

Implementation of Neuromorphic Systems: from discrete components to analog VLSI chips (testing and communication issues)

Vittorio Dante, Paolo Del Giudice and Maurizio Mattia
Physics Laboratory, Istituto Superiore di Sanità - Roma, Italy

Per eventuale corrispondenza: dr. Paolo Del Giudice
Laboratorio di Fisica – Istituto Superiore di Sanità
Viale Regina Elena, 299 – 00161 Roma

Tel.: 06 49902245
Fax: 06 49387075
e-mail: paolo@ibmteo.iss.infn.it

Riassunto

realizzazione di sistemi neuromorfi: da dispositivi a componenti discreti a chip analogici VLSI (problemi di test e comunicazione)

In questo lavoro si passano in rassegna dispositivi elettronici volti ad emulare in qualche misura la struttura e le funzioni di semplici sistemi neuronali, con particolare enfasi sui problemi di comunicazione. Viengono innanzitutto ricapitolate le caratteristiche generali di tali dispositivi, cosiddetti 'neuromorfi', e i problemi connessi ai 'test' effettuati su di essi. Si passa quindi agli sviluppi direttamente legati alla attività svolta presso l'ISS: una prima rete neuronale elettronica che realizza un semplice classificatore, in grado di sviluppare in modo autonomo delle rappresentazioni interne degli stimoli ricevuti; una rete di output, che raccoglie informazioni dal classificatore precedente e ne estrae la parte rilevante da trasmettere all'osservatore; un chip neuronale analogico VLSI, che realizza una rete ricorrente di neuroni impulsanti e sinapsi plastiche, con il sistema di test associato; una scheda progettata per interfacciare il bus standard PCI di un PC con un bus asincrono ideato specificamente per la comunicazione tra chip neuromorfi; un sistema di prova, orientato ad una applicazione, che si basa sulla precedente infrastruttura di comunicazione.

Parole chiave: sistemi neuromorfi, reti neuronali, apprendimento on-chip, comunicazione asincrona

Abstract

We review a series of implementations of electronic devices aiming at imitating to some extent structure and function of simple neural systems, with particular emphasis on communication issues. We first provide a short overview of general features of such 'neuromorphic' devices and the implications of setting up 'tests' for them. We then review the developments directly related to our work at the ISS: a pilot electronic neural network implementing a simple classifier, autonomously developing internal representations of incoming stimuli; an output network, collecting information from the previous classifier and extracting the relevant part to be forwarded to the observer; an analog, VLSI (Very Large Scale Integration) neural chip implementing a recurrent network of spiking neurons and plastic synapses, and the test setup for it; a board designed to interface the standard PCI (Peripheral Component Interconnect) bus of a PC with a special purpose, asynchronous bus for communication among neuromorphic chips; a short and preliminary account of an application-oriented device, taking advantage of the above communication infrastructure.

Key-words: neuromorphic systems, neural networks, on-chip learning, asynchronous communication

Introduction

Electronic devices trying to directly mimic the operation of some neural subsystem (be it a sensory function, or a higher cognitive function) are sometimes termed "neuromorphic", after Carver Mead pioneered their development in the early 90s [1,2].

While there is long-standing scientific tradition in the pursuit of a satisfactory modeling of the brain function, ranging from detailed, single neuron models up to models of higher cognitive functions at a very high level of abstraction, neuromorphic systems are meant to expose the implications of a physical realization of interacting neural and synaptic elements.

Discussing how and why this approach can be radically different from the ones based on simulating networks of interacting neurons on a general (or special) purpose computer would bring us beyond the scope of the present paper. It is worth mentioning, however, that the biological neural network is a *large, massively parallel* intricate web of *noisy, extremely low power, quite slow* elements, that manage to exploit the above features such as to outperform by far man-made machines designed for fast and reliable computation. Those features are shared to a large extent by networks of analog silicon neurons operating in a sub-threshold (low power consumption) regime [1], and open the way to a serious attempt at duplicating some key features of the biological system. In a sense, this amounts to take a step back from the abstract level of implementation-independent approach to intelligent computation that was the hallmark of artificial intelligence, hoping that by designing devices as close as possible to the choices that nature made we might have a chance to be naturally let to build efficient architectures for neural networks.

In principle, interesting neuromorphic systems are able to acquire "natural" external stimuli, somehow process them in view of specific needs and provide a usable response in real time. The above steps will be performed in general by different VLSI devices. So we are presented with a major issue: how to efficiently (and "naturally") manage the communication among neuromorphic devices that must operate in real time.

Some proposals in this direction appeared as early as 1993 [3], but real progress in this respect has been somewhat slow. We summarize in the final part of present paper some contributions that our group recently made in this area.

Besides communicating among themselves, neuromorphic chips should also be able to accept data from the experimenter, and return information about their state. This brings us to the issue of *testing* neuromorphic devices.

Conventional digital VLSI chips, complicated as they can be, are 'logically' easy to test; the chip implements a deterministic automata, and testing it, once the design is checked, means getting rid of fabrication defects, spurious effects which the simulators could not predict in detail and the like. But, as we will briefly discuss in the following, analog, VLSI neuromorphic devices we are interested in are *noisy*, to the point that instead of fighting against noise, it will be incorporated in some cases as a crucial dynamical ingredient of the system, and *complex*, meaning that emergent, collective properties appear as very indirect consequences of the single elements' structure. So *testing* such devices is actually nearer to an *experiment* than to a test, and the appropriate hardware and software have to be set up on purpose. Starting with simple, exploratory electronic neural networks implemented in discrete elements, up to analog, full custom

VLSI neural chips with $O(100)$ neurons and $O(1000)$ synapses, our group developed a series of associated setups for experimenting with those devices. We provide in this paper a short account of a series of 'test-oriented' workbenches, mostly developed at the Physics Laboratory of the ISS, and some lessons we learnt from performing experiments on a series of neural devices, whose purpose and main features we summarize along the paper.

Most developments we will mention took place in the frame of a wider group involving the physics departments of the Rome universities "La Sapienza" and "Tor Vergata", through a time span of several years; we acknowledge this by quoting the appropriate references along the paper.

LANN27: first steps towards an autonomous classifier

Our first steps towards proper, analog VLSI implementations of neuromorphic chips, were taken by designing and building simple, bulky pilot devices composed of discrete components and printed circuits, meant as a technologically simple arena for testing ideas and design principles.

