



(51) International Patent Classification:

G06N 3/063 (2006.01)

(21) International Application Number:

PCT/GB2017/051734

(22) International Filing Date:

14 June 2017 (14.06.2017)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

1611857.2

07 July 2016 (07.07.2016)

GB

(71) Applicant: **ARM LTD** [GB/GB]; 110 Fulbourn Road,
Cambridge CB1 9NJ (GB).

(72) Inventors: **DAS, Shidhartha**; 110 Fulbourn Road, Cam-
bridge CB1 9NJ (GB). **HOLM, Rune**; 110 Fulbourn Road,
Cambridge CB1 9NJ (GB).

(74) Agent: **TLIP LTD**; Leeds Innovation Centre, 103 Claren-
don Road, Leeds Yorkshire LS2 9DF (GB).

(81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,
AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ,
CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO,
DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN,
HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP,
KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME,
MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ,
OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA,
SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN,
TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every
kind of regional protection available): ARIPO (BW, GH,
GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ,
UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ,
TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK,
EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV,
MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,
TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,
KM, ML, MR, NE, SN, TD, TG).

(54) Title: AN ARTIFICIAL NEURAL NETWORK




FIGURE 2

(57) Abstract: P00271WO0 43 An artificial neural network ABSTRACT Broadly speaking, embodiments of the present techniques provide a reconfigurable hardware-based artificial neural network, wherein weights for each neural network node of the artificial neural network are obtained via training performed external to the neural network. [FIGURE 2]

Published:

— *with international search report (Art. 21(3))*

An artificial neural network

The present invention generally relates to artificial neural networks, and in particular to a non-volatile memory circuit for implementing an artificial neural network.

Non-volatile memory (NVM) is a class of memory in which the memory cell or element does not lose its state after power supplied to the cell/element is removed. Examples of non-volatile memory devices include flash memory, magnetic computer storage devices (e.g. hard disks, floppy disks, and magnetic tape), and optical discs. Other examples of NVM devices include phase change memory (PCM/PCRAM) devices, resistive random access memory (RRAM/ReRAM) devices, and Correlated Electron Material (CEM/CeRAM) devices. In a phase change memory device, programming (writing) involves heating the memory cell to melt it, (which causes a change in resistance) and cooling to either a conductive crystalline state or a non-conductive amorphous state. In a resistive RAM device, programming involves changing the resistance across a dielectric solid-state material. A CeRAM device is formed (wholly or in part) from a Correlated Electron Material (CEM), which exhibits an abrupt conductive or insulative state transition arising from electron correlations rather than solid state structural phase changes.

NVM elements may be used for neuromorphic computing.

Broadly speaking, the techniques described herein implement a dot-product computation of weights and input values in an artificial neural network (ANN). An ANN is formed of layers of interconnected artificial neurons (or neural network nodes), where the connections (or edges) are associated with a number/weight. The term "dot-product" used herein refers to the multiplication of a weight associated with each in-edge of, or connection into, a neural network node with the magnitude of the signal (e.g. voltage, current, etc.) along that in-edge. In some implementations, the weights and/or inputs may have negative magnitudes. The computation performed at each neural network node may occur in two distinct steps. The first step comprises summing the products of the weights and the input magnitudes (dot-product or the multiply-accumulate). The second step comprises using this sum as an input to an activation function of the neural network node.

The activation function defines the output of the node given particular inputs, and could be binary (e.g. "fire" or "don't fire"). The present techniques provide hardware/circuitry for a low-power and energy-efficient implementation of the dot-product (multiply-and-accumulate) computation within an ANN.

5

According to a first aspect of the present technique, there is provided an apparatus for implementing a reconfigurable artificial neural network, the apparatus comprising: a crosspoint array comprising: at least one row signal line and at least one column signal line; at least one neural network node, wherein
10 each neural network node is provided at an intersection of a row signal line and a column signal line, and wherein each neural network node has an associated weight; and a plurality of programmable non-volatile memory elements provided at each neural network node to store the associated weight, wherein each non-volatile memory element is switchable between a first impedance state and a
15 second impedance state; a communication module configured to receive data specifying the weight for each neural network node of the crosspoint array; and control circuitry configured to switch the plurality of non-volatile memory elements provided at each neural network node into the first or the second impedance state to implement (or program) the weight at the neural network node.

20

According to a second aspect of the present technique, there is provided method of implementing a reconfigurable artificial neural network in a crosspoint array comprising at least one neural network node located at an intersection of a row signal line and a column signal line, the method comprising: receiving data
25 specifying a weight for each neural network node of the crosspoint array; switching a plurality of non-volatile memory elements provided at each neural network node into a required impedance state to implement the specified weight at each neural network node; and storing, using the non-volatile memory elements, the weight for each neural network node of the crosspoint array.

30

According to a third aspect of the present technique, there is provided a system of implementing a reconfigurable artificial neural network, the system comprising: an apparatus as described herein; and a remote server configured to train an artificial neural network model; wherein the apparatus is configured to

receive data specifying a weight for each neural network node of the apparatus from the remote server.

5 The techniques are diagrammatically illustrated, by way of example, in the accompanying drawings, in which:

Figure 1a shows a schematic of an example neural network for image recognition;

10 Figure 1b shows a schematic diagram of a neural network inference function;

Figure 2 shows a block diagram of a system to implement a neural network in an apparatus;
15

Figure 3 shows a circuit diagram of a neural network implemented using a crosspoint array, where each neural network node comprises a crosspoint array module formed of a plurality of non-volatile memory elements;

20 Figure 4a shows a circuit diagram of an example 2x2 crosspoint array module provided at each neural network node;

Figure 4b shows an equivalent circuit representation of the 2x2 crosspoint array module of Figure 4a;
25

Figure 5a shows a circuit diagram of an example 4x4 crosspoint array module provided at each neural network node;

30 Figure 5b shows an expanded representation of the 4x4 crosspoint array module of Figure 5b;

Figure 6 shows a circuit diagram of an example 4x4 crosspoint array module comprising write circuitry;

Figure 7 shows a circuit diagram of an example 3x3 crosspoint array, where each neural network node comprises a crosspoint array module formed of parallel non-volatile memory elements;

5 Figure 8 is a circuit diagram illustrating how negative weights are applied in the neural network;

 Figure 9 is a flow diagram illustrating example steps to implement (or apply/program) weights in a neural network formed of non-volatile memory
10 elements arranged in series;

 Figure 10 is a flow diagram illustrating example steps to implement (or apply/program) weights in a neural network formed of non-volatile memory elements arranged in parallel; and
15

 Figure 11a shows a circuit diagram of a tiered artificial neural network implemented using crosspoint arrays, and Figure 11b shows a schematic version of Figure 11a.

20 Broadly speaking, embodiments of the present techniques provide a reconfigurable hardware-based artificial neural network, wherein weights for each neural network node of the artificial neural network are obtained via training performed external to the neural network. The present techniques may be particularly useful for implementing artificial neural networks in resource-
25 constrained devices, such as Internet of Things (IoT) devices or embedded devices which typically have limited processing and/or memory capabilities. Such constrained devices may not have the capabilities to train an artificial neural network. Thus, the present techniques provide ways of implementing an artificial neural network in a device (particularly constrained devices) without requiring the
30 device to perform the training in situ. Weights for the artificial neural network are obtained by training a neural network remote to the device, such as in the cloud or a remote server having sufficient processing capabilities.

 The concept of artificial neural networks is not described in detail herein.
35 Nevertheless, a brief explanation is provided to aid understanding of the present

techniques. Generally speaking, an artificial neural network (ANN) is a processing device (implemented in software, hardware, or a combination of both) that contains hundreds, thousands or even millions of processing units called "artificial neurons" (which are like the neurons in a human brain). The artificial neurons work in unison to solve specific problems. An ANN is typically organised in layers that comprise a number of interconnected artificial neurons. A basic ANN consists of three layers: a layer of "input" units that are connected to a layer of "hidden" units, where the hidden nodes are connected to a layer of "output" units. (In some cases, the layer of hidden units may not be present). Figure 1a shows a schematic of an example artificial neural network. The input units of the input layer receive raw information/data that is to be processed by the ANN (or to be learnt/recognised during training of an ANN). The output units in the output layer are located on the opposite side of the ANN and signal how the ANN has processed the raw information. In between the input units and the output units are one or more layers of hidden units, which perform the majority of an ANN. (In Figure 1a, two hidden layers are shown). In a fully connected ANN, each hidden unit and each output unit is connected to every unit in the layers either side, but this is merely one example type of ANN. The connections between one unit and another unit are represented by a number/weight. The weight may be positive (e.g. if one unit excites another unit), or negative (e.g. if one unit suppresses/inhibits another unit). Generally, the higher the weight, the more influence one unit has on another.

ANNs are often used for image and speech recognition tasks which are usually difficult tasks to solve using ordinary rule-based programming. In the example ANN shown in Figure 1a, the ANN is used to perform image recognition, and in particular, to determine whether or not an image is of a dog. Thus, the output layer provides an output indicating whether or not a received image is (or contains) an image of a dog. The output may be a binary output e.g. "yes" is 1 and "no" is 0, or the output may indicate how likely it is that the received image is an image of dog. In the example of Figure 1a, the input layer receives an image of a cat, and this is processed by the ANN to determine if the image is of a dog. ANNs may be used for alphanumeric character recognition, for example, and there may be an output for each possible character (e.g. a to z, and 0 to 9).

