DD

TRITA FYS 9012
KTH-Frescati

# REVIEW OF HARDWARE NEURAL NETWORKS: A USER'S PERSPECTIVE *



Clark S. Lindsey and Thomas Lindblad

*Physics Dept. - Frescati, Royal Institute of Technology*
*Frescativägen 24, 104 05 Stockholm, Sweden*
*lindsey@particle.kth.se,lindblad@vana.physto.se*



**Today one can choose from a wide range of neural network hardware. The most important benefit of such hardware is the great increase in speed over conventional sequential processors. The review here surveys a sample of neural network VLSI chips, accelerator boards, and multi-board neurocomputers. We look at the hardware from the potential users viewpoint and discuss some systems developed for high energy physics applications.**

## 1. Introduction

While many neural network applications, such as optical character recognition programs, run well enough on conventional von Neumann processors, some applications, such as in high energy physics, require the speed of hardware implementations[1]. The architectures of neural networks, many simple processors connected together, allow for fast parallel processing. On the other hand, designing hardware with both a large number of processors and high connectivity can be quite difficult. The human brain has $\sim 10^{12}$ neurons with 1000 synapses each, while current VLSI chips typically have 50-100 interconnected neurons. To build larger hardware networks, multi-chip boards and even mult-board systems can provide a few thousand interconnected neurons but require at least some sequential chip and board communication.

We review here a sample of the neural network hardware available commercially or via beta testing or collaboration with the manufacturer. Although the growth has been somewhat slow, there is now a fairly wide range of products available. Such hardware includes digital and analog hardware chips, PC accelerator boards, and multi-board neurocomputers. While the range is wide, it is fairly thin. For a given architecture and technology, e.g. an analog chip with 50-100 neurons running in less than $10\mu s$, there may be only one or two products available. See references 2-6 for further information on hardware neural networks.

## 2. Hardware Specifications

Basic specifications of a neural network include the network architecture (e.g. feedforward multi-layer, radial basis functions (RBF), etc.), number of external inputs/outputs, numbers of neurons and synapses per neuron, number of layers, etc. For a hardware implementation,

*Plenary talk given at the *Third Workshop on Neural Networks: From Biology to High Energy Physics*, Marciana Marina, Isola d'Elba, Italy, Sept. 26-30, 1994.

specifications include the technology used (analog, digital, or hybrid), the precision (in numbers of bits) of the input/outputs, of the weights, and of the accumulators, etc. Various figures of merit indicate the hardware performance. The most common performance rating is the Connection-Per-Sec (CPS), which is defined as the rate of multiplication and accumulate operations during recall processing. The Connection-Update-Per-Second (CUPS) value indicates the rate of weight changes during learning. There are no widely accepted benchmark tests for hardware neural networks, although the learning time for NetTalk is often given[7].

Because of the wide variety of network architectures and hardware implementations, no one or two numbers can give a true picture of the hardware capabilities. Normalizing the CPS value by the number of weights on the chip (CPSPW, or CPS per weight) was suggested as a better way to indicate the processing power of the chip[3]. Also, the Connection-Primitives-Per-Sec value, $CPPS = b_{in} \times b_w \times CPS$, includes the precision in the processing performance[5].

Similar ratings could be defined for CUPS. Normally, CUPS refers to back-propagation learning but the value is often given for other algorithms as well. An algorithm such as Boltzmann learning may only need a few passes through the training set as compared to perhaps 1000's of epochs for back-prop. So a Boltzmann chip may have a lower CUPS value than a back-prop chip yet still accomplish the learning in a shorter time. Further, the CPS and CUPS values do not apply to radial basis function networks where pattern presentation rates are the most relevant performance parameters.

## 3. Neural Networks in VLSI

We divide the neural network VLSI world into three broad categories: digital, analog, and hybrids. Within these categories are included the various network and VLSI architectures, on-chip learning, etc. Table I shows a list of some neural network chips.

### 3.1. *Digital*

The digital neural network category encompasses many sub-categories including slice architectures, SIMD and systolic array devices, and RBF architectures. For the designer, digital technology has the advantages of mature fabrication techniques, weight storage in RAM, and arithmetic operations exact within the number of bits of the operands and accumulators. From the users viewpoint, digital chips are easily embedded into most applications. However, digital operations are usually slower than in analog systems, especially in the *weight* $\times$ *input* multiplication, and analog inputs must first be converted to digital.

