



(51) International Patent Classification:

G06N 99/00 (2010.01) G06F 5/01 (2006.01)

(21) International Application Number:

PCT/US2017/031477

(22) International Filing Date:

06 May 2017 (06.05.2017)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/333,214 07 May 2016 (07.05.2016) US  
15/588,558 05 May 2017 (05.05.2017) US

(71) Applicant: 1026 LABS, INC. [US/US]; 1711 E. Olive Way, #407, Seattle, Washington 98102 (US).

(72) Inventors: BRUESTLE, Jeremy; c/o 1026 Labs, Inc., 1711 E. Olive Way, #407, Seattle, Washington 98102 (US).  
NG, Choong; c/o 1026 Labs, Inc., 1711 E. Olive Way, #407, Seattle, Washington 98102 (US).

(74) Agent: CHEN, Elliott Y.; HAN SANTOS, PLLC, 1411 4th Avenue, Suite 910, Seattle, Washington 98101 (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, 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,

(54) Title: AN APPARATUS FOR HARDWARE ACCELERATED MACHINE LEARNING




Figure 1

(57) Abstract: An architecture and associated techniques of an apparatus for hardware accelerated machine learning are disclosed. The architecture features multiple memory banks storing tensor data. The tensor data may be concurrently fetched by a number of execution units working in parallel. Each operational unit supports an instruction set specific to certain primitive operations for machine learning. An instruction decoder is employed to decode a machine learning instruction and reveal one or more of the primitive operations to be performed by the execution units, as well as the memory addresses of the operands of the primitive operations as stored in the memory banks. The primitive operations, upon performed or executed by the execution units, may generate some output that can be saved into the memory banks. The fetching of the operands and the saving of the output may involve permutation and duplication of the data elements involved.

SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR,  
TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

**Declarations under Rule 4.17:**

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

**Published:**

- *with international search report (Art. 21(3))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

# **AN APPARATUS FOR HARDWARE ACCELERATED MACHINE LEARNING**

## **CROSS REFERENCE TO RELATED APPLICATION**

**[0001]** This patent application claims priority to U.S. Provisional Patent Application Serial No. 62/333,214, entitled “Memory and Processing Architecture for Hardware Accelerated Machine Learning,” filed May 7, 2016, which is hereby incorporated by reference herein in its entirety.

## **BACKGROUND**

**[0002]** Machine learning is an emerging field of computer science that has recently attracted much attention and interests. Simply put, machine learning is an automated way of analyzing data and drawing conclusions or making predictions based on the data. Unlike solving an analytical problem, where logical relationship between input (i.e., the data) and output (i.e., the conclusions or predictions) are known and formulated into an algorithm which produces the output based on the input, machine learning takes an approach that is more like how human brain may process information. Specifically, machine learning algorithms “learns” the relationship, which may or may not be logical, between the input and the output by processing a certain amount of sample data. In other words, a programmer does not need to provide with a machine learning algorithm any logical, analytical and/or cognitive knowledge between the input and the output. The machine learning algorithm, which is a quite generic one at the onset of the training or learning process, will simply figure out a relationship between the input and the output by studying the sample data.

**[0003]** For example, a machine learning algorithm may be fed with 1,000 sample pictures each having a crocodile in it, as well as another 1,000 sample pictures each not having a crocodile in it. In addition, the machine learning algorithm is told which pictures have or have not a crocodile therein. No other information needs to be programmed into the algorithm regarding any specific details of a crocodile, such as its typical color, size, general shape, usual habitat, or the like. The pictures having a crocodile may have all kinds of crocodiles, such as a real crocodile, a cartoon crocodile, a wild crocodile, a farmed crocodile, a crocodile at rest by a river bank, a crocodile swimming in the water, an eating crocodile, a crocodile showing its belly, etc. Through processing the 2,000 sample pictures, possibly repeatedly for many times, the machine learning algorithm

modifies itself and gradually tunes the originally generic algorithm into a more complex one that is able to process any picture and make a prediction whether the picture contains a crocodile of any kind, with a high rate of success.

[0004] An unsurprising aspect of machine learning, just as equally true in human learning, is that a machine learning algorithm does not “learn very well” if it is not provided with a sufficiently large amount of sample data, and/or if it processes the sample data for only a few times. As demonstrated in the example above, a large quantity of sample or training data is generally needed, with sufficient times of repetition, for a machine learning algorithm to make the “learning” successful. Depending on the complexity of the problem to solve and the successful rate to achieve, a machine learning algorithm may take hours, days or even months to modify and fine-tune itself to become “well learned”. Therefore, a solution for accelerating the learning process is required.

#### **BRIEF DESCRIPTION OF THE DRAWINGS**

[0005] The Detailed Description is set forth with reference to the accompanying figures.

[0006] Figure 1 is a top-level context diagram for Hardware Accelerated Machine Learning according to the present disclosure.

[0007] Figure 2 is a context diagram for a Machine Learning Acceleration Hardware according to the present disclosure.

[0008] Figure 3 is a block diagram of an example architecture of an apparatus for Hardware Accelerated Machine Learning according to the present disclosure.

[0009] Figure 4 is a block diagram of an example multicast network as coupled to memory banks and execution units according to the present disclosure.

[0010] Figure 5 is a flow chart for an exemplary operation of an indexing module according to the present disclosure.

[0011] Figure 6 is a block diagram of an example forward modulo permutation network according to the present disclosure.

[0012] Figure 7 is a block diagram of an example reverse modulo permutation network according to the present disclosure.

[0013] Figure 8 is a block diagram of an example dynamic random access memory (DRAM) bank read unit according to the present disclosure.

[0014] Figure 9 is a block diagram of an example dynamic random access memory (DRAM) tile read unit according to the present disclosure.

[0015] Figure 10 is a block diagram of an example DRAM bank write unit according to the present disclosure.

### DETAILED DESCRIPTION

[0016] As stated above, machine learning is an automated way for a generic algorithm to learn, by analyzing training data, to draw certain conclusions or make certain predictions based on the data. To achieve a satisfactory result of machine learning, it typically requires the algorithm to process a large quantity of training data, which may take a prolong period of time and/or a lot of computation resources. Therefore, a solution for accelerating the learning process is required.

[0017] Figure 1 is a top-level context diagram for Hardware Accelerated Machine Learning according to the present disclosure. A machine learning (ML) acceleration hardware is usually employed by, or embedded in, a computer device, such as the computer with ML acceleration hardware 200 shown in Figure 1. Computer 200 takes in training data generated by cross-compiler 160 and TILE generator 180 of Figure 1. The reason that cross-compiler 160 is needed in this context is because the training data may originally be encoded using different ML languages, as explained below.

[0018] For some applications, it may be feasible to collect or otherwise prepare the training data in advance, such as the pictures with or without a crocodile in the example above. However, for some applications, the training data may be collected in real-life situations by various equipment or technical platforms each having a specific machine learning language. For example, a ML algorithm may aim to diagnose whether a patient may have cirrhosis based on some abdominal ultrasound images, and the algorithm would require many abdominal ultrasound images as training data. The ultrasound images may be collected from many patients located at various parts of the world by different medical equipment platforms, each employing a different ML language. The different ML languages readily post a problem to the ML algorithm, as the ML algorithm is required to take in training data employing various ML languages.

[0019] For this reason, the context diagram of Figure 1 includes multiple language front ends 140(1) – 140(N) for various ML languages 120(1) – 120(N). Each of language front ends 140(1) – 140(N) feeds a respective one of ML languages 120(1) – 120(N) to cross-compiler 160, which compiles the training data encoded in different ML languages 120(1) – 120(N) such that the training data is encoded with an intermediate language (IL)

called TILE. TILE generator 180 then formulate the output of cross-compiler 160 into the training data that can be taken in by computer 200 which has ML acceleration hardware.

**[0020]** The training of a ML algorithm, as well as the operation of the ML algorithm after it is trained, usually involve heavily so-called “tensor operations”, or computational operations of multidimensional tensors. In its simplest definition, a multidimensional tensor is a multidimensional array of real numbers. Most of the tensor operations involved in a ML algorithm fall in a category called “tensor contraction”, which takes two tensors as input and applies operations such as multiplication and accumulation to the two tensors, resulting in an output tensor.

**[0021]** The tensor contraction operations for machine learning may be performed by general purpose hardware such as a central processing unit (“CPU”). For achieving a better performance, a graphical processing unit (“GPU”) or array of GPUs may be employed which is designed to process vectors of image workloads more efficiently than a CPU. However, certain characteristics of tensor contraction operations for machine learning, such as consistent data flow, large number of multiplications, and specialized nature of the non-multiplication operations, allow for more significant acceleration via hardware by designing the hardware to be specifically tailored for performing such tensor contraction operations for machine learning.

**[0022]** Specifically, the acceleration hardware for machine learning may be implemented in a form of a Hardware Accelerated Machine Learning apparatus termed a “Tensor Processing Unit (TPU)” herein, that is able to efficiently execute a restricted set of primitive instructions for performing the tensor contraction operations for machine learning. In particular, a TPU is a semiconductor device that contains hardware for performing operations optimally for tensor operations and other operations specific to machine learning. For example, where a GPU may have hardware specific to quaternion operations, such operations, while related to linear algebra are not typically used in ML tensor operations. In contrast permutations and noise generation are generally not implemented in GPUs, but are commonly used in ML. In addition to hardware support for ML specific operations, the hardware is implemented in a way to take advantage of pipelining and other hardware/chip design techniques to lower the number of clock cycles used to perform those ML specific operations.

**[0023]** Figure 2 is a context diagram for a Machine Learning Acceleration Hardware according to the present disclosure, wherein a TPU 214 is embedded in the computer 200. In addition to performing tensor contractions, TPU 214 may aid in accelerating other

tensor-related auxiliary operations for machine learning, such as element-wise nonlinear operations and addition of random noise. In some embodiments, computer 200 may include more than one TPU 300 and thus support multi-thread processing for highly parallel applications such as machine learning.

**[0024]** Computer 200 may be any computing device including a standalone computer, a networked server, or an embedded system specific to machine learning. Computer 200 generally has a CPU 210 that is any processor that executes computer readable instructions ranging from general purpose CPUs to embedded controllers.

**[0025]** CPU 210 is communicatively coupled via a bus 212 to TPU 214 and a general-purpose memory 220. The general-purpose memory 220 may be any computer readable media that stores an operating system and other applications each comprised of computer readable instructions. General-purpose memory 220 may include storage media 222 and communications media 224. Computer storage media 222 may include volatile and non-volatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program modules, or other data. Typically, storage media 222 may be comprised of dynamic random access memory (DRAM), but this need not always be the case. In fact, storage media 222 may also be, but is not limited to, static random access memory (SRAM), ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other non-transmission medium that can be used to store information for access by a computing device. In contrast, communication media 224 may embody computer readable instructions, data structures, program modules, or other data in a modulated data signal, such as a carrier wave, or other transmission mechanism. As defined herein, computer storage media does not include communication media.

**[0026]** Computer 200 of Figure 2 may further include a communication peripheral 240 that is also communicatively coupled with TPU 214, CPU 210 and general-purpose memory 220 via bus 212. Communication peripheral 240 may be comprised of input/output (I/O) interface 242 and Network interface 244. I/O interface 242 may be any controller card, such as a universal asynchronous receiver/transmitter (UART) used in conjunction with a standard I/O interface protocol such as RS-232 and/or Universal Serial Bus (USB). In the case of highly parallel applications such as machine learning, the I/O interface may facilities one or more I/O channels and/or parallel I/O channels. For

example, operational control with TPU 214 may be effected over a Peripheral Component Interconnect (PCI) or Peripheral Component Interconnect – Express (PCI-E) bus. Memory controllers may interface via a parallel memory bus e.g. a DRAM bus.

**[0027]** On the other hand, network interface 244 may potentially work in concert with I/O interface 242 and may be a network interface card supporting Ethernet and/or Wi-Fi and/or any number of other physical and/or datalink protocols. Again, in the case of highly parallel applications such as machine learning, multiple network interfaces 244 may be supported by the computer 200.

**[0028]** Figure 3 is a block diagram of an example architecture of an apparatus for Hardware Accelerated Machine Learning, i.e., TPU 300, according to the present disclosure. An overall introduction of TPU 300 is given immediately below, while each component thereof will be separately described in detail at later parts of the present disclosure.

**[0029]** Firstly, TPU 300 includes specific-purpose memory 310 that may be comprised of multiple banks of memory (such as SRAM) that are local to TPU 300 (as opposed to general-purpose memory 220 of Figure 2 which is non-local to TPU 300). The SRAM banks of local memory 310 support concurrent random access. In particular, they support very wide memory access (i.e., multiple byte access within a single computer instruction) for accessing vectors and larger sets of values in a single computer instruction, which is a typical case for machine learning acceleration hardware.

**[0030]** Local SRAM banks 310 holds tensor data, or operands, that can be accessed by a plurality of execution units 320, a second element of TPU 300, which collectively perform tensor-related operations (or, “tensor operations”, in short) for machine learning via executing a series of instructions in an instruction set restricted to TPU 300. Namely, the instruction set contains a set of primitive instructions via which various tensor operations are performed by execution units 320 on the operands stored in local memory 310. Execution units 320 may store some output of the tensor operations to local memory 310 or even non-local memory 220 (of Figure 2), as instructed by some instructions of the instruction set. In general, execution units 320 are identical to each other and work concurrently, although each execution unit 320 may receive different operands from SRAM banks 310, and collectively perform an instruction.

**[0031]** As shown in Figure 3, TPU 300 includes a third element, one or more multicast networks 330, via which execution units 320 access data stored in one or more of SRAM banks 310 so as to perform tensor operations. For some operations, execution units 320

may generate some output to be written back to one or more of SRAM banks 310, also via multicast network 330. Multicast network 330 generally is composed of a series of layers each comprised of a plurality of switch nodes, and data stored in local memory 310 will propagate through the series of layers to reach execution units 320 to be processed. The data fetched from, or written to, SRAM banks 310 typically includes a vector of data elements, usually one element from/to each of SRAM banks 310. Some of the data elements may duplicate or otherwise multiply while propagating through the series of layers of the multicast networks 330, and the order of the data elements may be rearranged or permuted during the propagation. The multiplication and permutation of the data element may be specific to each instruction operation. Namely, the multicast networks 330 may have a generic design that is adaptive based on the specific instruction(s) being performed by execution units 320. In one embodiment, the multicast networks 330 are Beneš networks.

