



US 20070022063A1

(19) **United States**

(12) **Patent Application Publication**
Lightowler

(10) **Pub. No.: US 2007/0022063 A1**

(43) **Pub. Date: Jan. 25, 2007**

(54) **NEURAL PROCESSING ELEMENT FOR USE
IN A NEURAL NETWORK**

Publication Classification

(75) **Inventor: Neil Lightowler, Old Aberdeen (GB)**

Correspondence Address:

**DAVID S. RESNICK
100 SUMMER STREET
NIXON PEABODY LLP
BOSTON, MA 02110-2131 (US)**

(51) **Int. Cl.**

G06N 3/02 (2006.01)

G06F 15/18 (2006.01)

G06E 1/00 (2006.01)

(52) **U.S. Cl. 706/15**

(73) **Assignee: Axeon Limited, Bridge of Don (GB)**

(21) **Appl. No.: 11/445,484**

(22) **Filed: Jun. 1, 2006**

Related U.S. Application Data

(63) **Continuation of application No. 09/890,816, filed on
Aug. 1, 2001, now Pat. No. 7,082,419, filed as 371 of
international application No. PCT/GB00/00277, filed
on Feb. 1, 2000.**


(30) **Foreign Application Priority Data**

Feb. 1, 1999 (GB) 9902115.6

(57)

ABSTRACT

A neural processing element for use in a modular neural network is provided. One embodiment provides a neural network comprising an array of autonomous modules (300). The modules (300) can be arranged in a variety of configurations to form neural networks with various topologies, for example, with a hierarchical modular structure. Each module (300) contains sufficient neurons (100) to enable it to do useful work as a stand alone system, with the advantage that many modules (300) can be connected together to create a wide variety of configurations and network sizes. This modular approach results in a scaleable system that meets increased workload with an increase in parallelism and thereby avoids the usually extensive increases in training times associated with unitary implementations.






