Computing with Spikes

Wolfgang Maass

Institut fuer Grundlagen der Informationsverarbeitung
Technische Universitaet Graz
Inffeldgasse 16b
A-8010 Graz, Austria
maass@igi.tu-graz.ac.at
http://www.igi.TUGraz.at/maass

A frightening thought for a computer scientist is that there might be completely different ways of designing computing machinery, that we may miss by focusing on incremental improvements of current designs. In fact, we *know* that there exist much better design strategies, since the human brain has information processing capabilities that are in many aspects superior to our current computers. Furthermore the brain does not require expensive programming, debugging or replacement of failed parts, and it consumes just 10-20 Watts of energy. Unfortunately, most information processing strategies of the brain are still a mystery. In spite of many positive reports in the media, even the most basic questions concerning the organization of the computational units of the brain and the way in which the brain implements learning and memory, have not yet been answered. They are waiting to be unraveled by concerted efforts of scientists from many disciplines. Computer science is one of the disciplines from which substantial contributions are expected, and in fact other countries have established already hundreds of research facilities in the new hybrid discipline *Computational Neuroscience*[1], which is dedicated to the investigation of computational principles in the brain. Computer scientists are contributing to these efforts through their experience in the evaluation of real and hypothetical systems that compute, as well as experience with robots and other machines that learn, move around, and explore their

[1] http://home.earthlink.net/~perlewitz/

environment with their "senses" (vision, touch, hearing, etc.). In addition, an interesting theory called *computational complexity theory* has emerged in computer science during the last two decades. This theory provides analytical tools for judging whether a hypothetical model for the organization of computing and learning in the brain scales up from a "heuristic scale", where one draws wiring diagrams for a few dozens neurons and tests this model on simple computational tasks, to the real scale of the brain where more than 10^{10} neurons process in real-time a continuous stream of information from 10^{8} sensory neurons, with instant (well, in some cases almost instant) access to a giant reservoir of stored information which is continuously updated. These experiences, tools, and results from computer science start to broaden the empirical basis for investigating terms from neuroscience and psychology such as learning, perception, and cognition. We are now in a position where we can test for example certain theories about learning by building a computer system based on this theory, and by evaluating how it performs. In fact, computer science provides even faster tools that allow us to test various theoretical ideas regarding cognition and learning through computer simulations on standardized benchmark tasks within a few hours or days. Obviously, having more opportunities for falsifying theories about the organization of information processing in the human brain is very desirable in the light of Popper's theory of science [15], and it creates new chances for approaching age-old questions about cognition and learning with more precise methods.[2]

Our institute contributes to research on neural computation through theoretical work based on complexity theory and computational learning theory, and through its interdisciplinary collaboration with Prof. Henry Markram's experimental neurobiology laboratory, which is moving this year from the Weizmann Institute in Israel to the new Brain Mind Institute at the Ecole Polytechnique Federale De Lausanne. Very important for our institute is also the ongoing exchange of ideas and methods with Prof. Rodney Douglas and

[2] Apart from the current work on computational models for neural systems that is sketched in this article, there also exists the area of *(artificial) neural networks*. This area had emerged several decades ago as a first attempt to capture aspects of neural computation in formal mathematical models. Neural networks have provided fruitful new ideas for machine learning and pattern recognition, but they are no longer viewed as an up-to-date way of modeling information processing in biological neural systems. Their computational units and organizational structure are too different, more precisely they reflect more ideas from computer science than from biology.

the other experts for neurotechnology at the Institut für Neuroinformatik an der ETH Zürich.[3] Currently our joint work concentrates on three fundamental questions:

- What exactly is the computational function of neural microcircuits, which apparently represent the basic building blocks ("chips") from which our brains are built?

- What are the principles by which learning and memory are organized in these neural microcircuits?

- What are the organizational principles by which millions of such neural microcircuits can communicate and collaborate in our brains?

