

VLSI DESIGN OF AN ASSOCIATIVE MEMORY BASED ON DISTRIBUTED STORAGE OF INFORMATION

U. RÜCKERT

INTRODUCTION

In the last decade there has been an increasing interest in the use of artificial neural networks in various applications. One of these application areas, where the analysis of the performance of a neural network approach is comparatively advanced, is associative memory. Artificial neural networks (ANNs) are well suited for the implementation of associative memories at least because the processing elements (artificial neurons) in an ANN operate in a highly parallel way and thus a considerable gain in speed is to be expected. This parallelism is one of the major reasons for investigating new computational models inspired by neurophysiological processing principles. The idea is that information is stored in terms of synaptic connectivities between artificial neurons while the activities of the neurons represent the stored patterns. Thus the information is stored distributed over many units and not anywhere in particular. Each unit participates in the encoding of several informations (patterns). ANNs combine both storing and processing information whereas conventional computers are based on sequential processors operating on the contents of a passive memory in which information is accessed by finding the right place in memory.

Many different models have been discussed in literature under such names as "Lernmatrix", "Correlation Matrix", "Associative Memory" etc. [1,2]. A bottleneck in both theory and applications of ANNs seems to be the current lack of suitable parallel hardware exploiting the system inherent parallelism. Thus, development of parallel computing architectures is a fundamental requirement in making some of these models a realistic alternative to more conventional forms of information processing. On the other hand, there are purely technological reasons for studying parallel systems, because they may be the best way to increase the speed and the power of computation in the future. In this respect, ANNs are well adapted for VLSI system design, especially those models using simple processing units and regular interconnection schemes. The highly regular and modular architecture of certain ANN models is an attractive property that circuit designers want to transfer to parallel VLSI hardware. In system design this aspect tends to play a more important part, as Computer Aided Manufacturing enhances VLSI process to higher integration levels and Computer Aided Design leads to easy-to-handle design of smart chips.

There are two different approaches for supporting these models on parallel VLSI hardware [3]: *General-Purpose Neurocomputers* for emulating a wide range of neural network models and *Special-Purpose VLSI-Systems* which are dedicated to a specific neural network model. This paper is devoted to a special-purpose hardware implementation of a very simple associative memory based on neural networks. The memory has a simple matrix structure with binary elements (connections, synapses) and performs a pattern mapping or completion of binary input/output vectors. To the authors knowledge, this comparatively simple model of a distributed associative memory was first discussed by Willshaw et. al. in 1969 [4]. However, similar structures have been more generally discussed, e.g. by Kohonen and Palm [1,2]. The characteristics of the implemented model are described in the 2nd section.

The important aspect for VLSI implementation of this simple memory model is the close relationship to conventional memory structures. Hence, it can be densely integrated and large scale memories with several thousands of columns (model neurons) can be realized with current technologies already. Furthermore, the regular topology results in a rigorous modularization of the system indispensable for a successful management of the design and test complexity of VLSI systems. In this respect a pure digital and a hybrid analog/digital VLSI architecture will be described in the 3rd section and discussed in the 4th section.

THE ASSOCIATIVE MEMORY CONCEPT

In general the basic operation of an associative memory is a certain mapping between two finite sets X and Y . In a more abstract sense these two sets may be regarded as questions and answers or stimuli and responses, both coded as vectors of numbers (Figure 1).




Figure 1 Basic operations of an associative memory: heteroassociation (a) and autoassociation (b).

The associative memory should respond with y^h to the input x^h for every pair (x^h, y^h) stored in the associative memory. The paired associates can be selected freely, independently of each other. This operation is often called *pattern mapping* or *heteroassociative recall* [1,2]. Further it would

be convenient if the associative memory responds with y^h not only to the complete input x^h but also to sufficiently large parts of it. In other words the mapping should be fault-tolerant to incomplete or noisy versions of the input pattern. A special case of this functionality is the *autoassociative memory* where the stored pairs look like (x^h, x^h) . Given a sufficiently large part of x^h the memory responds with the whole pattern (*pattern completion*). Besides the discussion of the fuzzy term "sufficiently large" the input can be any part of the stored pattern and it even can be a noisy version of x^h . This operation is called the best match search in terms of pattern recognition.

