EMULATION OF HOPFIELD NETWORKS WITH SPIKING NEURONS IN TEMPORAL CODING*

Wolfgang Maass and Thomas Natschläger

Institute for Theoretical Computer Science, Technische Universität Graz,
Klosterwiesgasse 32/2, A-8010 Graz, Austria
E-mail: {maass, tnatschl}@igi.tu-graz.ac.at

ABSTRACT

A theoretical model for analog computation 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 analog computations with temporal coding in networks of detailed compartmental neuron models. One arrives in this way 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. A corresponding *layered* architecture yields a refinement of the synfire-chain model that can assume a fairly large set of different firing patterns for different inputs.

INTRODUCTION

We present a new model for associative memory in networks of spiking neurons that is based on analog computation with *temporal coding*. We consider network inputs which encode an analog pattern $\mathbf{x} \in [-1, 1]^n$ (the "stimulus") through relative *time delays* of the initial firing of all (or almost all) neurons in the network. In such a code any $\langle x_1, \ldots, x_n \rangle \in [-1, 1]^n$ can be represented by a firing pattern where the i-th neuron fires at time $T - c \cdot x_i$ (for suitable constants c and T). We refer to this scheme as *temporal coding* in the following. This type of neural code is discussed for example by Hopfield[1], Thorpe and Imbert[2] and Kjaer *et al*[3]. It has been argued that such neural code is especially useful for *very fast* neural computations with low firing rates.

We exhibit both through a rigorous mathematical result and through simulations of compartmental neuron models in GENESIS 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.

*Research was supported by the "Fonds zur Förderung der wissenschaftlichen Forschung (FWF), Austrian Science Fund", project number P12153.

This article also addresses an important general question regarding the modeling of neural computation: 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?

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*, respectively the somewhat more general *spike response model* [4, 5]. The only specific assumption that is needed for the construction of $\mathcal{S}_\mathcal{H}$ in theorem 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 1 *Let $\mathcal{H}$ be an arbitrary given Hopfield net with graded response[6] 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 ($\sigma(x) = -1$ if $x < -1$, $\sigma(x) = x$ if $-1 \leq x \leq +1$ and $\sigma(x) = +1$ if $x > +1$).*

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 $\langle x_1, \dots, x_n \rangle \in [-1, 1]^n$ 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 $k\,T - c\,\tilde{x}_i$, where $|x_i - \tilde{x}_i|$ can be made arbitrarily small. The reference time points $k\,T$ for $k = 0, 1, \dots$ are defined by $O(1)$ periodically firing auxiliary neurons in $\mathcal{S}_\mathcal{H}$. Any fixpoint of $\mathcal{H}$ corresponds to a stable periodic firing pattern of $\mathcal{S}_\mathcal{H}$.*

The *proof* of theorem 1 is based on a construction of Maass[5] for leaky integrate-and-fire neurons an is given in more detail elsewhere[7]. It relies on the observation that for a certain range of parameters the firing delay of a neuron v can be expressed as a weighted sum $\sum_{i=1}^{n} w_i x_i$ of the firing delays $c \cdot x_i \in [-c, c]$ of n presynaptic neurons (where the parameter w_i describes the efficacy of the synapse from the i-th presynaptic neuron, and c is some constant). The weights w_{ij} in $\mathcal{S}_\mathcal{H}$ are the same (properly scaled) as for the corresponding edges in $\mathcal{H}$.

SIMULATIONS WITH GENESIS

In the remainder of this article we take this theoretical construction as basis for a case study. We want to find out to what extent 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[8].