Fig. 1a




Fig. 1b




Fig. 2




Fig. 3




Fig. 4a




Fig. 4b




Fig. 4c

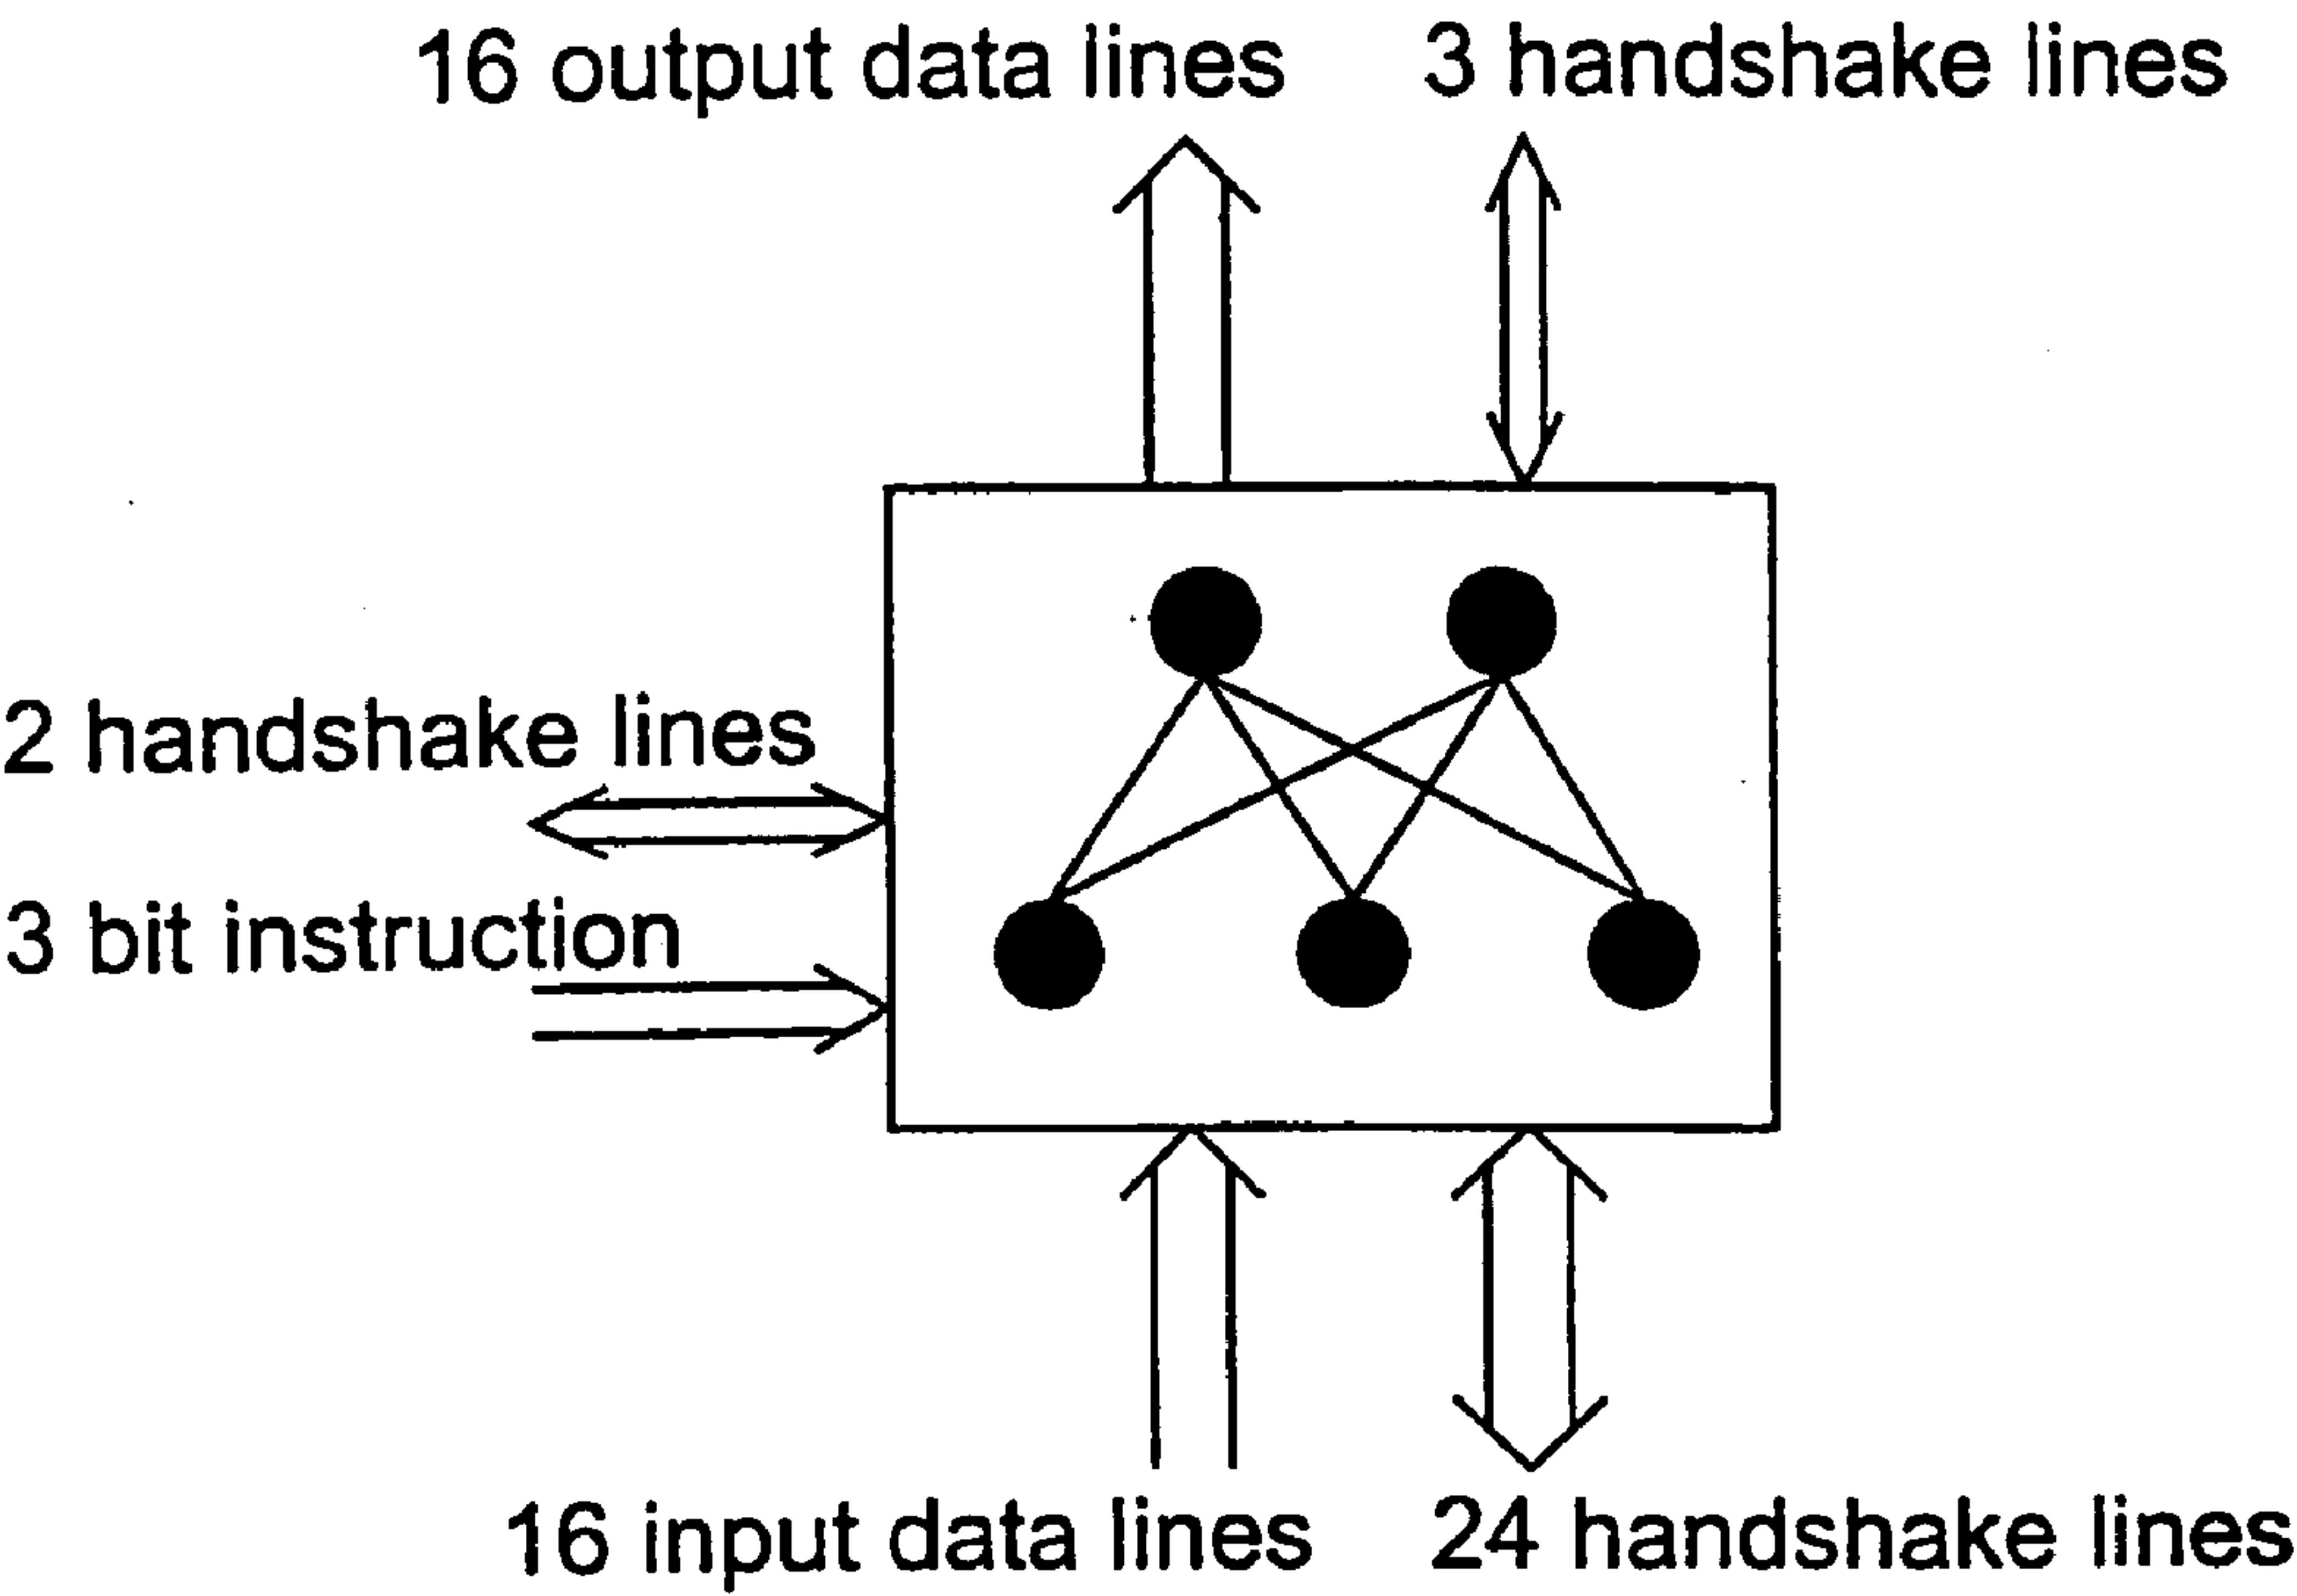


Fig. 5




Fig. 6




Fig. 7




Fig. 8




Fig. 9




Fig. 10




Fig. 11




Fig. 12




Fig. 11a




Fig. 11b

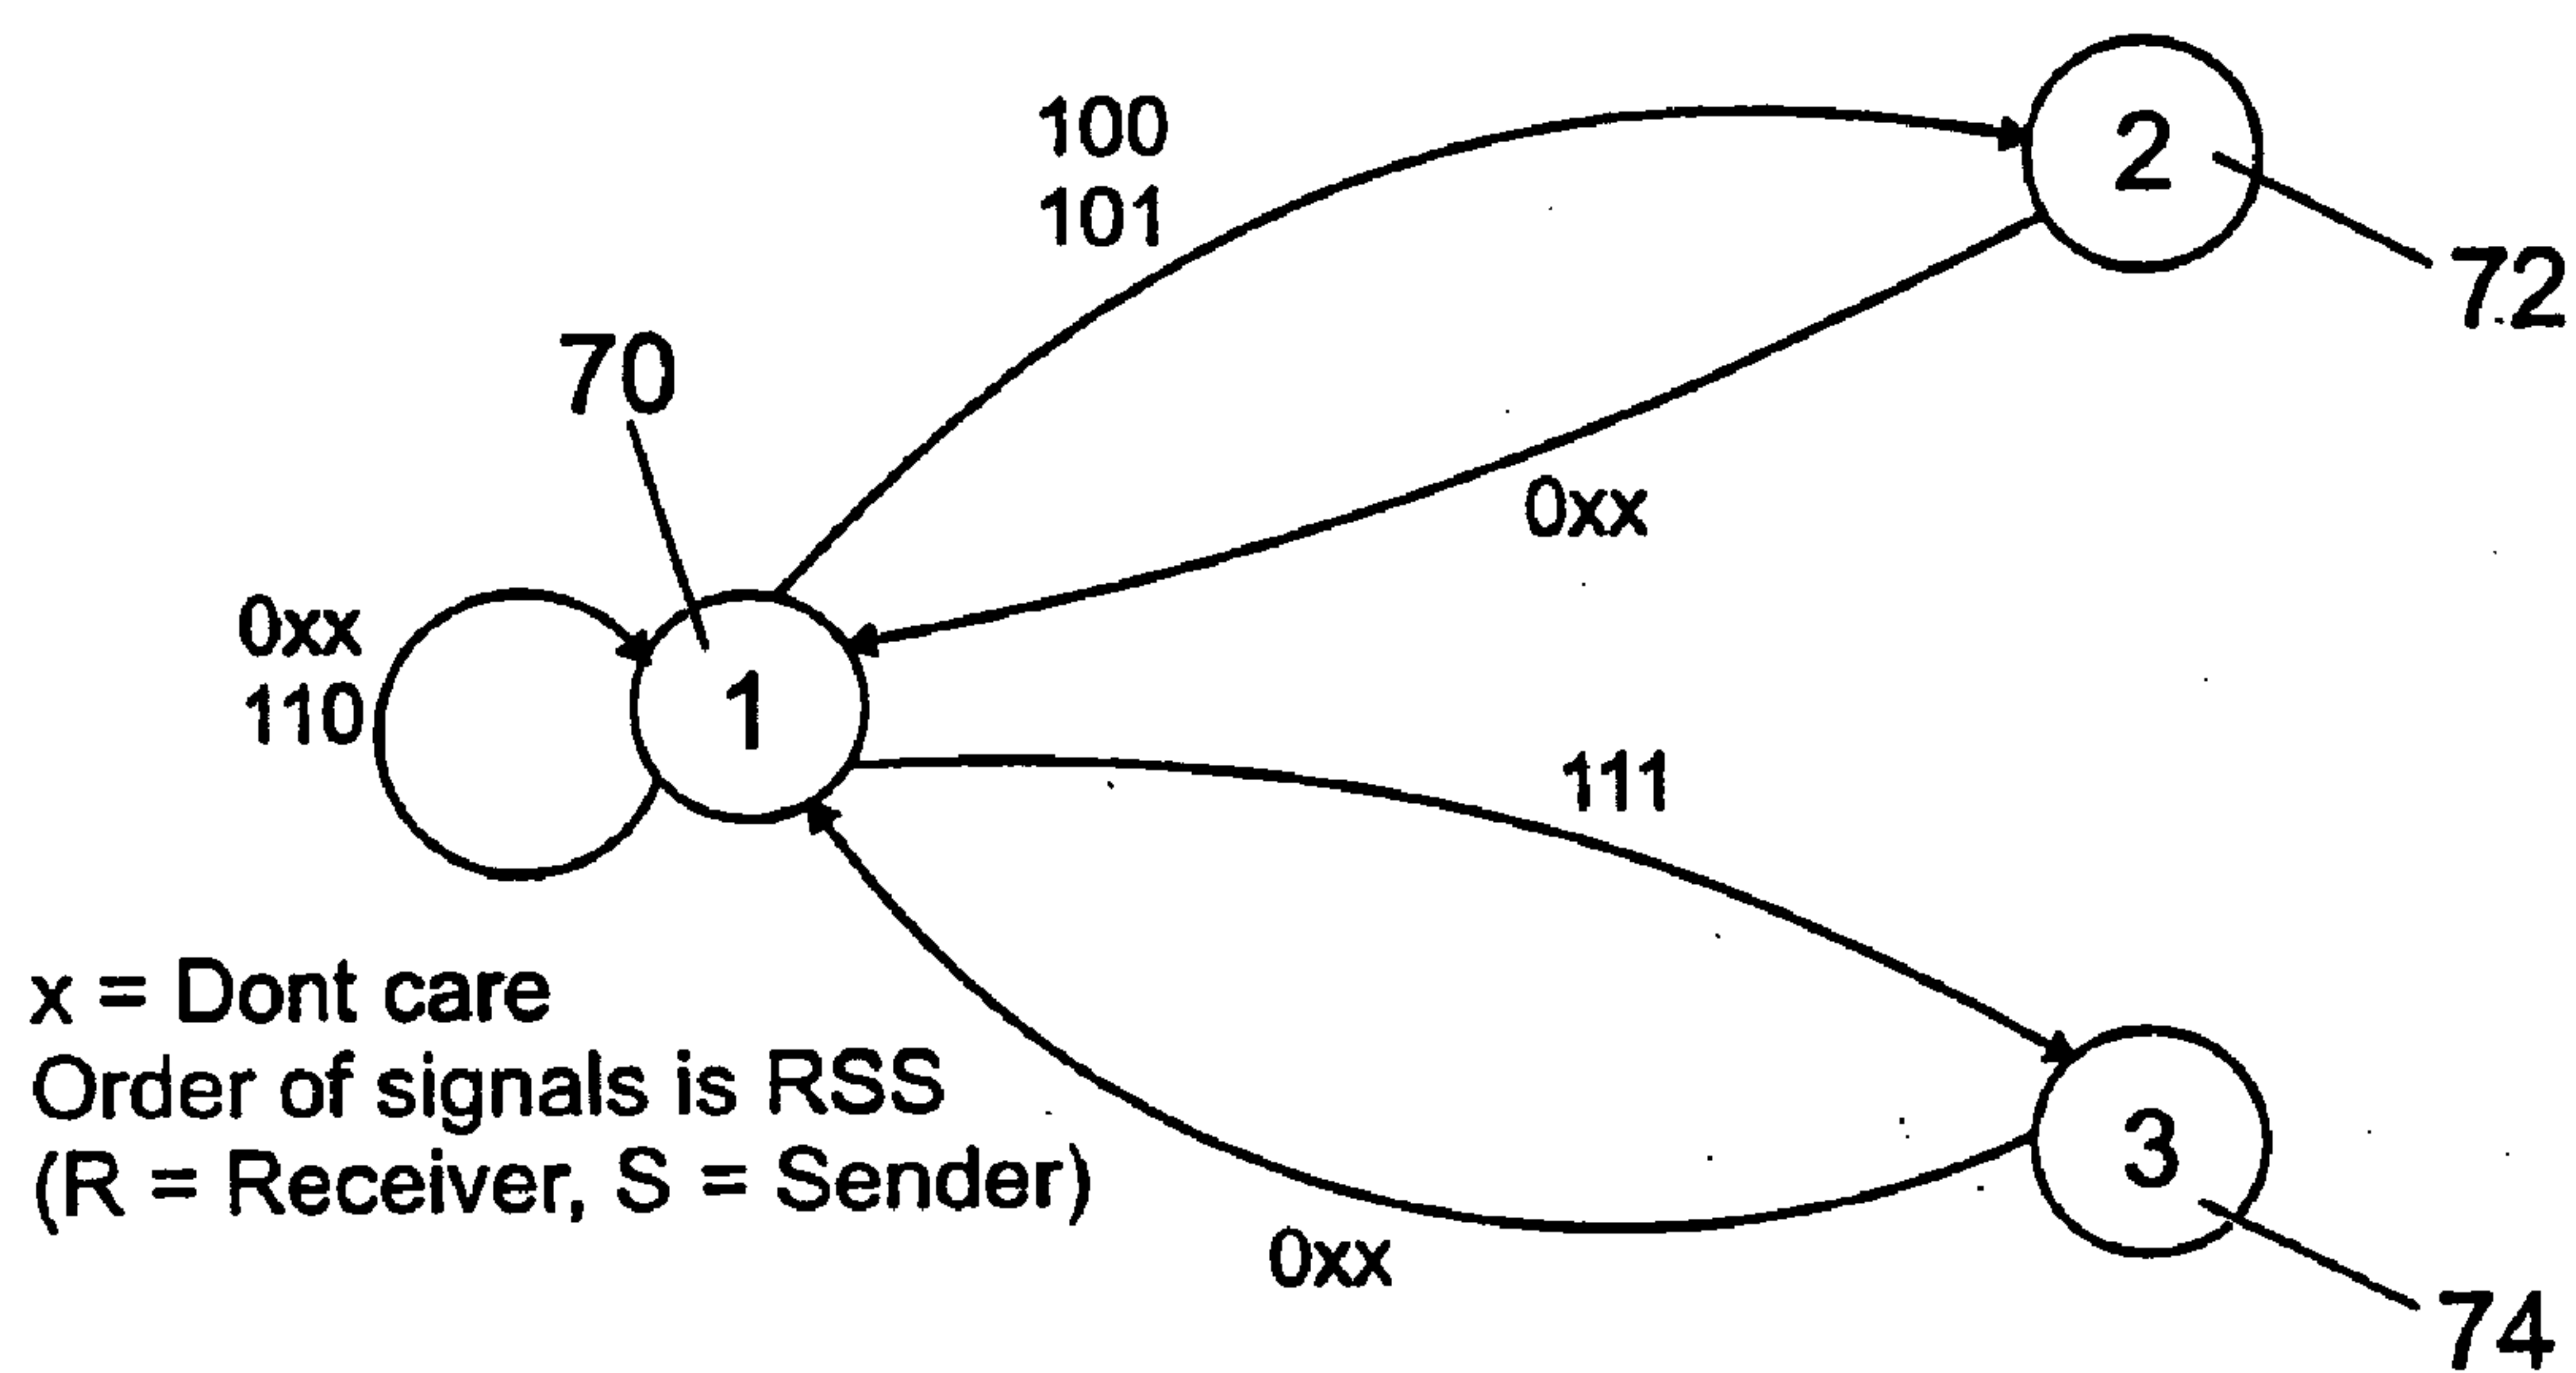


Fig. 13

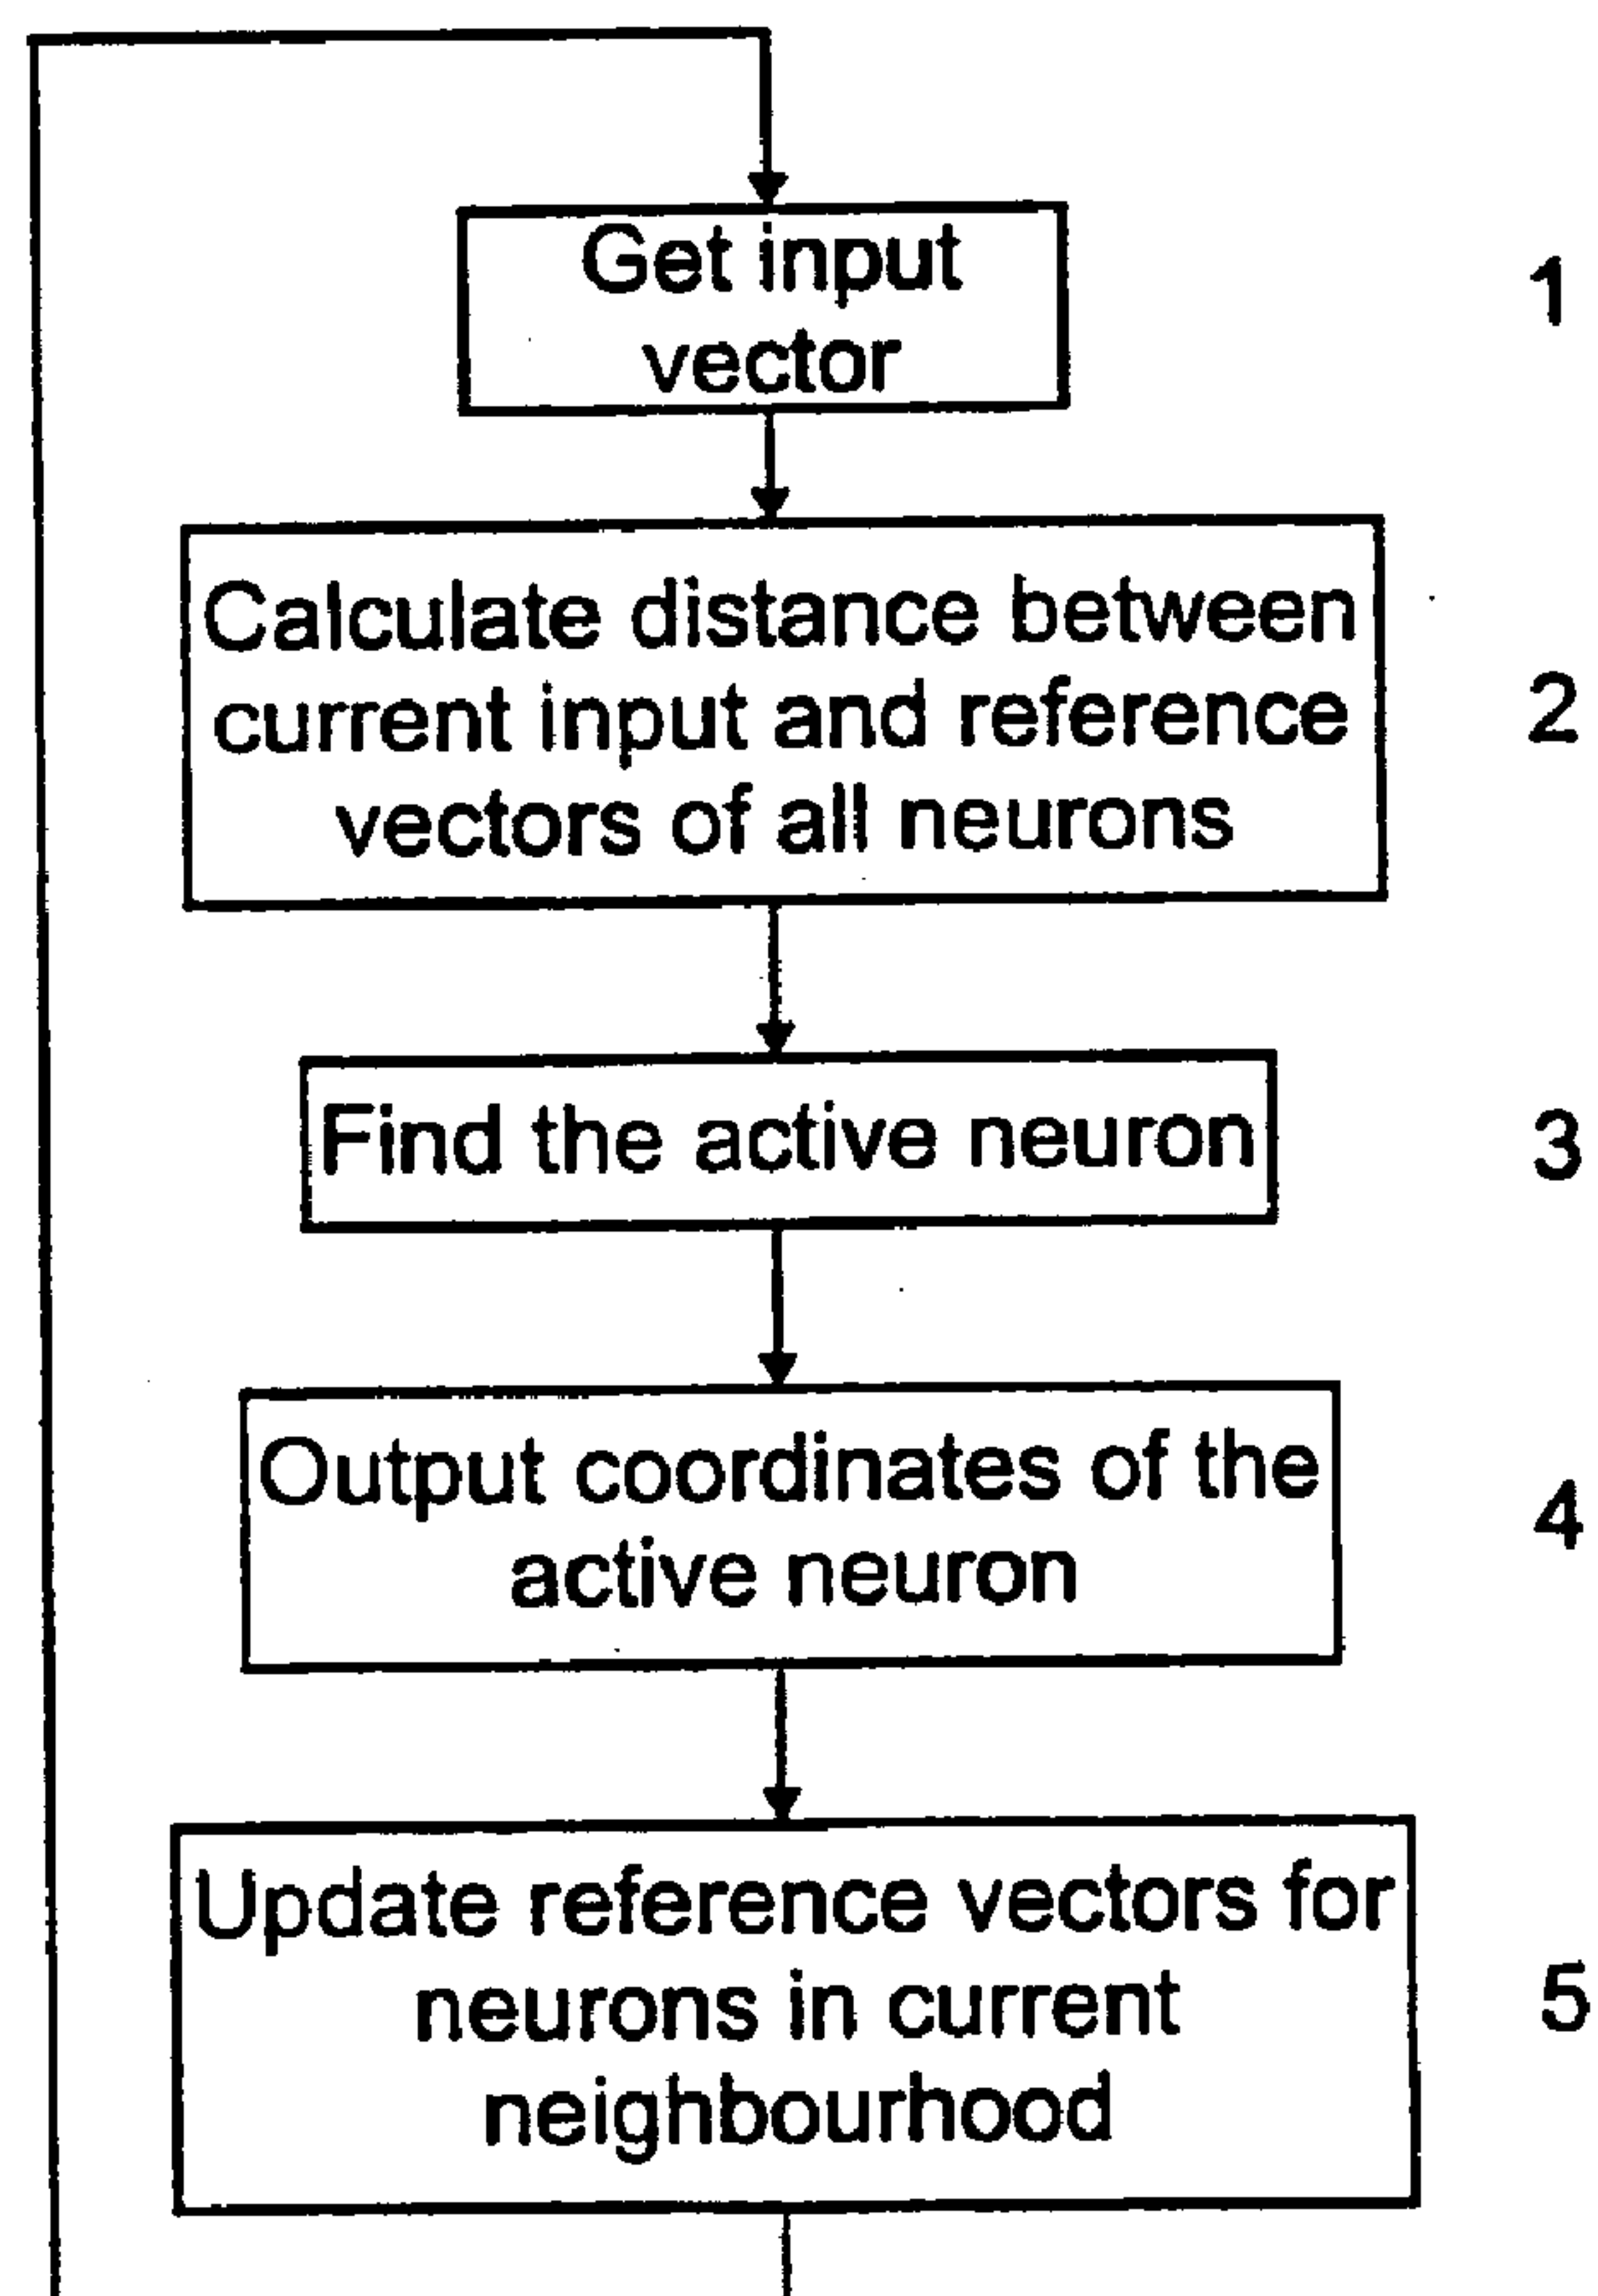


Fig. 14




Fig. 15




Fig. 16




Fig. 17



Fig. 18




Fig. 19a




Fig. 19b

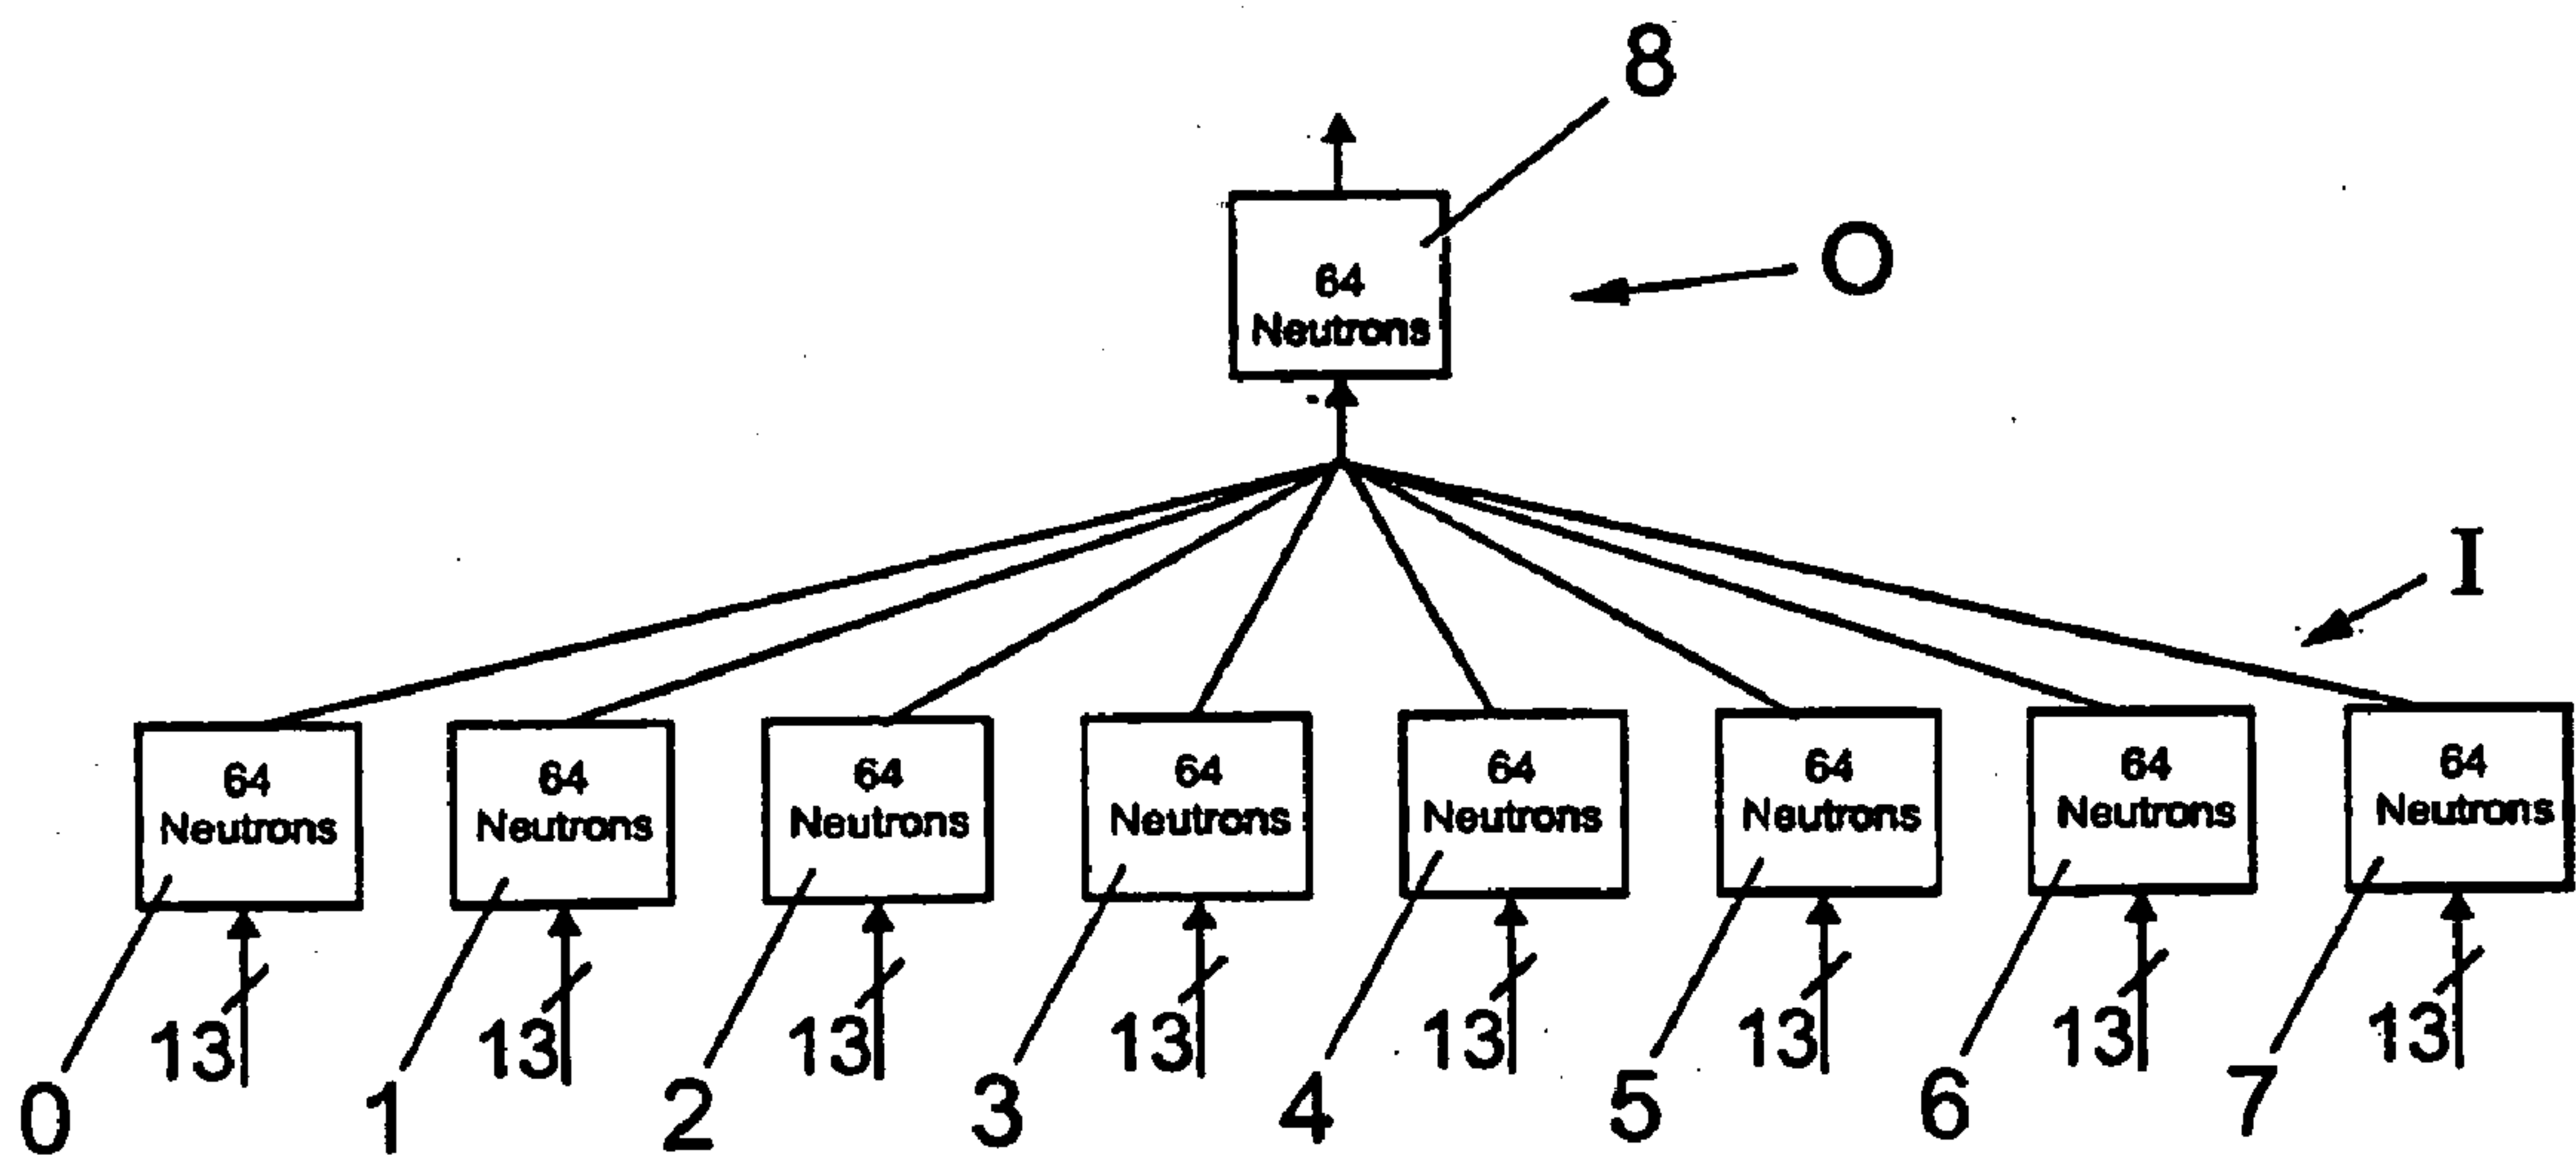


Fig. 20

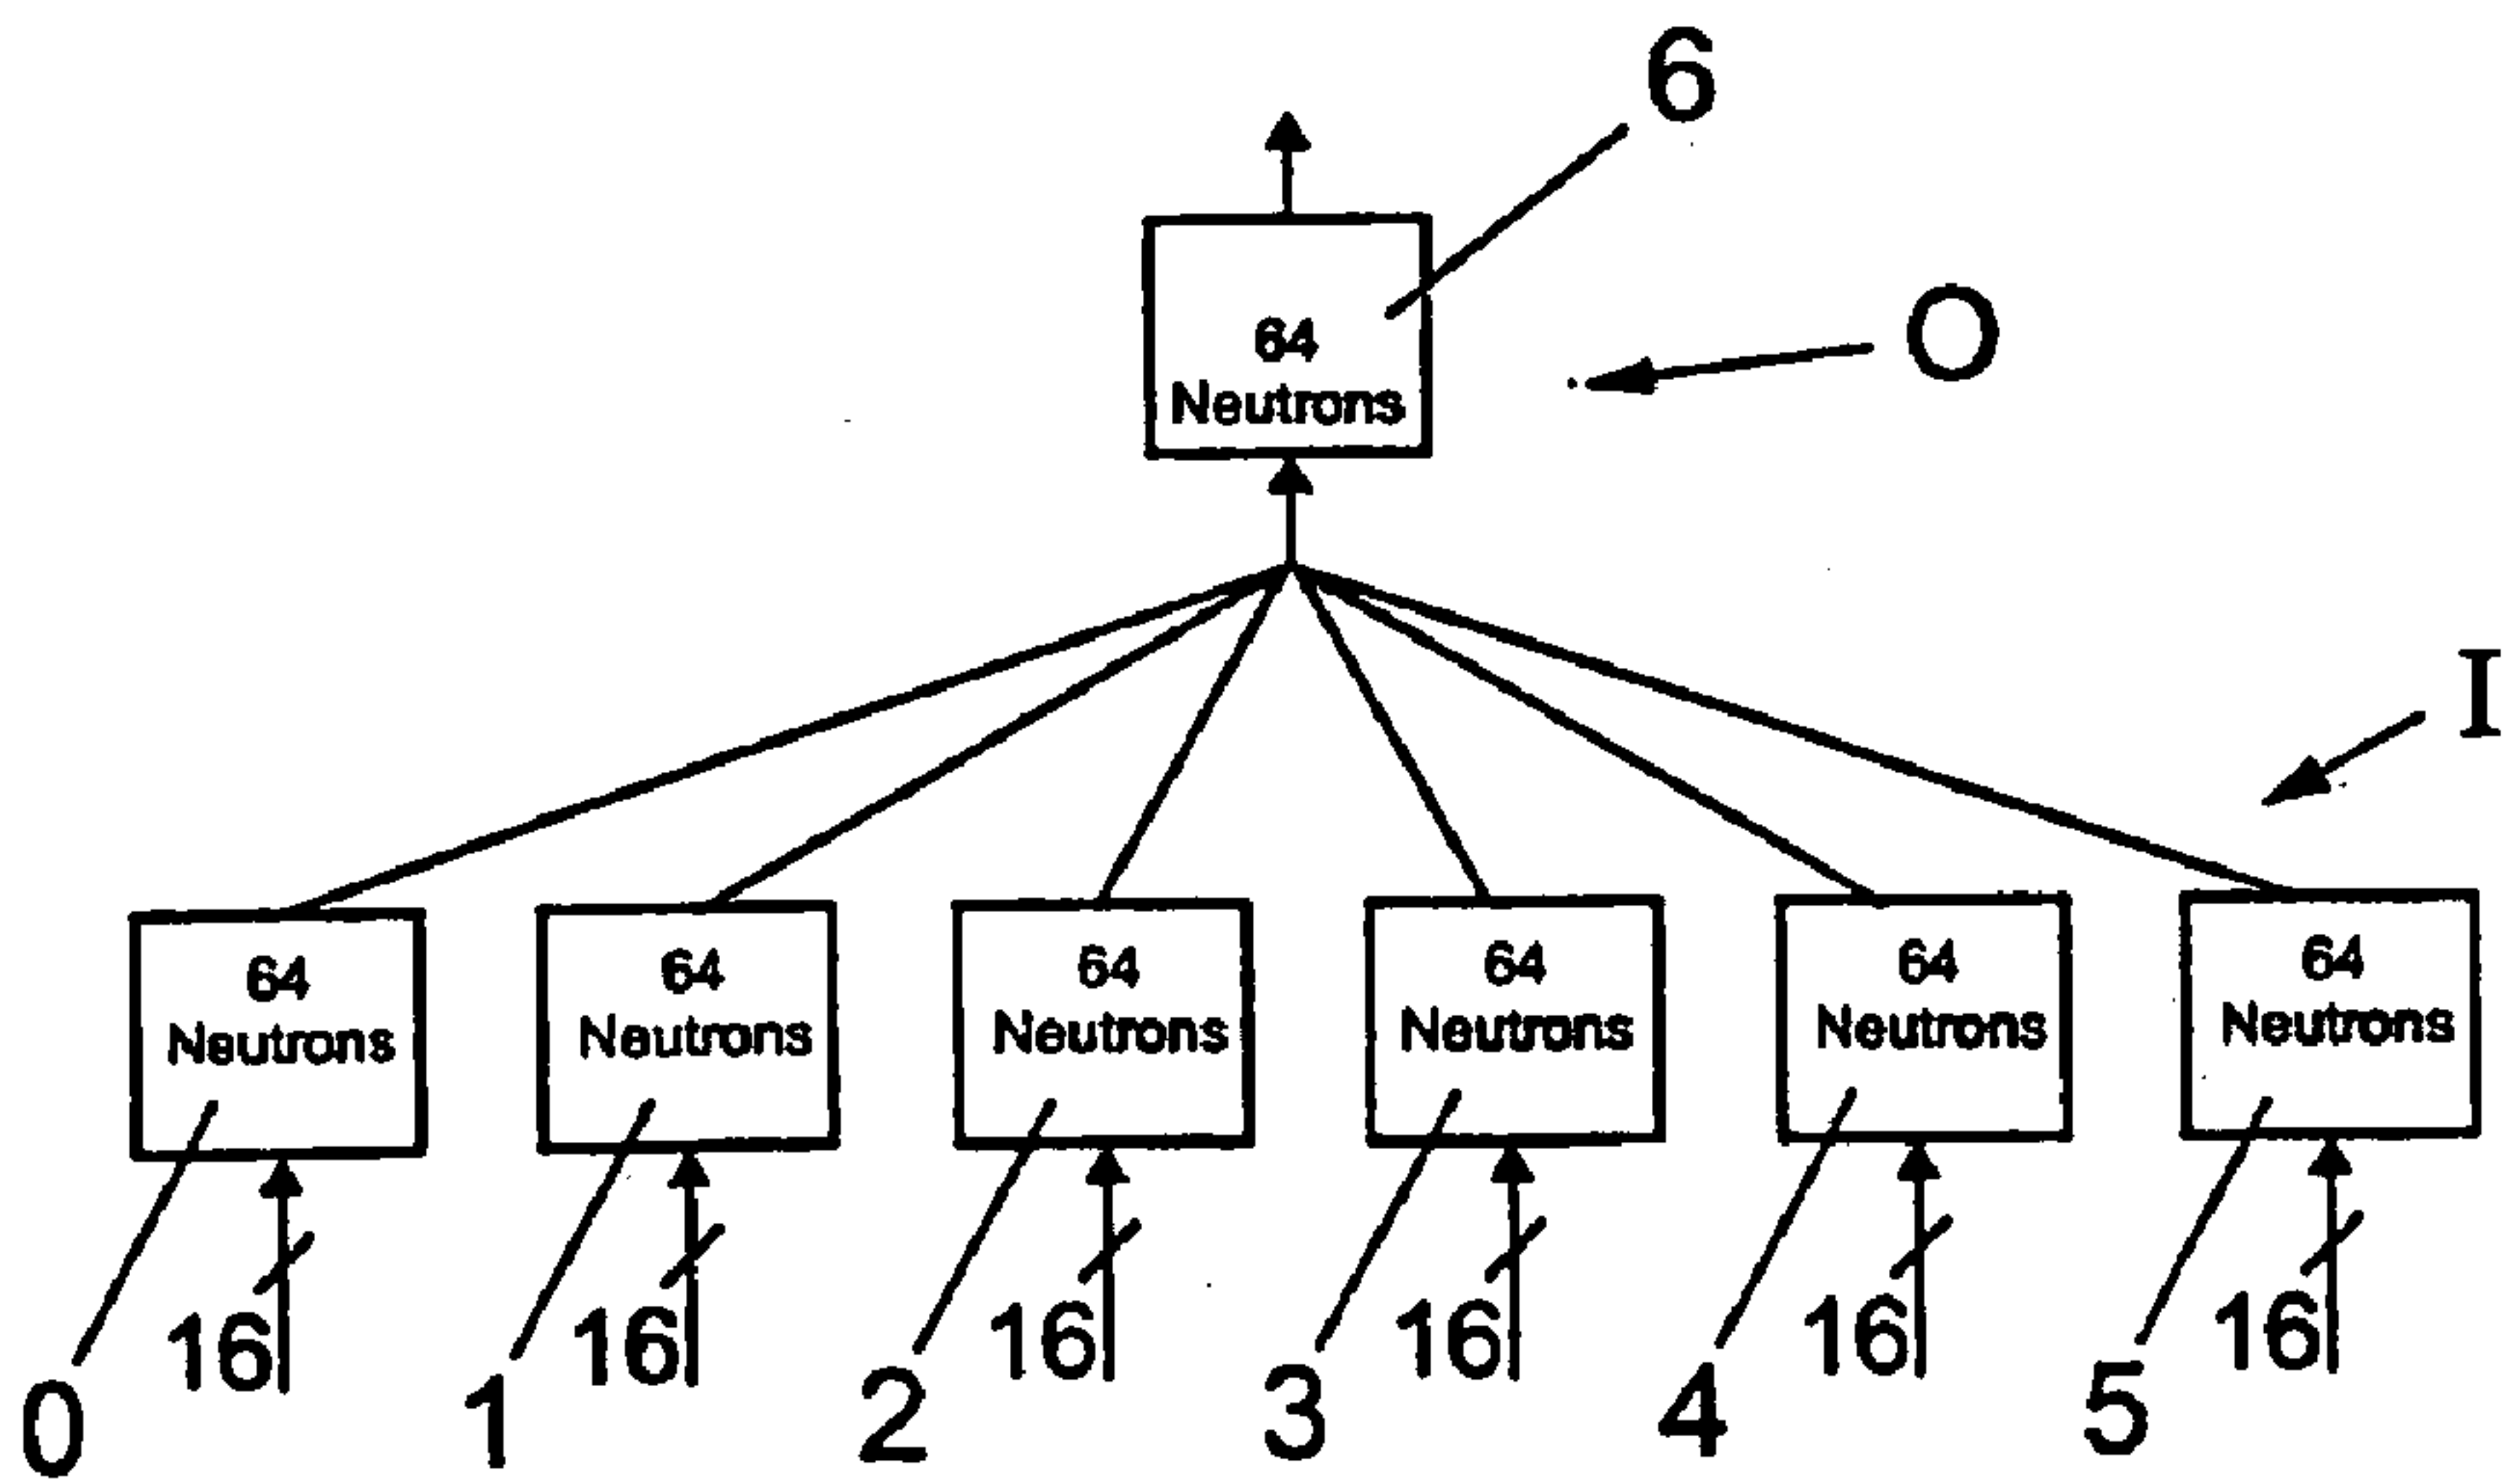


Fig. 21




Fig. 22




Fig. 23

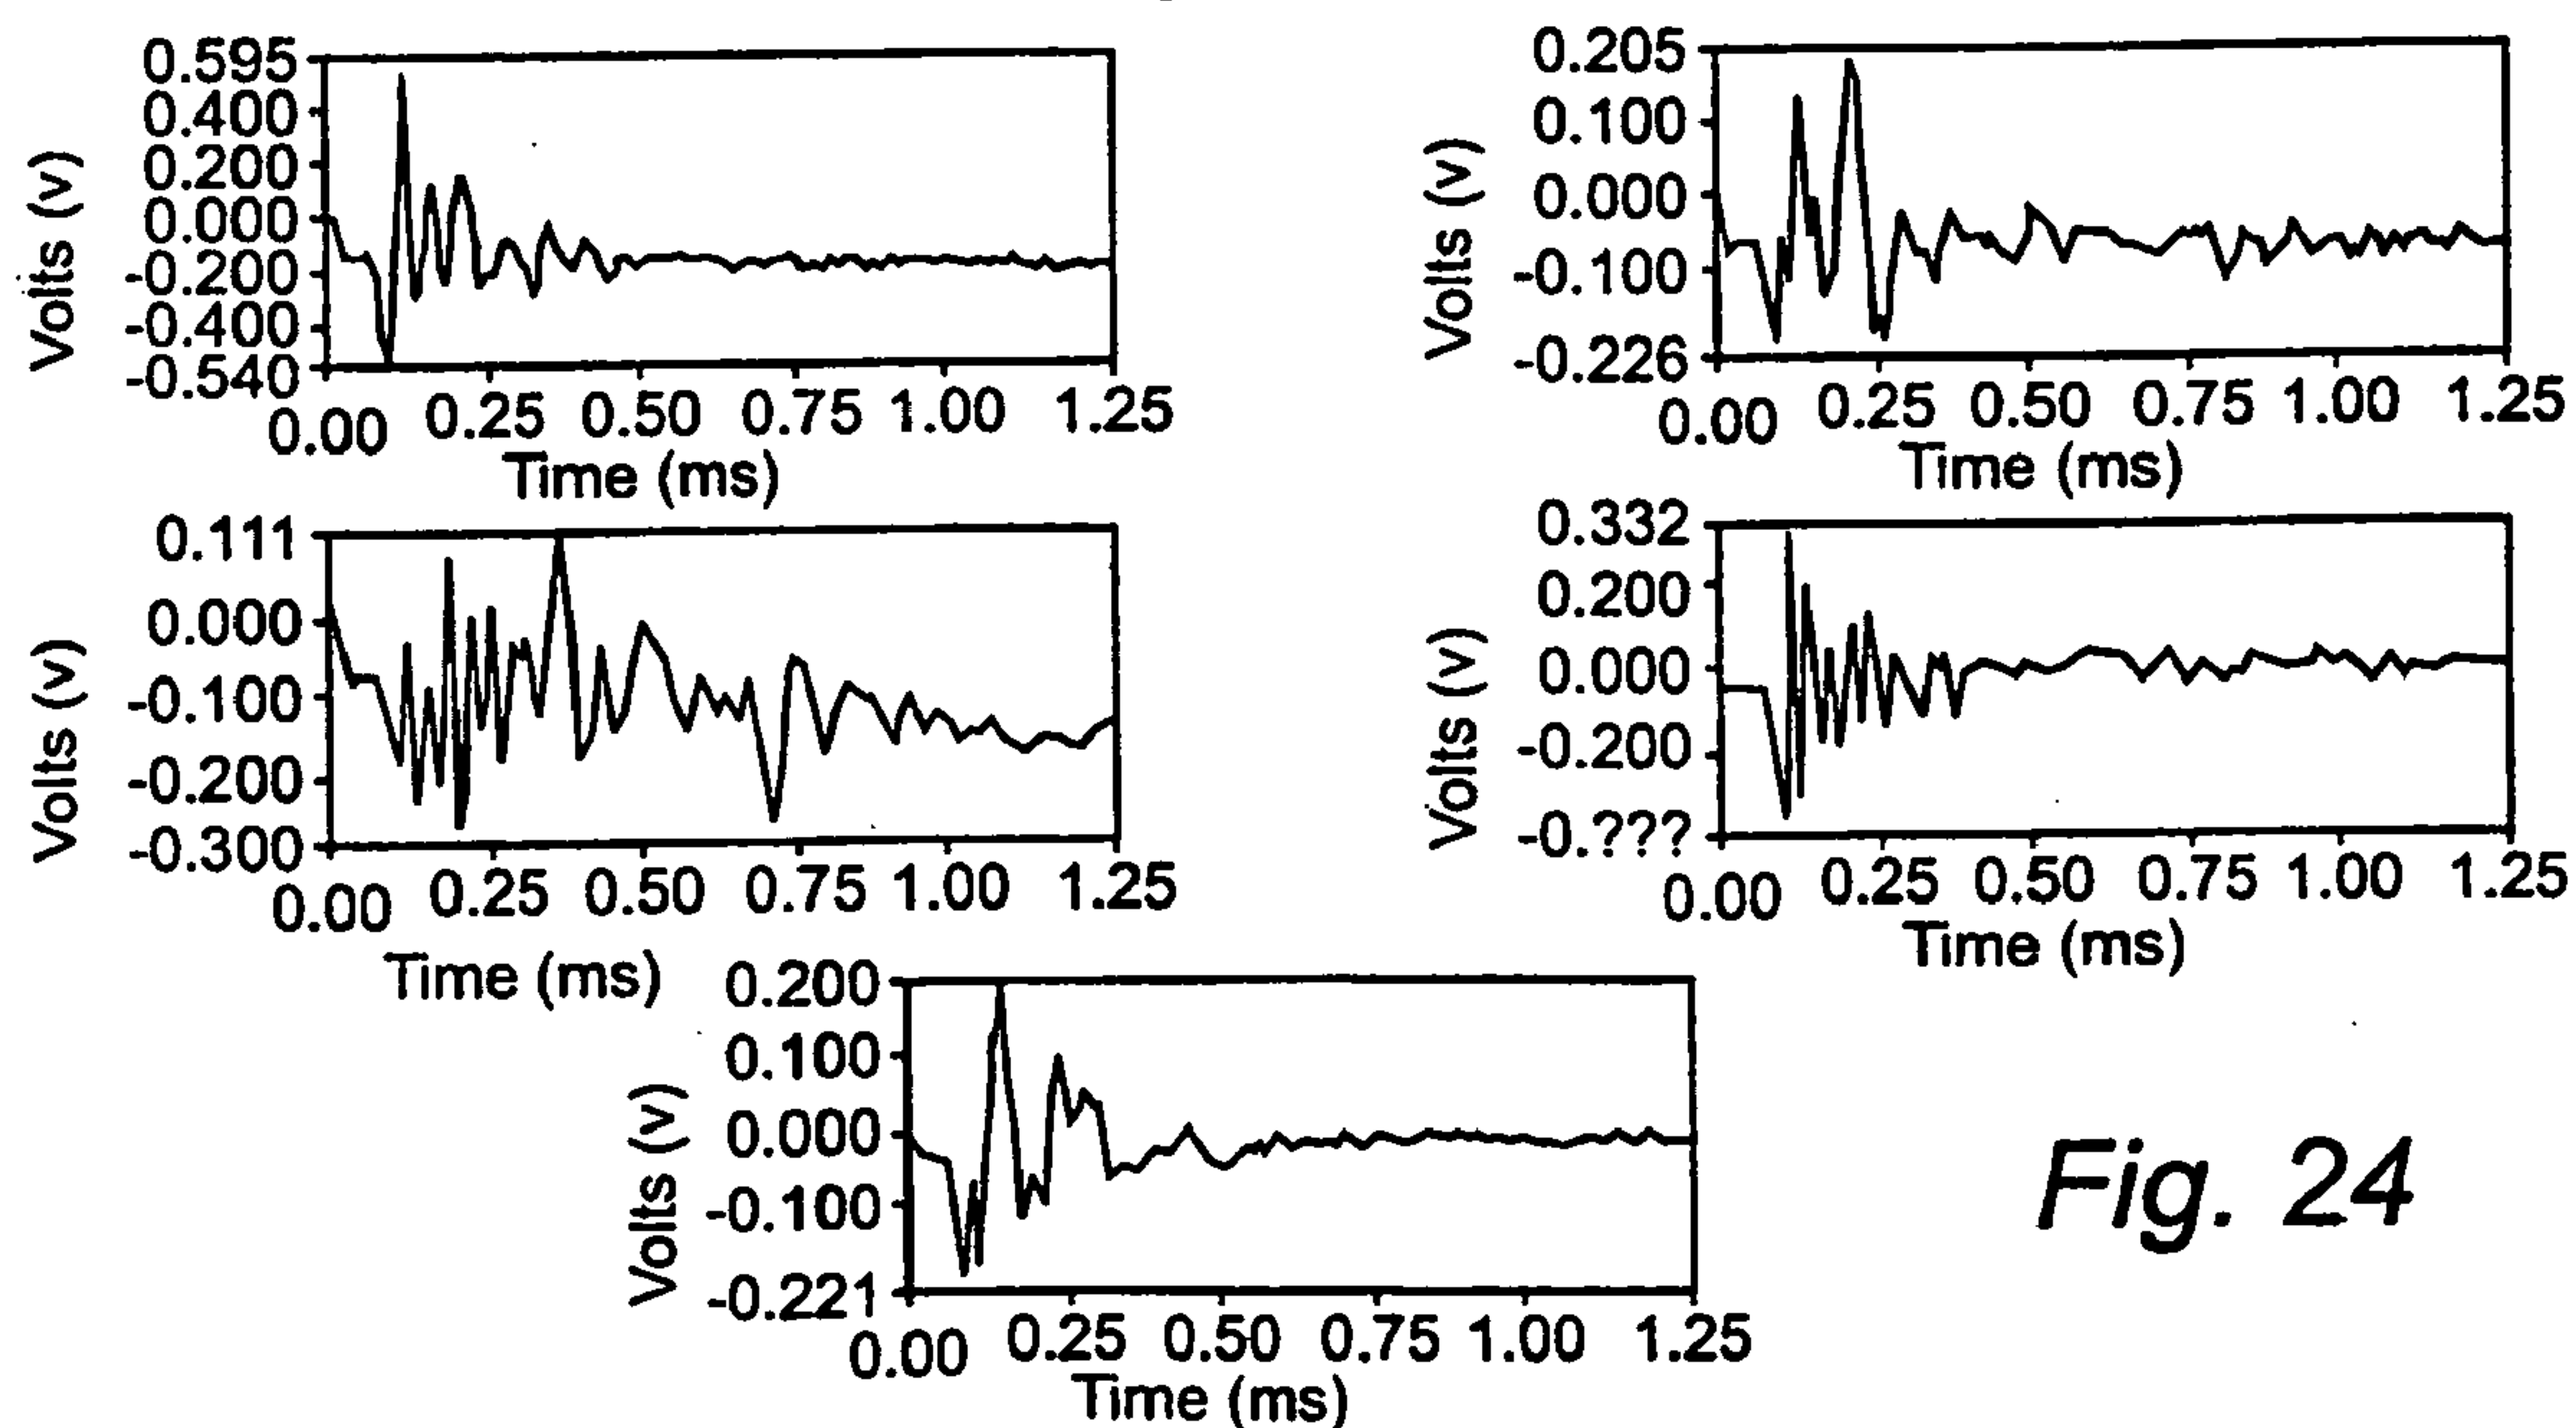


Fig. 24

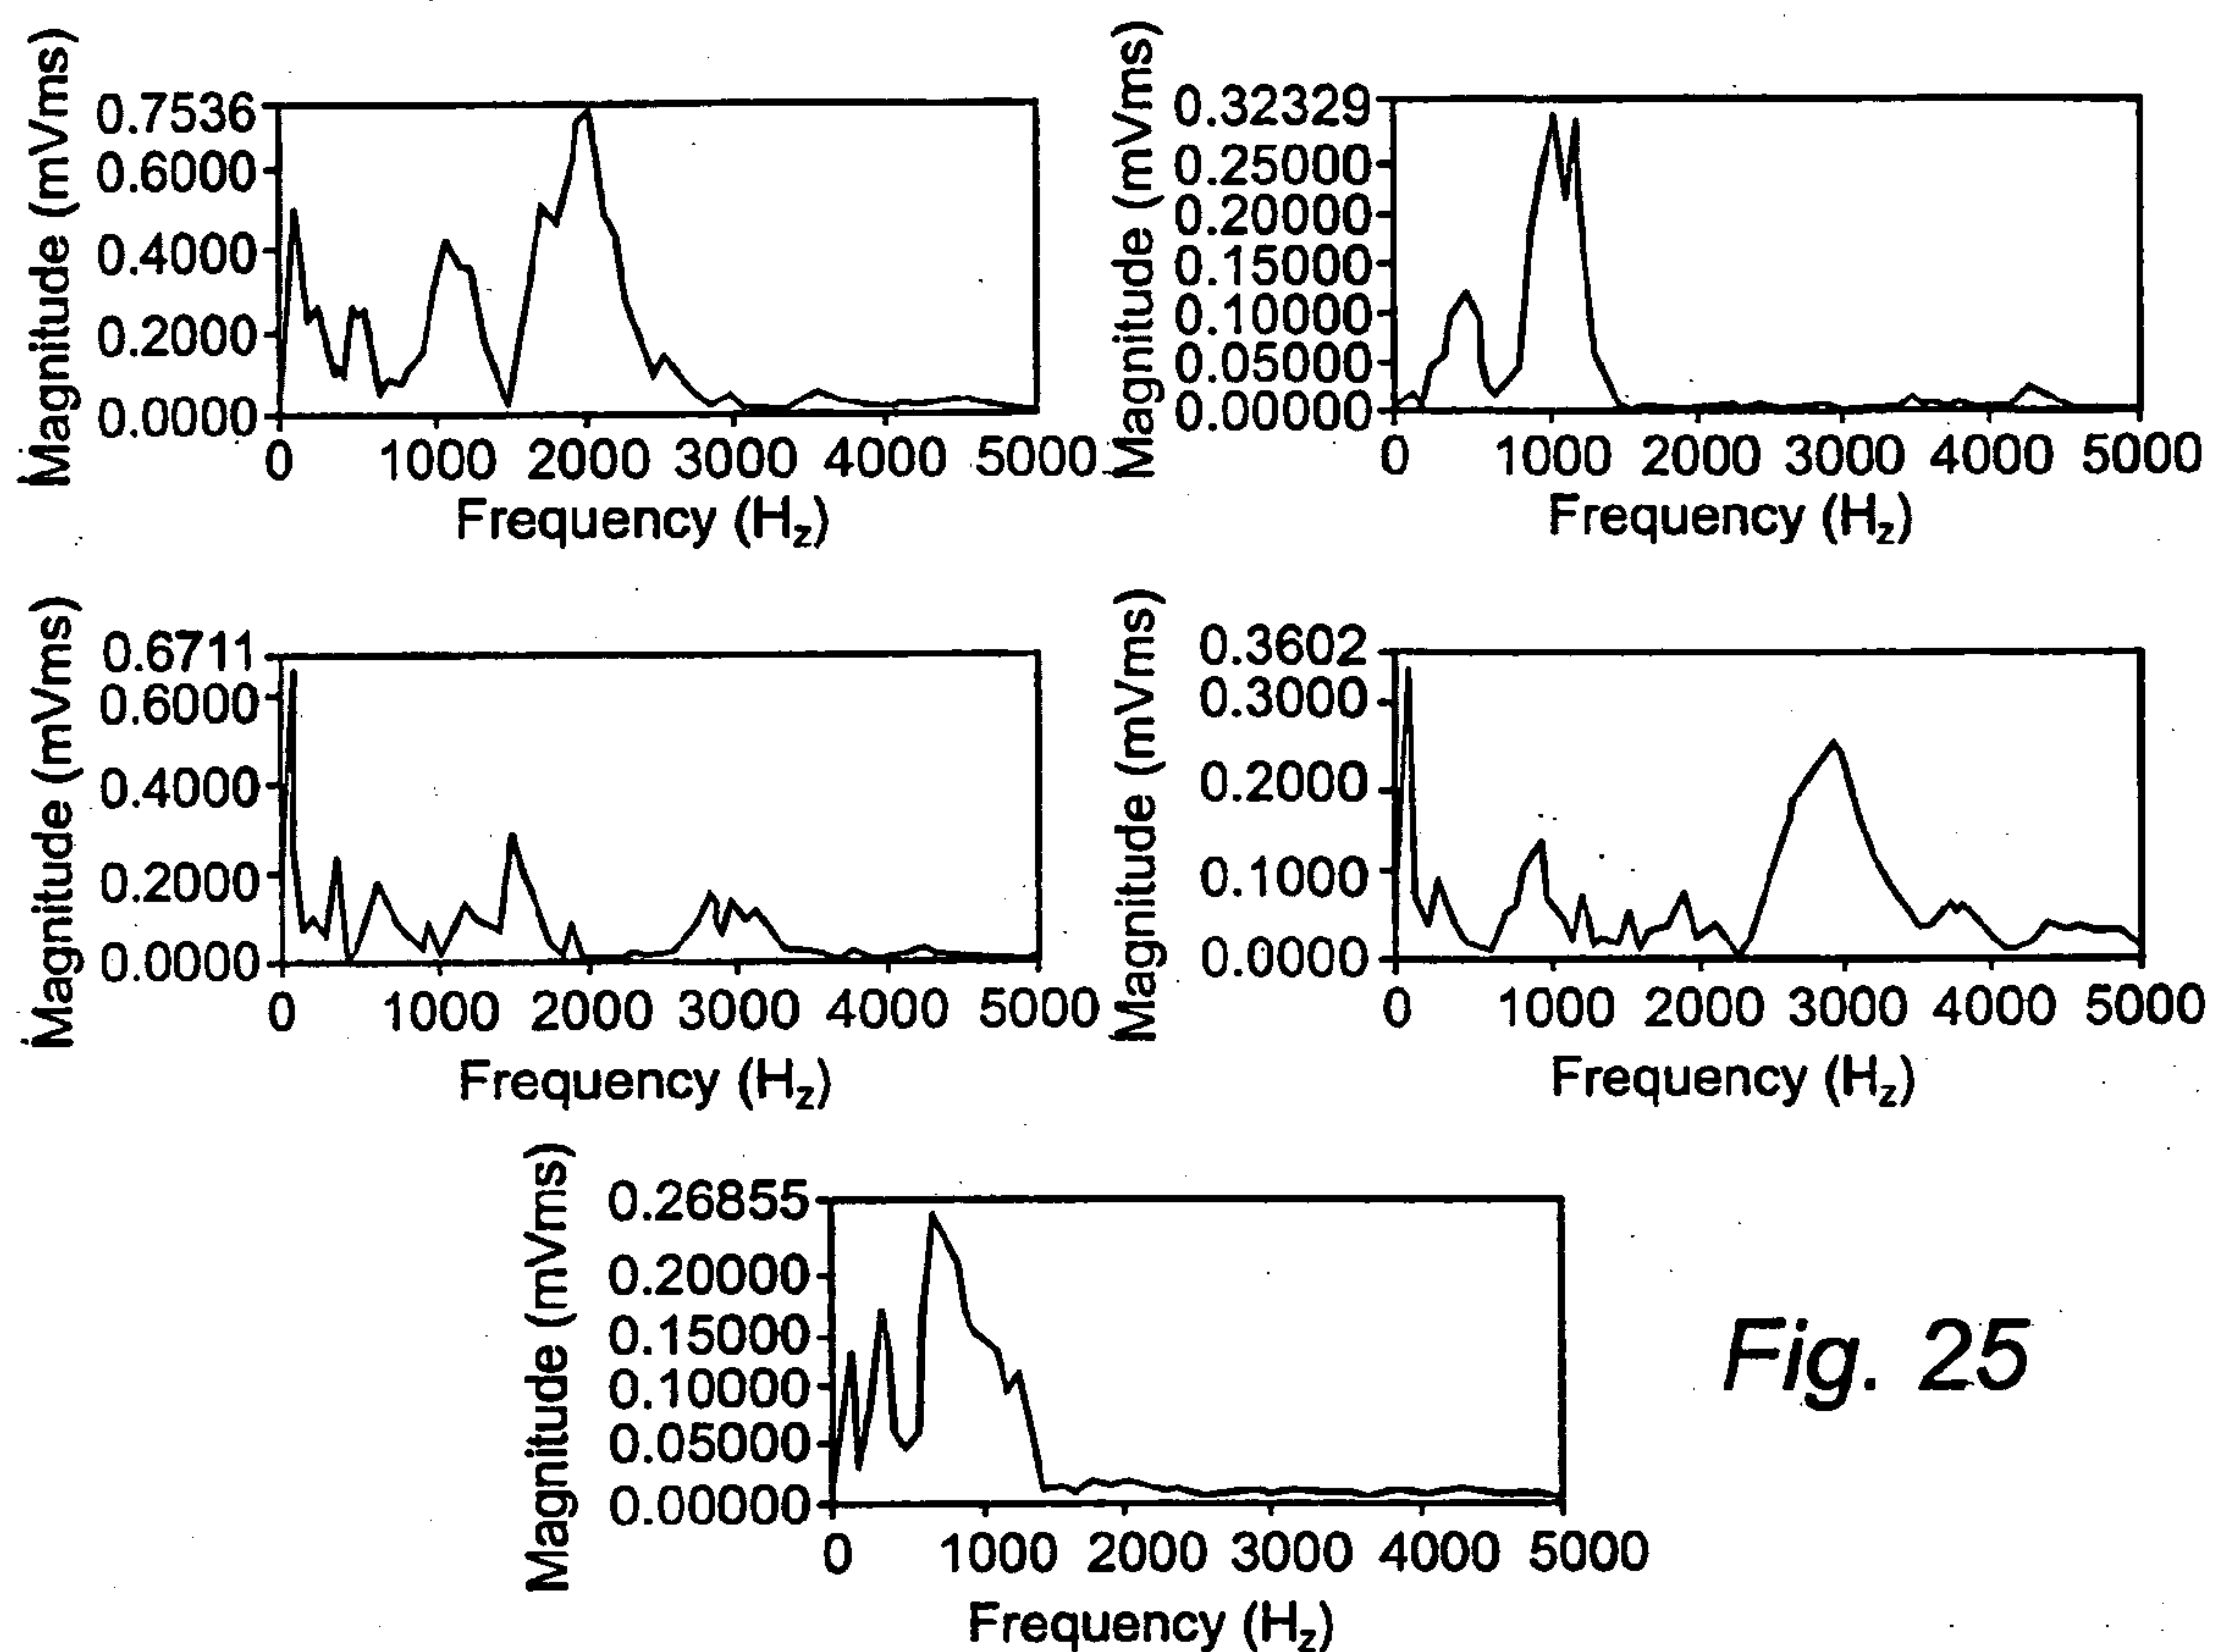


Fig. 25

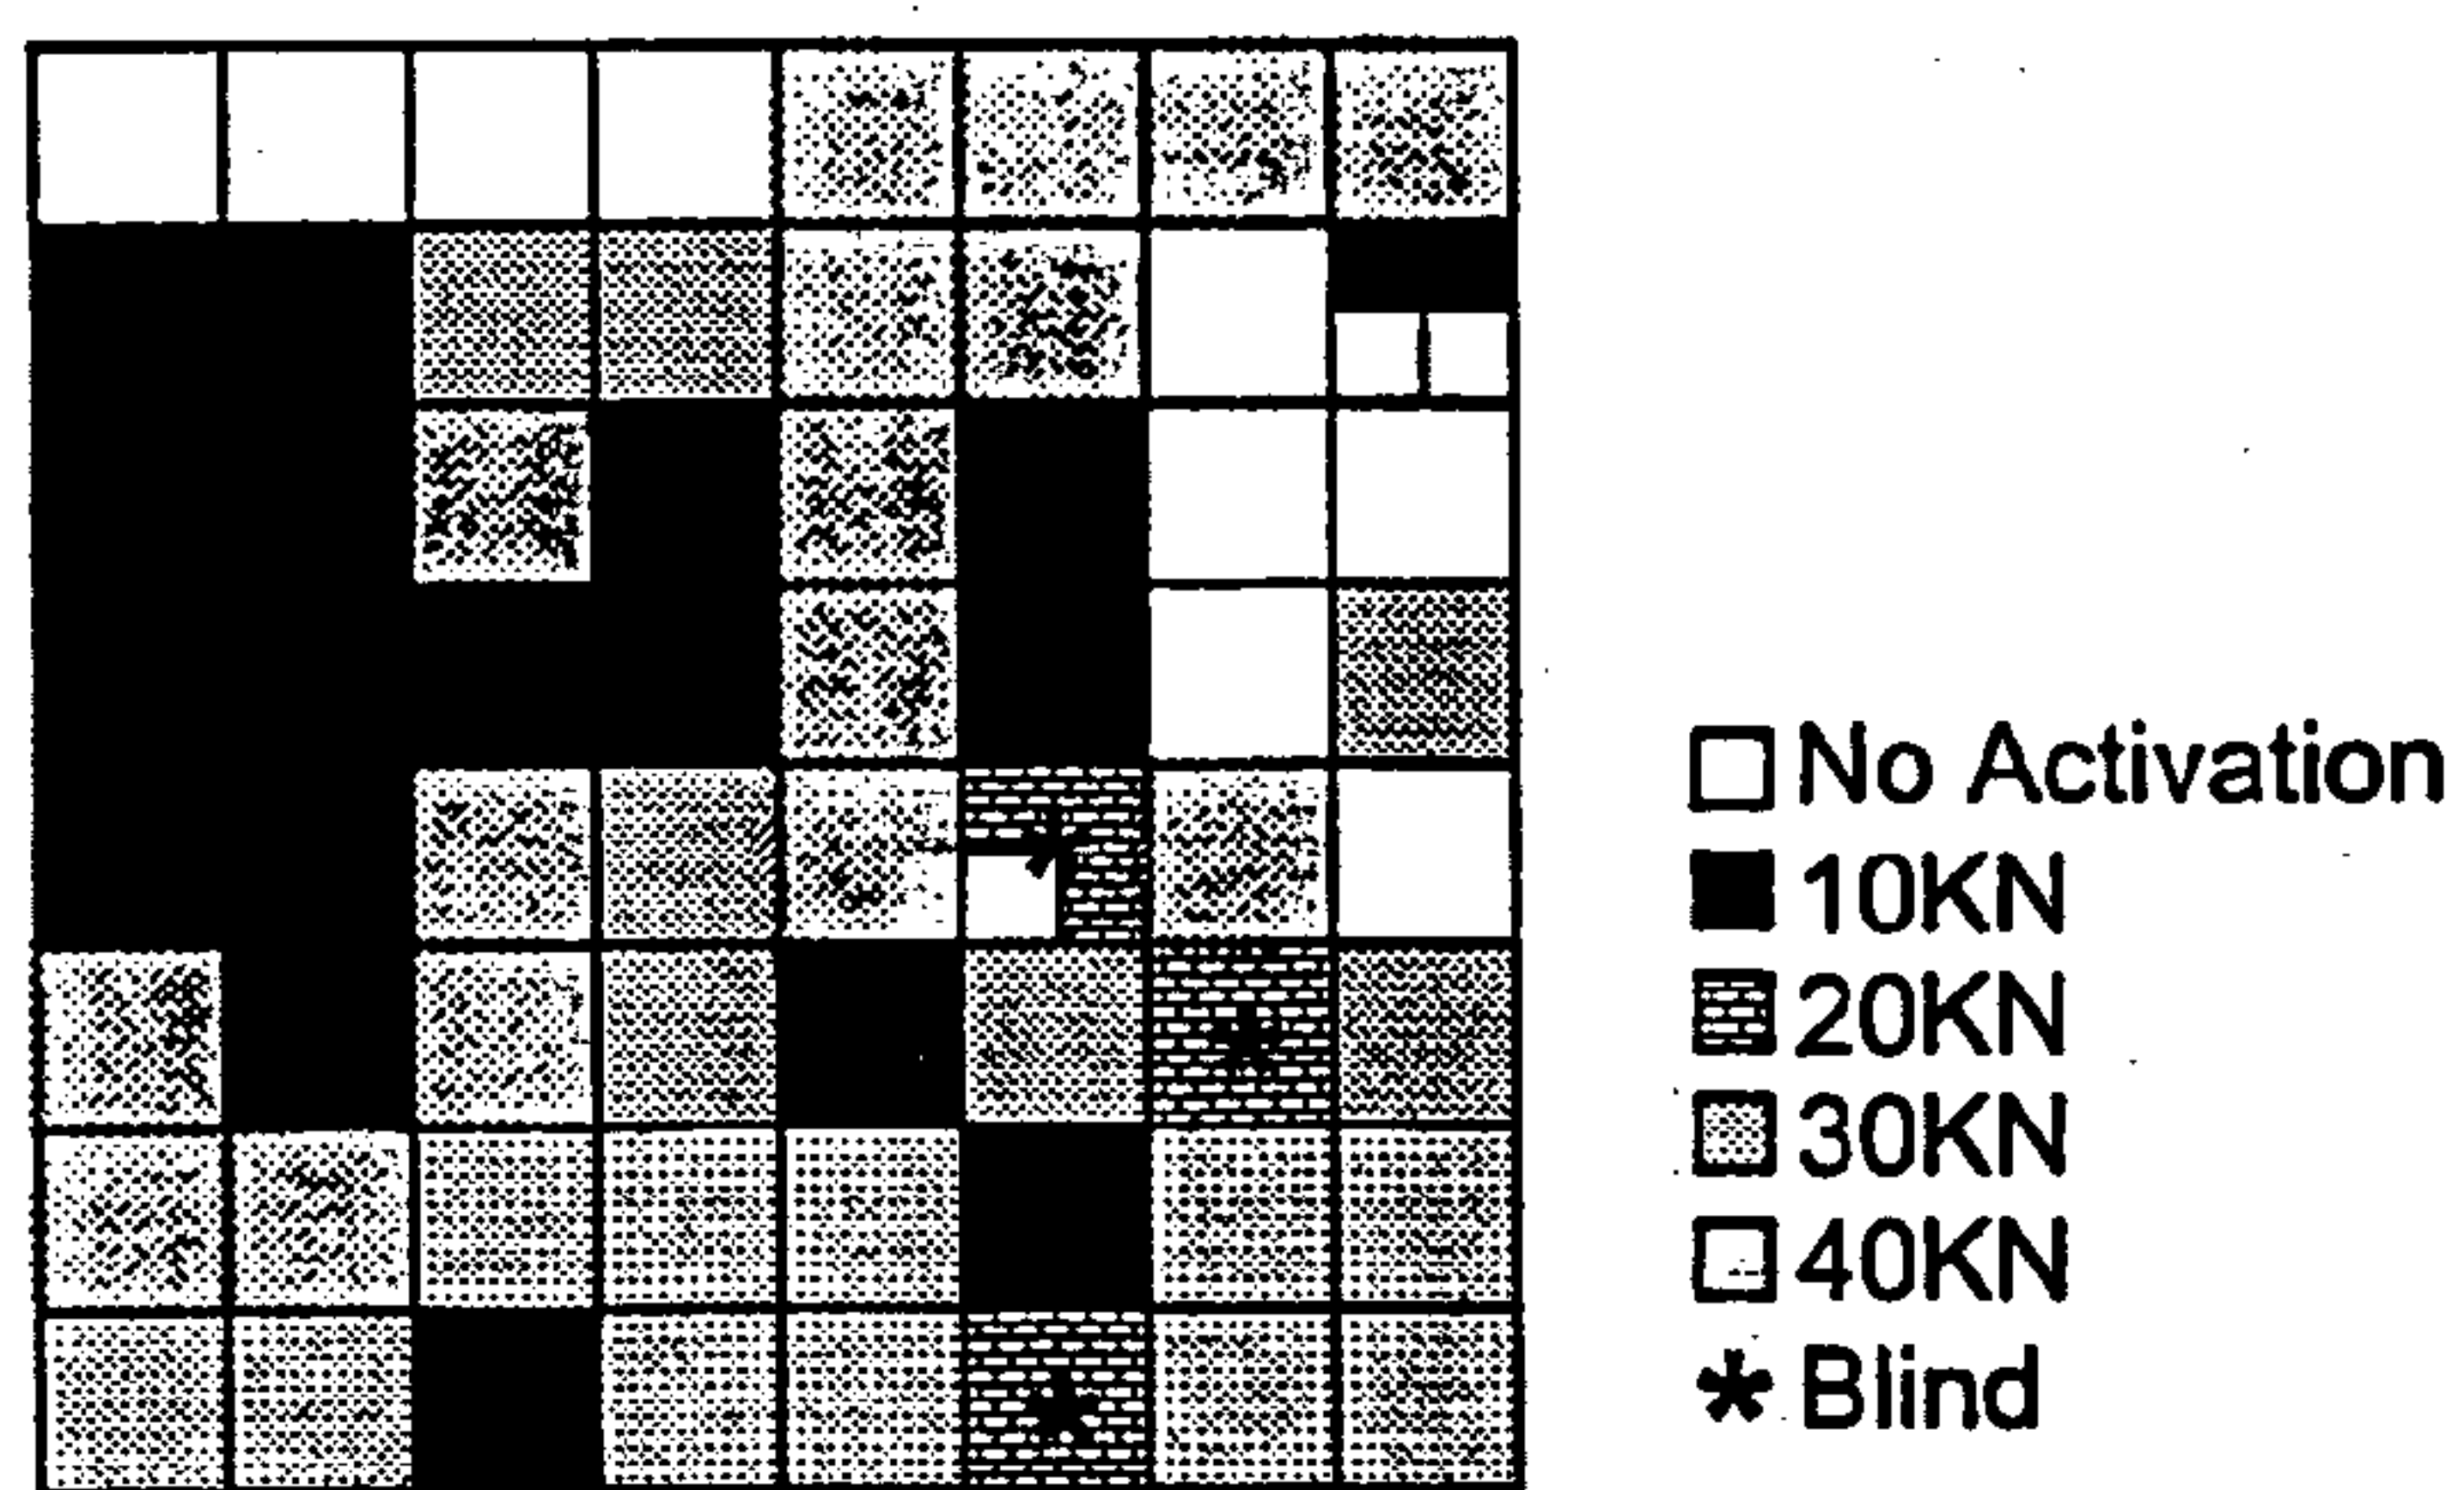


Fig. 26

Fig. 27

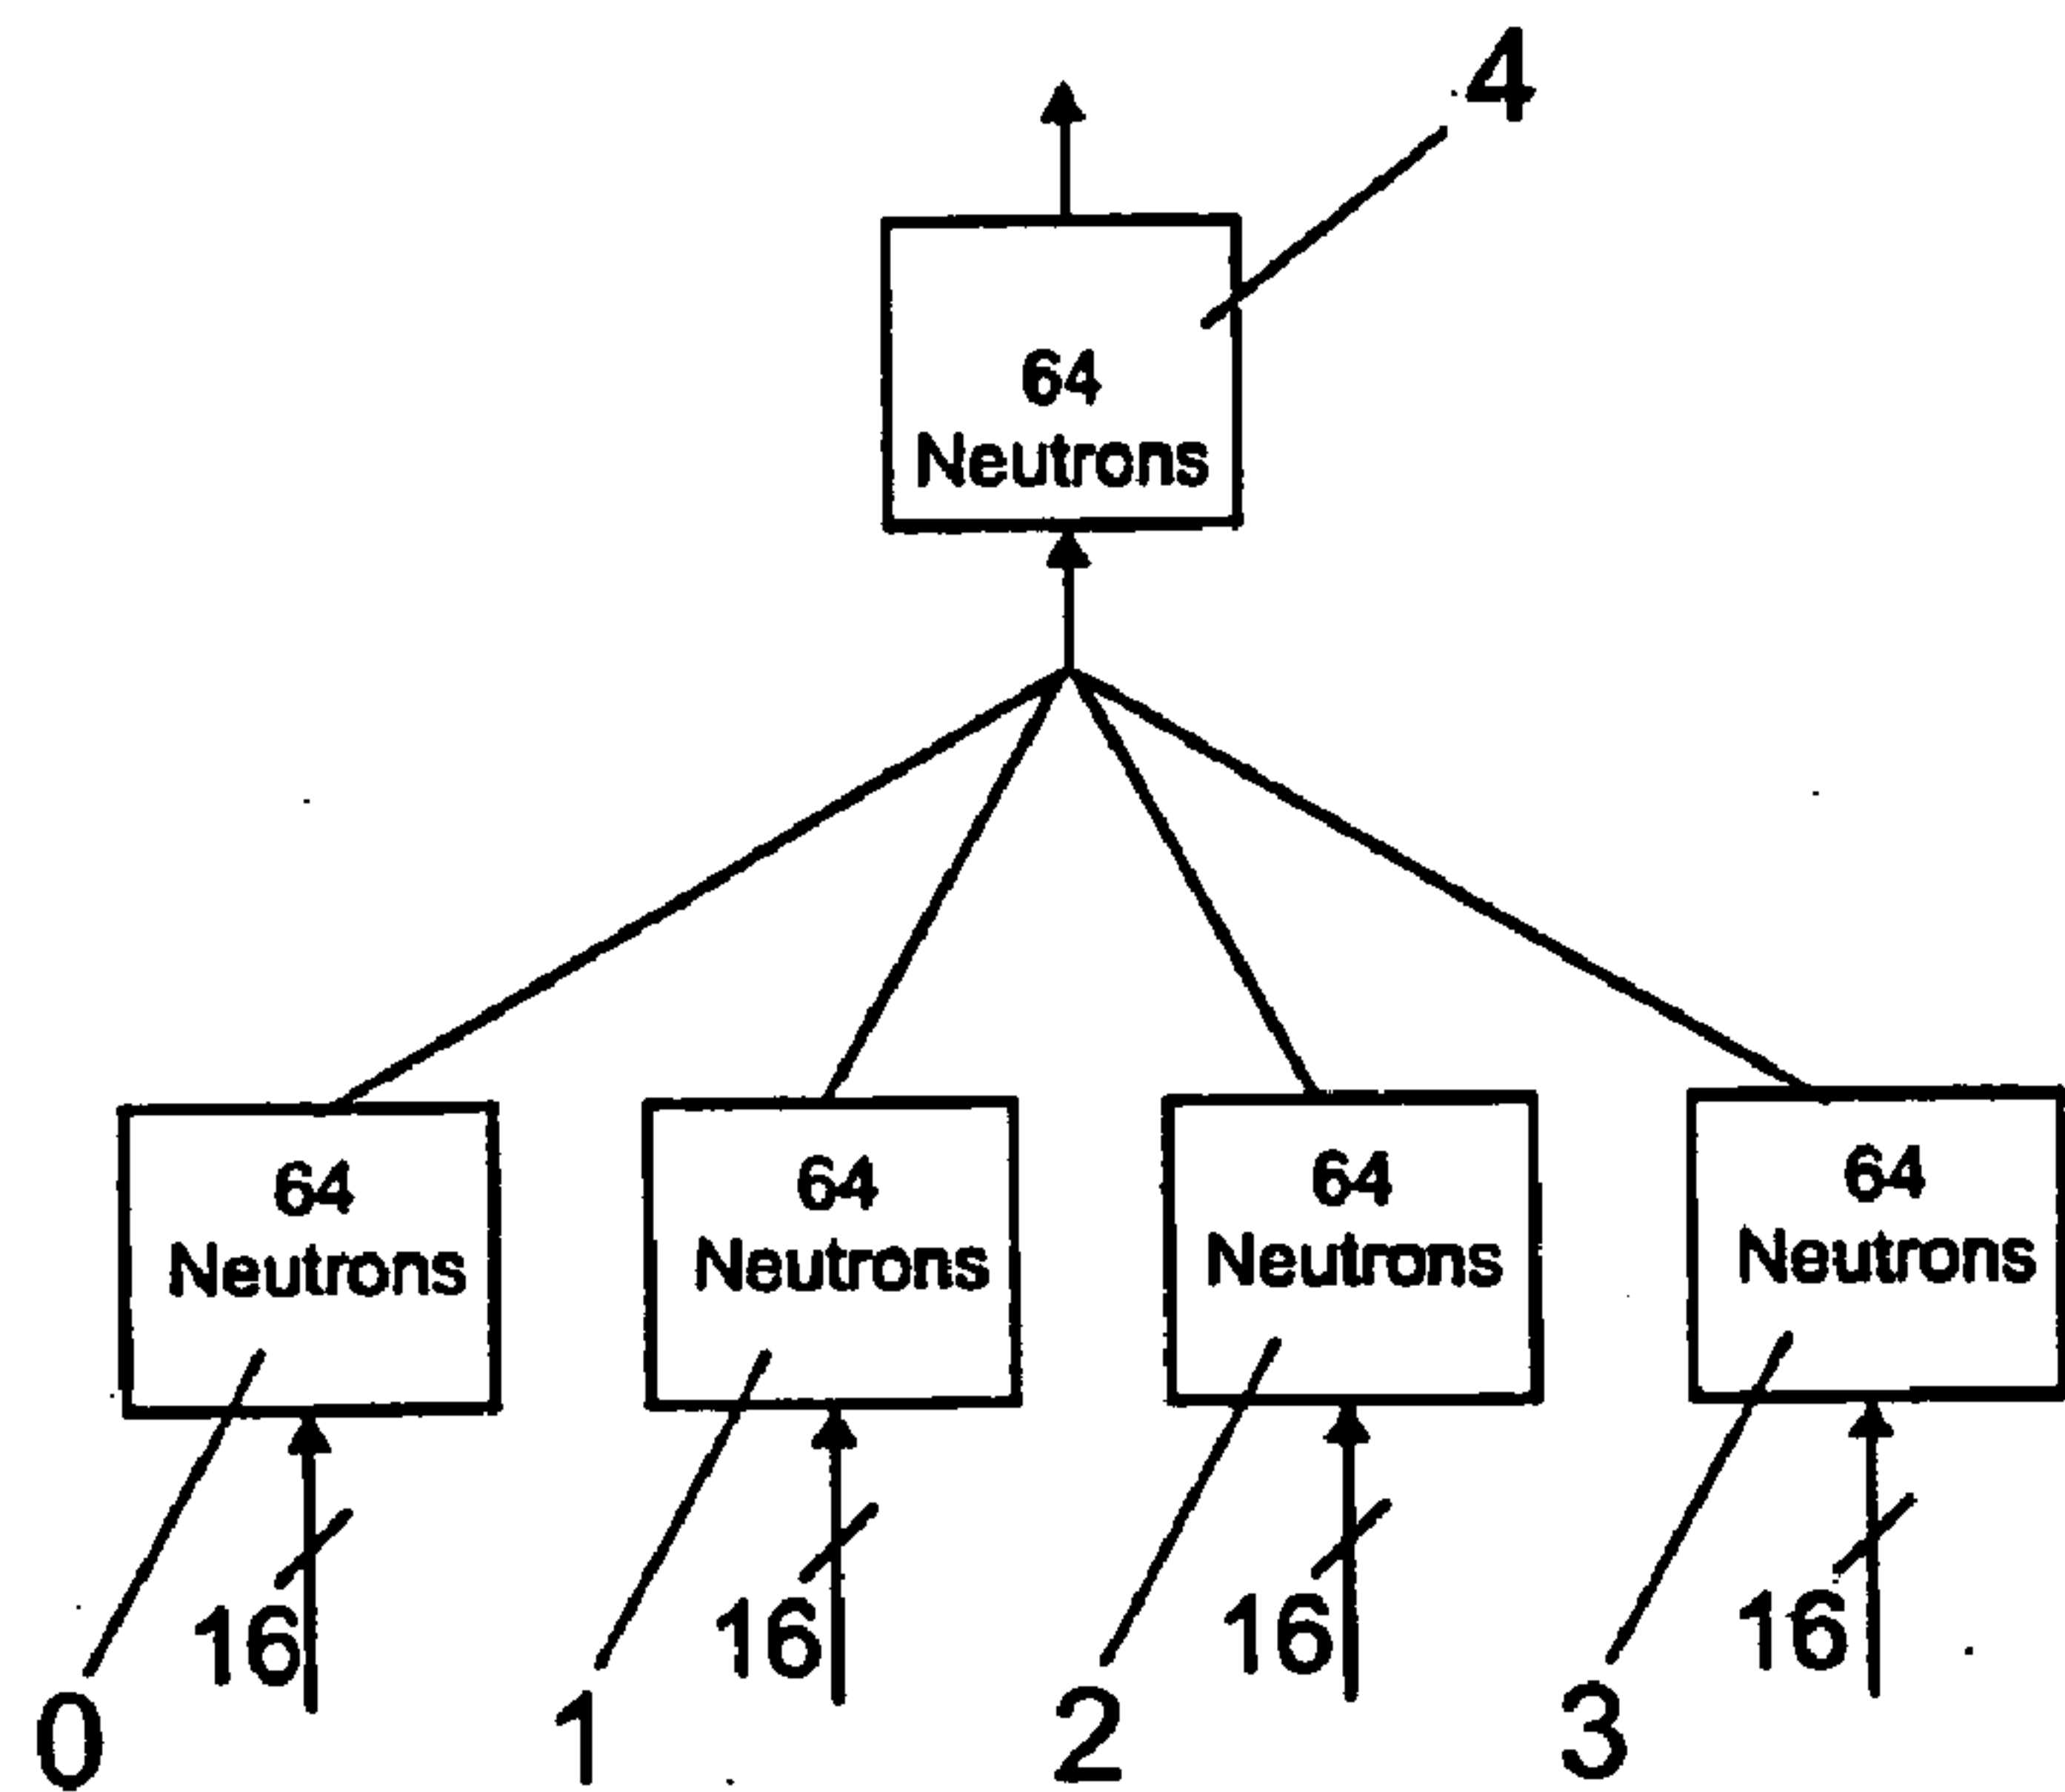


Fig. 28

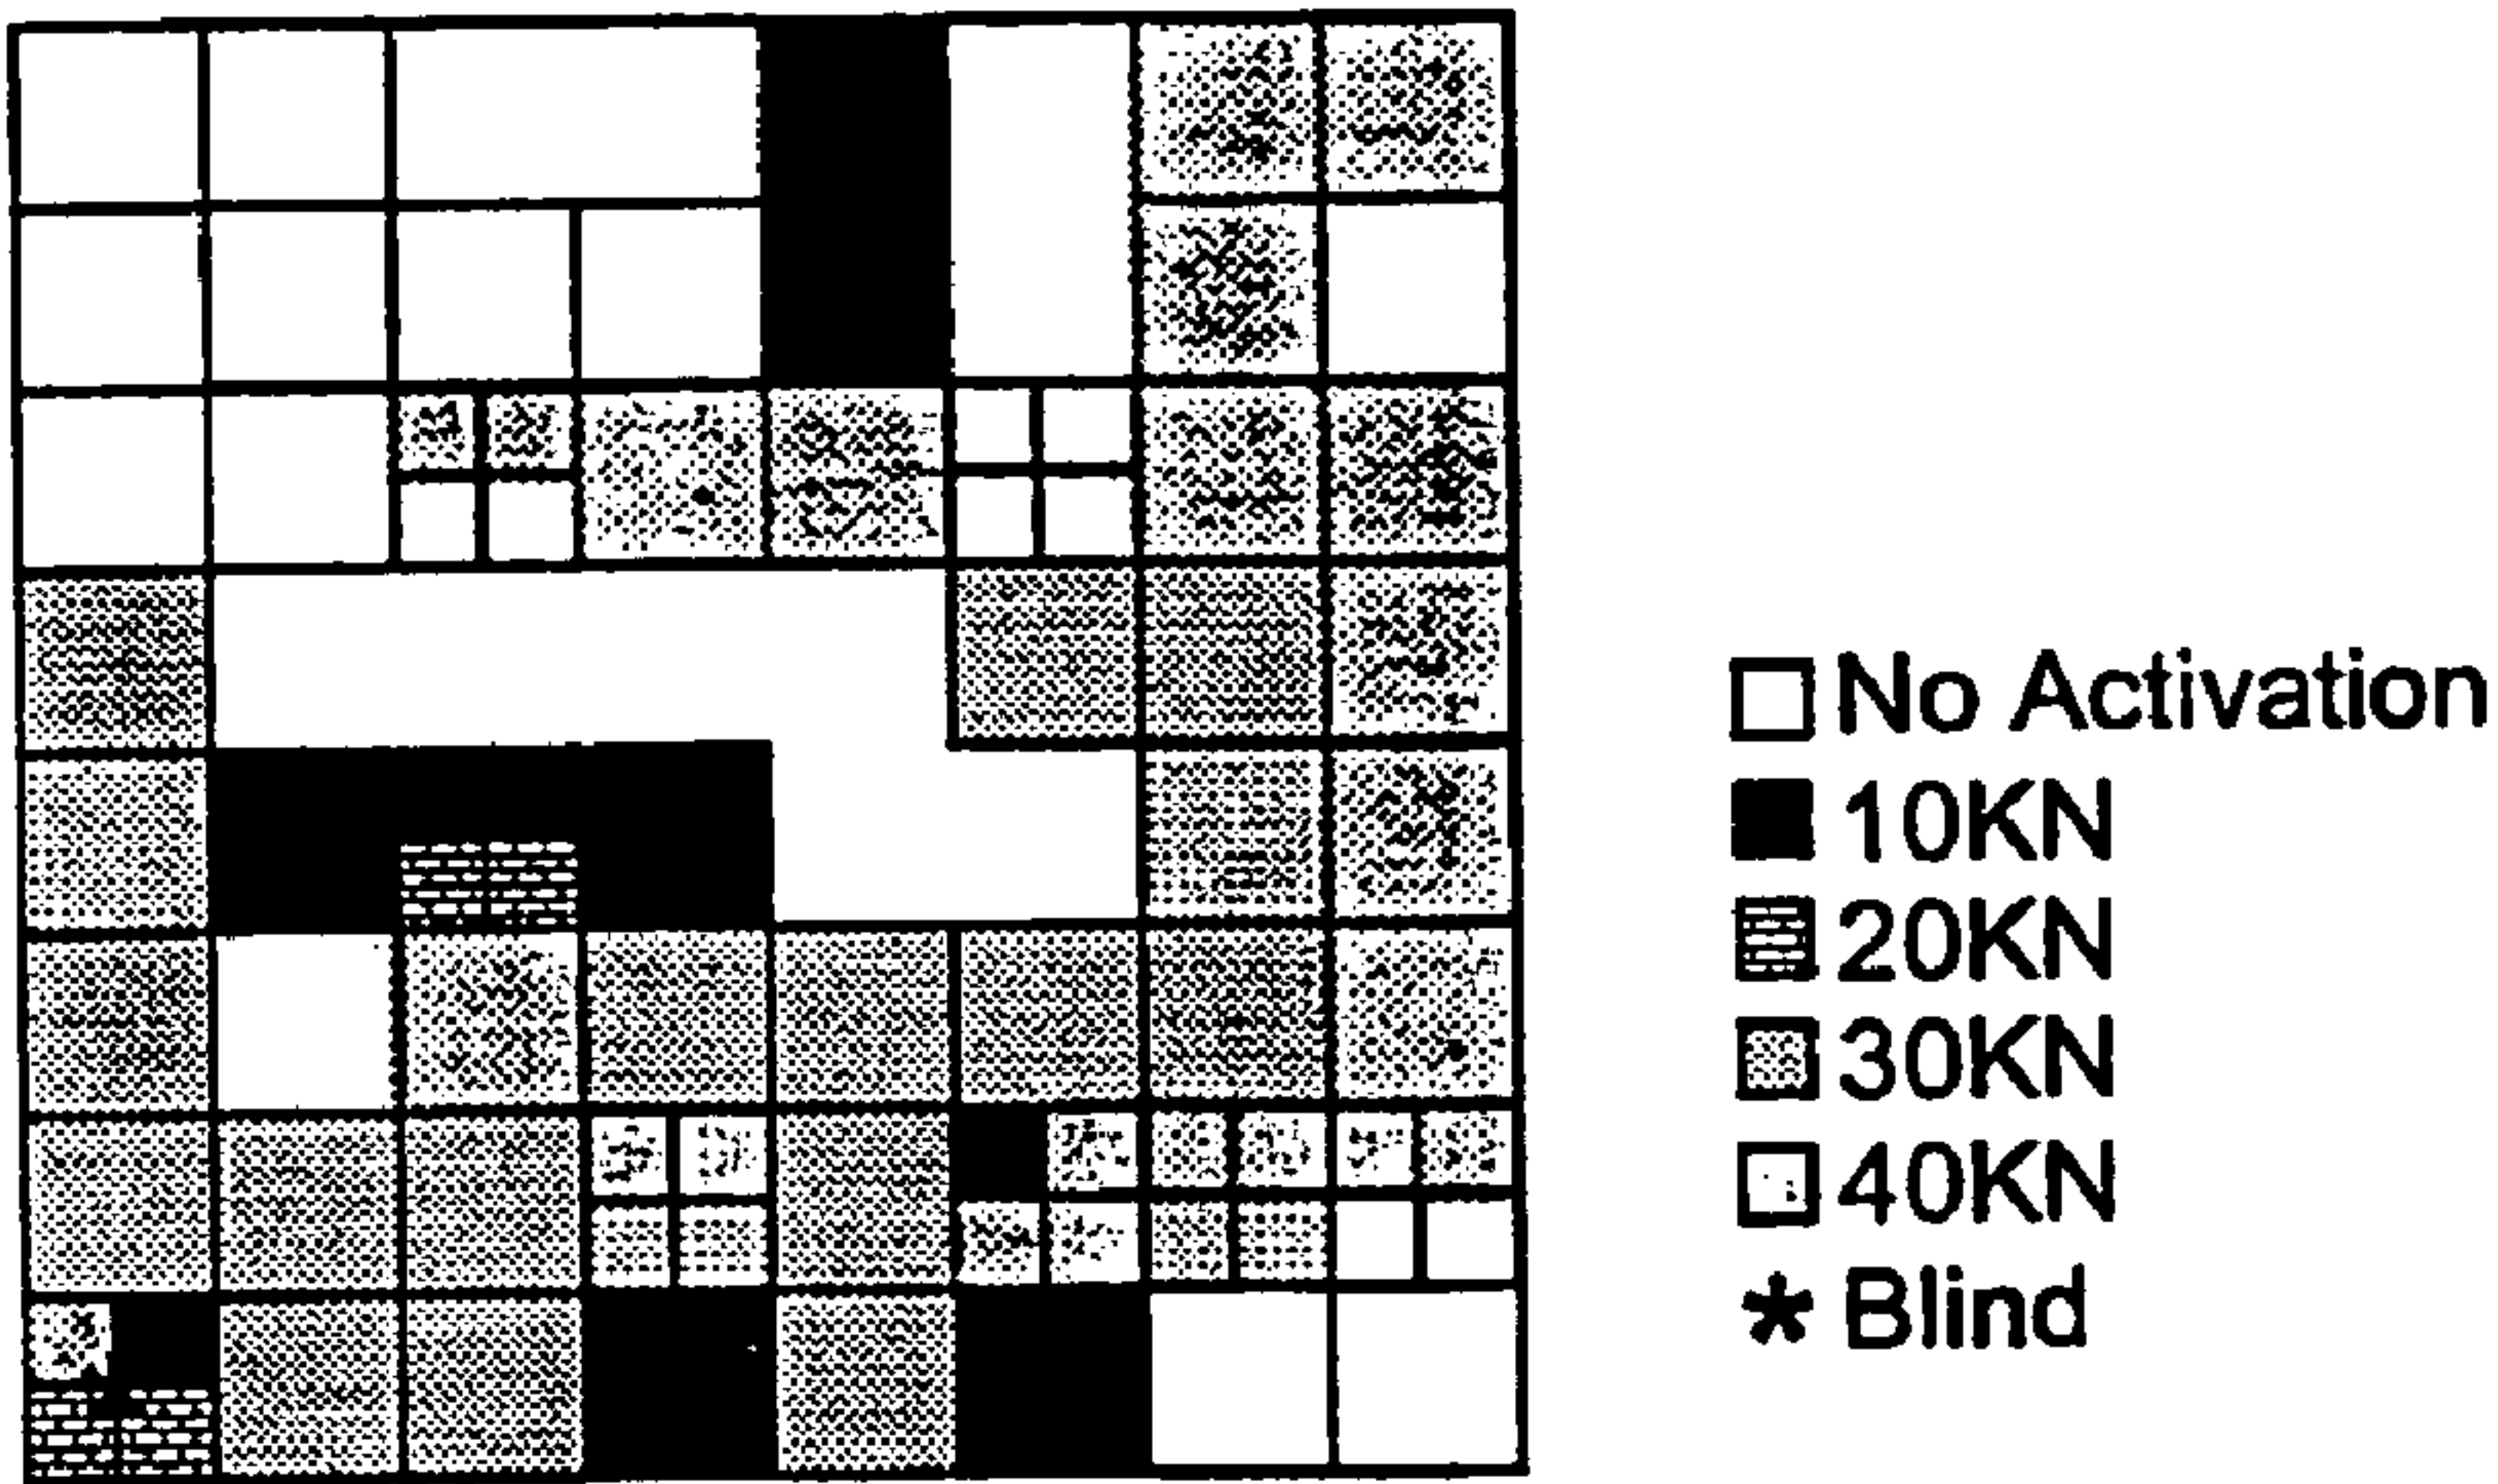


Fig. 29

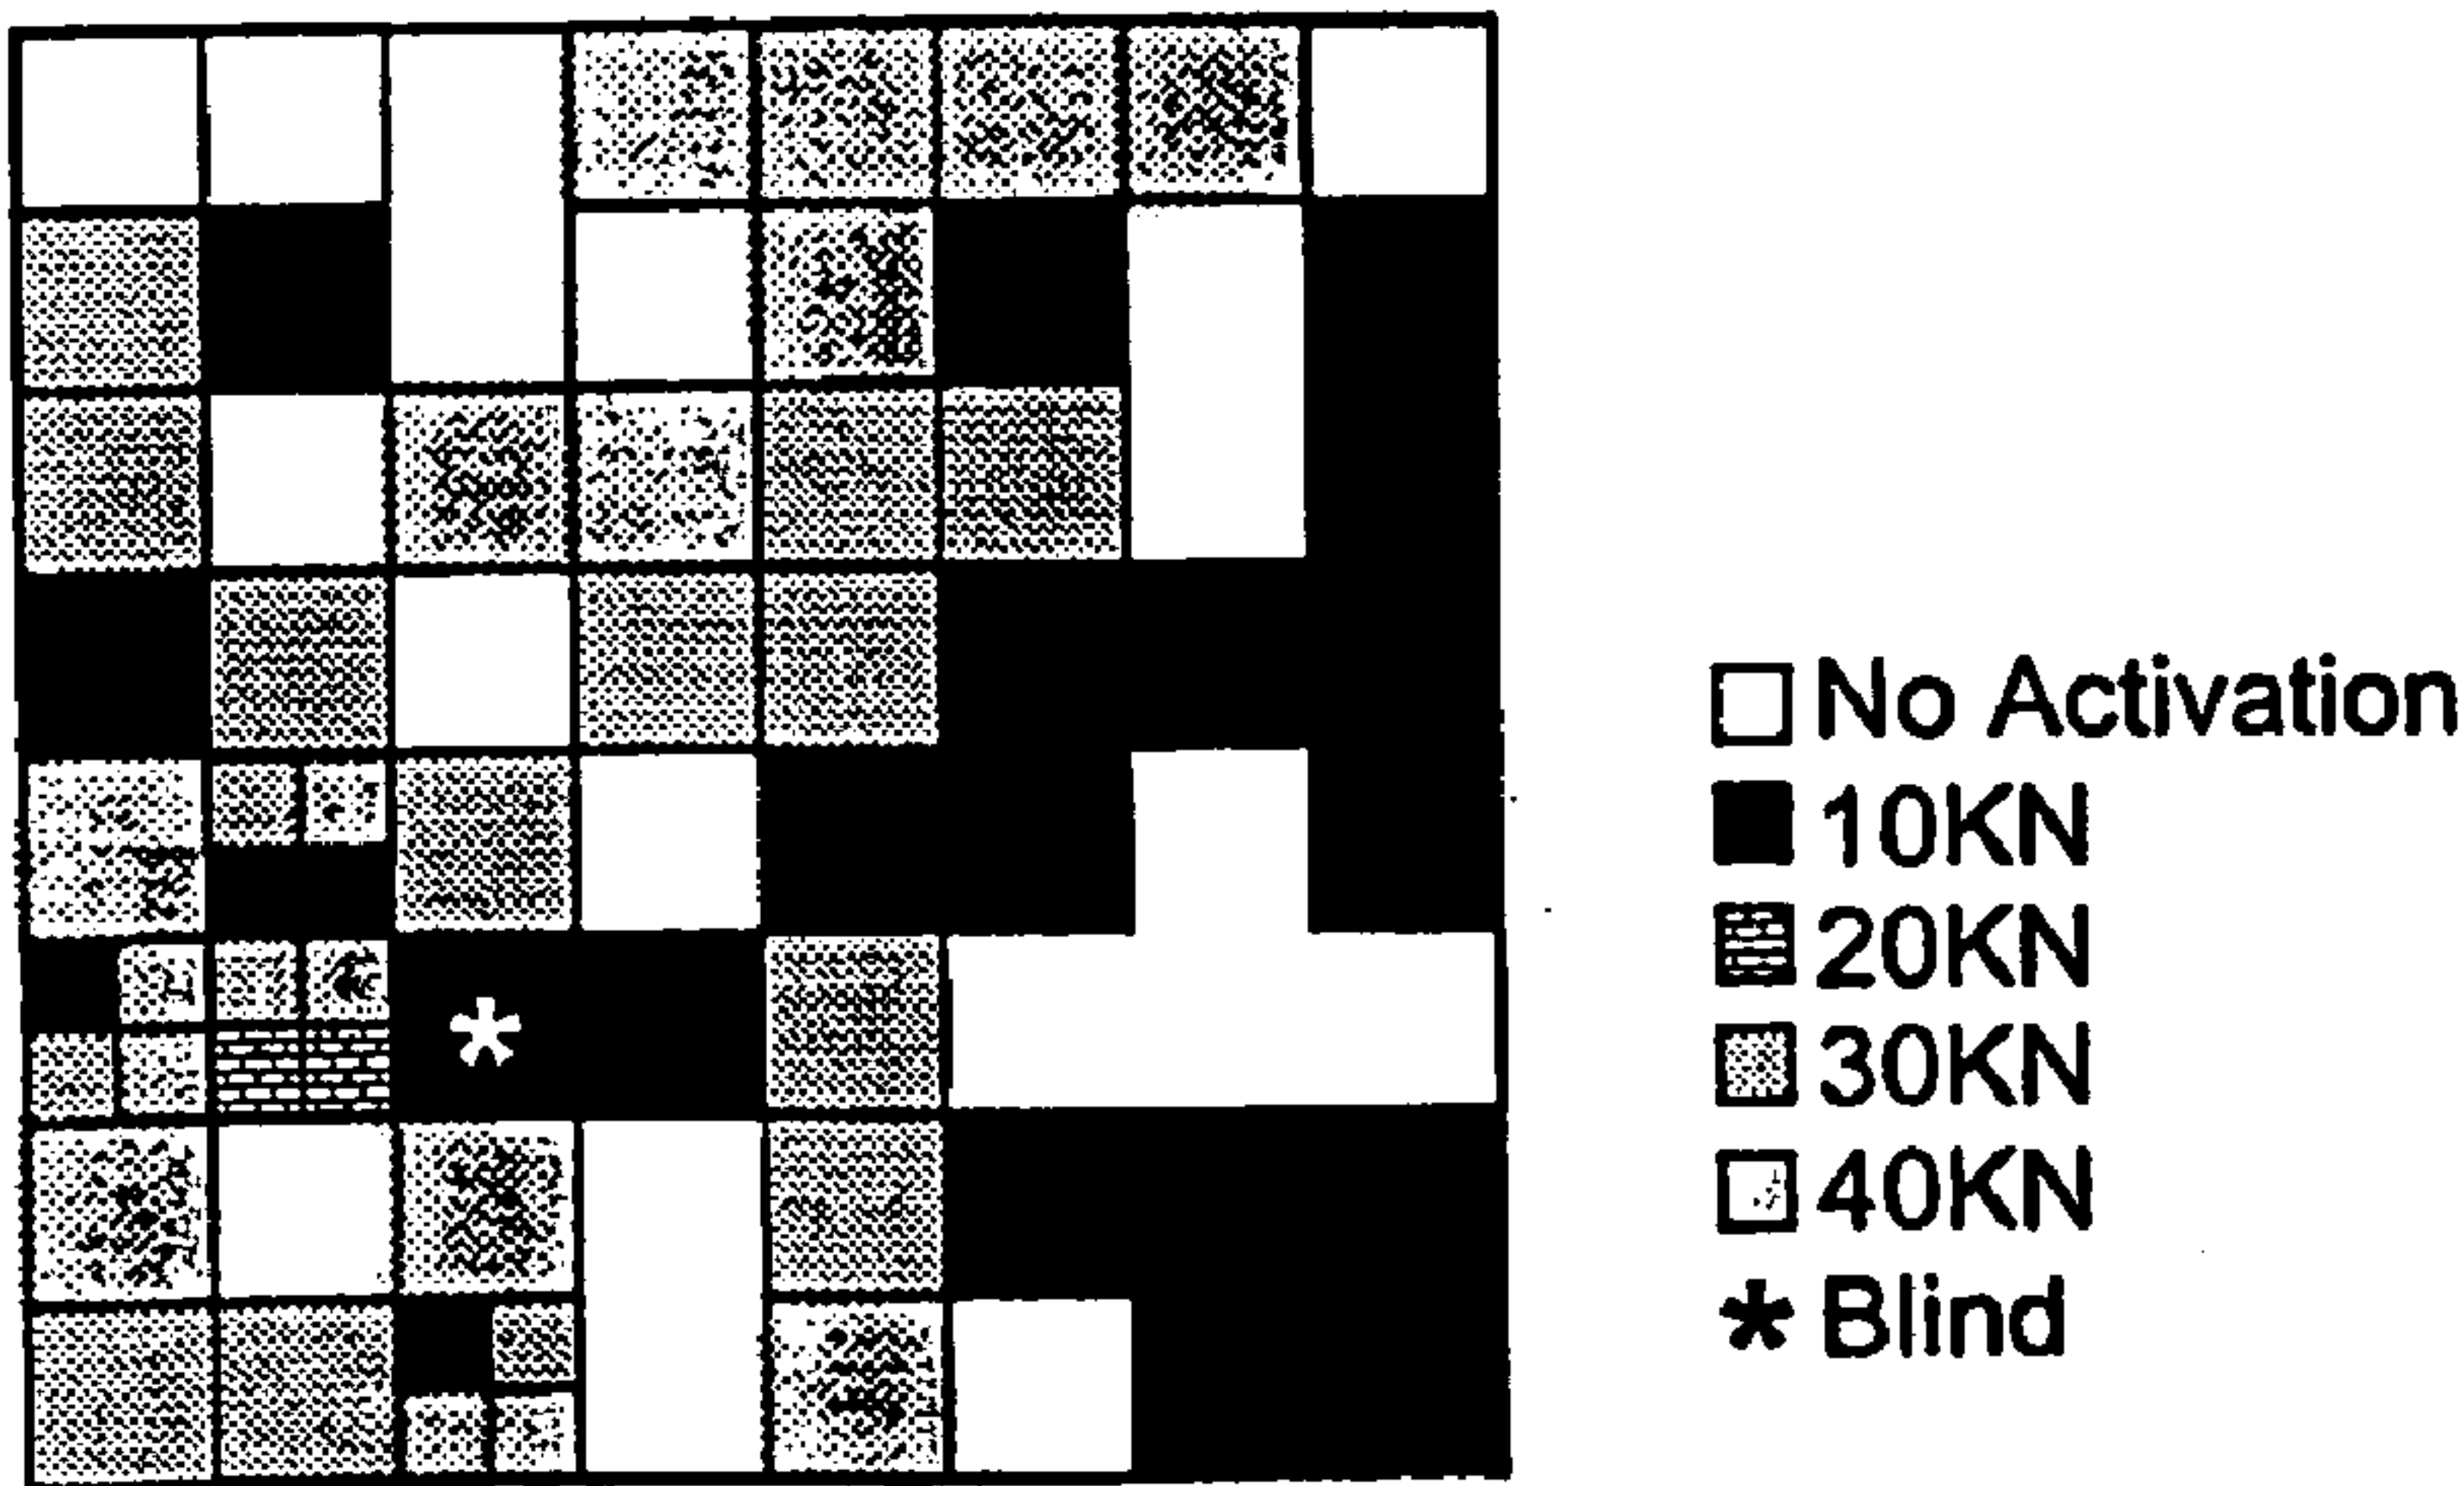




Fig. 30




- No Activation
- 10KN
- ▨ 20KN
- ▩ 30KN
- ░ 40KN
- * Blind

Fig. 31



- No Activation
- 10KN
- ▨ 20KN
- ▩ 30KN
- ░ 40KN
- * Blind

Fig. 32



- No Activation
- 10KN
- ▨ 20KN
- ▩ 30KN
- ░ 40KN
- * Blind

NEURAL PROCESSING ELEMENT FOR USE IN A NEURAL NETWORK

[0001] The present invention relates to neural networks and more particularly, but not exclusively, to an apparatus for creating, and a method of training, a neural network.