In embodiments of the present techniques, the artificial neural network may be configured to receive voltages at the input layer, and to output currents at the output layer. In Figure 1a, the graphs (connections) in the ANN represent weights and the nodes represent the sum of the weighted inputs passed through a sigmoid function. The weighted sum of input nodes is a linear operation whereas the sigmoid function is essentially non-linear. Each node of an artificial neural network has in-edges (i.e edges pointing into the node) and out-edges (i.e. edges pointing of the node). Each edge pointing into a node comes either from another node, or from an input into the artificial neural network. Similarly, each artificial neural network (ANN), or each layer within a tiered ANN, has an in-edge and an out edge. Figure 1b shows a schematic diagram of a neural network inference function. In the example shown in Figures 1a and 1b, both the input and the output nodes can be thought of data values upon which the following computation is undertaken (see Equation 1):

$$z = \sum_{i=1}^n w_i x_i$$

$$y = H(z)$$

Equation 1

In matrix formulation, the linear computation of the above computation (Equation 1) can be written as follows:

$$Z = W^T X$$

Equation 2

In the present techniques, if the input nodes (X) of an ANN are expressed in terms of voltages V and the output nodes (Z) are expressed in terms of currents I, then the matrix W is equivalent to an admittance matrix. As such to relate output current and input voltage matrices, it is possible to rewrite Equation 2 as:

$$I = Y^T V$$

Equation 3

Memristor or ReRAM-based crosspoint array architectures may be used for neuromorphic computing. Generally speaking, the architectures may be used to convert input nodes representing voltages into output nodes representing currents, and/or to convert input nodes representing currents into output nodes representing voltages. In the following, it will be understood that reference to converting current to voltage (or vice versa) is merely used as an illustrative example, and that the present techniques may equally be used to convert voltage to current (or vice versa). The output currents are usually determined by passing the input voltages through a conductance matrix implemented in the crosspoint array. The conductance matrix is implemented by taking advantage of the "analogue" nature of particular non-volatile memory elements such as memristors and ReRAM. Analogue non-volatile memory elements are filamentary in nature. Consequently, depending upon the application of a suitable voltage, the actual resistance value provided by an analogue non-volatile memory element can be controlled in fine steps between a high resistance state (HRS) or a low resistance state (LRS). However, there are usually inherent variations in the resistance values of these filamentary non-volatile memory elements. Thus, crosspoint arrays formed from filamentary non-volatile memory elements tend to require careful and precise tuning *in situ*, on an element-by-element basis, to minimise any error between the target admittance and a measured admittance. Hence, this approach is inherently "analogue" by nature.

Accordingly, ANNs which are formed of crosspoint array architectures comprising filamentary resistive elements cannot take advantage of cloud-based or remote neural network training, because these ANNs necessarily have to be trained *in situ* to take into account variations between each resistive element. Thus, they are unsuitable for use in Internet of Things (IoT) devices where the constrained processing, energy and memory capacity of the devices make *in situ* training difficult or impossible. Furthermore, these ANNs may not be suitable for re-use for different applications (e.g. character recognition, speech recognition, etc.) since each application will require its own specific weights, and switching the filamentary, analogue ANN from one application to another would require the ANN to be retrained and tuned *in situ*.

Due to their analogue nature, such approaches are susceptible to variations in the resistance values. The variations in the intermediate resistance values may be even greater than the overall difference between the HRS and the LRS.

5 The artificial neural network of the present techniques is formed using a plurality of non-filamentary non-volatile memory elements to store weights. The non-filamentary non-volatile memory elements overcome the problems described above and may enable artificial neural networks to be implemented in constrained devices such as IoT devices without requiring in situ training. This may enable
10 artificial neural networks to be implemented in IoT devices quickly and efficiently, as each IoT device may simply download weights for a specific application. The crosspoint arrays of the present techniques which are used to encode weights and perform a dot-product computation are all-digital, which enables the ANN to take advantage of cloud-based training. (However, as explained below, the ANN may
15 comprise circuitry to compute an activation function, and in embodiments, this circuitry may not be all-digital). Broadly speaking, data (e.g. digital words) specifying weights for a specific application may be downloaded by an IoT device and programmed into a crosspoint array provided within the IoT device. This may eliminate the need for expensive and time-consuming ANN training in each IoT
20 device. The all-digital approach may also provide greater immunity against resistance variability and hence, provides an ANN suitable for realistic applications.

 In embodiments of the present techniques, each artificial neural network node of the crosspoint array may comprise a correlated electron switch (CES)
25 element. A CES element is uniquely differentiated from memristors and ReRAMs due to its "non-filamentary" nature, which allows the existence of (at least) two distinct and fixed impedance states after fabrication. Since a CES element is restricted to distinct impedance states, replacing the memristors and ReRAMs with CES elements in known crosspoint array based ANNs will not give rise to the
30 advantages of the present techniques, because these known ANNs utilize the filamentary nature of memristors to provide a variety of resistance states (and thereby, a variety of weights). In the present techniques, a spectrum of resistance/impedance states may be created through the usage of multiple CES elements.

35

As mentioned above, each neural network node of the crosspoint array of the present techniques may comprise a correlated electron switch (CES) element. In this context, a CES element may exhibit an abrupt conductor/insulator transition from a conductive state (a low impedance and low capacitance state) to an insulative (high impedance and high capacitance state) state, or vice versa, arising from electron correlations in the CEM rather than solid state structural phase changes (e.g., crystalline/amorphous changes as occurs in phase change memory (PCM) devices or filamentary formation as occurs in resistive RAM devices).

In embodiments, the abrupt conductor/insulator transition in a CEM may be responsive to a quantum mechanical phenomenon, in contrast to melting/solidification or filament formation. Such a quantum mechanical transition between conductive and insulative states in a CES element may be understood in any one of several aspects.

In embodiments, the quantum mechanical transitions between an insulative and conductive states may be understood in terms of Mott transitions. For example, a CEM may switch from an insulative state to a conductive state if a first Mott transition condition is met and may switch from a conductive state to insulative state if a second Mott transition condition is met. The different Mott transition conditions will be described in further detail below.

In particular implementations of embodiments of the present techniques, a memory circuit may comprise: a memory array including at least one CES element; a write circuit for placing the CES element in a first memory state or a second memory state depending on signals provided to the memory array (or to a memory block within the memory array), wherein an impedance of the CES element is higher in the second memory state than in the first memory state; and a read circuit for sensing the memory state of the memory array (or block) and providing an electrical signal corresponding to the sensed memory state of the memory array (or block).

In one aspect, the impedance of the CES element in the second impedance state may be more than 100 times greater than the impedance of the CES element

in the first impedance state.

The CES element may function as a CES random access memory (CeRAM). The CES element comprises a material that may transition between or among a plurality of predetermined detectable memory states based, at least in part, on at least a portion of the CEM transitioning between memory states (e.g. from an insulative to a conductive state) utilizing, for example, a Mott transition. In this context, a "*memory state*" means a detectable state of the CES element that is indicative of a value, symbol, parameter or condition, just to provide a few examples.

In one particular implementation, a memory state of the CES element may be detected based, at least in part, on the value of a signal detected on one or more the terminals of the CES element in a "read operation".

In another particular implementation, the CES element may be placed in a particular memory state to represent or store a particular value, symbol, parameter or condition etc. by controlling the bias parameters (e.g. voltage) applied to the CES element and/or by controlling the current parameters (e.g. current density $J(\text{A}/\mu\text{m}^2)$ driven through the CES element.

In an embodiment, the CEM comprises a switching region (S) arranged in electrical contact with (e.g. sandwiched between) conductive regions (C), in a (C/S/C) arrangement on a semiconductor (e.g. in a stack arrangement). The conductive region (C) may comprise any suitable material, and may comprise a transition metal oxide, and may further include any material comprising metal-like characteristics, including, for example, polysilicon or a doped semiconductor. In embodiments, the conductive region (C) may additionally or alternatively comprise a material selected from a group comprising aluminium, cadmium, chromium, cobalt, copper, gold, iron, manganese, mercury, molybdenum, nickel, palladium, rhenium, ruthenium, silver, tin, titanium, vanadium, and zinc (which may be linked to a cation such as oxygen or other types of ligands), or combinations thereof.

According to an embodiment, the switching region (S) may comprise any

transition metal oxide (TMO) material, such as, for example, perovskites, Mott insulators, charge exchange insulators, and Anderson disorder insulators.

5 In particular implementations, the switching region (S) material may include switching materials such as nickel oxide, cobalt oxide, iron oxide, yttrium oxide, and perovskites such as Cr doped strontium titanate, lanthanum titanate, and the manganate family including praeysidium calcium, manganate, and praeysidium lanthanum manganite, just to provide a few examples. In particular, oxides incorporating elements with incomplete **d** and **f** orbital shells may exhibit
10 sufficient resistive switching properties for use as a switching region (S) material. In an embodiment, the CES element may be prepared without electroforming.