The device LANN27 (Learning Attractor Neural Network) [4,5,6] was the first attempt towards the implementation of a neural network with high feedback, able to dynamically adjust the connections among its neurons ('synapses') such as to *learn* to operate as an associative classifier of a stereotyped, but essentially unconstrained, flux of stimuli. Learning in this context, and in the whole paper, is meant as the process by which each synapse, connecting a pair of neurons, undergoes changes in its *efficacy* (the strength with which the 'pre-synaptic' neuron affects the 'post-synaptic' one), reflecting the average activity of those neurons. In particular, synaptic potentiation is meant to implement the Hebb's suggestion [7,8], that for simultaneously active pre- and post-synaptic neurons the synapse connecting them gets strengthened (this was the key idea in the Hopfield model [9,10] and in many variations of it). Since subsets of neurons are supposed to be driven to high activity when receiving extra afferent current because of a stimulation, covariance of activation of pre- and post-synaptic neurons is a mean to leave a trace of the structure of external stimuli in the pattern of synaptic efficacies. Subsets of neurons reciprocally connected by strengthened synapses will tend to fire together at higher rate, thus foreshadowing a dynamical scenario of stimulus-selective activation of sub-populations of neurons. Even a cursory account of the issue of synaptic plasticity is beyond the scope of the present paper; we only mention that in the LANN27, and in most neuromorphic devices we will discuss, synapses implement a *stochastic* version of the Hebbian idea [11], meaning that changes in the synaptic state are only probabilistically driven by the pre- and post-synaptic states. The biologically meaningful dynamical scenario which underlies what follows includes: high feedback networks of spiking neurons, some of which reacting with enhanced firing rates to a stimulation, thereby provoking some changes in a subset of synapses; other stimuli provoke other, possibly conflicting, synaptic changes. As some stimuli start becoming 'familiar' (meaning that traces of their structure are embedded in the synaptic couplings), the network reacts to them going into a *recognition* state, that is the sub-population of neurons selective for a given stimulus, after its removal does not relax into the pre-learning, low activity, *spontaneous* state, but jumps into an *attractor*, reverberating state of enhanced activity expressing the selective response. The idea behind the associative memory behavior of the network is that whole classes of similar stimuli make the network develop an archetypical internal representation (all the stimuli belonging to that class elicit essentially the same attractor, post-stimulus network response), by extracting the prototype representation of that class.

It is to be stressed that in this scenario the whole process by which the synaptic couplings acquire a stimuli-related structure is totally *unsupervised*, that is no computational paradigm is forced upon the system from outside, and learning comes as a real self-organization process.

Stochasticity enters the synaptic model as a way to remove serious limitations in terms of *capacity* of the network (the number of different classes of stimuli simultaneously embedded in the synaptic couplings), which are encountered if one sticks to the (reasonable) hypothesis that the allowed values expressing the synaptic efficacies in the model form a discrete and finite set and at the same time synaptic changes are *deterministic*.

The LANN27 is a very simple pilot implementation of some of the above ideas. The device includes 27 binary neurons (only capable of quiescent or maximally active states) and 351 plastic, stochastic synapses. Its electronic realization, in discrete components, is fully analog and asynchronous. The stochastic synaptic dynamics is supported by a noise generator associated with each synapse: the source of noise is therefore *logically* external with respect to the system's dynamics.

As mentioned in the Introduction, testing such a device is a complex process and requires hardware and software tools designed on purpose. Testing the LANN27 required:

- the electronic setup and the appropriate software for managing communication to and from the network, e.g. for setting the network's parameters and for reading out the network's state.
- defining and implementing sequences of input stimuli such as to expose different features of the learning process.
- the characterization of the collective behavior of the neurons in the network, as shaped by the learning process: devising a compact and evocative description of the development of selective attractors (not a minor point, with such a high number of degrees of freedom even for such a small network); defining suitable observables for assessing the network's memory capacity; devising a compact representation of the network's state space.

While we refer the interested reader to [5] for details and results, we briefly mention here that, naive and small as it was, LANN27 provided a number of interesting lessons, both with its virtues and its limitations. The main lesson has to do with the noise generation related to the stochastic synaptic dynamics. The size and power consumption of the noise generators proved to be the main obstacle on the way to a VLSI implementation of an improved version of LANN27. It was therefore proposed [12] that, going from naive, binary neurons to the so-called integrate-and-fire, spiking neurons (see later), a *distributed* source of randomness could be found in the variability of spikes emission by the neurons in the network. Trying to solve a technical problem brought about a major theoretical improvement in the model, which is one example of the fruitful interaction and reciprocal fertilization which takes place between modeling and electronic implementation of neuromorphic devices.

From the point of view of communication issues, directly relevant for the present paper, the test setup for the LANN27 was quite simple: an home-made interface board handled serial communication with the network under the supervision of an Intel microcontroller. The related firmware and higher level software, besides taking care of parameters setting and basic communication functions, could manage multiple readout modes of the network's state. Fig. (1) shows the LANN27 rack, the hierarchical organization of the network and the interface board managing the PC- LANN27 communication.

Figure 1 about here

Reading the outcome of the classification process

Before describing the analog VLSI implementation of spiking networks, we mention a further development which stemmed from our experience with the LANN27. Whatever state the network goes into as a result of the stimulation, checking its properties entails reading out and inspecting the states of all its neurons for a suitable time interval. This is an unnatural feature, even for the simple pilot device. To overcome these limitations we undertook to implement an "output network", taking care of capturing the LANN27 state and providing the experimenter with a compact representation of the 'recognition state', if any, expressing the attractor network's reaction to some stimulation. In other words, the output network performs a kind of compression of the information provided by the LANN27. While such a behavior can be obtained using well established *supervised* strategies for the output network, it is a non trivial task to achieve this, again, through a totally unsupervised learning dynamics.

Figure 2 about here

Figure 3 about here

The neurons of the output network are connected by a set of plastic synapses to the 27 neurons of the LANN27. Furthermore, neurons in the output network inhibit each other¹, such that if for any reason one of them is driven to higher activity it will tend to shut off all the others, remaining alone to signal in an unambiguous way whatever agent caused its greater input (the so called "winner take all" model).

"Forward" synapses connecting the LANN27 neurons to the ones in the output network are plastic and are a variation of the stochastic synapses in the LANN27. If an attractor, reverberant state of the LANN27 happens to provide higher input to one of the output neurons (the initial state of the forward synapses is random), it will light up as the candidate 'representation' of that attractor, shutting off all the other neurons; furthermore, those forward synapses get strengthened, so as to make more likely that the same neuron will be activated by the same attractor. The model is such that [13,14,15] representations of different LANN27 attractors in the output network are orthogonal, so the compressed representation of the attractor structure of the LANN27 amounts to a number of output neurons equal to the capacity of LANN27, each of them lighting up to signal the LANN27 having relaxed in the corresponding attractor in response to a stimulus. Although some problems were left open by the implemented output network, this device successfully managed to autonomously develop the above compressed representation in simple situations.

The output network too was implemented in discrete components, and communication with the LANN27 relied upon the system sketched in Figs. (2,3). It is based on a board (standard VME) hosted by the output network's rack (to be called MASTER in the following) and a number of smaller, "SLAVE" boards taking care of the forward synapses readout. The latter boards communicate with the MASTER via an I²C serial bus. The MASTER board manages the serial communication with the PC, e.g. for parameters setting and forwarding of the information on the synaptic states. The SLAVE boards read out the synaptic state, and locally store it to later fulfill possible requests from the MASTER. The choice of the I²C serial bus was appropriate to balance the need of multiple and flexible readout mechanisms (e.g. timed or interrupt-driven readings) with that of avoiding cumbersome multiple connections.

