

Networks of spiking neurons can emulate arbitrary Hopfield nets in temporal coding

Wolfgang Maass[†] and Thomas Natschläger[‡]

Institute for Theoretical Computer Science, Technische Universität Graz, Klosterwiesgasse 32/2, A-8010 Graz, Austria

Received 6 March 1997

Abstract. A theoretical model for analogue computation in networks of spiking neurons with temporal coding is introduced and tested through simulations in GENESIS. It turns out that the use of multiple synapses yields very noise robust mechanisms for analogue computations via the timing of single spikes in networks of detailed compartmental neuron models.

In this way, one arrives at a method for emulating arbitrary Hopfield nets with spiking neurons in temporal coding, yielding new models for associative recall of spatio-temporal firing patterns. We also show that it suffices to store these patterns in the efficacies of *excitatory* synapses.

A corresponding *layered* architecture yields a refinement of the synfire-chain model that can assume a fairly large set of different stable firing patterns for different inputs.

1. Introduction

Recent experimental results from neurophysiology have shown that in many biological neural systems not only the firing rate, but also the spatio-temporal pattern of neuronal firing carries important information.

Other recent experimental results indicate that it is in fact questionable whether biological neural systems are *able* to carry out analogue computation with analogue variables represented as firing rates. Due to ‘synaptic depression’ the amplitude of postsynaptic potentials arising from a presynaptic neuron u tends to scale as $1/f$ where f is the firing rate of u (see, e.g., Abbott *et al* 1997). Therefore both slowly firing neurons and rapidly firing neurons u inject roughly the same amount of current into a postsynaptic neuron during a given time window. This suggests that both a McCulloch–Pitts neuron and a sigmoidal neuron model overestimate the computational capability of a biological neuron for rate coding.

In addition, it has been argued that in view of the rather low firing rates of cortical neurons analogue computations in multi-layer neural systems with intermediate variables represented as firing rates would be much too slow to achieve the experimentally observed computation speed of concrete cortical neural systems (Rolls and Tovee 1994, Thorpe and Imbert 1989, Perret *et al* 1982).

Temporal coding with single spikes has been suggested as an alternative neural code, that would not be affected by synaptic depression. In addition, with this neural code one can

[†] E-mail: maass@igi.tu-graz.ac.at

[‡] E-mail: tnatschl@igi.tu-graz.ac.at

in principle achieve very high computation speed with biologically realistic low firing rates. In temporal coding with single spikes one can encode analogue variables in the firing times of neurons relative to the stimulus onset (Kjaer *et al* 1996), relative to some oscillation (Hopfield 1995), or relative to the firing times of other neurons (O'Keefe and Reece 1993, Thorpe and Imbert 1989).

In this paper we explore models for associative memory in networks of spiking neurons with the latter two types of temporal coding. In this regard our models differ from previous models for associative memory in networks of spiking neurons, which were based on rate coding (Fransén 1996, Simmen *et al* 1995, Ritz *et al* 1994, Gerstner and van Hemmen 1992, Lansner and Fransén 1992, Herz *et al* 1991). Through a rigorous mathematical result and through simulations of compartmental neuron models in GENESIS we demonstrate a computational mechanism that allows us to emulate any given Hopfield net with a network of spiking neurons with temporal coding. This mechanism can be implemented in a surprisingly noise robust way both with regard to some underlying background oscillation as in Hopfield (1995) and in a self-excitatory manner without any external 'clock'.

If one applies our method to a feedforward network instead of a recurrent net, one arrives at a new version of the synfire chain, a model which was introduced in Abeles (1991). This new version of the synfire chain can have a large set of different stable states ('memory patterns'), in contrast to the traditional synfire chain models that are based on just two states ('active' and 'inactive'). Furthermore, the occurrence of precisely timed firing patterns in the cortex (see Abeles *et al* 1993) can be explained just as well with this new version of the synfire chain as with the traditional model.

Apart from these specific results, this paper also addresses an important general question regarding the modelling of neural computation, namely, to what extent are mechanisms and results that have been demonstrated analytically for networks of relatively simple mathematical models for leaky integrate-and-fire neurons also valid for networks of substantially more complex compartmental neuron models?

Section 2 of this paper presents a theoretical result regarding the emulation of Hopfield nets in networks of spiking neurons. In section 3 we discuss the results of our GENESIS simulations of related networks of compartmental neuron models. Section 4 contains some conclusions.

2. Theoretical result

In this section we outline the construction of a network $\mathcal{S}_{\mathcal{H}}$ of spiking neurons which approximates the computation of an arbitrary given Hopfield network \mathcal{H} . As a model for a spiking neuron we take the common model of a *leaky integrate-and-fire neuron with noise*, and the somewhat more general *spike response model* of Gerstner and van Hemmen (1992), respectively. For details on the model see Maass (1997). The only specific assumption needed for the construction of $\mathcal{S}_{\mathcal{H}}$ in theorem 2.1 is that both the beginning of the rising part of an EPSP and the beginning of the descending part of an IPSP can be described by a linear function.

Theorem 2.1. Let \mathcal{H} be an arbitrary given Hopfield net with graded response (as in Hopfield 1984) and synchronous update. We assume that \mathcal{H} consists of n sigmoidal neurons u_i for $i \in \{1, \dots, n\}$ with arbitrary weights $w_{ij} \in \mathbb{R}$ for $i, j \in \{1, \dots, n\}$ and a piecewise linear activation function σ as indicated in figure 1. Then one can approximate any computation of \mathcal{H} by a recurrent network $\mathcal{S}_{\mathcal{H}}$ of n spiking neurons (with $O(1)$ auxiliary spiking neurons) in *temporal coding*. An input, internal state, or output of \mathcal{H} of the form $\langle x_1, \dots, x_n \rangle \in [-1, 1]^n$




Figure 1. Piecewise linear activation function of the Hopfield net \mathcal{H} with graded response.

is represented in $\mathcal{S}_{\mathcal{H}}$ by temporal coding, i.e. by a firing pattern of $\mathcal{S}_{\mathcal{H}}$ in which its i th neuron fires at time $kT - c\tilde{x}_i$, where $|x_i - \tilde{x}_i|$ can be made arbitrarily small. The reference time points kT for $k = 0, 1, \dots$ are defined by $O(1)$ periodically firing auxiliary neurons in $\mathcal{S}_{\mathcal{H}}$. Any fixed point of \mathcal{H} corresponds to a stable periodic firing pattern of $\mathcal{S}_{\mathcal{H}}$.

Proof. In $\mathcal{S}_{\mathcal{H}}$ we will have spiking neurons v_1, \dots, v_n which simulate the output of the sigmoidal neurons u_1, \dots, u_n in \mathcal{H} through the *time* at which they fire. We assume that there exists a synapse from neuron v_j to neuron v_i for any $i \neq j$. The spiking neurons v_1, \dots, v_n in $\mathcal{S}_{\mathcal{H}}$ fire periodically, for the k th time during the time interval $[kT - c, kT + c]$, $k \in \mathbb{N}$. When the neurons v_1, \dots, v_n in $\mathcal{S}_{\mathcal{H}}$ fire during the k th time interval at times $kT - c\tilde{x}_i(k)$, $i \in \{1, \dots, n\}$, then this temporal firing pattern will correspond to an output $\mathbf{x}(k) = \langle x_1(k), \dots, x_n(k) \rangle \in [-1, 1]^n$ of the n sigmoidal neurons of \mathcal{H} after their k th parallel computation step. These values are updated in \mathcal{H} according to the equation

$$x_i(k+1) = \sigma\left(\sum_{j=1}^n w_{ij} x_j(k)\right) \quad (2.1)$$

where σ is the activation function of \mathcal{H} (see figure 1) and w_{i1}, \dots, w_{in} are the weights of u_i in \mathcal{H} . The vector $\langle x_1(0), \dots, x_n(0) \rangle \in [-1, 1]^n$ is the network input.

Thus to simulate the $(k+1)$ th parallel computation step of \mathcal{H} in $\mathcal{S}_{\mathcal{H}}$, we need to make sure that each spiking neuron v_i fires at time $t_i(k+1) = (k+1)T - c\tilde{x}_i(k+1)$ with $|\tilde{x}_i(k+1) - x_i(k+1)|$ as small as possible.

