

RECURRENT NEURAL NETWORK CELL ACTIVATION TO PERFORM A PLURALITY OF OPERATIONS IN A SINGLE INVOCATION

BACKGROUND

[0001] One or more aspects relate, in general, to facilitating processing within a computing environment, and in particular, to improving such processing.

[0002] In order to enhance processing in computing environments that are data and/or computational-intensive, co-processors are utilized, such as artificial intelligence accelerators (also referred to as neural network processors or neural network accelerators). Such accelerators provide a great deal of compute power used in performing, for instance, involved computations, such as computations on matrices or tensors.

[0003] Tensor computations, as an example, are used in complex processing, including deep learning, which is a subset of machine learning. Deep learning or machine learning, an aspect of artificial intelligence, is used in various technologies, including but not limited to, engineering, manufacturing, medical technologies, automotive technologies, computer processing, etc.

[0004] Tensors and tensor computations enable large amounts of data and/or detailed data to be input to deep learning processing. However, an accelerator is limited by data bandwidth to/from the accelerator. Previously, in an effort to address this limitation, data locality and data re-use at the accelerator are employed. Advancements in the use of tensors and/or processing using such tensors will improve technologies that use machine learning, including computer processing.

SUMMARY

[0005] Shortcomings of the prior art are overcome, and additional advantages are provided through the provision of a computer program product for facilitating processing within a computing environment. The computer program product includes one or more computer readable storage media and program instructions collectively stored on the one or more computer readable storage media to perform a method. The method includes executing an instruction to perform a recurrent neural network cell activation. The executing includes performing a plurality of operations of the recurrent neural network cell

activation to provide a result of the recurrent neural network cell activation. The plurality of operations is performed in a single invocation of the instruction.

[0006] Using a single invocation of the instruction to perform a plurality of operations reduces complexity, reduces use of system resources and increases system performance.

[0007] In one example, the plurality of operations includes one or more sigmoid functions and one or more tangent functions. In one example, the plurality of operations includes tensor element-wise add and tensor element-wise multiplication operations.

[0008] As an example, the plurality of operations includes one or more sigmoid functions, one or more tangent functions, one or more tensor element-wise add operations and one or more tensor element-wise multiplication operations.

[0009] In one example, one or more inputs to the instruction include one or more concatenated tensors. A concatenated tensor may be directly used by an instruction executing on, e.g., an accelerator performing a cell activation of a recurrent neural network. The concatenated tensor may be accessed in one operation saving processing time and increasing processing speed. Further, there are fewer tensor pointers to be managed and there is a reduction in the copying or reorganizing of tensor data between invocations of the accelerator, improving processing speed.

[0010] In one example, the result is an output tensor, and the output tensor is an input to another invocation of the instruction, as an example.

[0011] As examples, the recurrent neural network cell activation includes a long short-term memory cell activation or a gated recurrent unit cell activation.

[0012] In one example, the performing the plurality of operations of the recurrent neural network cell activation is performed by an accelerator and produces intermediate computation data. The intermediate computation data is stored in the accelerator, as an example.

[0013] In one example, the performing the plurality of operations includes performing the plurality of operations on spatially close input data.

[0014] Computer-implemented methods and systems relating to one or more aspects are also described and claimed herein. Further, services relating to one or more aspects are also described and may be claimed herein.

[0015] Additional features and advantages are realized through the techniques described herein. Other embodiments and aspects are described in detail herein and are considered a part of the claimed aspects.

BRIEF DESCRIPTION OF THE DRAWINGS

[0016] One or more aspects are particularly pointed out and distinctly claimed as examples in the claims at the conclusion of the specification. The foregoing and objects, features, and advantages of one or more aspects are apparent from the following detailed description taken in conjunction with the accompanying drawings in which:

FIG. 1A depicts one example of a computing environment to incorporate and use one or more aspects of the present invention;

FIG. 1B depicts further details of a processor of FIG. 1A, in accordance with one or more aspects of the present invention;

FIG. 2A depicts one example of a result tensor, in accordance with one or more aspects of the present invention;

FIG. 2B depicts one example of multiplying concatenated weights by an input feature to provide an intermediate result used in accordance with one or more aspects of the present invention;

FIG. 2C depicts one example of biases added to the intermediate result of FIG. 2B to provide the result tensor of FIG. 2A, in accordance with one or more aspects of the present invention;

FIG. 2D depicts one example of a concatenated output tensor, in accordance with one or more aspects of the present invention;

FIG. 3A depicts one example of a 2D-tensor, in accordance with one or more aspects of the present invention;

FIGS. 3B-3C depict one example of processing used in creating tensors of select dimensions, in accordance with one or more aspects of the present invention;

FIG. 4A depicts one example of a long short-term memory cell activation, in accordance with one or more aspects of the present invention;

FIG. 4B depicts one example of a gated recurrent unit cell activation, in accordance with one or more aspects of the present invention;

FIGS. 5A-5B depict one example of a long short-term memory cell activation using chaining, in accordance with one or more aspects of the present invention;

FIG. 6A depicts one example of a format of a Neural Network Processing Assist instruction, in accordance with one or more aspects of the present invention;

FIG. 6B depicts one example of a general register used by the Neural Network Processing Assist instruction, in accordance with one or more aspects of the present invention;

FIG. 6C depicts examples of function codes supported by the Neural Network Processing Assist instruction, in accordance with one or more aspects of the present invention;

FIG. 6D depicts one example of another general register used by the Neural Network Processing Assist instruction, in accordance with one or more aspects of the present invention;

FIG. 6E depicts one example of a parameter block used by a query function of the Neural Network Processing Assist instruction, in accordance with one or more aspects of the present invention;

FIG. 6F depicts one example of a parameter block used by one or more non-query functions of the Neural Network Processing Assist instruction, in accordance with one or more aspects of the present invention;

FIG. 6G depicts one example of a tensor descriptor used by the Neural Network Processing Assist instruction, in accordance with one or more aspects of the present invention;

FIG. 7 depicts one example of a format of a Neural Network Processing (NNP)-data-type-1 data type, in accordance with one or more aspects of the present invention;

FIGS. 8A-8C depict examples of an input data layout used by the Neural Network Processing Assist instruction, in accordance with one or more aspects of the present invention;

FIGS. 9A-9C depict example output corresponding to the input data layout of FIGS. 8A-8C, in accordance with one or more aspects of the present invention;

FIGS. 10A-10B depict one example of facilitating processing within a computing environment, in accordance with one or more aspects of the present invention;

FIG. 11A depicts another example of a computing environment to incorporate and use one or more aspects of the present invention;

FIG. 11B depicts one example of further details of a memory of FIG. 11A, in accordance with one or more aspects of the present invention;

FIG. 11C depicts another example of further details of a memory of FIG. 11A, in accordance with one or more aspects of the present invention;

FIG. 12A depicts yet another example of a computing environment to incorporate and use one or more aspects of the present invention;

FIG. 12B depicts further details of the memory of FIG. 12A, in accordance with one or more aspects of the present invention;

FIG. 13 depicts one embodiment of a cloud computing environment, in accordance with one or more aspects of the present invention; and

FIG. 14 depicts one example of abstraction model layers, in accordance with one or more aspects of the present invention.

