



(51) International Patent Classification:

G06N 3/063 (2006.01)

G06N 3/04 (2006.01)

G06N 3/08 (2006.01)

(21) International Application Number:

PCT/US2021/055611

(22) International Filing Date:

19 October 2021 (19.10.2021)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

17/092,033

06 November 2020 (06.11.2020) US

(71) Applicant: MICRON TECHNOLOGY, INC. [US/US];
8000 South Federal Way, Boise, Idaho 83707 (US).(72) Inventors: ZAIDY, Aliasger Tayeb; 50 University Street,
Apt 1210, Seattle, Washington 98101 (US). VITEZ,
Marko; Loc. Trebiciano 269, 34149 Trieste (IT). CULUR-
CIELLO, Eugenio; 923 Windsor Drive, West Lafayette,
Indiana 47906 (US). CUMMINS, Jaime; 782 Cherry Av-
enue, Bainbridge Island, Washington 98110 (US). MING
CHANG, Andre Xian; 50 University St., Apt. 704, Seattle,
Washington 98101 (US).(74) Agent: WARD, John P. et al.; Greenberg Traurig, LLP,
77 West Wacker Drive, Suite 3100, Chicago, Illinois 60601
(US).(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, IT, 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, ST, SV, SY, TH, TJ, TM, TN,
TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, 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).

Published:

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


(54) Title: DISCOVERY OF HARDWARE CHARACTERISTICS OF DEEP LEARNING ACCELERATORS FOR OPTIMIZATION
VIA COMPILER

FIG. 13

(57) **Abstract:** Systems, devices, and methods related to a Deep Learning Accelerator and memory are described. For example, an integrated circuit device may be configured to execute instructions with matrix operands and configured with random access memory. A computing device running a compiler can interact and/or probe an integrated circuit device to identify hardware characteristics of the integrated circuit device in performing matrix computations. The compiler can generate and optimize a result of compilation from a description of an artificial neural network based at least in part on the hardware characteristics of the integrated circuit device. The result of compilation can include first data representative of parameters of the artificial neural network and second data representative of instructions executable by the integrated circuit device to generate an output of the artificial neural network based on the first data and an input to the artificial neural network.

DISCOVERY OF HARDWARE CHARACTERISTICS OF DEEP LEARNING
ACCELERATORS FOR OPTIMIZATION VIA COMPILER

RELATED APPLICATION

[0001] The present application claims priority to U.S. Pat. App. Ser. No. 17/092,033, filed Nov. 6, 2020 and entitled "DISCOVERY OF HARDWARE CHARACTERISTICS OF DEEP LEARNING ACCELERATORS FOR OPTIMIZATION VIA COMPILER," the entire disclosure of which is hereby incorporated herein by reference.

TECHNICAL FIELD

[0002] At least some embodiments disclosed herein relate to compilers in general and more particularly, but not limited to, compilers to generate instructions executable by accelerators for Artificial Neural Networks (ANNs), such as ANNs configured through machine learning and/or deep learning.

BACKGROUND

[0003] An Artificial Neural Network (ANN) uses a network of neurons to process inputs to the network and to generate outputs from the network.

[0004] Deep learning has been applied to many application fields, such as computer vision, speech/audio recognition, natural language processing, machine translation, bioinformatics, drug design, medical image processing, games, etc.

BRIEF DESCRIPTION OF THE DRAWINGS

[0005] The embodiments are illustrated by way of example and not limitation in the figures of the accompanying drawings in which like references indicate similar elements.

[0006] **FIG. 1** shows an integrated circuit device having a Deep Learning Accelerator and random access memory configured according to one embodiment.

[0007] **FIG. 2** shows a processing unit configured to perform matrix-matrix operations according to one embodiment.

[0008] **FIG. 3** shows a processing unit configured to perform matrix-vector

operations according to one embodiment.

[0009] FIG. 4 shows a processing unit configured to perform vector-vector operations according to one embodiment.

[0010] FIG. 5 shows a Deep Learning Accelerator and random access memory configured to autonomously apply inputs to a trained Artificial Neural Network according to one embodiment.

[0011] FIG. 6 shows a technique to generate instructions executable by a Deep Learning Accelerator to implement an Artificial Neural Network according to one embodiment.

[0012] FIGS. 7 and 8 illustrate techniques to map a compilation result for a generic Deep Learning Accelerator into instructions executable by a specific Deep Learning Accelerator to implement an Artificial Neural Network according to one embodiment.

[0013] FIG. 9 shows another technique to generate instructions executable by a Deep Learning Accelerator to implement an Artificial Neural Network according to one embodiment.

[0014] FIG. 10 shows an integrated circuit device having a Deep Learning Accelerator with configurable hardware capabilities with random access memory configured according to one embodiment.

[0015] FIG. 11 illustrates different hardware configurations of a processing unit of a Deep Learning Accelerator that are configurable via options stored in register according to one embodiment.

[0016] FIG. 12 illustrates a technique to generate instructions executable by a Deep Learning Accelerator with optimized hardware configuration to implement an Artificial Neural Network according to one embodiment.

[0017] FIG. 13 shows a technique to discover hardware characteristics of a Deep Learning Accelerator according to one embodiment.

[0018] FIG. 14 illustrates a technique to generate instructions executable by a Deep Learning Accelerator and optimized according hardware characteristics of the Deep Learning Accelerator according to one embodiment.

[0019] FIG. 15 shows a method to compile instructions to implement an Artificial Neural Network on a Deep Learning Accelerator based on its hardware characteristics according to one embodiment.

[0020] FIG. 16 shows a block diagram of an example computer system in which

embodiments of the present disclosure can operate.

DETAILED DESCRIPTION

[0021] At least some embodiments disclosed herein provide integrated circuits to implement the computation of Artificial Neural Networks (ANNs) with reduced energy consumption and computation time. Such an integrated circuit device is programmable. A compiler can be used to generate, from a description of an Artificial Neural Network (ANN), instructions executable in the integrated circuit device. When executed in the device, the instructions cause the integrated circuit device to perform the computation of the Artificial Neural Network (ANN). The compiler can discover hardware characteristics of the Deep Learning Accelerator (DLA), such as capabilities and behaviors, and use the discovered hardware characteristics in optimizing the instructions generated to implement the computation of an Artificial Neural Network (ANN) on the Deep Learning Accelerator (DLA). For example, the integrated circuit device can include a Deep Learning Accelerator (DLA) and random access memory. The random access memory is configured to store parameters of an Artificial Neural Network (ANN) and instructions having matrix operands. The instructions stored in the random access memory are executable by the Deep Learning Accelerator (DLA) to implement matrix computations according to the Artificial Neural Network (ANN).

[0022] For example, the compiler can automatically exercise the hardware of the Deep Learning Accelerator (DLA) to discover features, options, behaviors, performances, latency, and/or limitations of the Deep Learning Accelerator (DLA). For example, the compiler can load test sets of instructions into the random access memory to cause the Deep Learning Accelerator (DLA) to execute the instructions, inspect the execution result, and/or measure the performance of the execution of the instructions to discover the characteristics of the Deep Learning Accelerator (DLA) in capability, performance, and/or behavior. Based on the discovered characteristics of different hardware platforms of Deep Learning Accelerator (DLA), the compiler automatically customizes and/or optimizes implementations of a same Artificial Neural Network (ANN) on different hardware platforms.

[0023] For example, each neuron in the network receives a set of inputs. Some of the inputs to a neuron may be the outputs of certain neurons in the network; and

some of the inputs to a neuron may be the inputs provided to the neural network. The input/output relations among the neurons in the network represent the neuron connectivity in the network.

[0024] For example, each neuron can have a bias, an activation function, and a set of synaptic weights for its inputs respectively. The activation function may be in the form of a step function, a linear function, a log-sigmoid function, etc. Different neurons in the network may have different activation functions.

[0025] For example, each neuron can generate a weighted sum of its inputs and its bias and then produce an output that is the function of the weighted sum, computed using the activation function of the neuron.

[0026] The relations between the input(s) and the output(s) of an ANN in general are defined by an ANN model that includes the data representing the connectivity of the neurons in the network, as well as the bias, activation function, and synaptic weights of each neuron. Based on a given ANN model, a computing device can be configured to compute the output(s) of the network from a given set of inputs to the network.

[0027] For example, the inputs to an ANN network may be generated based on camera inputs; and the outputs from the ANN network may be the identification of an item, such as an event or an object.

[0028] In general, an ANN may be trained using a supervised method where the parameters in the ANN are adjusted to minimize or reduce the error between known outputs associated with or resulted from respective inputs and computed outputs generated via applying the inputs to the ANN. Examples of supervised learning/training methods include reinforcement learning and learning with error correction.

[0029] Alternatively, or in combination, an ANN may be trained using an unsupervised method where the exact outputs resulted from a given set of inputs is not known before the completion of the training. The ANN can be trained to classify an item into a plurality of categories, or data points into clusters.

[0030] Multiple training algorithms can be employed for a sophisticated machine learning/training paradigm.

[0031] Deep learning uses multiple layers of machine learning to progressively extract features from input data. For example, lower layers can be configured to identify edges in an image; and higher layers can be configured to identify, based on

the edges detected using the lower layers, items captured in the image, such as faces, objects, events, etc. Deep learning can be implemented via Artificial Neural Networks (ANNs), such as deep neural networks, deep belief networks, recurrent neural networks, and/or convolutional neural networks.

[0032] A typical Deep Learning Accelerator (DLA) can include a set of programmable hardware computing logic that is specialized and/or optimized to perform parallel vector and/or matrix calculations, including but not limited to multiplication and accumulation of vectors and/or matrices.

[0033] Further, the Deep Learning Accelerator can include one or more Arithmetic-Logic Units (ALUs) to perform arithmetic and bitwise operations on integer binary numbers.

[0034] The Deep Learning Accelerator is programmable via a set of instructions to perform the computations of an Artificial Neural Network (ANN).

[0035] The granularity of the Deep Learning Accelerator operating on vectors and matrices corresponds to the largest unit of vectors/matrices that can be operated upon during the execution of one instruction by the Deep Learning Accelerator. During the execution of the instruction for a predefined operation on vector/matrix operands, elements of vector/matrix operands can be operated upon by the Deep Learning Accelerator in parallel to reduce execution time and/or energy consumption associated with memory/data access. The operations on vector/matrix operands of the granularity of the Deep Learning Accelerator can be used as building blocks to implement computations on vectors/matrices of larger sizes.

[0036] The implementation of a typical/practical Artificial Neural Network involves vector/matrix operands having sizes that are larger than the operation granularity of the Deep Learning Accelerator. To implement such an Artificial Neural Network using the Deep Learning Accelerator, computations involving the vector/matrix operands of large sizes can be broken down to the computations of vector/matrix operands of the granularity of the Deep Learning Accelerator. The Deep Learning Accelerator can be programmed via instructions to carry out the computations involving large vector/matrix operands. For example, atomic computation capabilities of the Deep Learning Accelerator in manipulating vectors and matrices of the granularity of the Deep Learning Accelerator in response to instructions can be programmed to implement computations in an Artificial Neural Network.

[0037] In some implementations, the Deep Learning Accelerator lacks some of

the logic operation capabilities of a typical Central Processing Unit (CPU). However, the Deep Learning Accelerator can be configured with sufficient logic units to process the input data provided to an Artificial Neural Network and generate the output of the Artificial Neural Network according to a set of instructions generated for the Deep Learning Accelerator. Thus, the Deep Learning Accelerator can perform the computation of an Artificial Neural Network with little or no help from a Central Processing Unit (CPU) or another processor. Optionally, a conventional general purpose processor can also be configured as part of the Deep Learning Accelerator to perform operations that cannot be implemented efficiently using the vector/matrix processing units of the Deep Learning Accelerator, and/or that cannot be performed by the vector/matrix processing units of the Deep Learning Accelerator.