We achieve this by exploiting a mechanism in $\mathcal{S}_{\mathcal{H}}$ for computing a weighted sum in temporal coding which is described in Maass (1997). This simulation is based on the observation that in the presence of some other excitation which moves the membrane potential close to the firing threshold, individual EPSPs, IPSPs (or volleys of synchronized PSPs) are able to *shift* the firing time of a neuron. This mechanism is particularly easy to analyse if we work in a range where all PSPs can be approximated well by linear functions. For this range one can show that the resulting firing time is linearly related to the *weighted sum* of the firing times of the presynaptic neurons, with the weights corresponding to the *efficacies* ('strengths') of the synapses involved.

To be precise, the firing time $t_i(k+1)$ of neuron v_i during the simulation of the $(k+1)$ th parallel computation step of \mathcal{H} can be expressed as follows:

$$t_i(k+1) = kT + d + \frac{\theta}{\sum_{j=1}^n \alpha_{ij}} - \frac{\sum_{j=1}^n \alpha_{ij}(kT - t_j(k))}{\sum_{j=1}^n \alpha_{ij}} \quad (2.2)$$

Here θ is the threshold of the spiking neuron v_i , d the propagation delay (the time from the generation of a spike at neurons v_j to the onset of the PSP at the soma at neuron v_i) and α_{ij} is the slope of the linear increasing (decreasing) section of the corresponding EPSP (IPSP) at the soma of v_i . Equation (2.2) holds rigorously if neuron v_i fires at a time $t_i(k+1)$ so

that each postsynaptic potential from the previous round of firing of v_1, \dots, v_n is at time $t_i(k+1)$ in its initial linearly rising phase at the trigger zone of v_i (in the case of an EPSP) or in its initial linearly declining phase (in the case of an IPSP).

The term $\sum_{j=1}^n \alpha_{ij}$ in the denominator of equation (2.2) normalizes the sum of the effective weights $\alpha'_{ij} = \alpha_{ij} / \sum_{j=1}^n \alpha_{ij}$ in the weighted sum of the terms $kT - t_j(k)$ to 1. To simulate networks \mathcal{H} with weights that do not sum up to 1 one is forced to use an additional neuron v_0 in $\mathcal{S}_{\mathcal{H}}$. If v_0 fires at times kT , $k \in \mathbb{N}$, one can then satisfy equation (2.3) for the next firing time $t_i(k+1)$ of neuron v_i . One sets $T := d + \theta/\lambda$ and one chooses values α_{ij} for $j = 0, \dots, n$ so that $\sum_{j=0}^n \alpha_{ij} = \lambda$ and $\alpha_{ij} = \lambda w_{ij}$, where $\lambda > 0$ can be chosen arbitrarily. Since $\tilde{x}_j(k)$ is defined through the equation $kT - t_j(k) = c \tilde{x}_j(k)$, one obtains

$$t_i(k+1) = (k+1)T - c \sum_{j=1}^n w_{ij} \tilde{x}_j(k). \quad (2.3)$$

Hence the firing time $t_i(k+1) = (k+1)T - c \tilde{x}_i(k+1)$ of v_i encodes the value $\tilde{x}_i(k+1) = \sum_{j=1}^n w_{ij} \tilde{x}_j(k)$ in temporal coding. In other words, according to equation (2.3) neuron v_i computes in *temporal* coding a weighted sum of the values $\tilde{x}_j(k)$. The above only holds for values of $\sum_{j=1}^n w_{ij} \tilde{x}_j(k)$ in the linear range $[-1, 1]$ of the activation function σ . In order to achieve the result that v_i fires for any $(\tilde{x}_1(k), \dots, \tilde{x}_n(k)) \in [-1, 1]^n$ approximately at time $(k+1)T - c \sigma(\sum_{j=1}^n w_{ij} \tilde{x}_j(k))$, in the formal construction of $\mathcal{S}_{\mathcal{H}}$ one can employ some auxiliary neurons that prevent a firing of v_i during the intermediate time interval $(kT + c, (k+1)T - c)$ and which make sure that v_i definitely fires once during the time interval $[(k+1)T - c, (k+1)T + c]$.

In our formal construction these auxiliary neurons in $\mathcal{S}_{\mathcal{H}}$ have an undesired side effect on the firing time $t_i(k+1)$ if $\sum_{j=1}^n w_{ij} \tilde{x}_j(k) \in [-1, 1]$ and the value of this sum is close to -1 or $+1$. The resulting firing time $t_i(k+1)$ of v_i in $\mathcal{S}_{\mathcal{H}}$ is shifted due the PSPs generated by this auxiliary neuron. This is the reason why the computation of \mathcal{H} cannot be simulated *precisely* in $\mathcal{S}_{\mathcal{H}}$. However, the amount of this shift can be made arbitrarily small. For further details concerning these auxiliary neurons, as well as noise and refractory behaviour, the reader is referred to Maass (1997).

3. Simulations with GENESIS

In the remainder of this paper we take the preceding theoretical construction as the basis for a case study. We want to find out to what extent the mechanisms for computations with temporal coding that have been verified theoretically for integrate-and-fire neurons correspond to stable computational mechanisms for substantially more detailed compartmental models of biological neurons simulated in GENESIS (Bower and Beeman 1995).

Surprisingly, it turns out that the most essential computational mechanism underlying the proof of theorem 2.1 (and the construction of Maass (1997)) works *even better* in the more detailed neuron model of GENESIS. The previously described computational mechanism for computing a weighted sum $\sum_{j=1}^n w_{ij} x_j(k)$ is theoretically sound as long as during each wave of firing the range $[0, 2c]$ of the differences in firing times of the n presynaptic neurons v_1, \dots, v_n is so small that there is a time point at which (despite their temporal differences) the resulting postsynaptic potentials are at the soma of each v_i *all* in their initial *linear* phase. However, since non-NMDA EPSPs rise very fast, the theory would suggest that the length $2c$ of this interval would have to be chosen to be about 1 ms. This is so small that one has to be concerned about the effect of various sources of temporal jitter on the

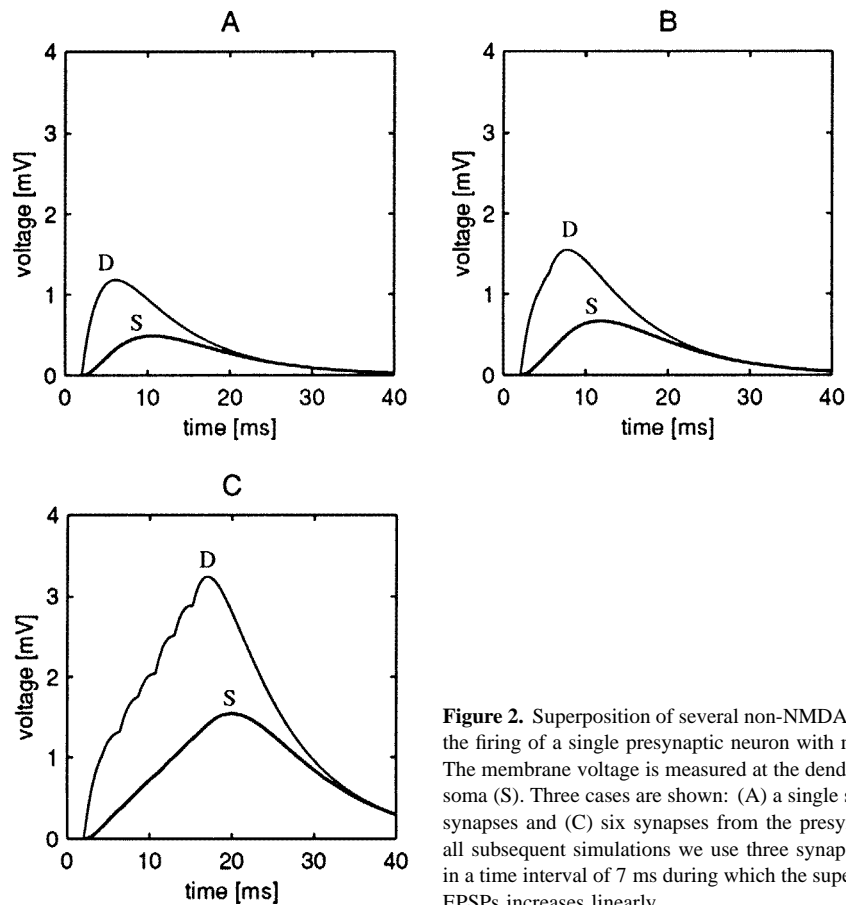


Figure 2. Superposition of several non-NMDA EPSPs caused by the firing of a single presynaptic neuron with multiple synapses. The membrane voltage is measured at the dendrite (D) and at the soma (S). Three cases are shown: (A) a single synapse, (B) three synapses and (C) six synapses from the presynaptic neuron. In all subsequent simulations we use three synapses, which results in a time interval of 7 ms during which the superposition of these EPSPs increases linearly.

precision of this temporal coding. Figure 2 shows that for GENESIS neurons the value $2c$ can be chosen to be much larger.