**[0032]** TPU 300 also includes a fourth element, instruction decoder 390. Each ML instruction contains various information for the operation of TPU 300, such as primitive operations for execution units 320 to operate, as well as memory addresses of the operand data for the primitive operations. Instruction decoder 390 decodes a ML instruction to reveal these information, and accordingly uses these information to orchestrate other elements of TPU 300 and/or computer 200 to work in concert and perform the ML operation as instructed. An instruction dictates what tensor operation (or, in some cases, a non-operation) execution units 320 are to perform for each clock cycle. In addition, an instruction may contain a number of address references which indicate the location of data to fetch from local memory 310 or non-local memory 220 as operand(s), and/or the location of local memory 310 or non-local memory 220 an operation output is supposed to store at. Also, an instruction may specify a next instruction to be performed by execution 320.

**[0033]** As shown in Figure 3, TPU 300 may include a fifth element, indexing module 340. The major function of indexing module 340 is to support activities of TPU 300 for looping over tensor indexes and computing memory addresses or other variables. Indexing module 340 may include one or more counter units and one or more addressing units, which will be described later in the present disclosure.

**[0034]** In some embodiments, TPU 300 may also include a sixth element, memory transfer engine 380, which functions to move tensor data from one memory space to another. For example, memory transfer engine 380 may facilitate moving tensor data

within local memory 310 of Figure 3, within non-local memory 220 of Figure 2, or between local memory 310 and non-local memory 220. Memory transfer engine 380 may include components such as index module 384, DRAM controller 386 and SRAM controller 387.

**[0035]** The six elements of TPU 300, as described above and shown in Figure 3, may include one or more of the following hardware components, modules or submodules:

- Execution Units 320
- Multicast Network 330
- Indexing Module 340. An index module may also be included in each of DRAM controller 386 and SRAM controller 387.
- Bank Permutation Network, which is included in memory transfer engine 380. Specifically, each of DRAM controller 386 and SRAM controller 387 may have one or more bank permutation networks.
- DRAM controller 386
- SRAM controller 387
- Memory Transfer Engine 380
- Instruction Decoder 390

Note that the DRAM controller 386 and the SRAM controller 387 are specifies of a memory controller that accesses computer readable memory. In general, we will refer to the DRAM controller 386 and the SRAM controller 387 generically as a banked asynchronous memory controller, or more specifically a banked asynchronous memory read unit and/or a banked asynchronous memory write unit. These hardware components are individually described in greater detail below.

### **Execution Units**

**[0036]** Collectively as a whole, the execution units (EUs) form the primary computational element of the TPU. The computational element performs the multiply-accumulate functionality used during tensor contractions. It also provides a method to perform piecewise linear element-wise operations and other auxiliary functions.

**[0037]** Each of the execution units is able to perform instructions from a limited instruction set, and for each clock cycle, initiates one operation. The EUs may be controlled by a single instruction decoder such that for each clock cycle all EUs of the TPU performs a same operation.

**[0038]** Each execution unit contains three logical registers, A, S, and R. The A register is a traditional accumulator. The S register is a status register which allows conditional execution of instructions. The R register holds a PRNG (Pseudo Random Number Generator) state, which will be disclosed below.

**[0039]** An execution unit may receive two primary operands from a memory system that includes a “B-side memory” and a “C-side memory”. The two operands may be called “B” and “C”, each retrieved from one side of the memory system, respectively. The two operands may be used in different ways (or even ignored, in some cases) by the various instructions of the instruction set. In addition, an instruction thereof may include bits used to determine whether to execute the instruction based on the status register and/or how to update the status register.

### *Design Parameters*

**[0040]** A set of design parameters, such as the parameters shown below along with their respective default values, determine the overall size and performance of the design of each EU. The default values are exemplary and may be subject to change depending on design requirements.

| Name        | Default Value | Description/Formula                |
|-------------|---------------|------------------------------------|
| ELEM_WIDTH  | 16            | Number of bits per operand         |
| ACCUM_WIDTH | 48            | Number of bits per accumulator     |
| EOP_BITS    | 3             | Number of bits for EU opcode       |
| COND_REGS   | 2             | Number of conditional registers    |
| COND_POW    | 4             | $\text{pow}(2, \text{COND\_REGS})$ |

### *Pinouts*

**[0041]** Each EU may have the following set of external pins via which the EU is communicatively coupled to and interact with other one or more EU and/or other components of the TPU.

| Name    | Width      | Direction | Description             |
|---------|------------|-----------|-------------------------|
| CLOCK   | 1          | IN        | Global Clock            |
| RESET   | 1          | IN        | Reset                   |
| IN_OP   | EOP_BITS   | IN        | Operation specification |
| IN_B    | ELEM_WIDTH | IN        | Input B                 |
| IN_C    | ELEM_WIDTH | IN        | Input C                 |
| IN_COND | COND_POW   | IN        | Conditional execution   |

|         |            |     |                               |
|---------|------------|-----|-------------------------------|
| IN_CSET | COND_REGS  | IN  | Conditional set mask          |
| IN_CRST | COND_REGS  | IN  | Conditional reset mask        |
| OUT_AL  | ELEM_WIDTH | OUT | Accumulator lower bits output |

[0042] In the above table, IN\_OP is the instruction opcode, as defined in the instruction set table below. IN\_B and IN\_C are two signed integer operands which come from the B side and C side of the memory system, respectively, each ELEM\_WIDTH wide. IN\_COND, IN\_CSET, and IN\_CRST are related to conditional execution of an instruction by the EU, as described below. OUT\_AL on a given clock represents the saturated low ELEM\_WIDTH bits of the accumulator. This output may have a fixed pipeline delay relative to the inputs.

#### *Instruction Set*

[0043] The table below shows an example list of opcodes each corresponding to an instruction of the instruction set executable by the EUs. The implementation is only exemplary and one may reorder or add additional instructions as needed.

| IN_OP | Mnemonic | Description            | Effect                         |
|-------|----------|------------------------|--------------------------------|
| 0     | MUL      | Multiply               | $A := B * C$                   |
| 1     | MAC      | Multiply Accumulate    | $A := A + B * C$               |
| 2     | ASR      | Arithmetic Shift Right | $A := A \gg B$                 |
| 3     | LTR      | Linear Transform       | $A := (AL - B) * C$            |
| 4     | CMP      | Compare                | $B \leq A \leq C$              |
| 5     | ARN      | Add Random Noise       | $A := A + R(B)$                |
| 6     | SRN      | Seed Random            | $R := (B \ll ELEM\_WIDTH)   C$ |

[0044] Details of the instructions are provided as follows.

[0045] **MUL** The multiply instruction takes B and C as ELEM\_WIDTH signed numbers, multiplies them at full width (i.e., generate a 2\*ELEM\_WIDTH wide output) and overwrites A register with the result. It is worth noting that, by multiplying 1 on either side, multiply can be used to perform a “SET” function.

[0046] **MAC** The multiply instruction takes B and C as ELEM\_WIDTH signed numbers, multiplies them at full width (i.e., generate a 2\*ELEM\_WIDTH wide output) and adds the result to the A register.

[0047] **ASR** The instruction performs an arithmetic right shift of accumulator by the number specified in B, with proper sign extension. If the value of B is negative, no shift is

performed. If the value of B is larger than ACCUM\_WIDTH, the shift is equivalent to a shift of ACCUM\_WIDTH.

**[0048] LTR** The instruction performs a linear transform on A (or more accurately, a register AL that is logically related to register A; see below). Specifically, it performs the function:  $A := (AL - B) * C$ , where AL, B, and C are taken as signed ELEM\_WIDTH bit numbers, and the multiplication is performed at full width.

**[0049] CMP** The instruction evaluates the expression  $(B \leq AL \text{ and } AL \leq C)$ . If the instruction is executed and the expression is true (i.e., the IN\_COND test is true), then the status updates are applied. Otherwise, the status updates are not applied.

**[0050] ARN** The instruction adds a B bit random number to the accumulator. Specifically, as a first step, B is capped to the valid range. That is, if  $B > 32$ , B is treated as 32; if  $B < 1$ , B is treated as 1. Secondly, the lower order B bits are retained, and the upper bit is extended to all higher bits (up to ACCUM\_WIDTH bits in total), resulting in a ACCUM\_WIDTH wide signed number that is in the range of a B bit signed integer. This value is then added to the accumulator. Following that, the PRNG state is updated as described below in the PRNG section.

**[0051] SRN** This instruction sets the state of the PRNG based on B and C. For this operation, B and C are considered as unsigned integers of ELEM\_WIDTH in length, and are then concatenated with one another, with B being the high order bits. If  $2 * \text{ELEM\_WIDTH}$  is greater than 32, the bits higher than 32 are ignored. If  $2 * \text{ELEM\_WIDTH}$  is less than 32, the high bits are 0 padded. This 32 bit number is then assigned to the R register.

### *Registers*

**[0052]** As previously stated, each EU contains a plurality of registers. The most important register is the accumulator (A) register, which is the primary destination of most instructions. This accumulator is a signed integer which is ACCUM\_WIDTH bits wide.

**[0053]** In addition to the physical accumulator register, there is a logical register AL, which is the saturated ELEM\_WIDTH version of the accumulator. It is a signed integer ELEM\_WIDTH in length. If the current value of the accumulator is representable by a signed ELEM\_WIDTH value, then AL will be that value. If the current value of the accumulator is lower than the smallest signed ELEM\_WIDTH value, AL will be the lowest valid ELEM\_WIDTH signed integer. Conversely, if the value of the accumulator is greater than the largest signed ELEM\_WIDTH integer, AL will be set to the largest ELEM\_WIDTH signed integer.

[0054] Furthermore, each execution unit contains a status register, S, which consists of COND\_REGS independent bits. The S register is used for conditional execution as described below.

[0055] Finally, each execution unit has a R register holding a PRNG state, which is described in a later section of the present disclosure regarding PRNG.

#### *Reset Behavior*

[0056] If the reset pin, RESET, is held high during a clock, a reset is initiated and the values of all other pins are ignored. In addition, the reset clears all the bits in all registers to 0. Note that since 0 is a fixed point for the PRNG, this means that all numbers produced by the PRNG will be 0 until the PRNG is properly seeded.

#### *Conditional Execution*

[0057] The conditional execution mechanism is designed to allow conditional logic without data dependent flow control. The mechanism is designed primarily for the use in non-linear transformation, but may also be used by the programmer for any other reason. Based on the state of the COND\_REGS bits in the status register and the input IN\_COND, the instruction may be skipped. Namely, the instruction may be transformed into a no-op.

[0058] Specifically, interpreting the current state of the status register as a COND\_REGS bit number S, if the the S'th bit of IN\_COND is set, instruction is executed; otherwise the instruction becomes a no-op. Thus, if IN\_COND is all 1 bits, the instruction is executed unconditionally; if IN\_COND is all 0 bits, the instruction is always skipped.

[0059] In addition, if an instruction is executed, the status bits are updated based on IN\_CSET and IN\_CRST, such that bits specified in IN\_CSET are set, and bits in IN\_CRST are reset. More specifically,  $S := (S \& \text{IN\_CRST}) \mid \text{IN\_CSET}$ . In the case of the "CMP" instruction, these changes are made only if the instruction is executed (via IN\_COND) and, additionally, if the comparison is true. If the instruction is skipped due to IN\_COND, or the comparison for CMP fails, the status bits are left unmodified.

#### *PRNG*

[0060] The state of the PRNG, as held in the R register, is 32 bits (treated as unsigned internally). When used to generate a random number, the current value of the R register is utilized, and the internal state is updated as:

$$r' = (r \ll 1) \mid (1 \& ((r \gg 31) \wedge (r \gg 29) \wedge (r \gg 25) \wedge (r \gg 24)));$$

*Timing*

**[0061]** In general, the execution unit may pipeline operations. This means that the value of AL in OUT\_AL may follow the instruction stream with some fixed delay. In addition, the instructions that use AL internally (e.g., LTR and CMP) may see an older version of AL due to internal pipelining, possibly necessitating the insertion of NOP, or no-op, instructions (in the case when IN\_COND = 0) after instructions that modify A before the use of AL (e.g., by LTR or CMP). However, changes to registers A, R, and S will always appear to complete in a single clock cycle.

**Multicast Network**

**[0062]** Figure 4 illustrates a block diagram of an example multicast network 400 as coupled to memory banks 312 and 314 and execution units 320 according to the present disclosure. At a high level, multicast network 400 moves data from a fixed SRAM bus of SRAM 312 and 314 to a fixed set of execution units 320, with support for reordering and/or duplicating values. In an embodiment where there are two independent SRAM banks in the memory system (e.g., the B side and the C side as shown in Figure 4), two independent multicast Beneš networks 442 and 444 are employed, one for each of the SRAM banks. The two multicast Beneš networks 442 and 444 may have a slight asymmetry, which makes them appropriate for row/column use cases.

**[0063]** As shown in Figure 4, multicast Beneš networks 442 and 444 are employed to transfer data from SRAM 312 and 314 into the fanout networks (i.e., the “Modulo Fanout” network 462 and the “Grouped Fanout” network 464 as shown in Figure 4). These Beneš networks 442 and 444 are fully general and may be arranged to set or realize any arbitrary pattern. From each of SRAM B 312 and SRAM C 314, data propagates through one of the multicast Beneš networks 442 and 444, and then a fixed fanout network (either the “Modulo Fanout” network 462 or the “Grouped Fanout” network 464) expands or otherwise converts the number of data elements from the SRAM bus width to the full width of the execution unit bus. The degree of fanout is fixed, and is generally a power of two. As shown in Figure 4, on the B side, the output from SRAM B 312 is tiled onto the execution units 320, whereas on the C side, each of the elements from SRAM C 314 is duplicated onto the execution units 320. The fanout networks 462 and 464 are referred to as the modulo and grouped fanout units, respectively.