DETAILED DESCRIPTION

[0017] In accordance with one or more aspects of the present invention, a capability is provided to create tensors of a selected data layout format for use in recurrent neural networks, such as recurrent neural networks on a long short-term memory (LSTM) architecture and/or a gated recurrent unit (GRU) architecture. As an example, the selected data layout format includes concatenated input and/or output formats used in, for instance, long short-term memory cell activations and/or gated recurrent unit cell activations.

[0018] Long short-term memory is an artificial recurrent neural network architecture that typically includes, e.g., a cell, which remembers state, and a plurality of gates that control the flow of information into and out of the cell. The gates include, for instance, an input gate, an output gate and a forget gate. A gated recurrent unit is another recurrent neural network architecture. It is similar to the long short-term memory architecture but may have fewer parameters and does not have an output gate. Each network uses timesteps, in which for each timestep, operations are performed on an input producing an output. The output of one timestep may be an input to a next timestep. For each timestep, a number of activations (e.g., sigmoid, tanh) and other operations (e.g., addition, multiplication) are applied to a hidden state (H), input and a cell (c) state. While each of these small steps (e.g., activations, operations) may be efficiently executed locally on a processor, calling an accelerator for each of these steps may be detrimental to the overall performance of the recurrent neural network and/or system due to, e.g., the start up time of the accelerator. Thus, in accordance with one or more aspects of the present invention, the individual activations and operations (e.g., for a timestep) are combined and performed as part of a single invocation of an instruction. This significantly increases processing speed and provides efficiencies since, e.g., there is only one invocation; intermediate computation data can be stored in the accelerator instead of being written back to memory; the SIMD (single instruction, multiple data) width and the pipelined nature of the accelerator can be used to do more computation in parallel with less cycle per compute; and higher precision can be used for intermediate results, leading to enhanced accuracy and higher stability for LSTM/GRU operations.

[0019] Further, in one or more aspects, the single instruction uses the selected data layout formats, providing spatially close input and/or output data, reducing address translation requests and improving processing speed. The selected data layout formats provide efficiencies in which, for instance, operations, such as a cell activations of a recurrent neural network, are able to be chained without requiring the general-purpose processor to inspect/rearrange the data for each timestep of the cell activation.

[0020] One example of a selected data layout format is a concatenated input format, in accordance with one or more aspects of the present invention. To provide such a format, in one example, weight tensors used by, e.g., recurrent neural network cells are transformed to reformatted weight tensors of a select dimension (e.g., 2D-reformatted tensors), which are, e.g., concatenated in a linear way to form a larger concatenated tensor. This enables, for instance, activations and other operations of a cell activation performed on resulting concatenated tensors to be executed in one single instruction invocation executed on, e.g., the accelerator. The resulting concatenated tensor is a selected input format used directly by the instruction on, e.g., the accelerator, which is executing a cell activation on the recurrent neural network.

[0021] A further example of a selected data layout format, in accordance with one or more aspects of the present invention, is a concatenated output format, such as a 2D-output tensor. The format is chosen such that, for instance, for each timestep, an output tensor can be accessed as a memory-contiguous sub-tensor that can be fed to the next timestep of, e.g., a computation. The timesteps remain adjacent in memory to return the final result consisting of the timesteps as one memory adjacent tensor.

[0022] One or more aspects of the present invention include reformatting tensors to provide reformatted tensors (which may also be referred to as sub-tensors) of a select dimension (e.g., 2D-tensors) that represent an original tensor. This optimizes processing, including, but not limited to, memory address calculation, load/store operations and/or prefetching. As an example, a tensor is reformatted such that the reformatted tensor starts on a boundary of a memory unit (e.g., memory page) and information of the original tensor is rearranged to fit within the reformatted-tensor (a.k.a., tiles) of a select dimension (e.g., 2D). The reformatted tensors have easily computable addresses and may be block loaded and/or stored (e.g., loaded/stored in one operation) providing efficiencies in using the reformatted tensors.

[0023] One example of an instruction to use concatenated input/output data formats provided in accordance with one or more aspects of the present invention and/or to combine multiple operations (e.g., activations and/or other operations) of a recurrent neural network cell activation is a Neural Network Processing Assist instruction, which is a single instruction (e.g., a single architected hardware machine instruction at the hardware/software interface) configured to perform multiple functions. Each of the functions is configured as part of the single instruction (e.g., the single architected instruction), reducing use of system resources and complexity, and improving system performance.

[0024] The instruction may be part of a general-purpose processor instruction set architecture (ISA), which is dispatched by a program on a processor, such as a general-purpose processor. It may be executed by the general-purpose processor and/or one or more functions of the instruction may be executed by a special-purpose processor, such as a co-processor or accelerator configured for certain functions, that is coupled to or part of the general-purpose processor. Other variations are also possible.

[0025] One embodiment of a computing environment to incorporate and use one or more aspects of the present invention is described with reference to FIG. 1A. As an example, the computing environment is based on the z/Architecture[®] instruction set architecture, offered by International Business Machines Corporation, Armonk, New York. One embodiment of the z/Architecture instruction set architecture is described in a publication entitled, "z/Architecture Principles of Operation," IBM Publication No. SA22-7832-12, Thirteenth Edition, September 2019, which is hereby incorporated herein by reference in its entirety. The z/Architecture instruction set architecture, however, is only one example architecture; other architectures and/or other types of computing environments of International Business Machines Corporation and/or of other entities may include and/or use one or more aspects of the present invention. z/Architecture and IBM are trademarks or registered trademarks of International Business Machines Corporation in at least one jurisdiction.

[0026] Referring to FIG. 1A, a computing environment 100 includes, for instance, a computer system 102 shown, e.g., in the form of a general-purpose computing device. Computer system 102 may include, but is not limited to, one or more general-purpose processors or processing units 104 (e.g., central processing units (CPUs)), at least one special-purpose processor, such as a neural network processor 105, a memory 106 (a.k.a.,

system memory, main memory, main storage, central storage or storage, as examples), and one or more input/output (I/O) interfaces 108, coupled to one another via one or more buses and/or other connections. For instance, processors 104, 105 and memory 106 are coupled to I/O interfaces 108 via one or more buses 110, and processors 104, 105 are coupled to one another via one or more buses 111.

[0027] Bus 111 is, for instance, a memory or cache coherence bus, and bus 110 represents, e.g., one or more of any of several types of bus structures, including a memory bus or memory controller, a peripheral bus, an accelerated graphics port, and a processor or local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include the Industry Standard Architecture (ISA), the Micro Channel Architecture (MCA), the Enhanced ISA (EISA), the Video Electronics Standards Association (VESA) local bus, and the Peripheral Component Interconnect (PCI).