Surprisingly, it turns out that the most essential computational mechanism that underlies the proof of theorem 1 (and the construction of Maass[5]) works in the more detailed neuron model of GENESIS *even better*. The previously mentioned computational mechanism for computing a weighted sum $\sum_{i=1}^{n} w_i x_i$ is theoretically sound as long as the range $[0, 2c]$ of the differences in firing times of the n presynaptic neurons $a_1, \dots, a_n$ is so small that there is a time point at which — in spite of their temporal differences

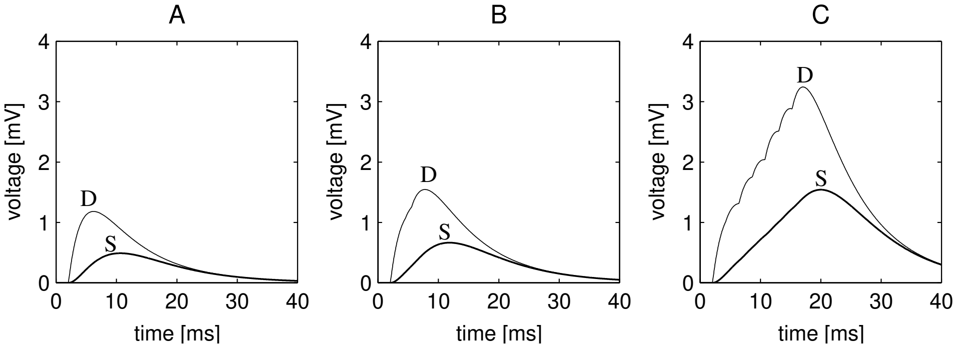

Figure 1. Superposition of several non-NMDA EPSP's caused by a single spike from a presynaptic neuron v_j. The membrane voltage of neuron v_i is measured at the dendrite (D) and at the soma (S). Three cases are shown: a single synapse (A), three synapses (B) and six synapses (C) from the presynaptic neuron. In all subsequent simulations we use three synapses, which results in a time interval of 7 ms during which the superposition of these EPSP's increases linearly.

— the resulting postsynaptic potentials are at the soma of v *all* in their initial *linear* phase. However, since non-NMDA EPSP's rise very fast, the theory would suggest that the length $2c$ of this interval would have to be chosen around 1 ms. This is so small that one has to be concerned about the effect of various sources of temporal jitter on the precision of this temporal coding. Figure 1 shows that for GENESIS-neurons the value $2\,c$ can be chosen 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 EPSP's 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[9, 10] (time from generation of a spike at neuron v_j to the onset of the PSP at the soma at neuron v_i).

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[11] so that 9 arbitrarily chosen vectors from $\{-1, 1\}^{60}$ (see Figure 2) are fixpoints of $\mathcal{H}$. We have simulated the network $\mathcal{S}_{\mathcal{H}}$ of spiking neurons (that simulates $\mathcal{H}$ according to theorem 1) in GENESIS with a network $\mathcal{G}_{\mathcal{H}}$ of 61 neurons. In $\mathcal{G}_{\mathcal{H}}$ each neuron is modeled with 122 compartments, and there are 3 synapses between each pair of neurons (whose delays differ by up to 3.5 msec). $\mathcal{G}_{\mathcal{H}}$ 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 modeled by inhibitory synapses.

In the protocols of our simulations (Figures 3 and 4) of $\mathcal{G}_{\mathcal{H}}$ the bars to the left ("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

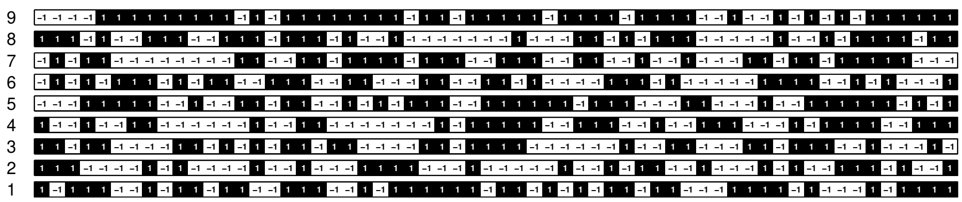

Figure 2. 9 randomly chosen "memory patterns" from $\{-1, +1\}^{60}$.