[0038] A typical Artificial Neural Network can be described/specified in a standard format (e.g., Open Neural Network Exchange (ONNX)). A compiler can be used to convert the description of the Artificial Neural Network into a set of instructions for the Deep Learning Accelerator to perform calculations of the Artificial Neural Network. The compiler can optimize the set of instructions to improve the performance of the Deep Learning Accelerator in implementing the Artificial Neural Network.

[0039] The Deep Learning Accelerator can have local memory, such as registers, buffers and/or caches, configured to store vector/matrix operands and the results of vector/matrix operations. Intermediate results in the registers can be pipelined/shifted in the Deep Learning Accelerator as operands for subsequent vector/matrix operations to reduce time and energy consumption in accessing memory/data and thus speed up typical patterns of vector/matrix operations in implementing a typical Artificial Neural Network. The capacity of registers, buffers and/or caches in the Deep Learning Accelerator is typically insufficient to hold the entire data set for implementing the computation of a typical Artificial Neural Network. Thus, a random access memory coupled to the Deep Learning Accelerator is configured to provide an improved data storage capability for implementing a typical Artificial Neural Network. For example, the Deep Learning Accelerator loads data and instructions from the random access memory and stores results back into the random access memory.

[0040] The communication bandwidth between the Deep Learning Accelerator and the random access memory is configured to optimize or maximize the utilization

of the computation power of the Deep Learning Accelerator. For example, high communication bandwidth can be provided between the Deep Learning Accelerator and the random access memory such that vector/matrix operands can be loaded from the random access memory into the Deep Learning Accelerator and results stored back into the random access memory in a time period that is approximately equal to the time for the Deep Learning Accelerator to perform the computations on the vector/matrix operands. The granularity of the Deep Learning Accelerator can be configured to increase the ratio between the amount of computations performed by the Deep Learning Accelerator and the size of the vector/matrix operands such that the data access traffic between the Deep Learning Accelerator and the random access memory can be reduced, which can reduce the requirement on the communication bandwidth between the Deep Learning Accelerator and the random access memory. Thus, the bottleneck in data/memory access can be reduced or eliminated.

[0041] Optionally, the compiler can be configurable to support different hardware platforms of Deep Learning Accelerators. Specifically, the compiler can generate different sets of instructions for different Deep Learning Accelerators based on a same description of an Artificial Neural Network. For example, Deep Learning Accelerators can be implemented using different technologies, such as Field-Programmable Gate Array (FPGA) or Application Specific Integrated circuit (ASIC). For example, Deep Learning Accelerators can have different hardware capabilities in implementing matrix operations, have different numbers of parallel processing units operable to perform matrix operations concurrently, and/or have different computation granularity where processing units can have different capacities in processing matrices of different sizes in execution of an instruction having matrix operands. The compiler can initially apply generic, platform agnostic optimization to the description of the Artificial Neural Network to generate a generic computing model according to common characteristics of computations implemented using different Deep Learning Accelerators. Then, the compiler maps the compilation result of the generic computing model to different hardware platforms/implementations of Deep Learning Accelerators. Optionally, the compiler can further optimize the compilation result for individual types of Deep Learning Accelerators to reduce energy consumption and/or computation time.

[0042] In some implementations, a Deep Learning Accelerator (DLA) is designed

to have multiple configurable hardware options. The different hardware options can be optimal in different scenarios of artificial neural network computations. During compilation and optimization of an Artificial Neural Network, a compiler is configured to optimize instructions generated for execution by the Deep Learning Accelerator. The compiler optimization can include the selection of hardware options to improve the overall performance of an implementation of an Artificial Neural Network (ANN) in the Deep Learning Accelerator. Thus, the compiler can optimize and/or customize the circuit configurations of the Deep Learning Accelerator itself in implementing a particular Artificial Neural Network.

[0043] FIG. 1 shows an integrated circuit device (101) having a Deep Learning Accelerator (103) and random access memory (105) configured according to one embodiment.

[0044] The Deep Learning Accelerator (103) in FIG. 1 includes processing units (111), a control unit (113), and local memory (115). When vector and matrix operands are in the local memory (115), the control unit (113) can use the processing units (111) to perform vector and matrix operations in accordance with instructions. Further, the control unit (113) can load instructions and operands from the random access memory (105) through a memory interface (117) and a high speed/bandwidth connection (119).

[0045] The integrated circuit device (101) is configured to be enclosed within an integrated circuit package with pins or contacts for a memory controller interface (107).

[0046] The memory controller interface (107) is configured to support a standard memory access protocol such that the integrated circuit device (101) appears to a typical memory controller in a way same as a conventional random access memory device having no Deep Learning Accelerator (103). For example, a memory controller external to the integrated circuit device (101) can access, using a standard memory access protocol through the memory controller interface (107), the random access memory (105) in the integrated circuit device (101).

[0047] The integrated circuit device (101) is configured with a high bandwidth connection (119) between the random access memory (105) and the Deep Learning Accelerator (103) that are enclosed within the integrated circuit device (101). The bandwidth of the connection (119) is higher than the bandwidth of the connection (109) between the random access memory (105) and the memory controller

interface (107).

[0048] In one embodiment, both the memory controller interface (107) and the memory interface (117) are configured to access the random access memory (105) via a same set of buses or wires. Thus, the bandwidth to access the random access memory (105) is shared between the memory interface (117) and the memory controller interface (107). Alternatively, the memory controller interface (107) and the memory interface (117) are configured to access the random access memory (105) via separate sets of buses or wires. Optionally, the random access memory (105) can include multiple sections that can be accessed concurrently via the connection (119). For example, when the memory interface (117) is accessing a section of the random access memory (105), the memory controller interface (107) can concurrently access another section of the random access memory (105). For example, the different sections can be configured on different integrated circuit dies and/or different planes/banks of memory cells; and the different sections can be accessed in parallel to increase throughput in accessing the random access memory (105). For example, the memory controller interface (107) is configured to access one data unit of a predetermined size at a time; and the memory interface (117) is configured to access multiple data units, each of the same predetermined size, at a time.

[0049] In one embodiment, the random access memory (105) and the integrated circuit device (101) are configured on different integrated circuit dies configured within a same integrated circuit package. Further, the random access memory (105) can be configured on one or more integrated circuit dies that allows parallel access of multiple data elements concurrently.

[0050] In some implementations, the number of data elements of a vector or matrix that can be accessed in parallel over the connection (119) corresponds to the granularity of the Deep Learning Accelerator operating on vectors or matrices. For example, when the processing units (111) can operate on a number of vector/matrix elements in parallel, the connection (119) is configured to load or store the same number, or multiples of the number, of elements via the connection (119) in parallel.

[0051] Optionally, the data access speed of the connection (119) can be configured based on the processing speed of the Deep Learning Accelerator (103). For example, after an amount of data and instructions have been loaded into the local memory (115), the control unit (113) can execute an instruction to operate on

the data using the processing units (111) to generate output. Within the time period of processing to generate the output, the access bandwidth of the connection (119) allows the same amount of data and instructions to be loaded into the local memory (115) for the next operation and the same amount of output to be stored back to the random access memory (105). For example, while the control unit (113) is using a portion of the local memory (115) to process data and generate output, the memory interface (117) can offload the output of a prior operation into the random access memory (105) from, and load operand data and instructions into, another portion of the local memory (115). Thus, the utilization and performance of the Deep Learning Accelerator are not restricted or reduced by the bandwidth of the connection (119).

[0052] The random access memory (105) can be used to store the model data of an Artificial Neural Network and to buffer input data for the Artificial Neural Network. The model data does not change frequently. The model data can include the output generated by a compiler for the Deep Learning Accelerator to implement the Artificial Neural Network. The model data typically includes matrices used in the description of the Artificial Neural Network and instructions generated for the Deep Learning Accelerator (103) to perform vector/matrix operations of the Artificial Neural Network based on vector/matrix operations of the granularity of the Deep Learning Accelerator (103). The instructions operate not only on the vector/matrix operations of the Artificial Neural Network, but also on the input data for the Artificial Neural Network.

[0053] In one embodiment, when the input data is loaded or updated in the random access memory (105), the control unit (113) of the Deep Learning Accelerator (103) can automatically execute the instructions for the Artificial Neural Network to generate an output of the Artificial Neural Network. The output is stored into a predefined region in the random access memory (105). The Deep Learning Accelerator (103) can execute the instructions without help from a Central Processing Unit (CPU). Thus, communications for the coordination between the Deep Learning Accelerator (103) and a processor outside of the integrated circuit device (101) (e.g., a Central Processing Unit (CPU)) can be reduced or eliminated.

[0054] Optionally, the logic circuit of the Deep Learning Accelerator (103) can be implemented via Complementary Metal Oxide Semiconductor (CMOS). For example, the technique of CMOS Under the Array (CUA) of memory cells of the random access memory (105) can be used to implement the logic circuit of the Deep

Learning Accelerator (103), including the processing units q(111) and the control unit (113). Alternatively, the technique of CMOS in the Array of memory cells of the random access memory (105) can be used to implement the logic circuit of the Deep Learning Accelerator (103).

[0055] In some implementations, the Deep Learning Accelerator (103) and the random access memory (105) can be implemented on separate integrated circuit dies and connected using Through-Silicon Vias (TSV) for increased data bandwidth between the Deep Learning Accelerator (103) and the random access memory (105). For example, the Deep Learning Accelerator (103) can be formed on an integrated circuit die of a Field-Programmable Gate Array (FPGA) or Application Specific Integrated circuit (ASIC).

[0056] Alternatively, the Deep Learning Accelerator (103) and the random access memory (105) can be configured in separate integrated circuit packages and connected via multiple point-to-point connections on a printed circuit board (PCB) for parallel communications and thus increased data transfer bandwidth.

[0057] The random access memory (105) can be volatile memory or non-volatile memory, or a combination of volatile memory and non-volatile memory. Examples of non-volatile memory include flash memory, memory cells formed based on negative- and (NAND) logic gates, negative-or (NOR) logic gates, Phase-Change Memory (PCM), magnetic memory (MRAM), resistive random-access memory, cross point storage and memory devices. A cross point memory device can use transistor-less memory elements, each of which has a memory cell and a selector that are stacked together as a column. Memory element columns are connected via two lays of wires running in perpendicular directions, where wires of one lay run in one direction in the layer that is located above the memory element columns, and wires of the other lay run in another direction and are located below the memory element columns. Each memory element can be individually selected at a cross point of one wire on each of the two layers. Cross point memory devices are fast and non-volatile and can be used as a unified memory pool for processing and storage. Further examples of non-volatile memory include Read-Only Memory (ROM), Programmable Read-Only Memory (PROM), Erasable Programmable Read-Only Memory (EPROM) and Electronically Erasable Programmable Read-Only Memory (EEPROM) memory, etc. Examples of volatile memory include Dynamic Random-Access Memory (DRAM) and Static Random-Access Memory (SRAM).

[0058] For example, non-volatile memory can be configured to implement at least a portion of the random access memory (105). The non-volatile memory in the random access memory (105) can be used to store the model data of an Artificial Neural Network. Thus, after the integrated circuit device (101) is powered off and restarts, it is not necessary to reload the model data of the Artificial Neural Network into the integrated circuit device (101). Further, the non-volatile memory can be programmable/rewritable. Thus, the model data of the Artificial Neural Network in the integrated circuit device (101) can be updated or replaced to implement an update Artificial Neural Network, or another Artificial Neural Network.

[0059] The processing units (111) of the Deep Learning Accelerator (103) can include vector-vector units, matrix-vector units, and/or matrix-matrix units. Examples of units configured to perform for vector-vector operations, matrix-vector operations, and matrix-matrix operations are discussed below in connection with **FIGS. 2 – 4**.

[0060] **FIG. 2** shows a processing unit configured to perform matrix-matrix operations according to one embodiment. For example, the matrix-matrix unit (121) of **FIG. 2** can be used as one of the processing units (111) of the Deep Learning Accelerator (103) of **FIG. 1**.