[0002] Artificial Neural Networks (ANNs) are parallel information processing systems inspired by what is known about the brain and the way it functions. They offer a computing mechanism that differs significantly from the conventional serial computer systems, not simply because they process information in a parallel manner but because they do not require explicit information about the problems they are required to tackle; instead they learn by example. However, rather than being designed and built as computing platforms, they are predominantly simulated on conventional serial computing systems in software. For small networks this approach is generally sufficient, especially when considering the improvement in processing speed that has been achieved in recent years. However, when real-time systems and large networks are required, the computational burden often requires other approaches.

[0003] The basic neuron does very little computation on its own but when large numbers of neurons are used, the total computation is often such that even the fastest of serial computers is unable to train a network in a reasonable time scale. The problem is exacerbated because, the larger the network, the more training steps are required and, consequently, the amount of computation required increases exponentially with increasing network size. There is also the added problem of inter-neuron communication, which also increases with increasing network size and must be taken into account when attempting to implement networks on parallel systems, because this communication can become a bottleneck, preventing substantial speedups for parallel implementations.

[0004] When considering parallel implementation of ANNs, it is important to consider how the system is to be parallelised. This is dependent not only on the underlying architecture/technology but also the algorithm and sometimes on the intended application itself. However, there is often more than one approach for any particular architecture and an understanding of the consequences of partitioning strategies is of great value. When using multi-processor systems, there are two basic approaches to parallelising the Self-Organising Map (SOM) algorithm; either the functionality of the network can be partitioned such that one processor may perform only one aspect of the functionality of a neuron but performs this function for a large number of neurons, or the network can be partitioned so that a set of neurons (a set typically consists of one or more neurons) is implemented on each processor in the system.