**[0064]** Most of the complexity of the multicast network lies in the Beneš networks 442 and 444. Details of Beneš networks, Beneš switch units, and the fixed fanout networks are provided below.

*Beneš Network (benes\_net)*

**[0065]** A multicast Beneš network allows data read from sequential memory to be reordered and/or duplicated. It supports fully general multicast patterns, allowing all possible multicast pattern to be realized. It accepts new data every clock cycle, and outputs new data every clock cycle. It is internally pipelined with a fixed delay, and supports reconfiguration in 2 clock cycles. Reconfiguration data arrives over the normal input lines, and is signaled via control lines.

**[0066]** The multicast Beneš network is internally composed of a series of permutation layers (such as layers 4427 of Beneš network 442 shown in Figure 4) and a series of switching layers (such as layers 4423 of Beneš network 442 shown in Figure 4), with two kinds of layers alternately disposed. Each of the permutation layers is configured to realize a fixed permutation of the bus elements. The switching layers contain BENES\_WIDTH/2 2x2 switching nodes, or switching units, such as numeral 4421 if Figure 4, and each of the switching nodes holds two updatable bits containing configuration data that determines the switching function. The switching nodes are described in further detail in the Beneš switch module disclosed below.

**[0067]** A set of design parameters, such as the parameters shown below along with their respective default values, determine the overall size and performance of the design of each Beneš network. The default values are exemplary and may be subject to change depending on design requirements.

| Name        | Default Value | Description/Formula          |
|-------------|---------------|------------------------------|
| ELEM_WIDTH  | 16            | Number of bits per element   |
| BENES_PO2   | 11            | Power of two of switches     |
| BENES_DEPTH | 23            | $2 * BENES\_PO2 + 1$         |
| BENES_UNITS | 2048          | $POW(2, BENES\_PO2)$         |
| BENES_WIDTH | 4096          | $2 * BENES\_UNITS$           |
| BUS_WIDTH   | 65536         | $ELEM\_WIDTH * BENES\_WIDTH$ |

**[0068]** Each multicast Beneš network may have the following set of external pins via which the multicast Beneš network is communicatively coupled to and interact with other components of the multicast network, and other components of the TPU (such as SRAM).

| Name     | Width | Direction | Description                 |
|----------|-------|-----------|-----------------------------|
| CLOCK    | 1     | IN        | Global Clock                |
| PASSTHRU | 1     | IN        | Ignore config, pass through |

|           |           |     |                                    |
|-----------|-----------|-----|------------------------------------|
| BROADCAST | 1         | IN  | Ignore config, broadcast element 0 |
| SET_0     | 1         | IN  | Update 0th entry of each node      |
| SET_1     | 1         | IN  | Update 1th entry of each node      |
| DATA_IN   | BUS_WIDTH | IN  | Input data                         |
| DATA_OUT  | BUS_WIDTH | OUT | Output data                        |

**[0069]** The Beneš network is composed of  $BENES\_DEPTH * BENES\_UNITS$  switching nodes, or switching units. Each switching unit has two inputs and two outputs. These are arranged in  $BENES\_DEPTH$  layers, each layer having  $BENES\_UNITS$  units across. Between layers, switches are connected to one another in a recursive pattern via the permutation layers. Specifically, the switching layers may be labeled or otherwise identified with an index, from  $-BENES\_PO2$  for the layer right after the input, to  $+BENES\_PO2$  for the final layer before the output. Call input  $i \in \{0, 1\}$  of unit  $j \in [0, BENES\_UNIT\_S)$  on layer  $l \in [-BENES\_PO2, BENES\_PO2]$  as  $I(l, j, i)$ , and the associated output as  $O(l, j, i)$ . We have:

$$W(l) = 1 > 0 ? 2^l : 2^{-(l+1)}$$

$$B(l, j) = \text{floor}(j / W(l)) * W(l)$$

$$J(l, j) = j - B(l, j)$$

$$D(l, j) = \text{floor}(J(l, j)/2)$$

$$H(l, j) = \text{floor}(2 * J(l, j) / W(l))$$

if ( $l > 0$ ):

$$I(l, j, i) = O(l-1, B(l, j) + D(l, j) + i * W(l)/2, j \% 2)$$

else:

$$I(l, j, i) = O(l-1, B(l, j) + (2 * j + i) \% W(l), H(l, j))$$

**[0070]** The global control data (PASSTHRU, BROADCAST, SET\_0, SET\_1) is passed to each switching node, with proper pipelining delays, so that all of the control for a single flow is passed along at the same rate as the data. In addition, each switching layer is parameterized slightly differently. Specifically, if we number each switching layer, with 0 being the layer connected directly to input, and  $2 * BENES\_PO2$  being the layer connected directly to output, we have:

$$CFG\_ELEM = (l < ELEM\_WIDTH)$$

$$CFG\_BIT = (l \% ELEM\_WIDTH)$$

**[0071]** The operation of the Beneš network is largely a function of the topology of the Beneš network itself, the structure of the Beneš switch (i.e., Beneš \_switch, as shown below), and the pipelining.

*Beneš Switch Units (benes\_switch)*

**[0072]** A set of design parameters, such as the parameters shown below along with their respective default values, determine the overall size and performance of the design of the Beneš switch units. The default values are exemplary and may be subject to change depending on design requirements.

| Name       | Default Value | Description/Formula              |
|------------|---------------|----------------------------------|
| ELEM_WIDTH | 16            | Number of bits per element       |
| CFG_ELEM   | 0             | Which element to get config from |
| CFG_BIT    | 0             | Which bit to get config from     |

**[0073]** Each multicast Beneš switching unit may have the following set of external pins via which the multicast Beneš switching unit is communicatively coupled to and interact with other components of the multicast network, and/or other components of the TPU.

| Name      | Width      | Direction | Description                        |
|-----------|------------|-----------|------------------------------------|
| CLOCK     | 1          | IN        | Global Clock                       |
| PASSTHRU  | 1          | IN        | Ignore config, pass through        |
| BROADCAST | 1          | IN        | Ignore config, broadcast element 0 |
| SET_0     | 1          | IN        | Update 0th entry of each node      |
| SET_1     | 1          | IN        | Update 1th entry of each node      |
| IN_0      | ELEM_WIDTH | IN        | Input entry 0                      |
| IN_1      | ELEM_WIDTH | IN        | Input entry 1                      |
| OUT_0     | ELEM_WIDTH | OUT       | Output entry 0                     |
| OUT_1     | ELEM_WIDTH | OUT       | Output entry 1                     |

**[0074]** Each multicast Beneš switching unit may have a plurality of registers such as ones shown below.

| Name  | Width | Description                  |
|-------|-------|------------------------------|
| CFG_0 | 1     | Which input goes to output 0 |
| CFG_1 | 1     | Which input goes to output 1 |

**[0075]** Each multicast Beneš switching unit may operate in an always operation mode as described by the virtual code below:

```

if PASSTHRU:
    OUT_0 = IN_0
    OUT_1 = IN_1
else if BROADCAST:
    OUT_0 = IN_0
    OUT_1 = IN_0
else:
    OUT_0 = (CFG_0 ? IN_1 : IN_0)
    OUT_1 = (CFG_1 ? IN_1 : IN_0)
cfg_bit = (CFG_ELEM ? IN_1 : IN_0)[CFG_BIT]

```

**[0076]** Each multicast Beneš switching unit may further operate in a clocked operation mode as described by the virtual code below:

```

if SET_0:
    CFG_0' = cfg_bit
if SET_1:
    CFG_1' = not cfg_bit

```

**[0077]** It is worth noting that, during normal configuration operation, PASSTHRU is held while SET\_0 and then SET\_1 as each set. If the network is being set to a pure permutation, SET\_0 and SET\_1 can both be set, resulting in a single cycle update. Also, passthru behavior or broadcast behavior can be activated without updating the network.

#### *Modulo Fanout (mod\_fanout)*

**[0078]** A set of design parameters, such as the parameters shown below along with their respective default values, determine the overall size and performance of the design of the modulo fanout network 462. The default values are exemplary and may be subject to change depending on design requirements.

| Name       | Default Value | Description/Formula        |
|------------|---------------|----------------------------|
| ELEM_WIDTH | 16            | Number of bits per element |
| IN_ELEMS   | 4096          | Number of input elements   |
| IN_WIDTH   | 65536         | IN_ELEMS * ELEM_WIDTH      |
| OUT_ELEMS  | 131072        | Number of output elements  |
| OUT_WIDTH  | 2097152       | OUT_ELEMS * ELEM_WIDTH     |

[0079] The modulo fanout network may have the following set of external pins via which the modulo fanout network 462 is communicatively coupled to and interact with EUs, a multicast Beneš network, and/or other components of the TPU.

| Name | Width     | Direction | Description |
|------|-----------|-----------|-------------|
| IN   | IN_WIDTH  | IN        | Input data  |
| OUT  | OUT_WIDTH | OUT       | Output data |

[0080] Basically, the mod\_fanout module is pure wiring. That is:

$$\text{OUT}[i] = \text{IN}[i \% \text{IN\_ELEMS}]$$

*Grouped Fanout (group\_fanout)*

[0081] The grouped fanout network 464 is identical to the modulo fanout network 462, except the output elements of the grouped fanout network are defined as:

$$\text{OUT}[i] = \text{IN}[\text{floor}(i / (\text{OUT\_ELEMS} / \text{IN\_ELEMS}))]$$

### **Indexing: Counter Unit and Addressing Unit**

[0082] The process of looping over indexes (usually dimensions of a tensor) and computing a memory address or other variable based on the values of the tensor is a common process of many TPU components. To provide support for this activity, we define here two units: a multidimensional counter unit and a multidimensional addressing unit. Specifically, the counter unit tracks the current value of the set of indexes in use, while the addressing unit keeps up-to-date addresses including a modulo offset. Details of the two units are given below.

*Multidimensional Counter Unit (mdim\_count)*

[0083] The counter unit provides the actual loop iteration mechanism used by other indexing components. A set of design parameters, such as the parameters shown below along with their respective default values, determine the overall size and performance of the design of the counter unit. The default values are exemplary and may be subject to change depending on design requirements.

| Name      | Default Value | Description/Formula                      |
|-----------|---------------|------------------------------------------|
| IDX_COUNT | 16            | Number of indexes in the counter         |
| IDX_WIDTH | 16            | Number of bits per index                 |
| IDX_SEL   | 4             | $\text{ceil}(\log_2(\text{IDX\_COUNT}))$ |
| IDX_BUS   | 256           | $\text{IDX\_COUNT} * \text{IDX\_WIDTH}$  |

**[0084]** The counter unit may have the following set of external pins via which the counter unit is communicatively coupled to and interact with other components of the TPU.

| Name   | Width   | Direction | Description                             |
|--------|---------|-----------|-----------------------------------------|
| CLK    | 1       | IN        | Global clock                            |
| RST    | 1       | IN        | Reset                                   |
| CFG    | 1       | IN        | Configure counter sizes                 |
| STEP   | 1       | IN        | Step enable                             |
| CFG_IN | IDX_BUS | IN        | Configuration input                     |
| OVER   | 1       | OUT       | Would the next step wrap the 0'th index |
| NEXT   | IDX_SEL | OUT       | Which index will decrement next step    |

**[0085]** The operation of the counter unit is described as follows. Internally, the multidimensional counter unit retains `IDX_COUNT` counters, each `IDX_WIDTH` bits long. Call counter  $i$ ,  $C[i]$ . In addition, the multidimensional counter unit contains `IDX_COUNT` sizes, each `IDX_WIDTH` bits long, called  $S[i]$ . At each clock, the following logic occurs (where  $C[i]$  and  $S[i]$  represent the pre-clock value of the counters and the sizes, respectively, and  $C'[i]$  and  $S'[i]$  represent the new value thereof).

if RST is high:

For all  $i$ ,  $C'[i] = 0$  and  $S'[i] \leq 0$

else if CFG is high:

For all  $i$ ,  $C'[i] = \text{CFG\_IN}[i]$ ,  $S'[i] = \text{CFG\_IN}[i]$

else if STEP is high:

For all  $i$ :

if  $i > \text{NEXT}$  or  $\text{OVER} == 1$ ,  $C'[i] = S[i]$

else if  $i == \text{NEXT}$ ,  $C'[i] = C[i] - 1$

if for all  $i$ ,  $C'[i] == 0$ :

$\text{OVER}' = 1$

$\text{NEXT}' = 0$

else:

$\text{OVER}' = 0$

$\text{NEXT}' = \max(i) \text{ such that } C'[i] \neq 0$

**[0086]** For example, given `IDX_COUNT` of 4, the following table illustrates the behavior of the multidimensional counter unit, where each row represents one clock cycle. The first 4 columns represent the input values prior to the clock, the next 8 columns represent the updated value of the registers ( $C'$ ,  $S'$ ) after the clock, and the final 2 columns

represent the outputs after the clock edge. Signal names have been abbreviated to reduce column width (e.g., R is the abbreviation of RST, C is the abbreviation of CFG, S is the abbreviation of STEP, O is the abbreviation of OVER, and N is the abbreviation of NEXT).