[0028] As examples, one or more special-purpose processors (e.g., neural network processors) may be separate from but coupled to one or more general-purpose processors and/or may be embedded within one or more general-purpose processors. Many variations are possible.

[0029] Memory 106 may include, for instance, a cache 112, such as a shared cache, which may be coupled to local caches 114 of processors 104 and/or to neural network processor 105 via, e.g., one or more buses 111. Further, memory 106 may include one or more programs or applications 116 and at least one operating system 118. An example operating system includes a z/OS[®] operating system, offered by International Business Machines Corporation, Armonk, New York. z/OS is a trademark or registered trademark of International Business Machines Corporation in at least one jurisdiction. Other operating systems offered by International Business Machines Corporation and/or other entities may also be used. Memory 106 may also include one or more computer readable program instructions 120, which may be configured to carry out functions of embodiments of aspects of the invention.

[0030] Moreover, in one or more embodiments, memory 106 includes processor firmware 122. Processor firmware includes, e.g., the microcode or millicode of a processor. It includes, for instance, the hardware-level instructions and/or data structures used in implementation of higher level machine code. In one embodiment, it includes, for instance, proprietary code that is typically delivered as microcode or millicode that includes trusted

software, microcode or millicode specific to the underlying hardware and controls operating system access to the system hardware.

[0031] Computer system 102 may communicate via, e.g., I/O interfaces 108 with one or more external devices 130, such as a user terminal, a tape drive, a pointing device, a display, and one or more data storage devices 134, etc. A data storage device 134 may store one or more programs 136, one or more computer readable program instructions 138, and/or data, etc. The computer readable program instructions may be configured to carry out functions of embodiments of aspects of the invention.

[0032] Computer system 102 may also communicate via, e.g., I/O interfaces 108 with network interface 132, which enables computer system 102 to communicate with one or more networks, such as a local area network (LAN), a general wide area network (WAN), and/or a public network (e.g., the Internet), providing communication with other computing devices or systems.

[0033] Computer system 102 may include and/or be coupled to removable/non-removable, volatile/non-volatile computer system storage media. For example, it may include and/or be coupled to a non-removable, non-volatile magnetic media (typically called a "hard drive"), a magnetic disk drive for reading from and writing to a removable, non-volatile magnetic disk (e.g., a "floppy disk"), and/or an optical disk drive for reading from or writing to a removable, non-volatile optical disk, such as a CD-ROM, DVD-ROM or other optical media. It should be understood that other hardware and/or software components could be used in conjunction with computer system 102. Examples, include, but are not limited to: microcode or millicode, device drivers, redundant processing units, external disk drive arrays, RAID systems, tape drives, and data archival storage systems, etc.

[0034] Computer system 102 may be operational with numerous other general-purpose or special-purpose computing system environments or configurations. Examples of well-known computing systems, environments, and/or configurations that may be suitable for use with computer system 102 include, but are not limited to, personal computer (PC) systems, server computer systems, thin clients, thick clients, handheld or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputer systems, mainframe computer systems,

and distributed cloud computing environments that include any of the above systems or devices, and the like.

[0035] In one example, a processor (e.g., processor 104 and/or processor 105) includes a plurality of functional components (or a subset thereof) used to execute instructions. As depicted in FIG. 1B, these functional components include, for instance, an instruction fetch component 150 to fetch instructions to be executed; an instruction decode unit 152 to decode the fetched instructions and to obtain operands of the decoded instructions; one or more instruction execute components 154 to execute the decoded instructions; a memory access component 156 to access memory for instruction execution, if necessary; and a write back component 158 to provide the results of the executed instructions. One or more of the components may access and/or use one or more registers 160 in instruction processing. Further, one or more of the components may, in accordance with one or more aspects of the present invention, include at least a portion of or have access to one or more other components used in providing concatenated input and/or output data formats, in combining multiple operations of a cell activation function, in tensor processing (including, but not limited to, creation and/or use of reformatted tensors), and/or in neural network processing assist processing of, e.g., a Neural Network Processing Assist instruction (or other processing that may use one or more aspects of the present invention), as described herein. The one or more other components may include, for instance, one or more combining/concatenation components 170, a tensor component 171 and/or a neural network processing assist component 172 (and/or one or more other components).

[0036] In accordance with one or more aspects of the present invention, processing within a computing environment is facilitated by providing improved data formats for use by a processor, such as a special-purpose processor (e.g., neural network processor 105). For instance, a concatenated input data format layout is provided, in which a plurality of tensors of a select dimension, such as a plurality of 2D-tensors, are concatenated to create a concatenated tensor. Similarly, in one example a concatenated output data format is provided, in which multiple output tensors are concatenated. Further details regarding concatenated input/output data layout formats are described with reference to FIGS. 2A-2D. In the figures, t refers to timestep, N_{mb} refers to batch size, s refers to size and l is length of an input feature.

[0037] Referring to FIG. 2A, one example of a concatenated tensor input (also referred to herein as a result tensor) 200 is depicted. In this example, multiple 2D-tensors 202, having a size s , are concatenated (e.g., linearly) to create a larger concatenated tensor 200 having a size $4s$. In one example, concatenated tensor 200 includes a plurality of (e.g., four) concatenated weight tensors (e.g., W_f , W_i , W_c , W_o) that have been multiplied by a feature input X (X_i). For instance, as shown in FIG. 2B, a feature input X (210) is multiplied 212 by a concatenated tensor of weights 214 providing an intermediate result (e.g., result tensor), which, referring to FIG. 2C, is added to a tensor of biases 220 to produce a result, which is, for instance, concatenated input tensor 200. In a neural network, a feature is a representation of what is to be observed (e.g., a next word in a sentence, a particular picture, etc.), a weight is a learnable parameter, and a bias is an offset, as examples. In one example, the multiplication and addition are performed as a Neural Network Processing Assist-Matrix Multiplication Operation-Broadcast²³ (e.g., NNPA-MATMUL-OP-BCAST²³) operation of a Neural Network Processing Assist instruction, an example of which is described below.

[0038] In one example, each weight tensor of FIG. 2B is a reformatted 2D-tensor provided to facilitate processing of tensors. The weight tensors are independently transformed to 2D-reformatted tensors, which are concatenated to provide a large tensor. The resulting tensor is an input format used, in accordance with an aspect of the present invention, directly by an instruction (e.g., the Neural Network Processing Assist instruction) on an accelerator (e.g., processor 105) executing a cell activation of a recurrent neural network. It allows matrix multiplications of the cell activation to be executed across timesteps in one single instruction executed on the accelerator. Each reformatted 2D-tensor starts, in accordance with an aspect of the present invention, on a boundary of a memory unit (e.g., memory page boundary) and information of the original tensor is rearranged in the reformatted tensor. The dimensions of the tensor in the dimensions of the reformatted tensor are rounded up to the next full tile in that dimension (e.g., padding is provided to create fixed sized tensors, e.g., 2D-tensors). For instance, row padding 216 and/or page padding 218 are provided, as described herein, to create fixed size tensors. This allows access to each tensor on a memory unit boundary (e.g., page boundary) and facilitates computation of the addresses of any 2D-tensor. By providing alignment on a memory unit boundary, address translation requests are reduced, and the data transfer rate is increased. Further, in one example, each 2D-tensor may be loaded via a direct memory access (DMA)-

like operation accessing one memory unit (e.g., page) in the accelerator memory at once. This significantly increases the bandwidth.