#### 3.1.1. *Slice Architectures*

Following the bit slice concept of conventional digital processors, the neural network slice chips provide building blocks to construct networks of arbitrary size and precision. Such chips typcially cost only about \$50/chip, perform at moderate speeds, and are without on-chip learning.

The Micro Devices MD1220 was probably the first commercial neural network chip[8]. Each chip has eight neurons with hard-limit thresholds and eight 16-bit synapses with 1-bit inputs. With bit-serial multipliers in the synapse, the chip provides about 9MCPS. Bigger networks and networks with higher bit inputs can be constructed with multiple chips. A 16-bit accumulator limits the total number of inputs because of overflows.

A similar chip is the Neuralogix NLX-420 *Neural Processor Slice*[9], which has has 16 process-

Table 1. Neural Network Chips. Architectures include feedforward (FdFwd), multlayer (ML), general processor (GP), floating point (FP), radial basis function (RBF), fully connected and recurrent (FCC). Precision refers to the number of bits (b) for the neuron values and weights (the *effective* bits for analog). Here, *na* indicates the information is either not available or not applicable.

| Type | Name | Architecture | Learn | Precision | Neurons | Synapses | Speed |
|---|---|---|---|---|---|---|---|
| Analog | Intel ETANN[27] | FdFwd, ML | no | 6b x 6b | 64 | 10280 | 2GCPS |
| | Synaptics Silicon Retina[25] | Neuromorphic | no | na | 48x48 | Resistive net | na |
| Digital | NeuraLogix NLX-420[9] | FdFwd, ML | no | 1-16b | 16 | Off-chip | 300CPS |
| | HNC 100-NAP[13,14] | GP,SIMD,FP | program | 32b | 100 PE | 512K off-chip | 250MCPS 64MCUPS |
| | Hitachi WSI[24] | Wafer, SIMD | Hopfield | 9b x 8b | 576 | 32k | 138MCPS |
| | Hitachi WSI[6,24] | Wafer, SIMD | BP | 9b x 8b | 144 | na | 300MCUPS |
| | Inova N64000[11,12] | GP,SIMD,Int | program | 1-16b | 64 PE | 128K | 870MCPS 220MCUPS |
| | IBM ZISC036[18] | RBF | ROI | 8b | 36 | 64x64 | 250k pat/s |
| | MCE MT19003[23] | FdFwd, ML | no | 13b | 8 | off-chip | 32MCPS |
| | Micro Devices MD-1220[8] | FdFwd, ML | no | 1b x 16b | 1 PE | 8 | 8.9MCPS |
| | Nestor/Intel NI1000[19] | RBF | RCE,PNN | 5b | 1 PE | 256x1024 | 40k pat/s |
| | Philips Lneuro-1[10] | FdFwd, ML | no | 1-16b | 16 PE | 64 | 26MCPS |
| | Siemens MA-16[15,16] | matrix ops | no | 16b | 16 PE | 16x16 | 400MCPS |
| Hybrid | AT&T ANNA[29] | FdFwd, ML | no | 3b x 6b | 16-256 | 4096 | 2.1GCPS |
| | Bellcore CLNN-32[30] | FCR | Boltzmann | 6b x 5b | 32 | 992 | 100MCPS 100MCUPS |
| | Mesa Reseach Neuroclassifier[32] | FdFwd, ML | no | 6b x 5b | 6 | 426 | 21GCPS |
| | Ricoh RN-200[34] | FdFwd, ML | BP | na | 16 | 256 | 3.0GCPS |

ing elements (PE). A common 16-bit input is multiplied by a weight in each PE in parallel. New weights are read from off-chip. The 16-bit weights and inputs can be user selected as 16 1-bit, 4 4-bit, 2 8-bit or 1 16-bit value(s). The 16 neuron sums are multiplexed through a user-defined piece-wise continuous threshold function to produce a 16-bit output. Internal feedback allows for multi-layer networks. Multiple chips can build large networks .