| R | C | S | CFG_IN  | S0 | S1 | S2 | S3 | C0 | C1 | C2 | C3 | O | N |
|---|---|---|---------|----|----|----|----|----|----|----|----|---|---|
| 1 | X | X | X       | X  | X  | X  | X  | X  | X  | X  | X  | X | X |
| 0 | 0 | 0 | X       | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 0  | 1 | 0 |
| 0 | 0 | 1 | X       | 0  | 0  | 0  | 2  | 0  | 0  | 0  | 0  | 1 | 0 |
| 0 | 1 | 0 | 3,1,0,2 | 3  | 1  | 0  | 2  | 3  | 1  | 0  | 2  | 0 | 3 |
| 0 | 0 | 1 | X       | 3  | 1  | 0  | 2  | 3  | 1  | 0  | 1  | 0 | 3 |
| 0 | 0 | 1 | X       | 3  | 1  | 0  | 2  | 3  | 1  | 0  | 0  | 0 | 1 |
| 0 | 0 | 1 | X       | 3  | 1  | 0  | 2  | 3  | 0  | 0  | 2  | 0 | 3 |
| 0 | 0 | 1 | X       | 3  | 1  | 0  | 2  | 3  | 0  | 0  | 1  | 0 | 3 |
| 0 | 0 | 1 | X       | 3  | 1  | 0  | 2  | 3  | 0  | 0  | 0  | 0 | 0 |
| 0 | 0 | 1 | X       | 3  | 1  | 0  | 2  | 2  | 1  | 0  | 2  | 0 | 3 |
| 0 | 0 | 1 | X       | 3  | 1  | 0  | 2  | 2  | 1  | 0  | 1  | 0 | 3 |
| 0 | 0 | 1 | X       | 3  | 1  | 0  | 2  | 2  | 1  | 0  | 0  | 0 | 1 |
| 0 | 0 | 1 | X       | 3  | 1  | 0  | 2  | 2  | 0  | 0  | 2  | 0 | 3 |
| . | . | . | .       | .  | .  | .  | .  | .  | .  | .  | .  | . | . |
| . | . | . | .       | .  | .  | .  | .  | .  | .  | .  | .  | . | . |
| . | . | . | .       | .  | .  | .  | .  | .  | .  | .  | .  | . | . |
| 0 | 0 | 1 | X       | 3  | 1  | 0  | 2  | 0  | 1  | 0  | 1  | 0 | 3 |
| 0 | 0 | 1 | X       | 3  | 1  | 0  | 2  | 0  | 1  | 0  | 0  | 0 | 1 |
| 0 | 0 | 1 | X       | 3  | 1  | 0  | 2  | 0  | 0  | 0  | 2  | 0 | 3 |
| 0 | 0 | 1 | X       | 3  | 1  | 0  | 2  | 0  | 0  | 0  | 1  | 0 | 3 |
| 0 | 0 | 1 | X       | 3  | 1  | 0  | 2  | 0  | 0  | 0  | 0  | 1 | 0 |
| 0 | 0 | 1 | X       | 3  | 1  | 0  | 2  | 3  | 1  | 0  | 2  | 0 | 3 |

**[0087]** Figure 5 is a flow chart showing the operation process of the example above. To clarify the flow, the logical order will be described here, but multiple independent steps may be scheduled on the same clock cycle, and some steps may take more than one clock cycle, with internal pipeline latches. As shown in Figure 5, process 500 may start from block 510.

**[0088]** At block 510, process 500 may involve resetting counter registers C[0] – C[3] to ‘0’ and size registers S[0] – S[3] also to ‘0’. Process 500 may proceed from block 510 to block 520.

**[0089]** At block 520, process 500 may involve setting OVER register to ‘1’ and NEXT register to ‘0’. Process 500 may proceed from block 520 to block 530.

**[0090]** At block 530, process 500 may involve configuring counter registers C[0] – C[3] and size registers S[0] – S[3] are configured according to configuration input CFG\_IN, which is ‘3, 1, 0, 2’ in the example above. Process 500 may proceed from block 530 to block 540.

**[0091]** At block 540, process 500 may involve determining whether any of counter registers C[0] – C[3] contains a non-zero value. In response to determining that at least one of counter registers C[0] – C[3] contains a non-zero value, process 500 may proceed from block 540 to block 550. On the other hand, in response to determining that each of counter registers C[0] – C[3] contains a zero value, process 500 may proceed from block 540 to block 510.

**[0092]** At block 550, process 500 may involve setting OVER register to ‘0’. Process 500 may proceed from block 550 to block 560.

**[0093]** At block 560, process 500 may involve updating the NEXT register based on a counter register of a highest sequence that contains a non-zero value. For example, in the 4<sup>th</sup> row of the table above, the NEXT register is updated as ‘3’, as the highest-sequence counter register that contains a non-zero value is register C[3], which contains ‘2’, a non-zero value. As another example, in the 6<sup>th</sup> row of the table above, the NEXT register is updated as ‘1’, as the highest-sequence counter register that contains a non-zero value is register C[1] (noting that each of C[2] and C[3] contains ‘0’), which contains ‘1’, a non-zero value. Process 500 may proceed from block 560 to block 570.

**[0094]** At block 570, process 500 may involve decrementing the non-zero value of the highest-sequenced counter register by 1. For example, in the 5<sup>th</sup> row of the table above, C[3] is decremented from ‘2’ (as shown in the 4<sup>th</sup> row) to ‘1’. Process 500 may proceed from block 570 to block 580.

**[0095]** At block 580, process 500 may involve updating counter registers having a higher sequence than specified by the NEXT register with the corresponding size registers. For example, given that in the 6<sup>th</sup> row of the table above the NEXT register has been updated as ‘1’, in the 7<sup>th</sup> row of the table, both counter registers C[2] and C[3] are updated according to size registers S[2] and S[3], respectively. That is, in the 7<sup>th</sup> row of the table,

C[2] is set to have the same value of S[2] (which is '0'), and C[3] is set to have the same value of S[3] (which is '2'), respectively. Process 500 may loop back from block 580 to block 540.

**[0096]** As process 500 goes on following the flow chart of Figure 5, process 500 may repeat blocks 540, 550, 560, 570 and 580 for many times in a loop, until every counter register contains '0', as shown in the second-to-last row of the table above. As can be seen in the last row of the table above, the tensor indexes are "looped over" in the process, as the last row becomes identical to the 4<sup>th</sup> row of the table. As the process continues, the indexes stored in the counter registers may continue changing, i.e., the multidimensional counter unit may loop over the indexes for many times.

#### *Multidimensional Address Unit (mdim\_addr)*

**[0097]** The multidimensional address unit adjusts an address based on the index which is changing due to operation of the multidimensional counter as described above. It allows an address to be an arbitrary polynomial of it's indexes, and tracks both the quotient and the remainder when divided by a fixed modulus. This operation is performed without the need for multiplication or division circuits, since each step can be done via addition.

**[0098]** A set of design parameters, such as the parameters shown below along with their respective default values, determine the overall size and performance of the design of the address unit. The default values are exemplary and may be subject to change depending on design requirements.

| Name       | Default Value | Description/Formula                      |
|------------|---------------|------------------------------------------|
| IDX_COUNT  | 16            | Number of indexes                        |
| IDX_SEL    | 4             | $\text{ceil}(\log_2(\text{IDX\_COUNT}))$ |
| DIVISOR    | 4192          | 131 (# of banks) * 32 (bank size)        |
| REM_WIDTH  | 13            | $\text{ceil}(\log_2(\text{DIVISOR}))$    |
| QUOT_WIDTH | 19            | Number of bits per address quotient      |
| ADDR_WIDTH | 32            | $\text{QUOT\_WIDTH} + \text{REM\_WIDTH}$ |
| DELTA_BUS  | 512           | $\text{IDX\_COUNT} * \text{ADDR\_WIDTH}$ |

**[0099]** The address unit may have the following set of external pins via which the address unit is communicatively coupled to and interact with other components of the TPU.

| Name | Width | Direction | Description  |
|------|-------|-----------|--------------|
| CLK  | 1     | IN        | Global clock |

|          |            |     |                                                                                                   |
|----------|------------|-----|---------------------------------------------------------------------------------------------------|
| RST      | 1          | IN  | Reset                                                                                             |
| CFG      | 1          | IN  | Configure address deltas                                                                          |
| STEP     | 1          | IN  | Step enable                                                                                       |
| IDX      | IDX_SEL    | IN  | Which index to step                                                                               |
| QUOT_IN  | QUOT_WIDTH | IN  | Quotient of initial address during                                                                |
| REM_IN   | REM_WIDTH  | IN  | config                                                                                            |
| DELTA_IN | DELTA_BUS  | IN  | Remainder of initial address during                                                               |
| QUOT     | QUOT_WIDTH | OUT | config Delta Configuration input                                                                  |
| REM      | REM_WIDTH  | OUT | The address quotient for this index<br>state<br><br>The address remainder for this index<br>state |

**[00100]** The operation of the address unit is described as follows. Internally, the multidimensional address unit retains address components in registers QUOT and REM which are also the outputs. It also holds deltas, which are composed of a quotient delta  $DQ[i]$ , and a remainder delta  $DR[i]$ , for each of  $0 \leq i < IDX\_COUNT$  indexes. Based on the selected index, it updates the address by adding the deltas. At each clock, the following logic occurs (where  $X$  represents the pre-clock value of  $X$ , and  $X'$  represents the new value).

if RST is high:

For all  $i$ ,  $DQ'[i] = 0$ ,  $DR'[i] = 0$

$QUOT' = 0$ ,  $REM' = 0$

else if CFG is high:

For all  $i$ ,  $DQ'[i] = DELTA\_IN[i].Q$ ,  $DR'[i] = DELTA\_IN[i].R$

$QUOT' = QUOT\_IN$ ,  $REM' = REM\_IN$

else if STEP is high:

$Q\_TMP = QUOT + DQ[IDX]$

$R\_TMP = REM + DR[IDX]$

$OVFLOW = (R\_TMP > DIVISOR)$

$QUOT' = OVFLOW ? Q\_TMP + 1 : Q\_TMP$

$REM' = OVFLOW ? R\_TMP - DIVISOR : R\_TMP$

$ADDR' = ADDR + D[IDX]$

**[00101]** Notably, if the desired clock speed and address width preclude single cycle operation, operations may be pipelined via the use of methods such as carry-save. This

will result in the QUOT and REM outputs to update every cycle, but with a fixed pipeline delay on the output.

### **Bank Permutation Network**

[00102] When accessing (that is, writing to memory or retrieving from memory) a vector of N data elements, the addresses of the N data elements often take the form:

$$A[i] = O + M * i$$

[00103] Where O is some base address offset, M is a multiplier, and i is an index  $0 \leq i < N$ . When the memory is divided into D banks, and each element is of full bank width and properly aligned, the banks being read from for each i are:

$$B[i] = (O + M * i) \% D$$

[00104] If the number of banks, D, is relatively prime to M, and  $N \leq D$ , there will be no bank conflicts. If D is prime, so long as M is not a strict multiple of D, D and M will be relatively prime. It is presumed hereafter that D is in fact prime relatively to M.

[00105] Given such an access pattern, a bank permutation network, such as bank permutation network 600 of Figure 6 or bank permutation network 700 of Figure 7, is employed to provide a way to permute the request addresses (for memory write operations) such that each request is routed to the proper memory bank so that each entry is in the order of i. Alternatively, the bank permutation network may be configured to permute the addresses such that the read results (for memory read operations) are routed back, so that each entry is in the order of i.

[00106] It is to be noted that, given  $o = O \% D$  and  $m = M \% D$ , the permutation is completely defined where O is the offset, M is the multiplier, and D is the number of banks

[00107] In general, the addition component (called x) can be induced by a simple shift (or rotate). For prime D, the multiplication permutation can be induced by a fixed permutation, a shift (or rotate), and another fixed permutation. This is because for a prime number, the multiplicative group is isomorphic to the cyclic group D-1. The first permutation maps each entry to a cyclic order, and the second permutation undoes that mapping. Also, it is to be noted that a cyclic shift, or a shift in a cyclic space, is sometimes referred to as “rotation”, which is typically performed by a barrel shifter/rotator hardware, such as barrel rotators 630 and 650 of Figure 6 and barrel rotators 730 and 750 of Figure 7.

[00108] To define these permutations, we must choose a generator over the multiplication prime field in question. A generator for a group is an element that, by

repeated applications, produces all the elements of the group. For example, for the prime field over  $D = 7$ , 3 is a multiplicative generator:

$$3^1 \pmod{7} = 3$$

$$3^2 \pmod{7} = 2$$

$$3^3 \pmod{7} = 6$$

$$3^4 \pmod{7} = 4$$

$$3^5 \pmod{7} = 5$$

$$3^6 \pmod{7} = 1$$

**[00109]** The chosen generator is denoted as  $g$ . Note that  $D$  and  $g$  are fixed at the time of design. A discrete log with respect to  $g$ ,  $\log_g(x)$ , can be defined as the value  $y$  such that  $g^y \pmod{D} = x$ . For example, for  $g = 3$ ,  $D = 7$ , we compute  $\log_g(6) = 3$ .

**[00110]** Since the rotation to handle the multiplicative part of the permutation happens in the cyclic space, it is required to compute the discrete log to determine the amount to rotate, which is complex to perform in hardware. In practical implementations, one may presume that  $M$ , and thus  $m$ , and  $\log_g(m)$  are known in advance. This allows a compiler to perform the appropriate computations and provide a fixed constant for the rotation. Specifically, to determine the necessary rotation, we compute:

$$n = (\log_g(m) + 1) \% (D - 1)$$

**[00111]** Namely, to specify a transform,  $o$  and  $r$  are provided at the time of permutation. It is to be noted that the permutation network may be configured to one of two different variants. The first variant is called the forward modulo permutation network, or `fwd_mod_perm` 600 as shown in Figure 6, which maps each  $i < D$  to a correspondingly appropriate bank position  $b = (O + M * i) \% D$  given the correct  $o$  and  $n$ . This permutation is used to send address and data to the memory banks (i.e., for memory write operations). The second variant is called the reverse modulo permutation network, or `rev_mod_perm` 700 as shown in Figure 7, which simply performs the inverse mapping of `fwd_mod_perm`, and is used to appropriately reorder the read data elements due to memory read operations.