[0061] In **FIG. 2**, the matrix-matrix unit (121) includes multiple kernel buffers (131 to 133) and multiple the maps banks (151 to 153). Each of the maps banks (151 to 153) stores one vector of a matrix operand that has multiple vectors stored in the maps banks (151 to 153) respectively; and each of the kernel buffers (131 to 133) stores one vector of another matrix operand that has multiple vectors stored in the kernel buffers (131 to 133) respectively. The matrix-matrix unit (121) is configured to perform multiplication and accumulation operations on the elements of the two matrix operands, using multiple matrix-vector units (141 to 143) that operate in parallel.

[0062] A crossbar (123) connects the maps banks (151 to 153) to the matrix-vector units (141 to 143). The same matrix operand stored in the maps bank (151 to 153) is provided via the crossbar (123) to each of the matrix-vector units (141 to 143); and the matrix-vector units (141 to 143) receives data elements from the maps banks (151 to 153) in parallel. Each of the kernel buffers (131 to 133) is connected to a respective one in the matrix-vector units (141 to 143) and provides a vector operand to the respective matrix-vector unit. The matrix-vector units (141 to 143) operate concurrently to compute the operation of the same matrix operand, stored in the maps banks (151 to 153) multiplied by the corresponding vectors stored in the

kernel buffers (131 to 133). For example, the matrix-vector unit (141) performs the multiplication operation on the matrix operand stored in the maps banks (151 to 153) and the vector operand stored in the kernel buffer (131), while the matrix-vector unit (143) is concurrently performing the multiplication operation on the matrix operand stored in the maps banks (151 to 153) and the vector operand stored in the kernel buffer (133).

[0063] Each of the matrix-vector units (141 to 143) in **FIG. 2** can be implemented in a way as illustrated in **FIG. 3**.

[0064] **FIG. 3** shows a processing unit configured to perform matrix-vector operations according to one embodiment. For example, the matrix-vector unit (141) of **FIG. 3** can be used as any of the matrix-vector units in the matrix-matrix unit (121) of **FIG. 2**.

[0065] In **FIG. 3**, each of the maps banks (151 to 153) stores one vector of a matrix operand that has multiple vectors stored in the maps banks (151 to 153) respectively, in a way similar to the maps banks (151 to 153) of **FIG. 2**. The crossbar (123) in **FIG. 3** provides the vectors from the maps banks (151) to the vector-vector units (161 to 163) respectively. A same vector stored in the kernel buffer (131) is provided to the vector-vector units (161 to 163).

[0066] The vector-vector units (161 to 163) operate concurrently to compute the operation of the corresponding vector operands, stored in the maps banks (151 to 153) respectively, multiplied by the same vector operand that is stored in the kernel buffer (131). For example, the vector-vector unit (161) performs the multiplication operation on the vector operand stored in the maps bank (151) and the vector operand stored in the kernel buffer (131), while the vector-vector unit (163) is concurrently performing the multiplication operation on the vector operand stored in the maps bank (153) and the vector operand stored in the kernel buffer (131).

[0067] When the matrix-vector unit (141) of **FIG. 3** is implemented in a matrix-matrix unit (121) of **FIG. 2**, the matrix-vector unit (141) can use the maps banks (151 to 153), the crossbar (123) and the kernel buffer (131) of the matrix-matrix unit (121).

[0068] Each of the vector-vector units (161 to 163) in **FIG. 3** can be implemented in a way as illustrated in **FIG. 4**.

[0069] **FIG. 4** shows a processing unit configured to perform vector-vector operations according to one embodiment. For example, the vector-vector unit (161) of **FIG. 4** can be used as any of the vector-vector units in the matrix-vector unit (141)

of FIG. 3.

[0070] In FIG. 4, the vector-vector unit (161) has multiple multiply-accumulate units (171 to 173). Each of the multiply-accumulate units (e.g., 173) can receive two numbers as operands, perform multiplication of the two numbers, and add the result of the multiplication to a sum maintained in the multiply-accumulate unit.

[0071] Each of the vector buffers (181 and 183) stores a list of numbers. A pair of numbers, each from one of the vector buffers (181 and 183), can be provided to each of the multiply-accumulate units (171 to 173) as input. The multiply-accumulate units (171 to 173) can receive multiple pairs of numbers from the vector buffers (181 and 183) in parallel and perform the multiply-accumulate (MAC) operations in parallel. The outputs from the multiply-accumulate units (171 to 173) are stored into the shift register (175); and an accumulator (177) computes the sum of the results in the shift register (175).

[0072] When the vector-vector unit (161) of FIG. 4 is implemented in a matrix-vector unit (141) of FIG. 3, the vector-vector unit (161) can use a maps bank (e.g., 151 or 153) as one vector buffer (181), and the kernel buffer (131) of the matrix-vector unit (141) as another vector buffer (183).

[0073] The vector buffers (181 and 183) can have a same length to store the same number/count of data elements. The length can be equal to, or the multiple of, the count of multiply-accumulate units (171 to 173) in the vector-vector unit (161). When the length of the vector buffers (181 and 183) is the multiple of the count of multiply-accumulate units (171 to 173), a number of pairs of inputs, equal to the count of the multiply-accumulate units (171 to 173), can be provided from the vector buffers (181 and 183) as inputs to the multiply-accumulate units (171 to 173) in each iteration; and the vector buffers (181 and 183) feed their elements into the multiply-accumulate units (171 to 173) through multiple iterations.

[0074] In one embodiment, the communication bandwidth of the connection (119) between the Deep Learning Accelerator (103) and the random access memory (105) is sufficient for the matrix-matrix unit (121) to use portions of the random access memory (105) as the maps banks (151 to 153) and the kernel buffers (131 to 133).

[0075] In another embodiment, the maps banks (151 to 153) and the kernel buffers (131 to 133) are implemented in a portion of the local memory (115) of the Deep Learning Accelerator (103). The communication bandwidth of the connection (119) between the Deep Learning Accelerator (103) and the random access memory

(105) is sufficient to load, into another portion of the local memory (115), matrix operands of the next operation cycle of the matrix-matrix unit (121), while the matrix-matrix unit (121) is performing the computation in the current operation cycle using the maps banks (151 to 153) and the kernel buffers (131 to 133) implemented in a different portion of the local memory (115) of the Deep Learning Accelerator (103).

[0076] FIG. 5 shows a Deep Learning Accelerator and random access memory configured to autonomously apply inputs to a trained Artificial Neural Network according to one embodiment.

[0077] An Artificial Neural Network (201) that has been trained through machine learning (e.g., deep learning) can be described in a standard format (e.g., Open Neural Network Exchange (ONNX)). The description of the trained Artificial Neural Network (201) in the standard format identifies the properties of the artificial neurons and their connectivity.

[0078] In FIG. 5, a Deep Learning Accelerator compiler (203) converts trained Artificial Neural Network (201) by generating instructions (205) for a Deep Learning Accelerator (103) and matrices (207) corresponding to the properties of the artificial neurons and their connectivity. The instructions (205) and the matrices (207) generated by the DLA compiler (203) from the trained Artificial Neural Network (201) can be stored in random access memory (105) for the Deep Learning Accelerator (103).

[0079] For example, the random access memory (105) and the Deep Learning Accelerator (103) can be connected via a high bandwidth connection (119) in a way as in the integrated circuit device (101) of FIG. 1. The autonomous computation of FIG. 5 based on the instructions (205) and the matrices (207) can be implemented in the integrated circuit device (101) of FIG. 1. Alternatively, the random access memory (105) and the Deep Learning Accelerator (103) can be configured on a printed circuit board with multiple point to point serial buses running in parallel to implement the connection (119).

[0080] In FIG. 5, after the results of the DLA compiler (203) are stored in the random access memory (105), the application of the trained Artificial Neural Network (201) to process an input (211) to the trained Artificial Neural Network (201) to generate the corresponding output (213) of the trained Artificial Neural Network (201) can be triggered by the presence of the input (211) in the random access memory (105), or another indication provided in the random access memory (105).

[0081] In response, the Deep Learning Accelerator (103) executes the instructions (205) to combine the input (211) and the matrices (207). The matrices (207) can include kernel matrices to be loaded into kernel buffers (131 to 133) and maps matrices to be loaded into maps banks (151 to 153). The execution of the instructions (205) can include the generation of maps matrices for the maps banks (151 to 153) of one or more matrix-matrix units (e.g., 121) of the Deep Learning Accelerator (103).

[0082] In some embodiments, the inputs to Artificial Neural Network (201) is in the form of an initial maps matrix. Portions of the initial maps matrix can be retrieved from the random access memory (105) as the matrix operand stored in the maps banks (151 to 153) of a matrix-matrix unit (121). Alternatively, the DLA instructions (205) also include instructions for the Deep Learning Accelerator (103) to generate the initial maps matrix from the input (211).

[0083] According to the DLA instructions (205), the Deep Learning Accelerator (103) loads matrix operands into the kernel buffers (131 to 133) and maps banks (151 to 153) of its matrix-matrix unit (121). The matrix-matrix unit (121) performs the matrix computation on the matrix operands. For example, the DLA instructions (205) break down matrix computations of the trained Artificial Neural Network (201) according to the computation granularity of the Deep Learning Accelerator (103) (e.g., the sizes/dimensions of matrices that loaded as matrix operands in the matrix-matrix unit (121)) and applies the input feature maps to the kernel of a layer of artificial neurons to generate output as the input for the next layer of artificial neurons.

[0084] Upon completion of the computation of the trained Artificial Neural Network (201) performed according to the instructions (205), the Deep Learning Accelerator (103) stores the output (213) of the Artificial Neural Network (201) at a pre-defined location in the random access memory (105), or at a location specified in an indication provided in the random access memory (105) to trigger the computation.

[0085] When the technique of **FIG. 5** is implemented in the integrated circuit device (101) of **FIG. 1**, an external device connected to the memory controller interface (107) can write the input (211) into the random access memory (105) and trigger the autonomous computation of applying the input (211) to the trained Artificial Neural Network (201) by the Deep Learning Accelerator (103). After a period of time, the output (213) is available in the random access memory (105); and

the external device can read the output (213) via the memory controller interface (107) of the integrated circuit device (101).

[0086] For example, a predefined location in the random access memory (105) can be configured to store an indication to trigger the autonomous execution of the instructions (205) by the Deep Learning Accelerator (103). The indication can optionally include a location of the input (211) within the random access memory (105). Thus, during the autonomous execution of the instructions (205) to process the input (211), the external device can retrieve the output generated during a previous run of the instructions (205), and/or store another set of input for the next run of the instructions (205).

[0087] Optionally, a further predefined location in the random access memory (105) can be configured to store an indication of the progress status of the current run of the instructions (205). Further, the indication can include a prediction of the completion time of the current run of the instructions (205) (e.g., estimated based on a prior run of the instructions (205)). Thus, the external device can check the completion status at a suitable time window to retrieve the output (213).

[0088] In some embodiments, the random access memory (105) is configured with sufficient capacity to store multiple sets of inputs (e.g., 211) and outputs (e.g., 213). Each set can be configured in a predetermined slot/area in the random access memory (105).

[0089] The Deep Learning Accelerator (103) can execute the instructions (205) autonomously to generate the output (213) from the input (211) according to matrices (207) stored in the random access memory (105) without helps from a processor or device that is located outside of the integrated circuit device (101).

[0090] In a method according to one embodiment, random access memory (105) of a computing device (e.g., integrated circuit device (101)) can be accessed using an interface (107) of the computing device to a memory controller. The computing device can have processing units (e.g., 111) configured to perform at least computations on matrix operands, such as a matrix operand stored in maps banks (151 to 153) and a matrix operand stored in kernel buffers (131 to 133).

[0091] For example, the computing device, implemented using the integrated circuit device (101) and/or other components, can be enclosed within an integrated circuit package; and a set of connections can connect the interface (107) to the memory controller that is located outside of the integrated circuit package.

[0092] Instructions (205) executable by the processing units (e.g., 111) can be written into the random access memory (105) through the interface (107).