Among the many different implementations of an associative memory in the field of neural networks the following simplest type is very attractive as well as effective in regard to VLSI implementation. The *Associative Matrix (AM)* is a $n \times m$ matrix of binary storage elements w_{ij} , the connection weights (Figure 2). The input vectors x^h as well as the output vectors y^h take a binary form.

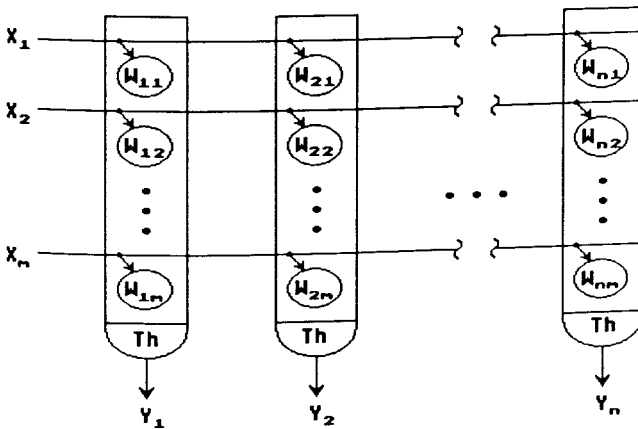


Figure 2 Structure of an Associative Matrix (AM)

The mapping is build up in the following way: The input vector x^h as well as the output vector y^h of every pair which should be stored in the AM ($h = 1, \dots, z$) are applied to the matrix simultaneously. At the beginning all storage elements in the matrix are zero. Each storage element at the crosspoint of an activated row and column ($x^h = y^h = 1$) will be switched on, whereas all the other storage elements remain unchanged. This clipped Hebb-like rule [2] programs the connection matrix and the information is stored in a distributed way (Figure 3 a,b):

$$w_{ij}^h = w_{ij}^{h-1} \vee (x_i^h \wedge y_j^h), w_{ij}^0 = 0, h = 1..z \quad (1)$$




Figure 3: Programming two pattern pairs into an Associative Matrix (a,b) and pattern mapping of an incomplete input pattern (c).

The recall of the constructed mapping is done by applying an input vector to the rows of the matrix. For each column i we sum up the products of the input components x_j^h and the corresponding connection weights w_{ij} :

$$S_i = \sum_{j=1}^m x_j^h \cdot w_{ij} \tag{2}$$

The associated binary output vector is obtained by the following threshold operation (Figure 3 c):

$$y_i^h = \begin{cases} 1, & \text{if } S_i \geq Th \\ 0, & \text{otherwise} \end{cases}, \quad Th \in \mathbb{N} \tag{3}$$

Obviously, because of the above mentioned programming rule the memory matrix gets more and more filled (the connections will never be switched off). Consequently, the output might contain more '1's than the desired output pattern. The chance that this kind of error will occur increases with the number z of stored pairs. This fact causes the following quantitative questions:

- i) How many patterns can be stored in an AM?
- ii) How many bits of information can be stored in an AM?

Both questions were answered, e.g. by Palm 1980 [5]. Summarizing his results, an AM has its optimal storage capacity C for sparsely coded input/output patterns.

Therefore, the number $l(k)$ of active components ('1') in the input (output) patterns should be logarithmically related to the pattern length (n,m) . Asymptotically, the optimal storage capacity for hetero-association is given in [5]:

$$C(m, n, k, l, z) \rightarrow \ln 2 \cdot m \cdot n$$

$$\text{for } n, m \rightarrow \infty \text{ and parameters: } k = \log n, l = \log m, z < \ln 2 \frac{m \cdot n}{k \cdot l} \quad (4)$$

Hence, the storage capacity C is proportional to the number of storage elements $n \cdot m$ and the number of patterns z that can be stored is much larger than the number of columns (artificial neurons). For example, an optimum of $C = 593,000$ bits for $m, n = 1000, z = 34,780$ ($k = 2, l = 9$) can be stored in the AM under the constraint that on the average 90% of the information of the output vector of each pair is stored [5]. Figure 4 shows the storage capacity of an AM as a function of the number of stored patterns (z) and Table 1 as a function of the number of activated components in the input (l) and output (k) patterns, respectively.