[0039] Similarly, in one example, bias tensor 220 is a concatenated bias tensor that includes a plurality of bias tensors 222. Each bias tensor is of a select fixed size, and therefore, row padding 224 and/or page padding 226 are provided, as described herein.

[0040] In addition to a concatenated input tensor, in accordance with one or more aspects of the present invention, concatenated output tensors are provided, an example of which is depicted in FIG. 2D. As shown in FIG. 2D, a concatenated output tensor 250 includes, for each input, a hidden state (h) tensor 260 concatenated to an internal cell state (c) tensor 270. In one example, each tensor 260, 270 is a reformatted tensor of a select dimension (e.g., 2D) and of a select size. To provide a tensor of the select size, row padding 280 and/or page padding 282 are provided, as described herein. The concatenated output tensor is, e.g., a concatenated 2D-reformatted output tensor. A concatenated output tensor is accessible as a memory contiguous sub-tensor that can be fed to the next timestep of the computation, while all timesteps, as an example, remain adjacent in memory to return the final results consisting of all timesteps as one memory adjacent tensor. As with the input tensors, the dimensions of the tensor in the dimensions of the reformatted tensor are rounded up to the next full tile in that dimension (e.g., padding is provided to create fixed sized tensors, e.g., 2D-tensors).

[0041] Further details relating to a 2D-tensor, in accordance with one or more aspects of the present invention, are described with reference to FIG. 3A. As shown, a 2D-tensor 300 starts on a memory boundary and uses a plurality of memory units, such as a plurality of 4K pages (e.g., pages 0-11 numbered in the 2D-tensor). Each page includes a pre-selected number of rows (e.g., 32) 302, and each row includes a preselected number of elements, e.g., 64 elements. If a row has less data than the pre-selected number of elements, it is padded 304 with, e.g., a pre-specified value, such as zeros or spaces, etc. Further, if there is not enough data to provide the pre-selected number of rows, additional padding 306 (e.g., unpredictable data, existing data, any value, etc.) is provided to add additional padded rows, as shown in FIG. 3A.

[0042] The architected data format of the 2D-tensor provides easily computable addresses and memory-wise adjacent tensor units, which allows a reduction in overhead of multiple and complicated address calculations. This assists hardware supported block-

load/store operations and prefetching engines, significantly increasing the effective data bandwidth (e.g., 2x-5x) to the accelerator (e.g., neural network processor 105).

[0043] Further details related to creating a 2D-tensor, in accordance with an aspect of the present invention, are described with reference to FIGS. 3B-3C. In one example, the processing creates tensors (e.g., 2D, 3D, 4D and/or other tensors) based on a 4D-feature data layout, described herein. As an example, this processing is performed by a processor, such as general-purpose processor 104. This processing is capable of producing 2D, 3D or 4D-tensors, as examples, but not limited to such examples.

[0044] Referring to FIG. 3B, in one example, an $e2_limit$ is set (352) equal to a $ceil(E2/32) * 32$ indicating the 2D-tensor being created has, e.g., 32 rows and E2 refers to a dimension-2-index-size. Further, an $e1_limit$ is set (354) equal to $ceil(E1/64) * 64$ indicating the 2D-tensor being created has, e.g., 64 elements per row and E1 refers to a dimension-1-index-size. An index $e4x$ is initialized to zero 356.

[0045] Subsequent to the initializations, a determination is made as to whether $e4x$ is less than E4 (358), E4 refers to a dimension-4-index-size. If $e4x$ is not less than E4, then processing ends 360; otherwise, processing continues with initializing an index $e3x$ to zero 362. A determination is made as to whether $e3x$ is less than E3 (364), E3 refers to a dimension-3-index-size. If $e3x$ is not less than E3, then the processing iterates in which $e4x$ is incremented by, e.g., 1 (366), and processing continues to 358. However, if $e3x$ is less than E3, then an index $e2x$ is initialized to zero 368. A determination is made as to whether $e2x$ is less than $e2_limit$ (370). If $e2x$ is not less than $e2_limit$, then the processing iterates in which $e3x$ is incremented by, e.g., 1 (372), and processing continues to 364. If $e2x$ is less than $e2_limit$, then an index $e1x$ is initialized to zero 374.

[0046] Referring to FIG. 3C, a determination is made as to whether $e1x$ is less than $e1_limit$ (376). If $e1x$ is not less than $e1_limit$, then the processing iterates in which $e2x$ is incremented by, e.g., 1 ($e2x = e2x + 1$) 378, and processing continues to 370 (FIG. 3B). If $e1x$ is less than $e1_limit$, then arr_pos (e.g., position in a row) is set equal to $(E3 * e2_limit * e1_limit * e4x) + (e2_limit * e3x * 64) + (e2x * 64) + ([e1x/64] * e2_limit * E3 * 64) + (e1x \bmod 64)$, where $[]$ is a ceil function 382.

[0047] A determination is made as to whether $e2x$ is greater than or equal to E2 (384). If $e2x$ is less than E2, a further determination is made as to whether $e1x$ is greater than or

equal to E1 (386). If e1x is less than E1, then value is set equal to input_array [e4x] [e3x] [e2x] [e1x] 388; and if e1x is greater than or equal to E1, then value = E1pad 390 (the row is padded). Further, if e2x is greater than or equal to E2 (384), then value = E2pad 392 (additional rows that are padded are added to the 2D-tensor). Subsequent to setting value, OutputTensor [arr_pos] = value. Further, index e1x is incremented by, e.g., 1 (e1x = e1x + 1) 394, and processing continues to 376.

[0048] As a further example, tensors may be created based on a 4D_kernel layout, described herein. To create 2D, 3D, 4D and/or other tensors, the processing of FIGS. 3B-3C is used, except 382 is replaced with kern_pos = ($\lfloor e1x/64 \rfloor * E4 * E3 * e2_limit * 64$) + ($e4x * E3 * e2_limit * 64$) + ($e3x * e2_limit * 64$) + ($e2x * 64$) + ($e1x \bmod 64$); and 394 is replaced with OutputTensor [kern_pos] = value.