[0093] Matrices (207) of an Artificial Neural Network (201) can be written into the random access memory (105) through the interface (107). The matrices (207) identify the parameters, the property and/or the state of the Artificial Neural Network (201).

[0094] Optionally, at least a portion of the random access memory (105) is non-volatile and configured to store the instructions (205) and the matrices (07) of the Artificial Neural Network (201).

[0095] First input (211) to the Artificial Neural Network can be written into the random access memory (105) through the interface (107).

[0096] An indication is provided in the random access memory (105) to cause the processing units (111) to start execution of the instructions (205). In response to the indication, the processing units (111) execute the instructions to combine the first input (211) with the matrices (207) of the Artificial Neural Network (201) to generate first output (213) from the Artificial Neural Network (201) and store the first output (213) in the random access memory (105).

[0097] For example, the indication can be an address of the first input (211) in the random access memory (105); and the indication can be stored a predetermined location in the random access memory (105) to cause the initiation of the execution of the instructions (205) for the input (211) identified by the address. Optionally, the indication can also include an address for storing the output (213).

[0098] The first output (213) can be read, through the interface (107), from the random access memory (105).

[0099] For example, the computing device (e.g., integrated circuit device (101)) can have a Deep Learning Accelerator (103) formed on a first integrated circuit die and the random access memory (105) formed on one or more second integrated circuit dies. The connection (119) between the first integrated circuit die and the one or more second integrated circuit dies can include Through-Silicon Vias (TSVs) to provide high bandwidth for memory access.

[0100] For example, a description of the Artificial Neural Network (201) can be converted using a compiler (203) into the instructions (205) and the matrices (207). The combination of the instructions (205) and the matrices (207) stored in the random access memory (105) and the Deep Learning Accelerator (103) provides an

autonomous implementation of the Artificial Neural Network (201) that can automatically convert input (211) to the Artificial Neural Network (201) to its output (213).

[0101] For example, during a time period in which the Deep Learning Accelerator (103) executes the instructions (205) to generate the first output (213) from the first input (211) according to the matrices (207) of the Artificial Neural Network (201), the second input to Artificial Neural Network (201) can be written into the random access memory (105) through the interface (107) at an alternative location. After the first output (213) is stored in the random access memory (105), an indication can be provided in the random access memory to cause the Deep Learning Accelerator (103) to again start the execution of the instructions and generate second output from the second input.

[0102] During the time period in which the Deep Learning Accelerator (103) executes the instructions (205) to generate the second output from the second input according to the matrices (207) of the Artificial Neural Network (201), the first output (213) can be read from the random access memory (105) through the interface (107); and a further input can be written into the random access memory to replace the first input (211), or written at a different location. The process can be repeated for a sequence of inputs.

[0103] The Deep Learning Accelerator (103) can include at least one matrix-matrix unit (121) that can execute an instruction on two matrix operands. The two matrix operands can be a first matrix and a second matrix. Each of two matrices has a plurality of vectors. The matrix-matrix unit (121) can include a plurality of matrix-vector units (141 to 143) configured to operate in parallel. Each of the matrix-vector units (141 to 143) are configured to operate, in parallel with other matrix-vector units, on the first matrix and one vector from second matrix. Further, each of the matrix-vector units (141 to 143) can have a plurality of vector-vector units (161 to 163) configured to operate in parallel. Each of the vector-vector units (161 to 163) is configured to operate, in parallel with other vector-vector units, on a vector from the first matrix and a common vector operand of the corresponding matrix-vector unit. Further, each of the vector-vector units (161 to 163) can have a plurality of multiply-accumulate units (171 to 173) configured to operate in parallel.

[0104] The Deep Learning Accelerator (103) can have local memory (115) and a control unit (113) in addition to the processing units (111). The control unit (113) can

load instructions (205) and matrix operands (e.g., some of the matrices (207)) from the random access memory (105) for execution by the processing units (111). The local memory can cache matrix operands used by the matrix-matrix unit. The connection (119) can be configured with a bandwidth sufficient to load a set of matrix operands from the random access memory (105) to the local memory (115) during a time period in which the matrix-matrix unit performs operations on two other matrix operands. Further, during the time period, the bandwidth is sufficient to store a result, generated by the matrix-matrix unit (121) in a prior instruction execution, from the local memory (115) to the random access memory (105).

[0105] At least some embodiments disclosed herein provides a compiler that can convert a same description of an Artificial Neural Network into different sets of instructions executable on different hardware platforms of Deep Learning Accelerators.

[0106] Deep Learning Accelerators can be implemented using different integrated circuit technologies, such as Field-Programmable Gate Array (FPGA) or Application Specific Integrated circuit (ASIC). Further, Deep Learning Accelerators can have different hardware capabilities in implementing matrix operations.

[0107] For example, different hardware implementations of Deep Learning Accelerators can have different numbers of parallel processing units operable to perform matrix operations concurrently.

[0108] For example, different hardware implementations of Deep Learning Accelerators can have different matrix computation granularity. An instruction can be used to perform a predefined matrix operation on matrix operands. However, the dimensional sizes of the matrix operands of the instruction can vary from one Deep Learning Accelerator to another.

[0109] In one embodiment, a compiler is configured to initially perform platform-agnostic compilation and optimization for a generic Deep Learning Accelerator. The hardware capability of the generic Deep Learning Accelerator is predefined to capture the common characteristics of a number of different Deep Learning Accelerators. The compilation result for the generic Deep Learning Accelerator can be mapped into compilation results for different Deep Learning Accelerators. Thus, the same description of the Artificial Neural Network can be compiled into different sets of instructions executable on different Deep Learning Accelerators that are implemented using different integrated circuit technologies (e.g., FPGA or ASIC)

and/or with different granularity and parallel execution capabilities. Optionally, the compiler can further optimize the compilation result for individual types of Deep Learning Accelerators to further reduce energy consumption and/or computation time.

[0110] FIG. 6 shows a technique to generate instructions executable by a Deep Learning Accelerator to implement an Artificial Neural Network according to one embodiment.

[0111] In FIG. 6, an ANN description (221) identifies the parameters of an Artificial Neural Network (201), including the behavior models of artificial neurons and the connectivity of the artificial neurons in the network. For example, the parameters can include the identifications of activation functions, biases, and/or states of the artificial neurons. For example, the parameters can include synaptic weights for connections among the artificial neurons. The description (221) in a standard format (e.g., Open Neural Network Exchange (ONNX)) can be provided as an input to a DLA compiler (203).

[0112] The DLA compiler (203) can perform compilation and optimization (223) according to a generic DLA specification (225). The generic DLA specification (225) identifies the computation capability of a generic Deep Learning Accelerator.

[0113] For example, the generic Deep Learning Accelerator can have common hardware features of many Deep Learning Accelerators that may be implemented using different technologies, with different granularity, and with different capacities.

[0114] For example, the generic Deep Learning Accelerator can be constructed as a virtual Deep Learning Accelerator to be implemented on a particular hardware platform of Deep Learning Accelerators.

[0115] For example, the generic Deep Learning Accelerator can be a platform agnostic characterization of a class of Deep Learning Accelerators that can be implemented via ASIC, FPGA, or another technology.

[0116] The DLA compiler (203) generates generic result (227) through compilation and optimization (223) for the generic Deep Learning Accelerator. For example, the generic result (227) can include the instructions for implementing the matrix computations of the Artificial Neural Network (201) on a generic or virtual Deep Learning Accelerator that is in compliance with the generic DLA specification (225).

[0117] The DLA compiler (203) can further perform DLA mapping (233) that maps

the generic result (227) into a compiler output (237) for a specific hardware platform of Deep Learning Accelerators. A specific DLA specification (235) identifies the hardware capabilities of the specific hardware platform of Deep Learning Accelerators. The compiler output (237) includes DLA instructions (205) executable on a Deep Learning Accelerator (103) that is in compliance with the specific DLA specification (235). The compiler output (237) further includes DLA matrices (207) that are representative of the parameters of the Artificial Neural Network (201).

[0118] Optionally, some aspects of the generic Deep Learning Accelerator can be parameterized, such as the number of processing units of a predetermined type operable to process data in parallel, the processing granularity of the processing units, etc. Thus, such aspects of the generic Deep Learning Accelerator can be configured for the compilation and optimization (223) to generate a generic result (227) for an optimized result in matching with the specific DLA specification (235) through the DLA mapping (233).

[0119] DLA compiler (203) can map the generic results (227) compiled for a generic Deep Learning Accelerator into the compiler output (237) for a specific platform of Deep Learning Accelerator by implementations of instructions and/or routines of the generic Deep Learning Accelerator using instructions and routines of the specific platform.

[0120] FIGS. 7 and 8 illustrate techniques to map a compilation result for a generic Deep Learning Accelerator into instructions executable by a specific Deep Learning Accelerator to implement an Artificial Neural Network according to one embodiment.

[0121] FIG. 7 illustrates a technique of using DLA routines (e.g., 243) to map the instructions of a generic Deep Learning Accelerator to DLA instructions (205) executable on a hardware platform specified or identified by the specific DLA specification (235).

[0122] For example, a generic DLA instruction (241) can be implemented using a DLA routine (243) executable in a specific hardware platform. The use of the generic DLA instruction (241) in the generic result (227) of compilation can be replaced with the use of the DLA routine (243) configured according to the specific DLA specification (235) of the specific hardware platform.

[0123] For example, the DLA routine (243) can be pre-optimized for the implementation of the generic DLA instruction (241) on the hardware platform having

the specific DLA specification (235).

[0124] In **FIG. 8**, a generic routine (245) implemented using instructions according to the generic DLA specification (225) is mapped to a DLA routine (247) implemented using instructions according to the specific DLA specification (225). The DLA routine (247) can be pre-optimized to improve the performance of the overall task performed by the routines, such that the performance of the DLA routine (247) is better than replacing the corresponding generic DLA instructions (e.g., 241) in the generic routine (245) with corresponding DLA routines (e.g., 243).

[0125] In general, different routines or combinations of instructions in the generic result (227) can have different weights in their contribution to the performance of the generic result (227) of compilation in implementing the computation of the Artificial Neural Network (201). Routines or instruction combinations that have larger shares of computation workloads can be mapped to optimized DLA routines (e.g., 247) to improve the performance of the compiler output (237).

[0126] Optionally, after the DLA mapping (233), the DLA compiler (203) can further perform further optimization to improve the performance of the compiler output (237), as illustrated in **FIG. 9**.

[0127] **FIG. 9** shows another technique to generate instructions executable by a Deep Learning Accelerator to implement an Artificial Neural Network according to one embodiment.

[0128] In **FIG. 9**, the DLA compiler (203) can perform initial compilation and optimization (223) of the Artificial Neural Network (201) based on the ANN description (221) and the generic DLA specification (225), in a way similar to **FIG. 6**. Further, the DLA compiler (203) can perform the DLA mapping (233) to convert the generic result (227) of compilation into a mapped result (229) for implementation according to a specific DLA specification (235). The DLA mapping (233) can be performed using the techniques of **FIGS. 7 and 8**.

[0129] After the DLA mapping (233), the DLA compiler (203) can further perform optimization (231) of the mapped result (229) of compilation to generate the compiler output (237). For example, the DLA compiler (203) can transform the mapped result (229) to reduce energy consumption and/or computation time for the implementation of the ANN description (221) on a platform identified by the specific DLA specification (235).

[0130] In a method according to one embodiment, a compiler converts a

description of an Artificial Neural Network into instructions for implementation on a Deep Learning Accelerator. For example, the method can be implemented on a computing device to generate DLA instructions (205) and DLA matrices (207) for the implementation of the matrix computations of Artificial Neural Network (201) in an integrated circuit device (101) illustrated in **FIG. 1** or a system illustrated in **FIG. 5**.

[0131] After the computing device receives a description (221) of an artificial neural network (201), the computing device generates a first result of compilation from the description (221) of the artificial neural network (201) according to a specification of a first device.