Figure 4 Storage capacity C in bits and the expected number E of additional '1's in the output pattern as a function of the number of stored patterns z ($m, n = 1000$): a) $k = l = 3$; b) $k = l = 4$.

Furthermore, it turns out that the AM works for pattern mapping applications in a more economic way compared to conventional methods (e.g. hashing) and other neural network models, if the number of patterns is large and their individual information content small [6].

These results encourage a hardware implementation in VLSI of this simple associative memory model in situations where such a mapping is a more natural way of storing information than a listing. Especially, because the AM works the more effective the larger the matrix is.

$m = n$	l	z	C/mn
1024	10	23313	0.565
2048	11	96012	0.591
4096	12	318353	0.604
8192	13	1182378	0.613
16384	14	4407707	0.621

Table 1: Number of patterns z that can be stored with low error probability and the corresponding storage capacity C as a function of parameters n, m, l ($k = 3$)

VLSI IMPLEMENTATION

The Associative Matrix can be handled most flexible as a simulation program on a conventional computer (workstation), of course. It could be shown that even serial simulations of the AM would have to perform less operations than a conventional implementation in terms of bitwise mask operations. For the special case $m \ll n$ and a sparse matrix the serial implementation is even for a large number of patterns fast enough for certain applications [7]. But in general, a serial implementation has a poor performance, especially if applied to very large matrices.

It is quite obvious that operation can be speeded up considerably by parallel processing. The implementation by means of multiprocessor architectures (SIMD machines) is a promising compromise between flexible modelling - the system is still program controlled - and a complete parallel processing of large matrices. In fact, at least two research groups have already designed a parallel associative computer (SIMD) based on a set of conventional microprocessors communicating via a common bus [8,9].

Consequently, the highest degree of parallelism is achieved by task-dedicated VLSI systems. It is well in the range of current technologies to implement an AM effectly on VLSI chips. So far two different special-purpose VLSI architectures have been designed for an AM at the University of Dortmund: a digital and a hybrid analog/digital implementation. Both of them will be discussed in this paper.

The system architecture in both cases is split up vertically into "slices"; each slice manages an equal number of columns. The slices are controlled by a conventional microprocessor (system control, Figure 5), distributing input data in an appropriate way to the slices and collecting output data

from the slices. In consequence of the sparsely coded input/output patterns the microprocessor transfers and collects the patterns optimally by means of the addresses of the activated components. Hence, a transfer operation of a m -bit pattern takes only $\log(m)$ cycles and address lines in the serial case.




Figure 5 Partition of an $m \times n$ Associative Matrix into slices.

Digital Implementation

In the case of digital implementation the columns of an AM are controlled by a special slice chip comprising several very simple processing units (PUs). Each PU controls one column of the matrix and computes serially the weighted sum (2) of the input pattern and the respective column. Because the input/output signals as well as the connection elements are binary, the basic building blocks of a PU are a counter and a comparator (Figure 6a).

The programming algorithm (1) for the connection matrix is realized by a simple OR-logic-Block and is incorporated on the chip, too. The connection matrix can be build up by conventional RAMs (Figure 6b).

In order to transfer the output pattern to the system control the addresses of the '1's in the pattern are generated locally in the slice chips by an additional priority-encoder. All slice chips are connected to a common bus and the access to the bus can be controlled by daisy-chaining or by an additional priority- encoder-logic (Figure 7). In case of the daisy-chaining method the time for transferring the output pattern is proportional to the number of slice chips. In the other case the time is proportional to the number of '1's in the output pattern. In both cases the transfer of the associated output pattern and the calculation of the weighted sum of inputs can be pipelined.