Other implementations may employ other transition metal compounds as the switching region (S) material without deviating from claimed subject matter.
15 For example, $\{M(\text{chxn})_2\text{Br}\}\text{Br}_2$ where M may comprise Pt, Pd, or Ni, and chxn comprises 1R, 2R-cyclohexanediamine, and other such metal complexes may be used without deviating from claimed subject matter.

In one aspect, the CEM may comprise materials that are transition metal
20 oxide (TMO) variable resistance materials, though it should be understood that these are exemplary only, and are not intended to limit claimed subject matter.

Particular implementations may employ other variable resistance materials as well. Whilst nickel oxide, NiO, is disclosed as one particular TMO above, NiO
25 materials may also be doped with extrinsic ligands, which may stabilize variable resistance properties. In particular, NiO variable resistance materials disclosed herein may include a carbon containing ligand, which may be indicated by NiO(Cx). Here, one skilled in the art may determine a value of x for any specific carbon containing ligand and any specific combination of carbon containing ligand with
30 NiO simply by balancing valences. In another particular example, NiO doped with extrinsic ligands may be expressed as NiO(Lx), where Lx is a ligand element or compound and x indicates a number of units of the ligand for one unit of NiO. One skilled in the art may determine a value of x for any specific ligand and any specific combination of ligand with NiO or any other transition metal simply by balancing
35 valences.

Furthermore, the terminals may comprise the same materials as the conductive regions (C), such that the conductive regions (C) may function as the terminals. In embodiments where the material for the terminals and the conductive regions (C) are not the same, it is preferable that the different materials are matched to reduce the effects of border interactions or surface defects which may otherwise affect performance of the CES element. For example, the materials should be matched to ensure the electrical properties (e.g. conductivities) and/or chemical properties (e.g. coefficients of thermal expansion match) of the CEM and terminals do not negatively impact the operation of the CES element.

By applying a specific voltage (V_{applied}) between the terminals and driving a specific current density (J_{applied}) through the CEM, the CES element may transition from a conductive (low impedance and low capacitance) memory state, hereinafter "LIS", to an insulative (high impedance and high capacitance) memory state, hereinafter "HIS", or may transition from a HIS to an LIS in accordance with a different applied voltage (V_{applied}) and current density (J_{applied}).

In this context, it should be understood that LIS and HIS are relative terms and not specific to any particular quantity or value for conductance, impedance, or capacitance. For example, whilst in a first memory state the CES element may be more conductive (or less insulative), than while the CES element is in a second memory state.

Furthermore, CES elements are "born on", whereby when fabricated (e.g. deposited), the CES element is in a LIS.

Turning now to Figure 2, this shows a block diagram of a system 100 to implement a neural network in an apparatus 102. As explained above, the artificial neural network is provided in apparatus 102 and is formed using at least one crosspoint array 104. The crosspoint array 104 comprises at least one row signal line and at least one column signal line, at least one neural network node (neuron) is provided at an intersection of a row signal line and column signal line. The neural network node has an associated weight, which is stored using a

plurality of programmable non-volatile memory elements provided at each neural network node. The weight for each neural network node is obtained via neural network training performed in a remote server 110. The weights may be provided to, or downloaded by, a communication module 110 of the apparatus 102. The apparatus 102 may use control circuitry 112 to determine how to program the received weights into the crosspoint array 104 in order to implement the artificial neural network for a particular application (e.g. speech recognition, character/image recognition, etc.) When the crosspoint array 104 has been programmed based on the received weights, an input 106 (e.g. voltages, currents, or other signals representing an image, sound, etc.) may be input into the apparatus 102. The input 106 is processed by the crosspoint array 104 and at least one output 108 is generated.

Accordingly, there is provided an apparatus for implementing a reconfigurable artificial neural network, the apparatus comprising: a crosspoint array comprising: at least one row signal line and at least one column signal line; at least one neural network node, wherein each neural network node is provided at an intersection of a row signal line and a column signal line, and wherein each neural network node has an associated weight; and a plurality of programmable non-volatile memory elements provided at each neural network node to store the associated weight, wherein each non-volatile memory element is switchable between a first impedance state and a second impedance state; a communication module configured to receive data specifying the weight for each neural network node of the crosspoint array; and control circuitry configured to switch the plurality of non-volatile memory elements provided at each neural network node into the first or the second impedance state to apply implement the weight to at the neural network node.

The crosspoint array may be used to perform the first step, i.e. summing the dot-product computation of weights and input values/signals, for each neural network node of an artificial neural network (ANN). In embodiments, the apparatus may comprise further circuitry to perform the second step, the computation of the activation function, which uses the sum from the first step as input into an activation function for the neural network node.

The term "non-volatile memory element" is used interchangeably herein with the terms "NVM element", "non-volatile memory device", "non-volatile random access memory", "non-volatile memory cell", "non-volatile memory array", or "non-volatile memory". In embodiments, the non-volatile memory element may be a non-polar NVM element, in which switching and conduction can take place in any direction (e.g. bipolar switching), and the term "non-volatile memory element" is used interchangeably herein with the term "non-polar NVM element". In embodiments, the non-volatile memory element may be a memristor, a resistive random access memory (ReRAM), a phase-change memory (PCM), and/or spin-transfer torque magnetic random-access memory (STT-MRAM or STT-RAM), and the term "non-volatile memory element" is used herein to refer to any of these types of memory. In embodiments, the non-volatile memory element is a correlated electron switch (CES), and the term "non-volatile memory element" is used interchangeably with the terms "Correlated Electron Material memory", "CeRAM", "Correlated Electron Switch", "CES", or "bipolar switching non-volatile memory element".

In embodiments, the communication module of the apparatus is configured to receive data specifying the weight for each neural network node and is configured to: obtain at least one data packet specifying the weights for each neural network node; and extract, from the data packet, the weight for each neural network node; and wherein the control circuitry is configured to: determine, using the weight for each neural network node, whether each non-volatile memory element at each neural network node is to be switched into the first impedance state or the second impedance state; and generate, responsive to the determination, a programming signal for each non-volatile memory element.

The control circuitry may comprise write circuitry configured to: receive the generated programming signal for each non-volatile memory element; and drive each non-volatile memory element dependent on the generated programming signal.

Figure 3 shows a circuit diagram of an artificial neural network (ANN) implemented using at least one crosspoint array 10. The crosspoint array 10 comprises at least one row signal line 18 and at least one column signal line 20.

In the illustrated example, the crosspoint array comprises a plurality of row signal lines 18 and a plurality of column signal lines 20. As mentioned earlier, an ANN comprises an input layer and therefore, the crosspoint array 10 comprises one or more input nodes 14. Each input node 14 is coupled to a row signal line 18. An ANN comprises an output layer and therefore, the crosspoint array comprises one or more output nodes 16, where each output node is coupled to a column signal line 18. In embodiments (now shown), the crosspoint array 10 may comprise a plurality of input nodes 14 and a single output node 16. The number of output nodes 16 may depend on the application being implemented by the ANN. For example, many output nodes 16 may be required for an ANN that is implementing alphanumeric character recognition (i.e. one for each character). In another example, a single output node 16 may be sufficient if the ANN is implementing a specific image recognition application, e.g. facial recognition.

The crosspoint array 10 comprises a neural network node 12 provided at each intersection of a row signal line 18 and a column signal line 20. In the illustrated example, the crosspoint array 10 is a MxN array and comprises a plurality of neural network nodes 12, labelled R11 to Rmn. In embodiments, M and N are different integer numbers, and in alternative embodiments, M and N are equal. A plurality of programmable non-volatile memory elements is provided at each neural network node 12 to store a weight associated with the neural network node 12. Each non-volatile memory element is switchable between a first impedance state and a second impedance state, and which state is selected is dependent on the weight to be stored in the neural network node 12. In the case of the non-volatile memory element being a correlated electron switch, there may, in embodiments, be only two possible impedance states: a high impedance state (HIS) or a low impedance state (LIS). Each neural network node 12 comprises a neural network node input 22 (also referred to herein as an "input edge" or "in-edge") and a neural network node output 24 (also referred to herein as an "output edge" or "out-edge").

Thus, in embodiments the crosspoint array may further comprise: at least one input terminal to receive an input signal for the crosspoint array; and at least one output terminal to provide an output signal generated by passing the input signal through at least one neural network node of the crosspoint array. The

neural network node may receive the input signal, multiply the associated weight of the neural network node with the input signal, and output a weighted signal.

Alternatively, the crosspoint array may comprise: a plurality of row signal
5 lines and a plurality of column signal lines; a plurality of neural network nodes, wherein each neural network node is provided at an intersection of a row signal line and a column signal line, and wherein each neural network node has an associated weight; and a plurality of programmable non-volatile memory elements provided at each neural network node, wherein each non-volatile memory element
10 is switchable between a first impedance state and a second impedance state.