[0005] Partitioning functionality of the network is an approach that has been used with transputer systems and, normally results in an architecture known as a systolic array. The basic principle of the systolic array is that the traditional single processing element is replaced by an array of processing elements with inputs and outputs only occurring at each end of the array. The processing that would traditionally be carried out by a single processor is then divided amongst the processor array. Normally, each processor would perform some of the functionality of the network and that function would only be performed by that processor.

The array then acts as a pipeline of processors, with data flowing in at one end and results flowing out of the other. Unfortunately, this approach is generally only appropriate for moderately sized networks because the inter-processor communication overheads become unmanageable very quickly and adding more processors does little or nothing to alleviate the problem.

[0006] When partitioning the SOM wherein one or more neurons are implemented on an individual processor, the communication overhead is lessened when compared to approaches that partition functionality but can still become a bottleneck as network size increases. Coarse grain parallelism is the term generally associated with a number of neurons implemented on each processor whereas fine grain parallelism is the term used when only a single neuron is implemented on individual processors. The communication overhead tends to become more prominent as the number of neurons per processor is reduced because traditional processors are implemented on separate devices and communication between devices has much greater overheads than communication amongst neurons on the same device. Fine grain parallelism normally results in a Single Instruction stream Multiple Data stream (SIMD) system and is suited to massively parallel architectures such as the Connection Machine.

[0007] If the implementation medium is to be in hardware such as very large scale integration (VLSI) or similar, then it may be possible to increase the level of parallelism to the extent of implementing each weight in parallel. However, this approach does little to improve overall parallelism of the system because only part of the functionality is performed at the weight level and consequently, such an approach does not lead to the most effective use of resources. The approach adopted is fine grain parallelism with a single processing element performing the functionality of a single neuron. To overcome some of the inter-processor communication problems it is suggested that several processors be implemented on a single device with strong short range communications.

Neural Network Implementations

[0008] In an attempt to overcome the limitations of general purpose parallel computing platforms some researchers attempted to develop specialised neural network computers. Such approaches attempt to develop architectures best suited to neural networks but are normally based on the traditional parallel architectures listed above. Modifications to these basic architectural approaches have often been used in an attempt to overcome some of the traditional problems such as inter-processor communication. Others have attempted to modify existing parallel systems such as the Connection Machine to improve their usefulness as neurocomputing architectures. Some have even considered reconfigurable neurocomputer systems based on Field Programmable Gate Array Technology (FPGA) but most neurocomputer systems, while useful for investigating the possibilities of ANNs, are normally too large and expensive to be used for many applications.

[0009] Driven mainly by the application domain researchers undertook to investigate direct hardware implementation of ANNs, and as biological neural systems appear to be analogue, there was a bias towards analogue implementation. Indeed, analogue implementation of ANNs appears to be beneficial in some ways, e.g. very little hardware is

required for the memory elements of such a system. However, there are also many problems with analogue implementation of ANNs because the fundamental building block of such systems is the capacitor. Due to the shortcomings of the capacitor, such as its tendency to suffer from leakage, a variety of schemes were developed to overcome these weaknesses.

[0010] Macq et al proposed an analogue approach to implementation of the SOM based on the use of currents to represent weight values. Such an approach may provide a mechanism for generating high density integration due to the small number of transistors required for each neuron, but this approach uses analogue synaptic weights based on current copiers, the principle component of which is the capacitor, which is prone to leakage. These leakage currents continuously modify the value stored by the capacitor thereby necessitating some form of refreshment to maintain reasonable precision of weight values. The main cause of this leakage is the reverse biased junction. Their proposed method of refreshment uses a converter to periodically refresh each synaptic weight. This is achieved by reading the current memorised by each cell using successive approximation and then writing back to the cell the next upper reference current. It is claimed that this approach allows for on chip learning. However, for the gain factor to reduce with time, as prescribed by Kohonen, adjustments need to be made to the reset signal, and for the neighbourhood to reduce with time the period of one of the timing circuit clocks must be adjusted. The impression given is that these changes would require manual intervention. The leakage current of capacitors also appears to be the main factor that would restrict the maximum number of memory cells in this design.

[0011] A charge based approach to implementation was suggested in "A Charge-Based On-Chip Adaptation Kohonen Neural Network" which claims that such an approach would lead to low power dissipation and compact device configurations. The approach uses switched capacitor circuits to store the weights and the adaptive weight synapses used utilises parasitic capacitances between two adjacent gates of the switched capacitor circuit to determine the learning rate. This will give a fixed learning rate, which will be different for each device manufactured due to the difficulties in manufacturing such components to exactly the same parameters from device to device. Weight integrity is also a potential problem area because, as with most analogue implementations of neural networks, weight values are stored by capacitors which have difficulty maintaining the charge held, and consequently the weight value. The authors of this paper attempt to address this issue but, for weights not being updated during a cycle, they simply regarded it as a forget effect. Unfortunately, as the number of neurons on the device increases, so too does the common node parasitic capacitance. This will require the size of the storage electrode of each neuron to be increased as network size increases to compensate.

[0012] Perhaps the most successful analogue implementations are those which utilise a pulse stream approach. It has long been known that biological neural systems use pulses to communicate between cells and simple oscillating circuits can be implemented in VLSI relatively easily. Unfortunately, the problem of analogue memory still overshadows such approaches. The main advantage of pulse stream

approaches is that hardware requirements for the arithmetic units are very low compared to the equivalent digital implementation; in particular multipliers which can be implemented in an analogue fashion using only three transistors require many gates for digital systems.

[0013] The problems of implementing digital multipliers and storing weight values provide two reasons that most digital implementations of the SOM have been restricted to small network sizes and are often only coprocessors rather than fully parallel implementations. The other main factor that has made a significant contribution to limiting network size is the inter-neuron communication overhead which increases exponentially with network size. Consequently, most fully digital implementations of the SOM require some modification to Kohonen's original algorithm, e.g. Ienne et al suggest two alternative modifications to the SOM algorithm for digital implementation. Van den Bout et al also propose an all digital implementation of the SOM and investigate a rapid prototyping approach towards neural network hardware development. This is facilitated by the use of Xilinx field programmable gate arrays (FPGAs) which provide a flexible platform for such endeavours and speed up construction time compared to VLSI development. Their approach uses stochastic signals to allow pseudo-analogue computation to be carried out using space efficient digital logic. A Markovian learning algorithm is used to simplify that suggested by Kohonen and the Manhattan distance metric is used in place of Euclidean distance to simplify distance calculations. Their approach towards the implementation of the SOM is later reiterated when they describe their VLSI implementation, TInMann.

[0014] Saarinen et al propose a fully digital approach to the implementation of Kohonen's SOM in order to create a neural coprocessor for PC based systems. Their approach uses three Xilinx XC3090 FPGAs to create 16 processing elements, and RAM to store both weight and input vector values. The host computer initialises the random weight values, loads up the input vector values and sets the network parameters (i.e. network size, number of inputs, gain factor and number of training steps). After the host computer has set these parameters the coprocessor system then trains the network according to the pre-specified parameters until training is complete. The architecture of the system consists of three main elements; a distance and update unit (DUU), a distance comparator unit (DCU) and an address control unit (ACU), each implemented on a separate FPGA which is clearly a partitioning of the network functionality and is not likely to be scaleable due to the communication overheads. In addition, this implementation does not implement the standard SOM but, a rather limited, one dimensional version.

[0015] While more obvious than many of the digital implementation approaches used, that of Saarinen is rather typical in that it partitions functionality. Most digital implementations appear to do the same, but they maintain the whole system on a single device. The rationale behind this is that when using digital multipliers, vast resources are normally required to implement them, so it is often more effective to have a limited number but to make them fast. To avoid using excessive resources for the Modular Map implementation, very limited reduced instruction set computers (RISC) processors are suggested that use an alternative approach to multiplication which will only require a fraction

of the resources needed to implement a traditional digital multiplier. In addition, while minor modifications to Kohonen's algorithm are made, its basic operation and two dimensional nature are maintained.

[0016] The paper by Ruping et al presented simultaneously with the paper by Lightowler et al presents a fully digital hardware implementation of the SOM which incorporates some of the same ideas as does the Modular Map design. To facilitate hardware implementation Ruping et al also use Manhattan distance instead of Euclidean distance and the gain factor is restricted to negative powers of two. A system comprising 16 devices is outlined and performance information is presented in terms of the operating speed of the system etc. Each of their devices implements 25 neurons as separate processing elements and allows for network size to be increased by using several devices. However, these devices only contain neurons; there is no local control for the neurons on a device. An external controller is required to interface with these devices and control the actions of their constituent neurons. Consequently, these devices are not autonomous as are Modular Maps and only lateral expansion which creates a Single Instruction stream Multiple Data stream (SIMD) architecture has been considered as an approach towards creating larger network sizes.

[0017] There have also been some commercial hardware implementations of ANNs, the number of which has been steadily growing over the last few years. They generally offer a speedup of around an order of magnitude compared to implementation on a PC alone but are predominantly coprocessors rather than stand alone systems and are not normally scaleable. However, while some of these implementations are only able to implement a single ANN paradigm, most use digital signal processing (DSP) chips, transputers or standard microprocessors, thereby allowing the system to be programmable to some extent and implement a range of standard ANNs.

[0018] The commercially available approach to implementation, (i.e. accelerator cards) offers the slowest speedup of the main implementation approaches but can still offer a significant speedup compared to simulation on standard PC systems and the growing number available on the market suggests that they are useful for a range of applications. General purpose multiprocessor systems offer a further speedup but large scale systems normally have significant communication overheads. Some researchers have attempted to modify standard multiprocessor architectures to improve their application to ANNs and have increased achievable speedup by doing so but while these systems have been useful in ANN research, they are not fully scaleable and require significant financial outlay. The greatest speedups for ANN implementations have been achieved by dedicated neural network chips but the problem again has been that these systems are limited to relatively small scale systems. As an approach towards developing scaleable neural network systems, there have been some attempts at developing modular systems.

Modular System

[0019] There is considerable evidence to suggest that biological neural systems have a modular organisation at various levels. At a macroscopic level, for example, it has been found that some people have no connection between the left and right hemispheres of the brain, which does bring

with it certain problems, but they are still able to function in a near to normal way, which shows that each hemisphere is able to function independently. However, it has also been noted that, while each hemisphere is almost identical physiologically, they specialise in functionality. When one begins to look closer at the cerebral hemisphere one finds that different functionality is found at different regions, even though these regions show a modular organisation and are made up of geometrically defined repetitive units. Research by Murre and Sturdy also supports this view of a modular organisation in their attempt at a quantitative analysis of the brain's connectivity. It is of interest that this modularity is also seen in relation to the topological maps formed in the neo-cortex, e.g. somatosensory maps for different parts of the body are found at different parts of the cerebral cortex and similar maps for other senses such as sound (tonotopic maps) are found in different regions again. Such evidence suggests that while the concept of topological maps which form the basis for Kohonen's self organising map is valid, it also suggests that the brain contains many of these maps. Consequently, it is reasonable to suggest that when attempting to develop scaleable, and particularly when trying to develop large scale implementations of the SOM, that a modular approach should be considered.

[0020] Researchers such as Happel and Murre have approached neural network design as an evolutionary process using genetic algorithms to determine network architectures. Their investigations into the design of modular neural networks using the CALM module are intended as a study to assist with understanding of the relationship between structure and functionality in the brain but they present some findings that may also assist with the development of information processing systems. They found that the best performing network architectures derived with their approach reproduced characteristics of the vision system with the is organisation of coarse and fine processing of stimuli in different pathways. They also present a range of evidence that supports the belief that the brain is highly organised and modular in its architecture.

[0021] The basic premise on which modular neural network systems are developed is that the computation performed by the network is decomposed into two or more separate modules which operate as individual entities. Not only can such approaches improve scaleability but considerable savings can be made on the learning times required for large networks, which are often rather slow. In addition, the generalisation abilities of large networks are often poor, whereas systems composed of several modules do not appear to suffer from this drawback. Research carried out by Jacobs et al using modules composed of Multi Layer Perceptrons (MLPs) used competition to split the input space into overlapping regions. Their work found that the modular approach had much improved training times compared to single large networks and gave better performance, especially where there were discontinuities within classes in the original input space. They also found, when building hierarchies of such systems, an architecture they refer to as a hierarchical mixture of experts, that the results yielded a probabilistic approach to decision tree modelling. Others, such as Hansen and Salamon, have considered ensembles of neural networks as a means of improving classification. Essentially the ensemble approach involves training several networks on the same task to achieve a more reliable output.

[0022] A modular approach to implementation of the SOM is a valid alternative to the more traditional approaches which attempt to create single networks. Other authors such as Helge Ritter have also presented research supporting a modular approach for the SOM. There also appears to be a sound basis for modularity in biological systems and, while no attempt is being made to replicate biological systems, they are nevertheless the initial inspiration for artificial neural networks. It is also pertinent to consider that, while Man has only been attempting to develop computing systems for a matter of centuries, natural evolution had produced a range of biological computers long before Man was on this earth. Even with the latest of modern technology, Man is unable to create computers that surpass the computing abilities of biological systems, so it is suggested that Man should continue to learn from nature.

[0023] According to a first aspect of the present invention, there is provided a neuron for use in a neural network, the neuron comprising

- [0024] an arithmetic logic unit;
- [0025] a shifter mechanism;
- [0026] a set of registers;
- [0027] an input port;
- [0028] an output port; and
- [0029] control logic.

[0030] According to a second aspect of the present invention, there is provided a module controller for controlling the operation of at least one neuron, the controller comprising

- [0031] an input port;
- [0032] an output port;
- [0033] a programmable read-only memory;
- [0034] an address map;
- [0035] an input buffer; and
- [0036] at least one handshake mechanism.

[0037] According to a third aspect of the present invention, there is provided a neuron module, the module comprising

- [0038] at least one neuron; and
- [0039] at least one module controller.

[0040] Preferably, the at least one neuron and the at least one module controller are implemented on one device. The device is typically a field programmable gate array (FPGA) device. Alternatively, the device may be a full-custom very large scale integration (VLSI) device, a semi-custom VLSI or an application specific integrated circuit (ASIC).

[0041] According to a fourth aspect of the present invention there is provided a neural network, the network comprising

- [0042] at least two neuron modules coupled together.

[0043] Typically, the neuron modules are coupled in a lateral expansion mode. Alternatively, the neuron modules may be coupled in a hierarchical mode. Optionally, the neuron modules may be coupled in a combination of lateral expansion modes and hierarchical modes.

[0044] In lateral expansion mode, the at least two neuron modules are typically connected on a single plane. Data is preferably input to the modules in the network only once. Thus, the modules forming the network are synchronised to facilitate this. The modules are preferably synchronised using a two-line handshake mechanism. The two-line mechanism typically has two states. The two states typically comprise a wait state and a data ready state. The wait state typically occurs where a sender and/or a receiver is not ready for the transfer of data from the sender to the receiver or vice versa. The data ready state typically occurs when both the sender and receiver are ready for data transfer. Data transfer follows immediately the data ready state occurs.

[0045] The neuron modules typically comprise at least one neuron, and at least one module controller.

[0046] Typically, the number of neurons in a module is a power of two. The number of neurons in a module is preferably 256. Any number of neurons may be used in a module, but the number of neurons is preferably a power of two.

[0047] A neuron typically comprises an arithmetic logic unit, a shifter mechanism, a set of registers, an input port, an output port, and control logic.

[0048] The arithmetic logic unit (ALU) typically comprises an adder/subtractor unit. The ALU is typically at least a 4-bit adder/subtractor unit, and preferably a 12-bit adder/subtractor unit. The adder/subtractor unit typically includes a carry lookahead adder (CLA).

[0049] The ALU typically includes at least two flags. A zero flag is typically set when the result of an arithmetic operation is zero. A negative flag is typically set when the result of an arithmetic operation is negative.

[0050] The ALU typically further includes at least two registers. A first register is typically located at one of the inputs to the ALU. A second register is typically located at the output from the ALU. The second register is typically used to store data until it is ready to be transferred eg stored.

[0051] The shifter mechanism typically comprises an arithmetic shifter. The arithmetic shifter is typically implemented using flip-flops. The shifter mechanism is preferably located in a data stream between the output of the ALU and the second register of the ALU. This location increases the flexibility of the neuron and increases the simplicity of the design.

[0052] The control logic typically comprises a reduced instruction set computer (RISC). The instruction set typically comprises thirteen different instructions.

[0053] The module controller typically comprises an input port, an output port, a programmable read-only memory, an address map, an input buffer, and at least one handshake mechanism.

[0054] The programmable read-only memory (PROM) typically contains the instructions for the controller and/or the subroutines for the at least one neuron.

[0055] The address map typically allows for conversion between a real address and a virtual address of the at least one neuron. The real address is typically the address of a neuron on the device. The virtual address is typically the address of the neuron within the network. The virtual

address is typically two 8-bit values corresponding to X and Y co-ordinates of the neuron on the single plane.

[0056] The at least one handshake mechanism typically includes a synchronisation handshake mechanism for synchronising data transfer between a sender and a receiver module. The synchronisation handshake mechanism typically comprises a three-line mechanism. The three-line mechanism typically has three states. The three states typically comprise a wait state, a no device state and a data ready state. The wait state typically occurs where a sender and/or a receiver is not ready for the transfer of data from the sender to the receiver or vice versa. The no device state is typically used where inputs are not present. Thus, reduced input vector sizes may be used. The no device state may also be used to prevent the controller from malfunctioning when an input stream(s) is temporarily lost or stopped. The data ready state typically occurs when both the sender and receiver are ready for data transfer. Data transfer follows immediately when the data ready state occurs. The three-line mechanism typically comprises two outputs from the receiver and one output from the sender. The advantage of the three-line mechanism is that no other device is required to facilitate data transmission between the sender and receiver or vice versa. Thus, the transmission of data is directly from point to point.

[0057] According to a fifth aspect of the present invention there is provided a method of training a neural network, the method comprising the steps of

- [0058] providing a network of neurons;
- [0059] reading an input vector applied to the input of the neural network;
- [0060] calculating the distance between the input vector and a reference vector for all neurons in the network;
- [0061] finding the active neuron;
- [0062] outputting the location of the active neuron; and
- [0063] updating the reference vectors for all neurons in a neighbourhood around the active neuron.

[0064] A distance metric is typically used to calculate the distance between the input vector and the reference vector. Preferably, the Manhattan distance metric is used. Alternatively, a Euclidean distance metric may be used.

[0065] Calculation of the Manhattan distance preferably uses a gain factor. The value of the gain factor is preferably restricted to negative powers of two.

[0066] The network of neurons typically comprises a neural network. The neural network typically comprises at least two neuron modules coupled together.

[0067] Typically, the neuron modules are coupled in a lateral expansion mode. Alternatively, the neuron modules may be coupled in a hierarchical mode. Optionally, the neuron modules may be coupled in a combination of lateral expansion modes and hierarchical modes.

[0068] In lateral expansion mode, the at least two neuron modules are typically connected on a single plane. Data is preferably input to the modules in the network only once. Thus, the modules forming the network are synchronised to facilitate this. The modules are preferably synchronised using a two-line handshake mechanism. The two-line

mechanism typically has two states. The two states typically comprise a wait state and a data ready state. The wait state typically occurs where the sender and/or the receiver is not ready for the transfer of data from the sender to the receiver or vice versa. The data ready state typically occurs when both the sender and receiver are ready for data transfer. Data transfer follows immediately the data ready state occurs.

[0069] The neuron modules typically comprise at least one neuron, and at least one module controller.

[0070] Preferably, the at least one neuron and the at least one module controller are implemented on one device. The device is typically a field programmable gate array (FPGA) device. Alternatively, the device may be a full-custom very large scale integration (VLSI) device, a semi-custom VLSI or an application specific integrated circuit (ASIC).

[0071] Typically, the number of neurons in a module is a power of two. The number of neurons in a module is preferably 256. Any number of neurons may be used in a module, but the number of neurons is preferably a power of two.

[0072] A neuron typically comprises an arithmetic logic unit, a shifter mechanism, a set of registers, an input port, an output port, and control logic.

[0073] The arithmetic logic unit (ALU) typically comprises an adder/subtractor unit. The ALU is typically at least a 4-bit adder/subtractor unit, and preferably a 12-bit adder/subtractor unit. The adder/subtractor unit typically includes a carry lookahead Adder (CLA).

[0074] The ALU typically includes at least two flags. A zero flag is typically set when the result of an arithmetic operation is zero. A negative flag is typically set when the result of an arithmetic operation is negative.

[0075] The ALU typically further includes at least two registers. A first register is typically located at one of the inputs to the ALU. A second register is typically located at the output from the ALU. The second register is typically used to store data until it is ready to be transferred eg stored.

[0076] The shifter mechanism typically comprises an arithmetic shifter. The arithmetic shifter is typically implemented using flip-flops. The shifter mechanism is preferably located in a data stream between the output of the ALU and the second register of the ALU. This location increases the flexibility of the neuron and increases the simplicity of the design.

[0077] The control logic typically comprises a reduced instruction set computer (RISC). The instruction set typically comprises thirteen different instructions.

[0078] The module controller typically comprises an input port, an output port, a programmable read-only memory, an address map, an input buffer, and at least one handshake mechanism.

[0079] The programmable read-only memory (PROM) typically contains the instructions for the controller and/or the subroutines for the at least one neuron.

[0080] The address map typically allows for conversion between a real address and a virtual address of the at least one neuron. The real address is typically the address of a neuron on the device. The virtual address is typically the

address of the neuron within the network. The virtual address is typically two 8-bit values corresponding to X and Y co-ordinates of the neuron on the single plane.

[0081] The at least one handshake mechanism typically includes a synchronisation handshake mechanism for synchronising data transfer between a sender and receiver module. The synchronisation handshake mechanism typically comprises a three-line mechanism. The three-line mechanism typically has three states. The three states typically comprise a wait state, a no device state and a data ready state. The wait state typically occurs where the sender and/or the receiver is not ready for the transfer of data from the sender to the receiver or vice versa. The no device state is typically used where inputs are not present. Thus, reduced input vector sizes may be used. The no device state may also be used to prevent the controller from malfunctioning when an input stream(s) is temporarily lost or stopped. The data ready state typically occurs when both the sender and receiver are ready for data transfer. Data transfer follows immediately when the data ready state occurs. The three-line mechanism typically comprises two outputs from the receiver and one output from the sender. The advantage of the three-line mechanism is that no other device is required to facilitate data transmission between the sender and receiver or vice versa. Thus, the transmission of data is directly from point to point.

[0082] Embodiments of the present invention shall now be described, with reference to the accompanying drawings in which:

[0083] FIG. 1a is a unit circle for a Euclidean distance metric;

[0084] FIG. 1b is a unit circle for a Manhattan distance metric;

[0085] FIG. 2 is a graph of gain factor against training time;

[0086] FIG. 3 is a diagram showing neighbourhood function;

[0087] FIGS. 4a-c are examples used to illustrate an elastic net principle;

[0088] FIG. 5 is a schematic diagram of a single Modular Map;

[0089] FIG. 6 is a schematic diagram of laterally combined Maps;

[0090] FIG. 7 is a schematic diagram of hierarchically combined Maps;

[0091] FIG. 8 is a scatter graph showing input data supplied to the network of FIG. 7;

[0092] FIG. 9 is a Voronoi diagram of a module in an input layer I of FIG. 7;

[0093] FIG. 10 is a diagram of input layer activation regions for a level 2 module with 8 inputs;

[0094] FIG. 11 is a schematic diagram of a Reduced Instruction Set Computer (RISC) neuron;

[0095] FIG. 12 is a schematic diagram of a module controller system;

[0096] FIG. 13 is a state diagram for a three-line handshake mechanism;

[0097] FIG. 14 is a flowchart showing the main processes involved in training a neural network;

[0098] FIG. 15 is a graph of activations against training steps for a typical neural net;

[0099] FIG. 16 is a graph of training time against network size using 16 and 99 element reference vectors;

[0100] FIG. 17 is a log-linear plot of relative training times for different implementation strategies for a fixed input vector size of 128 elements;

[0101] FIG. 18 is example greyscale representation of the range of images for a single subject used in a human face recognition application;

[0102] FIG. 19a is an example activation pattern created by the same class of data for a modular map shown in FIG. 23;

[0103] FIG. 19b is an example activation pattern created by the same class of data for a 256 neuron self-organising map (SOM);

[0104] FIG. 20 is a schematic diagram of a modular map (configuration 1);

[0105] FIG. 21 is a schematic diagram of a modular map (configuration 2);

[0106] FIG. 22 is a schematic diagram of a modular map (configuration 3);

[0107] FIG. 23 is a schematic diagram of a modular map (configuration 4);

[0108] FIGS. 24a to 24e are average time domain signals for a 10 kN, 20 kN, 30 kN, 40 kN and blind ground anchorage pre-stress level tests, respectively;

[0109] FIGS. 25a to 25e are average power spectrum for the time domain signals in FIGS. 24a to 24e respectively;

[0110] FIG. 26 is an activation map for a SOM trained with the ground anchorage power spectra of FIGS. 25a to 25e;

[0111] FIG. 27 is a schematic diagram of a modular map (configuration 5);

[0112] FIG. 28 is the activation map for module 0 in FIG. 27;

[0113] FIG. 29 is the activation map for module 1 in FIG. 27;

[0114] FIG. 30 is the activation map for module 2 in FIG. 27;

[0115] FIG. 31 is the activation map for module 3 in FIG. 27; and

[0116] FIG. 32 is the activation map for an output module (module 4) in FIG. 27.

[0117] As an approach to overcoming the constraints of unitary artificial neural networks a modular implementation strategy for the Self-organising Map (SOM) can be used. The basic building block of this system is the Modular Map which is itself a parallel implementation of the SOM. Kohonen's original algorithm has been maintained, except-

ing that parameters have been quantised and the Euclidean distance metric used as standard has been replaced by Manhattan distance. Each module contains sufficient neurons to enable it to do useful work as a stand alone system. However, the Modular Map design is such that many modules can be connected together to create a wide variety of configurations and network sizes. This modular approach results in a scaleable system that meets an increased workload with an increase in parallelism and thereby avoids the usually extensive increases in training times associated with unitary implementations.

BACKGROUND

[0118] An important premise on which the Modular Map has been developed is its ability to form topological maps of the input space, a phenomenon which has been likened to the 'neuronal maps' of the brain which are found in regions of the neo-cortex associated with various senses. The formation of such topology preserving maps occurs during the learning process defined for the Self Organising Map (SOM).

[0119] In the Modular Map implementation of the SOM the multidimensional Euclidean input space \mathfrak{R}^n , where \mathfrak{R} covers the range (0, 255) and ($0 < n \leq 16$), is mapped to a two dimensional output space \mathfrak{R}^2 (where the upper limit on \mathfrak{R} is variable between 8 and 255) by way of a non-linear projection of the probability density function. Each neuron in the network has a reference vector $m^i = [\mu_{i1}, \mu_{i2}, \dots, \mu_{in}] \in \mathfrak{R}^n$ where μ_{ij} are scalar weights, i is the neuron index and j the weight index.

[0120] An input vector $x = [\epsilon_1, \epsilon_2, \dots, \epsilon_n] \in \mathfrak{R}^n$ is presented to all neurons in the network where the closest matching reference vector (codebook vector) C is determined, i.e.

$$\sum_{j=0}^n |\xi_j - \mu_{ej}| = \min \left\{ \sum_{j=0}^n |\xi_j - \mu_{ij}| \right\}_{i=1}^k$$

where k =network size.

[0121] The neuron with minimum distance between its codebook vector and the current input (i.e. greatest similarity) becomes the active neuron. A variety of distance metrics can be used as a measure of similarity, the Euclidean distance being the most popular. However, it should be noted that the distance metric being used here is Manhattan distance, known to many as the L_1 metric of the family of Minkowski metrics, i.e. the distance between two points a and b is

$$L_p = (|a-b|^p / a-b/p)^{1/p}$$

[0122] Clearly, Euclidean distance would be the L_2 metric under Minkowski's scheme. An idea of these two distance functions can be gained by plotting the unit circle for both metrics. FIG. 1a shows the unit circle for the Euclidean metric, and FIG. 1b shows the unit circle for the Manhattan metric.

[0123] The Manhattan distance metric is both simple to implement and a reasonable alternative to the Euclidean distance metric which is rather expensive to implement in terms of hardware due to the need to calculate squares of the distances involved.

[0124] After the active neuron has been identified reference vectors are updated to bring them closer to the current input vector. The amount by which codebook vectors are changed is determined by their distance from the input, and the current gain factor $\alpha(t)$. If neurons are within the neighbourhood of the active neuron then their reference vectors are updated, otherwise no changes are made.

$$m_i(t+1) = m_i(t) + \alpha(t) \{x(t) - m_i(t)\} \text{ if } i \in N_c(t)$$

and

$$m_i(t+1) = m_i(t) \text{ if } i \notin N_c(t)$$

where $N_c(t)$ is the current neighbourhood and $t = 0, 1, 2, \dots$

[0125] Both the gain factor and neighbourhood size decrease with time from their original start-up values throughout the training process. Due to implementation considerations these parameters are constrained to a range of discrete values rather than the continuum suggested by Kohonen. However, the algorithms chosen to calculate values for gain and neighbourhood size facilitate convergence of codebook vectors in line with Kohonen's original algorithm.

[0126] The gain factor $\alpha(t)$ being used by the Modular Map is restricted to negative powers of two to simplify implementation. FIG. 2 is a graph of gain factor $\alpha(t)$ against training time when the gain factor $\alpha(t)$ is restricted to negative powers of two. By restricting the gain factor $\alpha(t)$ in this way it is possible to use a bit shift operation for multiplication rather than requiring an additional hardware multiplier which would clearly require more hardware and increase the complexity of the implementation. This approach does not unduly affect the performance of the algorithm and is suitable for simplifying hardware requirements.

[0127] A square, step function neighbourhood, one of several approaches suggested by Kohonen, could be defined by the Manhattan distance metric. This approach to defining the neighbourhood has the effect of rotating the square through 45 degrees and can be used by individual neurons to determine if they are in the current neighbourhood when given the index of the active neuron (see FIG. 3). FIG. 3 is a diagram showing the neighbourhood function when a square, step function neighbourhood is used. When all these parameters are combined to form the Modular Map it has the same characteristics as the self-organising map and gives comparable results when evaluated. The architecture of the Modular Map was also designed to allow for expansion by combining many such modules together to create larger maps while avoiding the usual communications bottleneck and maintaining self-organising map behaviour.

Stand Alone Maps

[0128] If, for visualisation purposes, a simplified case of the Modular Map is considered where only three dimensions are used as inputs, then a single map would be able to represent an input space enclosed by a cube and each dimension would have a possible range of values between 0 and 255. With only the simplest of pre-processing this cube could be placed anywhere in the input space \mathfrak{R}^n where \mathfrak{R} covers the range $(-\infty \text{ to } +\infty)$, and the codebook vector of each neuron within the module would give the position of a point somewhere within this feature space. The implemen-

tation suggested would allow each vector element to hold integer values within the given scale, so there are a finite number of distinct points which can be represented within the cube (i.e. 256^3). Each of the points given by the codebook vectors has an 'elastic' sort of bond between itself and the point denoted by the codebook vectors of neighbouring neurons so as to form an elastic net (FIG. 4).

[0129] FIGS. 4a to 4c shows a series of views of the elastic net when an input is presented to the network. The figures show the point position of reference vectors in three dimensional Euclidean space along with their elastic connections. For simplicity, reference vectors are initially positioned in the plane with $z=0$, the gain factor $\alpha(t)$ is held constant at 0.5 and both orthogonal and plan views are shown. After the input has been presented, the network proceeds to update reference vectors of all neurons in the current neighbourhood. In FIG. 4b, the neighbourhood function has a value of three. In FIG. 4c the same input is presented to the network for a second time and the neighbourhood is reduced to two for this iteration. Note that the reference points around the active neuron become close together as if they were being pulled towards the input by elastic bonds between them.