One can extend this length by replacing each synapse between a presynaptic neuron v_j and v_i by $l > 1$ synapses. In this way a single spike from v_j causes a superposition of l EPSPs at the soma of v_i . This superposition may have a substantially longer increasing phase if the signal pathways along the l synapses between v_j and v_i have a reasonable difference in their total delays (time from generation of a spike at neuron v_j to the onset of the PSP at the soma at neuron v_i). Besides the differences in propagation times in the dendritic tree (Zador *et al* 1995) the differences of the arrival times at various terminals of the axonal tree also contribute to the diversity of the total delays (in our simulations we assume a difference of 3.5 ms, see Manor *et al* (1991)). However, the simple models for leaky integrate-and-fire neurons predict that the resulting rising phase of a superposition of several EPSPs is quite 'bumpy', which limits its use for the here considered computational mechanism. In simulations with GENESIS the nonlinearity of this rising phase is clearly visible if one measures the resulting membrane voltages close to the synapses (see the curves labelled 'D' in figure 2). However, at the *soma*, where the shape of this superposition becomes *relevant* for the *computational mechanism* considered here, this *rising phase* of the superposition of several EPSPs is *almost perfectly linear* (see the curves labelled 'S' in figure 2). This observation allows us in our GENESIS simulations to stretch the length

$2c$ of the interval for temporal coding through the use of multiple synapses to $2c = 4$ ms (figure 2(C) shows that a much larger value can be chosen if more than three synapses are used and a difference of the delays of 10 ms is assumed).

Now we introduce a method for constructing for any given Hopfield net \mathcal{H} a network $\mathcal{G}_{\mathcal{H}}$ of biologically quite detailed spiking neurons so that $\mathcal{G}_{\mathcal{H}}$ has an equally rich set of attractors as \mathcal{H} . Note that each attractor of $\mathcal{G}_{\mathcal{H}}$ is a spatio-temporal firing pattern, whereas an attractor (fixed point) in \mathcal{H} has no temporal structure. Furthermore, we show that these firing patterns of $\mathcal{G}_{\mathcal{H}}$ in fact have attractor qualities: if one initiates $\mathcal{G}_{\mathcal{H}}$ with a firing pattern that differs from all of the stable firing patterns it moves to an attractor firing pattern within a few firing periods.

The ‘blueprint’ for the construction of $\mathcal{G}_{\mathcal{H}}$ is provided by the formal model $\mathcal{S}_{\mathcal{H}}$ that we have described in section 2. There are minor changes needed to move the construction of $\mathcal{S}_{\mathcal{H}}$ into a parameter range that is biologically more realistic.

One detail of our case study concerns the simulation of the *saturated* segments of the activation function σ (i.e. the required ‘nonlinearity’) of the given artificial neural net \mathcal{H} . Whereas auxiliary neurons are needed for that purpose in the theoretical construction of theorem 2.1 (similarly to Maass (1997)), the natural form of EPSPs and IPSPs suggests that these auxiliary neurons may be unnecessary in a practical context (see the discussion in remark 6 of Maass (1997)). This conjecture is supported through all our simulations of $\mathcal{G}_{\mathcal{H}}$, where no extra mechanism forces a neuron to fire once during each firing wave.




Figure 3. Nine randomly chosen ‘memory patterns’ from $\{-1, +1\}^{60}$.

For our GENESIS simulations we started out from a Hopfield net \mathcal{H} with $n = 60$ neurons, whose weights were computed with the projection method (Hertz *et al* 1991) so that nine randomly chosen vectors from $\{-1, 1\}^{60}$ (see figure 3) are fixed points of \mathcal{H} . We have simulated the network $\mathcal{S}_{\mathcal{H}}$ of spiking neurons (which simulates \mathcal{H} according to theorem 2.1) in GENESIS with a network $\mathcal{G}_{\mathcal{H}}$ of 61 neurons. In $\mathcal{G}_{\mathcal{H}}$ each neuron is modelled with 122 compartments, and there are three synapses between each pair of neurons (whose delays differ by up to 3.5 ms). $\mathcal{G}_{\mathcal{H}}$ has the same architecture and the same weights (properly scaled) as $\mathcal{S}_{\mathcal{H}}$ (30% of the weights were rounded to 0). Weights with negative values are modelled by inhibitory synapses.

In the protocols of our simulations of $\mathcal{G}_{\mathcal{H}}$ the bars on the left of the figures (‘i-error’) indicate the difference between the firing times of the *first* wave of firing and the closest memory pattern (both interpreted as vectors from $[-1, +1]^{60}$ via temporal coding). Correspondingly, the bars to the right (‘o-error’) indicate the difference between the firing times of the *last* shown firing wave and the same memory pattern. In both cases the non-firing of a neuron is treated like the firing of a neuron among the last in the firing wave. This is justified by the observation that both scenarios reflect equivalent ways of encoding the smallest possible analogue value ‘-1’ (see the discussions in Maass (1997)). A neuron v_j that fires *very late* during a firing wave has as negligible an impact on the firing times during the next firing wave as v_j not firing at all. It is interesting to note that




Figure 4. Membrane potential at the soma of neuron 5 in the simulation from figure 5 with and without noise. The diagram *with noise* reflects the additional internal noise underlying *all* our GENESIS simulations. We assume the existence of a random current at the soma of each neuron (Gaussian distribution with a variance of 10^{-18} A^2), and that each synapse fails with a probability of 15%.

in all our simulation both types of encoding ‘-1’ are present. During the first few firing waves ‘-1’ is encoded by a late firing, whereas after the network has reached an attracting firing pattern with very high preference the neurons choose to encode a negative value close to -1 by *not* firing at all. Thus a subsequent network is able to read out the result of the computation of \mathcal{G}_H in the more noise-robust form of a *rate code*.

In the space just before ‘o-error’ in figures 5–10 inclusive we have marked with a horizontal bar each neuron i of \mathcal{G}_H which does *not* satisfy the following property: neuron i fires during the last firing wave *if and only if* the i th component of the memory pattern is a ‘+1’. One can see that only relatively low ‘errors’ occur where a neuron does not satisfy this property.

To investigate the noise robustness of our model two internal sources of noise are considered (in addition to the noise by which the network inputs are perturbed): a noisy membrane potential and failing synapses (see figure 4). These two types of noise underlie all our simulations, which shows that our construction is robust against a substantial level of noise.