**[00112]** Before `fwd_mod_perm` and `rev_mod_perm` modules can be described in greater detail below, two simple wiring patterns are to be defined. The two wiring patterns are used to perform the mapping a logical group and a cyclic group. Specifically, a first wiring pattern, `map_to_cyclic` (such as numeral 620 of Figure 6 and 720 of Figure 7), is defined to take  $D - 1$  elements and map each entry  $i$  of the  $D - 1$  elements to entry  $\log_g(i)$ , for  $1 \leq i < D$ . A second wiring pattern, `map_from_cyclic` (such as numeral 640 of Figure 6 and 740 of Figure 7), is defined to do the opposite and map entry  $i$  to entry  $g^i \pmod{D}$ .

Since 0 is not a member of the multiplicative group, entry 0 of the elements is left unaltered by both the mapping and the rotations.

**[00113]** An example design for the forward and reverse permutation networks for  $D = 7$ ,  $g = 3$  is shown in Figures 6 and 7, respectively. The permutation networks are individually described in greater detail below.

*Forward Modulo Permutaton Network (fwd\_mod\_perm)*

**[00114]** The forward modulo permutation network 600 of Figure 6 operates to permute elements 610 with addresses of the form  $O + M*i$  to the proper memory banks. A set of design parameters, such as the parameters shown below along with their respective default values, determine the overall size and performance of the forward modulo permutation network 600. The default values are exemplary and may be subject to change depending on design requirements.

| Name      | Default Value | Description/Formula                   |
|-----------|---------------|---------------------------------------|
| DIVISOR   | 7             | Prime group                           |
| GENERATOR | 3             | Generator for prime group             |
| REM_WIDTH | 3             | $\text{ceil}(\log_2(\text{DIVISOR}))$ |
| WIDTH     | 32            | Data width for data to be permuted    |
| BUS_WIDTH | 224           | $D * \text{WIDTH}$                    |

**[00115]** The forward modulo permutation network 600 may have the following set of external pins via which the forward modulo permutation network 600 is communicatively coupled to and interact with other components of the TPU.

| Name | Width     | Direction | Description              |
|------|-----------|-----------|--------------------------|
| CLK  | 1         | IN        | Global clock             |
| RST  | 1         | IN        | Global Reset             |
| O    | REM_WIDTH | IN        | Offset modulo D          |
| N    | REM_WIDTH | IN        | Precomputed cyclic shift |
| IN   | BUS_WIDTH | IN        | Input lines              |
| OUT  | BUS_WIDTH | OUT       | Output lines             |

**[00116]** The input, IN, of the forward modulo permutation network is treated as DIVISOR entries,  $\text{IN}[i]$ , each of WIDTH elements wide. The output, OUT, is treated as DIVISOR entries,  $\text{OUT}[i]$ . The behavior of the network is as follows, where  $D = \text{DIVISOR}$ ,  $G = \text{GENERATOR}$ :

$$\text{TO\_CYC}[i] = \text{IN}[(G^{(i+1)}) \% D], \text{ for } 0 \leq i < D-1$$

$$\text{ROT\_CYC}[i] = \text{TO\_CYC}[(i + N) \% (D-1)], \text{ for } 0 \leq i < D-1$$

$$\text{TO\_MUL}[i] = \text{TO\_CYC}[\log G(i+1)], \text{ for } 0 \leq i < D-1$$

$$\text{OFF}[0] = \text{IN}[0]$$

$$\text{OFF}[i] = \text{TO\_MUL}[i-1], \text{ for } 1 \leq i < D$$

$$\text{OUT}[i] = \text{OFF}[(i + O) \% D], \text{ for } 0 \leq i < D$$

**[00117]** The above description of the forward modulo permutation network presumes that the logic is purely combinatorial and takes 0 clock cycles to complete. In practice, it is likely that the permutation network may require pipelining. Namely, the output follows the inputs with some fixed clock cycle delay. In this case, the reset line results in the output for the otherwise undefined clock cycles being all 0s.

*Reverse Modulo Permutaton Network (rev\_mod\_perm)*

**[00118]** The reverse modulo permutation network 700 operates to permute elements with addresses of the form  $O + M \cdot i$  from the proper memory banks back to the original logical order. A set of design parameters, such as the parameters shown below along with their respective default values, determine the overall size and performance of the reverse modulo permutation network 700. The default values are exemplary and may be subject to change depending on design requirements.

| Name      | Default Value | Description/Formula                            |
|-----------|---------------|------------------------------------------------|
| DIVISOR   | 7             | Prime group                                    |
| GENERATOR | 3             | Generator for prime group                      |
| REM_WIDTH | 3             | $\text{ceil}(\log_2(\text{DIVISOR}))$          |
| D_BITS    | 3             | Number of bits to represent size of banks<br>D |
| WIDTH     | 32            | Data width for data to be permuted             |
| BUS_WIDTH | 224           | $D * \text{WIDTH}$                             |

**[00119]** The reverse modulo permutation network 700 may have the following set of external pins via which the reverse modulo permutation network 700 is communicatively coupled to and interact with other components of the TPU.

| Name | Width  | Direction | Description              |
|------|--------|-----------|--------------------------|
| CLK  | 1      | IN        | Global clock             |
| RST  | 1      | IN        | Global Reset             |
| O    | D_BITS | IN        | Offset modulo D          |
| N    | D_BITS | IN        | Precomputed cyclic shift |

|     |           |     |              |
|-----|-----------|-----|--------------|
| IN  | BUS_WIDTH | IN  | Input lines  |
| OUT | BUS_WIDTH | OUT | Output lines |

**[00120]** The input, IN, of the reverse modulo permutation network is treated as DIVISOR entries, IN[i], each of WIDTH elements wide. The output, OUT, is treated as DIVISOR entries, OUT[i]. The behavior of the network is as follows, where  $D = \text{DIVISOR}$ ,  $G = \text{GENERATOR}$ :

$$\text{OFF}[i] = \text{IN}[(i - O) \% D], \text{ for } 0 \leq i < D$$

$$\text{TO\_CYC}[i] = \text{OFF}[(G^{(i+1)}) \% D], \text{ for } 0 \leq i < D-1$$

$$\text{ROT\_CYC}[i] = \text{TO\_CYC}[(i - N) \% (D-1)], \text{ for } 0 \leq i < D-1$$

$$\text{TO\_MUL}[i] = \text{TO\_CYC}[\log G(i+1)], \text{ for } 0 \leq i < D-1$$

$$\text{OUT}[0] = \text{OFF}[0]$$

$$\text{OUT}[i] = \text{TO\_MUL}[i-1], \text{ for } 1 \leq i < D$$

**[00121]** The above description of the reverse modulo permutation network presumes that the logic is purely combinatorial and takes 0 clock cycles to complete. In practice, it is likely that the permutation network may require pipelining. Namely, the output follows the inputs with some fixed clock cycle delay. In this case, the reset line results in the output for the otherwise undefined clock cycles being all 0s.

**[00122]** The forward and reverse modulo permutation networks 600 and 700 described above may work together with a modulo address generator to complete the task of address permutation. The modulo address generator is described below.

#### *Modulo Address Generator (mod\_addr)*

**[00123]** The goal of the modulo address generator (mod\_addr) is to produce and properly “bank” a vector addresses, given an offset and a multiplier, as well as vector size. In some embodiments, the memory of the TPU may have a number of memory banks,  $D$ , where  $D$  is prime, with a logical layout of modulo  $D$ . When accessing the memory, it is desired to read a vector with a fixed stride from logical memory. Specifically, for an offset  $O$ , a multiplier  $M$ , and a length  $L \leq D$ , a set of addresses may be defined as:

$$A[i] = O + M * i, 0 \leq i < L$$

**[00124]** In addition, each address must be send to an appropriate corresponding bank, where the corresponding bank for  $A[i]$  is  $A[i] \% D$ , and the internal address within that bank is  $A[i]/D$ . The goal of the mod\_addr unit is to produce, for each bank, the bank internal address and an enable flag for the respective bank.

**[00125]** Because in the common use case,  $M$  remains fixed for a large number of cycles while  $O$  and  $L$  change, and because changes to  $M$  require multiplication, the `mod_addr` module has a configuration mechanism to change  $M$ , while  $O$  and  $L$  can be changed on a per clock basis. To avoid a need for division,  $O$  is typically provided in the format of a quotient and a remainder relative to  $D$ , which the various indexing units produce. That is,  $O = Q * D + R$ , where  $0 \leq R < D$ . The multiplier  $M$  is also provided in a quotient-remainder form (e.g.,  $QM$  and  $RM$ ), as this limits the size of division required to  $\text{ceil}(\log_2(D))$  bits, and also allows an implementation using serial addition. That is, one may define  $M = QM * D + RM$ , where again,  $0 \leq R < D$ . In addition, as introduced previously, a precomputed cyclic shift based on  $M$ ,  $\text{CSHIFT} = (\log_2(M) + 1) \% (D - 1)$ , may be defined for generator  $g$ .

**[00126]** It is worth noting that `mod_addr` employs mechanism used by `fwd_mod_perm` and `rev_mod_perm` defined above.

**[00127]** A set of design parameters, such as the parameters shown below along with their respective default values, determine the overall size and performance of the modulo address generator. The default values are exemplary and may be subject to change depending on design requirements.

| Name      | Default Value | Description/Formula                       |
|-----------|---------------|-------------------------------------------|
| DIVISOR   | 7             | The divisor, $D$ , a prime                |
| GENERATOR | 3             | Generator for prime group                 |
| REM_WIDTH | 3             | $\text{ceil}(\log_2(\text{DIVISOR} + 1))$ |
| PHY_WIDTH | 16            | Physical (per bank) address bit width     |
| ADDR_BUS  | 112           | $\text{DIVISOR} * \text{PHY\_WIDTH}$      |

**[00128]** The modulo address generator may have the following set of external pins via which the counter unit is communicatively coupled to and interact with other components of the TPU.

| Name   | Width     | Direction | Description                   |
|--------|-----------|-----------|-------------------------------|
| CLK    | 1         | IN        | Global clock                  |
| RST    | 1         | IN        | Global Reset                  |
| CFG    | 1         | IN        | Configuration enable          |
| CFG_OK | 1         | OUT       | Configuration complete        |
| MULT_Q | PHY_WIDTH | IN        | Quotient of $M/D$ for config  |
| MULT_R | REM_WIDTH | IN        | Remainder of $M/D$ for config |

|        |           |     |                                     |
|--------|-----------|-----|-------------------------------------|
| CSHIFT | REM_WIDTH | IN  | Precomputed cyclic shift for config |
| OFF_Q  | PHY_WIDTH | IN  | Quotient of O/D                     |
| OFF_R  | REM_WIDTH | IN  | Remainder of O/D                    |
| L      | REM_WIDTH | IN  | Length                              |
| OUT    | ADDR_BUS  | OUT | Per bank addressing                 |
| OUT_EN | DIVISOR   | OUT | Per bank enable                     |

**[00129]** The mod\_addr module contains two submodules of type fwd\_mod\_perm, named “addr\_perm” and “enable\_perm”, respectively. The DIVISOR and GENERATOR parameters of these modules are set to the same values as the mod\_addr module. The WIDTH parameter of the “addr\_perm” module comes from the PHY\_WIDTH parameter, and the WIDTH of the “enable\_perm” module is 1. Since both submodules always receive the same N and O values, it is also possible to use a single fwd\_mod\_perm submodule with a WIDTH of PHY\_WIDTH + 1, but for the sake of clarity, two submodules are presumed hereafter.

**[00130]** The mod\_addr module contains internal registers that hold a quotient-remainder representation of  $i \cdot M$  for each  $i$ ,  $0 \leq i < D$ . The registers are referred to as MO\_Q[i] and MO\_R[i], respectively. The addition of these values with the current input offset are stored in address registers, A[i], each of PHY\_WIDTH. The enable lines live in enable registers, E[i]. In addition, the cyclic shift is saved in the N register, and current offset lives in the O register. Finally, since the configuration may take multiple clock cycles to complete, there is a CFG\_CYCLE counter which counts down till configuration completion.

**[00131]** The operation of the modulo address generator may be described by the virtual code below. In the virtual code, X' represents the value of X after the clock, and any registers not modified explicitly are assumed to remain the same.

always:

```

*_perm.CLK = CLK
*_perm.RST = RST
*_perm.N = N
*_perm.O = O
addr_perm.IN = A[i]
enable_perm.IN = E[i]
OUT = addr_perm.OUT
OUT_EN = enable_perm.OUT

```

```

    CFG_OK = CFG_CYCLE == 0
if RST is high:
    For all i, MO_Q[i] = 0
    For all i, MO_R[i] = i
    For all i, E[i] = 0
    For all i, A[i] = 0
    N' = 0
    CFG_CYCLE' = 0
    enable_perm.IN = all 0's
else if CFG is high and CFG_CYCLE = 0:
    CFG_CYCLE' = config time
    For all i, E[i] = 0
    For all i, A[i] = 0
    Begin configuration, eventually resulting in:
        MO_Q'[i] = MULT_Q*i + floor(MULT_R*i/DIVISOR)
        MO_D'[i] = MULT_R*i%DIVISOR
        N' = SHIFT
        enable_perm.IN = all 0's
else if CFG_CYCLE != 0:
    CFG_CYCLE' = CFG_CYCLE - 1
    enable_perm.IN = all 0's
else:
    O' = OFF_R
    For all i
        E[i]' = (i < L)
        A[i]' = OFF_Q[i]+MO_Q[i]+(OFF_R[i]+MO_R[i] >= DIVISOR:1:0)

```

**[00132]** The output values may lag the input values by a pipeline depth induced by the internal registers of the mod\_addr module as well as the pipeline delay induced by the fwd\_mod\_perm module.

### **DRAM Controller**

**[00133]** DRAM reads and writes are initiated from the memory transfer engine, which is described in a later part of the present disclosure. The memory transfer engine utilizes an abstraction which performs reads of two-dimensional tensor slices. The DRAM read of two-dimensional tensor slices is logically divided into two problems. First, it is desired to

reads up to N reads of full bank width at some fixed stride and variable offset, where N is less than or equal to the number of banks. Second, it is desired to read up to the full bank width of banks elements, so that the resulting tensor is square, and thus amenable to transposition. In addition, it is required to deal with the non-deterministic nature of DRAM for the reads and writes, with the assumption that each bank operates independently. Therefore, the employment of a queuing technique may be beneficial.