The Philips Lneuro 1.0 chip[10], which is designed to be easily interfaced to Transputers, also has 16-bit processing in which the neuron values can be interpreted as 8 2-bit, 4 4-bit, etc., sub-values. Unlike the NLX-420, there is a sizable (1kByte) on chip cache to hold weights. The transfer function is done off-chip, which allows for multiple chips to provide synapse-input products to the neurons to build very large networks.

### 3.1.2. *Multi-processor Chips*

A far more elaborate approach is to put many small processors on a chip. Two architectures dominate such designs: single instruction with multiple data (SIMD) and systolic arrays. For SIMD design, each processor executes the same instruction in parallel but on different data. In systolic arrays, a processor does one step of a calculation (always the same step) before passing it's result on to the next processor in a pipelined manner.

SIMD chips include the Inova N64000 and the HNC 100 NAP. The Adaptive Solutions CNAPS systems uses the Inova N64000 to build a SIMD array. The chip contains 64 PE's, with each PE possessing a 9x16 bit integer multiplier, 32-bit accumulator, and 4KBytes of on-chip memory for weight storage[11,12]. All chips execute the same instruction and common control and data buses allow for multiple chips to be combined. The Hecht-Nielson Computers 100 NAP (Neurocomputer Array Processor) contains only 4 PE's but each PE performs true 32-bit floating point arithmetic[13,14]. Weights are stored in off-chip memory and multiple chips can be cascaded.

A systolic array system can be built with the Siemens MA-16[15,16]. The MA-16 provides for fast matrix-matrix operations (mult, sub, or add) of 4x4 matrices with 16-bit elements. The multipler outputs and accumulators have 48-bit precision. Weights are stored off-chip and neuron transfer functions are off-chip via lookup tables. Multiple chips can be cascaded.

### 3.1.3. *Radial Basis Functions*

RBF networks provide fast learning and straight-forward interpretation[17]. The comparison of input vectors to stored training vectors can be done quickly if non-Euclidian distances, such as the Manhatten block norm (sum of element differences), are calculated with no multiplication operations. Two commercial RBF products are now available: the IBM ZISC036 (Zero Instruction Set Computer) chip[18] and the Nestor Ni1000 chip[19]. The ZISC036 contains 36 prototype-vector neurons, where the vectors have 64 8-bit elements, and can be assigned to categories from 1 to 16383. Multiple chips can be easily cascaded to provide additional prototypes. The distance norm is selectable between Manhatten block and the largest element difference. The chip implements a Region of Influence learning algorithm [20] using signum basis functions with radii of 0 to 16383. Recall is according to the ROI identification or via nearest neighbor readout. Recall processing takes $4\mu s$ for a 250k/sec pattern presentation rate. The Nestor Ni1000, developed jointly by Intel and Nestor, contains 1024 prototypes of 256 5-bit elements. The chip has two on-chip learning algorithms, RCE[21] and PNN[22], and other algorithms

can be microcoded. The processing rate is about 40k patterns/sec with a 40MHz clock.

#### 3.1.4. *Other Digital Designs*

Some digital neural network chips don't quite fit into the above three sub-categories. Examples include the Micro Circuit Engineering MT19003 NISP *Neural Instruction Set Processor*[23] and the Hitachi Wafer Scale Integration chips [24]. The NISP is basically a very simple RISC processor with seven instructions, optimized for implementation of multi-layer networks, and loaded with small programs to direct the processing. Feed-forward processing reaches 40MCPS. At the other end of the complexity scale are the Hitachi Wafer Scale Integration chips. Both Hopfield and back-propagation wafers have been built. A neurocomputer with 8 of the back-prop wafers, each with 144 neurons, achieved 2.3GCUPS [6].

### 3.2. *Analog*

Analog hardware networks can exploit physical properties to do network operations and thereby obtain high speed and densities. A common output line, for example, can sum current outputs from synapses to sum the neuron inputs. However, analog design can be very difficult because of the need to compensate for variations in manufacturing, in temperature, etc. Creating an analog synapse involves the complications of analog weight storage and the need for a multiplier linear over a wide range. While many designs use analog techniques to carry out conventional architectures like multi-layer feedforward networks, *neuromorphic* designs, such as the Synaptics Silicon Retina [25], emulate biological functions as closely as possible [28].