Together we have discovered last year a new way of thinking about neural computation, see [11] and the articles by Henry Markram and Natschlaeger et al. in this issue, that differs radically both from traditional ways of designing computers, and from preceding models for neural computation (which obviously are strongly inspired by current computer designs). This new approach takes into account that neural circuits acquire through their numerous feedback connections ("recurrent connectivity") and the rich temporal dynamics of their components a computational life of their own, which cannot be understood by just looking at their components, neurons and synapses. Since this setup is so different from computational systems that can be analyzed with models related to Turing machines, we have developed a new conceptual framework. Based on this approach we can now for the first time carry out complex real-time information processing tasks on quite realistic computer models of neural microcircuits (see [11], [5]). In order to understand the ideas behind this work, one first needs to get a rough idea about neurons and synapses, the components of neural microcircuits. Therefore the remainder of this article is devoted to a quick introduction to *wetware*.

If you pour water over your PC, the PC will stop working. This is because very late in the history of computing – which started about 500 million years ago[4] in and near to the sea – the

[3] In contrast to our country, Switzerland is expecting major scientific and technological benefits from research in these areas – it has established during the past 10 years several new institutes in the areas Computational Neuroscience and Neuroinformatics, with well over a hundred researchers.

[4] According to current scientific knowledge about 500 million years ago, during heavy evolutionary pressure on all living organisms caused by drastic climatic changes, the first organisms with a nervous system emerged. But

PC and other devices for information processing were developed that require a dry environment. We still carry an echo of this history of computing in our heads: the neurons in our brain are embedded into an artificial sea-environment, the salty aqueous extracellular fluid which surrounds the neurons in our brain. The close relationship between the wetware in our brain, and the wetware in evolutionary much older organisms that still live in the sea, is

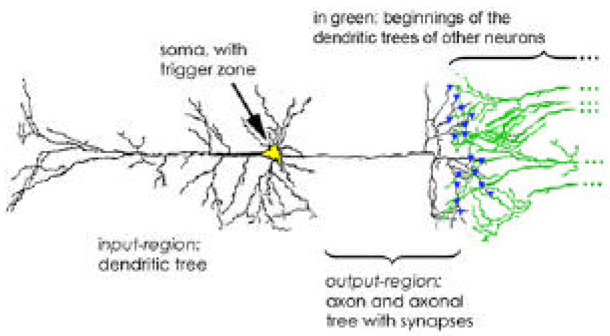

Figure 1: Simplified drawing of a neuron, with input region, the cell body or soma (the trigger zone lies at its right end of the soma, just where the axon begins), and output region. Synapses are indicated by blue triangles.

actually quite helpful for research. Neurons in the squid are 100 to 1000 times larger than the neurons in our brain, and therefore easier to study. Nevertheless the mathematical equations that Hodgkin and Huxley derived to model the dynamics of the neuron that controls the escape reflex of the squid (for which they received the Nobel prize in 1963), also apply to the neurons in the human brain.

One of the technical problems that nature had to solve for enabling computation in wetware was how to communicate intermediate results from the computation of one neuron to other neurons, or to output-devices such as muscles. In a PC one sends streams of bits over copper wires. But copper wires were not available a few hundred million years ago, and they also do not work so well in a sea-environment. The solution that nature found was the so-called action potential or *spike.* The spike plays in a brain a similar role as a bit in a digital computer: it is the common unit of information in wetware. A spike is a sudden voltage increase (see Figure 2) for about 1 ms (1 ms = 1/1000 second) that is created at the cell body (soma) of a neuron,

one could also argue that the history of computing started somewhat earlier, even before there existed any more precisely at its *trigger zone*, and propagated along a lengthy fiber (called axon) that extends from the cell body. This axon corresponds to an insulated copper wire in hardware. It contains active zones (the nodes of Ranvier), where the shape of the propagated spikes is restored. The gray matter of the human brain contains a large amount of such axons: about 4