Figure 6 Basic building blocks of a digital implementation of an Associative Matrix column (a) and slice (b).




Figure 7 Transfer scheme for the generated addresses of '1's in the output pattern.

Up to now several standard-cell-designs comprising 32 PUs (e.g. $2\mu\text{m}$ CMOS, 38mm^2 , $\approx 20,000$ transistors, 53 pads, 10MHz) and a full-custom-design comprising 128 PUs ($2\mu\text{m}$ CMOS, 1cm^2 , 48,000 transistors, Figure 8) of a slice chip have been finished. Instead of realizing a whole chip in silicon we have first fabricated and tested successfully a single PU of the

full-custom-design at the University of Dortmund (Figure 9). The tested PU is able to perform the calculations for the equations (1) - (3) at least at a clock rate of 12 MHz. Optimization in respect to speed hasn't been done yet. Instead, we have concentrated on a modular and testable design. As





Figure 8 Layout and floorplan of a slice chip comprising 128 processing units ($2\mu\text{m}$ CMOS, one metal and one poly layer).

can be seen from the microphotograph in Figure 9, the PU is built up by six bit-slices. The bit-slices can be configured to a scan-path for test reasons and an extension to a 8 or 16 bit PU is easily achievable by duplicating the bit-slices.

Based on this facts, a 8192x8192-AM can be build up by a 64 MBit-RAM and 64 slice chips, each comprising 128 PUs, for example. This AM stores more than one million sparsely coded patterns with low error probability (Table 1), which corresponds to a storage capacity of 40 Mbits. Such an implementation performs a pattern mapping within $10\mu\text{s}$. The association time is proportional to the number of '1's in the input/output patterns ($\log(m) + \log(n)$), hence independent of the number z of stored pairs.




Figure 9 Microphotograph of the test circuit for a single 6-bit processing unit ($2\mu\text{m}$ CMOS, 350 transistors, $1.1\text{mm} \times 0.4\text{mm}$).

With current submicron technologies it is possible to integrate 512 and more PUs on a single slice chip. Hence, the limiting factor for the number of PUs per chip is the pin count. Nevertheless, it is possible to have 256 PUs on a slice chip at least. The important disadvantage of this digital approach up to now is that most RAM chips have a 1, 4 or 8 bit organization whereas a longer word length (>16) is more appropriate for this approach. In order to get the highest degree of parallelism the optimal memory organization is $m \times u$ ($m = \text{number of matrix rows}$, $u = \text{number of PUs per slice chip}$). For the above mentioned $8\text{k} \times 8\text{k}$ -RAM,

1MBit-RAM chips with a $8k \times 128$ organization are required, for example. Therefore, the system architecture has to be slightly modified in order to make effective use of currently available memory chips. Work on this topic is carried out in the moment.

Digital / Analog Implementation

The largest computational load implementing an AM is incurred by the weighted sum of input signals (2). Using analog circuit techniques [10], this sum can be effectively computed by summing analog currents or charge packets, for example. In Figure 10 a simple circuit concept is proposed in CMOS technology. The matrix operation is calculated by current summing and the threshold operation is done by an analog voltage comparator.




Figure 10 Analog implementation of an Associative Matrix column.

The accuracy of analog circuits is not as high as for digital circuits, but they can be built much more compactly and are more appropriate for the highly parallel signal transfer operations immanent in neural networks. Nonlinearities or parasitic effects of the devices allow us to realize complex functions, as an exponential or square-root function [11]. Note, however, that analog circuits are not so densely integrated as it may seem at first glance. They demand large area transistors to assure an acceptable precision and to provide good matching of functional transistor pairs, as used in current mirrors or differential stages. Furthermore, analog circuits are influenced by device mismatches due to the fabrication process and it is very difficult to control offset voltages, for example. Consequently, analog implementations should be applied to artificial neural networks requiring only modest precision. One example for such a network is the AM because there are only $\log(m)$ terms contributing to the weighted sum S_i . The required accuracy of an AM is only about 4 to 5 bits even for large matrices ($m, n > 10,000$) and in the range of analog circuit techniques.