*DRAM Bank Read Unit (dram\_bank\_read)*

**[00134]** The overall design of the DRAM bank read unit is shown in Figure 8. Logic is shown in rectangles, and queues are shown in ovals therein. At a very high level, vector address requests are generated and bank permuted appropriately, and put into queues for DRAM to fetch from. DRAM fetches are performed independently (and possibly asynchronously) per bank, and the results are enqueued until all data is available, at which point the permutation is undone and the data gathered. This results in a VALID signal (as shown in Figure 8) going high. Up to MAX\_QUEUE requests may exist with the dram\_bank\_read unit simultaneously, and each request's response is returned in identical order. However, the specific number of cycles each request takes is subject to queueing and DRAM non-determinacy.

**[00135]** A set of design parameters, such as the parameters shown below along with their respective default values, determine the overall size and performance of the design of the DRAM bank read unit 800. The default values are exemplary and may be subject to change depending on design requirements.

| Name       | Default Value | Description/Formula                                                         |
|------------|---------------|-----------------------------------------------------------------------------|
| DIVISOR    | 11            | Number of banks, i.e., bank divisor, must be prime                          |
| GENERATOR  | 2             | Generator for multiplicative group of DIVISOR                               |
| BANK_WIDTH | 32            | Width of single bank (in elements)                                          |
| PHY_WIDTH  | 28            | Physical (per bank) address bit width                                       |
| ELEM_WIDTH | 16            | Bit width of an element                                                     |
| MAX_QUEUE  | 48            | Maximum queue size                                                          |
| REM_WIDTH  | 4             | $\text{ceil}(\log_2(\text{DIVISOR}))$                                       |
| LEN_WIDTH  | 4             | $\text{ceil}(\log_2(\text{DIVISOR} + 1))$                                   |
| DBUS_WIDTH | 5632          | $\text{DIVISOR} * \text{BANK\_WIDTH} * \text{ELEM\_WIDTH}$ (Data bus width) |

**[00136]** The DRAM bank read unit 800 may have the following set of external pins via which the DRAM bank read unit 800 is communicatively coupled to and interact with other components of the TPU.

| Name     | Width      | Direction | Description                         |
|----------|------------|-----------|-------------------------------------|
| CLK      | 1          | IN        | Global clock                        |
| RST      | 1          | IN        | Global Reset                        |
| CFG      | 1          | IN        | Configuration enable                |
| CFG_OK   | 1          | OUT       | Configuration complete              |
| MULT_Q   | PHY_WIDTH  | IN        | Quotient of multiplier for config   |
| MULT_R   | REM_WIDTH  | IN        | Remainder of multiplier for config  |
| CSHIFT   | REM_WIDTH  | IN        | Precomputed cyclic shift for config |
| OFF_Q    | PHY_WIDTH  | IN        | Quotient of starting offset         |
| OFF_R    | REM_WIDTH  | IN        | Remainder of starting offset        |
| L        | LEN_WIDTH  | IN        | Length                              |
| VALID    | 1          | OUT       | Is there data to read               |
| DATA_OUT | DBUS_WIDTH | OUT       | The data                            |

**[00137]** The dram\_bank\_read module 800 has a mod\_addr submodule 820 and a rev\_mod\_perm module 880. All inputs to the dram\_bank\_read module 800 are sent to the inputs of the mod\_addr module 820 with the same names. CLK, and RST are sent directly to the rev\_mod\_perm. In addition, a register maintaining CSHIFT and set by CFG is passed to the N input of rev\_mod\_perm. The CFG\_OK output of mod\_addr is wired directly to the CFG\_OK output of the dram\_bank\_read module itself. Also, the OUT bus from rev\_mod\_perm is connected to DATA\_OUT.

**[00138]** Additionally, the dram\_bank\_read module 800 includes a set of per-bank DRAM queues 840. Each queue holds up to MAX\_QUEUE entries. There is no internal logic to verify queue flow control. Instead, each queue can hold the same number of entries, and the circuit utilizing the dram\_bank\_read 800 must ensure that no more than MAX\_QUEUE read operations are in the current pipeline, perhaps via a counter. This prevents any queue from being full.

**[00139]** The first type of queues are the address queues (shown as AQ in Figure 8). They hold bank specific DRAM addresses, and there is one such queue per DRAM bank. The second type of queue is the data queues (shown as DQ in Figure 8). They hold full width DRAM bank outputs, and again, are one per bank. The final queue is the control

queue (CQ) in Figure 8. It holds control data, specifically the OFF\_R and a set of per-bank enable bits for the read operation. There is exactly one control queue.

**[00140]** The dram\_bank\_read module 800 has only one register, N, which holds the N value used by rev\_mod\_perm.

**[00141]** The operation of the DRAM bank read unit 800 is detailed below:

**[00142]** If RST is high, all components are reset, queue are reset to empty, and N is cleared. If CFG is high, CFG is passed to mod\_addr, and the N register is set to the value in CSHIFT. Otherwise, the normal operation follows. The behavior is described in terms of three separate components: enqueue\_cntl 830 which sets up reads, deque\_cntl 860 which combines banks to make a complete read, and the actual DRAM readers, represented in Figure 8 as D.

**[00143]** At a high level, every timestep a read may arrive, or alternately, the lack of a read may be represented by a length of zero. Only the banks actually read from have read addresses queued. If any banks are read from (length > 0), an additional control message is queued containing the offset modulus to undo the permutation in rev\_mod\_perm, as well as the OUT\_EN bits which were set. Each DRAM reader D then pulls addresses from it's input queue AQ and writes data to it's output data queue DQ independently. Finally, the deque\_cntl unit 860 examines the top of the control queue to determine which data queues must also be ready for the read results to be reconstructed. If all the required banks are ready, the deque\_cntl 860 removes the appropriate queue entries and sends the now combined reads through the rev\_mod\_perm module 880, using the stored offset modulus to permute the output. In addition, an appropriately delayed VALID signal is sent to arrive at the output on the same clock as DATA\_OUT.

**[00144]** The enqueue\_cntl logic 830 operates as follows: The OUT value from mod\_addr 820 is divided into banks and becomes the queue input for each address queue. The same bank's OUT\_EN value is used as the ENQUEUE signal for that bank's address queue. In addition, a reduction OR operation is performed on the entire OUT\_EN field, resulting in a signal ANY\_EN. The full OUT\_EN field along with an appropriately delayed copy of the OFF\_R value are given as the input to the control queue, and ANY\_EN serves as the queue's ENQUEUE signal.

**[00145]** The DRAM readers D operate as follows: Each DRAM reader is bank specific, and is bounded on both sides by a queue. The DRAM reader will pull from the address queue (if not empty), schedule a read, and write the results in order to the data queue, which should never be full due to the requirement of no more than MAX\_QUEUE entries

in flight in the entire dram\_bank\_read module 800. All DRAM reads will be full bank width. Beyond this, the DRAM reader implementation may use any method appropriate for the DRAM technology in use.

**[00146]** The deque\_cntl logic 860 operates as follows: The deque\_cntl module 860 will examine the EMPTY value of each data queue, as well as EMPTY value in the control queue, as well as the enable bits from the front entry of the control queue. When the control queue is not empty, and for each enable bit set in the front entry, the data queue for that bank is also not empty, VALID (with appropriate pipeline delay to match rev\_mod\_perm) will be set to true. In addition, in this case, the DEQUE signal will be asserted for each data queue for which an enable bit is set. Finally, the input to the rev\_mod\_perm module 880's IN bus will consist of the front entry for all data queues for which the associated enable bit is set, or if the bit is not set, an all zero value.

*DRAM Tile Read Unit (dram\_tile\_read)*

**[00147]** A diagram of a DRAM tile read unit 900 is shown in Figure 9. At a very high level, the goal of the DRAM tile read unit 900 is to read tiles of fixed width (i.e., bank size) and a variable height (i.e., up to bank size). The resulting elements are zero padded to a square array of bank size elements. Logic is shown in rectangles, and the queue is shown as an oval in Figure 9. Internally, the tile read unit 900 uses the bank read unit 800 of Figure 8 to perform one or more banked reads of the actual data, and then composes the results into a final output.

**[00148]** A set of design parameters, such as the parameters shown below along with their respective default values, determine the overall size and performance of the design of the DRAM tile read unit 900. The default values are exemplary and may be subject to change depending on design requirements.

| Name       | Default Value | Description/Formula                                |
|------------|---------------|----------------------------------------------------|
| DIVISOR    | 11            | Number of banks, i.e., bank divisor, must be prime |
| GENERATOR  | 2             | Generator for multiplicative group of DIVISOR      |
| BANK_WIDTH | 32            | Width of single bank (in elements)                 |
| PHY_WIDTH  | 28            | Physical (per bank) address bit width              |
| ELEM_WIDTH | 16            | Bit width of an element                            |
| MAX_QUEUE  | 16            | Maximum queue size                                 |
| REM_WIDTH  | 4             | $\text{ceil}(\log_2(\text{DIVISOR}))$              |

|             |       |                                                                                    |
|-------------|-------|------------------------------------------------------------------------------------|
| LEN_WIDTH   | 6     | $\text{ceil}(\log_2(\text{DIVISOR} + 1))$                                          |
| DBUS_WIDTH  | 16384 | $\text{BANK\_WIDTH} * \text{BANK\_WIDTH} * \text{ELEM\_WIDTH}$<br>(Data bus width) |
| MAX_STEPS   | 3     | $\text{ceil}(\text{BANK\_WIDTH} / \text{DIVISOR})$                                 |
| CTR_SIZE    | 2     | $\text{ceil}(\log_2(\text{MAX\_STEPS}))$                                           |
| INNER_QUEUE | 48    | $\text{MAX\_STEPS} * \text{MAX\_QUEUE}$                                            |

**[00149]** The DRAM tile read unit 900 may have the following set of external pins via which the DRAM tile read unit 900 is communicatively coupled to and interact with other components of the TPU.

| Name     | Width      | Direction | Description                         |
|----------|------------|-----------|-------------------------------------|
| CLK      | 1          | IN        | Global clock                        |
| RST      | 1          | IN        | Global Reset                        |
| CFG      | 1          | IN        | Configuration enable                |
| CFG_OK   | 1          | OUT       | Configuration complete              |
| MULT_Q   | PHY_WIDTH  | IN        | Quotient of multiplier for config   |
| MULT_R   | REM_WIDTH  | IN        | Remainder of multiplier for config  |
| CSHIFT   | REM_WIDTH  | IN        | Precomputed cyclic shift for config |
| OFF_Q    | PHY_WIDTH  | IN        | Quotient of starting offset         |
| OFF_R    | REM_WIDTH  | IN        | Remainder of starting offset        |
| L        | LEN_WIDTH  | IN        | Length                              |
| VALID    | 1          | OUT       | Is there data to read               |
| DATA_OUT | DBUS_WIDTH | OUT       | The data                            |

**[00150]** The dram\_tile\_read module wraps a dram\_bank\_read module. The non-computed parameters of the dram\_bank\_read module 920 follow those of the dram\_tile\_read module 900, except for MAX\_QUEUE, which is set to INNER\_QUEUE. The CLK, RST, CFG, MULT\_Q, MULT\_R, and CSHIFT inputs of the dram\_tile\_read module 900 are directly wired to the dram\_bank\_read 920, in addition to being used by the rest of the module 900. Only the OFF\_Q, OFF\_R and L inputs to dram\_bank\_read 920 are set by addr\_splitter 940.

**[00151]** The full\_read\_queue 960 contains an address (Q + R portions) and a length, L. It holds MAX\_QUEUE entries. The control\_queue (“CQ” in Figure 9) holds a single bit, and consists of INNER\_QUEUE entries.

**[00152]** The operation of the DRAM tile read unit 900 is detailed below:

[00153] When L is > 0, the OFF\_Q, OFF\_R and L values are enqueued into full\_read\_queue 960. Otherwise, the operation of the dram\_tile\_read 900 is logically divided into two subcomponents, addr\_splitter 940, and data\_joiner 980, which are defined below.

[00154] The address splitter addr\_splitter 940 operates as follows: The addr\_splitter component 940 holds 3 registers:

| Name   | Width     | Description                     |
|--------|-----------|---------------------------------|
| STEP_Q | PHY_WIDTH | Amount to add to OFF_Q per step |
| ADD_Q  | PHY_WIDTH | Amount to add to OFF_Q now      |
| L_OFF  | LEN_WIDTH | How far into L                  |

[00155] When RST is high, all registers are set to 0. The L value of dram\_bank\_read 920 is set to zero.

[00156] When CFG is high, STEP\_Q is set to  $MULT\_Q * DIVISOR + MULT\_R$ , and all other registers are set to 0. CFG\_OK will be set to high when this operation is complete, and CFG\_OK from dram\_bank\_read 920 is also high. The L value of dram\_bank\_read 920 is set to zero.

[00157] Otherwise:

if full\_read\_queue.EMPTTY:

dram\_bank\_read.L = 0

full\_read\_queue.DEQUEUE = 0

control\_queue.ENQUEUE = 0

else:

L\_REM = full\_read\_queue.OUTPUT.L - L\_OFF

IS\_DONE = (L\_REM <= DIVISOR)

dram\_bank\_read.OFF\_Q = full\_read\_queue.OUTPUT.Q + ADD\_Q

dram\_bank\_read.OFF\_R = full\_read\_queue.OUTPUT.R + ADD\_Q

dram\_bank\_read.L = (IS\_DONE ? L\_REM : DIVISOR)

full\_read\_queue.DEQUEUE = IS\_DONE

control\_queue.INPUT = IS\_DONE

control\_queue.ENQUEUE = 1

ADD\_Q' (IS\_DONE ? 0 : ADD\_Q + STEP\_Q)

L\_OFF' (IS\_DONE ? 0 : L\_OFF + DIVISOR)

**[00158]** The data joiner component, data\_joiner 980, holds 2 registers, DATA\_OUT, which also acts as the output value of the same name, and a counter C, which holds CTR\_SIZE bits.