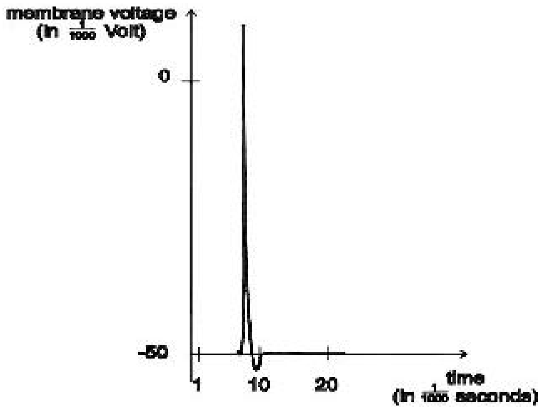

Figure 2: Time cause of an action potential at the soma of a neuron.

km in every cubic millimeter. Axons have numerous branching points (see the axonal tree on the right hand side of Figure 1), where most spikes are duplicated, so that they can enter each branch of the axonal tree. In this way a spike from a single neuron can be transmitted to a few thousand other neurons. But in order to move from one neuron to another, the spike has to pass a rather complicated information transmission device, a so-called *synapse* (marked by a blue triangle in Figure 1, and shown in more detail in Figure 3). When a spike enters a synapse, it is likely to trigger a complex chain of events that are indicated in Figure 3[5]: a small vesicle filled with special molecules ("neurotransmitter") is fused with the cell membrane of the presynaptic terminal, thereby releasing its content, the neurotransmitter, into the extracellular fluid. Whenever a neurotransmitter molecule reaches a particular molecular arrangement (a "receptor") in the cell membrane of the next neuron, it will open a channel in that cell membrane through which charged particles (ions) can enter the next cell. This causes an increase or decrease (depending on the type of channel that is opened and the types of ions that this channel lets through) of the membrane voltage by a few millivolt (1 millivolt = 1/1000 volt). One calls these potential changes EPSPs (excitatory postsynaptic potentials) if they increase the membrane voltage, and IPSPs (inhibitory postsynaptic potentials) if they

nervous systems: 3 to 4 billion years ago when nature discovered information processing and storage via RNA.

[5] See http://www.wwnorton.com/gleitman/ch2/tutorials/2tut5.htm for an online animation.

decrease the membrane potential. In contrast to the spikes, which all look alike, the size and shape of these postsynaptic potentials depends very much on the particular synapse that causes it. In fact it will also depend on the current "mood" and the recent "experiences" of this

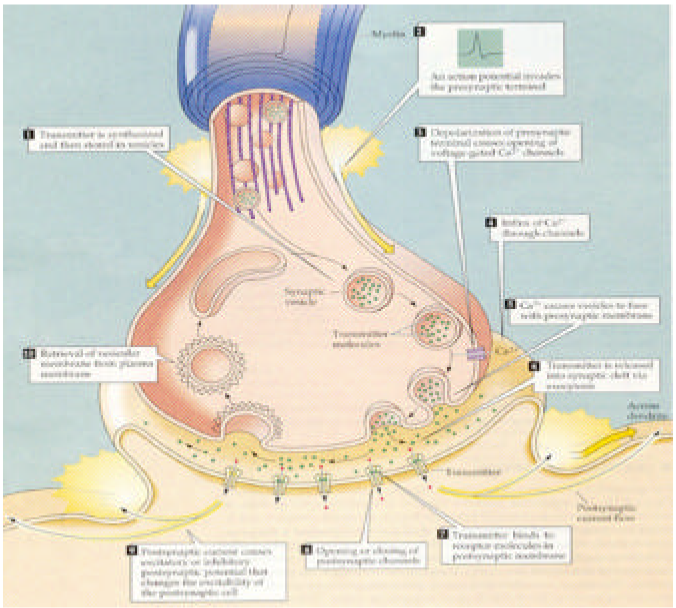

Figure 3: The spike enters the presynaptic terminal, which is an endpoint of the axonal tree of the preceding neuron, shown on the top. It may cause a vesicle filled with neurotransmitter to fuse with the cell membrane, and to release its neurotransmitter molecules into the small gap (synaptic cleft) to the cell membrane of the next neuron (called postsynaptic neuron in this context), which is shown at the bottom. If the neurotransmitter reaches a receptor of the postsynaptic neuron, a channel will be opened that lets charged particles pass into the postsynaptic neuron. Empty vesicles are recycled by the presynaptic terminal.