[0130] Inputs are presented to the network in the form of multi-dimensional vectors denoting positions within the feature space. When an input is received, all neurons in the network calculate the similarity between their codebook vectors and the input using the Manhattan distance metric. The neuron with minimum Manhattan distance between its codebook vector and the current input, (i.e. greatest similarity) becomes the active neuron. The active neuron then proceeds to bring its codebook vector closer to the input, thereby increasing their similarity. The extent of the change applied is proportional to the distance involved, this proportionality being determined by the gain factor $\alpha(t)$, a time dependent parameter.

[0131] However, not only does the active neuron update its codebook vector, so too do all neurons in the current neighbourhood (i.e. neurons topographically close to the active neuron on the surface of the map up to some geometric distance defined by the neighbourhood function) as though points closely connected by the elastic net were being pulled towards the input by the active neuron. This sequence of events is repeated many times throughout the learning process as the training data is fed to the system. At the start of the learning process the elastic net is very flexible due to large neighbourhoods and gain factor, but as learning continues the net stiffens up as these parameters become smaller. This process causes neurons close together to form similar codebook values.

[0132] During this learning phase, the codebook vectors tend to approximate various distributions of input vectors with some sort of regularity and the resulting order always reflects properties of the probability density function $P(x)$ (ie the point density of the reference vectors becomes proportional to $[P(x)]^{1/3}$). A similar effect is found in biological neural systems where the number of neurons within regions of the cortex corresponding to different sensory modalities appear to reflect the importance of the corresponding feature set. The importance of a feature set is related to the density of receptor cells connected to that feature as would be expected. However, there also appears to be a strong relationship between the number of neurons representing a

feature and the statistical frequency of occurrence of that feature. The scale of this relationship is often loosely referred to as the magnification factor. While the reference vectors are tending to describe the density function of inputs, local interactions between neurons tend to preserve continuity on the surface of the map. A combination of these opposing forces causes the vector distribution to approximate a smooth hyper-surface in the pattern space with optimal orientation and form that best imitates the overall structure of the input vector density. This is done in such a way as to cause the map to identify the dimensions of the feature space with greatest variance which should be described in the map. The initial ordering of the map occurs quite quickly and is normally achieved within the first 10% of the training phase, but convergence on optimal reference vector values can take a considerable time. The trained network provides a non-linear projection of the probability density function $P(x)$ of the high-dimensional input data x onto a 2-dimensional surface (i.e. the surface of neurons).

[0133] FIG. 5 is a schematic representation of a single modular map. At start-up time the Modular Map needs to be configured with the correct parameter values for the intended arrangement. All the 8-bit weight values are loaded into the system at configuration time so that the system can have either random weight values or pre-trained values at start-up. The index of all individual neurons, which consist of two 8-bit values for the X and Y coordinates, are also selected at configuration time. The flexibility offered by allowing this parameter to be set is perhaps more important for situations where several modules are combined, but still offers the ability to create a variety of network shapes for a stand alone situation. For example, a module could be configured as a one or two dimensional network. In addition to providing parameters for individual neurons at configuration time the parameters that apply to the whole network are also required (i.e. the number of training steps, the gain factor and neighbourhood start values). Intermediate values for the gain factor and neighbourhood size are then determined by the module itself during run time using standard algorithms which utilise the current training step and total number of training steps parameters.

[0134] After configuration is complete, the Modular Map enters its operational phase and data are input. 16 Bits (i.e. two input vector elements) at a time. The handshake system controlling data input is designed in such a way as to allow for situations where only a subset of the maximum possible inputs is to be used. Due to tradeoffs between data input rates and flexibility the option to use only a subset of the number of possible inputs is restricted to even numbers (i.e. 14, 12, 10 etc). However, if only say 15 inputs are required then the 16th input element could be held constant for all inputs so that it does not affect the formation of the map during training. The main difference between the two approaches to reducing input dimensionality is that when the system is aware that inputs are not present it does not make any attempt to use their values to calculate the distance between the current input and the codebook vectors within the network, thereby reducing the workload on all neurons and consequently reducing propagation time of the network.

[0135] After all inputs have been read by the Modular Map the active neuron is determined and its X,Y coordinates are output while the codebook vectors are being updated. As the training process has the effect of creating a topological

map (such that neural activations across the network have a meaningful order as though a feature coordinate system were defined over the network) the X,Y coordinates provide meaningful output. By feeding inputs to the map after training has been completed it is straightforward to derive an activation map which could then be used to assign labels to the outputs from the system.

Lateral Maps

[0136] As many difficult tasks require large numbers of neurons the Modular Map has been designed to enable the creation of networks with up to 65,536 neurons on a single plane by allowing lateral expansion. Each module consists of, for example, 256 neurons and consequently this is the building block size for the lateral expansion of networks. Each individual neuron can be configured to be at any position on a 2-dimensional array measuring up to 256^2 but networks should ideally be expanded in a regular manner so as to create rectangular arrays. The individual neuron does in fact have two separate addresses; one is fixed and refers to the neuron's location on the device and is only used locally; the other, a virtual address, refers to the neuron's location in the network and is set by the user at configuration time. The virtual address is accommodated by two 8-bit values denoting the X and Y coordinates; it is these coordinates that are broadcast when the active neuron on a module has been identified.

[0137] When modules are connected together in a lateral configuration, each module receives the same input vector. To simplify the data input phase it is desirable that the data be made available only once for the whole configuration of modules, as though only one module were present. To facilitate this all modules in the configuration are synchronised so that they act as a single entity. The mechanism used to ensure this synchronism is the data input handshake mechanism. By arranging the input data bus for lateral configurations to be inoperative until all modules are ready to accept input, the modules will be synchronised. All the modules perform the same functionality simultaneously, so they can remain in synchronisation once it has been established, but after every cycle new data is required and the synchronisation will be reinforced.

[0138] All modules calculate the local 'winner' by using all neurons on the module to simultaneously subtract one from their calculated distance value until a neuron reaches a value of zero. The first neuron to reach a distance of zero is the one that initially had the minimum distance value and is therefore the active neuron for that module. The virtual coordinates of this neuron are then output from the module, but because all modules are synchronised, the first module to attempt to output data is also the module containing the 'global winner' (i.e. the active neuron for the whole network). The index of the 'global winner' is then passed to all modules in the configuration. When a module receives this data it supplies it to all its constituent neurons. Once a neuron receives this index it is then able to determine if it is in the current neighbourhood in exactly the same way as if it were part of a stand alone module. Some additional logic external to modules is required to ensure that only the index which is output from the first module to respond is forwarded to the modules in the configuration (see FIG. 6). In FIG. 6, logic block A accepts as inputs the data ready line from each module in the network. The first module to set this

line contains the "global winner" for the network. When the logic receives this signal it is passed to the device ready input which forms part of the two line handshake used by all modules in lateral expansion mode. When all modules have responded to the effect that they are ready to accept the coordinates of the active neuron the module with these coordinates is requested by logic block A to send the data. When modules are connected in this lateral manner they work in synchronisation, and act as though they were a single module which then allows them to be further combined with other modules to form larger networks.

[0139] Once a network has been created in this way it acts as though it were a stand alone modular map and can be used in conjunction with other modules to create a wide range of network configurations. However, it should be noted that as network size increases the number of training steps also increases because the number of training steps required is proportional to the network size which suggests that maps are best kept to a moderate size whenever possible.

Hierarchical Maps

[0140] The Modular Map system has been designed to allow expansion by connecting maps together in different ways to cater for changes in network size, and input vector size, as well as providing the flexibility to enable the creation of novel neural network configurations. This modular approach offers a mechanism that maintains an even workload among processing elements as systems are scaled up, thereby providing an effective parallelism of the Self Organising Map. To facilitate expansion in order to cater for large input vectors, modules are arranged in a hierarchical manner which also appears plausible in terms of biological systems where, for example, layers of neurons are arranged in a hierarchical fashion in the primary visual system with layers forming increasingly complex representations the further up the hierarchy they are situated.

[0141] FIG. 7 shows an example of a hierarchical network, with four modules 10, 12, 14, 16 on the input layer I. The output from each of the modules 12, 14, 16, 18 on the input layer I is connected to the input of an output module 18 on the output layer O. Each of the modules 10, 12, 14, 16, 18 has a 16 bit input data bus, and the modules 10, 12, 14, 16 on the input layer I have 24 handshake lines connected as inputs to facilitate data transfer between them, as will be described hereinafter. The output module 18 has 12 handshake lines connected as inputs, three handshake lines from each of the modules 10, 12, 14, 16 in the input layer I.

[0142] As each Modular Map is limited to a maximum of 16 inputs it is necessary to provide a mechanism which will enable these maps to accept larger input vectors so they may be applied to a wide range of problem domains. Larger input vectors are accommodated by connecting together a number of Modular Maps in a hierarchical manner and partitioning the input data across modules at the base of the hierarchy. Each module in the hierarchy is able to accept up to 16 inputs, and outputs the X,Y coordinates of the active neuron for any given input; consequently there is a fan-in of eight modules to one which means that a single layer in such a hierarchy will accept vectors containing up to 128 inputs. By increasing the number of layers in the hierarchy the number of inputs which can be catered for also increases (i.e. Max Number of inputs $= 2 \cdot 8^n$ where n = number of layers in

hierarchy). From this simple equation it is apparent that very large inputs can be catered for with very few layers in the hierarchy.

[0143] By building hierarchical configurations of Modular Maps to cater for large input vectors the system is in effect parallelising the workload among many processing elements. This approach was preferred over the alternative of using more complex neurons which would be able to accept larger input vectors. There are many reasons for this, not least the problems associated with implementation which, in the main, dictate that hardware requirements increase with increasing input vector sizes catered for.

[0144] Furthermore, as the input vector size increases, so too does the workload on individual neurons which leads to considerable increases in propagation delay through the network. Hierarchical configurations keep the workload on individual neurons almost constant, with an increasing workload being met by an increase in neurons used to do the work. It should be noted that there is still an increase in propagation time with every layer added to the hierarchy.

[0145] To facilitate hierarchical configurations of modular maps it is necessary to ensure that communication between modules is not going to form a bottleneck which could adversely affect the operating speed of the system. To circumvent this, a bus is provided to connect the outputs from up to eight modules to the input of a single module on the next layer of the hierarchy (see FIG. 7). To avoid data collision and provide sequence control, each Modular Map has 16 input data lines plus three lines for each 16 bit input (two vector elements), i.e. 24 handshake lines which corresponds to a maximum of eight input devices.

[0146] Consequently, each module also has a three bit handshake and 16 bit data output to facilitate the interface scheme. One handshake line will be used to advise the receiving module that the sender is present; one line will be used to advise it that the sender is ready to transmit data; and the third line will be used to advise the sender that it should transmit the data. After the handshake is complete the sender will then place its data on the bus to be read by the receiver. The simplicity of this approach negates the need for additional interconnect hardware and thereby keeps to a minimum the communication overhead. However, the limiting factor with regard to these hierarchies and their speed of operation is that each stage in the hierarchy cannot be processed faster than the slowest element at that level, but there are circumstances under which the modules complete their classification at differing rates and thereby affect operational speed. For example, one module may be required to have greater than the 256 neurons available to a single Modular Map and would be made up of several maps connected together in a lateral type of configuration (as described above) which would slightly increase the time required to determine its activations, or perhaps a module has less than its maximum number of inputs thereby reducing its time to determine activations. It should also be noted that under normal circumstances (i.e. when all modules are of equal configurations) that the processing time at all layers in the hierarchy will be the same as all modules are carrying out equal amounts of work; this has the effect of creating a pipelining effect such that throughput is maintained constant even when propagation time through the system is dependent on the number of layers in the hierarchy.

[0147] As each Modular Map is capable of accepting a maximum of 16 inputs and generates only a 2-dimensional output, there is a dimensional compression ratio of 8:1 which offers a mechanism to fuse together many inputs in a way that preserves the essence of the features represented by those inputs with regard to the metric being used.

[0148] An ordered network can be viewed in terms of regions of activation surrounding the point positions of its reference vectors, a technique sometimes referred to as Voronoi sets. With this approach the whole of the feature space is partitioned by hyper-planes marking the boundaries of activation regions, which contain all points from the input space that are closer to the enclosed reference point than to any other point in the network. These regions normally meet each other in the same order as the topological arrangement of neurons within the network. As with most techniques applied to artificial neural networks, this approach is only suitable for visualisation in two or three dimensions, but can still be used to visualise what is happening within hierarchical configurations of Modular Maps. The series of graphs shown in FIGS. 8 to 10 emphasise some of the processes taking place in hierarchical configurations. Although a 2-D data set has been used for clarity, the processes identified here are also applicable to higher dimensional data.

[0149] A Modular Map containing 64 neurons configured in a square with neurons equally spaced within a 2-D plane measuring 256^2 was trained on 2000 data points randomly selected from two circular regions within the input space of the same dimensions (see FIG. 8). The trained network formed regions of activation as shown in the Voronoi diagram of FIG. 9. From the map shown in FIG. 9 it is clear that the point positions of reference vectors (shown as black dots) are much closer together (i.e. have a higher concentration) around regions of the input space with a high probability of containing inputs. It is also apparent that, although a simple distance metric (Manhattan distance) is being used by neurons, the regions of activation can have some interesting shapes. It should also be noted that the formation of regions at the outskirts of the feature space associated with the training data are often quite large and suggest that further inputs to the trained system considerably outwith the normal distribution of the training data could lead to spurious neuron activations. It was also-observed that three neurons of the trained network had no activations at all for this data, the reference vector positions of these three neurons (marked on the Voronoi diagram of FIG. 9 by *) fall between the two clusters shown and act as a divider between the two classes.

[0150] As an approach to identifying the processes involved in multidimensional hierarchies, the trained network detailed in FIG. 9 was used to provide several inputs to another network of the same configuration (except the number of inputs) in a way that mimicked a four into one hierarchy (i.e. four networks on the first layer, one on the second). After the module at the highest level in the hierarchy had been trained, it was found that the regions of activation for the original input space were as shown in FIG. 10. Comparison between FIGS. 9 and 10 shows that the same regional shapes have been maintained exactly, except that some regions have been merged together, showing that complicated non-linear regions can be generated in this way without affecting the integrity of classification. It can also be seen that the regions of activation being merged together are

normally situated where there is a low probability of inputs so as to make more efficient use of the resources available and provide some form of compression. It should be noted that there is an apparent anomaly because the activation regions of the three neurons of the first network, which are inactive after training, have not been merged together, the reason being that this region of inactivity is formed naturally between the two clusters during training due to the 'elastic net' effect outlined earlier and is consequently unaffected by the merging of regions. This combining of regions has also increased the number of inactive neurons to eight for the second layer network. The processes highlighted apply to higher dimensional data and suggest that such hierarchical configurations not only provide a mechanism for partitioning the workload of large input vectors, but can also provide a basis for data fusion of a range of data types, from different sources and input at different stages in the hierarchy.

[0151] When modules are connected together in a hierarchical manner there is still the opportunity to partition input data in various ways. The most obvious approach is to simply split the original high dimensional input data into vectors of 16 inputs or less, i.e. given the original feature space \mathfrak{R}^n , is partitioned into groups of 16 or less. When data is partitioned in this way, each module forms a map of its respective input domain, there is no overlap of maps, and a module has no interaction with other modules on its level in the hierarchy. However, it is also realistic to consider an approach where inputs to the system would span more than one module, thereby enabling some data overlap between modules. An approach of this nature can assist modules in their classification by providing them with some sort of context for the inputs; it is also a mechanism which allows the feature space to be viewed from a range of perspectives with the similarity between views being determined by the extent of the data overlap. Simulations have also shown that an overlap of inputs (i.e. feeding some inputs to two or more separate modules) can lead to an improved mapping and classification.

[0152] A similar approach to partitioning could also be taken to give better representation to the range of values in any dimension, i.e. \mathfrak{R} could be partitioned. Partitioning a single dimension of the feature space across several inputs should not normally be required, but if the reduced range of 256 which is available to the Modular Map should prove to be too restrictive for an application, then the flexibility of the Modular Map is able to support such a partitioning approach. The range of values supported by the Modular Map inputs should be sufficient to capture the essence of any single dimension of the feature space, but pre-processing is normally required to get the best out of the system.

[0153] Partitioning \mathfrak{R} is not as simple as partitioning n , and would require a little more pre-processing of input data, but the approach could not be said to be overly complex. However, when partitioning \mathfrak{R} , only one of the inputs used to represent each of the feature space dimensions will contain input stimuli for each input pattern presented to the system. Consequently, it is necessary to have a suitable mechanism to cater for this eventuality, and the possible solutions are to either set the system input to the min or max value depending on which side of the domain of this input the actual input stimuli is on, or do not use an input at all if it does not contain active input stimuli.

[0154] The design of the Modular Map is of such flexibility that inputs could be partitioned across the network system in some interesting ways, e.g. inputs could be taken directly to any level in the hierarchy. Similarly, outputs can also be taken from any module in the hierarchy, which may be useful for merging or extracting different information types. There is no compulsion to maintain symmetry within a hierarchy which could lead to some novel configurations, and consequently separate configurations could be used for specific functionality and combined with other modules and inputs to form systems with increasing complexity of functionality. It is also possible to introduce feedback into Modular Map systems which may enable the creation of some interesting modular architectures and expand possible functionality.

Neural Pathways and Hybrid Networks

[0155] Various types of sensory modalities such as light, sound and smell are mapped to different parts of the brain. Within each of these modalities specific stimuli, e.g. lines or corners in the visual system, act selectively on specific populations of neurons situated in different regions of the cortex. The number of neurons within these regions reflect the importance of the corresponding feature set. The importance of a feature set is related to the density of receptor cells connected to that feature. However, there is also a strong relationship between the number of neurons representing a feature and the statistical frequency of occurrence of that feature. The scale of this relationship is often loosely referred to as the magnification factor.

[0156] While the neocortex contains a great many neurons, somewhere in the region of 10^9 , it only contains two broad categories of neuron; smooth neurons and spiny neurons. All the neurons with spines (pyramidal cells and spiny stellates) are excitatory and all smooth neurons (smooth stellates) are inhibitory. The signals presented to neurons are also limited to two types of electrical message. The mechanisms by which these signals are generated are similar throughout the brain and the signals themselves cannot be endowed with special properties because they are stereotyped and much the same in all neurons. It seems that with such a limited range of components with stereotyped signals that the connections will have an important bearing on the capabilities of the brain.

[0157] It may be possible to facilitate dynamically changing context dependent pathways within Modular Map systems by utilising feedback and the concepts of excitatory and inhibitory neurons as found in nature. This prospect exists because the interface of a Modular Map allows for the processing of only part of the input vector, and supports the possibility of a module being disabled. The logic for such inhibitory systems would be external to the modules themselves, but could greatly increase the flexibility of the system. Such inhibition could be utilised in several ways to facilitate different functionality, e.g. either some inputs or the output of a module could be inhibited. If insufficient inputs were available a module or indeed a whole neural pathway could be disabled for a single iteration, or if the output of a module were to be within a specific range then parts of the system could be inhibited. Clearly, the concept of an excitatory neuron would be the inverse of the above with parts of the system only being active under specific circumstances.

[0158] When implementing ANNs in hardware difficulties are encountered as network size increases. The underlying reasons for this are silicon area, pin out considerations and inter-processor communications. By utilising a modular approach towards implementation, the inherent partitioning strategy overcomes the usual limitations on scalability. Only a small number of neurons are required for a single module and separate modules are implemented on separate devices.

[0159] The Modular Map design is fully digital and uses a fine grain implementation approach, i.e. each neuron is implemented as a separate processing element. Each of these processing elements is effectively a simple Reduced Instruction Set Computer (RISC) with limited capabilities, but sufficient to perform the functionality of a neuron. The simplicity of these neurons has been promoted by applying modifications to Kohonen's original algorithm. These modifications have also helped to minimise the hardware resources required to implement the Modular Map design.

Background

[0160] Essentially the Self-Organising Map (SOM) consists of a two dimensional array of neurons connected together by strong lateral connections. Each neuron has its own reference vector which input vectors are measured against. When an input vector is presented to the network, it is passed to all neurons constituting the network. All neurons then proceed to measure the similarity between the current input vector and their local reference vectors. This similarity is assessed by calculating the distance between the input vector and the reference vector, generally using the Euclidean distance metric. In the Modular Map implementation Euclidean distance is replaced by Manhattan distance because Manhattan distance can be determined using only an adder/subtractor unit whereas calculations of Euclidean distances require determination of the squares of differences involved and would therefore require a multiplier unit which would use considerably greater hardware resources.

[0161] There are a range of techniques that could be utilised to perform the multiplication operations required to calculate Euclidean distance. These include multiple addition operations, which would introduce unacceptable time delays, or traditional multiplier units such as a Braun's multiplier, but compared to an adder/subtractor unit the resource requirements would be significantly increased. There would also be an increase in the time required to obtain the result of a multiplication operation compared to the addition/subtraction required to calculate Manhattan distance. Furthermore, when using multiplication, the number of bits in the result is equal to the number of bits in the multiplicand plus the number of bits in the multiplier, which would produce a 16 bit result for an 8 bit by 8 bit multiplication and would therefore require at least a 16 bit adder to calculate the sum of distances. This requirement would further increase the resource requirements for calculating Euclidean distance and, consequently, further increases the advantages of using the Manhattan distance metric.

[0162] Once all neurons in the network have determined their respective distances they communicate via strong lateral connections with each other to determine which amongst them has the minimum distance between its reference vector and the current input. The Modular Map implementation maintains strong local connections, but determi-

nation of the winner is achieved without the communications overhead suggested by Kohonen's original algorithm. All neurons constituting the network are used in the calculations to determine the active neuron and the workload is spread among the network as a result.

[0163] During the training phase of operation all neurons in the immediate vicinity of the active neuron update their reference vectors to bring them closer to the current input. The size of this neighbourhood changes throughout the training phase, initially being very large and finally being restricted to the active neuron itself. The shape of neighbourhood can take on many forms, the two most popular being a square step function and a gaussian type neighbourhood. The Modular Map approach again utilises Manhattan distance to measure the neighbourhood, which results in a square neighbourhood, but it is rotated through 45 degrees so that it appears to be a diamond shape (FIG. 3). This further assists the implementation because an adder/subtractor unit is still all that is required at this stage. However, additional hardware is required to update reference vector values because reference vectors are only updated by a proportion of the distance between the input and reference vectors. The proportionality of the update applied is determined by what is normally referred to as the gain factor $\alpha(t)$ which Kohonen specifies as a decreasing monotonic function. Consequently, a mechanism is required that will enable multiplication of distances by a suitable range of fractional values. This is achieved by restricting $\alpha(t)$ to negative powers of two. By restricting $\alpha(t)$ in this way it is possible to perform the required multiplication by using only an arithmetic shifter, which is considerably less expensive in terms of hardware resources than a full multiplier unit.

The Neuron

[0164] The Modular Map approach has resulted in a simple Reduced Instruction Set Computer (RISC) type architecture for neurons. The key elements of the neuron design which are shown in FIG. 11 are an adder/subtractor unit (ALU) 50, a shifter mechanism 52, a set of registers and control logic 54. The ALU 50 is the main computational component and by utilising an arithmetic shifter mechanism 52 to perform all multiplication functions, the ALU 50 requirements have been kept to a minimum.

[0165] All registers in a neuron are individually addressable as 8 or 12 bit registers although individual bits are not directly accessible. Instructions are received by the neuron from the module controller and the local control logic interprets these instructions and coordinates the operations of the individual neuron. This task is kept simple by maintaining a simple series of instructions that only number thirteen in total.

[0166] The adder/subtractor unit 50 is clearly the main computational element within a neuron. The system needs to be able to perform both 8 bit and 12 bit arithmetic, with 8 bit arithmetic being the most frequent. A single 4 bit adder/subtractor unit could be utilised to do both the 8 bit and 12 bit arithmetic, or an 8 bit unit could be used. However, there will be considerably different execution times for different sizes of data if a 12 bit adder/subtractor unit is not used (e.g. if an 8 bit unit is used it will take approximately twice as long to perform 12 bit arithmetic as it would 8 bit arithmetic because two passes through the adder/subtractor would be required). In order to avoid vari-

able execution times for the different calculations to be performed a 12 bit adder/subtractor unit is preferable.

[0167] A 12 bit adder/subtractor unit utilising a Carry Lookahead Adder (CLA) would require approximately 160 logic gates, and would have a propagation delay equal to the delay of 10 logic gates. The ALU 50 also has two flags and two registers directly associated with it. The two flags associated with the ALU 50 are a zero flag, which is set when the result of an arithmetic operation is zero, and a negative flag, which is set when the result is negative.

[0168] The registers associated with the ALU are both 12 bit; a first register 56 is situated at the ALU output; a second register 58 is situated at one of the ALU inputs. The first register 56 at the output from the ALU 50 is used to buffer data until it is ready to be stored. Only a single 12 bit register 58 is required at the input to the ALU 50 as part of an approach that allows the length of instructions to be kept to a minimum. The design is a register-memory architecture, and arithmetic operations are allowed directly on register values but the instruction length used for the neuron is too small to include an operation and the addresses of two operands in a single instruction. Thus, the second register 58 at one of the ALU inputs is used so that the first datum can be placed there for use in any following arithmetic operations. The address of the next operand can be provided with the operator code and, consequently, the second datum can be accessed directly from memory.

[0169] The arithmetic shifter mechanism 52 is only required during the update phase of operation to multiply the difference between input and weight elements by the gain factor value $\alpha(t)$. The gain factor $\alpha(t)$ is advantageously restricted to four values (i.e. 0.5, 0.25, 0.125 and 0.0625). Consequently, the shifter mechanism 52 is required to shift right by 0, 1, 2, 3 and 4 bits to perform the required multiplication. The arithmetic shifter 52 can typically be implemented using flip flops which is a considerable improvement on the alternative of a full multiplier unit which would require substantially more resources to implement.

[0170] It should be noted that, for the bit shift approach to work correctly, weight values are required to have as many additional bits as there are bit shift operations (i.e. given that a weight value is 8 bits, when 4 bit shifts are allowed, 12 bits need to be used for the weight value). The additional bits store the fractional part of weight values and are only used during the update operation to ensure convergence is possible; there is no requirement to use this fractional part of weight values while determining Manhattan distance.

[0171] For simplicity with flexibility the arithmetic shifter 52 is positioned in the data stream between the output of the ALU 50 and its input register 58, but is only active when the gain value is greater than zero. This approach was regarded as a suitable approach to limiting the number of separate instructions because the gain factor values are supplied by the system controller at the start of the update phase of operations and can be reset to zero at the end of this operational phase.

[0172] The data registers of these RISC neurons require substantial resources and must hold 280 bits of data. The registers must be readily accessible by the neuron, especially the reference vector values which are accessed frequently. In

order for the system to operate effectively access to weight values is required either 8 or 12 bits at a time for each neuron, depending on the phase of operation. This requirement necessitates on-chip memory because there are a total of 64 neurons attempting to access their respective weight values simultaneously. This results in a minimum requirement of 512 bits rising to 768 bits (during the update phase) that need to be accessed simultaneously. Clearly, this would not be possible if the weight values were stored off chip because a single device would not have enough I/O pins to support this in addition to other I/O functions required of a Modular Map. There are ways of maximising data access with limited pin outs but, a bottleneck situation could not be entirely avoided if memory were off chip.

[0173] The registers are used to hold reference vector values (16*12 bits), the current distance value (12 bits), the virtual X and Y coordinates (2*8 bits), the neighbourhood size (8 bits) and the gain value $\alpha(t)$ (3 bits) for each neuron. There are also input and output registers (2*8bits), registers for the ALU (2*12), a register for the neuron ID (8 bit) and a one bit register for maintaining an update flag. Of these registers all can be directly addressed except for the output register and update flag, although the neuron ID is fixed throughout the training and operational phases, and like the input register is a read only register as far as the neuron is concerned.

[0174] At start up time all registers except the neuron ID are set to zero values before parameter values are provided by an I/O controller. At this stage the initial weight values are provided by the controller to allow the system to start from either random weight values or values previously determined by training a network. While 12 bit registers are used to hold the weight values, only 8 bits are used for determining a neuron's distance from an input, and only these 8 bits are supplied by the controller at start up; the remaining 4 bits represent the fractional part of the weight value, are initially set to zero, and are only used during weight updates.

[0175] The neighbourhood size is also supplied by the controller at start up but, like the gain factor $\alpha(t)$, it is a global variable that changes throughout the training process requiring new values to be effected by the controller at appropriate times throughout training. The virtual coordinates are also provided by the controller at start up time, but are fixed throughout the training and operational phases of the system and provide the neuron with a location from which to determine if it is within the current neighbourhood. Because virtual addresses are used for neurons, any neuron can be configured to be anywhere within a 256^2 array which provides great flexibility when networks are combined to form systems using many modules. It is advantageous for the virtual addresses used in a network to maximise the virtual address space (i.e. use the full range of possible addresses in both the X and Y dimensions). For example, if a 64 neuron module is used, the virtual addresses of neurons along the Y axis should be 0,0 0,36 0,72 etc. In this way the outputs from a module will utilise the maximum range of possible values, which in this instance will be between 0 and 252. Simulations found that classification results were poor when this practice was not adopted.

[0176] It should also be noted that, because there is a requirement to use mixed sizes of data, an update flag is used

as a switch mechanism for the data type to be used. This mechanism was found to be necessary because when 8 bit values and 12 bit values are being used there are differing requirements at different phases of operation. During the normal operational phase only 8 bit values are necessary but they are required to be the least significant 8 bits, e.g. when calculating Manhattan distance. However, during the update phase of operation both 8 bit and 12 bit values are used. During this update phase all the 8 bit values are required to be the most significant 8 bits and when applying changes to reference vectors the full 12 bit value is required. By using a simple flag as a switch the need for duplication of instructions is avoided so that operations on 8 and 12 bit values can be executed using the same instruction set.

[0177] The control logic within a neuron is kept simple and is predominantly just a switching mechanism. All instructions are the same size, i.e. 8 bits, but there are only thirteen distinct instructions in total. While an 8 bit instruction set would in theory support 256 separate instructions, one of the aims of the neuron design has been to use a reduced instruction set. In addition, separate registers within a neuron need to be addressable to facilitate all the operations required of them and, where an instruction needs to refer to a particular register address, that address effectively forms part of the instruction.

[0178] The instruction length has been set at 8 bits because the data bus is only 8 bits wide which sets the upper limit for a single cycle instruction read. There is also a requirement to address locations of operands for six of the instructions which necessitates the incorporation of up to 25 separate addresses into these instructions and will require 5 bits for the address of the operand alone. However, the total instruction length can still be maintained at 8 bits because instructions that do not require operand addresses can use some of these bits as part of their instruction and, consequently, there is room for expansion of the instruction set within the instruction space.

[0179] All instructions for neuron operations are 8 bits in length and are received from the controller. The first input to a neuron is always an instruction, normally the reset instruction to zero all registers. The instruction set is as follows:

[0180] RDI: (Read Input) will read the next datum from its input and write to the specified register address. This instruction will not affect arithmetic flags.

[0181] WRO: (Write arithmetic Output) will move the current data held at the output register 56 of the ALU to the specified register address. This instruction will overwrite any existing data in the target register and will not affect the systems arithmetic flags.

[0182] ADD: Add the contents of the specified register address to that already held at the ALU input. This instruction will affect arithmetic flags and, when the update register is zero all 8 bit values will be used as the least significant 8 bits of the possible 12, and only the most significant 8 bits of weight vectors will be used (albeit as the least significant 8 bits for the ALU) when the register address specified is that of a weight whereas, when the update register is set to one, all 8 bit values will be set as the most significant bits and all 12 bits of weight vectors will be used.

[0183] SUB: Subtract the value already loaded at the ALU input from that at the specified register address. This instruc-

tion will affect arithmetic flags and will treat data according to the current value of the update register as detailed for the add command.

[0184] BRN: (Branch if Negative) will test the negative flag and will carry out the next instruction if it is set, or the next instruction but one if it is not.