**[00159]** The operation of the data joiner 980 is as follows:

```

if dram_bank_read.VALID:
    DATA_OUT[C*DIVISOR : C*DIVISOR + DIVISOR - 1] =
dram_bank_read.DATA_OUT
    control_queue.DEQUEUE = 1
    if control_queue.OUTPUT:
        C' = 0
        VALID = 1
    else:
        C' = C + 1
        VALID = 0
else:
    VALID = 0
    control_queue.DEQUEUE = 1

```

#### *DRAM Bank Write Unit (dram\_bank\_write)*

**[00160]** An illustrative figure of the DRAM bank write unit 1000 is shown in Figure 10. The design of DRAM bank write unit 1000, dram\_bank\_write, follows the design of the DRAM bank read unit 800 of Figure 8, dram\_bank\_read, very closely. For example, DRAM bank write unit 1000 also includes a mod\_addr submodule 1020, just as DRAM bank read unit 800 has mod\_addr submodule 820. In addition, DRAM bank write unit 1000 also includes a set of per-bank DRAM queues 1040, just as DRAM bank read unit 800 has a set of per-bank DRAM queues 840. The primary difference is that in addition to the addressing information, it also receives the data to write, and the output is simply an acknowledgment of completion.

**[00161]** The enqueue\_cntl logic 1030 is nearly identical to that of dram\_bank\_read 800, as is the dequeue\_cntl logic 1060. The design parameters are exactly identical to those of the DRAM bank read unit 800 presented above. The pinouts differ slightly, being:

| Name | Width | Direction | Description  |
|------|-------|-----------|--------------|
| CLK  | 1     | IN        | Global clock |
| RST  | 1     | IN        | Global Reset |

|         |            |     |                                     |
|---------|------------|-----|-------------------------------------|
| CFG     | 1          | IN  | Configuration enable                |
| CFG_OK  | 1          | OUT | Configuration complete              |
| MULT_Q  | PHY_WIDTH  | IN  | Quotient of multiplier for config   |
| MULT_R  | REM_WIDTH  | IN  | Remainder of multiplier for config  |
| CSHIFT  | REM_WIDTH  | IN  | Precomputed cyclic shift for config |
| OFF_Q   | PHY_WIDTH  | IN  | Quotient of starting offset         |
| OFF_R   | REM_WIDTH  | IN  | Remainder of starting offset        |
| L       | LEN_WIDTH  | IN  | Length                              |
| DATA_IN | DBUS_WIDTH | IN  | The data                            |
| ACK     | 1          | OUT | Acknowledge one write               |

[00162] The key change is DATA\_OUT becomes DATA\_IN, and VALID becomes ACK. Operationally, the following changes exist:

[00163] A fwd\_mod\_perm 1080 is applied to the data. Specifically, fwd\_mod\_perm.O is set via OFF\_R, and fwd\_mod\_perm.N is set via CSHIFT. The output may require pipelining delay to assure that the address information generated by a write arrives at the queues at the same time as the data.

[00164] The queue before the DRAM operation (D) is now a write queue (WQ), containing address, enable, and data as it's entries. That is, data is added. The queue after the DRAM operation is now simply an acknowledgment (AC), which means that the queue can be replaced by a counter, such that EMPTY is defined as cnt==0, and ENQUEUE is defines as cnt=cnt+1, DEQUEUE is defined as cnt=cnt-1. The counter may be referred to as the acknowledge counter.

[00165] The actual DRAM operation now performs a write of it's input queue WQ, and when the write is acknowledged, increments the acknowledge counter AC.

[00166] The deque\_cntl logic 1060 is identical to that of the DRAM bank read unit 800 of Figure 8, except that the deque\_cntl logic 1060 of Figure 10 does not generate data, and VALID is renamed ACK.

#### *DRAM Tile Write Unit (dram\_tile\_write)*

[00167] The DRAM tile write unit, as shown in Figure XX, is very similar to that of the DRAM tile read unit described above, except that the pinouts are slightly changed to:

| Name | Width | Direction | Description  |
|------|-------|-----------|--------------|
| CLK  | 1     | IN        | Global clock |
| RST  | 1     | IN        | Global Reset |

|         |            |     |                                     |
|---------|------------|-----|-------------------------------------|
| CFG     | 1          | IN  | Configuration enable                |
| CFG_OK  | 1          | OUT | Configuration complete              |
| MULT_Q  | PHY_WIDTH  | IN  | Quotient of multiplier for config   |
| MULT_R  | REM_WIDTH  | IN  | Remainder of multiplier for config  |
| CSHIFT  | REM_WIDTH  | IN  | Precomputed cyclic shift for config |
| OFF_Q   | PHY_WIDTH  | IN  | Quotient of starting offset         |
| OFF_R   | REM_WIDTH  | IN  | Remainder of starting offset        |
| L       | LEN_WIDTH  | IN  | Length                              |
| DATA_IN | DBUS_WIDTH | IN  | The data                            |
| ACK     | 1          | OUT | Acknowledge one write               |

**[00168]** In addition, the queue of operations (now called `full_write_queue`) includes the full tile data to write. The `dram_bank_read` is replaced by a `dram_bank_write`, and the data to the `DATA_IN` of the `dram_bank_write` is the `DIVISOR` banks of elements from the `full_write_queue` starting at `L_OFF`. It is to be noted that, since `L_OFF` can only be a multiple of `DIVISOR`, it may be useful to add an additional counter register so this slicing can be done via simple muxing. Finally, there is no output generated by `data_joiner`, and `VALID` is replaced by `ACK`.

### **SRAM Controller**

**[00169]** SRAM is composed of a prime number of banks, each multiple elements wide, and is dual port. One port of each bank of SRAM is attached to the execution subsystem, and support simple contiguous 'linear' addressing. The other port of SRAM reads and writes tiles on behalf of the memory transfer engine, generally reading from and writing to DRAM.

#### *RAM Tile Read Unit (`sram_tile_read`)*

**[00170]** At a very high level, the goal of the SRAM tile read unit is to read tiles of fixed width (i.e., bank size) and a variable height (i.e., up to bank size). The resulting elements are zero padded to a square array of bank size elements. Components include: *SRAM Bank Write Unit (`sram_bank_write`) and SRAM Tile Write Unit (`sram_tile_write`)*

### **Memory Transfer Engine**

**[00171]** As described earlier, the goal of memory transfer engine 380 of Figure 3 is to move tensor slices from one memory space to another. Since each tensor is arranged in such a way that all but the final dimensions have strides which are multiples of bank width, and the final dimension is packed (i.e., stride of 1), all dimensions can be treated

equally except the final one. Given that a transfer has a source and a destination, the final dimension of each means that at a minimum, two dimensional tiles must be considered during transfer.

**[00172]** The memory transfer engine acts as a generic mechanism which includes a tile reading circuit and a tile writing circuit. However, in general for each transfer case (i.e., from SRAM to DRAM, and from DRAM to SRAM) a separate instance of the entire transfer module exists. That is, it is generic at design time.

**[00173]** A set of design parameters, such as the parameters shown below along with their respective default values, determine the overall size and performance of the design of the memory transfer engine. The default values are exemplary and may be subject to change depending on design requirements.

| Name          | Default Value | Description/Formula                                      |
|---------------|---------------|----------------------------------------------------------|
| IDX_COUNT     | 6             | Number of indexes                                        |
| IDX_WIDTH     | 16            | Index length in bits                                     |
| IDX_BUS       | 96            | $IDX\_COUNT * IDX\_WIDTH$                                |
| IDX_SEL       | 3             | $\text{ceil}(\log_2(IDX\_COUNT))$                        |
| MAX_QUEUE     | 32            | Maximum queued tile reads + writes                       |
| RD_DIVISOR    | 11            | Number of read banks, i.e., bank divisor, must be prime  |
| RD_REM_WIDTH  | 4             | $\text{ceil}(\log_2(DIVISOR))$                           |
| RD_PHY_WIDTH  | 28            | Physical (per bank) address bit width                    |
| RD_ADDR_WIDTH | 32            | $RD\_REM\_WIDTH + RD\_PHY\_WIDTH$                        |
| RD_POLY_WIDTH | 192           | $RD\_ADDR\_WIDTH * IDX\_COUNT$                           |
| WR_DIVISOR    | 131           | Number of write banks, i.e., bank divisor, must be prime |
| WR_REM_WIDTH  | 8             | $\text{ceil}(\log_2(DIVISOR))$                           |
| WR_PHY_WIDTH  | 10            | Physical (per bank) address bit width                    |
| WR_ADDR_WIDTH | 18            | $WR\_REM\_WIDTH + WR\_PHY\_WIDTH$                        |
| WR_POLY_WIDTH | 108           | $WR\_ADDR\_WIDTH * IDX\_COUNT$                           |
| BANK_WIDTH    | 32            | Maximum length of combined read/write                    |
| L_WIDTH       | 5             | $\text{ceil}(\log_2(MAX\_LENGTH))$                       |
| QUEUE_WIDTH   | 6             | $\text{ceil}(\log_2(MAX\_QUEUE + 1))$                    |

**[00174]** The memory transfer engine may have the following set of external pins via which the memory transfer engine is communicatively coupled to and interact with other components of the TPU.

| Name      | Width         | Direction | Description                 |
|-----------|---------------|-----------|-----------------------------|
| CLK       | 1             | IN        | Global clock                |
| RST       | 1             | IN        | Global Reset                |
| CFG       | 1             | IN        | Configuration enable        |
| CFG_OK    | 1             | OUT       | Configuration complete      |
| RUN       | 1             | IN        | Run enable                  |
| DONE      | 1             | OUT       | Run Complete                |
| IDX_LENS  | IDX_BUS       | IN        | Index sizing                |
| DO_TRANS  | 1             | IN        | Do transpose                |
| RD_BASE   | RD_ADDR_WIDTH | IN        | Base read address           |
| RD_POLY   | RD_POLY_WIDTH | IN        | Read polynomial             |
| RD_FMUL   | RD_ADDR_WIDTH | IN        | Multiplier for read index   |
| RD_FIDX   | IDX_SEL       | IN        | Index number for read index |
| RD_FLEN   | BANK_WIDTH    | IN        | Final length for read index |
| RD_CSHIFT | RD_REM_WIDTH  | IN        | CSHIFT for RD_FMUL          |
| WR_BASE   | RD_ADDR_WIDTH | IN        | Base read address           |
| WR_POLY   | WR_POLY_WIDTH | IN        | Read polynomial             |
| WR_FMUL   | WR_ADDR_WIDTH | IN        | Multiplier for read index   |
| WR_FIDX   | IDX_SEL       | IN        | Index number for read index |
| WR_FLEN   | BANK_WIDTH    | IN        | Final length for read index |
| WR_CSHIFT | WR_REM_WIDTH  | IN        | CSHIFT for WR_FMUL          |

**[00175]** The memory transfer engine includes the following submodules, as shown in the table below:

| Module Type | Module Name |
|-------------|-------------|
| mdim_count  | rd_cntr     |
| mdim_count  | wr_cntr     |
| mdim_addr   | rd_addr     |
| mdim_addr   | wr_addr     |
| tile_reader | rd_tile     |
| tile_writer | wr_tile     |

[00176] In the table above, each of rd\_cntr and wr\_cntr is a multidimensional counter unit that is described previously. Also, each of rd\_addr and wr\_addr is a multidimensional address unit that is described previously. In addition, tile\_reader is a tile reading module (such as dram\_tile\_read module described previously), while tile\_writer is a tile writing module (such as dram\_tile\_write module described previously). Most of the pins to each submodule consist of fixed connections either to the main modules or between submodules, as per the following table:

| Source Module | Pins      | Destination Module | Pins     |
|---------------|-----------|--------------------|----------|
| mem_tx        | CLK       | *                  | CLK      |
| mem_tx        | RST       | *                  | RST      |
| mem_tx        | CFG       | *                  | CFG      |
| mem_tx        | IDX_LENS  | rd_cntr            | CFG_IN   |
| mem_tx        | RD_BASE.Q | rd_addr            | QUOT_IN  |
| mem_tx        | RD_BASE.R | rd_addr            | REM_IN   |
| mem_tx        | RD_POLY   | rd_addr            | DELTA_IN |
| mem_tx        | RD_FMUL.Q | rd_tile            | MULT_Q   |
| mem_tx        | RD_FMUL.R | rd_tile            | MULT_R   |
| mem_tx        | RD_CSHIFT | rd_tile            | CSHIFT   |
| rd_cntr       | NEXT      | rd_addr            | IDX      |
| mem_tx        | IDX_LENS  | wr_cntr            | CFG_IN   |
| wr_addr       | QUOT      | wr_tile            | OFF_Q    |
| wr_addr       | REM       | wr_tile            | OFF_R    |
| mem_tx        | WR_BASE.Q | wr_addr            | QUOT_IN  |
| mem_tx        | WR_BASE.R | wr_addr            | REM_IN   |
| mem_tx        | WR_POLY   | wr_addr            | DELTA_IN |
| mem_tx        | WR_FMUL.Q | wr_tile            | MULT_Q   |
| mem_tx        | WR_FMUL.R | wr_tile            | MULT_R   |
| mem_tx        | WR_CSHIFT | wr_tile            | CSHIFT   |
| rd_cntr       | NEXT      | wr_addr            | IDX      |
| wr_addr       | QUOT      | wr_tile            | OFF_Q    |
| wr_addr       | REM       | wr_tile            | OFF_R    |

[00177] In the table above, \* refers to all modules which have a pin by the proper name. Also, .Q and .R refer to the quotient and remainder part of combined buses. In addition, mem\_tx. CFG\_OK is just the AND of rd\_tile.CFG\_OK and wr\_tile.CFG\_OK.

