associative memories. In: Caianiello ER (ed) Physics of cognitive processes. World Scientific Publishing, Singapore, pp 380-442
- Rubin N, Sompolinsky H (1989) Neural networks with low firing rates. Europhys Lett 10:465-470
- Skarda CA, Freeman WJ (1987) How brains make chaos to make sense of the world. Brain Behav Sci 10:161-195
- Vaadia E, Ahissar E, Bergman H, Lavner Y (1991) Correlated activity of neurons: a neural code for higher brain functions? In: Krüger J (ed) Neuronal Cooperativity. Springer, Berlin Heidelberg New York, pp 249-279