[0185] BRZ: (Branch if Zero) will test the zero flag and will carry out the next instruction if it is set. If the flag is zero the next but one instruction will be executed.

[0186] BRU: (Branch if Update) will test the update flag and will carry out the next instruction if it is set, or the next instruction but one if it is not.

[0187] OUT: Output from the neuron the value at the specified register address. This instruction does not affect the arithmetic flags.

[0188] MOV: Set the ALU input register to the value held in the specified address. This instruction will not affect the arithmetic flags.

[0189] SUP: Set the update register. This instruction does not affect the arithmetic flags.

[0190] RUP: Reset the update register. This instruction does not affect the arithmetic flags.

[0191] NOP: (No Operation) This instruction takes no action for one instruction cycle.

[0192] MRS: Master reset will reset all registers and flags within a neuron to zero.

The Module Controller

[0193] FIG. 12 shows a schematic representation of a module controller for controlling the operation of a number of RISC neurons, one of which is shown in FIG. 11. The Module Controller is required to handle all device input and output in addition to issuing instructions to neurons within a module and synchronising their operations. To facilitate these operations the controller system comprises the I/O ports 60, 62; a programmable read-only-memory (PROM) 64 containing instructions for the controller system and subroutines for the neural array; an address map 66 for conversion between real and virtual neuron addresses; an input buffer 68 to hold incoming data; and a number of handshake mechanisms (see FIG. 12).

[0194] The controller handles all input for a module which includes start-up data during system configuration, the input vectors 16 bits (two vector elements) at a time during normal operation, and also the index of the active neuron when configured in lateral expansion mode. Outputs from a module are also handled exclusively by the controller. The outputs are limited to a 16 bit output representing Cartesian coordinates of the active neuron during operation and parameters of trained neurons such as their weight vectors after training operations have been completed. To enable the above data transfers a bi-directional data bus is required between the controller and the neural array such that the controller can address either individual neurons or all neurons simultaneously; there is no requirement to allow other groups of neurons to be addressed but the bus must also carry data from individual neurons to the controller.

[0195] While Modular Map systems are intended to allow modules to operate asynchronously from each other, except

when in lateral expansion mode it is necessary to synchronise data communication in order to simplify the mechanism required. When two modules have a data connection linking them together a handshake mechanism is used to synchronise data transfer from the module transmitting the data (the sender) to the module receiving the data (the receiver). The handshake is implemented by the module controllers of the sender and receiver modules, only requires three handshake lines and can be viewed as a state machine with only three possible states:

[0196] 1) wait (Not ready for input)

[0197] 2) No Device (No input stream for this position)

[0198] 3) Data Ready (Transfer data)

[0199] The handshake system is shown as a simple state diagram in FIG. 13. With reference to FIG. 13, the wait state 70 occurs when either the sender or receiver (or both) are not ready for data transfer. The no device state 72 is used to account for situations where inputs are not present so that reduced input vector sizes can be utilised. This mechanism could also be used to facilitate some fault tolerance when input streams are out of action so that the system did not come to a halt. The data ready state 74 occurs when both the sender and the receiver are ready to transfer data and, consequently, data transfer follows immediately this state is entered. This handshake system makes it possible for a module to read input data in any sequence. When a data source is temporarily unavailable the delay can be minimised by processing all other input vector elements while waiting for that datum to become available. Individual neurons could also be instructed to process inputs in a different order but, as the controller buffers input data there is no necessity for neurons to process data in the same order it is received. The three possible conditions of this data transfer state machine are determined by two outputs from the sender module and one output from the receiving module. The three line handshake mechanism allows the transfer of data direct to each other wherein no third party device is required, and data communication is maintained as point to point.

[0200] Similarly, data is also output 16 bits at a time, but as there are only two 8 bit values output by the system, only a single data output cycle is required, with the three line handshake mechanism used to synchronise the transfer of data, three handshake connections are also required at the output of a module. However, the inputs are intended to be received from up to eight separate sources, each one requiring three handshake connections thereby giving a total of 24 handshake connections for the input data. This mechanism will require 24 pins on the device but, internal multiplexing will enable the controller to use a single three line handshake mechanism internally to cater for all inputs.

[0201] To facilitate reading the coordinates for lateral expansion mode, a two line handshake system is used. The mechanism is similar to the three line handshake system, except the 'device not present' state is unnecessary and has therefore been omitted.

[0202] The module controller is also required to manage the operation of neurons on its module. To facilitate such control there is a programmable read-only memory (PROM) 64 which holds subroutines of code for the neural array in addition to the instructions it holds for the controller. The

program is read from the PROM and passed to the neural array a single instruction at a time. Each instruction is executed immediately when received by individual neurons. When issuing these instructions the controller also forwards incoming data and processes outgoing data. There are four main routines required to support full system functionality plus routines for setting up the system at start up time and outputting reference vector values etc. at shutdown. The start up and shutdown routines are very simple and only require data to be written to and read from registers using the RDI and OUT commands. The four main routines are required to enable the calculation of Manhattan distance (calcdist); find the active neuron (findactive); determine which neurons are in the current neighbourhood (nbhood); and update reference vectors (update). Each of these procedures will be detailed in turn.

[0203] The most frequently used routine (calcdist) is required to calculate the Manhattan distance for the current input. When an input vector is presented to the system it is broadcast to all neurons an element at a time, (i.e. each 8 bit value) by the controller. As neurons receive this data they calculate the distance between each input value and its corresponding weight value, adding the results to the distance register. The controller reads the routine from the program ROM, forwards it to the neural array and forwards the incoming data at the appropriate time. This subroutine is required for each vector element and will be as follows:

MOV (W_i)	/*Move weight (W_i) to the ALU input register.*/
SUB (X_i)	/*Subtract the value at the ALU register from the next input.*/
MOV (R_i)	/*Move the result (R_i) to the ALU input register.*/
BRN	/*If the result was negative*/
SUB dist	/*distance = distance - R_i */
ADD dist	/*Else distance = distance + R_i */
WRO dist	/*Write the new distance to its register.*/

[0204] Once all inputs have been processed and neurons have calculated their respective Manhattan distances the active neuron needs to be identified. As the active neuron is simply the neuron with minimum distance and all neurons have the ability to make these calculations the workload can be spread across the network. This approach can be implemented by all neurons simultaneously subtracting one from their current distance value repeatedly until a neuron reaches a zero distance value, at which time it would poll the controller to notify it that it was the active neuron. Throughout this process the value to be subtracted from the distance is supplied to the neural array by the controller. On the first iteration this will be zero to check if any neuron has a match with the current input vector (i.e. distance is already zero) thereafter the value forwarded will be one. The subroutine findactive defines this process as follows:

MOV input	/*Move the input to the ALU input register.*/
SUB dist	/*Subtract the next input from the current distance value.*/

-continued

BRZ	/*If result is zero.*/
OUT ID	/*out put the neuron ID.*/
NOP	/*Else do nothing.*/

[0205] On receiving an acknowledge signal from one of the neurons in the network, by way of its ID, the controller would output the virtual coordinates of the active neuron. The controller uses a map (or lookup table) of these coordinates which are 16 bits so that neurons can pass only their local ID (8 bits) to the controller. It is important that the controller outputs the virtual coordinates of the active neuron immediately they become available because when hierarchical systems are used the output is required to be available as soon as possible for the next layer to begin processing the data, and when modules are configured laterally it is not possible to know the coordinates of the active neuron until they have been supplied to the input port of the module.

[0206] When modules are connected together in a lateral manner, each module is required to output details of the active neuron for that device before reference vectors are updated because the active neuron for the whole network may not be the same as the active neuron for that particular module. When connected together in this way, modules are synchronised and the first module to respond is the one containing the active neuron for the whole network only the first module to respond will have its output forwarded to the inputs of all the modules constituting the network. Consequently, no module is able to proceed with updating reference vectors until the coordinates of the active neuron have been supplied via the input of the device because the information is not known until that time. When a module is in 'lateral mode' the two line handshake system is activated and after the coordinates of the active neuron have been supplied the output is reset and the coordinates broadcast to the neurons on that module.

[0207] When coordinates of the active neuron are broadcast, all neurons in the network determine if they are in the current neighbourhood by calculating the Manhattan distance between the active neurons virtual address and their own. If the result is less than or equal to the current neighbourhood value, the neuron will set its update flag so that it can update its reference vector at the next operational phase. The routine for this process (nbhood) is as follows:

MOV Xcoord	/*Move the virtual X coordinate to the ALU input register.*/
SUB input	/*Subtract the next input (X coord) from value at ALU.*/
WRO dist	/*Write the result to the distance register.*/
MOV Ycoord	/*Move the virtual Y coordinate the ALU.*/
SUB input	/*Subtract the next input (Y coord) from value at ALU.*/
MOV dist	/*Move the value in distance register to ALU.*/
ADD result	/*Add the result of the previous arithmetic to the value at ALU input.*/

-continued

MOV result	/*Move the result of the previous arithmetic to the ALU input.*/
SUB input	/*Subtract the next input (neighbourhood val) from value at ALU.*/
BRN	/*If the result is negative.*/
SUP	/*Set the update flag.*/
BRZ	/*If the result is zero.*/
SUP	/*Set the update flag.*/
NOP	/*Else do nothing.*/

[0208] All neurons in the current neighbourhood then go on to update their weight values. To achieve this they also have to recalculate the difference between input and weight elements, which is inefficient computationally as these values have already been calculated in the process of determining Manhattan distance. However, the alternative would require these intermediate values to be stored by each neuron, thereby necessitating an additional 16 bytes of memory per neuron. To minimise the use of hardware resources these intermediate values are recalculated during the update phase. To facilitate this the module controller stores the current input vector and is able to forward vector elements to the neural array as they are required. The update procedure is then executed for each vector element as follows:

RDI gain	/*Read next input and place it in the gain register.*/
MOV W_i	/*Move weight value (W_i) to ALU input.*/
SUB input	/*Subtract the input from value at ALU.*/
MOV result	/*Move the result to the ALU.*/
ADD W_i	/*Add weight value (W_i) to ALU input.*/
BRU	/*If the update flag is set.*/
WRO W_i	/*Write the result back to the weight register.*/
NOP	/*Else do nothing.*/

[0209] After all neurons in the current neighbourhood have updated their reference vectors the module controller reads in the next input vector and the process is repeated. The process will then continue until the module has completed the requested number of training steps or an interrupt is received from the master controller. The term 'master controller' is used to refer to any external computer system that is used to configure Modular Maps. The master controller is not required during normal operation as Modular Maps operate autonomously but is required to supply the operating parameters and reference vector values at start up time, set the mode of operation and collect the network parameters after training is completed. Consequently, the module controller receives instructions from the master controller at these times. To enable this, modules have a three bit instruction interface exclusively for receiving input from the master controller. The instructions received are very basic and the total master controller instruction set only comprises six instructions which are as follows:

[0210] RESET: This is the master reset instruction and is used to clear all registers etc. in the controller and neural array

[0211] LOAD: Instructs the controller to load in all the setup data for the neural array including details of the gain

factor and neighbourhood parameters. The number of data items to be loaded is constant for all configurations and data are always read in the same sequence. To enable data to be read by the controller the normal data input port is used with a two line handshake (the same one used for lateral mode), which is identical to the three line handshake described earlier, except that the device present line is not used.

[0212] UNLOAD: Instructs the controller to output network parameters from a trained network. As with the LOAD instruction the same data items are always output in the same sequence. The data are output from the modules data output port.

[0213] NORMAL: This input instructs the controller to run in normal operational mode

[0214] LATERAL: This instructs the controller to run in lateral expansion mode. It is necessary to have this mode separate to normal operation because the module is required to read in coordinates of the active neuron before updating the neural arrays reference vectors and reset the output when these coordinates are received.

[0215] STOP: This is effectively an interrupt to advise the controller to cease its current operation.

The Module

[0216] An individual neuron is of little use on its own, the underlying philosophy of neural networks dictates that they are required in groups to enable parallel processing and perform the levels of computation necessary to solve computationally difficult problems. The minimum number of neurons that constitute a useful group size is debatable and is led more by the problem to be addressed (i.e. the application) than by any other parameters. It is desirable that the number of neurons on a single module be small enough to enable implementation on a single device. Another consideration was motivated by the fact that Modular Maps are effectively building blocks that are intended to be combined to form larger systems. As these factors are interrelated and can affect some network parameters such as neighbourhood size, it was decided that the number of neurons would be a power of 2 and the network size which best suited these requirements was 256 neurons per module.

[0217] As the Modular Map design is intended for digital hardware there are a range of technologies available that could be used, e.g. full custom very large scale integration (VLSI), semi-custom VLSI, application specific integrated circuit (ASIC) or Field Programmable Gate Arrays (FPGA). A 256 neuron Modular Map constitutes a small neural network and the simplicity of the RISC neuron design leads to reduced hardware requirements compared to the traditional SOM neuron.

[0218] The Modular Map design maximises the potential for scaleability by partitioning the workload in a modular fashion. Each module operates as a Single Instruction Stream Multiple Data stream (SIMD) computer system composed of RISC processing elements, with each RISC processor performing the functionality of a neuron. These modules are self contained units that can operate as part of a multiple module configuration or work as stand alone systems.

[0219] The hardware resources required to implement a module have been minimised by applying modifications to

the original SOM algorithm. The key modification being the replacement of the conventional Euclidean distance metric by the simpler and easier to implement Manhattan distance metric. The modifications made have resulted in considerable savings of hardware resources because the modular map design does not require conventional multiplier units. The simplicity of this fully digital design is suitable for implementation using a variety of technologies such as VLSI or ASIC.

[0220] A balance has been achieved between the precision of vector elements, the reference vector size and the processing capabilities of individual neurons to gain the best results for minimum resources. The potential speedup of implementing all neurons in parallel has also been maximised by storing reference vectors local to their respective neurons (i.e. on chip as local registers). To further support maximum data throughput simple but effective parallel point to point communications are utilised between modules. This Modular Map design offers a fully digital parallel implementation of the SOM that is scaleable and results in a simple solution to a complex problem.

[0221] One of the objectives of implementing Artificial Neural Networks (ANNs) in hardware is to reduce processing time for these computationally intensive systems. During normal operation of ANNs significant computation is required to process each data input. Some applications use large input vectors, sometimes containing data from a number of sources and require these large amounts of data processed frequently. It may even be that an application requires reference vectors updated during normal operation to provide an adaptive solution, but the most computationally intensive and time consuming phase of operation is network training. Some hardware ANN implementations, such as those for the multi-layer perceptron, do not implement training as part of their operation, thereby minimising the advantage of hardware implementation. However, Modular Maps do implement the learning phase of operation and, in so doing, maximise the potential benefits of hardware implementation. Consequently, consideration of the time required to train these networks is appropriate.

Background

[0222] The modular approach towards implementation results in greater parallelism than does the equivalent unitary network implementation. It is this difference in parallelism that has the greatest effect on reducing training times for Modular Map systems. Consideration was given to developing mathematical models of the Modular Map and SOM algorithms for the purpose of simulating training times of the two systems.

[0223] The Modular Map and SOM algorithms have the same basic phases of operation, as depicted in the flowchart of FIG. 14. When considering an implementation strategy in terms of partitioning the workload of the algorithm and employing various scales of parallelism, the potential speedup of these approaches should be considered in order to minimise network training time. Of the five operational phases shown in FIG. 14, only two are computationally intensive and therefore significantly affected by varying system parallelism. These two phases of operation involve the calculation of distances between the current input and the reference vectors of all neurons constituting the network,

and updating the reference vectors of all neurons in the neighbourhood of the active neuron (i.e. phases 2 and 5 in FIG. 14).

[0224] To facilitate investigation into the potential speedup of Modular Map systems over the alternative unitary networks and serial implementation, the model used was based on the two computationally intensive phases of operation mentioned above. This allows assessment of the trends in training times while varying parameters such as network size and vector size, and facilitating an understanding of the relative training times for different implementation strategies.

Training Times for Parallel Implementation

[0225] A simplified mathematical model of the Modular Map can be constructed for the purpose of assessing training times. The starting point for this model will be the neuron, as it is the fundamental building block of the Modular Map. When the neuron is presented with an input vector $x = [\epsilon_1, \epsilon_2, \dots, \epsilon_n] \in \mathbb{R}^n$ it proceeds to calculate the distance between its reference vector $m_i = [\mu_{i1}, \mu_{i2}, \dots, \mu_{in}] \in \mathbb{R}^n$ and the current input vector x . The distance calculation used by the Modular Map is the Manhattan distance, i.e.

$$\text{Distance} = \sum_{j=0}^n |\xi_j - \mu_j|$$

where n =vector size.

[0226] The differences between vector elements are calculated in sequence as while all neurons are implemented in parallel, vector elements are not. To implement the system utilising this level of parallelism is not practical because it would require either 16 separate processors per neuron, or a vector processor for each neuron, so that the distances between all vector elements could be calculated simultaneously. The resources required to process all vector elements in parallel would be substantially greater than the requirements of the RISC neuron (FIG. 11) and would greatly reduce the chances of implementing a Modular Map on a single device. Consequently, when n dimensional vectors are used, n separate calculations are required.

[0227] If the time required by a neuron to determine the distance for one dimension is taken to be t_d seconds and there are n dimensions, then the total time taken to calculate the distance between input and reference vectors (d) will be nt_d seconds i.e. $d = nt_d$ (seconds). The summation operation is carried out as the distance between each element is determined and is therefore a variable overhead dependent on the number of vector elements, and does not affect the above equation for distance calculation time. However, the value for t_d will reflect the additional overhead of this summation operation, as it will all variable overheads proportional to vector size for this calculation. The reason being that the distance calculation time (t_d) is the fundamental timing unit used in this model. It has no direct relationship to the time an addition or subtraction operation will take for any particular device; it is the time required to calculate the distance for a single element of a reference vector including all variable overheads associated with this operation.

[0228] As all neurons are implemented in parallel the total time required for all neurons to calculate Manhattan distance will be equal to the time it takes for a single neuron to calculate its Manhattan distance. Once neurons have calculated their Manhattan distances the active neuron has to be identified before any further operations can be carried out. This process involves all neurons simultaneously subtracting one from their current distance value until one neuron reaches a value of zero. As this process only continues until the active neuron has been identified, (the neuron with minimum distance) relatively few subtraction operations are required.

[0229] Data generated during the training of Modular Maps for the GRANIT application (discussed later) was used to evaluate the overheads involved in finding the active neuron. FIG. 15 is a graph of the activation values (Manhattan distances) of the active neuron for the first 100 training steps. The data was generated for a 64 neuron Modular Map with 16 inputs using a starting neighbourhood covering 80% of the network. The first few iterations of the training phase (less than 10) have a high value for their Manhattan distances as can be seen from FIG. 15. However, after the first 10 iterations there is little variation for the distances between the reference vector of the active neuron and the current input. Thus, the average activation value after this initial period is only 10, which would require only 10 subtraction operations to find the active neuron. Consequently, there is a substantial overhead for the first few iterations, but these will be similar for all networks and can be regarded as a fixed overhead which is not accounted for in the simple timing model used. Throughout the rest of the training phase the overhead of calculating the active neuron is insubstantial and will be assumed to be negligible for the sake of simplicity.

[0230] During the training phase of operation, reference vectors are updated after the distances between the current input and the reference vectors of all neurons have been calculated. This process again involves the calculation of differences between vector elements as detailed above. Computationally this is inefficient because these values have already been calculated during the last operational phase. However, to have used the previously calculated values would have required an additional 16 bytes of local memory for each neuron to store these values and to avoid the additional resource overhead these values are recalculated. After the distance between each element has been calculated these intermediate results are then multiplied by the gain factor. The multiplication phase is carried out by an arithmetic shifter mechanism which is placed within the data stream and therefore does not require any significant additional overhead (see FIG. 11). The addition of these values to the current reference vector will have an impact on the update time for a neuron approximately equivalent to the original summation operation carried out to determine the differences between input and reference vectors. Consequently, the time taken for a neuron to update its reference vector is approximately equal to the time it takes to calculate the Manhattan distance, i.e. d (seconds), because the processes involved are the same (i.e. difference calculations and addition). The number of neurons to have their reference vectors updated in this way varies throughout the training period, often starting with approximately 80% of the network and reducing to only one by the end of training. However, the time a Modular Map takes to update a single

neuron will be the same as it requires to update all its neurons because the operations of each neuron are carried out in parallel.

[0231] Kohonen states that the number of training steps required to train a single network is proportional to network size. So let the number of training steps (s) be equal to the product of the proportionality constant (k) and the network size (N) (i.e. Number of training steps required (s)=kN). From this simplified mathematical model it can be seen that the total training time (T_{par}) will be the product of the number of training steps required (s), the time required to process each input vector (d), and the time required to update each reference vector (d) i.e. Total training time (T_{par})=2ds (seconds), but $d=nt_d$ and $s=kN$, so substituting and rearranging gives:

$$T_{\text{par}}=2Nnkt_d \quad \text{Equation 1.1}$$

[0232] This simplified model is suitable for assessing trends in training times and shows that the total training time will be proportional to the product of the network size and the vector size, but the main objective is to assess relative training times. In order to assess relative training times consider two separate implementations with identical parameters, excepting that different vector sizes, or network sizes, are used between the two systems such that vector size n_2 is some multiple (y) of vector size n_1 . If $T_1=2Nn_1kt_d$ and $T_2=2Nn_2kt_d$, then by rearranging the equation for T_1 , $n_1=T_1/(2Nkt_d)$ but, $n_2=yn_1=y(T_1/(2Nkt_d))$. By substituting this result into the above equation for T_2 it follows that:

$$T_2=2Ny(T_1/(2Nkt_d))kt_d=yT_1 \quad \text{Equation 1.2}$$

[0233] The consequence of this simple analysis is that a module containing simple neurons with small reference vectors will train faster than a network of more complex neurons with larger reference vectors. This analysis can also be applied to changes in network size where it shows that training time will increase with increasing network size. Consequently, to minimise training times both networks and reference vectors should be kept to a minimum as is done with the Modular Map.

[0234] This model could be further expanded to consider hierarchical configurations of Modular Maps. One of the advantages of building a hierarchy of modules is that large input vectors can be catered for without significantly increasing the system training time.

[0235] This situation arises because the training time for a hierarchy is not the sum of training times for all its constituent layers, but the total training time for one layer plus the propagation delays of all the others. The propagation delay of a module (T_{prop}) is very small compared to its training time and is approximately equal to the time taken for all neurons to calculate the distance between their input and reference vectors. This delay is kept to a minimum because a module makes its output available as soon as the active neuron has been determined, and before reference vectors are updated. A consequence of this type of configuration is that a pipelining effect is created with each successive layer in the hierarchy processing data derived from the last input of the previous layer.

$$T_{\text{prop}}=nt_d \quad \text{Equation 1.3}$$

[0236] All modules forming a single layer in the hierarchy are operating in parallel and a consequence of this parallelism is that the training time for each layer is equal to the

training time for a single module. When several modules form such a layer in a hierarchy the training time will be dictated by the slowest module at that level which will be the module with the largest input vector (assuming no modules are connected laterally). As a single Modular Map has a maximum input vector size of 16 elements and under most circumstances at least one module on a layer will use the maximum vector size available, then the vector size for all modules in a hierarchy (n_h) can be assumed to be 16 for the purposes of this timing model. In addition, each module outputs only a 2-dimensional result which creates an 8:1 data compression ratio so the maximum input vector size catered for by a hierarchical Modular Map configuration will be 2×8^l (where l is the number of layers in the hierarchy). Consequently, large input vectors can be accommodated with very few layers in a hierarchical configuration and the propagation delay introduced by these layers will, in most cases, be negligible. It then follows that the total training time for a hierarchy (T_h) will be:

$$T_h=2Nn_hkt_d+(l-1)n_hkt_d \approx 2Nn_hkt_d \quad \text{Equation 1.4}$$

[0237] By following a similar derivation to that used for equation 1.2 it can be seen that:

$$T_{\text{par}} \approx yT_h \quad \text{Equation 1.5}$$

[0238] Where the scaling factor $y=n/n_h$.

[0239] This modular approach meets an increased workload with an increase in resources and parallelism which results in reduced training times compared to the equivalent unitary network and, this difference in training times is proportional to the scaling factor between the vector sizes (i.e. y).

Training Times for Serial Implementation

[0240] The vast majority of ANN implementations have been in the form of simulations on traditional serial computer systems which effectively offer the worst of both worlds because a parallel system is being implemented on a serial computer. As an approach to assessing the speedup afforded by parallel implementation the above timing model can be modified. In addition, the validity of this model can be assessed by comparing predicted relative training times with actual training times for a serial implementation of the Modular Map.

[0241] The main difference between parallel and serial implementation of the Modular Map is that the functionality of each neuron is processed in turn which will result in a significant increase in the time required to calculate the Manhattan distances for all neurons in the network compared to a parallel implementation. As the operations of neurons are processed in turn there will also be a difference between the time required to calculate Manhattan distances and update reference vectors. The reason for this disparity with serial implementation is that only a subset of neurons in the network have their reference vectors updated, which will clearly take less time than updating all neurons constituting the network when each reference vector is updated in turn.

[0242] The number of neurons to have their reference vectors updated varies throughout the training period, starting with 80% and reducing to only one by the end of training. As this parameter varies with time it is difficult to incorporate into the timing model, but as the neighbourhood

size is decreasing in a regular manner the average neighbourhood size over the whole training period covers approximately 40% of the network. The time required to update each reference vector is also approximately equal to the time required to calculate the distance for each reference vector, and consequently the time spent updating reference vectors for a serial implementation will average 40% of the time spent calculating distances. In order to maintain simplicity of the model being used, the workload of updating reference vectors will be evenly distributed among all neurons in the network and, consequently, the time required for a neuron to update its reference vectors will be 40% of the time required for it to calculate the Manhattan distance, i.e. update time=0.4d (seconds).

[0243] In this case equation 1.1 becomes:

$$T_{\text{serial}} = 1.4N^2nkt_d(\text{seconds}) \quad \text{Equation 1.6}$$

[0244] This equation clearly shows that for serial implementation the training time will increase in proportion to the square of the network size. Consequently, the training time for serial implementation will be substantially greater than for parallel implementation. Furthermore, comparison of equation 1.1 and 1.6 shows that $T_{\text{serial}} = 0.7 NT_{\text{par}}$, i.e. the difference in training time for serial and parallel implementation will be proportional to the network size.

[0245] A series of simulations were carried out using a single processor on a PowerXplorer system to assess the trends and relationships between training times for serial implementation of Modular Maps and provide some evidence to support the model being used. The simulations used a Modular Map simulator (MAPSIM) to train various Modular Maps with a range of network and vector sizes. As the model does not take account of data input and output overheads these were not used in the determination of training times, although the training times recorded did include the time taken to find the active neuron.

[0246] Some assumptions and simplifications have been incorporated into this model, but have been incorporated in such a way as to facilitate a good approximation of timing behaviour. The simulations that were run to help evaluate this model showed that trends in training time did follow those prescribed by equation 1.6 (see FIG. 16). FIG. 16 shows that the range of training time required for a 99 element vector increases substantially for increased network size, whereas for a 16 element vector, the increase in training time is not so substantial. When the actual training time is known for one configuration, the training times for other configurations can be calculated using equation 1.2 and all predicted times using this approach were within 10% of the actual training time measured on the PowerXplorer.

[0247] The three main implementation strategies are serial implementation, fine grain parallelism for a unitary network and fine grain parallelism for a modular network. FIG. 17 is a graph which has been constructed to show the theoretical differences in training times for these three strategies. The training times presented for serial implementation have been derived from actual training times measured on the PowerXplorer and the other plots have been calculated relative to these values using the model. FIG. 17 clearly indicates that a modular approach to implementation which utilises fine grain parallelism offers considerably reduced training times compared to the other strategies considered.

[0248] The model has been developed from the two computationally intensive phases of operation that involve the calculation of distances and updating of reference vectors, as shown in FIG. 14. These are the phases of operation that will be most affected by increasing system parallelism and offer a good approximation of timing behaviour.

[0249] Consideration could also be given to the overheads of data input and output for these implementation strategies although the impact of these overheads will be minimal compared to the time required for the computationally intensive phases of operation mentioned above. The data output operation involves outputting the XY coordinates of the active neuron for the Modular Map. This approach could also be used for the other implementation approaches considered here. The Modular Map design allows the output to be made available as soon as the coordinates of the active neuron have been determined. Both output values are maintained at the output of the device until they are read, but once the output has been made available the other processes continue, leaving the data transfer to be handled by an autonomous handshake system. The same approach could be adopted by a unitary network system, but serial implementation would have to output the X and Y coordinates separately and all other processing would have to stop while these operations were being carried out. This would result in the serial implementation taking more time to perform data output than the other two approaches, but the impact on overall training time would be minimal.

[0250] The data input phase of operation requires more time than does data output, but again the Modular Map design aims to minimise the overheads involved. The Modular Map will require a maximum of eight read cycles per input vector because input vectors have a maximum of 16 elements and two of these elements are read on each cycle. In addition, the inputs for Modular Maps are buffered and most of these read cycles can be carried out while previously read data is being processed by the neural array. If the same approach were used for a unitary network with larger input vectors, the overheads would be similar because the neural array would be processing previously read data while new data was being input to the data buffer. Again it is the serial implementation strategy that will suffer the greatest overhead for this phase of operation because each vector element has to be read in separately, and while data is being input no other processing is able to proceed. Consequently, serial implementation will suffer a data input overhead proportional to the vector size.

Applications

[0251] Modular Maps offer a versatile implementation of Kohonen's Self-Organising Map (SOM) that is suitable for use in a wide variety of problem domains. Two possible applications have been used as examples of the applications for which Modular Maps are suited; human face recognition and ground anchorage integrity testing. The applications have little in common other than their ill-defined nature but, Modular Maps offer possible solutions in both domains. The SOM is also applied to these problems to provide a benchmark for the Modular Map approach.

[0252] Human face recognition is an ill-defined problem that is difficult to tackle using conventional computing techniques but has aspects that make it amenable to solution by neural network systems. There are many approaches to

the face recognition problem that have been attempted over the years utilising a range of techniques including statistical and genetic algorithm approaches. However, the aim here is to assess Modular Maps as an alternative to the traditional SOM. Consequently, comparisons are only made between the SOM and Modular Map solutions.

[0253] As the SOM is the basis for the Modular Map design, the classification and clustering of the two systems are further compared in the application domain of ground anchorage integrity testing (GRANIT). This is also an application that is difficult to tackle using conventional computing techniques, but its ill-defined nature and high noise levels make it a suitable application for a neural network solution. The application is currently being developed at the University of Aberdeen to provide an easy to use mechanism to replace the current conventional test procedures used within the civil engineering industry which are time consuming, expensive and often destructive.

Human Face Recognition

[0254] Human face recognition is generally regarded as a very difficult task for computing systems to undertake. There are databases containing face images available via the Internet, e.g. the Olivetti web site but, like many Internet resources, there is no standardisation from one site to another. Consequently, it is difficult to obtain a data set of face images in a usable format containing sufficient variations and instances of each face to enable training of ANN systems. However, at the University of Aberdeen, Dr Ian Craw of the Department of Mathematics has been working in the field of face recognition for some time and has built several face databases. Access to some of this data was arranged, along with permission to use it as part of the evaluation of Modular Map systems, which avoided the problems of loading large data files from the Internet.

[0255] The data base used for evaluation of Modular Maps was derived from photographs of human faces taken by a colour CCD camera connected to a framegrabber which digitised colour at a resolution of 576×768 pixels. A total database of 378 images made up from 14 photographs of 27 different subjects was created in this way. The photographs were taken over a period of weeks with varying intervals between shots using differing lighting conditions and a variety of orientations of the subject. FIG. 18 shows a typical example of the types of images used in greyscale. Excessive variation was avoided to prevent potential matches based on condition rather than subject. None of the photographs included faces with glasses or beards but the clothing worn by subjects changed throughout their series of photographs.