The first analog commercial chip was the Intel 80170NW ETANN (Electrically Trainable Analog Neural Network) that contains 64 neurons and 10280 weights[27]. The non-volatile weights are stored as charge on floating gates and a Gilbert multipler provides 4-quadrant multiplication. A flexible design, including internal feedback and division of the weights into two 64x80 banks (including 16 biases), allows for multiple configurations including 3-layers of 64 neurons/layer, and 2-layers with 128 inputs and 64 neurons. No on-chip training was provided so a *chip-in-the-loop* mode with a PC is necessary.

### 3.3. *Hybrid*

Hybrid designs attempt to combine the best of analog and digital techniques. Typically, the external inputs/outputs are digital to facilitate integration into digital systems, while internally some or all of the processing is analog. The AT&T ANNA *Artificial Neural Network ALU* [29], for example, is externally digital but uses capacitor charge, periodically refreshed by DAC's, to store the weights. Similarly, the Bellcore CLNN-32 chip has 5-bit weights loaded digitally but the processing of the network with Boltzmann style annealing is done in analog[30].

The NeuroClassifier from the Mesa Research Institute at the Univ. of Twente has 70 analog inputs, 6 hidden and 1 analog output with 5-bit digital weights[32]. The feed-forward processing rate is an astounding 20ns, representing 20GCPS. The final output is without a squashing function so that multiple chips can be added to increase the number of hidden units.

The use of pulse rates or pulse widths is another method to emulate nets in hardware. The first commercial implementation was the Neural Semiconductor chip set with the SU3232 synapse unit and the NU32 neuron unit[33]. The Ricoh Company has reported a pulse chip with a special back-propagation algorithm implemented on-chip[34]. The RN-100 contained only a

single neuron with 8 inputs and 8 outputs. An array of 12 RN-100's learned to balance a 2-D pendulum in just 30s. A later chip, RN-200, has 16 neurons each with 16 synapses[34].

## 4. PC Accelerator Cards and Neurocomputers

Very large networks (e.g. 1000's of neurons) may only be practical with specialized neural network hardware. While large general purpose parallel machines (e.g. Connection Machine) can certainly provide sufficient performance, cheaper alternatives are available with co-processor, or accelerator, cards for PC. There are also more elaborate *neurocomputers* with multiple boards in separate enclosures, e.g. VME, with extensive software environments. Such neurocomputers may be expensive (e.g. > $20USD) but are still much cheaper than the big parallel mainframes.

Table 2 lists some of the accelerator cards and neurocomputers now available. Several of the cards simply use fast RISC chips (e.g. Intel i860) or DSP's as coprocessors to speed up the network processing. The boards usually come with software that includes several neural network algorithms. Other boards, like the IBM ZISC board, take advantage of neural network chips to optimize network performance. A disadvantage with many such co-processor cards is that they do not allow signals to enter directly to the card but must come over the slow PC bus. This reduces the advantage of using such cards for real-time processing.

Three neurocomputer systems are listed. The Adaptive Solutions CNAPS uses the Inova N64000 chip on on VME boards in a custom cabinet run from a UNIX host[12]. Boards come with 1 to 4 chips and two boards can process the same network to give a total of 512 PE's. The software includes a C-language library, assembler, compiler, and a package of NN algorithms. Similarly, the HNC SNAP Neurocomputer comes with typically include 2 VME boards, each with four NAP 100 chips, providing 32 PE's total[14]. The boards are controlled from a PC by the HNC Balboa accelerator card. The Siemens SYNAPSE-1 uses a systolic array of 8 MA-16 chips in a custom cabinet with a Unix host[16].