The results of the simulation of \mathcal{G}_H with a *very* noisy version of memory pattern 1 as

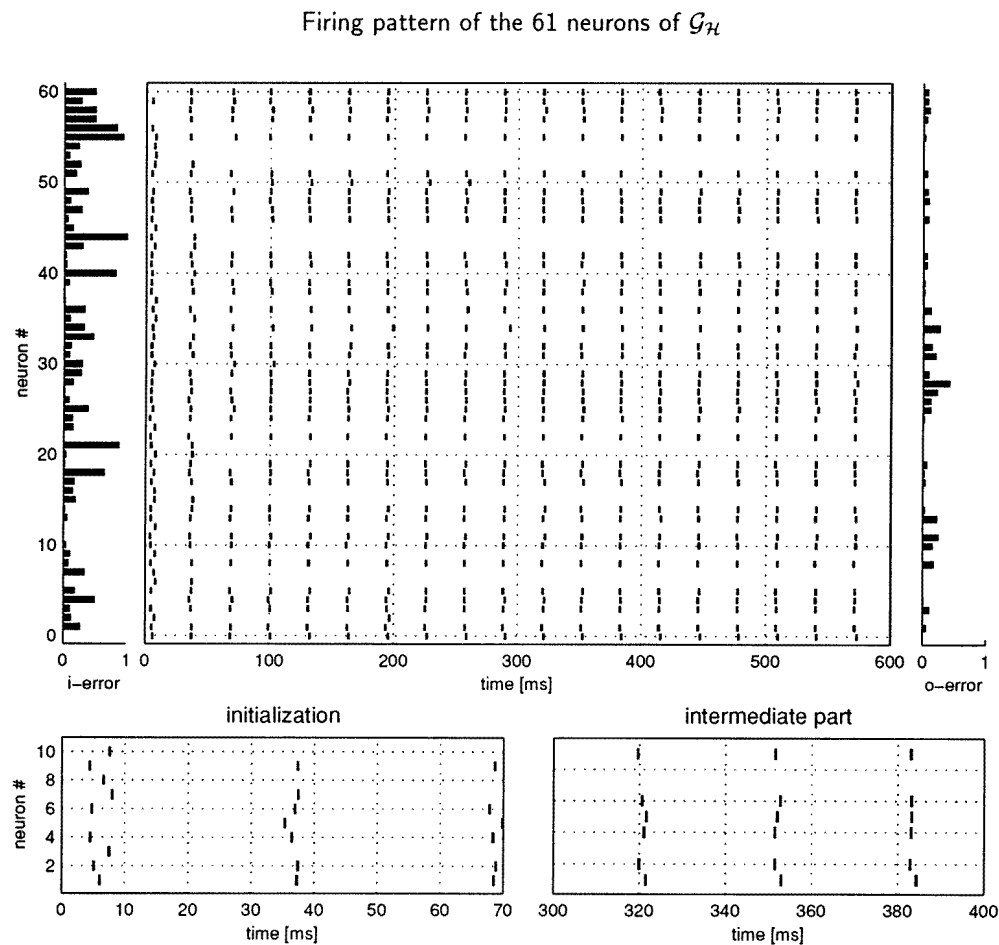


Figure 5. GENESIS simulation of $\mathcal{G}_{\mathcal{H}}$ with a corrupted version x of memory pattern 1 as input (see text for details). The two lower diagrams show the firing times of the first 10 neurons at a higher temporal resolution. The resulting effective difference between the *input* vector $x \in [-1, +1]^{60}$ that is given to $\mathcal{G}_{\mathcal{H}}$ in *temporal coding* and memory pattern 1 is indicated on the left ('i-error'). For each neuron i an error bar of length in $[0, 1]$ indicates the size of the deviation between the i th component of x and the i th component of memory pattern 1. The difference between the *output* vector $y \in [-1, +1]^{60}$ in the same temporal coding (with non-firing viewed as an encoding of -1) and memory pattern 1 is indicated on the right ('o-error').

input is shown in figure 5. The corrupted input x was constructed as follows from memory pattern 1: a randomly chosen 15% of the components of memory pattern 1 were multiplied with -1 . On top of this, the value of each component of x was moved by a random value from $[-0.4, 0.4]$. This input x was presented to $\mathcal{G}_{\mathcal{H}}$ in *temporal coding*, although not precisely: a randomly chosen 10% of the neurons in $\mathcal{G}_{\mathcal{H}}$ were prevented from firing (corresponding to an input value -1). Further diagrams (figures 8, 9 and 10 presenting simulation results with all nine stored patterns shown in figure 3) show that in fact *all arbitrarily chosen fixed points of \mathcal{H}* can be associatively recalled by $\mathcal{G}_{\mathcal{H}}$. Figures 5 and 10 show that after 10 firing waves each memory pattern is recalled in the 'digitized' output code (firing/non-firing) *without any error*. Furthermore, even after five firing waves almost




Figure 6. Here the input to $\mathcal{G}_{\mathcal{H}}$ is a linear combination of memory pattern 1 with factor 0.5 and of memory pattern 4 with factor 0.3, presented in temporal coding, with noise added as in figure 5. The error bars to the left indicate the difference between the actual input and memory pattern 1, which is the stronger component. The error bars to the right indicate the difference between the output and memory pattern 1 in temporal coding.

no ‘digital’ errors in associative recall in terms of firing/non-firing occur (see figures 8 and 9).

It turns out that in contrast to the theoretical model $\mathcal{S}_{\mathcal{H}}$, no separate ‘oscillator’ is needed in $\mathcal{G}_{\mathcal{H}}$ to define the reference times kT for $k = 0, 1, 2, \dots$. Instead, neuron 0 of $\mathcal{G}_{\mathcal{H}}$, whose firing times provide the reference times kT , is a neuron like all the others, and it receives inputs from all other neurons in $\mathcal{G}_{\mathcal{H}}$ through synapses with equal weight. It also has outgoing synapses to all other neurons in $\mathcal{G}_{\mathcal{H}}$. Thus its firing contributes to the triggering of the next firing wave.

The network $\mathcal{G}_{\mathcal{H}}$ is not only capable of an associative recall of memory patterns if the input consists of a noisy version of one of the stored patterns. It also has the ability to find the stronger component if the stimulus is a combination of two of the stored memory patterns. This fact is demonstrated through simulations. The results are shown in figure 6. This is in contrast to a pattern-recognition system comprised of spiking neurons proposed by Hopfield. Hopfield’s construction (Hopfield 1995) is based on ‘grandmother neurons’ which each encode one spatio-temporal firing pattern through the locations of synapses on their dendritic tree.

In $\mathcal{G}_{\mathcal{H}}$ we employ large transmission delays of 25 ms in order to avoid refractory effects. This yields ‘oscillations’ of $\mathcal{G}_{\mathcal{H}}$ of about 40 Hz, a value that is very common in biological neural systems. In another version we have replaced the completely connected network $\mathcal{G}_{\mathcal{H}}$ by a layered network $\mathcal{G}_{\mathcal{H}}^{\text{synfire}}$, with a transmission delay of 5 ms between neurons in successive layers. One may argue that this is a neurophysiologically realistic value (Manor *et al* 1991). $\mathcal{G}_{\mathcal{H}}^{\text{synfire}}$ consists of three layers of 61 neurons each, with 70% of the connections from one layer to the next (and from layer 3 back to layer 1) with a synaptic failure probability of 15% as in $\mathcal{G}_{\mathcal{H}}$, and weights corresponding to those in $\mathcal{G}_{\mathcal{H}}$. Our simulations (see figure 7) show that the layered network, which is closely related to a synfire chain (Abeles 1991), has the same ability as $\mathcal{G}_{\mathcal{H}}$ for associative recall of any of the nine given memory patterns. Furthermore, it needs just 50 ms for almost error-free recall. Its firing