[00178] The memory transfer engine includes the following registers:

| Name     | Width       | Description                |
|----------|-------------|----------------------------|
| PEND     | QUEUE_WIDTH | How many reads are pending |
| RD_FIDX  | IDX_SEL     | Stored value of RD_FIDX    |
| RD_FLEN  | IDX_SEL     | Stored value of RD_FLEN    |
| WR_FIDX  | IDX_SEL     | Stored value of WR_FIDX    |
| WR_FLEN  | IDX_SEL     | Stored value of WR_FLEN    |
| DO_TRANS | 1           | Stored value of DO_TRANS   |

[00179] The operation of the memory transfer engine is as follows:

[00180] If RST is held high, all modules are reset, and the registers are zeroed.

[00181] If CFG is held high, PEND is zeroed, and RD\_FIDX, RD\_FLEN, WR\_FIDX, and WR\_FLEN registers are set to the current input values. When all submodules report CFG\_OK, the transfer module reports CFG\_OK.

[00182] Always, the value of rd\_tile.DATA\_OUT is passed to wr\_tile.DATA\_IN, either transposed or not, based on the value of the DO\_TRANS register. In the case where the tile elements are transposed, the value of rd\_tile.DATA\_OUT is considered as a BANK\_WIDTH by BANK\_WIDTH tile, and the tile elements are transposed when passed to wr\_tile.DATA\_IN.

[00183] The value of wr\_tile.L, wr\_cntr.STEP, and wr\_addr.STEP are set as follows:

if (RST or CFG or !wr\_tile.VALID):

    wr\_cntr.STEP = 0

    wr\_addr.STEP = 0

    wr\_tile.L = 0

else:

    wr\_cntr.STEP = 1

    wr\_addr.STEP = 1

    if (wr\_cntr != WR\_FIDX):

        wr\_tile.L = BANK\_WIDTH

    else:

        wr\_tile.L = WR\_FLEN + 1

**[00184]** The value of rd\_tile.L, rd\_cntr.STEP and rd\_addr.STEP are set as follows:

```

if (RST or CFG or rd_cntr.OVER):
    rd_cntr.STEP = 0
    rd_addr.STEP = 0
    rd_tile.L = 0
else:
    rd_cntr.STEP = 1
    rd_addr.STEP = 1
    if (wr_cntr != RD_FIDX):
        rd_tile.L = BANK_WIDTH
    else:
        rd_tile.L = RD_FLEN + 1

```

**[00185]** PEND is increased by one if a read is started, and decreased by one if a write is complete (i.e., wr\_tile.ACK is true). Both of these may happen in the same cycle. Finally, DONE is set when rd\_cntr.DONE is true, wr\_cntr.DONE is true, and PEND is 0.

### **Conclusion**

**[00186]** Although the subject matter has been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to the specific features or acts described above. Rather, the specific features and acts described above are disclosed as example forms of implementing the claims.

## CLAIMS

WHAT IS CLAIMED IS:

1. An execution unit of hardware acceleration for machine learning, the execution unit running on a clock and configured to execute one instruction of an instruction set per clock cycle of the clock, the execution unit comprising:

- a clock input receiving the clock;
- an operation specification input receiving an opcode specifying the one instruction;
- a first operand input receiving a first operand;
- a second operand input receiving a second operand;
- an accumulation register storing a current accumulation value; and
- an output,

wherein:

- the accumulation register is updated, within the clock cycle, with a new accumulation value based on the opcode and one or more of the first operand, the second operand and the current accumulation value, and

- the output comprises saturated lower bits of the new value.

2. The execution unit of claim 1, wherein the instruction set comprises instructions comprising:

- a multiply operation;
- a multiply accumulate operation;
- an arithmetic shift right operation;
- a linear transform operation;
- a compare operation;
- an add random noise operation; and
- a seed random number generator operation.

3. The execution unit of claim 1, further comprising:

- a random number register storing a current random number,

- wherein the one instruction comprises a seed random number generator operation, and wherein the random number register is updated, within the clock cycle, with a new random number based on the first operand, the second operand and the current random number.

4. The execution unit of claim 1, further comprising:  
a conditional execution input receiving a conditional execution,  
wherein the one instruction is executed or skipped as determined by the conditional execution.
5. The execution unit of claim 4, further comprising:  
a conditional set mask receiving a conditional set mask;  
a conditional reset mask receiving a conditional reset mask; and  
a status register,  
wherein, when the one instruction is executed as determined by the conditional execution, bits of the status register are updated, within the clock cycle, based on corresponding bits of the conditional set mask and corresponding bits of the conditional reset mask.
6. The execution unit of claim 5, wherein the one instruction comprises a compare operation, and wherein bits of the status register are updated only when the compare operation returns a true result.
7. An indexing module for looping over a plurality of indexes each corresponding to one of multiple dimensions of a tensor, the indexing module comprising:  
a counter unit comprising:  
a plurality of counter registers each storing a respective one of the plurality of indexes; and  
a plurality of size registers each corresponding to one of the counter registers; and  
an address unit,  
wherein:  
the counter unit updates the plurality of indexes stored in the counter registers by repeating a method to loop over the plurality of indexes, and  
the address unit generates, for each of the indexes stored in the counter registers, an address that is a polynomial of the respective index.
8. The indexing module of claim 7, wherein the address unit further generates a quotient and a remainder of the address by dividing the address with a fixed modulus.

9. The indexing module of claim 8, wherein the address unit receives, for each of the indexes stored in the counter registers, a delta comprising a quotient delta and a reminder delta with respective to the fixed modulus, and wherein the address unit updates of the address by adding the quotient delta to the quotient of the address and the reminder delta to the reminder of the address.
10. The indexing module of claim 7, wherein:  
the indexing module is driven by a clock,  
the counter unit further comprises an OVER register and a NEXT register, and  
the method comprising:  
    resetting the counter registers and the size registers to 0;  
    setting the OVER register to 1 and the NEXT register to 0;  
    configuring the counter registers with a configuration input;  
    configuring the size registers with the configuration input;  
    setting the OVER register to 0 if any counter register contains a non-zero value;  
    updating the NEXT register based on a counter register of a highest sequence of the counter registers that contains a non-zero value;  
    decrementing the non-zero value of the counter register by 1 and updating counter registers having a higher sequence than specified by the NEXT register with the corresponding size registers; and  
    repeating the updating the NEXT register, the decrementing the non-zero value and the updating the counter registers until each of the counter registers contains 0.
11. A permutation network for transferring data elements between a vector of logical addresses and multiple banks of random access memory (RAM), the permutation network comprising:  
    a first wiring pattern configured to map a logical space to a cyclic space;  
    a first barrel shifter;  
    a second wiring pattern configured to map the cyclic space to the logical space;  
and  
    a second barrel shifter,  
wherein:

a quantity number of the data elements is less than or equal to a quantity number of the banks of RAM,

the vector of logical addresses is represented by an address offset and a multiplier,

the multiplier is relatively prime to the quantity number of the banks of RAM,

the first barrel shifter rotates the data elements in the cyclic space based on the multiplier, and

the second barrel shifter rotates the data elements in the logical space based on the address offset.

12. The permutation network of claim 11, performing on the data elements addressed by the vector of logical addresses a forward modulo permutation suitable for saving the data elements to the banks of RAM, via:

mapping, by the first wiring pattern, the vector of logical addresses to the cyclic space;

rotating right, by the first barrel shifter, the data elements in the cyclic space based on the multiplier;

mapping, by the second wiring pattern, the data elements after the rotating right by the first barrel shifter to the logical space; and

rotating right, by the second barrel shifter, the data elements in the logical space based on the address offset.

13. The permutation network of claim 12, further comprising:

a reset mechanism for maintaining output bits of the second barrel shifter to be 0 until the permutation network finishes pipelining operations.

14. The permutation network of claim 11, performing on the data elements addressed by the vector of logical addresses a reverse modulo permutation for reordering the data elements read from the banks of RAM, via:

rotating left, by the second barrel shifter, the data elements in the logical space based on the address offset;

mapping, by the first wiring pattern, the data elements after the rotating left by the second barrel shifter to the cyclic space;

rotating left, by the first barrel shifter, the data elements in the cyclic space based on the multiplier; and

mapping, by the second wiring pattern, the data elements after the rotating left by the first barrel shifter from the cyclic space to the logical space.

15. The permutation network of claim 14, further comprising:

a reset mechanism for maintaining output bits of the second wiring pattern to be 0 until the permutation network finishes pipelining operations.




Figure 1




Figure 2




Figure 3




Figure 4

5/10




Figure 5




Figure 6




Figure 7

8/10




Figure 8

9/10




Figure 9

10/10




Figure 10

## INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 17/31477

A. CLASSIFICATION OF SUBJECT MATTER  
 IPC(8) - G06N 99/00 (2017.01), G06F 5/01 (2017.01)  
 CPC - G06F5/01, G06F2205/00, G06N99/005

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)

See Search History Document

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

See Search History Document

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

See Search History Document

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages                                                                                                                                                                                                                                                                                                                                                                                                                                | Relevant to claim No. |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| Y         | LOZITO et al. "FPGA Implementations of Feed Forward Neural Network by using Floating Point Hardware Accelerators." In: THEORETICAL AND APPLIED ELECTRICAL ENGINEERING, VOLUME: 12, NUMBER: 1, March 2014 [online] [retrieved on 20 September 2017 (20.09.2017)] Retrieved from the Internet <URL: <a href="http://advances.utc.sk/index.php/AEEE/article/view/831">http://advances.utc.sk/index.php/AEEE/article/view/831</a> >, entire document, especially Abstract; Figure 1; page 1, col 1-2; page 2, col 1-2 | 1-6                   |
| Y         | US 6,571,268 B1 (Giacalone et al.), 27 May 2003 (27.05.2003), entire document, especially Abstract; col 3, ln 25-35; col 5, ln 25-52; col 8, ln 25-52; col 21, ln 45-67                                                                                                                                                                                                                                                                                                                                           | 1-6                   |
| Y         | US 2014/0136583 A1 (Hyde et al.), 15 May 2014 (15.05.2014), Para [0055]-[0058], [0085]-[0089], [0147]-[0148]                                                                                                                                                                                                                                                                                                                                                                                                      | 2-3                   |
| A         | US 2015/0199963 A1 (GOOGLE INC.), 16 July 2015 (16.07.2015), entire document                                                                                                                                                                                                                                                                                                                                                                                                                                      | 1-6                   |
| A         | US 2014/0344194 A1 (The Trustees of Princeton University), 20 November 2014 (20.11.2014), entire document                                                                                                                                                                                                                                                                                                                                                                                                         | 1-6                   |
| P,A       | US 2016/0379137 A1 (Microsoft Technology Licensing), 29 December 2016 (29.12.2016), entire document                                                                                                                                                                                                                                                                                                                                                                                                               | 1-6                   |

☐ Further documents are listed in the continuation of Box C.

☐ See patent family annex.

\* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

20 September 2017

Date of mailing of the international search report

02 OCT 2017

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents  
 P.O. Box 1450, Alexandria, Virginia 22313-1450  
 Facsimile No. 571-273-8300

Authorized officer:

Lee W. Young

PCT Helpdesk: 571-272-4300  
 PCT OSP: 571-272-7774

# INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 17/31477

## Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:  
because they relate to subject matter not required to be searched by this Authority, namely:
  
2. ☐ Claims Nos.:  
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
  
3. ☐ Claims Nos.:  
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

## Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

This application contains the following inventions or groups of inventions which are not so linked as to form a single general inventive concept under PCT Rule 13.1. In order for all inventions to be examined, the appropriate additional examination fees must be paid.

Group I: Claims 1-6 are directed to an execution unit of hardware acceleration for machine learning.

Group II: Claims 7-10 are directed to an indexing module for looping over a plurality of indexes each corresponding to one of multiple dimensions of a tensor.

Group III: Claims 11-15 are directed to a permutation network for transferring data elements between a vector of logical addresses and multiple banks of random access memory.

--- ( See Continuation in Supplemental Box ) ---

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying additional fees, this Authority did not invite payment of additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☒ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:  
1-6

### Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☐ No protest accompanied the payment of additional search fees.

## INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 17/31477

Continuation of:

Box III. Observations where unity of invention is lacking

The inventions listed as Groups I-III do not relate to a single inventive concept under PCT Rule 13.1 because under PCT Rule 13.2 they lack the same or corresponding technical features for the following reasons:

### Special Technical Features

Group I includes the special technical feature of a clock input receiving the clock; a first operand input receiving a first operand; a second operand input receiving a second operand; an accumulation register storing a current accumulation value; and an output, the accumulation register is updated, within the clock cycle, with a new accumulation value based on the opcode and one or more of the first operand, the second operand and the current accumulation value, and the output comprises saturated lower bits of the new value that is not required in any other groups.

Group II includes the special technical feature of a counter unit comprising: a plurality of counter registers each storing a respective one of the plurality of indexes; a plurality of size registers each corresponding to one of the counter registers; an address unit, the counter unit updates the plurality of indexes stored in the counter registers by repeating a method to loop over the plurality of indexes, and the address unit generates, for each of the indexes stored in the counter registers, an address that is a polynomial of the respective index that is not required in any other groups.

Group III includes the special technical feature of a first wiring pattern configured to map a logical space to a cyclic space; a first barrel shifter; a second wiring pattern configured to map the cyclic space to the logical space; a second barrel shifter, wherein: a quantity number of the data elements is less than or equal to a quantity number of the banks of RAM, the vector of logical addresses is represented by an address offset and a multiplier, the multiplier is relatively prime to the quantity number of the banks of RAM, the first barrel shifter rotates the data elements in the cyclic space based on the multiplier, and the second barrel shifter rotates the data elements in the logical space based on the address off set that is not required in any other groups.

### Common Technical Features

No common technical features between Groups I-III that would otherwise unify the groups.

Therefore, Groups I, II and III lack unity under PCT Rule 13.