The crosspoint array may further comprise: a plurality of input terminals (or input nodes), each input terminal configured to receive an input signal for the crosspoint array; and at least one output terminal (or output node) to provide an
15 output signal generated by passing the input signals through the neural network nodes of the crosspoint array. The crosspoint array may comprise a plurality of output terminals, each output terminal providing an output signal.

Each neural network node may receive two or more input signals, multiply
20 the associated weight with each input signals to generate weighted signals, and output a summed signal representing the summed weighted signals. The multiplication may be a dot-product calculation of the weight and the input signal. The at least one output terminal of the crosspoint array may be configured to receive a summed signal from one or more of the plurality of neural network
25 nodes, and to combine the summed signals to provide the output signal.

The crosspoint array 10 may be configured to receive voltages (voltage values) at each input node 14, e.g. V_1 to V_m , and may be configured to output currents (current values) at each output node 16, e.g. I_1 to I_n . In alternative
30 embodiments, the crosspoint array 10 may be configured to receive currents and to output voltages.

In embodiments, each neural network node 12 of the crosspoint array 10 may comprise a crosspoint array module (not as shown in Figures 4a and 4b,3
35 and described in more detail below). The crosspoint array module comprises an

array of N row single lines and M column signal lines, with a non-volatile memory element provided at each intersection of the row and column signal lines. Within each crosspoint array module, there is a direct path from the neural network node input 22 (in-edge) to the neural network node output 24 (out edge), and at least one indirect path from the neural network node input 22 to the neural network node output 24. The direct path passes through a single, or a minimum number of, NVM elements or, in other words, through a selected cell/portion of a crosspoint array. The indirect path is a sneak leakage path through unselected cells or portions of a crosspoint array. In this embodiment, the direct and sneak leakage paths in the crosspoint array module are exploited to create a range of resistance (R) values that can be encoded digitally.

In embodiments, each neural network node may comprise a crosspoint array module, and the plurality of programmable non-volatile memory elements provided at each neural network node are arranged in the crosspoint array module.

Thus, in embodiments, the present techniques provide a "tiered", fully-connected crosspoint array architecture that may achieve a multi-bit digital encoding of a resistance matrix. The sneak leakage paths across the crosspoint array may provide a linear increase in equivalent resistance values that may enable an all-digital encoding of a resistance matrix. This embodiment is described in more detail with reference to Figures 4a to 7 below.

In alternative embodiments, each neural network node 12 of the crosspoint array 10 may comprise a crosspoint array module (not shown in Figure 3) which comprises two or more non-volatile memory elements arranged in a parallel connection between the neural network node input 22 to the neural network node output 24. In this embodiment, the parallel connection of non-volatile memory elements in each neural network node 12 is used to create a range of conductance (G) values that can be encoded digitally. Thus, in embodiments, the present techniques provide a multi-element parallel-connected crosspoint array, where the parallel connection of the non-volatile memory elements may provide an all-digital encoding of an admittance/conductance matrix. This embodiment is described in more detail with reference to Figures 8 and 9 below.

In embodiments, the weight associated with each neural network node is a negative weight. In embodiments the first impedance state is a high impedance state, and the second impedance state is a low impedance state.

5

The non-volatile memory element may be a bipolar switching non-volatile memory element. In embodiments, the non-volatile memory element is any one of: a memristor, a resistive random access memory (ReRAM), a phase-change memory (PCM), and a spin-transfer torque magnetic random-access memory (STT-MRAM). Preferably, the non-volatile memory element is a correlated electron switch (CES).

10

The communication module of the apparatus may be configured to receive data specifying the weight for each neural network node of the crosspoint array from an artificial neural network model trained external to the apparatus, e.g. from a remote server or the cloud.

15

As the ANN of the present techniques is reconfigurable for different applications, the communication module may be configured to receive further data specifying a further weight for each neural network node of the crosspoint array, wherein the further weight is obtained from a machine learning technique or statistical model trained external to the apparatus. The further weight may be obtained from a further ANN model trained external to the apparatus, or from other statistical models, such as a hidden Markov model (HMM).

20

25

The apparatus may comprise a data store coupled to the communication module and configured to store the received data. Storing the received data may enable the apparatus to switch between different applications more readily, as the weights for each ANN application may be retrievable from a local data store when required.

30

The apparatus may be configured to receive instructions indicating a particular artificial neural network is to be implemented in the apparatus. Responsive to the instructions, the control circuitry may be configured to: retrieve, from the data store, the data corresponding to the artificial neural network

35

indicated in the received instructions; and use the retrieved data to extract a weight for each neural network node of the crosspoint array.

All-Digital Multi-Bit Resistance Matrix

5

As mentioned above, each neural network node 12 of the crosspoint array 10 may comprise a crosspoint array module. Figure 4a shows a circuit diagram of an example crosspoint array module provided at each neural network node 12. In this embodiment, the crosspoint array architecture 10 is used to generate a linear scale of resistance values, by taking advantage of sneak leakage paths in the crosspoint array. Generating a linear scale of resistance values may enable the actual resistance values (or weightings) for each artificial neural network to be digitally encoded and programmed into the crosspoint array.

15

The crosspoint array module shown in Figure 4a comprises a 2x2 array of row signal lines 18' and column signal lines 20', with a non-volatile memory element 26a – 26d provided at each intersection of the row and column signal lines. The crosspoint array module shown in Figure 4a is merely exemplary, and in embodiments, may comprise more than two row signal lines and/or more than two column signal lines. In embodiments, the number of row signal lines and column signal lines may not be equal (i.e. the crosspoint array module may comprise a non-square array or row and column signal lines). In embodiments, the crosspoint array module may comprise an array of "n" row and "m" column signal lines, where n and m are integers, and where n and m may be equal or different. We refer to such an array as an nxm dimensional crosspoint array.

25

Within the crosspoint array module of neural network node 12, there is a direct path from the neural network node input 22 to the neural network node output 24. The direct path is through non-volatile memory element 26a. In the crosspoint array module there is at least one indirect path (a sneak leakage path) from the neural network node input 22 to the neural network node output 24. The indirect, sneak leakage path is through non-volatile memory element 26b, then non-volatile memory element 26c, and then through non-volatile memory element 26d. In other words, the sneak leakage path is through the series connection of non-volatile memory elements 26b, 26c and 26d. In this embodiment, the direct

35

and sneak leakage paths in the crosspoint array module are exploited to create a range of resistance (R) values that can be encoded digitally.

In Figure 4a, the equivalent resistance through the crosspoint array module is measured between the neural network node input 22 and the neural network node output 24. In this crosspoint array module, only input 22 and output 24 are connected to other parts of the crosspoint array 10, but other arrangements comprising more inputs and outputs may be possible.

Figure 4b shows an equivalent circuit representation of the 2x2 crosspoint array module of Figure 4a. The equivalent resistance between neural network node input 22 and neural network node output 24 is labeled as R_{eq} in Figure 4b and is essentially a tuple. If a non-volatile memory element (e.g. a CES element) is in the low resistance state (LRS, or LIS), its resistance is labeled as "R". If a non-volatile memory element is in the high resistance state (HRS, or HIS), the resistance of the conducting path through the crosspoint array module is at least two orders of magnitude higher than the LRS and is therefore approximated as "infinity", or effectively open circuit.

When the NVM element used is a CES (Correlated Electron Switch), the HRS and LRS are separated by at least two orders of magnitude. In embodiments where other types of NVM elements are used, the separation between HRS and LRS may not be as large. In such a case, a series transistor may also be used as a proxy for "open-circuit" resistance. In the following equations, the variable "R" is used to denote the LRS resistance.

$$R_{eq} = R \text{ (when } R_{01} = LRS, \text{ and any or all of } R_{00}, R_{10}, R_{11} = HRS) \quad (a)$$

$$R_{eq} = 3R \text{ (when } R_{01} = HRS \text{ and } R_{00} = R_{10} = R_{11} = LRS) \quad (b)$$

$$R_{eq} = \frac{3R}{4} \text{ (when } R_{01} = LRS \text{ and } R_{00} = R_{10} = R_{11} = LRS) \quad (c)$$

$$\text{Hence, } R_{eq} = \{R, 3R, \infty\}$$

Equation 4

In embodiments of the present techniques, the aim is to achieve a linear mapping between the resistance/impedance values and the digitally-encoded

representation. Hence, some embodiments discussed below do not make use of the fractional weights (i.e. as shown in Equation 4(c) above). However, in alternative embodiments, where the linearity requirement is less stringent, fractional weights may also be used. The present techniques exploit the resultant linearity as a means of design and usage simplification (and may, in embodiments, also use fractional weights).

Thus, the direct and the sneak resistance paths have enabled a linear resistance scale. The resistances of a 2x2 crosspoint array module can be digitally encoded with a single bit. (Note that in embodiments, the non-linear resistance separation – $0.75R$ – may be discounted, as it may not be encoded digitally. However, in embodiments, the non-linear resistance separation(s) may be utilized).