[0132] For example, the specification of the first device can be a generic DLA specification (225); and the first result of compilation can be the generic result (227), illustrated in **FIGS. 6 – 9**, that is the result of DLA compiler (203) performing compilation and optimization (223) according to the generic DLA specification (225).

[0133] The first result can include first data representative of first instructions executable on the first device to implement matrix computations of the artificial neural network (201) in accordance with the specification of the first device.

[0134] For example, the first instructions executable on the first device can include generic DLA instructions (e.g., 241) and/or generic routines (e.g., 245) used in the generic results (227) to implement the computation of the Artificial Neural Network (201) on a generic Deep Learning Accelerator. The generic Deep Learning Accelerator can be a virtual device in accordance with the generic DLA specification (225), or a reference implementation of the generic DLA specification (225).

[0135] The computing device maps the first result of compilation into a second result according to a specification of a second device.

[0136] For example, the specification of the second device can be a specific DLA specification (235); and the second result can be the compiler output (237) illustrated in **FIG. 7**, or the mapped result (229) illustrated in **FIG. 8**. For example, the second device can be an integrated circuit device (101) of **FIG. 8** having matrix processing units illustrated in **FIGS. 2 – 4**.

[0137] The second result can include second data representative of second instructions executable on the second device to implement matrix computations of the artificial neural network (201).

[0138] For example, the second instructions can be the DLA instructions (205) in accordance with the specific DLA specification (235). The second instructions can

include DLA routines (e.g., 243 and/or 247).

[0139] The computing device can further generate, from the description (221) of the artificial neural network (201), third data representative of parameters of the artificial neural network (201).

[0140] For example, the third data representative of parameters of the artificial neural network (201) can include DLA matrices (207). Some of the DLA matrices (207) can be loaded into the kernel buffers (131 to 133) in a processing unit (111) of an integrated circuit device (101). Some of the DLA matrices (207) can be loaded into the maps bank (151 to 153) in the processing unit (111) of the integrated circuit device (101).

[0141] For example, the second device can be the integrated circuit device (101) of **FIG. 1** that has random access memory (105) configured to store the third data representative of the parameters of the artificial neural network and the second data representative of the second instructions. The integrated circuit device (101) of **FIG. 1** further includes at least one processing unit (111) configured to execute the second instructions to generate an output (213) of the artificial neural network (201) based on the third data representative of the parameters of the artificial neural network (201) and fourth data representative of an input (211) to the artificial neural network (201).

[0142] As illustrated in **FIGS. 7** and **8**, the mapping of the first result into the second result can include mapping an instruction executable by the first device in the first result into a routine executable by the second device in the second result. For example, a generic DLA instruction (241) in the generic result (227) can be mapped to a DLA routine (243) executable by a Deep Learning Accelerator (103) of a specific platform identified by the specific DLA specification (235). Preferably, the DLA routine (243) can be pre-optimized to perform a task defined by the generic DLA instruction (241).

[0143] As illustrated in **FIG. 8**, the mapping of the first result into the second result can include mapping a combination of instructions in the first result executable by the first device into a routine in the second result executable by the second device. For example, the combination of the instructions can be a generic routine (245) that is mapped to a corresponding DLA routine (247) during the operation of DLA mapping (233). Preferably, the corresponding DLA routine (247) can be pre-optimized to perform a task defined by the combination of the instructions (e.g., the generic

routine (245)).

[0144] Optionally, as illustrated in **FIG. 9**, the DLA compiler (203) can further transform the second result into a third result having fifth data representative of third instructions executable in the second device.

[0145] For example, the second result can include the mapped result (229) illustrated in **FIG. 9**; and the third result can be the compiler output (237) illustrated in **FIG. 9**. The DLA compiler (203) performs optimization (231) in the transformation such that, when executed in the Deep Learning Accelerator (103) that is in accordance with, or in compliance with, the specific DLA specification (235), the DLA instructions (205) compiled in the compiler output (237) has better performance than the instructions compiled in the mapped result (229).

[0146] Optionally, the computing device can store, into the random access memory (105) of the integrated circuit device (101), the third data representative of the parameters of the artificial neural network (201) and the second data representative of the second instructions (or the fifth data representative of third instructions). Further, the computing device, or another device, can store, into the random access memory (105) of the integrated circuit device (101), the fourth data representative of the input (211) to the artificial neural network (201) to cause the integrated circuit device (101) to execute the second instructions (or third instructions) and generate the output (213) of the artificial neural network (201).

[0147] **FIG. 10** shows an integrated circuit device having a Deep Learning Accelerator with configurable hardware capabilities with random access memory configured according to one embodiment.

[0148] In **FIG. 10**, the processing units (111) can be used in different configurations. Different configurations of the processing units (111) provide different tradeoffs in functionality, efficiency, performance, and/or energy consumption.

[0149] A set of registers (251) is provided to control the circuit configuration that is currently available to perform computations. A set of hardware options (253) specified in the registers (251) selects the circuit configuration used by the processing units (111) in executing instructions for matrix operations.

[0150] When a set of hardware options (253) is stored in the registers (251), the processing units (111) are configured to operate according to a selected one of multiple designs of circuit configurations during data processing. When another set

of hardware options (253) is stored in the registers (251), the processing units (111) are configured to operate according to another one in the multiple designs. Thus, at least some computation aspects of the processing units (111) can be configured or selectively used by specifying hardware options (253) in the registers.

[0151] For example, a Deep Learning Accelerator (103) of in one embodiment can have a hardware option configured to control the granularity of matrix computation.

[0152] For example, based on a hardware option specified in the registers (251), a vector-vector unit (e.g., 161) can be configured to compute the sum of the products of elements from its vector buffers, or the sum of the products of first half of elements from its vector buffers, or the sum of the products of second half of elements from its vector buffers, or a combination thereof.

[0153] For example, based on a hardware option specified in the registers (251), a matrix-vector unit (e.g., 141) can be configured to compute the product of a matrix (e.g., as stored in a set of maps banks (151, ..., 153) and a vector (e.g., as stored in a kernel buffer (131)), or the product of a portion of the matrix and a portion of the vector, or the product of another portion of the matrix and another portion of the vector, or a combination thereof.

[0154] For example, based on a hardware option specified in the registers (251), a matrix-vector unit (e.g., 141) can be configured to compute the product of a matrix (e.g., as stored in a set of maps banks (151, ..., 153) and a vector (e.g., as stored in a kernel buffer (131)), or the product of a portion of the matrix and a portion of the vector, or the product of another portion of the matrix and another portion of the vector, or a combination thereof.

[0155] For example, based on a hardware option specified in the registers (251), a matrix-matrix unit (e.g., 121) can be configured to compute the product of a matrix (e.g., as stored in a set of maps banks (151, ..., 153) and another matrix (e.g., as stored in a set of kernel buffers (131, ..., 133)), or the product of portions of the matrices, or the product of alternative portions of the matrices, or a combination thereof.

[0156] Thus, the hardware options (253) can be used to adjust the granularity level of the processing units (111) and organize the concurrent executions of parallel units, such as matrix-vector units (141, ..., 143) in a matrix-matrix unit (121), vector-vector units (161, ..., 163) in a matrix-vector unit (141), and/or multiply-accumulate

units (171, ..., 173) in a vector-vector unit (161).

[0157] When different sets of options are specified in the registers (251), the processing units (111) are effectively configured to have different hardware capabilities for matrix computation.

[0158] **FIG. 11** illustrates different hardware configurations of a processing unit of a Deep Learning Accelerator that are configurable via options stored in register according to one embodiment. For example, the processing unit of **FIG. 11** can be used in the Deep Learning Accelerator (103) of **FIGS. 1, 5**, and/or **10**.

[0159] In **FIG. 11**, a processing unit (255) is controlled by a register (257).

[0160] In configuration A (265), option A (261) is specified in the register (257). The option A (261) in the register (257) causes the processing unit (255) to function as processing unit A (259) in functionality and/or performance.

[0161] When the option A (261) in the register (257) is changed to option B (263), the combination of the processing unit (255) and the register (257) is in configuration B (267). The option B (263) in the register (257) causes the processing unit (255) to function as processing unit B (269) in functionality and/or performance.

[0162] The processing unit A (259) and the processing unit B (269) are different in functionality and/or performance. The use of the processing unit A (259) can be advantageous over the use of the processing unit B (269) in some computation tasks or scenarios, but not in other computation tasks or scenarios. The options (261 and 263) can be selectively stored in the register (257) to selectively configure or convert the processing unit (255) into the processing unit A (259) and the processing unit B (269). Thus, the processing units (259 and 269) can be selectively deployed in different configurations (265 and 267) for different computation tasks or scenarios.

[0163] **FIG. 12** illustrates a technique to generate instructions executable by a Deep Learning Accelerator with optimized hardware configuration to implement an Artificial Neural Network according to one embodiment.

[0164] A DLA compiler (203) initially converts an ANN description (221) into a generic result through compilation and optimization (223) according to a generic DLA specification.

[0165] For example, the ANN description (221) can identify various aspects of an Artificial Neural Network (201), including the behavior models of artificial neurons and the connectivity of the artificial neurons in the network. The parameters used in the ANN description (221) can include the identifications of activation functions,

biases, and/or states of the artificial neurons. Further, the parameters can include synaptic weights for connections among the artificial neurons. The description (221) can be specified in a standard format (e.g., Open Neural Network Exchange (ONNX)) and provided as an input to the DLA compiler (203).

[0166] The generic DLA specification (225) identifies the computation capability of a generic Deep Learning Accelerator. Thus, the compilation and optimization (223) is independent of hardware platforms or capabilities of Deep Learning Accelerators.

[0167] The generic result (227) can include the instructions for implementing the matrix computations of the Artificial Neural Network (201) on a generic or virtual Deep Learning Accelerator that is in compliance with the generic DLA specification (225).

[0168] Subsequently, the DLA compiler (203) can map the generic result (227) into a mapped result (229) through an operation of DLA mapping (233). The DLA mapping (233) is in accordance with a specific DLA specification (235) that identifies the hardware capabilities of a specific hardware platform of Deep Learning Accelerators.

[0169] In **FIG. 12**, a Deep Learning Accelerator according to the specific DLA specification (235) has configurable hardware options (253), as illustrated in **FIGS. 10 and 11**. The DLA compiler (203) uses a default set of options in performing the DLA mapping (233) to convert the generic result (227) into the mapped results (229). Thus, the instructions in the mapped results (229) are configured to use a Deep Learning Accelerator (103) with the default set of hardware options (253) stored in its registers (251).

[0170] For example, the DLA mapping (233) can be performed using the techniques of **FIGS. 7 and 8**. DLA compiler (203) can map the generic results (227), compiled for a generic Deep Learning Accelerator, into the mapped result (229) executable on a specific platform of Deep Learning Accelerators configured using the default set of hardware options (253) in their registers (251).

[0171] After the DLA mapping (233), the DLA compiler (203) can further perform optimization (231) of the mapped result (229) of compilation to generate the compiler output (237). During the optimization (231), the DLA compiler (203) can selectively adjust the hardware options (253) to improve the performance of the Deep Learning Accelerator in implementing the Artificial Neural Network (201) specified by the ANN description (221).

[0172] For example, the DLA compiler (203) can perform the optimization (231) by reducing energy consumption and/or computation time used in executing the DLA instructions (205) in the Deep Learning Accelerator (103). The optimized set of hardware options (253) can be used to optimize the hardware of Deep Learning Accelerator (103) specifically for the implementation of the particular Artificial Neural Network (201) specified by the ANN description (221). The hardware optimization can be implemented after the manufacturing of the integrated circuit device (101) by storing the optimized set hardware options (253) in the registers (251).

[0173] For example, the DLA instructions (205) can include instructions to store the optimized set of hardware options (253) into the registers (251) during an initialization operation to configure the Deep Learning Accelerator (103) for the execution of the remaining portion of the DLA instructions (205).