[0049] The created tensors (e.g., reformatted tensors created from a reformatting of an original tensor) may be used by one or more instructions. For instance, address information (e.g., beginning of a 4D-tensor or a 2D-tensor, as examples), dimensions of the tensor, etc., are forwarded from the general-purpose processor to a special-purpose processor (e.g., neural network 105) for use in loading/storing the data in a correct format (e.g., in correct locations in pages of memory) and for using the data (e.g., in tensor computations). In other embodiments, a general-purpose processor uses the created reformatted tensor(s). Other variations are possible.

[0050] In accordance with one or more aspects, multiple reformatted tensors are concatenated to provide concatenated input and/or output tensors. In one example, one or more concatenated input tensors are input to a recurrent neural network cell activation, such as a long short-term memory cell activation or a gated recurrent unit cell activation, which produces one or more concatenated output tensors. Further details regarding example cell activations are described with reference to FIGS. 4A-4B.

[0051] As an example, referring to FIG. 4A, a first input tensor 400a (e.g., input tensor 1) and a second input tensor 400b (e.g., input tensor 2) are input to a long short-term memory cell activation 401. For instance, first input tensor 400a and second input tensor 400b are concatenated tensors (e.g., result tensors), each including, e.g., a concatenation of, e.g., four individual tensors 400a1-400a4, and 400b1-400b4, respectively, each of which is input to an add operation of long short-term memory cell activation 401. As an example, input tensors 400a1, 400b1 are input to add operation 402a; input tensors 400a2, 400b2 are

input to add operation 402b; input tensors 400a3, 400b3 are input to add operation 402c; and input tensors 400a4, 400b4 are input to add operation 402d. Each add operation is, for instance, equivalent to an NNPA-ADD operation, an example of which is described herein. An output of add operation 402a is input to a sigmoid activation 404a; an output of add operation 402b is input to a sigmoid activation 404b; an output of add operation 402c is input to a tangent activation 406; and an output of an add operation 402d is input to a sigmoid activation 404c. Sigmoid activations 404a, 404b and 404c and tangent activation 406 are equivalent to, e.g., an NNPA-SIGMOID function and an NNPA-TANH function, respectively, examples of which are described herein. Outputs of sigmoid activation 404b and tangent activation 406 are input to a multiplication operation 408, which is equivalent to, e.g., an NNPA-MUL function, an example of which is described herein.

[0052] Outputs of sigmoid activation 404a and multiplication operation 408 are input to a combined operation 410, along with a third input tensor 400c (e.g., input tensor 3). In this example, input tensor 400c is not a concatenated tensor and is an output from a previous timestep. For instance, input tensor 400c is a cell state portion of a concatenated output tensor. Combined operation 410 is, for instance, a fused multiply add (FMA) operation, which is equivalent to, e.g., a NNPA-BATCHNORM function, an example of which is described herein. (In other examples, individual operations rather than a combined operation may be used.) In operation 410, the output from sigmoid activation 404a and input tensor 400c are multiplied to provide an intermediate result. The intermediate result is added to the output of multiplication operation 408 to provide another intermediate result. The other intermediate result (e.g., the result of combined operation 410) is input to tangent activation 412, which is equivalent to, e.g., a NNPA-TANH function, an example of which is described herein. An output of tangent function 412 and an output of sigmoid function 404c are input to a multiplication operation 414, which is equivalent to, e.g., an NNPA-MUL function, an example of which is described herein. An output of NNPA-MUL 414 is an output tensor 420a (e.g., output tensor 1). Further, in one example, an output of combined operation 410 is an output tensor 420b (e.g., output tensor 2). As an example, output tensors 420a and 420b are concatenated output tensors, such as described with reference to FIG. 2D.

[0053] Referring to FIG. 4B, an example of a gated recurrent unit cell activation is described. In one example, a first input tensor 450a (e.g., input tensor 1) and a second input tensor 450b (e.g., input tensor 2) are input to a gated recurrent unit cell activation 451. For

instance, first input tensor 450a and second input tensor 450b are concatenated tensors (e.g., result tensors), each including, e.g., a concatenation of, e.g., three individual tensors 450a1-450a3, and 450b1-450b3, respectively, each of which is input to an operation of gated recurrent unit cell activation 451. As an example, input tensors 450a1, 450b1 are input to add operation 452a; and input tensors 450a2, 450b2 are input to add operation 452b. Each add operation is, for instance, equivalent to an NNPA-ADD operation, an example of which is described herein. An output of add operation 452a is input to a sigmoid activation 454a; and an output of add operation 452b is input to a sigmoid activation 454b. Sigmoid activations 454a, and 454b are equivalent to, e.g., an NNPA-SIGMOID function, an example of which is described herein. Outputs of sigmoid activations 454a and 454b are input to a multiplication operation 456a and 456b, respectively, which are equivalent to, e.g., an NNPA-MUL function, an example of which is described herein. Another input to multiplication operation 456a is input tensor 450c. In this example, input tensor 450c is not a concatenated tensor and is an output from a previous timestep. For instance, input tensor 450c is a cell state portion of a concatenated output tensor. Further, another input to multiplication operation 456b is input tensor 450b3.

[0054] In one example, the output of sigmoid function 454a is also input to a subtraction operation 458, along with a numerical value of 1. One example of a subtraction operation is an NNPA-SUB function, an example of which is described herein.

[0055] An output of multiplication operation 456b and input tensor 450a3 are input to an addition operation 460, which is equivalent, e.g., to an NNPA-ADD function, an example of which is described herein. An output of addition operation 460 is input to a tangent activation 462, which is equivalent, e.g., to an NNPA-TANH function, an example of which is described herein. Outputs of subtraction operation 458 and tangent activation 462 are input to a multiplication operation 464, which is equivalent, e.g., to an NNPA-MUL function, an example of which is described herein. An output of multiplication operation 464 and an output of multiplication operation 456a are input to an addition operation 466, which is equivalent to, e.g., an NNPA-ADD function, an example of which is described herein. An output of addition operation 466 is an output tensor 468. As an example, output tensor 468 is a concatenated output tensor, such as described with reference to FIG. 2D.

[0056] As described above, multiple activations (e.g., sigmoid, tangent) and other operations (e.g., addition, subtraction and/or multiplication) are combined and performed as

part of one cell activation, which is executed (e.g., on an accelerator, such as neural network processor 105) based on invocation of a single instruction (e.g., a Neural Network Processing Assist instruction). The single instruction is implemented to combine the individual activations and other operations. This provides a higher accuracy due to the combining of, e.g., multiply and add operations together without a loss of precision on the intermediate results. Further, higher numeric accuracy can be achieved by saving intermediate computations in the accelerator in higher precision. Additionally, in accordance with one or more aspects of the present invention, the activations and other operations of the cell activation are separated from the matrix multiplication used to create the concatenated input tensors, reducing the complexity of single operations and allowing reuse of the basic blocks for other recurrent neural networks. That is, recurrent neural networks (e.g., on a long short-term memory architecture or a gated recurrent unit architecture) rely on several matrix multiplications between an input feature (e.g., X of FIG. 2B) and different weight tensors (e.g., unconcatenated, unreformatted weight tensors of FIG. 2B) followed by several activation functions (e.g., sigmoid, tangent of FIGS. 4A-4B) on the generated intermediate results. Typically, the matrix multiplications and activation functions are performed separately on independent tensor buffers, which results in several separate instructions to compute a recurrent neural network timestep and may require copying/reorganizing data between those individual instructions, significantly reducing performance. Advantages of, e.g., an on-chip accelerator (e.g., neural network processor 105) is significantly reduced if data manipulation on the general-purpose processor between accelerator operations is needed. This is due to lower bandwidth, required serialization and set-up time to start the accelerator. Therefore, in accordance with one or more aspects of the present invention, a data layout format (e.g., a reformatted concatenated tensor) is provided that is used directly by an instruction on the accelerator executing cell activations of the recurrent neural network.