Larger crosspoint array modules may be used to provide larger/wider resistance ranges and thereby, a multi-bit resistance code. Figure 5 shows a circuit diagram of an example 4x4 crosspoint array module provided at each neural network node 12. The crosspoint array module shown in Figure 5 comprises a 4x4 array of row signal lines 18' and column signal lines 20' and 21', with a non-volatile memory element provided at each intersection of the row and column signal lines. The 4x4 crosspoint array module may be constructed using the 2x2 crosspoint array module of Figure 4a as a building block. The dashed boxes shown four 2x2 crosspoint array modules 50a to 50d. Each 2x2 building block 50a to 50d operates in a similar way to the 2x2 crosspoint array module described above with reference to Figures 4a and 4b. For example, in building block 50b, there is a direct path through building block 50b through non-volatile memory element R01, and an indirect path through non-volatile memory elements R00, R10 and R11.

There are two direct paths through the crosspoint array module, from neural network node input 22 to the neural network node output 24. One direct path through the 4x4 crosspoint array module is through non-volatile memory element R01 in building block 50a, and the other direct path is through the non-volatile memory elements R00, R10 and R11 in building block 50a. Thus, the direct path may have values R or $3R$. There are many indirect, sneak leakage paths through

the 4x4 crosspoint array module, depending on how many building blocks are passed through. For instance, 3R may be achieved by passing through R00, R10 and R11 of building block 50b only, and 5R may be achieved when the equivalent resistances through building block 50b is R, 50c is 3R and 50d is R. Accordingly, the 4x4 crosspoint array module may be able to provide the following equivalent resistances between neural network node input 22 and neural network node output 24:

$$Req = \{R, 3R, 5R, 7R, 9R, \infty\} \quad \text{Equation 5}$$

As before, here "R" indicates the value of LRS resistance in an NVM element.

Thus, the direct and the sneak resistance paths in the 4x4 crosspoint array module have enabled a linear resistance scale, which can be digitally encoded with three bits. (Note that in embodiments, the non-linear resistance separations may be discounted, as they may not be readily encoded in a digital form. However, in embodiments, the non-linear resistance separation(s) may be utilized).

An 8x8 crosspoint array module may be required to provide a fully populated 3-bit encoding space. In this case, the resistance scale will be composed of the following values (where "R" indicates the value of LRS resistance in an NVM element):

$$Req = \{R, 3R, 5R, 7R, 9R, 11R, 13R, 15R, \infty\} \quad \text{Equation 6}$$

The composite 4x4 crosspoint array module of Figure 5 may be constructed using two metal layers (i.e. one metal layer for the column signal lines, and one metal layer for the row signal lines) without requiring further layers of metal. As shown in Figure 5, the row signal lines 18' and even column signal lines 20' are continuous, while the odd column signal lines 21' are discontinuous. The discontinuous column signal lines 21' are used to enable the writing of each non-volatile memory element in the crosspoint array module.

In embodiments, the crosspoint array module of each neural network node may comprise an array formed of N row signal lines and M column signal lines,

wherein one of the plurality of non-volatile memory elements is provided at each intersection of the N row signal lines and M column signal lines. Each crosspoint array module may comprise: a direct path from a neural network node input to a neural network node output; and at least one indirect path from the neural network node input to the neural network node output.

In particular embodiments, the crosspoint array modules may be used to digitally encode a linear scale of impedance values in the crosspoint array.

The at least one indirect path between the neural network node input to the neural network node output may be formed of two or more non-volatile memory elements in a series connection.

The crosspoint array may be configured to convert voltage to current. In such embodiments, each input terminal of the crosspoint array may be configured to receive a voltage and each output terminal of the crosspoint array may be configured to output a current.

Figure 6 shows a circuit diagram of an example 4x4 crosspoint array module comprising control circuitry to drive each non-volatile memory element into the required impedance state (to store the required weights for the artificial neural network). The control circuitry may comprise means (e.g. transistors 54, 56) to select a column signal line and means (e.g. shorting transistors 52a – 52d) to couple column signal lines together.

Writing the non-volatile memory elements coupled to odd (discontinuous) columns 21' is achieved by first writing the non-volatile memory elements of the even columns 20' into the HRS (or HIS), which effectively means they are in an "open circuit" state. Secondly, an odd column 21' is shorted to a neighbouring even column 20' (via shorting transistors 52a to 52d), and the non-volatile memory elements of the odd columns 21' are written through the even columns 20'. The shorting transistors 52a to 52d may achieve this with minimum voltage degradation. Once the non-volatile memory elements coupled to the odd columns 21' have been written/programmed into the required impedance state, the odd columns 21' are disconnected from the even columns 20'. The non-volatile

memory elements coupled to the even columns 20' are written again, this time into the required impedance state. (A particular non-volatile memory element is selected by selecting the appropriate row and column signal lines).

5 Writing is typically achieved by pulling all the column signal lines 20', 21' low, and biasing the row signal lines 18' with either the "set" or the "reset" voltage. Accidental writes to the sneak leakage paths is prevented due to the compliance current requirement which is limited by having multiple CES elements in series.

10 The shorting transistors 52a to 52d may be routed using a third metal layer in the crosspoint array module, such that they may exist within a crosspoint array "shadow" and not consume additional area. Hence, the bulk of the area may be dominated by the crosspoint array. This arrangement achieves an extremely area-efficient implementation of a CES-based crosspoint array.

15

 In embodiments, each non-volatile memory element may be coupled to a switching transistor 58. In embodiments, instead of driving a non-volatile memory element into HRS (or HIS), the switching transistor 58 may be used to disconnect a non-volatile memory element from the crosspoint array module, i.e. to
20 deactivate a particular path through the crosspoint array module. However, this requires additional circuitry.

 Thus, in embodiments, each of the N row signal lines of each crosspoint array module is formed of a continuous conductive channel which runs across the
25 crosspoint array module; and each of the M column signal lines of each crosspoint array module comprises: a plurality of continuous column signal lines, wherein each continuous column signal line is formed of a continuous conductive channel which runs across the crosspoint array module; and a plurality of discontinuous column signal lines, wherein each discontinuous column signal line is formed of
30 multiple spatially separated conductive channels. The multiple conductive channels that form each discontinuous column signal line (e.g. column signal lines 21' in Figure 6) may be formed in a single metal layer.

 The plurality of continuous column signal lines may be interlaced with the
35 plurality of discontinuous column signal lines in an alternating arrangement.

Control circuitry may be configured to drive each non-volatile memory element of each crosspoint array module that is in electrical communication with a continuous column signal line by: selecting a continuous column signal line; and
5 driving each non-volatile memory element connected to the selected continuous line dependent on the generated programming signal.

Control circuitry may be configured to drive each non-volatile memory element that is in electrical communication with a discontinuous column signal line
10 by: writing each non-volatile memory connected to a neighbouring continuous column signal line into a first impedance state; connecting the discontinuous column signal line to the neighbouring continuous column signal line; driving each non-volatile memory element connected to the discontinuous line dependent on the generated programming signal; and disconnecting the discontinuous column
15 signal line from the neighbouring continuous column signal line.

Figure 7 is a flow diagram illustrating example steps to implement or program the weights into a neural network formed of non-volatile memory elements arranged in series. The process begins by receiving weights for the crosspoint array 104 from a remote server 110 (step S90). Control circuitry 112
20 may determine, using the received weights, which impedance state (LIS or HIS) each non-volatile memory (NVM) element in the crosspoint array 104 is required to be in to implement the received weights at each neural network node (step S92). This may be performed individually for each neural network node, i.e. for the crosspoint array module at a neural network node. At step S94, all the even
25 columns in a crosspoint array module (of a neural network node) are written to a high impedance state. Each odd column in a crosspoint array module is connected to a neighbouring even column to short the odd columns (step S96). A row and an odd column are selected (step S98) and the non-volatile memory element
30 coupled to the selected row and column is written/programmed into the required impedance state (step S100). The control circuitry 112 may check if all of the non-volatile memory elements coupled to odd columns have been written (step S102) and if not, the selection and writing process is repeated until all of these NVM elements have been written. The odd columns are then disconnected from
35 the neighbouring even columns (step S104). The NVM elements coupled to the

even columns may now be written, one at a time, by selecting a row and an even column (step S106). The selected NVM element is written into the required impedance state (step S108). The control circuitry 112 may check if all of the non-volatile memory elements coupled to even columns have been written (step S110) and if not, the selection and writing process is repeated until all of these NVM elements have been written. The process then ends. Note that while the even columns are being written to, the NVM elements in the odd columns are not perturbed since they are effectively disconnected from the write circuitry.

The above-described techniques can be used to provide larger crosspoint arrays to increase the resistance scale. For example, each neural network node 12 may comprise an 8x8 crosspoint array module, which may increase the resistance scale to an 8-level scale (i.e. $R_{eq} = R, 3R, 5R, 7R, 9R, 11R, 13R, 15R$), thereby achieving fully populated 3-bit encoding. Similarly, a 32x32 crosspoint array module may achieve a 16-level scale and a fully populated 4-bit encoding. Each resistance value in a resistance scale is equally separated from its neighbouring resistance values by $2R$. This may enable any effects arising from CES element process variations to be mitigated. A crosspoint array may be constructed of any size and for the required encoding simply by using crosspoint array modules at each neural network node of the appropriate size. For example, to implement a 256x256 artificial neural network, a 256x256 crosspoint array is required. If a 3-bit resistance encoding is required, each neural network node comprises an 8x8 crosspoint array module. If a neural network comprises 256 inputs and 256 outputs, the resultant neural network will have a total of 2^{16} (256x256) neural network nodes. Each neural network node (as mentioned before) requires an 8x8 crosspoint array module in order to provide 3-bit resistance encoding. Hence, a neural network implemented using a 256x256 crosspoint array, and which provides 3-bit encoding, effectively comprises an array of 1024x1024 NVM elements (i.e. 1MB of memory).