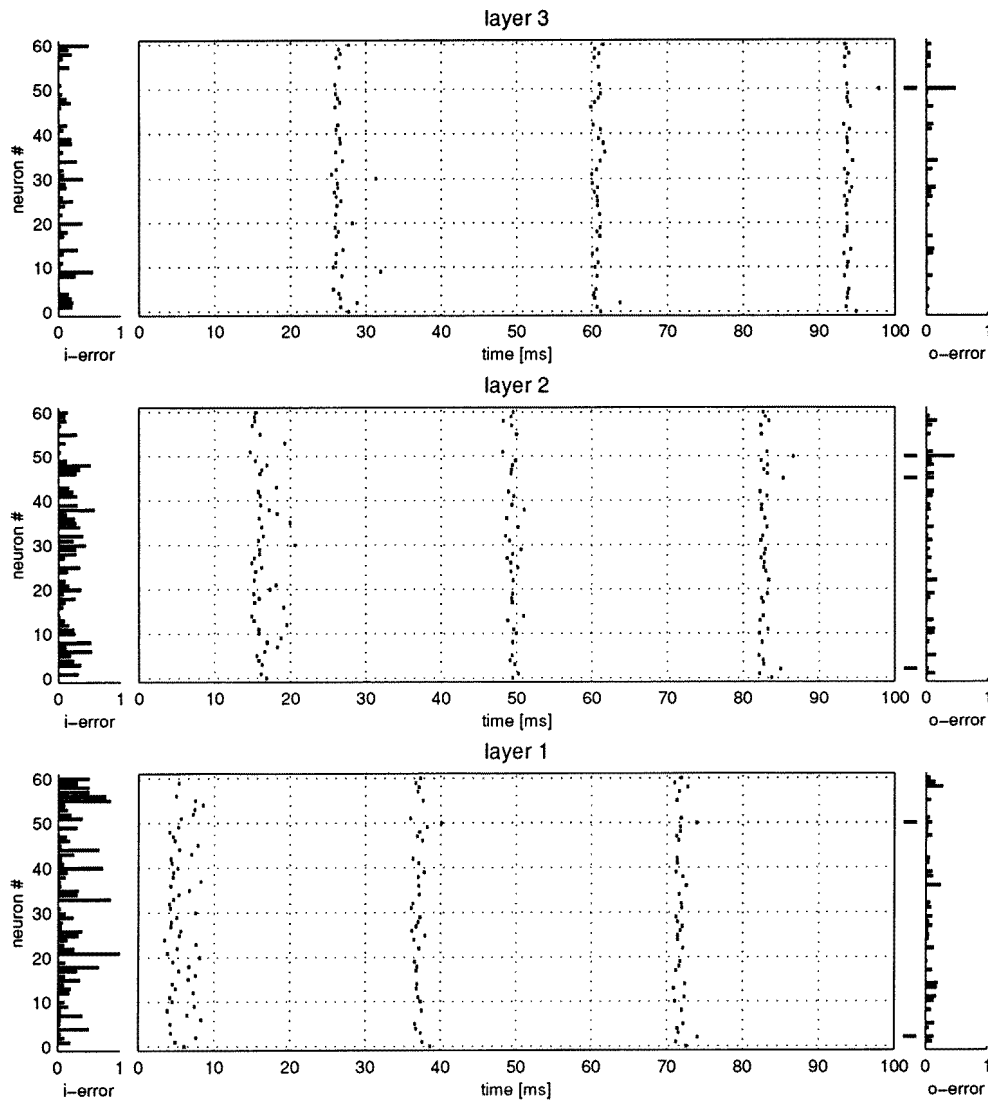


Figure 7. Simulation of $\mathcal{G}_{\mathcal{H}}^{\text{synfire}}$ for the same input as in figure 5. The firing behaviour is plotted separately for each of the three layers of $\mathcal{G}_{\mathcal{H}}^{\text{synfire}}$, with error bars indicating deviations from memory pattern 1 in its first and last firing wave. Thus the error reduction achieved during the first two firing rounds can be seen in the three plots of the ‘i-error’ on the left.

behaviour is consistent with the experimental data from Abeles *et al* (1993), just as the original model of a synfire chain. One could just as well work with a purely feedforward synfire chain. In fact our simulations show that a feedforward synfire chain with three to five layers would suffice for associative recall of a stored pattern by the synfire chain.

In $\mathcal{G}_{\mathcal{H}}$ and $\mathcal{G}_{\mathcal{H}}^{\text{synfire}}$ we modelled each negative weight w_{ij} from the network \mathcal{H} through separate inhibitory synapses. Therefore also the precise timing of IPSPs of the last wave of firing becomes relevant to the next firing time $t_i(k+1)$ of neuron i . A closer look at the structure of the inhibitory pathways found in many biological neural systems may

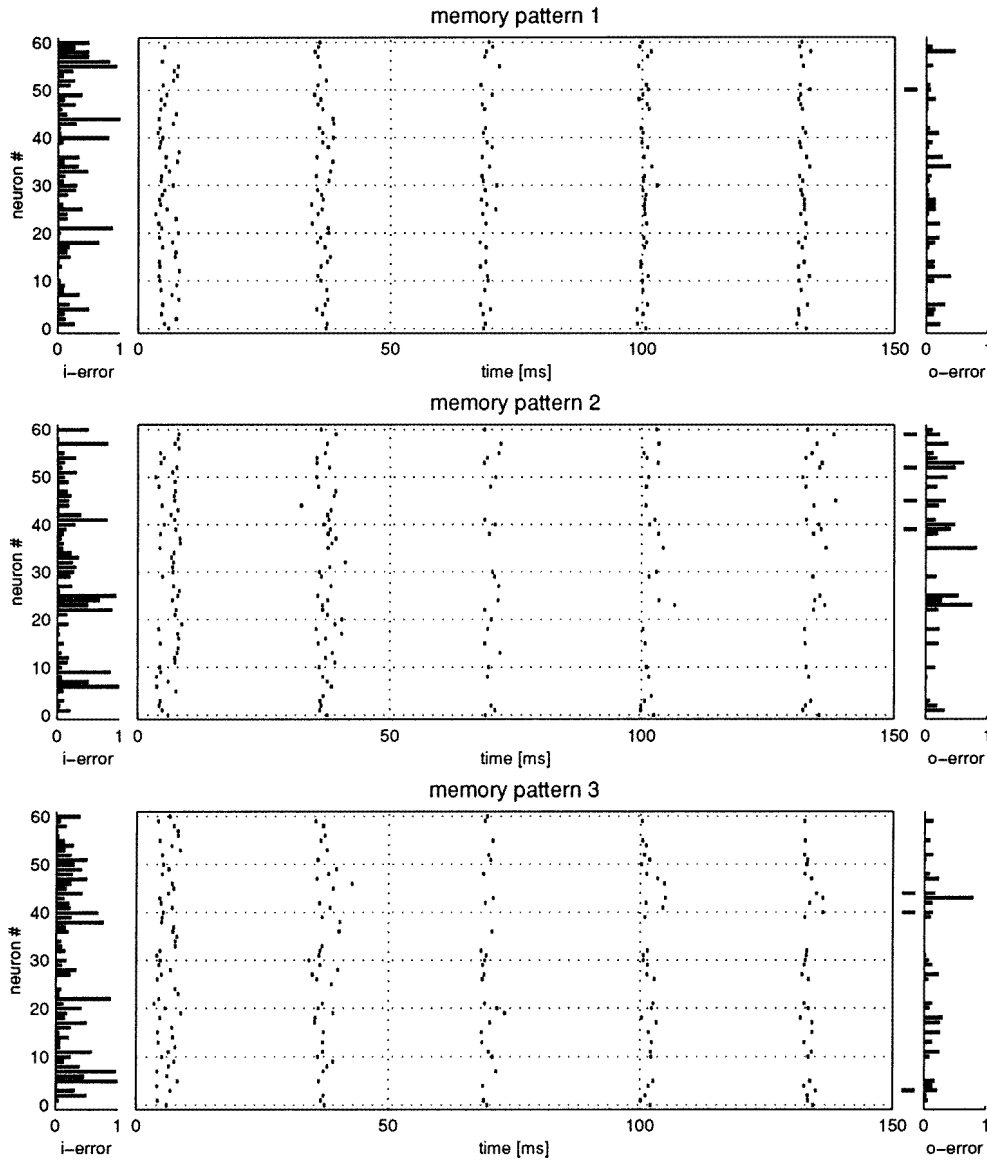


Figure 8. Simulations of the network $\mathcal{G}_{\mathcal{H}}$ as in figure 5, with corrupted versions of memory patterns 1, 2 and 3 as input. Note that almost error-free recall occurs already after five firing waves.

suggest that inhibition is not used in such a precise manner. One may argue that the role of inhibition is to regulate the overall activity in a neural system. To obtain results which are consistent with this point of view we investigate a third implementation of the network $\mathcal{S}_{\mathcal{H}}$ in GENESIS. In this network $\tilde{\mathcal{G}}_{\mathcal{H}}$ positive weights w_{ij} from the network \mathcal{H} are represented by excitatory synapses as in $\mathcal{G}_{\mathcal{H}}$, but negative weights w_{ij} are not modelled explicitly. Instead, an inhibitory bias (equal for all neurons $i \in \{1, \dots, n\}$ in $\tilde{\mathcal{G}}_{\mathcal{H}}$) is applied. This inhibition is produced by synapses formed between the additional neuron 0 and each neuron i . Thus all