[0057] Further, in accordance with one or more aspects, a data layout format is chosen in which a concatenated output tensor is generated based on computing a cell activation of a timestep that enables accelerator operations to be chained without the need for a general-purpose processor to inspect/rearrange the data. Further, the instruction provides spatially close input and output sources to reduce address translations. By locating the data adjacently in memory, less address translations are needed. This contributes to an overall increase in the speed of processing within an accelerator (e.g., neural network processor 105) and to an increase in higher precision.

[0058] One example of an overall chaining operation is described with reference to FIGS. 5A-5B. In FIGS. 5A-5B, Nmb is batch size, t is timestep, s is size, and l is a length of a feature. In this example, the cell activation to use chaining is a long short-term memory cell activation 500, an example of which is described herein with reference to FIG. 4A. However, in other examples, it may be other cell activations, including, but not limited to, gated recurrent unit cell activations, an example of which is described herein with reference to FIG. 4B, and/or other cell activations.

[0059] Referring to FIG. 5A, outputs of cell activation 500 include a history (h) tensor 502 and a cell state (c) tensor 504, which are used to provide a concatenated output tensor 510. The concatenated output tensor is then input to a next timestep of cell activation 500 (i.e., chaining). For instance, a history tensor 510a of concatenated tensor 510 is input to a matrix multiplication operation 520 and a cell state tensor 510b of concatenated tensor 510 is input to a combined operation 530 (e.g., a fused multiply add operation, such as a NNPA-BATCHNORM). In other examples, individual operations, rather than a combined operation may be used.

[0060] In one example, in matrix multiplication operation 520, history tensor 510a and concatenated weighted matrix 540 are multiplied to provide an intermediate result, which is added to a concatenated bias tensor 550 (FIG. 5B) to provide a concatenated tensor (e.g., input tensor 2), which is input to cell activation 500. Further, in one example, another concatenated tensor (e.g., input tensor 1) is also input to cell activation 500. Input tensor 1, as described herein and further described with reference to FIG. 5B, is created by concatenating a plurality of weight tensors 560 to provide a concatenated weight tensor 562. Concatenated weight tensor 562 is multiplied, using, e.g., a matrix multiplication broadcast operation 564 (e.g., NNPA-MATMUL-OP-BCAST23), by a feature input 566 to provide an intermediate result, which is added to a concatenated bias tensor 570 using, e.g., a matrix multiplication broadcast operation 564 to provide a resulting input tensor 1. Concatenated bias tensor 570 is created from a plurality of bias tensors 572, as described herein.

[0061] Concatenated weight tensor 562, concatenated bias tensor 570 and/or concatenated output tensor 510 (FIG. 5A) are, e.g., reformatted tensors, in accordance with one or more aspects of the present invention. A reformatted tensor, as described herein, starts on a memory boundary (e.g., a page boundary) and includes padding to complete a tensor of a select size. For instance, if a tensor is to include a select number of rows (e.g.,

32 rows) and the reformatted tensor has less rows, then padded rows are added until the tensor includes the selected number of rows. Additionally, and/or alternatively, in one example, each row is to include a select number of elements (e.g., 64 elements) and if a row has less elements than the row can include, padding is added to the row until the row includes the selected number of elements.

[0062] Layers of a concatenated tensor (e.g., individual tensors of the concatenated tensor) are selected as input to the cell activation. For instance, referring to FIG. 5A, individual input tensors of input tensor 1 are selected 525 to be input to particular operations. Other examples are possible.

[0063] In accordance with one or more aspects of the present invention, a single architected instruction is provided that supports a data layout format that enables the creation and/or use of reformatted tensors and/or concatenated tensors, and/or that combines activations and operations in a cell activation performed by a single invocation of the instruction. One example of such instruction is a Neural Network Processing Assist instruction. In one example, the instruction is initiated on a general-purpose processor (e.g., processor 104) and a function specified by the instruction is either executed on the general-purpose processor and/or a special-purpose processor (e.g., neural network processor 105) depending on the function. For instance, in one example, a query function of the Neural Network Processing Assist instruction is performed on the general-purpose processor and non-query functions are performed on the special-purpose processor. However, other variations are possible. If the function is to be performed on the special-purpose processor (e.g., it is a non-query function, or in another example, one or more selected functions), then information is provided, e.g., by the general-purpose processor to the special-purpose processor for use in executing the function, such as memory address information relating to tensor data to be used in neural network computations. The special-purpose processor obtains the information and performs the function. After execution of the function is complete, processing returns to the general-purpose processor, which completes the instruction. In other examples, the instruction is initiated, executed and completed on one or more general-purpose processors or one or more special-purpose processors. Other variations are possible.

[0064] In one example, referring to FIG. 6A, a Neural Network Processing Assist instruction 600 has an RRE format that denotes a register and register operation with an

extended operation code (opcode). As shown in FIG. 6A, in one example, Neural Network Processing Assist instruction 600 includes an operation code (opcode) field 602 (e.g., bits 0-15) indicating a neural network processing assist operation. In one example, bits 16-31 of the instruction are reserved and are to contain zeros. In the description herein of the instruction and/or functions of the instructions, specific locations, specific fields and/or specific sizes of the fields are indicated (e.g., specific bytes and/or bits). However, other locations, fields and/or sizes may be provided. Further, although the setting of a bit to a particular value, e.g., one or zero, may be specified, this is only an example. The bit, if set, may be set to a different value, such as the opposite value or to another value, in other examples. Many variations are possible.

[0065] In one example, the instruction uses a plurality of general registers implicitly specified by the instruction. For instance, Neural Network Processing Assist instruction 600 uses implied registers general register 0 and general register 1, examples of which are described with reference to FIGS. 6B and 6D, respectively.