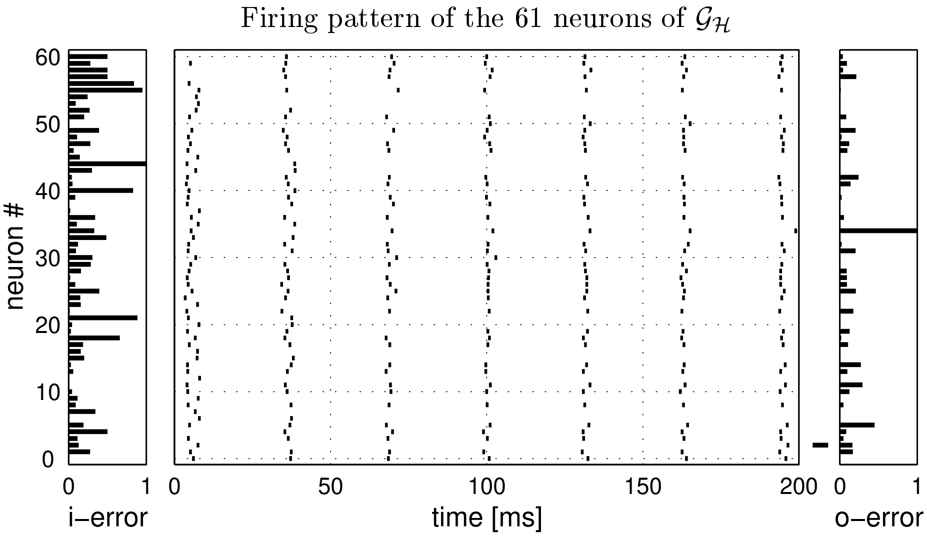

Figure 3. GENESIS-simulation of $\mathcal{G}_{\mathcal{H}}$ with a corrupted version $\mathbf{x}$ of memory pattern 1 as input (see text for details). The resulting effective difference between the *input* vector $\mathbf{x} \in [-1,+1]^{60}$ that is given to $\mathcal{G}_{\mathcal{H}}$ in temporal coding and memory pattern 1 is indicated to 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 $\mathbf{x}$ and the i-th component of memory pattern 1. The difference between the *output* vector $\mathbf{y} \in [-1,+1]^{60}$ in temporal coding and memory pattern 1 is indicated to the right ("o-error").

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 a non-firing of a neuron is treated like a firing at the very end of the firing wave. This is justified by the observation that both scenarios reflect equivalent ways of encoding the smallest possible analog value "-1"[5]. A neuron v_j that fires *very late* during a firing wave has an equally negligible impact an the firing times during the next firing wave as if v_j would not have fired at all. It is interesting to note that in all our simulation both types of encoding "-1" are present.

In the space right before "o-error" in Figures 3 and 4 we have marked with a horizontal bar each neuron i of $\mathcal{G}_{\mathcal{H}}$ which does *not* satisfy 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".

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 perturbated): a noisy membrane potential and failing synapses. These two types of noise are underlying all our simulations, which shows that our construction is robust against a substantial level of noise. We assume the existence of a random current at the soma of each neuron (Gaussian distribution with a variance of $10^{-18}\,\mathrm{A}^2$), and that each synapse fails with a probability of 15%.