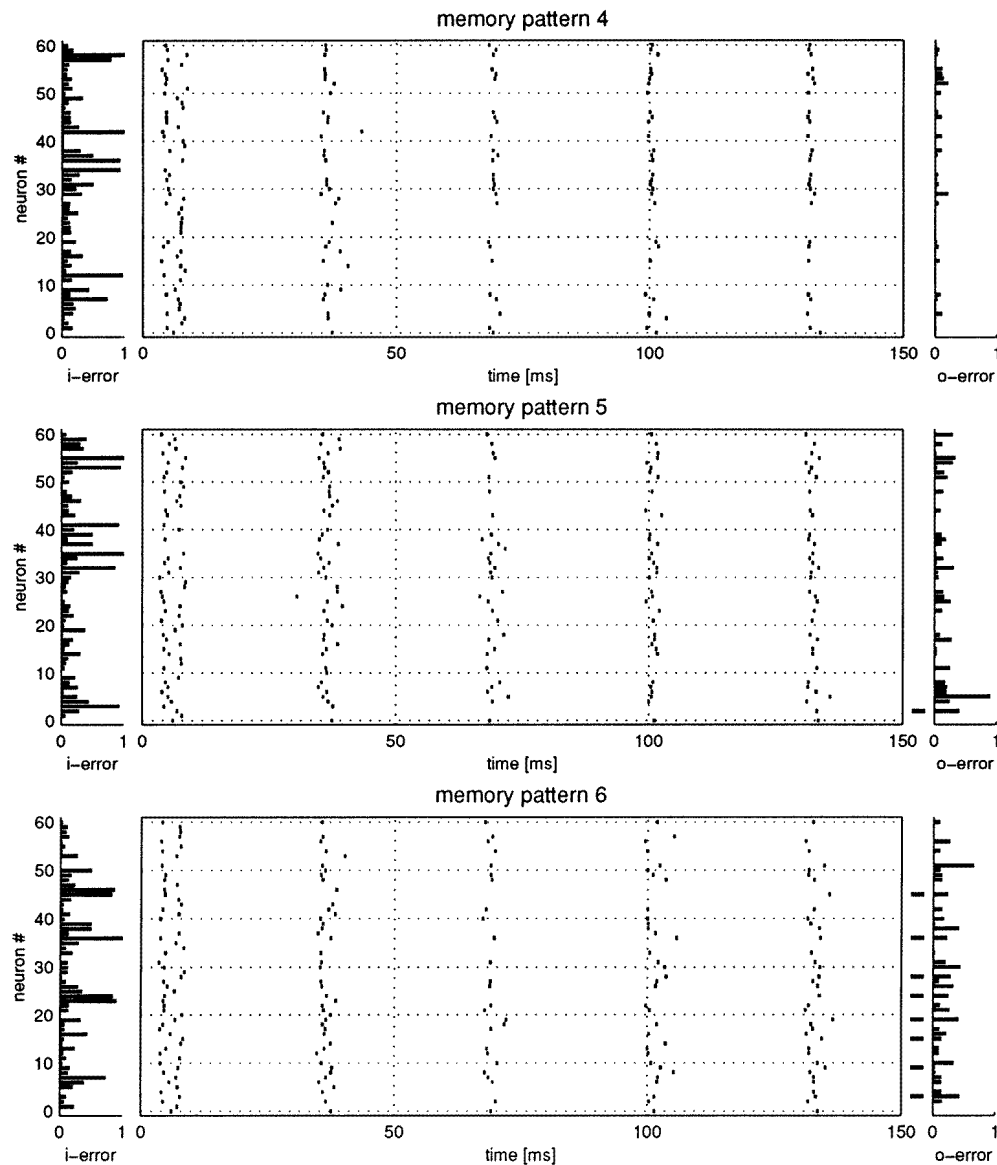


Figure 9. Simulations of the network $\tilde{\mathcal{G}}_{\mathcal{H}}$ as in figure 5, with corrupted versions of memory patterns 4, 5 and 6 as input.

information about the weights w_{ij} of the simulated Hopfield net \mathcal{H} is stored exclusively in the efficacies of *excitatory* synapses in $\tilde{\mathcal{G}}_{\mathcal{H}}$.

The construction of $\tilde{\mathcal{G}}_{\mathcal{H}}$ is based on a common heuristic. We set $\tilde{w}_{ij} = w_{ij}$ if $w_{ij} > 0$, $\tilde{w}_{ij} = 0$ if $w_{ij} < 0$, where the w_{ij} are the weights of a Hopfield net \mathcal{H} . We assume that the weights w_{ij} were chosen in order to store in \mathcal{H} p patterns ξ^1, \dots, ξ^p either with the standard Hebbian learning rule (Hopfield 1982) or the projection method (Hertz *et al* 1991). We consider the distribution of values $h_i^\mu = \sum_{j=1}^n \tilde{w}_{ij} \xi_j^\mu$, where ξ_j^μ denotes the j th component of ξ^μ . This distribution typically has two clusters: one with larger values of h_i^μ

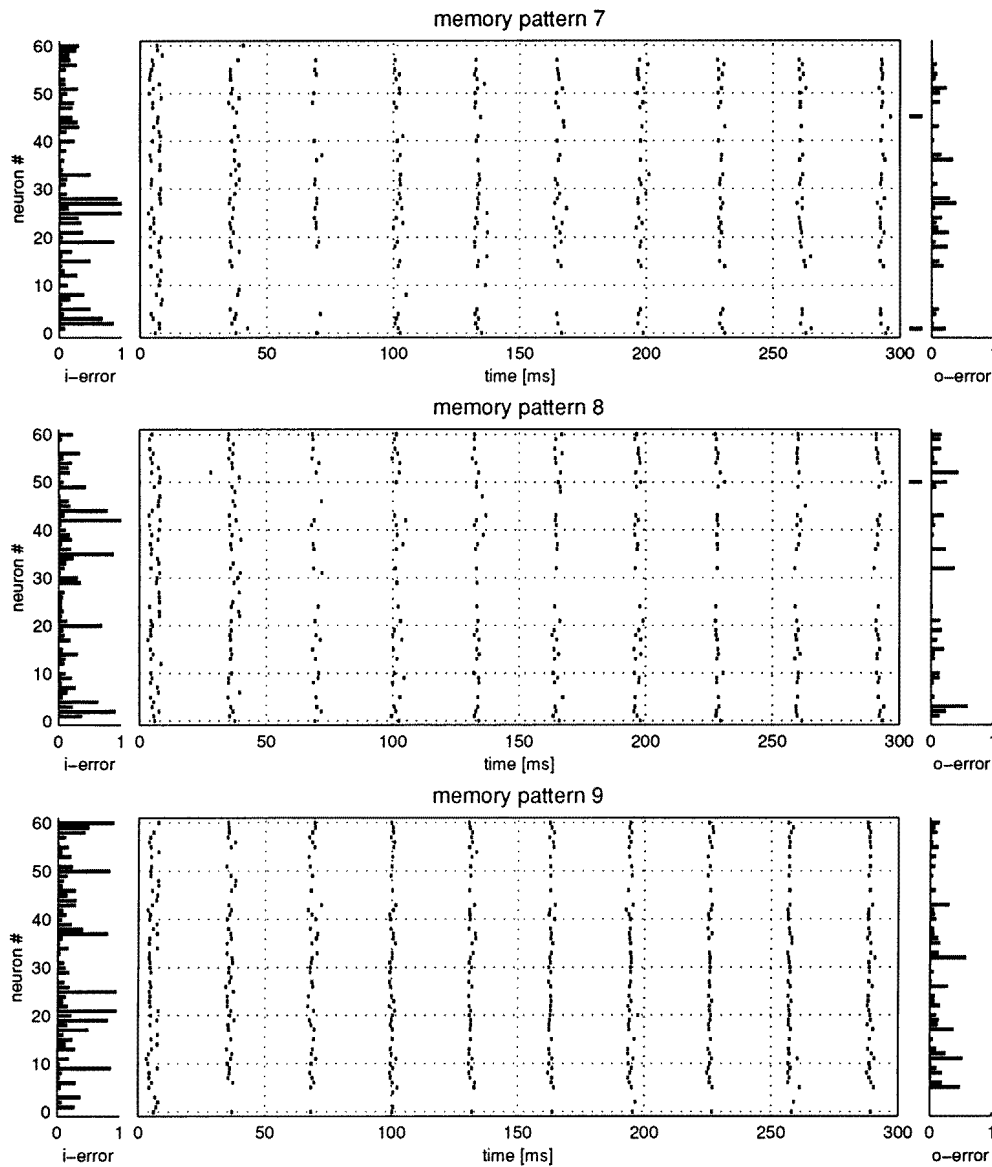


Figure 10. Simulations of the network $\mathcal{G}_{\mathcal{H}}$ as in figure 5, with corrupted versions of memory patterns 7, 8 and 9 as input.