The design of the connection element is based on conventional storage devices (ROM, RAM, EEPROM). For example, a conventional static memory cell has to be enlarged by two transistors (Figure 10), an EEPROM cell requires no additional transistors [12]. Hence, one million programmable connections can be integrated on one chip with current VLSI techniques. The 8192x8192-AM requires 64 of such chips each comprising 128 columns.

Even more limiting to the overall size of an AM slice than the area needed for the connections are the pin requirements of each slice. Taking advantage of the sparsely coded patterns, serial as well as parallel transfer of the patterns is suitable. The input/output organization in the serial case is similar to that of the digital implementation (Figure 11a). For a full parallel transfer of the patterns the m rows and n columns of the matrix are divided into g blocks of equal size. Under the assumption that at uppermost one component in each block is active, only m/g (1 of m/g) decoders are needed for a full parallel transfer of the input pattern to the AM (Figure 11b). We calculate the number of pins as:

$$p = \log_2 \left(\frac{m}{g} \right) \cdot g \quad (5)$$

For $g \approx \log(m)$, a 8192x128 AM slice requires less than 130 pins. Because of the full parallel operation a recall occurs within $1\mu\text{s}$. Two test chips in $2.5\mu\text{m}$ CMOS technology, a 64x64 AM according to Figure 11a ($7 \times 8 \text{mm}^2$, $\approx 40,000$ transistors) and a 96x16-AM-slice according to Figure 11b ($7 \times 5 \text{mm}^2$, $\approx 20,000$ transistors) [10], with programmable connections have been fabricated at the University of Dortmund. With both test chips the above mentioned functionality has been tested and verified. A

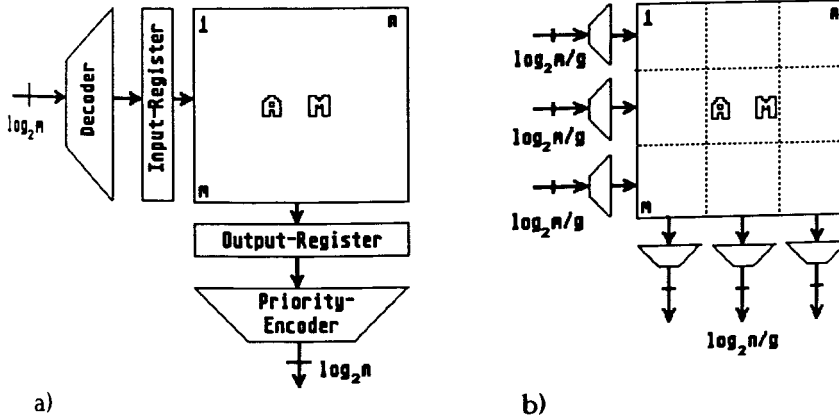


Figure 11 Input/output organization of a semi-parallel (a) and full parallel (b) analog implementation of an Associative Matrix.

microphotograph of the 64x64 AM test chip in CMOS technology is shown in Figure 12. In this case, data input is done through a buffer in order to reduce the pin count, whereas for test reasons the output of each column is attached to a separate pad. The accuracy of the proposed analog circuits turned out to be at least 4 bit. In other words, sparsely coded binary patterns with up to 16 activated components can be handled correctly by this simple implementation.

Because of the modular and regular structure of the architecture, the implementation of large AMs ($n, m > 10000$) is feasible. A further attractive feature of the AM is its fault tolerance to defective connections. Even in the presence of 5 % defects, up to 20,000 sparsely coded patterns ($l = 13, k = 3$) can be stored in a 1000x1000 AM with low error probability. Therefore, the AM will also be well adapted for the involving wafer-scale-integration technique.