All-Digital Multi-Bit Resistance Matrix

As mentioned above with reference to Figure 3, each neural network node 12 of the crosspoint array 10 may comprise a crosspoint array module which comprises two or more non-volatile memory elements arranged in a parallel

connection between the neural network node input 22 to the neural network node output 24. Figure 8 shows a circuit diagram of an example 3x3 crosspoint array 60, where each neural network node 12 comprises a crosspoint array module formed of parallel non-volatile memory elements. The parallel connection of non-volatile memory elements in each neural network node 12 is used to create a range of conductance (G) values that can be encoded digitally.

Thus, in embodiments each neural network node comprises a crosspoint array module, and the plurality of programmable non-volatile memory elements provided at each neural network node are arranged in the crosspoint array module. Each crosspoint array module comprises two or more non-volatile memory elements arranged in a parallel connection between a neural network node input to a neural network node output.

The crosspoint array modules may digitally encode a linear scale of conductance values in the crosspoint array. The crosspoint array may be configured to convert current to voltage. In this case, each input terminal of the crosspoint array may be configured to receive a current and each output terminal of the crosspoint array may be configured to output a voltage.

The crosspoint array 60 is formed of three column signal lines 20 and three row signal lines 18. At each neural network node 12 there are at least two non-volatile memory elements that are connected in parallel between the neural network node input 22 and neural network node output 24. In the illustrated example, there are two non-volatile memory elements 62a and 62b in each crosspoint array module at each neural network node 12. Each non-volatile memory element 62a, 62b is coupled to a switching transistor 64a, 64b. The switching transistors enable a particular non-volatile memory element to be connected or disconnected from the crosspoint array module of a neural network node 12.

To store weights for a particular ANN, each non-volatile memory element is addressed by programming both the row and the column voltages. Since all the non-volatile memory elements are in parallel, each row of the crosspoint array is coupled to an access device that enables selection of a single NVM element from

the plurality of NVM elements in a row, and thereby prevents accidental writes. The access devices are usually only used for NVM element selection during programming. In embodiments, the access device may be shared across multiple neural network nodes 12 in the same row 18.

5

Thus, in embodiments, at least one access device may be coupled to the or each row signal line of the crosspoint array, wherein the access device is configured to select one of the plurality of non-volatile memory elements coupled to the row signal line.

10

Control or write circuitry may be configured to drive a non-volatile memory by: selecting, using the access device, one of the plurality of non-volatile memory elements; driving the selected non-volatile memory element dependent on the generated programming signal.

15

In embodiments, each non-volatile memory element may be coupled to a switching device, and the write circuitry is configured to: select, using the access device, one of the plurality of non-volatile memory elements; and control the switching device coupled to the selected non-volatile memory element to activate a path through the selected non-volatile memory element.

20

The parallel arrangement of NVM elements in this 3x3 crosspoint array provides equivalent resistances that are $\{R, R/2, R/3..., R/8\}$ apart. This is a non-linear resistance scale. However, the equivalent conductance scale increases linearly as follows $\{G, 2G, 3G..., 8G\}$. Thus, in this embodiment, it is the conductivity that is digitally encoded in a 3-bit scale.

25

In the illustrated embodiment, the ANN is configured to convert voltages to currents, such that the input nodes 14 received voltage values and the output nodes 16 output current values.

30

Figure 9 is a flow diagram illustrating example steps to implement or program weights into a neural network formed of non-volatile memory elements arranged in parallel. The process begins by receiving weights for the crosspoint array 104 from a remote server 110 (step S200). Control circuitry 112 may

35

determine, using the received weights, which impedance state (LIS or HIS) each non-volatile memory (NVM) element in the crosspoint array 104 is required to be in to implement the received weights at each neural network node (step S202). This may be performed individually for each neural network node, i.e. for the crosspoint array module at a neural network node. At step S204, a row and a column of the crosspoint array 104 are selected. An access device is used to select one of the NVM elements of the plurality of parallel NVM elements in the selected row (step S206). In embodiments, if an NVM element is required to be programmed into an HRS state, it could either be programmed into HRS or it could be deselected/deactivated entirely. Thus, at step S208, the control circuitry may determine if the selected NVM element is to be written into HRS, and if so, may deactivate the path through the selected NVM element (step S210). Alternatively, the control circuitry may write the NVM element into the required state (LRS or HRS) (step S212). The control circuitry 112 may check if all of the non-volatile memory elements in the crosspoint array have been written (step S214) and if not, the selection and writing process is repeated until all of the NVM elements have been written or deactivated. The process then ends.

Figure 10 is a circuit diagram illustrating how negative weights are programmed into the artificial neural network implemented using parallel or series connected NVM elements in a crosspoint array. Here, for each negative weight implemented/programmed at a neural network node, a negative voltage is required. This may be done by using a negative reference voltage within a digital to analogue converter (DAC). For a negative weight at a neural network node, the column signal line at which the neural network node is provided is programmed such that a positive signal is programmed with infinity weight or HRS. A negative input is provided with the required weight. Thus, the negative input serves as a way of subtracting current away from the column.

Thus, in embodiments, the write circuitry may be configured to drive a non-volatile memory by: selecting, using the access device, one of the plurality of non-volatile memory elements; and driving the selected non-volatile memory element dependent on the generated programming signal. In embodiments, each non-volatile memory element may be coupled to a switching device, and the write circuitry may be configured to: select, using the access device, one of the plurality

of non-volatile memory elements; and control the switching device coupled to the selected non-volatile memory element to activate a path through the selected non-volatile memory element.

5 Thus, the present techniques use NVM elements that have two distinct resistance states to provide an artificial neural network. The above-described techniques may use memristors, ReRAMs and/or CES elements as the NVM elements. Using CES elements may be particularly useful to mitigate process variations that give rise to resistance values. Previous approaches are susceptible
10 to individual variations in the resistance values, whereas in the present techniques there is at least a 200% difference between successive resistance (or conductance) values.

 The present techniques provide an artificial neural network that can be used
15 for different applications easily, simply by obtaining the weights for each application and implementing them in the crosspoint array. The NVM elements may be programmed on the fly and therefore, it may be possible to reprogram the same neural network for different applications. For example, the ANN described herein may be used for low-power single keyword speed trigger; simple
20 face detection for security cameras; or a voice activated UI as an alternative to physical or virtual buttons.

 The present techniques provide weight storage (which may be a significant static power leakage source in neural networks) and provide neural network
25 processing by means of converting voltage to current. There is no active switching taking place during the processing, which results in an energy-efficient ANN.

 In embodiments, the ANN of the present techniques may be used to further tune the ANN weightings. For example, the crosspoint array may be coupled to a
30 hardware feedback module, which could potentially be used to optimise the weights of each neural network node. This may be achieved by measuring currents in the crosspoint array (e.g. at each neural network node output) and feeding the measured current back into the ANN to tune the weights stored in each neural network node.

35

In embodiments, the ANN of the present techniques may comprise a plurality of crosspoint arrays, where each crosspoint array provides one layer of the ANN. For example, a first crosspoint array may implement a conductance matrix which converts voltages to currents, and a second crosspoint array may
5 implement a resistance matrix which converts currents to voltages. The outputs of the first crosspoint array may provide the inputs of the second crosspoint array, such that the ANN converts voltage to current to voltage across the two layers.

In embodiments, the apparatus may comprise a plurality of crosspoint
10 arrays, wherein each crosspoint array provides one layer in a multi-layer (or tiered) artificial neural network. Figure 11a shows a circuit diagram of a tiered artificial neural network implemented using crosspoint arrays, and Figure 11b shows a schematic version of Figure 11a.

Broadly speaking, the techniques described herein implement a dot-product
15 computation of weights and input values in an artificial neural network (ANN). As mentioned above, an artificial neural network (ANN), or each ANN layer within a tiered ANN, has at least one in-edge (i.e. input into the ANN) and at least one out-edge (i.e. output of the ANN). Between the in-edge(s) and out-edge(s) of the ANN
20 are multiple neural network nodes. Each neural network node has its own in-edges and out-edges. The term "dot-product" used herein refers to the multiplication of the weight associated with each in-edge of an ANN with the magnitude of the input voltage that connects the particular edge to a neural network node. In general, each neural network node may have multiple such
25 edges driven by multiple such inputs. The computation at each neural network node may occur in two distinct steps. The first step comprises summing the products of the weights and the input magnitudes (dot-product or the multiply-accumulate), as shown in Figure 1b. The second step comprises using this sum as an input to an activation function of the neural network node. The activation
30 function is the non-linear part of the computation, and may be considered a "squashing function". The activation function of a neural network node defines the output of the node given particular inputs. In some cases, the activation function may be binary, such that the node can either fire or not fire. The activation or squashing function can take several forms, such as a Rectified Linear
35 Unit (ReLU), a hyperbolic tangent (sigmoid) function, a Heaviside step function, a

piece-wise linear function, etc. In the simplest case, the squashing function could be a binary function: if the sum of the products of the weights and the inputs is less than zero, then the neural network node outputs a "0" (or low voltage), else it outputs a "1" (or high voltage).