[0256] The background of the photographs was eliminated to leave images of 128×128 pixels, but the hair which is not invariant over time was left in the picture. Thirty-four landmarks were then found manually for each image to create a face model. The images are then scaled ('morphed') to minimise the error between landmark positions for individual images and a reference face; the reference face being used here is the average of the ensemble of faces. This process normalises the images for inter-ocular distance and ocular location (i.e. the faces are scaled and translated to put the centre of both eyes in the same X,Y location for all images). This normalisation process removes the effects of different camera locations and face orientations and offers an alternative to positioning subjects carefully before images

are acquired. The average image is calculated from the whole database and, in addition to being used as detailed above, is subtracted from each image resulting in a face subspace of $n-1$, where n was the original dimensionality of the images.

[0257] Principal Component Analysis (PCA) may then be performed separately on the shape-free face images and the shape vectors consisting of the X,Y location of the points on the original face image. The data used for the evaluations used the shape-free face images. The normalised images were considered as raster vectors and subjected to PCA where the eigenvalues and unit eigenvectors (eigenfaces of 99 elements) of the image cross-correlation matrix were obtained. PCA has the effect of reducing the dimensionality of the data by "transforming to a new set of variables (principal components) which are uncorrelated, and which are ordered so that the first few components retain most of the variation present in all of the original variables". While PCA is a standard statistical technique for reducing the dimensionality of data and attempting to preserve as much of the original information as possible it is difficult to give meaningful labels to individual components.

[0258] Hancock and Burton have investigated principal component representations of faces and suggest several correlations with PCA components of shape vectors and face features such as head size, nodding and shaking of the head and variations in face shape. However, little is suggested about the correlations between PCA components derived from the shape-free vectors and face features. It appears that individual PCA components derived from shape free face images do not normally correlate directly to individual face features, but the first two components of the eigenface are believed to be associated with the size of the face and lighting conditions. It is because of the application that these eigenvectors are often referred to as eigenfaces.

[0259] It was these eigenfaces that were made available for the Modular Map investigation. In ANN terms this database contained a very limited dataset and, normally many more than 14 instances of a class would be used to train a network. However, this still offered an improvement over other sources such as the Olivetti data base which only had 10 instances of each face. To facilitate both training and testing of ANN systems nine eigenfaces for each subject were used to train a network and the other five were used to test its classification. The test set was selected across the range of orientation and lighting conditions so that the training set would also cover the whole range of conditions.

[0260] The eigenface data consisted of double precision floating point values between minus one and plus one but Modular Maps only accept eight bit inputs. Consequently, the face data needed to be converted to suitable eight bit values before it could be used with Modular Map systems. This was achieved using some utility programs developed for use with Modular Map systems. This software was able to offset data values so that all values were positive, scale the data to cover the range 0 to 255 and convert it to integer (8 bit) values. The effects of this data manipulation do not change the relationships between vector elements as the same scaling and offset are applied to each element but, rounding does occur during the conversion process. It is also perhaps noteworthy that all data used in the training and testing of a network should use the same scaling factor and offset values to maintain its integrity.

[0261] To facilitate the training and testing of neural networks the eigenface data was split into nine training vectors and five test vectors for each face. To ensure that the networks were trained on the whole range of possible orientations and lighting conditions the first two and last two vectors in a class were always used for training. The rest of the data was selected as training vectors and test vectors alternately such that on one simulation eigenfaces 1, 2, 4, 6, 8, 10, 12, 13 and 14 were used to train the network while eigenfaces 3, 5, 7, 9 and 11 were used to test the network. The next simulation would then use eigenfaces 1, 2, 3, 5, 7, 9, 11, 13 and 14 to train the network and eigenfaces 4, 6, 8, 10 and 12 to test the network etc.

[0262] Using Kohonen's Self Organising Map to Classify Face Data

[0263] Simulations using Kohonen's Self Organising Map (SOM) were carried out to provide a benchmark for the Modular Map evaluation. The first of these simulations used the original double precision floating point data and a 64 neuron SOM, but the majority of vectors caused the activation of the same neuron. Investigation found that the problem was that the original data set actually covered a smaller range than had been expected and required excessive precision with regard to the ANN processes. Rather than the data covering the whole range between minus one and plus one, most vector elements had a maximum variance of less than 0.1 over the entire data set and the maximum variance found for any element was less than 0.7. Consequently, it was possible to have vectors originating from different faces with a Euclidean distance much less than one.

[0264] The SOM implementation used double precision values but, rounding errors within the mechanism resulted in problems with the original data set.

[0265] Due to the problems encountered with the original eigenfaces, the data was scaled to cover the range between 0 and 255 but, using floating point values rather than the 8 bit data required for Modular Maps. When the 135 test vectors were presented to the network this approach proved to offer much better results but, high classification error rates of 40% were still encountered (i.e. of the 135 test vectors presented to the network after training, only 81 (60%) were correctly identified). The reason for this poor performance was that each class of data caused the activation of several neurons and there were simply not enough neurons in the network for all activation regions to be distinct (i.e. a larger network was required). FIG. 19a is an example activation region for a modular map and FIG. 19b is an example activation map for a SOM. When the same data was used with a SOM network of 256 neurons the error rate dropped to 6%. When simulations were run using a quantised version of the data set (i.e. using integer values) the results were found to be identical thereby suggesting that the rounding errors within the data introduced by the quantisation process were not significant (see the error rate table (table 1 below).

TABLE 1

Summary Classification Error Rate Table. Figures quoted are mean classification errors with standard deviation. All figures are quoted to the nearest integer value.		
ANN type	Configuration Details	% Error
SOM	64 Neurons Floating point data (99 element vectors)	40 ± 12
SOM	64 Neurons Integer data (99 element vectors)	40 ± 12
SOM	256 Neurons Floating point data (99 element vectors)	6 ± 1
SOM	256 Neurons Integer data (99 element vectors)	6 ± 1
SOM	1024 Neurons Floating point data (99 element vectors)	6 ± 1
SOM	256 Neurons Floating point data Using overlap data (127 element vectors)	7 ± 1
Modular Map	Nine Module Hierarchy 7 with 13 inputs 1 with 8 inputs Output = 64 Neurons (configuration 1)	19 ± 3
Modular Map	Seven Module Hierarchy 6 with 16 inputs Output = 64 Neurons (configuration 2)	18 ± 3
Modular Map	Nine Module Hierarchy Using overlap data 7 with 16 inputs, 1 with 15 inputs Output = 64 Neurons (configuration 3)	11 ± 2
Modular Map	Nine Module Hierarchy Using overlap data 7 with 16 inputs, 1 with 15 inputs Output = 256 Neurons (configuration 4)	4 ± 1

Using Modular Maps to Classify Face Data

[0266] Modular Maps can be combined in different ways and use different data partitioning strategies. Four separate Modular Map configurations are used to outline the effects of using different approaches. The first approach to Modular Map solution of the eigenface classification problem presented is intended more as a 'how not to do' approach. This combination of modules, configuration 1, utilises nine Modular Map networks each with 64 neurons (see FIG. 20). The topology of the system is hierarchical with eight modules at the base of the hierarchy (the input layer I) and one at the output level (output layer O). The data was partitioned so that seven modules each had 13 inputs and one module had 8 inputs. This data partitioning strategy may result in poor classification because a module will give better results when the whole of the reference vector is utilised (i.e. when all 16 inputs are used).

[0267] The results from simulations using configuration 1 (FIG. 20) showed poor classification of the face data with an average classification error of 19% from the output module. It can also be seen from table 2 below that the error rate for

module 7, which only has eight inputs as opposed to the 13 used by all other networks at that level, are much higher than all other networks.

[0268] A factor contributing to this is that module 7 has much fewer inputs, which will naturally lead to poorer performance but, it should also be noted that there is a general trend of classification errors from modules at the base of the hierarchy which correlates to the importance of the elements of the eigenvectors (i.e. the first few PCA elements have most of the variation). However, the small number of vector elements used is the most prominent factor contributing to poor performance and this is highlighted by the results of configuration 2 (FIG. 21) which show considerably better classification results for most modules at the base of the hierarchy when all 16 inputs are used.

TABLE 2

Error Rate Table for Configuration 1 (FIG. 20)		
Module	No of Inputs	% Error
0	13	20
1	13	22
2	13	21
3	13	21
4	13	28
5	13	29
6	13	29
7	8	39
8	16	19

[0269] The second Modular Map configuration (configuration 2 shown in FIG. 21) used only seven modules in total; six on the input layer I of the hierarchy and one at the output layer O. The data was partitioned so that all modules at the base of the hierarchy had sixteen inputs, which gives a total of 96 input vector elements as opposed to the 99 in the original eigenfaces; the final three elements of the eigenfaces being the least significant ones and therefore omitted.

[0270] The results from this series of simulations showed an improved classification but, only an increase of 1% on the previous error rates for the output module were achieved (table 3 below). The overall performance increase is due in part to the fact that the output module is now only using 12 out of the 16 possible inputs. However, most modules had reduced error rates compared to the previous series of simulations and all modules had better classification rates than had been experienced for module 7 in configuration 1 (FIG. 20). An additional two modules could be added to the base of the hierarchy so that the output module would be using all of its inputs. One possible approach would be to simply present the first 16 elements of the eigenfaces to two modules. This type of approach is normally referred to as an ensemble and has been found to improve classification. There are no known dependencies between vector elements of the eigenfaces and there is no direct correlation between individual elements and particular face features so the data overlap approach was used to spread the data being used for two inputs across the whole vector rather than relying solely on any one block of 16 elements.

TABLE 3

Error Rate Table for Configuration 2 (FIG. 21)		
Module	No of Inputs	% Error
0	16	21
1	16	20
2	16	21
3	16	22
4	16	25
5	16	25
6	16	28
7	14	18

[0271] Utilising all inputs for modules at the base of the hierarchy improves classification. To maximise on this and the number of inputs to the next layer of the hierarchy, some of the input vector elements can be fed to more than one module. This 'data overlap' technique is where the data is split into groups of 16 element inputs, but the last few elements of one input vector are also used as inputs for the next module. This was accomplished by feeding vector elements 0 to 15 to module 0 and, elements 12 to 27 to module 1 etc. so that there was effectively an overlap of four vector elements between modules. In this way modules 0 to 6 all had 16 inputs but, module 7 only had 15 because when using the original 99 element vectors this was the closest to maximum input usage that could be achieved without using different strategies for different modules. This approach was chosen because it enables most modules at the base of the hierarchy to have 16 inputs and therefore helps to maximise the limited amount of training data.

[0272] As with the first configuration, a total of nine modules all with 64 neurons were used and were connected together in a hierarchical manner as shown in FIG. 22. The simulations carried out using this 'data overlap' approach showed a significant improvement over configurations 1 and 2 (FIGS. 20 and 21) because the classification error from the output module had been reduced to 11%. However, the classification errors for modules at the base of the hierarchy did not show any significant statistical difference to those found with configuration 2 (FIG. 21) (compare table 3 and table 4 below). This suggests that the improvement in classification is not due to the particular partitioning strategy used, but to the fact that more inputs to the hierarchy were used.

TABLE 4

Error Rate Table for Configuration 3 (FIG. 22)		
Module	No of Inputs	% Error
0	16	21
1	16	20
2	16	19
3	16	21
4	16	24
5	16	24
6	16	26
7	15	28
8	16	11

[0273] From the simulations performed using the SOM it was noted that the activation regions for the face data were such that a 256 neuron SOM was required to classify the

data with reasonable accuracy. The simulations carried out using Modular Maps for this data found that fewer neurons were active on the output module of a Modular Map hierarchy than for the SOM. This occurs because of the data compression being performed by successive layers in the hierarchy and results in a situation where fewer neurons are required in the output network of a hierarchy of Modular Maps than are required by a single SOM for the same problem. However, when only a two layer hierarchy is being used the compression is not sufficient for a 256 neuron module to be replaced by a 64 neuron module. In addition, Modular Maps can be combined both laterally and hierarchically to provide the architecture suitable for numerous applications.

[0274] Configuration 4 (FIG. 23) has 256 neurons at the output layer O of a Modular Map hierarchy but all other modules in the system were still maintained at 64 neurons. To create an array of 256 neurons, four Modular Maps are connected together in a lateral configuration and because modules connected in this way act as though they were a single Modular Map they can then be further combined to create hierarchies containing different sized networks.

[0275] For these simulations the input data and the eight base modules were identical to those detailed for configuration 3 (FIG. 22); the only change was to the size of the output module. The results of these simulations showed that the classification error at the output of the hierarchy had been reduced to 4% (the results from layer one being identical to those for configuration 3) which offered an improvement over all previous simulations, including the ones using the standard Kohonen network.

ANN Classification of Faces

[0276] The hardware required to provide the Modular Map solution for this face recognition problem would comprise 12 modules which could be implemented on twelve VLSI devices. The SOM solution, however, would require a network of 256 neurons, each capable of using reference vectors of 99 elements. The digital hardware requirements for a parallel implementation of such a SOM would not fit onto a single VLSI device and would require wafer scale integration for a monolithic implementation. Even when attempting to implement this SOM on several separate devices there are no known systems with a comparable level of parallelism to the Modular Map solution outside the realms of neuro-computers and super-computers. There are, of course, many other ways of implementing a SOM of this size, e.g. transputer systolic array, but at present the difficulties of implementing this comparatively small SOM network on a single device in digital hardware have been sufficient to prevent its occurrence.

[0277] The results of these simulations show that Modular Maps can be combined in a hierarchical and/or lateral configuration to good effect. It was also shown that to maximise the classification potential of Modular Map hierarchies all inputs to modules should be used. There are a variety of possible approaches to maximising inputs and in this case a 'data overlap' approach was used to maximise the limited training data available and thereby improve classification results.

[0278] It was also found that the Modular Map approach to classification of this face data offers slightly better clas-

sification than the traditional SOM (see the summary error rates table 1). In addition, the clustering on the surface of output modules was improved over that found on the SOM as can be seen from the activation maps presented in appendix A. When using a Modular Map hierarchy in configuration 4 (FIG. 23) the output module averaged 147 inactive neurons compared to 106 for the 256 neuron SOM, the reason being that the number of neurons active for individual classes is reduced (i.e. tighter clustering is found on the surface of the map). The clustering produced by the Modular Map systems is similar to that of the SOM, but was generally better defined. This can be seen when comparing the neural activations created by the same single class for the two systems, an example of which is presented in FIGS. 19a and 19b. This example corresponds to the activations for data class 3 in appendix A. These differences are due to the different architectures of the two systems. The SOM will only have a single reference vector (containing 99 elements in this case) while a Modular Map hierarchy results in reference vectors for the output neurons being constructed from a number of reference vectors from lower levels in the hierarchy (effectively providing 127 elements here). Because the reference vectors of the output layer of a Modular Map hierarchy are constructed from several lower level reference vectors it is possible to represent complex regions of the feature space with few neurons at the output.

[0279] The Modular Map solution to the face recognition problem requires more neurons than does the SOM solution, but the RISC neurons used by Modular Maps are much simpler which will result in a much reduced resource requirement when implemented in hardware as intended. It is the architecture of the Modular Map approach that has resulted in better classification rather than the number of neurons. This is emphasised by the failure of the SOM to improve over the previously stated classification results when network size is increased beyond 256 neurons. When a SOM containing 1024 neurons was trained on the same data detailed above for the face recognition problem, the classification of this data still resulted in a 6% error for the test data. Simulations were also carried out to check that the 'data overlap' approach used for the Modular Map hierarchy shown in configuration 4 (FIG. 23) was not giving the Modular Map solution an unfair advantage. These simulations used the same data as had been used for the Modular Map configuration except that the separate input vectors for modules were joined together to form 127 element vectors (i.e. $7 \times 16 + 1 \times 15$ vector elements). When a 256 neuron SOM was trained using these 127 element vectors equivalent to the 'data overlap' used for configuration 4 (FIG. 23), the classification results did not improve, but resulted in an additional 1% error compared to simulations using the 99 element vectors, i.e. classification error was 7% (see the summary error table 1).

[0280] In addition, the eigenface data used in the above face recognition were derived using Principal Component Analysis (PCA) which reduced the dimensionality of the original pictures by transforming the original variables into a new set of variables (the principal components) in a way that retains most of the variation present in the original data. The principal components are ordered so that the first few dimensions retain most of the variation present in all of the original variables. The data presented to the modular map array maintained this order such that module 0 in a hierarchy had the first few dimensions and the highest indexed module

on the lowest level had the last few dimensions etc. While the error rates of modules on the lowest layer in a hierarchy do not show a monotonic increase in error rate with increasing index, the general trend shows that error rates increase as the PCA components show decreasing variance.

[0281] When combining Modular Maps in hierarchical configurations, the error rates at the output network were less than those found for any modules at lower levels in the hierarchy (see tables 2, 3 and 4). Both classification and clustering improve moving up through subsequent layers in a Modular Map hierarchy as though higher layers in the hierarchy were performing some higher level functionality.

Ground Anchorage Integrity Testing

[0282] The Ground Anchorage Integrity Testing System (GRANIT) is being developed as a joint project between the Universities of Aberdeen and Bradford in collaboration with AMEC Civil Engineering Ltd. This work is built on the research of Prof. A. A. Rodger and Prof. G. S. Littlejohn into the effects of close proximity blasting to rock bolt behaviour.

[0283] As part of this development process, field trials were carried out at the Adlington site of AMEC Civil Engineering Ltd. Two test ground anchorages were installed by AMEC Civil Engineering Ltd for the purpose of these trials. The analysis pertains to a single strand anchor which has a diameter of 15.2 mm, a total length of 10 m and a bond length of 2 m. The drilling records for this anchorage show that the soil composition was weathered sandstone between 5 m and 5.8 m with strong sandstone between 5.8 m and 9.95 m. Using a pneumatic impact device to apply an impulse vibration was initiated within the anchorage system. An accelerometer affixed to the anchorage strand was then used to detect vibrations within the system.

[0284] The accelerometer output was fed, via a charge amplifier, to a notebook PC where the signals were sampled at 40 kSamples/Sec by a National Instruments DAQ 700 data acquisition card controlled by the GRANIT software developed at the University of Aberdeen. This software was developed using National Instruments Labwindows/CVI and the C programming language. The intricacies of data sampling and signal pre-processing are handled by the DAQ 700 software and Labwindows. However, laboratory tests using known signals were carried out to check that signals were being captured and processed as expected and no problems were identified.

[0285] Data was gathered for five pre-stress levels of the ground anchorage system; four of these levels were known to be 10 kN, 20 kN, 30 kN and 40 kN values, while the fifth level was initially unknown and used as a blind test to evaluate the potential predictive capacity of the GRANIT system. After results of the data analysis were presented to AMEC Civil Engineering the pre-stress value of the anchorage when the blind data were generated was revealed to be approximately 18 kN. Fifty (50) waveforms containing 512 samples were taken at each level. Throughout this evaluation process the blind test data were used only as a check; they were not taken into account when determining statistics of the main data set etc.

[0286] The time domain signals generated by the ground anchorage approximate a damped impulse response (see FIGS. 24a to 24e) and the envelope of these signals often provides an indication of the pre-stress level of the anchor-

age. FIGS. 24a to 24e show the average time domain signals for the 10 kN, 20 kN, 30 kN, 40 kN and blind tests respectively. However, the power spectra of these signals provides a better insight into varying pre-stress levels, and offers a significant compression of the data by transforming the original 512 dimensional time domain signals into their frequency components which, in this instance, resulted in 64 components. A 5th order Butterworth low pass filter with a threshold of 5 kHz was used to remove unwanted high frequency components. The power spectrum of these signals provides the average frequency components over the entire signal and shows that power spectra vary for varying pre-stress levels in the ground anchorage. Manual comparison of the power spectra can be difficult, but can be used to provide an approximation of pre-stress levels (see FIGS. 25a to 25e). FIGS. 25a to 25e show the average power spectrum for the 10 kN, 20 kN, 30 kN, 40 kN and blind tests respectively. Analysis utilising wavelet transforms could be used to provide a more detailed time-frequency analysis but the power spectra data offers considerable compression over the original input data and provided sufficient information for this analysis.

Classification of Ground Anchorage Pre-Stress Levels Using the Self-Organising Map

[0287] A 64 neuron SOM was trained using the 64 dimensional power spectra derived from response signals of the ground anchorage generated at known pre-stress levels. The activation map was then derived after training was complete by feeding test data to the network and noting which neuron was active for which class of data. However, this labelling process can be time consuming when carried out manually so a small utility program was developed which takes the output from the network and calculates the activation map automatically by correlating the original class of inputs with the resultant neuron activation. Once the activations on the surface of the map had been determined, the blind data set was fed to the SOM and the resultant activations were recorded and can be seen in FIG. 26. All 50 samples gathered during the blind field test caused the activation of neurons associated with the 20 kN data class.

[0288] The grouping of activations (clustering) on the surface of the SOM does not show a gradual transition from low to high pre-stress levels moving across the surface of the map (see FIG. 26). However, in most cases, there is a clear distinction between activations for different pre-stress levels, with very few neurons being active for two or more pre-stress values. There are regions of activation on the surface of the map that can be assigned to known pre-stress values of the anchorage but no individual pre-stress level has a single distinctive cluster of activations. There are several reasons for this, one of which is that data sets were not as consistent as would have been desired, especially the 30 and 40 kN cases. One factor that is responsible for these inconsistencies is that the impact applied to the anchorage varied slightly throughout the testing period. However, the activation map created from this data (FIG. 26) shows that the active neurons for the blind data set correspond to neurons which were active for the 20 kN data set. Consequently, it can be stated that the closest matching pre-stress value to the blind data set is 20 kN.

Classification of Ground Anchorage Pre-Stress Levels Using Modular Maps

[0289] A simple Modular Map configuration was used with the ground anchorage data detailed above to show that Modular Map hierarchies give improvements in classification and clustering moving up the hierarchy. A total of five modules were employed in a hierarchical configuration as shown in FIG. 27. As the data consisted of 64 dimensional vectors, each of the original vectors were partitioned into four separate vectors of 16 elements. The data were also scaled and quantised to fulfil the input requirements of Modular Maps but, in order to keep the configuration as simple as possible no attempts were made to create an optimal solution to the ground anchorage integrity testing problem and no data overlapping was used.

[0290] When the Modular Map system was trained on the same power spectra data of ground anchorage response signals as the SOM (see FIGS. 25a to 25e), the resultant activation maps for modules at the base of the hierarchy show poor classification and clustering of the blind data set (see FIGS. 28 to 31). The unknown pre-stress value could not be determined correctly from any individual one of these activation maps and, it is also unlikely that it could be identified by manual inspection of any combination of lower level maps.

[0291] However, all 50 samples of the blind test data set caused the activation of neurons associated with the 20 kN data on the output module of the hierarchy, as had occurred with the SOM (see FIG. 32) showing that classification does indeed improve moving up through a modular map hierarchy.

[0292] In addition, identification of each data class required fewer neurons in the output module of the hierarchy than had been required for the SOM. Instead of the three neurons that were active for the 20 kN data on the SOM (see FIG. 26). This class of data only resulted in two active neurons for the Modular Map. As the Modular Map system had fewer active neurons for each data class than did the SOM, there were 24 inactive neurons and, consequently, a 40 neuron module could have been used in place of the 64 neuron module. This effect was also found to increase as the depth of hierarchy increases such that the disparity between the number of neurons required by the SOM and the output module of a hierarchy increases with increasing depth of hierarchy. There are still similarities between the activations formed by the SOM and Modular Map for this data, with each class accounting for approximately the same percentage of activations for both systems, suggesting that the essential features of the data have been maintained. Overall the Modular Map also has fewer clusters (regions of activation) per class, than does the SOM, thereby reducing the disjoint nature of activation sets. For example, on the SOM the 30 kN case has three separate clusters and the 40 kN case has four separate clusters but, the Modular Map has two and three clusters for this data respectively.

[0293] The Modular Map approach to face recognition results in a hierarchical modular architecture which utilises

a 'data overlap' approach to data partitioning. When compared to the SOM solution for the face recognition problem, Modular Maps offer better classification results. This improvement in classification is achieved because a modular architecture is used. Modular Maps provide the basic building block for modular architectures and can be combined both laterally and hierarchically to good effect as has been shown.

[0294] When hierarchical configurations of Modular Maps are created the classification at the output layer offers an improvement over that of the SOM because the clusters of activations are more compact and better defined for modular hierarchies. This clustering and classification improves moving up through successive layers in a modular hierarchy such that higher layers, i.e. layers closer to the output, effectively perform higher, or more complex, functionality.

[0295] Application solutions using a modular approach based on the Modular Map will result in more neurons being used than would be required for the standard SOM. However, the RISC neurons used by Modular Maps require considerably less resources than the more complex neurons used by the SOM. The Modular Map approach is also scaleable such that arbitrary sized networks can be created whereas many factors impose limitations on the size of monolithic neural networks. In addition, as the number of neurons in a modular hierarchy increases, so does the parallelism of the system such that an increase in workload is met by an increase in resources to do the work. Consequently, network training time will be kept to a minimum and this will be less than would be required by the equivalent SOM solution, with the savings in training time for the Modular Map increasing with increasing workload.

[0296] Modifications and improvements may be made to the foregoing without departing from the scope of the present invention. Although the above description describes the preferred forms of the invention as implemented in special hardware, the invention is not limited to such forms. The modular map and hierarchical structure can equally be implemented in software, as by a software emulation of the circuits described above.

Appendix A

Sample Activation Maps

[0297] The activation maps presented in this appendix were derived from the application of human face recognition detailed in chapter 7. This application had 27 separate classes, i.e. there were pictures of 27 humans. Each square on the activation map represents a single neuron. When a neuron has activations for a particular class, the class number is denoted. Where no class number is denoted the neuron is not associated with any class, i.e. it has no activations.

4			15	15	15	11	16	16						13	13
4	4		15			11	11	16						13	13
4	6	6					23	23	10	10				13	
4		6	6					23		10	10			12	21
	6	6		9	9	9		23	10		12			12	21
5		19	19			9			2	2		12			21
5	5	5		19	25		25		2					21	21
20	20					25		12	12	2	2	7	7	7	
1	18		18			15			12				7	26	
1	1	18		18	18	15		14	14		14			26	
27	27	1	1	18		15		14	14	14	14			26	19
27	27	18	18	16	16	11							26	26	19
20	22	22	22	16			11		17					26	19
20		22	3		11	11		17					8		24
		9		3				17	17	8			8	7	24
9	9	9		3	3	3	17	17			8	8	7		24

Figure A.1: Example activation map for a 256 neuron SOM trained on eigenface data

24			3				21		21			4	4		
			3	3				21			4	4			
						7		7			4				
26				11			7		7						6
	26		11		11		17	17			8				6
19	26	26		12		13	13		17			8		6	
19			14	12	13	13		17			8			6	6
		14	14			13		17		12					
		14	14						12		12	15			
	22	22	22					5	5	5	15	15			
								11		5		15			
1			25	27	27	27	27		11			2	20		
1			25	9			19	23			2	2	2	20	20
	1	1	18		9		19	23	23				2	20	
		18	18	18	9	9		16		23	10	10			
	18		18					16	16		10	10	10		

Figure A.2: Example activation map for a Modular Map Hierarchy (Configuration 4) trained on eigenface data

1-23. (canceled)

24. A neural network module (300) comprising an array of neural processing elements (100) and at least one neural network controller (200), the neural processing elements comprising:

arithmetic logic means (50);

an arithmetic shifter mechanism (52);

data multiplexing means (115,125);

memory means (56,57,58,59);

data input means (110) including at least one input port;

data output means (120) including at least one output port;
and

control logic means (54);

and the controller (200) comprising

control logic means (270,280);

data input means (60) including at least one input port;

data output means (62) having at least one output port;

data multiplexing means (290,292,294);

memory means (64,68,280);

an address map (66); and

at least one handshake mechanism (210,220,230);

characterized in that the controller (200) is adapted to perform computations on data incoming to and outgoing from the neural processing elements.

25. The neural network module as claimed in claim 24 wherein the controller is further adapted to provide addressed and non-addressed instructions to the neural processing elements.

26. The neural network module as claimed in claim 24 wherein the memory means of the controller includes programmable memory means.

27. The neural network module as claimed in claim 24 wherein the memory means of the controller includes buffer memory associated with said data input means and/or said data output means.

28. The neural network module as claimed in claim 24 wherein the controller further comprises a collection of registers and a program counter.

29. The neural network module (300) as claimed in claim 24 wherein the number of processing elements (100) in the array is a power of two.

30. A modular neural network comprising:

one module (300) as claimed in claim 24, or at least two modules (300) as claimed in any of claims 24 to 29 coupled together.

31. The modular neural network as claimed in claim 30 further comprising arbitration logic adapted to ensure that during neural activity, only the index from a single module representing the active neuron is output to the processing elements.

32. The modular neural network as claimed in claim 31 wherein the arbitration logic is provided on each processing element.

33. The modular neural network as claimed in claim 31 wherein the arbitration logic is provided on each module.

34. The modular neural network as claimed in claim 31 wherein the arbitration logic comprises a binary tree.

35. The modular neural network as claimed in claim 31 wherein the arbitration logic provides a bus grant signal on the output of each processing element.

36. The modular neural network as claimed in claim 30 including synchronization means to facilitate data input to the neural network.

37. The modular neural network as claimed in claim 36, wherein said synchronization means enables data to be input only once when the modules (300) are coupled in hierarchical mode.

38. The modular neural network as claimed in claim 36 wherein the synchronization means is adapted to implement a two-line handshake mechanism.

39. A neural network device comprising a neural network as claimed in claim 30 wherein an array of processing elements (100) is implemented on the neural network device with at least one module controller (200).

40. The neural network device as claimed in claim 39, wherein the device is a field programmable gate array (FPGA) device.

41. The neural network device as claimed in claim 39, comprising one of the following: a full-custom very large scale integration (VLSI) device, a semi-custom VLSI device, or an application specific integrated circuit (ASIC) device.

42. A neural processing element (100) for use in a neural network, the processing element comprising:

arithmetic logic means (50);

an arithmetic shifter mechanism (52);

data multiplexing means (115,125);

memory means (56,57,58,59);

data input means (110) including at least one input port;

data output means (120) including at least one output port;
and

control logic means (54);

characterized in that the control logic means (54) is adapted to receive addressed and non-addressed instructions from a module controller.

43. The neural processing element (100) as claimed in claim 42, wherein each neural processing element (100) is a single neuron in the neural network.

44. The neural processing element as claimed in claim 42 further comprising data bit-size indicating means for enabling operations on different bit-size data values to be executed using the same instruction set.

45. A neural network controller (200) for controlling the operation of at least one neural processing element (100) as claimed in any of claims 42 to 44, the controller (200) comprising:

control logic means (270,280);

data input means (60) including at least one input port;

data output means (62) having at least one output port;

data multiplexing means (290,292,294);

memory means (64,68,280);

an address map (66); and

at least one handshake mechanism (210,220,230);

characterized in that the controller (200) is adapted to provide addressed and non-addressed instructions to a neural processing element.

46. The neural network controller (200) as claimed in claim 45 wherein the memory means includes programmable memory means.

47. The neural network controller (200) as claimed in claim 45 wherein the memory means includes buffer memory associated with said data input means and/or said data output means.

48. A neural network module (300) comprising an array of neural processing elements (100) as claimed in claim 42; and at least one neural network controller (200) as claimed in claim 45.

49. A modular neural network comprising:

at least one module (300) comprising an array of neural processing elements (100) the neural processing elements comprising

arithmetic logic means (50);

an arithmetic shifter mechanism (52);

data multiplexing means (115,125);

memory means (56,57,58,59);

data input means (110) including at least one input port;

data output means (120) including at least one output port; and

control logic means (54);

characterized in that the module further comprising arbitration logic adapted to ensure that during neural activity, only the index from a single module representing the active neuron is output to the processing elements.

50. The modular neural network as claimed in claim 49 wherein the arbitration logic is provided on each processing element.

51. The modular neural network as claimed in claim 49 wherein the arbitration logic is provided on each module.

52. The modular neural network as claimed in claim 49 wherein the arbitration logic comprises a binary tree.

53. The modular neural network as claimed in claim 49 wherein the arbitration logic provides a bus grant signal on the output of each processing element.

54. A computer program which upon execution on a computer constitutes together with the computer upon which it is executed an apparatus according to claim 24.

55. A method of classifying data comprising:

using the apparatus of claim 24.

* * * * *