DISCUSSION

Though the AM system concept is comparatively simple, it has very attractive features in regard to other associative or neural network VLSI implementations:

- the asymptotic storage capacity is $0.69 \cdot n \cdot m$ bits
- the number of sparsely coded patterns that can be stored in an AM is much larger than the number of columns (artificial neurons)





Figure 12 Microphotograph and floorplan of the analog/digital implementation of the 64x64 Associative Matrix test chip.

- the number of operations during association is only $O(\log(n) \cdot m)$ instead of $O(m \cdot n)$
- the simpler circuit design requires less silicon area.

Because of the modular and regular structure of the proposed architectures, the implementation of very large AMs ($n, m > 10000$) is feasible. This aspect is very important for practical applications where the AM has to be extended to a useful number of storage elements. Work on possible applications of an associative memory of this type is done at the moment by different research groups, e.g. in the field of speech recognition, scene analysis and information retrieval.

Comparing both VLSI approaches presented above, we can call on efficient software tools for a fast, reliable and even complex digital system design. For the memory matrix we can use standard RAM chips employing the highest density in devices. In general, the matrix dimensions (n, m) can be extended by using additional RAM chips. An important disadvantage up to now is that most RAM chips have a 1 or 4 bit organization whereas a longer word length (> 16) is more appropriate for the digital implementation.

On the contrary, the design of analog circuits demands much more time, good theoretical knowledge about transistor physics and a heuristic experience of layout. Only a few process lines are characterized by analog circuits. The noise immunity and precision is low compared to digital circuits. The fixed matrix dimension in case of a special-purpose implementation is a further disadvantage. In their favour, we point out that analog circuits can be built much more compactly and are more appropriate for the highly parallel signal transfer operations immanent in neural networks. For example, a 1000×1000 AM can be integrated on one chip, whereas the digital concept requires several slice chips and at least one RAM chip. In conclusion, both approaches have their advantages and it remains to be seen which type of implementation will be more effective in certain applications.

REFERENCES

- [1] T. Kohonen, "Associative Memory: A System-Theoretical Approach", Springer Verlag, Berlin, 1977.
- [2] G. Palm, "Neural Assemblies", Springer Verlag, Berlin, 1982.
- [3] P.C. Treleaven, "Neurocomputers", Int. Journal of Neurocomputing, Vol. 89/1, pp. 4-31, 1989.

- [4] D.J. Willshaw, O.P. Bunemann, H.C. Longuet-Higgins, "A Non-Holographic Model of Associative Memory", *Nature* 222, No. 5197, pp. 960-962, 1969.
- [5] G. Palm, "On Associative Memory", *Biol. Cybernetics* 36, pp. 19-31, 1980.
- [6] G. Palm, "On Associative Memories", in: *Physics of Cognitive Processes*, E.R. Caianiello, Ed., World Science, pp. 380-420, 1986.
- [7] H.J. Bentz, M. Hagstroem, G. Palm, "Information Storage and Effective Data Retrieval in Sparse Matrices", *Neural Networks*, Vol. 2, No. 4, pp. 289-293, 1989.
- [8] G. Palm, T. Bonhoeffer, "Parallel Processing for Associative and Neural Networks", *Biol. Cybern.* 51, pp. 201-204, 1984.
- [9] U. Rueckert, J. Moschner, "A SIMD Architecture for Parallel Simulation of Associative Networks", to be published, 1990.
- [10] K. Goser, U. Hilleringmann, U. Rueckert, K. Schumacher, "VLSI Technologies For Artificial Neural Networks", *IEEE Micro*, Dec. 1989, pp. 28-44, 1989.
- [11] C. Mead, "Analog VLSI and Neural Systems", Addison-Wesley, 1989.
- [12] U. Rueckert, I. Kreuzer, K. Goser, "A VLSI Concept for an Adaptive Associative Matrix based on Neural Networks", *Proc. COMPEURO Hamburg*, May 1987, pp. 31-34, 1987.