"Neurophysiology on silicon": first experiments with neural chips of spiking neurons

Coming closer to biologically plausible (and computationally interesting) synthetic neural systems implies first of all endowing the neural network model with more flexible and realistic neural and synaptic elements. As for the former, the above mentioned integrate-and-fire [16] (hereafter IF) neuron has been, and still is, widely regarded as a useful compromise between simplicity and amenability to mathematical analysis on one side, and potential for expressing interesting collective behaviour on the other. The IF neuron is a model neuron whose 'membrane potential' integrates the afferent current (the succession of spikes impinging on it) and, if a threshold is crossed, emits of a spike. The main simplifications include: the lack of any spatial structure, the choice of membrane potential as the only state variable, the linearity of the dynamical equations, the *ad hoc* mechanism for spike emission. Yet, simple as they may seem, networks composed of such model neurons exhibit an incredibly rich behavior, and are still the main workhorse of modelers. Modeling the synaptic elements is a much harder task, since experiments are as yet much less effective in constraining the possible synaptic models than neural ones. Some general features, however, are shared by a large class of synaptic models, among which the above mentioned Hebbian-like mechanism for synaptic potentiation, and the fact that neural spikes trigger synaptic changes. We also anticipated that stochasticity of synaptic dynamics is a key ingredient in the neuromorphic devices we developed: indeed, a discrete and finite number of allowed states for the synapse (with its mentioned deadly effect on the capacity of the network, for deterministic synapse) is quite a natural assumption for any material device (biological or electronic). So, for the purpose of the present discussion we qualify our synaptic device as a *Hebbian and stochastic, spike driven* synapse with two allowed values for the efficacy (corresponding to a *potentiated* and a *depressed* state).

LANN21 [17] is a recurrent network composed of 21 neurons (14 excitatory neurons connected by 60 plastic synapses; 7 inhibitory neurons and 35 fixed synapses), implemented in full custom analog VLSI (CMOS² 1.2 μm , 1.5 x 1.38 mm), which incorporates the above features. Besides the recurrent current originated from the activity of the neurons in the network, neurons also receive external input from outside

¹ A further inhibitory neuron affecting all the neurons of the output network controls the overall level of activity.

² Complementary Metal Oxide Semiconductor.

the chip, meant to express, for example, incoming stimuli. These chips operate in the so called sub-threshold regime [1], which allows a very low power consumption.

Our experiments with this chip had the ambitious goal of reproducing in such a small dynamical system several aspects of the collective behavior expected in a large recurrent network of spiking neurons.

This of course required the design and the construction of a system for data acquisition, parameters setting, injection of external currents into the neurons in the chip. Furthermore, once one parameter has been set externally, it remains to be inferred how the setting translates into an effective parameter inside the chip. Let's say that we fixed the values of the higher and lower synaptic efficacies to some numbers; this implies in practice injecting an appropriate (and actually very small – nanoamperes - in most cases) current into selected pins of the chip; the current is then mirrored inside the chip, but many sources of inhomogeneities make the estimate of what will be circulating in the chip somewhat uncertain. So one has to resort to indirect methods to estimate the actual value of the parameter, the synaptic efficacies in this case, on the basis of the *dynamics* of the system: we parameterize the theoretical predictions on the neurons' behavior with the unknown jumps in their membrane potential associated with the efficacies to be estimated, in such a way that by extracting only information on the spikes emitted by the neurons in the network (the only observable in most neurophysiological experiments!) we can fit the *actual* value of the synaptic efficacies.

This is very much like what happens in a real experiment, and provides an example of the mentioned tricky nature of *testing* these devices.

As an example of more technical requirements that had to be met by the communication system, we mention that one usually needs to acquire data from multiple (digital and analog) channels, and common DAQ board are not suited for the purpose.

Fig.(4) illustrates the data acquisition and communication system for the LANN21 chip. A 'motherboard' hosts the pod for the neural chip, which is surrounded by an array of multi-purpose modules, which implement the needed analog and digital front-ends, for injecting currents needed for parameters setting, and also noisy currents (see later).

The motherboard (governed by a microcontroller) directly receives data from the PC via a serial line (also a parallel interface is implemented). Data acquisition is accomplished by means of an auxiliary device (lower right part of the figure), gathering spikes from the neural chip through 16 channels, accumulating them on a FIFO³ (plus some additional processing that we skim over) and sending them to a PC via a serial line upon request. The use of the 'slow' serial line for our purpose was not a serious limitation, in that for interesting collective dynamical regimes of the network the rate of events to be collected is not so high.

To make contact with theoretical predictions it proved useful to be able to inject uncorrelated, *noisy* currents into the neurons in the network: the behavior of an IF neuron driven by a stochastic current has peculiar features, relevant for modeling biologically plausible regimes of neurons' firing. The circuit sketched in Fig. (5) was incorporated in the modules surrounding the pod, and provided a good generator of a Gaussian current. It is based on a classical noise generator exploiting the properties of feedback shift registers [18,19]; the digital waveforms schematically drawn in the figure are usually filtered (e.g. by an RC low-pass circuit) to produce an analog, Gaussian signal, while in this case the IF neuron itself provides the filter.

In the implemented noise generator the length of the shift register is 20 (the period of the pseudo-random sequence is an exponential function of the length of the shift register).

One of the collective properties of IF neurons that we could check with the LANN21 chip has to do with the variability of the emission times of the spikes. Providing the network with steady external currents only, all neurons will fire regularly in the absence of interaction among them; when synaptic couplings are turned on (and dynamically calibrated, as explained before), it turns out that the sole inhomogeneity, associated with the (quenched) random graph of connections among neurons, suffices to endow the firing temporal pattern with a degree of irregularity which is comparable with the one induced by a significant amount of external noise. The latter statement can be made precise if one uses the mean field theory to make predictions about the network's behavior; despite the fact that mean field predictions are expected to provide good approximations for *very* large systems, it turns out that this 21 neurons feedback system was 'large' enough to match well mean field predictions [17]. 'Noise' therefore emerges, as anticipated, as a collective property of the interacting assembly of neurons, and this endogenous irregularity in the neurons' firing can serve as an effective drive for stochastic synaptic dynamics.

³ First In First Out.

Figure 4 about here

Figure 5 about here

Implementing a standard for neuromorphic communication: the AER bus and a PCI interface

Assemblies of IF neurons communicate via instantaneous spikes. Since the early developments of neuromorphic electronics, it seemed attractive to take advantage of this peculiar nature of neuronal communication to draw inspiration for the design of efficient communication infrastructures for neural systems. The AER proposal (for Address Event Representation) [3,20,21] looked particularly promising; the idea was to design a special purpose 'bus', which was actually very simple and effective: each neuron emitting a spike (thus generating an 'event') competes with its fellow neurons in the same chip to access the bus (an arbiter managing the priorities for bus access is a key ingredient of the design), and its address is put on the bus. From this point on, several scenarios can be outlined, and have been the subject of subsequent research: either the 'address event' is broadcast to all the candidate recipients, and a suitable decoding circuitry on each chip takes care of letting the message in if it contains the appropriate target (which implies having the dendritic tree of each neuron resident on the same chip hosting the neuron), or some programmed switchboard, holding information about the pattern of couplings among the neurons on different chips, acts as a gateway to forward the event. Further complications can accommodate transmission delays and additional features. Such an asynchronous bus is a most natural choice when all the messages exchanged are stereotyped, and only the pair sender-receiver matters. For any practical purpose, and to make neuromorphic hardware a really usable platform, a standard communication infrastructure is indispensable, and AER has long been the best candidate. Exchanging experience among groups active in the field, and building multi-modular systems capable of interesting computational abilities entails a down-to-earth and standardized implementation of AER. Our group, in collaboration with the Institute of Neuroinformatics (ETH, Zurich), could recently contribute to take a step in this direction, by designing and implementing a PCI-AER interface, allowing a number of AER-compliant neuromorphic devices to communicate with each other, based on a board coupling in a transparent way this special purpose bus to the standard PCI bus [22].