5

The present techniques provide a low-power and energy-efficient implementation of the dot-product (multiply-and-accumulate) computation within an ANN. Figures 11a and 11b show generalised schematic diagrams of an ANN comprising two layers: an input layer (comprising inputs $I_0 \dots I_3$) and an output
10 layer (comprising outputs $O_{20} \dots O_{23}$). In general, there may be multiple hidden layers of neural network nodes. Without loss of generality, Figures 11a and 11b show a two-layer ANN to illustrate how the crosspoint arrays of the present techniques may be used to provide multi-layer ANNs. The crosspoint array takes the input values (e.g. inputs $I_0 \dots I_3$), and performs a dot-product computation
15 using the weights stored by the non-volatile memory elements provided at each neural network node of the array.

The inputs into the ANN (e.g. inputs $I_0 \dots I_3$) may be provided as voltage values or as current values. For voltage inputs, the crosspoint array transforms
20 the inputs to the intermediate outputs (O^*), using the weights stored by the NVM elements provided at each node. For current inputs, the intermediate outputs (O^*) are voltage values. These outputs (O^*) are provided as inputs into circuitry that implements the activation function (i.e. the squashing function, or non-linear part of the computation), represented by the triangles in Figure 11a. In the simplest
25 case, this squashing or activation function may be a binary function, e.g. a step function, as explained above. The final outputs for the first layer in the ANN is denoted as $O_{<1,2,3,4>}$ in Figure 11a, and these outputs form the inputs to the second layer in the ANN, as shown.

30 In embodiments, a crosspoint array may form each layer of an artificial neural network. Alternatively, a crosspoint array may form one or more (but not necessarily all) layers of an artificial neural network. In each case, the outputs of one crosspoint array are fed into circuitry that implements the activation function computation, prior to being supplied as the inputs into the subsequent layer of the
35 artificial neural network.

Thus, in embodiments, the crosspoint array of the apparatus is a first crosspoint array forming a first layer in the artificial neural network, wherein the first crosspoint array comprises a plurality of inputs I_{1p} and a plurality of outputs O_{1q} . The apparatus may comprise a second crosspoint array forming a second layer in the artificial neural network, wherein the second crosspoint array comprises a plurality of inputs I_{2r} and at least one output O_{2s} . In this embodiment, the apparatus may comprise further circuitry between each layer of the ANN to implement the activation function for each output edge of each layer.

The outputs O_{1q} of the first crosspoint array may be coupled to the inputs I_{2r} of the second crosspoint array. In embodiments, the first crosspoint array digitally encodes a linear scale of conductance values, and the second crosspoint array digitally encodes a linear scale of impedance values. In embodiments, the first crosspoint array converts voltage to current, and the second crosspoint array converts current to voltage.

As mentioned earlier, the apparatus may be a constrained resource device such as an internet of things (IoT) device.

The present techniques provide a method of implementing a reconfigurable artificial neural network in a crosspoint array comprising at least one neural network node located at an intersection of a row signal line and a column signal line, the method comprising: receiving data specifying a weight for each neural network node of the crosspoint array; switching a plurality of non-volatile memory elements provided at each neural network node into a required impedance state to implement the specified weight at each neural network node; and storing, using the non-volatile memory elements, the weight for each neural network node of the crosspoint array.

The present techniques provide a system of implementing a reconfigurable artificial neural network, the system comprising: an apparatus as described herein; a remote server configured to train an artificial neural network model; wherein the apparatus is configured to receive data specifying a weight for each neural network node of the apparatus from the remote server.

Those skilled in the art will appreciate that while the foregoing has described what is considered to be the best mode and where appropriate other modes of performing present techniques, the present techniques should not be limited to the specific configurations and methods disclosed in this description of the preferred embodiment. Those skilled in the art will recognise that present techniques have a broad range of applications, and that the embodiments may take a wide range of modifications without departing from the any inventive concept as defined in the appended claims.

10

CLAIMS

1. An apparatus for implementing a reconfigurable artificial neural network, the apparatus comprising:

5 a crosspoint array comprising:

at least one row signal line and at least one column signal line;

at least one neural network node, wherein each neural network node is provided at an intersection of a row signal line and a column signal line, and wherein each neural network node has an associated weight; and

10 a plurality of programmable non-volatile memory elements provided at each neural network node to store the associated weight, wherein each non-volatile memory element is switchable between a first impedance state and a second impedance state;

a communication module configured to receive data specifying the weight
15 for each neural network node of the crosspoint array; and

control circuitry configured to switch each of the plurality of non-volatile memory elements provided at each neural network node into the first or the second impedance state to implement the weight at the neural network node.

20 2. The apparatus as claimed in claim 1 wherein the communication module configured to receive data specifying the weight for each neural network node is configured to:

obtain at least one data packet specifying the weights for each neural network node; and

25 extract, from the data packet, the weight for each neural network node; and

wherein the control circuitry is configured to:

determine, using the weight for each neural network node, whether each non-volatile memory element at each neural network node is to be
30 switched into the first impedance state or the second impedance state; and

generate, responsive to the determination, a programming signal for each non-volatile memory element.

3. The apparatus as claimed in claim 2 wherein the control circuitry comprises
35 write circuitry configured to:

receive the generated programming signal for each non-volatile memory element; and

drive each non-volatile memory element dependent on the generated programming signal.

5

4. The apparatus as claimed in any one of claims 1 to 3 wherein the crosspoint array further comprises:

at least one input terminal to receive an input signal for the crosspoint array; and

10

at least one output terminal to provide an output signal generated by passing the input signal through at least one neural network node of the crosspoint array.

5. The apparatus as claimed in claim 4 wherein the neural network node receives the input signal, multiplies the associated weight of the neural network node with the input signal, and outputs a weighted signal.

15

6. The apparatus as claimed in any one of claims 1 to 3 wherein the crosspoint array comprises:

20

a plurality of row signal lines and a plurality of column signal lines;

a plurality of neural network nodes, wherein each neural network node is provided at an intersection of a row signal line and a column signal line, and wherein each neural network node has an associated weight; and

a plurality of programmable non-volatile memory elements provided at each neural network node, wherein each non-volatile memory element is switchable between a first impedance state and a second impedance state.

25

7. The apparatus as claimed in claim 6 wherein the crosspoint array further comprises:

30

a plurality of input terminals, each input terminal configured to receive an input signal for the crosspoint array; and

at least one output terminal to provide an output signal generated by passing the input signals through the neural network nodes of the crosspoint array.

35

8. The apparatus as claimed in claim 7 wherein the crosspoint array comprises a plurality of output terminals, each output terminal providing an output signal.

9. The apparatus as claimed in claim 7 or 8 wherein each neural network node receives two or more input signals, multiplies the associated weight with each input signal to generate weighted signals, and outputs a summed signal representing the summed weighted signals.

10. The apparatus as claimed in claim 9 wherein the at least one output terminal is configured to receive a summed signal from one or more of the plurality of neural network nodes, and to combine the summed signals to provide the output signal.

11. The apparatus as claimed in any one of claims 3 to 10 wherein each neural network node comprises a crosspoint array module, and the plurality of programmable non-volatile memory elements provided at each neural network node are arranged in the crosspoint array module.

12. The apparatus as claimed in claim 11 wherein the crosspoint array module of each neural network node comprises an array formed of N row signal lines and M column signal lines, wherein one of the plurality of non-volatile memory elements is provided at each intersection of the N row signal lines and M column signal lines.

13. The apparatus as claimed in claim 12 wherein each crosspoint array module comprises:

a direct path from a neural network node input to a neural network node output; and

at least one indirect path from the neural network node input to the neural network node output.

14. The apparatus as claimed in claim 13 wherein the crosspoint array modules digitally encode a linear scale of impedance values.

15. The apparatus as claimed in claim 13 or 14 wherein the at least one indirect path between the neural network node input to the neural network node output is formed of two or more non-volatile memory elements in a series connection.

5 16. The apparatus as claimed in any one of claims 13 to 15 wherein the crosspoint array is configured to convert voltage to current, and each input terminal of the crosspoint array is configured to receive a voltage and each output terminal of the crosspoint array is configured to output a current.

10 17. The apparatus as claimed in any one of claims 13 to 15 wherein the crosspoint array is configured to convert current to voltage, and each input terminal of the crosspoint array is configured to receive a current and each output terminal of the crosspoint array is configured to output a voltage.