synapse, since the postsynaptic potentials have different sizes, depending on the pattern of spikes that have reached the synapse in the past, on the interaction of these spikes with the firing activity of the postsynaptic neuron, and also on other signals that reach the synapse in the form of various molecules (e.g. neurohormones) through the extracellular fluid. Hence a synapse is not a static signal transformer like a modem, but rather an intelligent preprocessor, that modulates the currently transmitted information in the light of experiences stretching from the recent past way back into the history of the organism. The human brain contains about 10^{15} of these synapses.

Sometimes people wonder whether it is possible to replace wetware by hardware, to replace for example parts of a brain by silicon chips. This is not so easy because wetware does not consist of fixed computational components, like a silicon chip, that perform the same operation in the same way on every day of its working life. Instead, the channels and receptors of neurons and synapses move around, disappear, and are replaced by new and

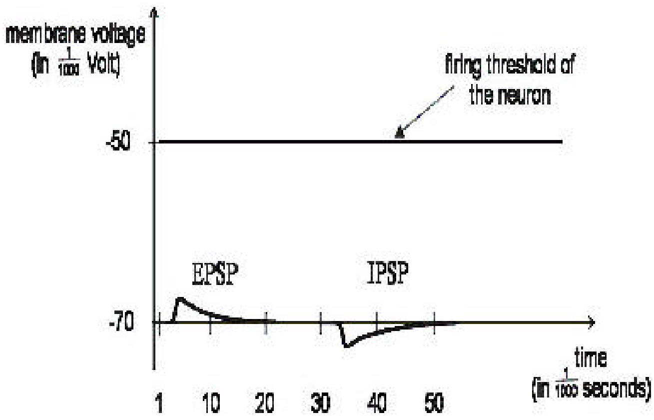

Figure 4: Postsynaptic potentials are either excitatory (EPSP) or inhibitory (IPSP). The membrane voltage is a sum of many such postsynaptic potentials. As soon as this sum reaches the firing threshold, the neuron fires.

possibly different receptors and channels that are continuously reproduced by a living cell in dependence of the individual "experience" of that cell (such as the firing patterns of the pre- and postsynaptic neuron, and the cocktail of biochemical substances that reach the cell through the extracellular fluid). This implies that next year a synapse in your brain is likely to perform its operations quite differently from today, whereas a silicon clone of your brain would be stuck with the "old" synapses from this year. Thus we would need new types of *adaptive* hardware to replace the function of neural circuits in wetware.

The postsynaptic potentials created by the roughly 10.000 synapses converging on a single neuron are transmitted by a tree of input wires ("dendritic tree", see Figure 1) to the *trigger zone* at the cell body of a neuron. Whenever the sum of these hundreds and thousands of continuously arriving voltage changes reaches the firing threshold at the trigger zone, the neuron will "fire" (a chain reaction orchestrated through the rapid opening of channels in the cell membrane that allow positively charged sodium ions to enter the neuron, thereby increasing the membrane voltage, which causes further channels to open) and send out a spike through its axon.[6] So we are back at our starting point, the spike. Altogether one might argue

[6] See http://www.wwnorton.com/gleitman/ch2/tutorials/2tut2.htm for an online animation.

that the all-or-nothing nature of a spike (since there are no half-spikes etc.) points to a *digital* type of computation in neural systems. But on the other hand in the absence of a global clock or synchronization device the time when a spike is sent is an important *analog* signal. Furthermore we have seen that numerous analog processes (especially at the synapses) modulate the effects that a spike has on other neurons, and thereby the time when these other neurons will send out their next spike. Thus neural computation turns out to be a hybrid mixture of analog and digital computing.