Figure 6 about here

Fig. (6) illustrates the operation of the AER-based communication. Two neurons in the chip named 'sender' (the roles are interchangeable, of course) emit spikes in rapid succession; they compete for accessing the bus, and the competition is arbitrated by the tree-like circuitry visible in the figure. The winning neuron is allowed access to the bus, while the path taken along the tree-like arbiter leaves a trace in an encoder, which memorizes the address of the emitting neuron. A simple handshake phase provides the go signal for putting the winning event on the bus. A suitable decoding circuitry on the 'receiver' chip directs the event to the appropriate recipients. After completing the transmission of the winning event, the one waiting to be dispatched on the sender chip is managed.

The present version of the PCI-AER board implements the following functions:

- The MAPPER, which essentially implements the programmable connectivity pattern between up to four sender chips, and up to four receiver chips.
- The MONITOR, which allows to tap the transactions on the AER bus, to attach time information to them and to forward the joint information to a PC via PCI.
- The SEQUENCER, that can generate events on the AER bus, emulating a further, virtual neural chip and/or a flux of external spikes to the neural chips.

One appropriate benchmark for this pilot communication system will be provided in the short term by the (already produced) LANN128 chip. The latter is an evolution of the above mentioned LANN21, composed of 128 neurons (88 excitatory neurons connected by 1443 plastic synapses, 40 inhibitory neurons and 1483 fixed synapses) in which some non trivial problems, dealing with the optimal routing strategy (which is an issue in such a crowded and interconnected chip) [23], have been solved. The number of I/O signals forbids any meaningful upgrade of the setup in Fig. (4) which was used for the LANN21, and AER-based communication provides a natural solution. Since the LANN128 design was not originally conceived to incorporate AER (which implies a mixed, analog and digital design), an intermediate solution is being set up, and is sketched in Fig. (7). In particular, the chip named TLANN128 in the scheme implements the AER interface.

Figure 7 about here

Tentative steps towards applications: a real-time tactile transducer of dynamic visual signals

To test the operation of multi-modular systems communicating via the AER bus, we also found convenient to turn to simpler devices than the feedback recurrent neural networks we have considered so far. The idea was to arrange a simple system relying on the PCI-AER board for real time, event driven communication. A small (16x16 pixels) neuromorphic retina (designed and built at the Institute of Neuroinformatics, [24]) gathers visual information, outputting a real time dynamical representation of the visual scene: pixels 'light up' (fast firing) when detecting a high luminosity gradient, but quickly adapt, getting dimmer and dimmer (low firing) if the luminosity pattern does not change (thus acting as a moving edge detector). The PCI-AER board gets spikes from the retina, and manages to forward them to the Braille transducer on the left. The latter is a modified version of a device built for Braille keyboards, and consists essentially of an array of small plastic pins which are driven up and down by piezoelectric crystals reacting to the input currents. The systems is therefore a primitive, but interesting, pilot device providing a tactile dynamic representation of a simple visual scene. Fig. (8) shows a snapshot of the Braille transducer while the retina is looking at a moving circle.

Figure 8 about here

References

1. Mead C. *Analog VLSI and neural systems*. Reading, MA: Addison Wesley 1989.
2. Douglas RJ, Mahowald M, Mead C. Neuromorphic analog VLSI. *Annu. Rev. Neurosci* 1995; 18: 255.
3. Lazzaro J et al. Silicon auditory processors as computer peripherals. *IEEE Trans. Neur. Net.* 1993; 4: 523.
4. Badoni D, Bertazzoni S, Buglioni S, Salina G, Amit DJ, Fusi S. Electronic implementation of a stochastic learning attractor neural network. *Network* 1995; 6: 125.
5. Del Giudice P, Fusi S, Badoni D, Dante V, Amit DJ. Learning attractors in an asynchronous, stochastic electronic neural network. *Network* 1998; 9: 183
6. Amit DJ, Del Giudice P, Fusi S. *Apprendimento dinamico della memoria di lavoro: una realizzazione in elettronica analogica* (in italian) in: *Frontiere della vita*. Rome: Istituto della Enciclopedia Italiana: 1999; 3: 599.
7. Hebb DO. *The organization of behaviour*. Wiley; 1949.
8. Amit DJ. The Hebbian paradigm reintegrated: Local reverberations as internal representations. *Behavioral and Brain Sciences*: 1995; 4: 617.

9. Hopfield JJ. Neural networks and physical systems with emergent selective computational abilities. *Proc. Natl. Acad. Sci. USA*; 198; 79: 2554.
10. Amit DJ. Modeling brain function. *Cambridge University Press*; 1989.
11. Amit DJ, Fusi S. Dynamic learning in neural networks with material synapses. *Neural Computation* 1994; 6: 957.
12. Fusi S, Annunziato M, Badoni D, Salomon A, Amit DJ. Spike-driven synaptic plasticity: Theory, simulation, VLSI implementation. *Neural Computation* 2000; 12: 2221.
13. Battaglia FP, Fusi S. Learning in neural networks with partially structured synaptic transitions. *Network* 1995; 6: 261.
14. Buglioni S. unpublished Thesis, Rome University 1996.
15. Cirenei M. unpublished Thesis, Rome University 1999.
16. Tuckwell HC. *Introduction to theoretical neurobiology*. Cambridge University Press; 1988. Vol. 2.
17. Fusi S, Del Giudice P, Amit DJ. *Neurophysiology of a VLSI spiking neural network: LANN21*. International Joint Conference on Neural Networks Como (Italy) 2000, available on line at <http://neural.iss.infn.it>
18. Horowitz P, Hill W. *The art of electronics*. Cambridge University Press; 1989.
19. Alspector J et al. A VLSI-efficient technique for generating multiple uncorrelated noise sources and its application to stochastic neural networks. *IEEE Trans. Circuits System* 1991; 38: 109.
20. Deiss SR, Douglas RJ, Whatley AM. *A pulse-coded communications infrastructure for neuromorphic systems* in *Pulsed neural networks*, ed. Maas W. and Bishop C.M. The MIT Press, Cambridge, MA 1999.
21. Douglas R, Mahowald M, Whatley A. *Strutture di comunicazione nei sistemi analogici neuromorfi* (in italian) in: *Frontiere della vita*. Rome: Istituto della Enciclopedia Italiana; 1999. 3: 599.
22. Dante V. 2000. Workshop on neuromorphic engineering, Telluride, CO. Documentation available on line at <http://neural.iss.infn.it>
23. Chicca E. unpublished Thesis, Rome University 1999.
24. Kramer J, Indiveri G. *Neuromorphic vision sensors and preprocessors in system applications*. In: *Proceedings of the AFPAEC98 (Advanced Focal Plane Arrays and Electronic Cameras)*. 1998.