Table 2. Neural network accelerator cards and neurocomputers. Here, *na* indicates the information is not available, *prop* indicates a proprietary chip.

| Type | Name | Chip | Performance |
|---|---|---|---|
| PC Accelerators | AND HNet Transputer 1.0 [35] | Transputer T400 | na |
| | BrainMaker Accel. [36] | TI TMS320C25 DSP | 40MCPS,500MFLOPS |
| | Current Tech. MM32k [37] | prop. 2048 PE/chip | 4.9MCPS, 2.5MCUPS |
| | HNC Balbo 860 [14] | Intel i860 | 80MFLOPS |
| | IBM ZISC ISA [38] | IBM ZISC036 | 800k pat/sec |
| | Neural Tech NT6000[39] | TI TMS320C20 DSP | 2MCPS |
| | NeurodynamX XR50[40] | Intel i860 | 45MCPS |
| | Nestor Ni1000 [41] | Nestor Ni1000 | 40k pat/sec |
| | Rapid Imaging 0491E1-ISA[42] | Intel ETANN | 2GCPS |
| | Telebyte 1000 NeuroEng.[43] | prop. | 140MCPS |
| | Vision Harvest NeuroSim.[44] | Intel i860 | 30MCPS,100MFLOPS |
| | Ward Sys. NeuralBoard[45] | 50MHZ RISC | 25MFLOPS |
| Neurocomputers | Adaptive Solutions CNAPS [12] | Inova N64000 | 5.70GCPS, 1.5GCUPS |
| | HNC SNAP [14] | HNC 100 NAP | 500MCPS, 128MCUPS |
| | Siemens SYNAPSE-1 [16] | Siemens MA-16 | 800MCPS |

Table 3. Neural network hardware systems developed for experimental high energy physics.

| Group | Experiment | NNW Purpose | NNW Description |
|---|---|---|---|
| Bologna [49] | CERN fixed target WA-92 expt. | Tag events with 2nd vertex | VME Card, ETANN |
| | | | VME Card, MA-16 |
| CDF/Michigan [47] | Fermilab pp collider, CDF | Tag events with Cal. shower ID | Fastbus card, ETANN |
| Max-Planck München & Dortmund[50] | DESY HERA ep collider, H1 | Level II trigger | 10 CNAPS VME cards, with 1 chip (64 PE) each |
| Royal Inst. of Technology[51,52,53] | prototypes | General purpose | VME card, ETANN, 68070 control |
| | | General purpose | ISA card, 2 ZISC036 |
| | | General purpose | VME card, 4 ZISC036 |
| | | General purpose | VME card, 2 NLX-420 |

## 5. Neural Network Hardware in High Energy Physics

In accelerator experiments the interaction rates may reach up to 40MHZ. In such experiments most of the data consists of background events that must be eliminated by on-line processing. The final rate of events to tape may be only a few tens of Hz. This reduction of $10^6$ is accomplished by a hierarchy of filters or triggers, each doing ever more complex analysis on the events that pass the preceeding trigger. The first level trigger, usually in pipeline style hardware, does simple cuts, such as calorimeter pulse height discrimination, within about $2\mu s$. A second level, also in hardware, runs in about $10-20\mu s$. The third and fourth levels, running in software each have a few hundred msecs.

Hardware neural networks have thus far been used for 2nd level triggers[1,2]. In table 3 we list some systems developed for high energy physics (HEP) experiments using some of the neural network hardware mentioned above.

The CDF experiment at CERN was the first experiment to use hardware neural networks (ETANN) in a HEP experiment[47]. While one ETANN was initially intended to identify b-jets with trained networks[48], this never got beyond the test stage. The other two ETANN boards did electromagnetic shower isolations but were not trained networks. The chips merely executed templates that compared the energies in the seed towers with the surrounding tower energies. The parallel architecture of the ETANN provided fast template processing. However, there are now plans to use these boards to identify tau particles with trained networks.

WA-92 experiment was an ongoing experiment that allowed a temporary neural network trigger to be added as a demonstration of such technology[49]. Only the ETANN board was ready for the 1993 run but the MA-16 board has since been benchtested with similar input data. The actual architecture implemented in the chips was a Fisher discriminate (FD) that required only one layer of processing. A simulated 3-layer network performed slightly better than the FD but the ETANN's limited precision eliminated this advantage. Since the discriminate planes were determined from example data, this is claimed the first true neural network trigger.

The H1 experiment at the HERA ep collider will implement the entire second level trigger with a set of neural network cards[50]. It will use the VME CNAPS boards, each with a single N64000 (64 PE). They found that three layer networks (e.g. 64-64-1) executed within $10\mu s$ with these boards, meeting the trigger time constraints. With the use of off-the-shelf boards, the

neural network hardware development was accomplished quickly and work could focus instead on optimizing the algorithms and on the interface electronics.