The question is now how a *network of neurons* can compute with spikes. Figure 5 presents an illustration of a tiny network consisting of just 3 neurons, which communicate via

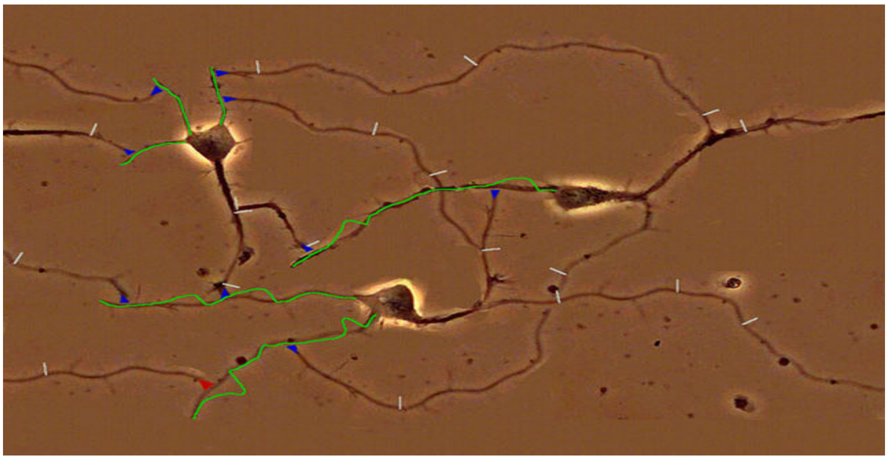

Figure 5: A simulated network of 3 neurons. Postsynaptic potentials in the input regions of the neurons (dendritic tree) are indicated by green curves. Spikes are indicated by white bars on the axons, and synapses by blue triangles. In the online available computer installation you can create your own input spike train, and watch the response of the network. You can also change the strength of the synapses, a process that may correspond to learning in your brain. See [7], (http://www.igi.TUGraz.at/maass/118/118.html) for more detailed information.

sequences of spikes (usually referred to as *spike trains*). It is taken from an animated computer installation which is online available[7]. It allows you to create your own spike train, and watch how the network responds to it. You can also change the strength of the synapses, and thereby simulate (in an extremely simplified manner) processes that take place when the

[7] See http://www.igi.TUGraz.at/demos/index.html. This computer installation was programmed by Thomas Natschlaeger and Harald Burgsteiner, with support from the Steiermaerkische Landesregierung. Detailed explanations and instructions are online available from http://www.igi.TUGraz.at/maass/118/118.html, see [7]. Further background information is online available from [14], [6], [8].

neural system "learns". But we still do not know how to transmit information via spikes, hence let us look at the protocol of a real computation in wetware.

In Figure 6 the spike trains emitted by 30 (randomly selected) neurons in the visual area of a monkey brain are shown for a period of 4 seconds. All information of your senses, all your ideas and thoughts are coded in a similar fashion by spike trains. If you would for example make a protocol of all the visual information which reaches your brain within 4 seconds, you would arrive at a similar figure, but with 1.000.000 rows instead of 30, because the visual information is transmitted from the retina of your eye to your brain by the axons of about 1.000.000 neurons.