The results of the simulation of $\mathcal{G}_{\mathcal{H}}$ with a *very* noisy version of memory pattern 1 as input is shown in Figure 3. The input $\mathbf{x}$ was constructed as follows from memory pattern 1: randomly chosen 15% of the components of memory pattern 1 were multiplied with -1. On top of this, the value of each component of $\mathbf{x}$ was moved by a random value from $[-0.4, 0.4]$. This input $\mathbf{x}$ was presented to $\mathcal{G}_{\mathcal{H}}$ in *temporal coding*, although not precisely: randomly chosen 10% of the neurons in $\mathcal{G}_{\mathcal{H}}$ were prevented from firing (corresponding to an input value -1). Further simulations[7] verify that in fact *all arbitrarily chosen fixpoints of $\mathcal{H}$* can be associatively recalled by $\mathcal{G}_{\mathcal{H}}$. Figure 3 shows that after 6 firing waves almost no "digital" errors in associative recall in terms

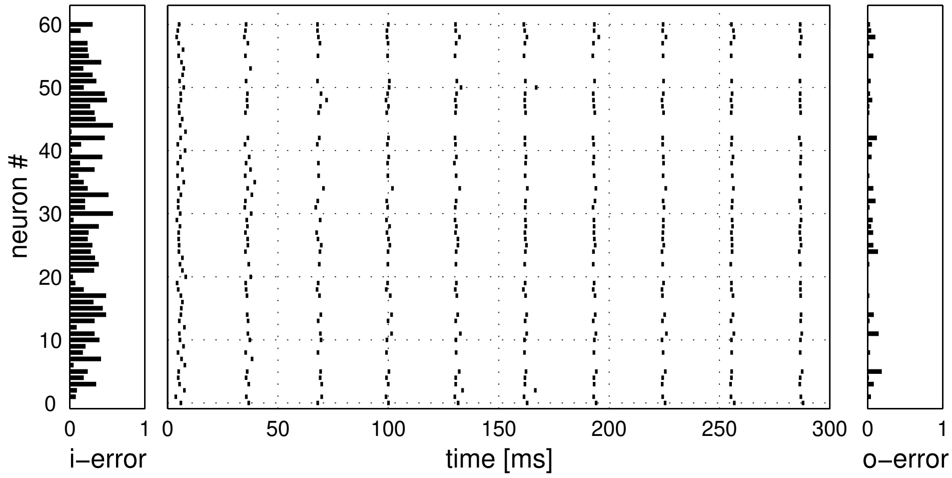

Figure 4. Here the input to $\mathcal{G}_{\mathcal{H}}$ is a linear combination of memory pattern 1 with factor 0.5 and of memory pattern 4 with factor 0.3, presented in temporal coding, with noise added as in Figure 3. The error bars to the left indicate the difference between the actual input and memory pattern 1 – which is its stronger component. The error bars to the right indicate the difference between the output and memory pattern 1 in temporal coding.

of firing/non-firing occur.

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 $k\,T$ for $k = 0, 1, 2, \ldots$. Instead, neuron 0 of $\mathcal{G}_{\mathcal{H}}$, whose firing times provide the reference times $k\,T$, 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 (see Figure 4). This is in contrast to a pattern recognition system comprised of spiking neurons proposed by Hopfield[1].

In $\mathcal{G}_{\mathcal{H}}$ we employ large transmission delays of 25 ms in order to avoid refractory effects and we modeled each negative weight w_{ij} form the network $\mathcal{H}$ through separate inhibitory synapses. One may argue that these two assumptions are biologically rather implausible. To overcome this we have replaced the completely connected network $\mathcal{G}_{\mathcal{H}}$ by a layered network $\mathcal{G}_{\mathcal{H}}^{synfire}$, with transmission delays of 5 ms between neurons in successive layers[7]. In $\mathcal{G}_{\mathcal{H}}^{synfire}$ negative weights w_{ij} are not modeled explicitly. Instead an inhibitory bias (equal for all neurons) is applied. This "global" inhibition is produced by synapses formed between the additional neuron 0 and each neuron i. Thus all information about the weights w_{ij} of the simulated Hopfield net $\mathcal{H}$ is stored exclusively in the efficacies of *excitatory* synapses in $\mathcal{G}_{\mathcal{H}}^{synfire}$. Simulations have shown that the layered network, has the same ability as $\mathcal{G}_{\mathcal{H}}$ for associative recall of any of the 9 given memory patterns[7].