15 18. The apparatus as claimed in any one of claims 13 to 17 wherein:
each of the N row signal lines of each crosspoint array module is formed of a continuous conductive channel which runs across the crosspoint array module;
and

each of the M column signal lines of each crosspoint array module
20 comprises:

a plurality of continuous column signal lines, wherein each continuous column signal line is formed of a continuous conductive channel which runs across the crosspoint array module; and

25 a plurality of discontinuous column signal lines, wherein each discontinuous column signal line is formed of multiple spatially separated conductive channels.

19. The apparatus as claimed in claim 18 wherein the plurality of continuous column signal lines are interlaced with the plurality of discontinuous column signal
30 lines in an alternating arrangement.

20. The apparatus as claimed in claim 19 when dependent on claim 3, wherein the control circuitry is configured to drive each non-volatile memory element of each crosspoint array module that is in electrical communication with a continuous
35 column signal line by:

selecting a continuous column signal line; and
driving each non-volatile memory element connected to the selected continuous line dependent on the generated programming signal.

5 21. The apparatus as claimed in claim 19, when dependent on claim 3, wherein the control circuitry is configured to drive each non-volatile memory element that is in electrical communication with a discontinuous column signal line by:

writing each non-volatile memory connected to a neighbouring continuous column signal line into a first impedance state;

10 connecting the discontinuous column signal line to the neighbouring continuous column signal line;

driving each non-volatile memory element connected to the discontinuous line dependent on the generated programming signal; and

15 disconnecting the discontinuous column signal line from the neighbouring continuous column signal line.

22. The apparatus as claimed in any one of claims 13 to 21 wherein each crosspoint array module digitally encodes a linear scale of impedance values.

20 23. The apparatus as claimed in claim 11 wherein each crosspoint array module comprises two or more non-volatile memory elements arranged in a parallel connection between a neural network node input to a neural network node output.

24. The apparatus as claimed in claim 23 wherein the crosspoint array modules
25 digitally encode a linear scale of conductance values.

25. The apparatus as claimed in claim 23 or 24 wherein the crosspoint array is configured to convert current to voltage, and each input terminal of the crosspoint array is configured to receive a current and each output terminal of the crosspoint
30 array is configured to output a voltage.

26. The apparatus as claimed in claim 23 or 24 wherein the crosspoint array is configured to convert voltage to current, and each input terminal of the crosspoint array is configured to receive a voltage and each output terminal of the crosspoint
35 array is configured to output a current.

27. The apparatus as claimed in any one of claims 23 to 26 further comprising at least one access device coupled to the or each row signal line of the crosspoint array, wherein the access device is configured to select one of the plurality of non-volatile memory elements coupled to the row signal line.

28. The apparatus as claimed in claim 27, when dependent on claim 3, wherein the write circuitry is configured to drive a non-volatile memory by:

selecting, using the access device, one of the plurality of non-volatile memory elements;

driving the selected non-volatile memory element dependent on the generated programming signal.

29. The apparatus as claimed in claim 28 wherein each non-volatile memory element is coupled to a switching device, and the write circuitry is configured to:

select, using the access device, one of the plurality of non-volatile memory elements; and

control the switching device coupled to the selected non-volatile memory element to activate a path through the selected non-volatile memory element.

30. The apparatus as claimed in any preceding claim wherein the weight associated with each neural network node is a negative weight.

31. The apparatus as claimed in any preceding claim wherein the first impedance state is a high impedance state.

32. The apparatus as claimed in any preceding claim wherein the first impedance state is a high impedance state and the second impedance state is a low impedance state.

33. The apparatus as claimed in any preceding claim wherein the non-volatile memory element is a non-polar non-volatile memory element.

34. The apparatus as claimed in any preceding claim wherein the non-volatile memory element is any one of: a memristor, a resistive random access memory (ReRAM), a spin-transfer torque magnetic random access memory (STT-MRAM), and a phase-change memory (PCM).

5

35. The apparatus as claimed in any preceding claim wherein the non-volatile memory element is a correlated electron switch (CES).

36. The apparatus as claimed in any preceding claim wherein the communication module is configured to receive data specifying the weight for each neural network node of the crosspoint array from an artificial neural network model trained external to the apparatus.

37. The apparatus as claimed in any preceding claim wherein the communication module is configured to receive further data specifying a further weight for each neural network node of the crosspoint array, wherein the further weight is obtained from a machine learning model trained external to the apparatus.

38. The apparatus as claimed in any preceding claim further comprising a data store coupled to the communication module and configured to store the received data.

39. The apparatus as claimed in claim 38 wherein the apparatus is configured to receive instructions indicating a particular artificial neural network is to be implemented in the apparatus.

40. The apparatus as claimed in claim 39 wherein the control circuitry is configured to:

retrieve, from the data store, the data corresponding to the artificial neural network indicated in the received instructions; and

use the retrieved data to extract a weight for each neural network node of the crosspoint array.

41. The apparatus as claimed in any preceding claim

wherein the crosspoint array is a first crosspoint array forming a first layer in the artificial neural network, wherein the first crosspoint array comprises a plurality of inputs I_{1p} and a plurality of outputs O_{1q} , and

5 wherein the plurality of outputs O_{1q} are provided as inputs into a second layer in the artificial neural network.

42. The apparatus as claimed in any preceding claim wherein the apparatus is an internet of things (IoT) device.

10 43. A method of implementing a reconfigurable artificial neural network in a crosspoint array comprising at least one neural network node located at an intersection of a row signal line and a column signal line, the method comprising:

receiving data specifying a weight for each neural network node of the crosspoint array;

15 switching a plurality of non-volatile memory elements provided at each neural network node into a required impedance state to implement the specified weight at each neural network node; and

storing, using the non-volatile memory elements, the weight for each neural network node of the crosspoint array.

20



44. A system of implementing a reconfigurable artificial neural network, the system comprising:

an apparatus of any one of claims 1 to 42; and


a remote server configured to train an artificial neural network model;

25 wherein the apparatus is configured to receive data specifying a weight for each neural network node of the apparatus from the remote server.

1 / 11

**FIGURE 1a****FIGURE 1b**

2 / 11

**FIGURE 2**




FIGURE 3

4 / 11




FIGURE 4a







FIGURE 4b





7 / 11

*FIGURE 7*





FIGURE 8

9 / 11

**FIGURE 9**




FIGURE 10

11 / 11




FIGURE 11a




FIGURE 11b

INTERNATIONAL SEARCH REPORT

International application No

PCT/GB2017/051734

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06N3/063
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06N G06F G11C G06G

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal, WPI Data, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	F. ALIBART ET AL: "Pattern classification by memristive crossbar circuits using ex situ and in situ training", NATURE COMMUNICATIONS, vol. 4, 2012, 25 June 2013 (2013-06-25), XP055196306, DOI: 10.1038/ncomms3072 pages 1-4	1-44
X	A. M. SHERI ET AL: "Neuromorphic character recognition system with two PCMO memristors as a synapse", IEEE TRANSACTIONS ON INDUSTRIAL ELECTRONICS, vol. 61, no. 6, 1 August 2013 (2013-08-01), pages 2933-2941, XP011534989, DOI: 10.1109/TIE.2013.2275966 sections II-IV	1-44
	----- -/--	



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

22 September 2017

Date of mailing of the international search report

02/10/2017

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Douarche, Nicolas

INTERNATIONAL SEARCH REPORT

International application No

PCT/GB2017/051734

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	S. PARK ET AL: "Electronic system with memristive synapses for pattern recognition", SCIENTIFIC REPORTS, vol. 5, 10123, 5 May 2015 (2015-05-05), XP055408514, DOI: 10.1038/srep10123 pages 1-5 -----	1-44
X	C. YAKOPCIC, T. M. TAHA: "Ex-situ programming in a neuromorphic memristor based crossbar circuit", PROCEEDINGS OF THE 2015 NATIONAL AEROSPACE AND ELECTRONICS CONFERENCE (NAECON'15), 15 June 2015 (2015-06-15), pages 300-304, XP032887199, DOI: 10.1109/NAECON.2015.7443087 sections I and II -----	1-44
X	G. W. BURR ET AL: "Experimental demonstration and tolerancing of a large-scale neural network (165 000 synapses) using phase-change memory as the synaptic weight element", IEEE TRANSACTIONS ON ELECTRON DEVICES, vol. 62, no. 11, 7 July 2015 (2015-07-07), pages 3498-3507, XP011587225, DOI: 10.1109/TED.2015.2439635 sections I and II -----	1-44
X	A. SHAFIEE ET AL: "ISAAC: a convolutional neural network accelerator with in-situ analog arithmetic in crossbars", PROCEEDINGS OF THE 21ST INTERNATIONAL CONFERENCE ON PROGRAM COMPREHENSION (ICPC'16), 18 June 2016 (2016-06-18), pages 14-26, XP032950645, DOI: 10.1109/ISCA.2016.12 sections I-III -----	1-44
A	I. VOURKAS, G. C. SIRAKOULIS: "Memristive crossbar-based nonvolatile memory", MEMRISTOR-BASED NANOELECTRONIC COMPUTING CIRCUITS AND ARCHITECTURES, 27 August 2015 (2015-08-27), pages 101-147, XP009500145, DOI: 10.1007/978-3-319-22647-7_5 sections 5.2 and 5.4 -----	1-44