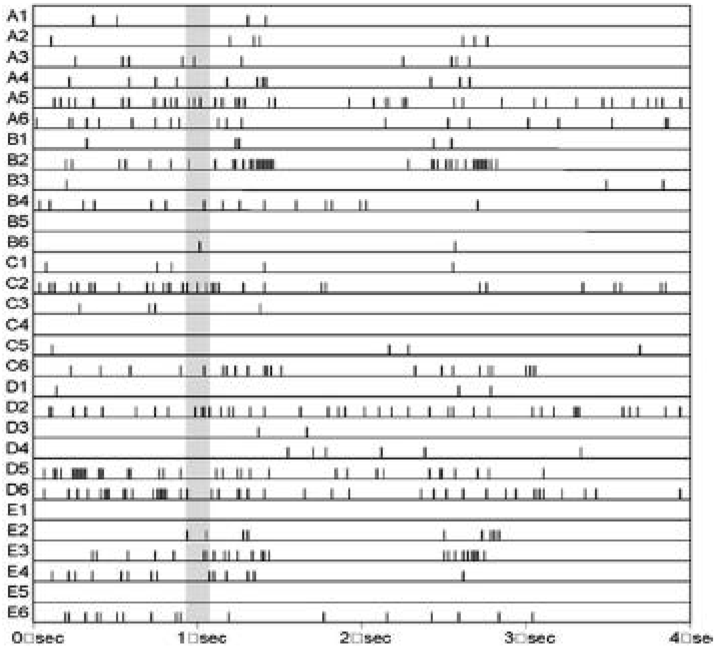

Figure 6: Recording of spike trains from 30 neurons in the visual cortex of a monkey over 4 seconds. The 30 neurons were labeled A1 – E6 (the labels are given on the left of the figure). The spike train emitted by each of the neurons is recorded in a separate row. Each point in time when a neuron fires and thereby emits a spike is marked by a small vertical bar. Hence one could read this figure like a music score.

The time needed for a typical fast computation in the human brain is marked by a vertical gray bar; these are 150ms. Within this time the human brain can solve quite complex information processing problems, like for example the recognition of a face. Currently our computers would need a substantially longer computation time for such tasks.

Researchers used to think that the only computationally relevant signal in the output of a neuron was the frequency of its firing. But you may notice in Figure 6 that the frequency of firing of a neuron tends to change rapidly, and that the temporal distances between the spikes are so irregular that you have a hard time estimating the average frequency of firing of a neuron by looking at just 2 or 3 spikes from that neuron. On the other hand the human brain can compute quite fast, in about 150 ms, with just 2 or 3 spikes per neuron. Hence other features of spike trains must be used by the brain for transmitting information. Recent experimental studies (see for example [17], [3], [16]) show that in fact the full spatial and temporal pattern of spikes emitted by neurons is relevant for the message which they are sending to other neurons. Hence it would be more appropriate to compare the output of a collection of neurons with a piece of music played by an orchestra. To recognize such piece of music it does not suffice to know how often each note is played by each musician. Instead we have to know how the notes of the musicians are embedded into the melody and into the pattern of notes played by other musicians. One assumes now that in a similar manner many groups of neurons in the brain code their information through the pattern in which each neuron fires relative to the other neurons in the group. Hence, one may argue that music is a code that is much closer related to the codes used by the human brain than the bit-stream code used by a PC.

The investigation of theoretical and practical possibilities to compute with such spatio-temporal patterns of pulses has lead to the creation of a new generation of artificial neural networks, so-called *pulsbased* neural networks (see [9] and [2] for surveys and recent research results). Such networks are appearing now also in the form of novel electronic hardware [12], [1], [13]. An interesting feature of these pulsbased neural networks is that they do not require a global synchronisation (like a PC, or a traditional artificial neural network). Therefore they can save a lot of energy[8], since no clock signal has to be transmitted at every moment to all components of the network, even if there is currently nothing to compute. Furthermore neurons can use *time* as a new dimension for *coding information*, for example by firing simultaneously or in a specific temporal pattern. One big open problem is the *organization of computation* in such systems, since the operating system of wetware is still unknown, even for the squid. Hence our current research, jointly with the neurobiologist Prof. Henry Markram, concentrates on the organization of computations in neural microcircuits, the lowest level of circuit architecture in the brain. Our approach is based on a novel theoretical model, the liquid

[8] Wetware consumes much less energy than any hardware that is currently available. Our brain, which has about as many computational units as a very large supercomputer, consumes just 10 to 20 Watt.

state machine, which strongly differs from Turing machines and other traditional computational models. More information on this recent development can be found in the subsequent two articles of this volume, as well as in the research reports [11], [10], [5].