[0066] Referring to FIG. 6B, in one example, general register 0 includes a function code field, and status fields which may be updated upon completion of the instruction. As an example, general register 0 includes a response code field 610 (e.g., bits 0-15), an exception flags field 612 (e.g., bits 24-31) and a function code field 614 (e.g., bits 56-63). Further, in one example, bits 16-23 and 32-55 of general register 0 are reserved and are to contain zeros. One or more fields are used by a particular function performed by the instruction. Not all fields are used by all of the functions, in one example. Each of the fields is described below:

[0067] Response Code (RC) 610: This field (e.g., bit positions 0-15) contains the response code. When execution of the Neural Network Processing Assist instruction completes with a condition code of, e.g., one, a response code is stored. When an invalid input condition is encountered, a non-zero value is stored to the response code field, which indicates the cause of the invalid input condition recognized during execution and a selected condition code, e.g., 1, is set. The codes stored to the response code field are defined, as follows, in one example:

[0068]	<u>Response Code</u>	<u>Meaning</u>
--------	----------------------	----------------

[0069]	0001	The format of the parameter block, as specified by the parameter block version number, is not supported by the model.
[0070]	0002	The specified function is not defined or installed on the machine.
[0071]	0010	A specified tensor data layout format is not supported.
[0072]	0011	A specified tensor data type is not supported.
[0073]	0012	A specified single tensor dimension is greater than the maximum dimension index size.
[0074]	0013	The size of a specified tensor is greater than the maximum tensor size.
[0075]	0014	The specified tensor address is not aligned on a 4 K-byte boundary.
[0076]	0015	The function-specific-save-area-address is not aligned on a 4 K-byte boundary.
[0077]	F000-FFFF	Function specific response codes. These response codes are defined for certain functions.

[0078] Exception Flags (EF) 612: This field (e.g., bit positions 24-31) includes the exception flags. If an exception condition is detected during execution of the instruction, the corresponding exception flag control (e.g., bit) will be set to, e.g., one; otherwise, the control remains unchanged. The exception flags field is to be initialized to zero prior to the first invocation of the instruction. Reserved flags are unchanged during the execution of the instruction. The flags stored to the exception flags field are defined as follows, in one example:

[0079]	<u>EF (Bit)</u>	<u>Meaning</u>
[0080]	0	Range Violation. This flag is set when a non-numeric value was either detected in an input tensor or stored to the output tensor. This flag is, e.g., only valid when the instruction completes with condition code, e.g., 0.

[0081] 1-7 Reserved.

[0082] Function Code (FC) 614: This field (e.g., bit positions 56-63) includes the function code. Examples of assigned function codes for the Neural Network Processing Assist instruction are depicted in FIG. 6C. All other function codes are unassigned. If an unassigned or uninstalled function code is specified, a response code of, e.g., 0002 hex and a select condition code, e.g., 1, are set. This field is not modified during execution.

[0083] As indicated, in addition to general register 0, the Neural Network Processing Assist instruction also uses general register 1, an example of which is depicted in FIG. 6D. As examples, bits 40-63 in the 24-bit addressing mode, bits 33-63 in the 31-bit addressing mode or bits 0-63 in the 64-bit addressing mode include an address of a parameter block 620. The contents of general register 1 specify, for instance, a logical address of a leftmost byte of the parameter block in storage. The parameter block is to be designated on a doubleword boundary; otherwise, a specification exception is recognized. For all functions, the contents of general register 1 are not modified.

[0084] In the access register mode, access register 1 specifies an address space containing the parameter block, input tensors, output tensors and the function specific save area, as an example.

[0085] In one example, the parameter block may have different formats depending on the function specified by the instruction to be performed. For instance, the query function has a parameter block of one format and other functions of the instruction have a parameter block of another format. In another example, all functions use the same parameter block format. Other variations are also possible.

[0086] One example of a parameter block used by a query function, such as the NNPA-Query Available Functions (QAF) operation, is described with reference to FIG. 6E. As shown, in one example, a NNPA-Query Available Functions parameter block 630 includes, for instance:

[0087] Installed Functions Vector 632: This field (e.g., bytes 0-31) of the parameter block includes the installed functions vector. In one example, bits 0-255 of the installed functions vector correspond to function codes 0-255, respectively, of the Neural Network Processing Assist instruction. When a bit is, e.g., one, the corresponding function is installed; otherwise, the function is not installed.

[0088] Installed Parameter Block Formats Vector 634: This field (e.g., bytes 32-47) of the parameter block includes the installed parameter block formats vector. In one example, bits 0-127 of the installed parameter block formats vector correspond to parameter block formats 0-127 for the non-query functions of the Neural Network Processing Assist instruction. When a bit is, e.g., one, the corresponding parameter block format is installed; otherwise, the parameter block format is not installed.

[0089] Installed Data Types 636: This field (e.g., bytes 48-49) of the parameter block includes the installed data types vector. In one example, bits 0-15 of the installed data types vector correspond to the data types being installed. When a bit is, e.g., one, the corresponding data type is installed; otherwise, the data type is not installed. Example data types include (additional, fewer and/or other data types are possible):

[0090] Bit Data Type

[0091] 0 NNP-data-type-1

[0092] 1-15 Reserved

[0093] Installed Data Layout Formats 638: This field (e.g., bytes 52-55) of the parameter block includes the installed data layout formats vector. In one example, bits 0-31 of the installed data layout formats vector correspond to data layout formats being installed. When a bit is, e.g., one, the corresponding data layout format is installed; otherwise, the data layout format is not installed. Example data layout formats include (additional, fewer and/or other data types are possible):

[0094] Bit Data Layout Format

[0095] 0 4D-feature tensor

[0096] 1 4D-kernel tensor

[0097] 2-31 Reserved

[0098] Maximum Dimension Index Size 640: This field (e.g., bytes 60-63) of the parameter block includes, e.g., a 32-bit unsigned binary integer that specifies a maximum number of elements in a specified dimension index size for any specified tensor. In another

example, the maximum dimension index size specifies a maximum number of bytes in a specified dimension index size for any specified tensor. Other examples are also possible.

[0099] Maximum Tensor Size 642: This field (e.g., bytes 64-71) of the parameter block includes, e.g., a 32-bit unsigned binary integer that specifies a maximum number of bytes in any specified tensor including any pad bytes required by the tensor format. In another example, the maximum tensor size specifies a maximum number of total elements in any specified tensor including any padding required by the tensor format. Other examples are also possible.

[00100] Installed-NNP-Data-Type-1-Conversions Vector 344: This field (e.g., bytes 72-73) of the parameter block includes the installed-NNP-Data-Type-1-conversions vector. In one example, bits 0-15 of the installed-NNP-Data-Type-1-conversions vector correspond to installed data type conversion from/to NNP-data-type-1 format. When a bit is one, the corresponding conversion is installed; otherwise, the conversion is not installed. Additional, fewer and/or other conversions may be specified.

[00101]	Bit	Data Type
[00102]	0	Reserved
[00103]	1	BFP tiny format
[00104]	2	BFP short format
[00105]	3-15	Reserved

[00106] Although one example of a parameter block for a query function is described with reference to FIG. 6E, other formats of a parameter block for a query function, including the NNPA-Query Available Functions operation, may be used. The format may depend, in one example, on the type of query function to be performed. Further, the parameter block and/or each field of the parameter block may include additional, fewer and/or other information.