[0174] In some implementations, the hardware options (253) can be adjusted during the execution of the DLA instructions (205). For example, a first portion of the DLA instructions (205) can be executed using a first set of hardware options (253) in the registers (251); and a second portion of the DLA instructions (205) can be executed using a second set of hardware options (253) in the register (251).

[0175] In some implementations, the DLA instructions (205) do not include instructions to change the content of the registers (251). In an operation to load the compiler output (237) into the Random Access Memory (105) of the Integrated Circuit Device (101) to configure the computation of the Artificial Neural Network (201), a host system of the Integrated Circuit Device (101) loads the hardware options (253) selected by the DLA compiler (203) into the registers (251). As a result, the Deep Learning Accelerator (103) is configured to execute the DLA instructions (205), optimized for the hardware options (253), to implement the computations of the Artificial Neural Network (201).

[0176] In a method according to one embodiment, a DLA compiler (203) is configured to generate and optimize instructions for a Deep Learning Accelerator having configurable hardware options. The method can be used to generate instructions and select hardware options for implementing the computation of an Artificial Neural Network (201) using a Deep Learning Accelerator (103) illustrated in **FIGS. 1, 5, and 10 – 11**.

[0177] In the method, a computing device receives a description (221) of an artificial neural network (201). Further, the computing device generates a first result

of compilation from the description (221) of the artificial neural network (201) according to a specification (235) of a first device.

[0178] For example, the first result can be the mapped result (229) illustrated in **FIG. 12** as a result of compilation and optimization (223) according to a generic DLA specification (225) and DLA mapping according to a specific DLA specification (235).

[0179] For example, the first device can at least one processing unit (e.g., 111, 141, or 255) configured to perform matrix computations and having hardware configurations (e.g., 265 and 267) selectable via at least one register (e.g., 257).

[0180] For example, a function of the processing unit (e.g., 111, 141, or 255) can be adjustable according to content stored in the at least one register (e.g., 257). As illustrated in **FIG. 11**, when a first set of hardware options (e.g., 261) is specified in the at least one register (e.g., 257), the processing unit (e.g., 111, 141, 255) is configured to perform a first function of a processing unit A (259); and when a second set of hardware options (e.g., 263) is specified in the at least one register (e.g., 257), the processing unit (e.g., 111, 141, 255) is configured to perform a second function of another processing unit B (259) that is different from the first function.

[0181] Further, the computing device transforms the first result of compilation into a second result to select hardware options (e.g., 253) of the first device.

[0182] For example, the second result can be the compiler output (237) illustrated in **FIG. 12**. The second result can include first data representative of parameters of the artificial neural network, such as the DLA matrices (207). The second result can further include second data representative of instructions executable by the at least one processing unit of the first device to generate an output (213) of the artificial neural network (201) responsive to third data representative of an input (211) to the artificial neural network (201). The second result can further include fourth data representative of the hardware options (e.g., 253) to be stored in the at least one register (e.g., 257) to configure the at least one processing unit (e.g., 111, 141, or 255).

[0183] In one implementation, the content in the at least one register (e.g., 251, 257) is updatable via execution of a portion of the instructions represented by the second data stored in the random access memory (105) connected to the Deep Learning Accelerator (103).

[0184] For example, at least one interface (e.g., 107) of the Integrated Circuit

Device (101) can be configured to receive the third data as the input (211) to the artificial neural network (201) and store the third data into the random access memory (105).

[0185] Content stored in the at least one register (251) can be updatable through the at least one interface (e.g., 107) prior to execution of the instructions (205) represented by the second data stored in the random access memory (105). Thus, the content in the at least one register (251) does not change during the execution of the DLA instructions (205) generated by the DLA compiler (203).

[0186] Alternatively, the content stored in the at least one register (251) can be dynamically updated via a portion of the Deep Learning Accelerator (103) instructions. For example, the processing unit (255) can operate on some of the DLA matrices (207) using the configuration A (265) and other DLA matrices (207) using the configuration B (267).

[0187] For example, a dimension of the two matrix operands of an instruction to be processed by the processing unit (255) can be configured according to the at least one register (257) for execution of the instruction in the processing unit (255).

[0188] Optionally, the DLA compiler (203) can optimize the instructions, generated to implement an Artificial Neural Network (201) in a Deep Learning Accelerator (103), based on not only hardware configuration options, but also hardware characteristics of the Deep Learning Accelerator (103) in general. Such hardware characteristics of the Deep Learning Accelerator (103) can include features, options, behaviors, performances, latency, and/or limitations of the Deep Learning Accelerator (DLA). The DLA compiler (203) can learn the hardware characteristics of a Deep Learning Accelerator (DLA) through interacting with the Deep Learning Accelerator (DLA). For example, the DLA compiler (203) can test the hardware of the Deep Learning Accelerator (103) based on the computing performance of executing different instructions, different combinations of instructions, different scenarios of instruction executions, in combination with different hardware configuration options. For example, the DLA compiler (203) can optionally exercise the hardware of a Deep Learning Accelerator (DLA) by running test sets of instructions to discover and/or measure the hardware capabilities, options, performance, latency, etc. Optionally, the DLA compiler (203) can use an Artificial Neural Network to determine features, options, behaviors, performances, latency, and/or limitations of the Deep Learning Accelerator (DLA) through machine learning.

The information about the hardware characteristics of the Deep Learning Accelerator (DLA) can be used in the optimization operation of the compiler in generating or customizing the instructions and/or a runtime for implementing an Artificial Neural Network (201) on the Deep Learning Accelerator (DLA).

[0189] FIG. 13 shows a technique to discover hardware characteristics of a Deep Learning Accelerator according to one embodiment.

[0190] In FIG. 13, a computing device (271) having a DLA compiler (203) is connected to an Integrated Circuit Device (101) to learn its hardware characteristics using test programs (273). The results of the execution of the test programs (273) can be retrieved from the Random Access Memory (105) as the responses (275) to the test programs. The responses (275) can be analyzed to identify the hardware characteristics (279) of the Deep Learning Accelerator (103).

[0191] For example, the DLA compiler (203) can generate one or more test programs (273) according to the DLA specification (235) of the Deep Learning Accelerator (103). Some of the test programs can be generated according to examples of Artificial Neural Networks (e.g., 201). Alternatively, or in combination, some of the test programs can be generated to exercise various processing units (111) of the Deep Learning Accelerator (103) without modeling a particular artificial neural network.

[0192] For example, the test programs (273) can be written to the Random Access Memory (105) in a way similar to the DLA instructions (205) and DLA matrices (207) of an implementation of an Artificial Neural Network (201) being written into the Random Access Memory (105) for execution by the Deep Learning Accelerator (103). Further, the execution of the test programs (273) can be triggered via writing an input to the Random Access Memory (105), in a way similar to providing the input (211) to the Artificial Neural Network (201) in the Random Access Memory (105) to trigger the execution of the DLA instructions (205) of an Artificial Neural Network (201).

[0193] The responses (275) can include timestamps of executions at various stages of the test programs (273). The timestamps can be used to measure the execution time of instructions between the time timestamps and thus the performance characteristics of executing particular instructions and/or particular sequences of instructions by the processing units (111) of the Deep Learning Accelerator (103).

[0194] In some implementations, the Integrated Circuit Device (101) is capable of generating timestamps during execution of instructions stored in the Random Access Memory (105); and the test programs can be configured by the DLA compiler (203) to include instructions to direct the Integrated Circuit Device (101) to generate the timestamp.

[0195] In some implementations, the computing device (271) generates a timestamp for triggering the execution of a test program (273) stored in the Random Access Memory (105) and a timestamp for receiving a response (275) resulting from the execution of the test program (273). The timestamps measure the performance of the Deep Learning Accelerator (103) in executing the test program (273). The performance measurements of different test programs (273) can be used to determine the performances of instruction differences in the test programs (273).

[0196] For example, two test programs can share a common set of instructions, while one of the test programs has an extra set of instructions. Thus, the difference in the execution times of the two test programs provides a performance measurement of the Deep Learning Accelerator (103) in executing the extra set of instructions. The extra set of instructions can be repeated multiple times in the corresponding test program to improve the accuracy of the performance measurement. Thus, multiple test programs can be constructed to measure the performance levels of various instructions executable in the processing units (111).

[0197] For example, when the Deep Learning Accelerator (103) have different hardware options (253), the DLA compiler (203) can generate different routines to perform a same task using different hardware options (253). The performance levels of the routines can be measured to identify the optimal routine and corresponding hardware options (253) for the task. The optimal routine can be configured in a runtime or library and used to implement an Artificial Neural Network (201) in the Deep Learning Accelerator (103).

[0198] Optionally, a DLA compiler (203) running in the computing device (271) is configured to identify the specification and/or the DLA hardware characteristics (279) of the Deep Learning Accelerator (103) connected to the computing device (271). For example, according to a set of specifications of known hardware platforms of Deep Learning Accelerators, the DLA compiler (203) can construct test programs (273) for execution by the Deep Learning Accelerator (103) connected to the computing device (271). Based on the computation results in the responses (275),

the DLA compiler (203) can determine which of the specifications the Deep Learning Accelerator (103) connected to the computing device (271) is in compliance with. Optionally, after the identification of the specification of the Deep Learning Accelerator (103), the DLA compiler (203) can construct a test program (273) to verify and confirm that functionality of the Deep Learning Accelerator (103) is currently in compliance with the specification.

[0199] In some instances, the hardware platform of the Deep Learning Accelerator (103) connected to the computing device (271) may be new to the DLA compiler (203). The Deep Learning Accelerator (103) is not manufactured according to any of the specifications known to the DLA compiler (203). In such a situation, the DLA compiler (203) may use the test programs to probe the functionality of the Deep Learning Accelerator (103). The probing operation can be used to discovery functionality that are similar to some of the known platforms. The result of the probing can generate a specification of at least a subset of the functionality of the Deep Learning Accelerator (103). When the subset is sufficient to implement the computation of an Artificial Neural Network (201), the DLA compiler (203) can generate a compiler output (237) for implementing the Artificial Neural Network (201), even though the full and/or official specification of the Deep Learning Accelerator (103) is not available.

[0200] **FIG. 14** illustrates a technique to generate instructions executable by a Deep Learning Accelerator and optimized according hardware characteristics of the Deep Learning Accelerator according to one embodiment.

[0201] In **FIG. 14**, DLA excises (277) are constructed and/or performed to determine DLA hardware characteristics (279) in supporting optimization (231) in compiling an ANN description (221) for a hardware platform of Deep Learning Accelerators identified by a specific DLA specification (235).

[0202] For example, the DLA excises (277) can include running test programs (273) in a Deep Learning Accelerator (103) according to the specification (235) and analyzing the responses (275) of the Deep Learning Accelerator (103) to the test programs (273), as illustrated in **FIG. 13**.

[0203] Optionally, the DLA hardware characteristics (279) obtained from the DLA excises (277) can be included in the specification (235) of the Deep Learning Accelerator (103).

[0204] In **FIG. 14**, a DLA compiler (203) initially converts an ANN description

(221) into a generic result (227) through compilation and optimization (223) according to a generic DLA specification (225).

[0205] The ANN description (221) can the behavior models of artificial neurons and the connectivity of the artificial neurons in the network. The ANN description (221) can include the identifications of activation functions, biases, and/or states of the artificial neurons. The ANN description (221) can include synaptic weights for connections among the artificial neurons.

[0206] The DLA compiler (203) converts the description (221) in a standard format (e.g., Open Neural Network Exchange (ONNX)) into the generic result (227) that can be implemented on a Deep Learning Accelerator in compliance with the generic DLA specification (225).

[0207] For example, the generic result (227) can include the DLA instructions and DLA matrices for implementing the computations of the Artificial Neural Network (201) on a virtual Deep Learning Accelerator configured according to the generic DLA specification (225). For example, the virtual Deep Learning Accelerator may be implemented via software running on a microprocessor and/or the computing device (271).