CONCLUSIONS

We have exhibited a new way of emulating a Hopfield net in temporal coding with spiking neurons, that differs in essential aspects from previous models for associative memory in networks of spiking neurons[12, 13, 14, 4, 15, 16]. The resulting networks of spiking neurons carry out their computation very fast, even with biologically realistic low firing rates. This emulation is based on a rigorous theoretical result, and our GENESIS-simulations show that it is quite noise-robust.

The layered version of the resulting network of spiking neurons provides a refinement of the synfire-chain model[17]. This refined version of a synfire chain overcomes one of the essential bottlenecks of the original synfire chain model: its low storage capacity. Whereas the original version of a synfire chain has only two states (active and inactive), the here considered layered networks (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[18]) also support the existence of our refined version in the same biological neural system.

REFERENCES

1. J. Hopfield. Pattern recognition computing using action potential timing for stimulus representation. *Nature*, 376:33–36, 1995.
2. S. T. Thorpe and M. Imbert. Biological constraints on connectionist modelling. In R. Pfeifer, Z. Schreter, F. Fogelman-Soulie, and L. Steels, editors, *Connectionism in Perspective.* Elsevier, 1989.
3. T. W. Kjaer, T. J. Gawne, and B. J. Richmond. Latency: Another potential code for feature binding in striate cortex. *Journal of Neurophysiology*, 76(2):1356–1360, 1996.
4. W. Gerstner and J. L. van Hemmen. Associative memory in a network of "spiking" neurons. *Network*, 3:139–164, 1992.
5. W. Maass. Fast sigmoidal networks via spiking neurons. *Neural Computation*, 9:279–304, 1997.
6. J. J. Hopfield. Neurons with graded response have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences, USA*, 81:3088–3092, 1984.
7. W. Maass and T. Natschläger. Networks of spiking neurons can emulate arbitrary Hopfield networks in temporal coding. *Network: Computation in Neural Systems*, 8(4):355–372, 1997.
8. J. M. Bower and D. Beeman. *The Book of GENESIS: Exploring Realistic Neural Models with the GEneral NEural SImulation System.* Springer-Verlag, Inc. Published by TELOS, New York, 1995.
9. A. Zador, H. Agmon-Snir, and I. Segev. The morphoelectronic transform: A graphical approach to dendritic function. *The Journal of Neuroscience*, 15(3):1669–1682, 1995.
10. Y. Manor, C. Koch, and I. Segev. Effect of geometrical irregularities on propagation delay in axonal trees. *Biophysical Journal*, 60:1424–1437, 1991.
11. J. Hertz, A. Krogh, and R. G. Palmer. *Introduction to the Theory of Neural Computation.* Addison-Wesley, 1991.
12. E. Fransén. *Biophysical Simulation of Cortical Associative Memory.* PhD thesis, Stockholm University, October 1996.
13. M. W. Simmen, E. T. Rolls, and A. Treves. Rapid retrieval in an autoassociative network of spiking neurons. In J. M. Bower, editor, *Computational Neuroscience*, pages 273–278. Academic Press, London, 1995.
14. R. Ritz, W. Gerstner, U. Fuentes, and J. L. van Hemmen. A biologically motivated and analytically soluble model of collective oscillations in the cortex. *Biological Cybernetics*, 71:349–358, 1994.
15. A. Lansner and E. Fransén. Modelling hebbian cell assemblies comprised of cortical neurons. *Network*, 3:105–119, 1992.

16. A. V. M. Herz, Z. Li, and J. L. van Hemmen. Statistical mechanics of temporal association in neural networks with transmission delays. *Physical Review Letters*, 66(10):1370–1373, 1991.
17. M. Abeles. *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge University Press, 1991.
18. M. Abeles, H. Bergman, E. Margalit, and E. Vaadia. Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *Journal of Neurophysiologie*, 70(4):1629–1638, October 1993.