Several prototype cards, with ETANN, ZISC036 and NLX-420 chips, have been built by our group at the Royal institute of Technology for possible HEP applications[51,52,53]. The boards have been tested with simulated data for tasks such as secondary vertex finding, Higgs identification, etc.

## 6. Discussion

Hardware neural networks have now passed the novelty stage. From chips to PC coprocessors to full scale neurocomputers, there are number of choices. However, one may find that for a particular application the choices are limited and involve various tradeoffs. The ETANN, for example, offers high speed analog processing but with low precision. The Neuroclassifier offers extremely high speed but also low precision and only a few hidden units. Neurocomputers like the CNAPS can offer high speed, large networks, and multiple net architectures, but are expensive and have elaborate software environments. The radial basis function networks can offer high performance, especially for learning, but may not have sufficient generalization powers for some applications. Accelerator boards can suffice for running very large networks in a *reasonable* time, but probably not in real-time, where the VME cards are most viable.

For high energy physics applications, there are now a number of options for the second level trigger, especially when the network sizes involve around 60 inputs/neurons per layer. Both digital and analog systems are available. So far the work has only involved feedforward layer networks but it would be interesting to investigate applications of the radial basis function networks. For the first level, the only real option is the Neuroclassifier, although several chips must be combined for networks with more than 6 hidden units.

## Acknowledgments

We would like to thank the help of B. Denby, Å. Eide, G. Sekhniaidze and other co-workers. This work was supported in part by grants from the Swedish Engineering Sciences Research Council (TFR) under contract Doss 240, Dnr: 93-400, from the Natural Sciences Research Council (NFR) and from the Carl Trygger Foundation. We also acknowledge the support of M. Holler and F. Martin of Intel, and J.-P. LeBouquin and M. Grandguillot of IBM-Essonnes.

## References

1. B. Denby, "The Use of Neural Networks in High Energy Physics", *Neural Computation*,5, 505-549, July 1993.
2. World Wide Web site: http://www1.cern.ch/NeuralNets/nnwInHep.html, http:/www1.cern.ch/NeuralNets/nnwInHepHard.html.
3. M. Holler, "VLSI Implementation of Learning and Memory Systems: A Review", *Advances in Neural Information Processing Systems 3*, Morgan Kaufmann, San Mateo, Ca. 1991.
4. H.P. Graf, E. Sackinger, L.D. Jackel, "Recent Developments of Electronic Neural Nets in North America", *Journal of VLSI Signal Processing*,5,19-31, 1993.
5. E. van Keulan, S. Colak, H. Withagen, H. Hegt, "Neural Network Hardware Performance Criteria", *Proc. of the IEEE Conference on Neural Networks*, **III**, 1885-1888, 1994.