Acknowledgement:

I would like to thank Margot Göttsberger and Dr. Thomas Natschläger for helpful comments on an earlier draft of this paper.

References

[1] Deiss, S. R., Douglas, R. J., and Whatley, A. M. (1999). *A pulse-coded communications infrastructure for neuromorphic systems*. In Maass, W., and Bishop, C., editors, Pulsed Neural Networks. MIT-Press, Cambridge.

[2] Grossberg, S., Maass, W., and Markram, H., editors (2001). Spiking Neurons in Neuroscience and Technology. *Special Issue of Neural Networks*, vol. 14 (6-7).

[3] Koch, C. (1999). *Biophysics of Computation: Information Processing in Single Neurons*. Oxford University Press (Oxford).

[4] Krüger, J., and Aiple, F. (1988). *Multielectrode investigation of monkey striate cortex: spike train correlations in the infragranular layers*. Neurophysiology, 60:798-828.

[5] Legenstein, R., Maass, W., and Markram, H. (2002). *Input Prediction and Autonomous Movement Analysis in Recurrent Circuits of Spiking Neurons*, submitted for publication.

[6] Maass, W. (2000a). *Das menschliche Gehirn – nur ein Rechner?*. In Burkard, R. E., Maass, W., and P. Weibel, editors, Zur Kunst des Formalen Denkens, pages 209-233. Passagen Verlag (Wien), 2000.
See http://www.igi.tugraz.at/maass/KFD/

[7] Maass, W. (2000b). *Spike trains – im Rhythmus neuronaler Zellen*. In Kriesche, R., and Konrad, H., editors, Katalog der steirischen Landesausstellung gr2000az, pages 36-42. Springer Verlag.
See # 118 on http://www.igi.tugraz.at/maass/publications.html

[8] Maass, W. (2002). *Paradigms for computing with spiking neurons*. In Leo van Hemmen, editor, Models of Neural Networks, volume 4. Springer (Berlin), to appear.
See # 110 on http://www.igi.tugraz.at/maass/publications.html

[9] Maass, W., and Bishop, C., editors (1999). *Pulsed Neural Networks.* MIT-Press (Cambridge, MA). Paperback (2001).
See http://www.igi.tugraz.at/maass/PNN.html

[10] Maass, W., and Markram, H. (2002) *On the Computational Power of Recurrent Circuits of Spiking Neurons*, submitted for publication.

[11] Maass, W., Natschlaeger, T., and Markram, H. (2001). *Real-time computing without stable states: a new framework for neural computation based on perturbations*, submitted for publication.

[12] Mead, C. (1989). *Analog VLSI and Neural Systems.* Addison-Wesley (Reading).

[13] Murray, A. F. (1999). *Pulse-based computation in VLSI neural networks.* In Maass, W., and Bishop, C., editors, Pulsed Neural Networks. MIT-Press (Cambridge, MA).

[14] Natschlaeger, T. (1996). *Die dritte Generation von Modellen für neuronale Netzwerke Netzwerke von Spiking Neuronen*. In: Jenseits von Kunst, Peter Weibel editor. Passagen Verlag.

See http://www.igi.tugraz.at/tnatschl/

[15] Popper, K. (1959). *The Logic of Scientific Discovery*. Hutchinson (London).

[16] Recce, M. (1999). *Encoding information in neuronal activity.* In Maass, W., and Bishop, C., editors, Pulsed Neural Networks. MIT-Press (Cambridge, MA).

[17] Rieke, F., Warland, D., Bialek, W., and de Ruyter van Steveninck, R. (1997). *SPIKES: Exploring the Neural Code*. MIT-Press (Cambridge, MA).