[0208] Subsequently, the DLA compiler (203) can map the generic result (227) into a mapped result (229) through an operation of DLA mapping (233). The DLA mapping (233) is in accordance with a specific DLA specification (235) that identifies the hardware capabilities of a specific hardware platform of Deep Learning Accelerators, such as a Deep Learning Accelerator (103) connected to the computing device (271) in which DLA compiler (203) is running.

[0209] In **FIG. 14**, a Deep Learning Accelerator according to the specific DLA specification (235) can optionally have configurable hardware options (253), as illustrated in **FIGS. 10** and **11**. The DLA compiler (203) uses a default set of options in performing the DLA mapping (233) to convert the generic result (227) into the mapped results (229). Thus, the instructions in the mapped results (229) are configured to use a Deep Learning Accelerator (103) with the default set of hardware options (253) stored in its registers (251).

[0210] For example, the DLA mapping (233) can be performed using the techniques of **FIGS. 7** and **8**. DLA compiler (203) can map the generic results (227), compiled for a generic Deep Learning Accelerator, into the mapped result (229) executable on a specific platform of Deep Learning Accelerators configured using

the default set of hardware options (253) in their registers (251).

[0211] After the DLA mapping (233), the DLA compiler (203) can further perform optimization (231) of the mapped result (229) of compilation to generate the compiler output (237). During the optimization (231), the DLA compiler (203) can selectively adjust the hardware options (253) and/or instruction sequences or routines according to the DLA hardware characteristics (279) to improve the performance of the Deep Learning Accelerator in implementing the Artificial Neural Network (201) specified by the ANN description (221).

[0212] For example, a computing task can be implemented using different sequences of instructions (optionally running using different hardware options). Based on the hardware characteristics (279), the DLA compiler (203) can calculate or estimate the time periods required to execute the different sequences and select a sequence that has the best performance among the different sequences. The DLA compiler (203) can optionally transform the instructions incrementally to reduce the calculated/estimated execution time during the optimization (231).

[0213] **FIG. 15** shows a method to compile instructions to implement an Artificial Neural Network on a Deep Learning Accelerator based on its hardware characteristics according to one embodiment.

[0214] For example, the method of **FIG. 15** can be implemented in a computing device (271) running a DLA compiler (203) illustrated in **FIG. 14**.

[0215] At block 301, a computing device (271) interacts with an integrated circuit device (101) to identify hardware characteristics (279) of the integrated circuit device (101) in performing matrix computations.

[0216] For example, the integrated circuit device (101) can have a Deep Learning Accelerator (103) illustrated in **FIGS. 1** or **10**. The Deep Learning Accelerator (103) can have one or more processing units (e.g., 111, 121, 141, 161) as illustrated in **FIGS. 1 – 4** to perform matrix computations.

[0217] For example, the hardware characteristics (279) of the integrated circuit device identify features, options, behaviors, performances, or latency of at least one processing unit (e.g., 111, 121, 141, 161) of the integrated circuit device (101), or any combination thereof.

[0218] For example, the computing device (271) can interact with the integrated circuit device (101) by generating and loading test programs into the integrated circuit device (101) and receiving responses (275) of executing the test programs in

the integrated circuit device (101), as illustrated in **FIG. 13**.

[0219] For example, the computing device (271) can determine the hardware characteristics (279) from the responses (275), including the computation results of the test programs (273) and/or timestamps associated with executions of the test programs (273) in the integrated circuit device (101).

[0220] For example, the computing device (271) can identify a specification of the integrated circuit device (101), among a plurality of predetermined specifications of integrated circuit devices, based on the responses (275) of executing the test programs in the integrated circuit device. For example, the identification can be made when the results of the test programs (273) match with the results generated by an integrated circuit device in compliance with the identified specification, but not the results generated by integrated circuit devices in compliance with other specifications.

[0221] Optionally, the computing device (271) can probe hardware options (253) of a processing unit (e.g., 111, 121, 141, 161) of the integrated circuit (101) through try and error and in view of the based on the responses (275) of executing the test programs (273) in the integrated circuit device (101).

[0222] For example, the computing device (271) may generate the test programs based specifications of a plurality of hardware platforms of integrated circuit devices configured to perform matrix computations. The computing device (271) can check responses (275) to determine which instructions worked and/or performed what functions with what latency, etc. in determining the hardware characteristics (279).

[0223] At block 303, the computing device (271) receives a description (221) of an artificial neural network (201).

[0224] At block 305, a DLA compiler (203) running in the computing device (271) generates a result of compilation from the description (221) of the artificial neural network (201) based at least in part on the hardware characteristics (279) of the integrated circuit device (101).

[0225] For example, the result of compilation can include first data representative of parameters of the artificial neural network (201), such as the DLA matrices (207). The result of compilation can further include second data representative of instructions (205) executable by the integrated circuit device (101) to generate an output (213) of the artificial neural network (201) based on the first data representative of parameters of the artificial neural network (201) and based on an

input (211) to the artificial neural network (201).

[0226] For example, the DLA compiler (203) can transform the result of compilation from the description (221) of the artificial neural network (201) to improve, based on the hardware characteristics (279) of the integrated circuit device (101), performance of the compilation result when compilation result is run or executed in the integrated circuit device (101). For example, as illustrated in **FIG. 14**, optimization (231) can be performed based on hardware options (253) and/or hardware characteristics (279).

[0227] For example, the computing device (271) can be implemented using a machine illustrated in **FIG. 16**.

[0228] **FIG. 16** illustrates an example machine of a computer system within which a set of instructions, for causing the machine to perform any one or more of the methodologies discussed herein, can be executed.

[0229] In some embodiments, the computer system of **FIG. 16** can implement a system of **FIG. 5** with an integrated circuit device (101) of **FIG. 1** having matrix processing units illustrated in **FIGS. 2 – 4**.

[0230] The computer system of **FIG. 16** can be used to perform the operations of a DLA Compiler (203) described with reference to **FIGS. 1 – 15** by executing instructions configured to perform the operations corresponding to the DLA Compiler (203).

[0231] In some embodiments, the machine can be connected (e.g., networked) to other machines in a Local Area Network (LAN), an intranet, an extranet, and/or the Internet. The machine can operate in the capacity of a server or a client machine in client-server network environment, as a peer machine in a peer-to-peer (or distributed) network environment, or as a server or a client machine in a cloud computing infrastructure or environment.

[0232] For example, the machine can be configured as a personal computer (PC), a tablet PC, a set-top box (STB), a Personal Digital Assistant (PDA), a cellular telephone, a web appliance, a server, a network router, a switch or bridge, or any machine capable of executing a set of instructions (sequential or otherwise) that specify actions to be taken by that machine. Further, while a single machine is illustrated, the term “machine” shall also be taken to include any collection of machines that individually or jointly execute a set (or multiple sets) of instructions to perform any one or more of the methodologies discussed herein.

[0233] The example computer system illustrated in **FIG. 16** includes a processing device (402), a main memory (404), and a data storage system (418), which communicate with each other via a bus (430). For example, the processing device (402) can include one or more microprocessors; the main memory can include read-only memory (ROM), flash memory, dynamic random access memory (DRAM), such as synchronous DRAM (SDRAM) or Rambus DRAM (RDRAM), static random access memory (SRAM), etc. The bus (430) can include, or be replaced with, multiple buses.

[0234] The processing device (402) in **FIG. 16** represents one or more general-purpose processing devices such as a microprocessor, a central processing unit, or the like. More particularly, the processing device can be a complex instruction set computing (CISC) microprocessor, reduced instruction set computing (RISC) microprocessor, very long instruction word (VLIW) microprocessor, or a processor implementing other instruction sets, or processors implementing a combination of instruction sets. The processing device (402) can also be one or more special-purpose processing devices such as an application specific integrated circuit (ASIC), a field programmable gate array (FPGA), a digital signal processor (DSP), a network processor, or the like. The processing device (402) is configured to execute instructions (426) for performing the operations discussed in connection with the DLA compiler (203). Optionally, the processing device (402) can include a Deep Learning Accelerator (103).

[0235] The computer system of **FIG. 16** can further include a network interface device (408) to communicate over a computer network (420).

[0236] Optionally, the bus (430) is connected to an integrated circuit device (101) that has a Deep Learning Accelerator (103) and Random Access Memory (105) illustrated in **FIG. 1** and/or **FIG. 10**. The compiler (203) can write its compiler output (237) into the Random Access Memory (105) of the integrated circuit device (101) to enable the Integrated Circuit Device (101) to perform matrix computations of an Artificial Neural Network (201) specified by the ANN description (221). Optionally, the compiler output (237) can be stored into the Random Access Memory (105) of one or more other integrated circuit devices (101) through the network interface device (408) and the computer network (420).

[0237] The data storage system (418) can include a machine-readable medium (424) (also known as a computer-readable medium) on which is stored one or more

sets of instructions (426) or software embodying any one or more of the methodologies or functions described herein. The instructions (426) can also reside, completely or at least partially, within the main memory (404) and/or within the processing device (402) during execution thereof by the computer system, the main memory (404) and the processing device (402) also constituting machine-readable storage media.

[0238] In one embodiment, the instructions (426) include instructions to implement functionality corresponding to a DLA Compiler (203), such as the DLA Compiler (203) described with reference to **FIGS. 5 – 15**. While the machine-readable medium (424) is shown in an example embodiment to be a single medium, the term “machine-readable storage medium” should be taken to include a single medium or multiple media that store the one or more sets of instructions. The term “machine-readable storage medium” shall also be taken to include any medium that is capable of storing or encoding a set of instructions for execution by the machine and that cause the machine to perform any one or more of the methodologies of the present disclosure. The term “machine-readable storage medium” shall accordingly be taken to include, but not be limited to, solid-state memories, optical media, and magnetic media.

[0239] The present disclosure includes methods and apparatuses which perform the methods described above, including data processing systems which perform these methods, and computer readable media containing instructions which when executed on data processing systems cause the systems to perform these methods.

[0240] A typical data processing system may include an inter-connect (e.g., bus and system core logic), which interconnects a microprocessor(s) and memory. The microprocessor is typically coupled to cache memory.

[0241] The inter-connect interconnects the microprocessor(s) and the memory together and also interconnects them to input/output (I/O) device(s) via I/O controller(s). I/O devices may include a display device and/or peripheral devices, such as mice, keyboards, modems, network interfaces, printers, scanners, video cameras and other devices known in the art. In one embodiment, when the data processing system is a server system, some of the I/O devices, such as printers, scanners, mice, and/or keyboards, are optional.

[0242] The inter-connect can include one or more buses connected to one another through various bridges, controllers and/or adapters. In one embodiment

the I/O controllers include a USB (Universal Serial Bus) adapter for controlling USB peripherals, and/or an IEEE-1394 bus adapter for controlling IEEE-1394 peripherals.

[0243] The memory may include one or more of: ROM (Read Only Memory), volatile RAM (Random Access Memory), and non-volatile memory, such as hard drive, flash memory, etc.

[0244] Volatile RAM is typically implemented as dynamic RAM (DRAM) which requires power continually in order to refresh or maintain the data in the memory. Non-volatile memory is typically a magnetic hard drive, a magnetic optical drive, an optical drive (e.g., a DVD RAM), or other type of memory system which maintains data even after power is removed from the system. The non-volatile memory may also be a random access memory.

[0245] The non-volatile memory can be a local device coupled directly to the rest of the components in the data processing system. A non-volatile memory that is remote from the system, such as a network storage device coupled to the data processing system through a network interface such as a modem or Ethernet interface, can also be used.

[0246] In the present disclosure, some functions and operations are described as being performed by or caused by software code to simplify description. However, such expressions are also used to specify that the functions result from execution of the code/instructions by a processor, such as a microprocessor.