6. M. Glesner and W. Pöchmüller, *Neurocomputers: An Overview of Neural Networks in VLSI*, Chapman & Hall, London, 1994.
7. T. Sejnowski and C. Rosenberg, "NetTalk: A Parallel Network that Reads Aloud", Tech. Rep. JHU/EECS-86/01, The Johns Hopkins Univ. Elect. Eng. and Comp. Sci. Dept. 1986.
8. MD1220 Data Sheet, March 1990, Micro Devices, 30 Skyline Dr., Lake Mary, Fl. 32746.
9. NLX420 Data Sheet, June 1992, Neurologix, Inc., 800 Charcot Av., Suite 112, San Jose, Ca.
10. N. Mauduit, M. Duranton, J. Gobert, "Lneuro 1.0: A Piece of Hardware LEGO for Building Neural Network Systems", *IEEE Trans. on Neural Networks*,**3**,414-422, May 1992.
11. D. Hammerstrom, "A VLSI Architecture for High-Performance, Low-Cost, On-chip Learning", *Proc. Int. Joint Conf. on Neural Networks*, June 1990,**II**, 537-544, San Diego, Ca.
12. Adaptive Solutions, Inc. 1400 N.W. Compton Dr., Suite 340, Beaverton, Or. 97006.
13. R.W. Means, L. Lisenbee, "Extensible Linear Floating Point SIMD Neurocomputer Array Processor", *Proc. Int. Joint Conf. on Neural Networks*, July 1991,**I**, 587-592, Seattle, Wash.
14. Hecht-Nielson Computers, Inc., 5501 Oberlin Dr., San Diego, Ca. 92121-1718.
15. J. Beichter, N. Bruels, E. Meister, U. Ramacher, H. Klar, "Design of General-purpose Neural Signal Processor", *2nd Int. Conf. on Microelectronics for Neural Networks*, 1991, 311-315, Münich, Ge.
16. U. Ramacher, "SYNAPSE - A Neurocomputer that Synthesizes Neural Algorithms on a Parallel Systolic Engine", *Journal of Parallel and Distributed Computing*,**14**, 306-318, 1992. Also, Siemens Nixdorf, Informationssysteme AG, D55 Sci. Computing, D-81730 Münich, Ge.
17. P. Wasserman, *Advanced Methods in Neural Computing*, Van Nostrand Reinhold, 1993.
18. J-P. LeBouquin, "IBM Microelectronics ZISC, Zero Instruction Set Computer", *Proc. of the World Congress on Neural Networks*, Supplement, San Diego, 1994.
19. M. Holler, C. Park, J. Diamond, U. Santoni, S.C. The, M. Glier, C.L. Scofield, L. Nunez, "A High Performance Adaptive Classifier Using Radial Basis Funcitons", *Proc. of Gov. Microcircuit Applications Conf.*, Las Vegas, Nevada, Nov. 1992.
20. M. Verleysen,J. Cabestany, Project ESPRIT ELENA, Ralisation VLSI de reseaux de neurones. VLAGO, ISSN 1243-4835 No. 94-1: les processeurs neuronaux,juin 1994
21. D.L. Reilly, L.N. Cooper, and C. Elbaum,"A neural model for category learning", *Biological Cybernetics*, **45**, pp. 35-41 (1982).
22. D.F. Specht, "Probablistic Neural Networks for Classification, Mapping, or Associative Memory", *Proc. of the IEEE Int. Conf. on Neural Networks*,**I**, 525-532, San Diego, Ca. 1988.
23. MT19003 Data Sheet, 1994, MCE, Alexander Way, Tewkesbury, Gloucestershire GL20 GTB.
24. M. Yasunga, N. Msuda, M. Yagyu, M. Asai, M. Yamada, A. Masaki, "Design, Fabrication and Evaluation of a 5-Inch Wafer Scale Neural Network LSI Composed of 576 Digital Neurons", *Proc. Int. Joint Conf. on Neural Networks IJCNN'90*, June 1990,**II**, 527-535, San Diego, Ca.
25. Synaptics, 2860 Zanker Rd., Suite 206, San Jose, Ca. 95134, USA.
26. C. Mead, Analog VLSI and Neural Systems, Addison Wesley, Menlo Park, Ca. 1989.
27. 80170NX ETANN, Data Sheet, Intel Corp., Santa Clara, CA. 1991.
28. C. Mead, Analog VLSI and Neural Systems, Addison Wesley, Menlo Park, Ca. 1989.
29. B. E. Boser, E. Sackinger, J. Bromley, Y. LeCun, L.D. Jackel, "Hardware Requirements for Neural Network Pattern Classifiers", *IEEE Micro*, pp. 32-40, Feb. 1992.
30. J. Alspector, T. Jayakumar, and S. Luna, "Experimental Evaluation of Learning in a Neural Microsystem",*Advances in Neural Information Processing Systems - 4*, ed. J. Moody et al., 871-878, Morgan-Kaufmann Pub., San Mateo, Ca. 1992.