Figure captions

Figure 1

Schematic representation of the LANN27 architecture. The rack hosting the board for neurons and synapses is visible at the top, together with the board managing the communication with the PC. The lower part of the figure shows the hierarchical organization of the system.

Figure 2

Block diagram of the acquisition and control system for the output network, based on I²C bus.

Figure 3

Main elements of the output network.

Figure 4

Acquisition and control system for the LANN21 chip. One of the modules for the current injection is visible on the right side of the motherboard, and a picture of the chip below it. The lower right picture refers to the device managing the spike acquisition.

Figure 5

Schematic diagram of the circuit for noise generation.

Figure 6

Sketch of the AER operation. See text for details.

Figure 7

Block diagram of the acquisition and control system for the LANN128 chip. The PCI-AER board is visible on the left.

Figure 8

The "Haptic" setup: The silicon retina is visible on the right; Its spikes are fed into the PCI-AER board, which drives the tactile transducer visible on the left.

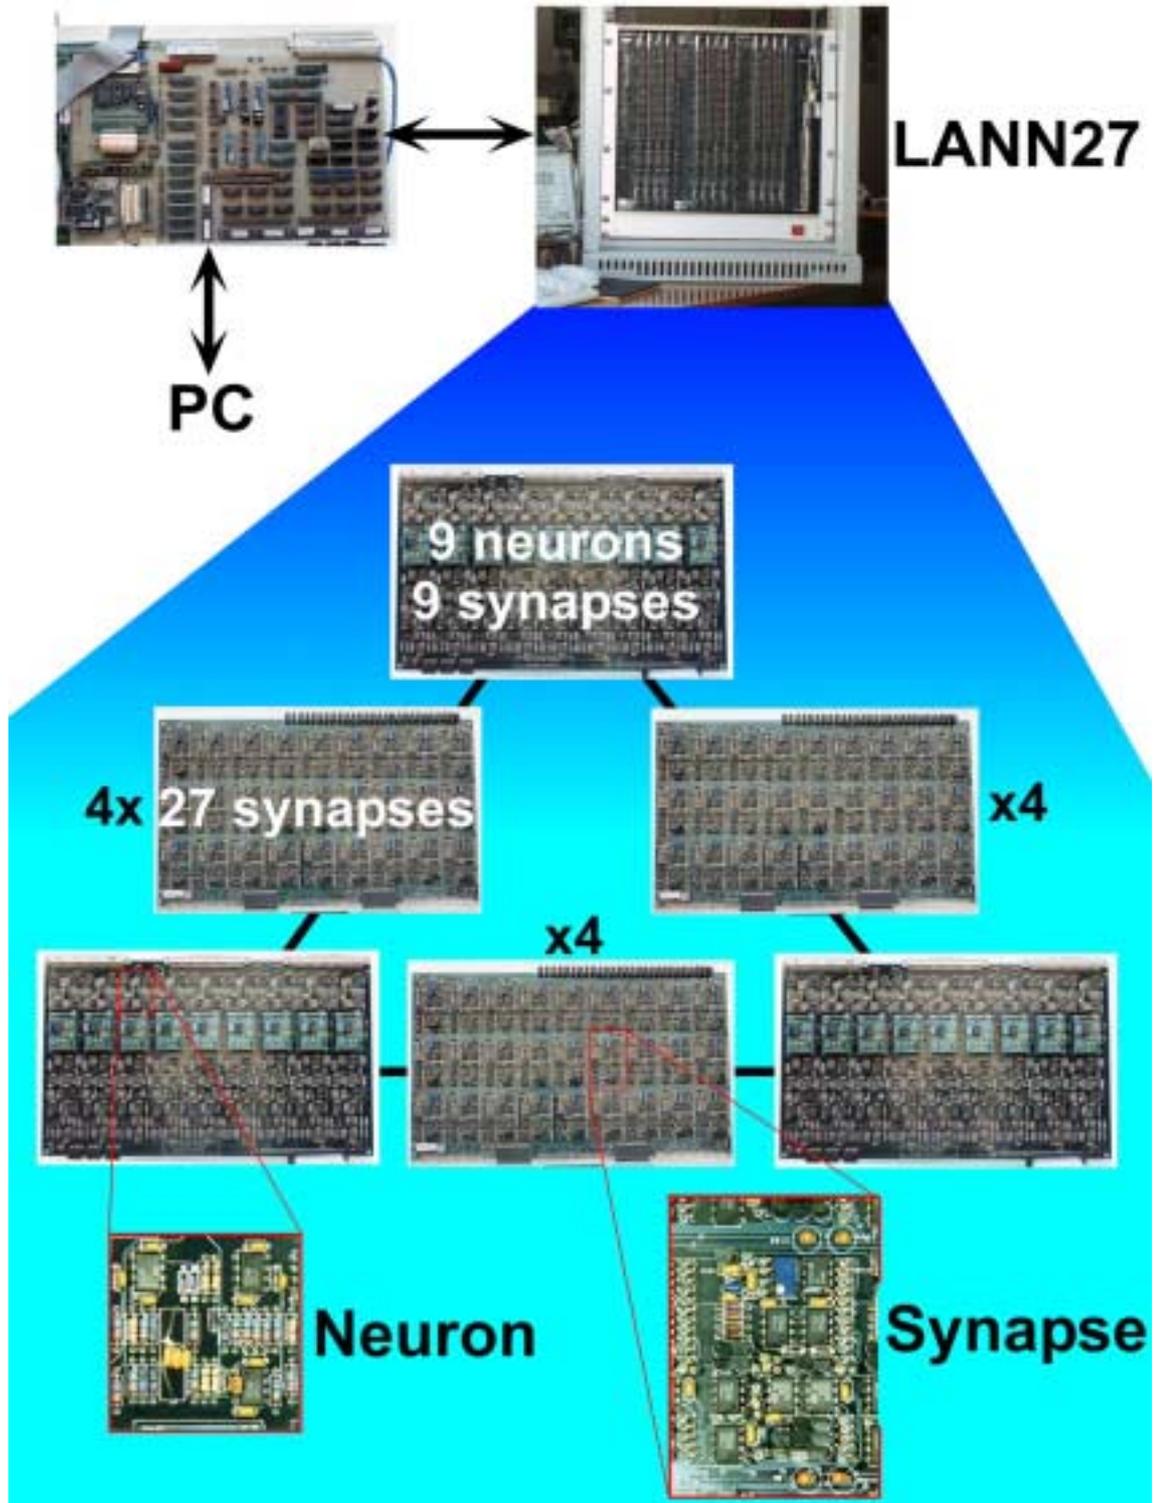


Fig. 1

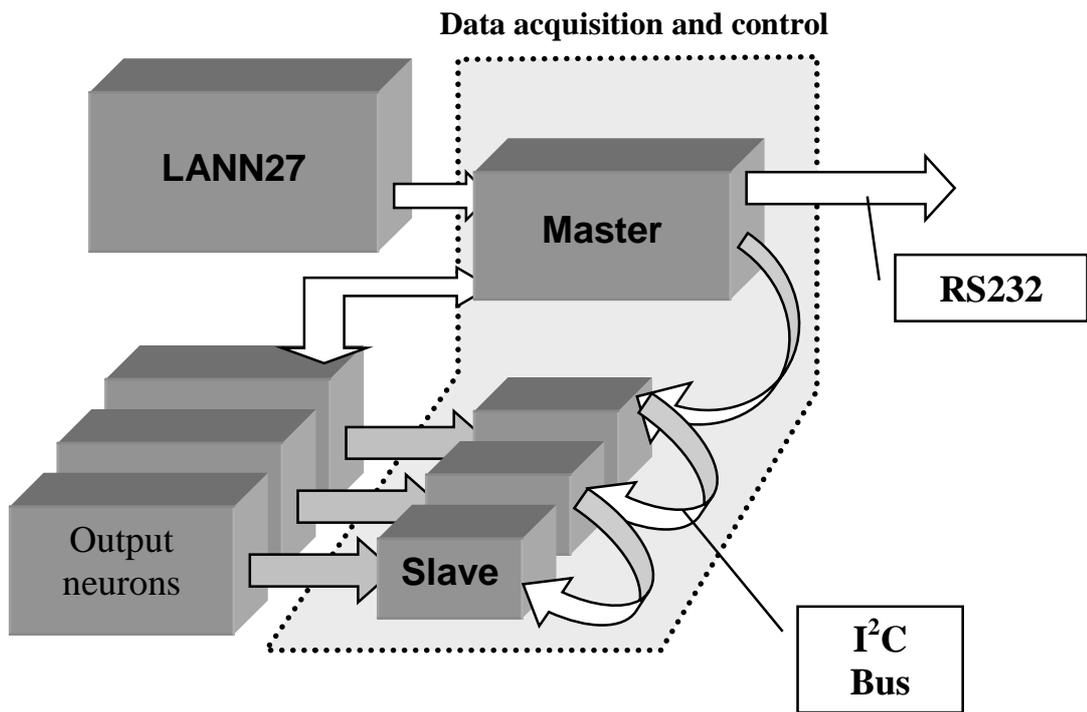


Fig. 2

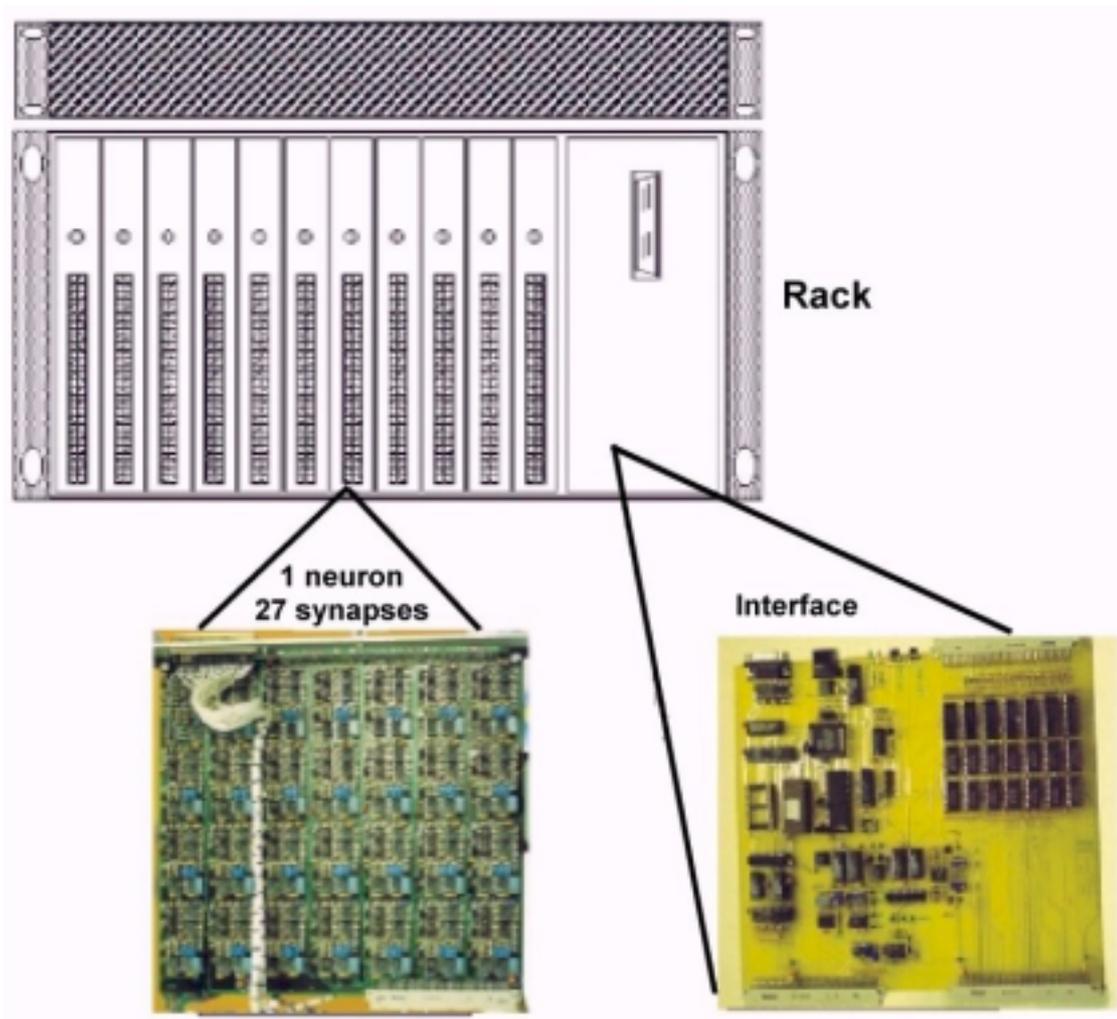


Fig. 3

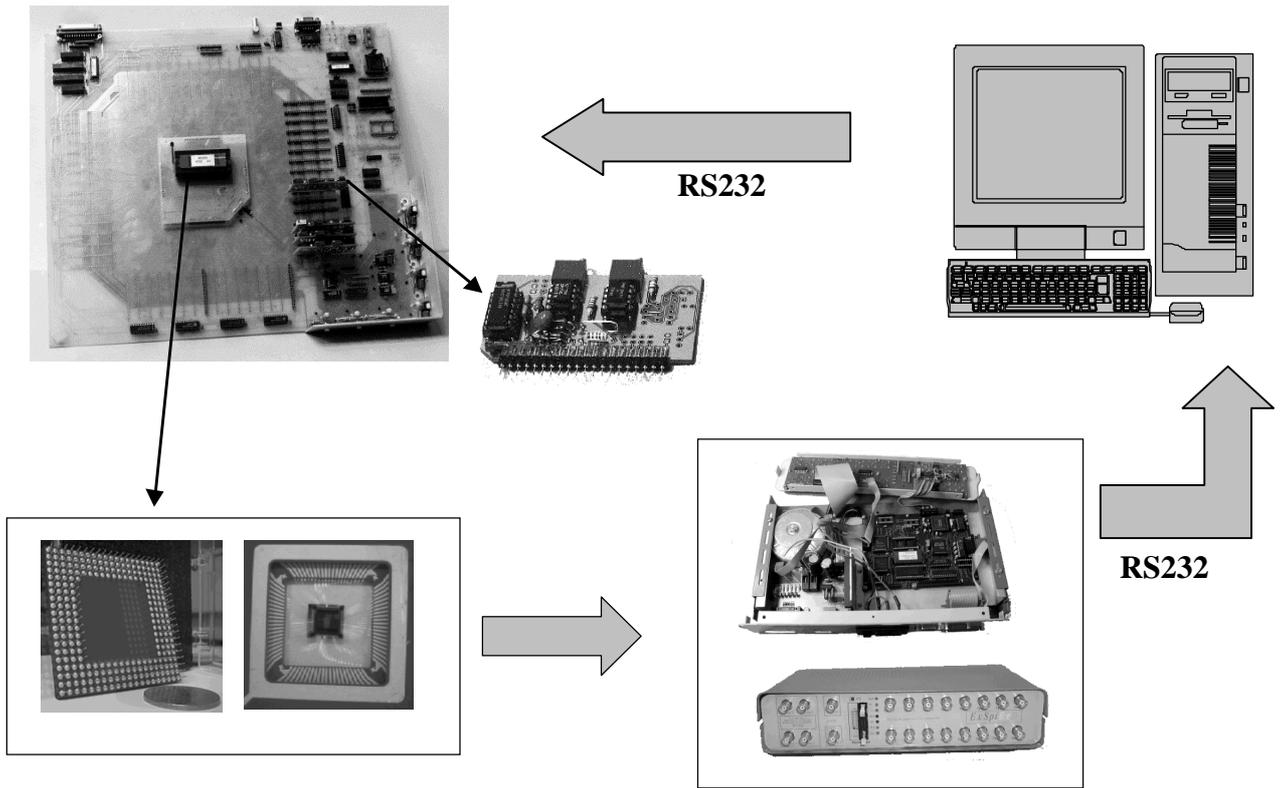


Fig. 4

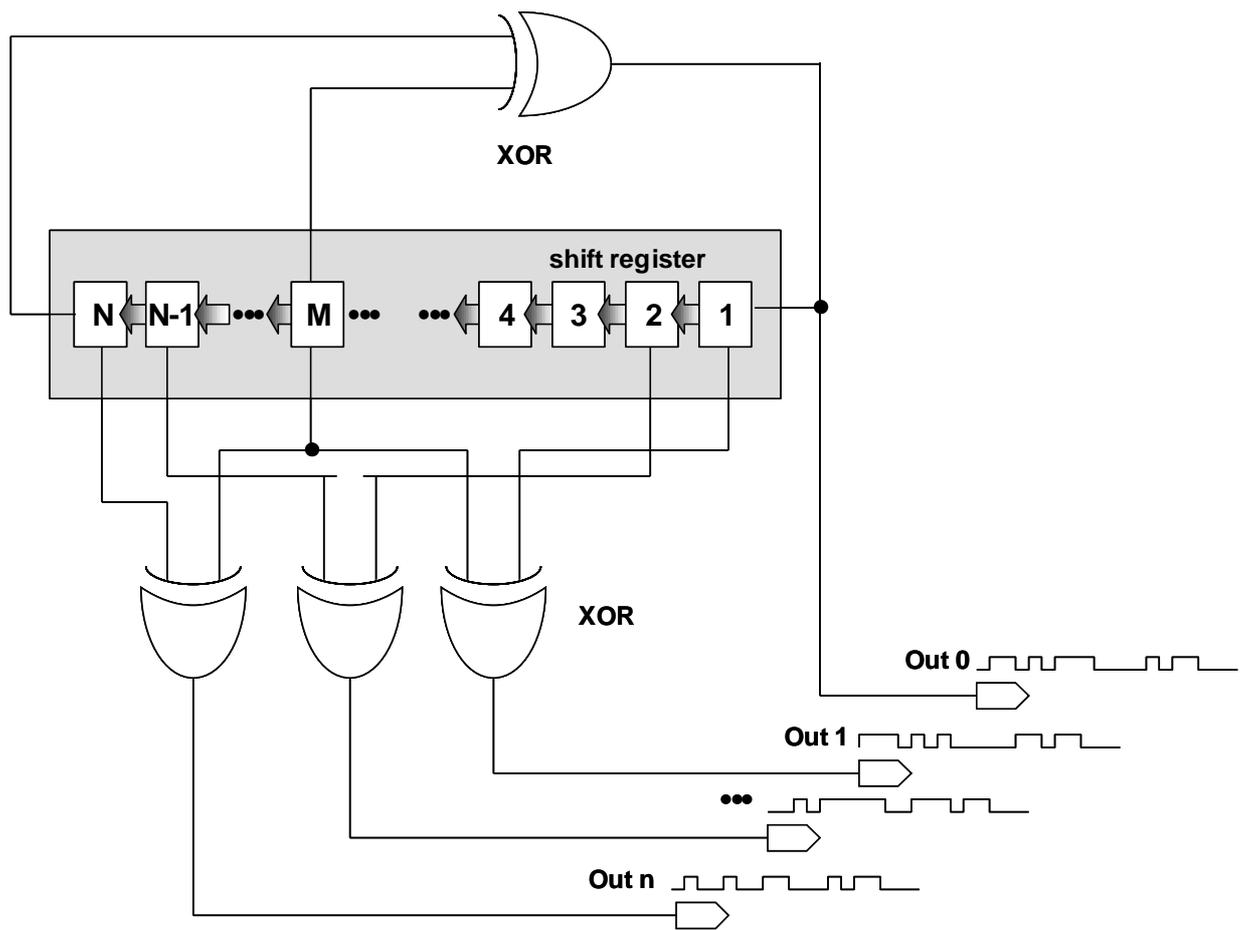


Fig. 5

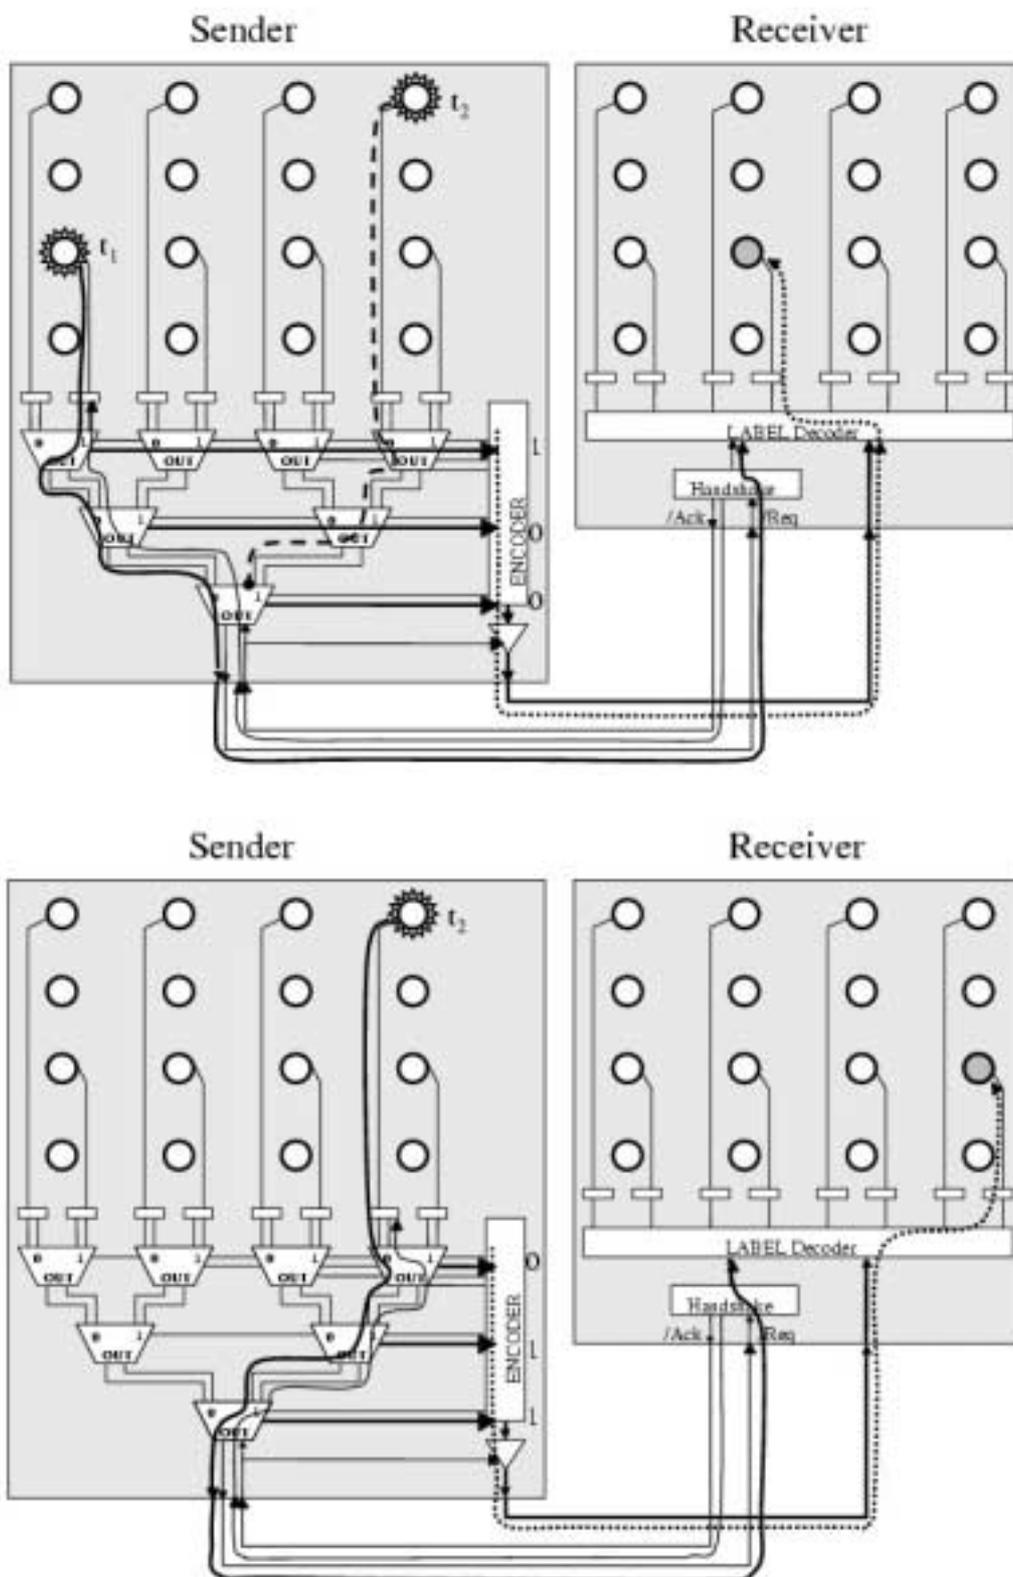


Fig. 6

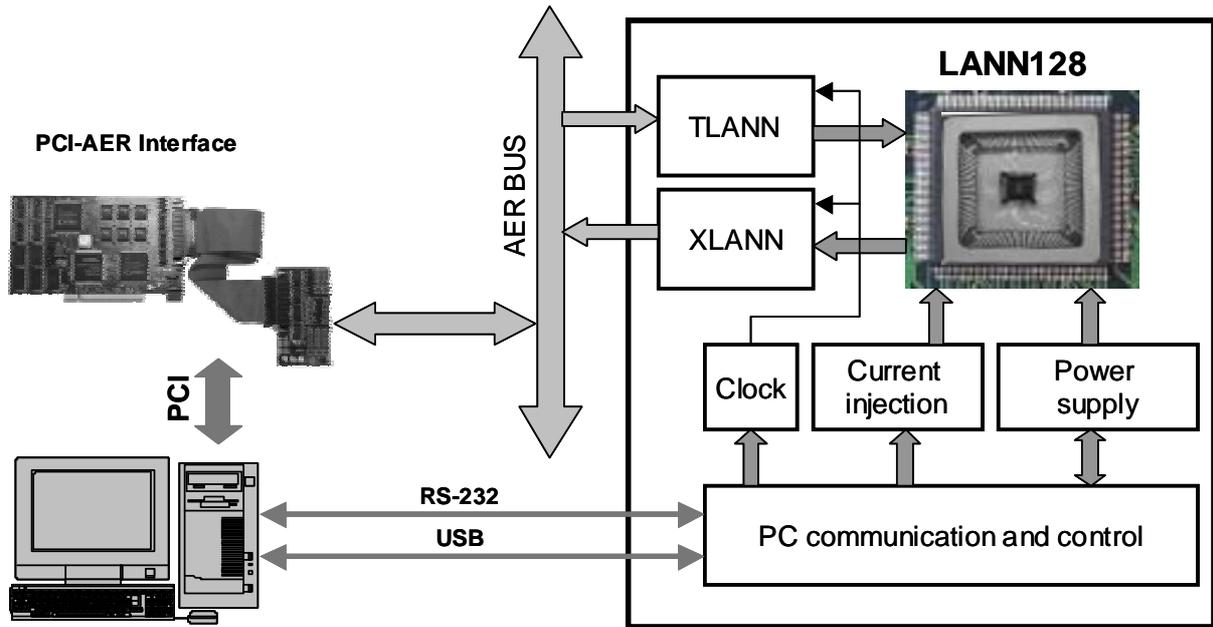


Fig. 7



Fig. 8