[0247] Alternatively, or in combination, the functions and operations as described here can be implemented using special purpose circuitry, with or without software instructions, such as using Application-Specific Integrated Circuit (ASIC) or Field-Programmable Gate Array (FPGA). Embodiments can be implemented using hardwired circuitry without software instructions, or in combination with software instructions. Thus, the techniques are limited neither to any specific combination of hardware circuitry and software, nor to any particular source for the instructions executed by the data processing system.

[0248] While one embodiment can be implemented in fully functioning computers and computer systems, various embodiments are capable of being distributed as a computing product in a variety of forms and are capable of being applied regardless of the particular type of machine or computer-readable media used to actually effect the distribution.

[0249] At least some aspects disclosed can be embodied, at least in part, in

software. That is, the techniques may be carried out in a computer system or other data processing system in response to its processor, such as a microprocessor, executing sequences of instructions contained in a memory, such as ROM, volatile RAM, non-volatile memory, cache or a remote storage device.

[0250] Routines executed to implement the embodiments may be implemented as part of an operating system or a specific application, component, program, object, module or sequence of instructions referred to as “computer programs.” The computer programs typically include one or more instructions set at various times in various memory and storage devices in a computer, and that, when read and executed by one or more processors in a computer, cause the computer to perform operations necessary to execute elements involving the various aspects.

[0251] A machine readable medium can be used to store software and data which when executed by a data processing system causes the system to perform various methods. The executable software and data may be stored in various places including for example ROM, volatile RAM, non-volatile memory and/or cache.

Portions of this software and/or data may be stored in any one of these storage devices. Further, the data and instructions can be obtained from centralized servers or peer to peer networks. Different portions of the data and instructions can be obtained from different centralized servers and/or peer to peer networks at different times and in different communication sessions or in a same communication session. The data and instructions can be obtained in entirety prior to the execution of the applications. Alternatively, portions of the data and instructions can be obtained dynamically, just in time, when needed for execution. Thus, it is not required that the data and instructions be on a machine readable medium in entirety at a particular instance of time.

[0252] Examples of computer-readable media include but are not limited to non-transitory, recordable and non-recordable type media such as volatile and non-volatile memory devices, Read Only Memory (ROM), Random Access Memory (RAM), flash memory devices, floppy and other removable disks, magnetic disk storage media, optical storage media (e.g., Compact Disk Read-Only Memory (CD ROM), Digital Versatile Disks (DVDs), etc.), among others. The computer-readable media may store the instructions.

[0253] The instructions may also be embodied in digital and analog communication links for electrical, optical, acoustical or other forms of propagated

signals, such as carrier waves, infrared signals, digital signals, etc. However, propagated signals, such as carrier waves, infrared signals, digital signals, etc. are not tangible machine readable medium and are not configured to store instructions.

[0254] In general, a machine readable medium includes any mechanism that provides (i.e., stores and/or transmits) information in a form accessible by a machine (e.g., a computer, network device, personal digital assistant, manufacturing tool, any device with a set of one or more processors, etc.).

[0255] In various embodiments, hardwired circuitry may be used in combination with software instructions to implement the techniques. Thus, the techniques are neither limited to any specific combination of hardware circuitry and software nor to any particular source for the instructions executed by the data processing system.

[0256] The above description and drawings are illustrative and are not to be construed as limiting. Numerous specific details are described to provide a thorough understanding. However, in certain instances, well known or conventional details are not described in order to avoid obscuring the description. References to one or an embodiment in the present disclosure are not necessarily references to the same embodiment; and, such references mean at least one.

[0257] In the foregoing specification, the disclosure has been described with reference to specific exemplary embodiments thereof. It will be evident that various modifications may be made thereto without departing from the broader spirit and scope as set forth in the following claims. The specification and drawings are, accordingly, to be regarded in an illustrative sense rather than a restrictive sense.

CLAIMS

What is claimed is:

1. A method, comprising:
transmitting, by a computing device, one or more commands to an integrated circuit device, each command associated with one or more requests for a response indicative of hardware characteristics of the integrated circuit device in performing matrix computations;
receiving, at the computing device and after determination of the hardware characteristics of the integrated circuit device, data representative of a description of an artificial neural network; and
generating, by the computing device, a result of compilation from the data representative of the description of the artificial neural network based at least in part on the hardware characteristics of the integrated circuit device, the result of compilation including first data representative of parameters of the artificial neural network and second data representative of instructions executable by the integrated circuit device to generate an output of the artificial neural network based on the first data and an input to the artificial neural network.
2. The method of claim 1, wherein the hardware characteristics of the integrated circuit device identify features, options, behaviors, performances, or latency of at least one processing unit of the integrated circuit device, or any combination thereof.
3. The method of claim 2, wherein the transmitting of the one or more commands to the integrated circuit device includes loading test programs into the integrated circuit device to receive responses of executing the test programs in the integrated circuit device.
4. The method of claim 3, further comprising:
determining the hardware characteristics from the responses of executing the test programs in the integrated circuit device.

5. The method of claim 3, further comprising:
determining the hardware characteristics from timestamps associated with
executions of the test programs in the integrated circuit device.
6. The method of claim 3, further comprising:
identifying a specification of the integrated circuit device, among a plurality of
predetermined specifications of integrated circuit devices, based on the
responses of executing the test programs in the integrated circuit
device.
7. The method of claim 3, further comprising:
probing hardware options of a matrix processing unit of the integrated circuit
based on the responses of executing the test programs in the
integrated circuit device.
8. The method of claim 3, further comprising:
generating the test programs based specifications of a plurality of hardware
platforms of integrated circuit devices configured to perform matrix
computations.
9. The method of claim 3, further comprising:
transforming the result of compilation from the description of the artificial
neural network to improve, based on the hardware characteristics of
the integrated circuit device, performance of the result when executed
in the integrated circuit device.
10. A computing device, comprising:
memory; and
at least one microprocessor coupled to the memory and configured to probe
an integrated circuit device to identify hardware characteristics of the
integrated circuit device in performing matrix computations, and
generate a result of compilation from data representative of a
description of an artificial neural network based at least in part on the
hardware characteristics of the integrated circuit device, the result of
compilation including first data representative of parameters of the
artificial neural network and second data representative of instructions

executable by the integrated circuit device to generate an output of the artificial neural network based on the first data and an input to the artificial neural network.

11. The computing device of claim 10, wherein the hardware characteristics of the integrated circuit device identify features, options, behaviors, performances, or latency of at least one processing unit of the integrated circuit device, or any combination thereof.
12. The computing device of claim 10, wherein the at least one microprocessor is further configured to generate test programs, load the test programs into the integrated circuit device for execution, and retrieve results of executing the test programs from the integrated circuit device; and wherein the hardware characteristics of the integrated circuit device are identified from the results.
13. The computing device of claim 12, wherein the at least one microprocessor is further configured to perform optimization of the result of compilation based on the hardware characteristics of the integrated circuit device.
14. The computing device of claim 13, further comprising the integrated circuit device coupled with the at least one microprocessor.
15. The computing device of claim 14, wherein the integrated circuit device comprises an integrated circuit die of a Field-Programmable Gate Array (FPGA) or Application Specific Integrated circuit (ASIC) implementing a Deep Learning Accelerator, the Deep Learning Accelerator comprising at least one processing unit configured to perform matrix operations.
16. The computing device of claim 15, wherein the at least one processing unit includes a matrix-matrix unit configured to operate on two matrix operands of an instruction;
wherein the matrix-matrix unit includes a plurality of matrix-vector units
configured to operate in parallel;
wherein each of the plurality of matrix-vector units includes a plurality of
vector-vector units configured to operate in parallel; and
wherein each of the plurality of vector-vector units includes a plurality of

multiply-accumulate units configured to operate in parallel.

17. A non-transitory computer storage medium storing instructions which when executed by a computing device cause the computing device to perform a method, the method comprising:
transmitting, by the computing device, one or more requests to an integrated circuit device to provide one or more responses indicative of hardware characteristics of the integrated circuit device in performing matrix computations;
receiving, in the computing device, data representative of a description of an artificial neural network; and
generating, by the computing device, a result of compilation from the data representative of the description of the artificial neural network based at least in part on the hardware characteristics of the integrated circuit device, the result of compilation including first data representative of parameters of the artificial neural network and second data representative of instructions executable by the integrated circuit device to generate an output of the artificial neural network based on the first data and an input to the artificial neural network.
18. The non-transitory computer storage medium of claim 17, wherein the hardware characteristics of the integrated circuit device identify features, options, behaviors, performances, or latency of at least one processing unit of the integrated circuit device, or any combination thereof.
19. The non-transitory computer storage medium of claim 18, wherein the method further comprises:
generating test programs according to predetermined specifications of a plurality of hardware platforms of devices configured to perform matrix computations;
loading the test programs into the integrated circuit device;
instructing the integrated circuit device to execute the test programs;
measuring performance levels of the integrated circuit device in executing instructions contained used in the test programs; and
determining the hardware characteristics of the integrated circuit device based

at least in part on the performance levels.

20. The non-transitory computer storage medium of claim 19, wherein the method further comprises:
- probing hardware options of a processing unit of the integrated circuit device based on results of executing the test programs in the integrated circuit device.




FIG. 1

2 / 16




FIG. 2

3 / 16




FIG. 3

4 / 16




FIG. 4

5 / 16




FIG. 5

6 / 16




FIG. 6

7 / 16




FIG. 7

8 / 16




FIG. 8

9 / 16




FIG. 9

10 / 16




FIG. 10

11 / 16




FIG. 11

12 / 16




FIG. 12

13 / 16




FIG. 13

14 / 16




FIG. 14

15 / 16




FIG. 15

16 / 16




FIG. 16

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2021/055611

A. CLASSIFICATION OF SUBJECT MATTER**G06N 3/063**(2006.01)i; **G06N 3/08**(2006.01)i; **G06N 3/04**(2006.01)i

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 3/063(2006.01); G06F 17/30(2006.01); G06F 9/48(2006.01); G06F 9/50(2006.01); G06N 3/08(2006.01)

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: hardware characteristics, description, artificial neural network (ANN), compilation, parameter

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2020-0184342 A1 (AMAZON TECHNOLOGIES, INC.) 11 June 2020 claims 1-4; and figures 1, 7	1-20
A	WO 2019-216404 A1 (PANASONIC INTELLECTUAL PROPERTY MANAGEMENT CO., LTD.) 14 November 2019 (2019-11-14) claims 1-7; and figures 2, 5	1-20
A	US 2020-0193278 A1 (SAMSUNG ELECTRONICS CO., LTD.) 18 June 2020 (2020-06-18) claims 1-9; and figures 7, 11	1-20
A	US 9501493 B1 (INTERNATIONAL BUSINESS MACHINES CORPORATION) 22 November 2016 (2016-11-22) claims 1-3; and figures 1, 5	1-20
A	US 2020-0293868 A1 (ROVIERO, INC.) 17 September 2020 (2020-09-17) claims 1-8; and figure 1	1-20



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

“D” document cited by the applicant in the international application

“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

09 February 2022

Date of mailing of the international search report

10 February 2022

Name and mailing address of the ISA/KR

Korean Intellectual Property Office
189 Cheongsu-ro, Seo-gu, Daejeon
35208, Republic of Korea

Authorized officer

YANG, JEONG ROK

Facsimile No. +82-42-481-8578

Telephone No. +82-42-481-5709

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.

PCT/US2021/055611

Patent document cited in search report			Publication date (day/month/year)	Patent family member(s)			Publication date (day/month/year)
US	2020-0184342	A1	11 June 2020	None			
WO	2019-216404	A1	14 November 2019	CN	112088383	A	15 December 2020
				JP	6799197	B2	09 December 2020
				US	2021-0056420	A1	25 February 2021
				WO	2019-216404	A1	14 November 2019
US	2020-0193278	A1	18 June 2020	CN	111325312	A	23 June 2020
				DE	102019120930	A1	18 June 2020
				KR	10-2020-0075185	A	26 June 2020
				SG	10201908709	A	29 July 2020
US	9501493	B1	22 November 2016	None			
US	2020-0293868	A1	17 September 2020	None			