31. Th. Lindblad, C. S. Lindsey, F., Block, T. Jayakumar, "Using Software and Hardware Neural Networks in a Higgs Search", *Nucl. Instr. and Meth.* in print.
32. P. Masa, K. Hoen, H. Wallinga, "A 20 Input, 20 Nanosecond Pattern Classifier",*Proc. Int. Joint Conf. on Neural Networks IJCNN'91*, July 1994, Orlando, Fla., USA.
33. M.S. Tomlinson, D.J. Walker, M.A. Sivilotti, "A Digital Neural Network Architecture for VLSI", *Proc. Int. Joint Conf. on Neural Networks*,II, 545-550, June 1990, San Diego, Ca.
34. H. Eguci, T. Furata, H. Horiguchi, S. Oteki, T. Kitaguchi, "Neural Network LSI Chip with On-chip Learning", *Proc. Int. Joint Conf. on Neural Networks IJCNN'90*, June 1990, San Diego, Ca. Also, priv. comm. from H. Eguci.
35. AND America Ltd., 2140 Winston Park Dr., Suite 202, Oakville, Ont. L6H5V5, Canada.
36. California Scientific Software, 10024 Newtown Rd., Nevada City, Ca. 95959-9794.
37. M. Glover, W.T. Miller, "A Massively-parallel SIMD Processor for Neural Network and Machine Vision Applications",*Adv. in Neural Info. Processing Systems - 6*, ed. J.D. Cowan et al.,Morgan-Kaufmann,1994. Also, Current Tech., 97 Madbury Rd., Durham, NH. 03824.
38. ZISC Neural Network Series, ZISC/ISA Accel. for PC (Users Manual), Prel., ver. 0 (11/1994), IBM-Essonnes, Comp. Dev. Lab., 224 Bd. John Kennedy, F-91105 Corbeil-Essonnes, France.
39. Neural Technologies, Ltd., Petersfield, UK, 730-260256.
40. NeuroDynamX, P.O. 323, Boulder, Co. 80306.
41. Nestor Inc., One Richmond Square, Providence, RI 02906.
42. Rapid Imaging Inc., 5955 T.G. Lee Blvd., Suite 150, Orlando, Fla. 32822.
43. Telebyte Technology, Inc., 270 Pulaski Rd., Greenlawn, NY 11740.
44. Vision Harvest, HCR Box 36, Hatch N.M. 87937.
45. Ward Systems Group, Executive Park West, 5 Hillrest Dr., Frederick, Md. 21702.
46. Intel Corp.,"Neural Network Solutions", Order Num:296961-002, Intel Corp., 2200 Mission College Blvd., Mail Stop RN3-17, Santa Clara, Ca. 95052, 1992.
47. W. Badgett, K. Burkett, M. Campbell, D.Y. Wu, B. Denby, C.S. Lindsey, R. Blair, S. Kuhlmann, J. Romano., "A Neural Network Calorimeter Trigger Used in CDF", *1992 Nucl. Sci. Symposium*, Orlando, Fla. 1992, Conf. Record p.825.
48. B. Denby, M. Campbell, F. Bedesch, N. Chriss, C. Bowers, and F. Nesti, "Neural Networks for Triggering", *IEEE Trans. Nucl. Sci.* **37**(2), 248-254.
49. R. Odorico, *IEEE Trans. on Nucl. Sci.*, **40**, (1993) 705.
50. C. Kiesling, J. Fent, W. Fröchtenicht, D. Glodner, T. Kobler, H. Kolanski, T. Krämerkämper,J. Möck, P. Ribarics, "Concept for a Fast Second Level Trigger Using a Neural Network Architecture for the H1-Experiment at HERA", 411-421, *New Computing Techniques in Physics Research III*, ed. K-H Becks & F. Perret-Gallix, World Sci., 1994.
51. J. Molnar, G. Székely, Th. Lindblad, C.S. Lindsey, B. Denby, S.R. Amendolia, Å. Eide, "A Flexible VME-module with a 68070 Computer for Embedded Applications of the 80170NX Analog Neural Network Chip", *ICFA Intstr. Bulletin*, No. 10, Dec. 1993.
52. Å. Eide, Th. Lindblad, C.S. Lindsey, M. Minerskjöld, G. Sekhniaidze, and G. Székely, "An Implementation of the Zero Instruction Set Computer (ZISC036) on a PC/ISA-bus Card", Proc. of the 1994 WNN/FNN, Washington DC, Dec. 1994, *SPIE* in print.
53. Th. Lindblad, C.S. Lindsey, M. Minerskjöld, G. Sekhniaidze, G. Székely, Å. Eide, "Implementing the New Zero Instruction Set Computer (ZISC036) from IBM", *NIM-A*, in print.