(for $\langle i, \mu \rangle$ with $\xi_i^\mu = +1$) and one with smaller values of h_i^μ (for $\langle i, \mu \rangle$ with $\xi_i^\mu = -1$). Therefore one can choose a bias b such that $h_i^\mu + b > 0$ for most $\langle i, \mu \rangle$ with $\xi_i^\mu = 1$, and $h_i^\mu + b < 0$ for most $\langle i, \mu \rangle$ with $\xi_i^\mu = -1$. In other words, through a careful choice of the bias b one can repair the damage resulting from the replacement of the weights w_{ij} by \tilde{w}_{ij} (see subsection 7.5.4 of Peretto (1992)).

Through simulations in GENESIS we show that $\tilde{\mathcal{G}}_{\mathcal{H}}$ does indeed have the capability of associative recall of stored patterns. For these simulations we started out from a Hopfield net \mathcal{H} with $n = 60$ neurons whose weights were computed with the standard Hebbian

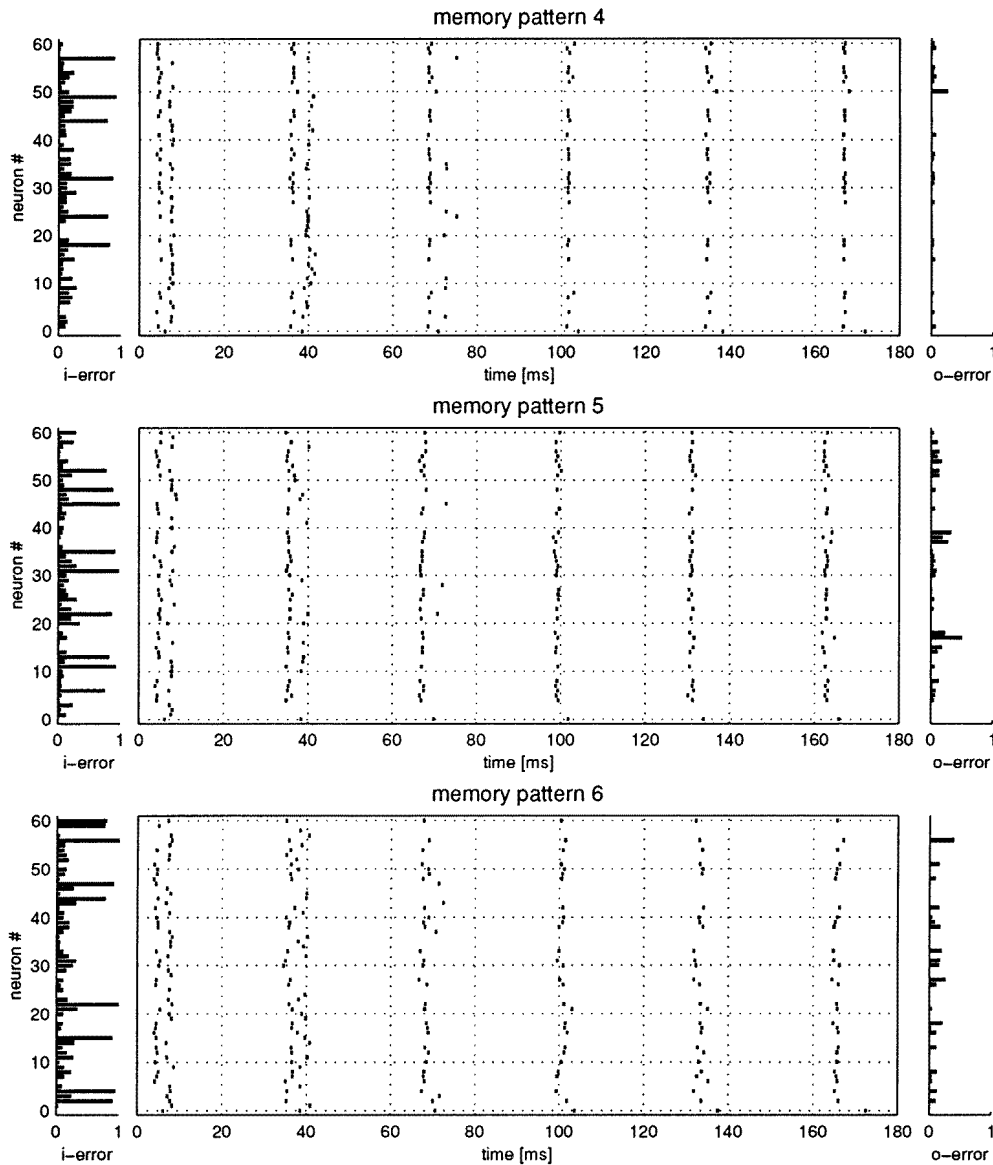


Figure 11. Simulations of the network $\tilde{\mathcal{G}}_{\mathcal{H}}$ with patterns stored exclusively in excitatory synapses. The simulation is carried out as in figure 5, with corrupted versions of memory patterns 4, 5 and 6 as input.

learning rule (Hopfield 1982) so that five of the nine vectors of figure 3 are fixed points of \mathcal{H} . We have chosen for these simulations the Hebb rule instead of the projection rule for the construction of \mathcal{H} in order to demonstrate that this makes no difference for our simulations with spiking neurons. Negative weights w_{ij} are set to 0. The common value of the efficacies of the inhibitory synapses was primarily chosen according to the heuristic for the choice of the bias b we have described above. However, it was necessary to fine-tune this common value to achieve the desired behaviour of $\tilde{\mathcal{G}}_{\mathcal{H}}$. Figures 11 and 12 show the

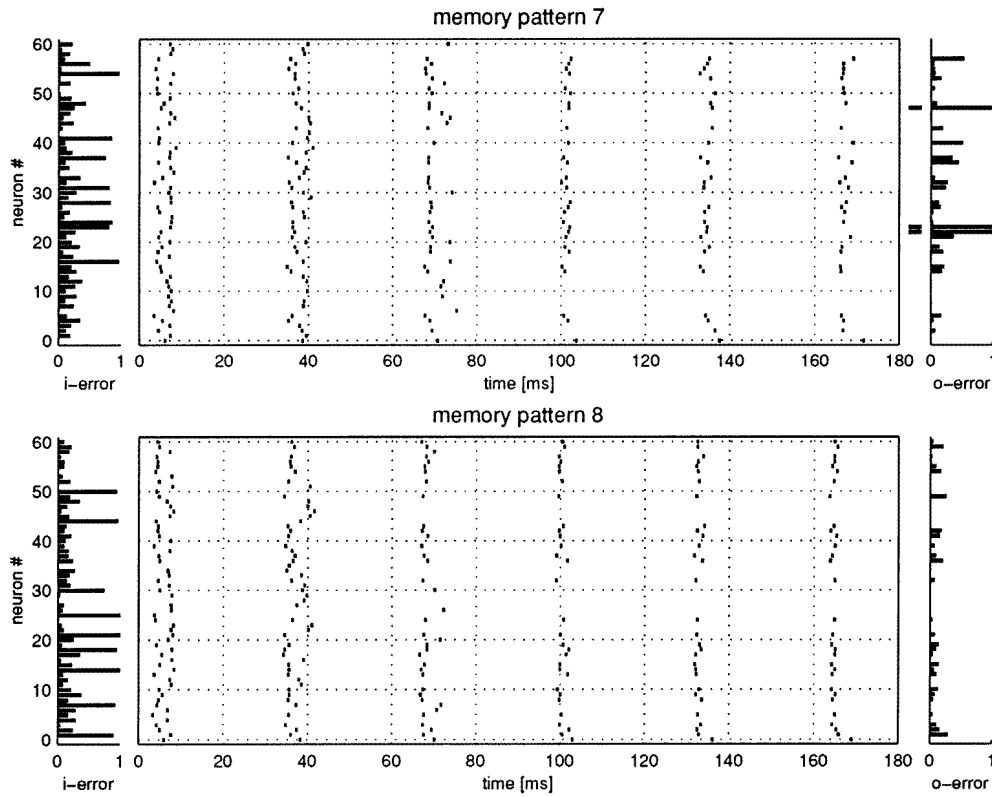


Figure 12. Simulations of the network $\tilde{\mathcal{G}}_{\mathcal{H}}$ with patterns stored exclusively in excitatory synapses. The simulation is carried out as in figure 5, with corrupted versions of memory patterns 7 and 8 as input.

protocols of the simulations of $\tilde{\mathcal{G}}_{\mathcal{H}}$ with noisy versions of the five stored memory patterns (i.e. patterns 4, 5, 6, 7 and 8 of figure 3) as inputs. These results clearly demonstrate that $\tilde{\mathcal{G}}_{\mathcal{H}}$ is able to associatively recall all stored memory patterns. In these simulations the same internal noise as in all previous simulations is present.

One can also extend the network $\tilde{\mathcal{G}}_{\mathcal{H}}$ to a multilayer network $\tilde{\mathcal{G}}_{\mathcal{H}}^{\text{synfire}}$ just as described above for $\mathcal{G}_{\mathcal{H}}^{\text{synfire}}$. Figures 11 and 12 suggest that five layers would suffice to provide almost error-free recall. The network $\tilde{\mathcal{G}}_{\mathcal{H}}^{\text{synfire}}$ now involves only excitatory synapses and a ‘global’ inhibition. Therefore the construction of $\tilde{\mathcal{G}}_{\mathcal{H}}^{\text{synfire}}$ is more closely related to the original version of a synfire chain as $\mathcal{G}_{\mathcal{H}}^{\text{synfire}}$. On the other hand $\tilde{\mathcal{G}}_{\mathcal{H}}^{\text{synfire}}$ has the same qualitative behaviour as $\mathcal{G}_{\mathcal{H}}^{\text{synfire}}$. In particular, it can recall many different stored firing patterns. However, we found that the number of memory patterns which can be stored in $\tilde{\mathcal{G}}_{\mathcal{H}}^{\text{synfire}}$ is slightly smaller than for $\mathcal{G}_{\mathcal{H}}^{\text{synfire}}$.

4. Conclusions

We have exhibited a new way of simulating Hopfield nets with graded response by recurrent networks of spiking neurons with temporal coding. The resulting networks of spiking neurons carry out their computation very fast, even with biologically realistic low firing rates.

We have first demonstrated this simulation method through a rigorous theoretical result. We have then chosen the resulting networks as a basis for a case study, where we check the essential computational mechanisms of a common simplified theoretical model through simulations in GENESIS. It turns out that our construction method yields surprisingly noise-robust mechanisms for analogue computations by compartmental neuron models with temporal coding. In particular the dynamic range that is available for temporal coding can be made quite large by using multiple synapses between neurons. In this way one can work in a range where the temporal coding scheme is not seriously affected by realistic values of temporal jitter in firing times.

Our simulations in GENESIS show that one can in fact simulate an arbitrarily given Hopfield net through a network of compartmental neuron models with temporal coding. The resulting network has even under the influence of rather realistic amounts of internal noise a rich set of different spatio-temporal firing patterns as attractors. These attractors mimic the set of attractors of the given Hopfield net.

The layered version of the resulting network of spiking neurons provides a refinement of the synfire-chain model (Abeles 1991). This refined version of a synfire chain overcomes one of the essential bottlenecks of the original synfire chain model: its low storage capacity. Whereas a synfire chain has only two states (active and inactive), the layered networks considered here (with a similar architecture, but a more complex weight-assignment) exhibit a fairly *large reservoir of different stable firing patterns*.

We show that this can even be achieved by storing all information about these ‘memory patterns’ in the efficacies of *excitatory* synapses. To the best of our knowledge all arguments that support the existence of synfire chains in higher cortical areas (on the basis of the consistency of their firing pattern with recorded data (Abeles *et al* 1993)) also support the existence of our refined version in the same biological neural system.

Acknowledgments

This work was supported by the ‘Fonds zur Förderung der wissenschaftlichen Forschung (FWF), Austrian Science Fund’, project number P12153.

References

- Abbott L F, Varela J A, Sen K and Nelson S B 1997 Synaptic depression and cortical gain control *Science* **275** 220–3
enhances information transfer between cortical neurons *Preprint*
- Abeles M 1991 *Corticonics: Neural Circuits of the Cerebral Cortex* (Cambridge: Cambridge University Press)
- Abeles M, Bergman H, Margalit E and Vaadia E 1993 Spatiotemporal firing patterns in the frontal cortex of behaving monkeys *J. Neurophysiol.* **70** 1629–38
- Bower J M and Beeman D 1995 *The Book of GENESIS: Exploring Realistic Neural Models with the GEneral NEural Simulation System* (New York: TELOS)
- Fransén E 1996 Biophysical simulation of cortical associative memory *PhD Thesis* Stockholm University
- Gerstner W and van Hemmen J L 1992 Associative memory in a network of ‘spiking’ neurons *Network* **3** 139–64
- Hertz J, Krogh A and Palmer R G 1991 *Introduction to the Theory of Neural Computation* (Reading, MA: Addison-Wesley)
- Herz A V M, Li Z and van Hemmen J L 1991 Statistical mechanics of temporal association in neural networks with transmission delays *Phys. Rev. Lett.* **66** 1370–3
- Hopfield J 1995 Pattern recognition computing using action potential timing for stimulus representation *Nature* **376** 33–6
- 1982 Neural networks and physical systems with emergent collective computational abilities *Proc. Natl Acad. Sci. USA* **79** 2554–8

- 1984 Neurons with graded response have collective computational properties like those of two-state neurons *Proc. Natl Acad. Sci. USA* **81** 3088–92
- Kjaer T W, Gawne T J and Richmond B J 1996 Latency: another potential code for feature binding in striate cortex *J. Neurophysiol.* **76** 1356–60
- Lansner A and Fransén E 1992 Modelling hebbian cell assemblies comprised of cortical neurons *Network: Comput. Neural Syst.* **3** 105–19
- Maass W 1997 Fast sigmoidal networks via spiking neurons *Neural Comput.* **9** 279–304
- Manor Y, Koch C and Segev I 1991 Effect of geometrical irregularities on propagation delay in axonal trees *Biophys. J.* **60** 1424–37
- O’Keefe J and Reece M L 1993 Phase relationship between hippocampal place units and the EEG theta rhythm *Hippocampus* **3** 3317–30
- Peretto P 1992 *An Introduction to the Modeling of Neural Networks* (Cambridge: Cambridge University Press)
- Perret D I, Rolls E T and Caan W C 1982 Visual neurons responsive to faces in the monkey temporal cortex *Exp. Brain Res.* **47** 329–42
- Ritz R, Gerstner W, Fuentes U and van Hemmen J L 1994 A biologically motivated and analytically soluble model of collective oscillations in the cortex *Biol. Cybern.* **71** 349–58
- Rolls E T and Tovee M J 1994 Processing speed in the cerebral cortex, and the neurophysiology of visual backward masking *Proc. R. Soc. B* **257** 9–15
- Simmen M W, Rolls E T and Treves A 1995 Rapid retrieval in an autoassociative network of spiking neurons *Computational Neuroscience* ed J M Bower (New York: Academic) pp 273–8
- Thorpe S T and Imbert M 1989 Biological constraints on connectionist modelling *Connectionism in Perspective* ed R Pfeifer, Z Schreter, F Fogelman-Soulie and L Steels (Amsterdam: Elsevier)
- Zador A, Agmon-Snir H and Segev I 1995 The morphoelectronic transform: a graphical approach to dendritic function *J. Neurosci.* **15** 1669–82