[00107] In addition to the parameter block for a query function, in one example, there is a parameter block format for non-query functions, such as non-query functions of the Neural-Network Processing Assist instruction. One example of a parameter block used by a non-

query function, such as a non-query function of the Neural Network Processing Assist instruction, is described with reference to FIG. 6F.

[00108] As shown, in one example, a parameter block 650 employed by, e.g., the non-query functions of the Neural Network Processing Assist instruction includes, for instance:

[00109] Parameter Block Version Number 652: This field (e.g., bytes 0-1) of the parameter block specifies the version and size of the parameter block. In one example, bits 0-8 of the parameter block version number are reserved and are to contain zeros, and bits 9-15 of the parameter block version number contain an unsigned binary integer specifying the format of the parameter block. The query function provides a mechanism of indicating the parameter block formats available. When the size or format of the parameter block specified is not supported by the model, a response code of, e.g., 0001 hex is stored in general register 0 and the instruction completes by setting a condition code, e.g., condition code 1. The parameter block version number is specified by the program and is not modified during the execution of the instruction.

[00110] Model Version Number 654: This field (e.g., byte 2) of the parameter block is an unsigned binary integer identifying the model which executed the instruction (e.g., the particular non-query function). When a continuation flag (described below) is one, the model version number may be an input to the operation for the purpose of interpreting the contents of a continuation state buffer field (described below) of the parameter block to resume the operation.

[00111] Continuation Flag 656: This field (e.g., bit 63) of the parameter block, when, e.g., one, indicates the operation is partially complete and the contents of the continuation state buffer may be used to resume the operation. The program is to initialize the continuation flag to zero and not modify the continuation flag in the event the instruction is to be re-executed for the purpose of resuming the operation; otherwise, results are unpredictable.

[00112] If the continuation flag is set at the beginning of the operation and the contents of the parameter block have changed since the initial invocation, results are unpredictable.

[00113] Function-specific-save-area-address 658: This field (e.g., bytes 56-63) of the parameter block includes the logical address of the function specific save area. In one example, the function-specific-save-area-address is to be aligned on a 4 K-byte boundary;

otherwise, a response code of, e.g., 0015 hex is set in general register 0 and the instruction completes with a condition code of, e.g., 1. The address is subject to the current addressing mode. The size of the function specific save area depends on the function code.

[00114] When the entire function specific save area overlaps the program event recording (PER) storage area designation, a PER storage alteration event is recognized, when applicable, for the function specific save area. When only a portion of the function specific save area overlaps the PER storage area designation, it is model-dependent which of the following occurs:

[00115] * A PER storage alteration event is recognized, when applicable, for the entire function specific save area.

[00116] * A PER storage alteration event is recognized, when applicable, for the portion of the function specific save area that is stored.

[00117] When the entire parameter block overlaps the PER storage area designation, a PER storage alteration event is recognized, when applicable, for the parameter block. When only a portion of the parameter block overlaps the PER storage area designation, it is model-dependent which of the following occurs:

[00118] * A PER storage alteration event is recognized, when applicable, for the entire parameter block.

[00119] * A PER storage alteration event is recognized, when applicable, for the portion of the parameter block that is stored.

[00120] A PER zero-address detection event is recognized, when applicable, for the parameter block. Zero address detection does not apply to the tensor addresses or the function-specific-save-area-address, in one example.

[00121] Output Tensor Descriptors (e.g., 1-2) 660/Input Tensor Descriptors (e.g., 1-3) 665: One example of a tensor descriptor is described with reference to FIG. 6G. In one example, a tensor descriptor 660, 665 includes:

[00122] Data Layout Format 682: This field (e.g., byte 0) of the tensor descriptor specifies the data layout format. Valid data layout formats include, for instance (additional, fewer and/or other data layout formats are possible):

[00123]	<u>Format</u>	<u>Description</u>	<u>Alignment (bytes)</u>
[00124]	0	4D-feature tensor	4096
[00125]	1	4D-kernel tensor	4096
[00126]	2-255	Reserved	--

[00127] If an unsupported or reserved data layout format is specified, the response code of, e.g., 0010 hex, is stored in general register 0 and the instruction completes by setting condition code, e.g., 1.

[00128] Data Type 684: This field (e.g., byte 1) specifies the data type of the tensor. Examples of supported data types are described below (additional, fewer and/or other data types are possible):

[00129]	<u>Value</u>	<u>Data Type</u>	<u>Data Size (bits)</u>
[00130]	0	NNP data-type-1	16
[00131]	1-255	Reserved	--

[00132] If an unsupported or reserved data type is specified, a response code of, e.g., 0011 hex is stored in general register 0 and the instruction completes by setting condition code, e.g., 1.

[00133] Dimension 1-4 Index Size 686: Collectively, dimension index sizes one through four (e.g., E4, E3, E2, E1) specify the shape of a 4D-tensor. Each dimension index size is to be greater than zero and less than or equal to the maximum dimension index size (640, FIG. 6E); otherwise, a response code of, e.g., 0012 hex is stored in general register 0 and the instruction completes by setting condition code, e.g., 1. The total tensor size is to be less than or equal to the maximum tensor size (642, FIG. 6E); otherwise, a response code, e.g., 0013 hex is stored in general register 0 and the instruction completes by setting condition code, e.g., 1.

[00134] In one example, to determine the number of bytes in a 4D-feature tensor with elements of NNPA-data-type-1 (i.e., total tensor size), the following is used: $\text{dimension-index-4} * \text{dimension-index-3} * \text{ceil}(\text{dimension-index-2}/32) * 32 * \text{ceil}(\text{dimension-index-1}/64) * 64 * 2$.

[00135] Tensor Address 688: This field (e.g., bytes 24-31) of the tensor descriptor includes a logical address of the leftmost byte of the tensor. The address is subject to the current addressing mode.

[00136] If the address is not aligned on the boundary of the associated data layout format, a response code of, e.g., 0014 hex, is stored in general register 0 and the instruction completes by setting condition code, e.g., 1.

[00137] In the access register mode, access register 1 specifies the address space containing all active input and output tensors in storage.

[00138] Returning to FIG. 6F, parameter block 650 further includes, in one example, function-specific-parameters 1-5 (670), which may be used by specific functions, as described herein.

[00139] Further, parameter block 650 includes, in one example, a continuation state buffer field 675, which includes data (or a location of data) to be used if operation of this instruction is to be resumed.

[00140] As an input to the operation, reserved fields of the parameter block should contain zeros. When the operation ends, reserved fields may be stored as zeros or remain unchanged.

[00141] Although one example of a parameter block for a non-query function is described with reference to FIG. 6F, other formats of a parameter block for a non-query function, including a non-query function of the Neural Network Processing Assist instruction, may be used. The format may depend, in one example, on the type of function to be performed. Further, although one example of a tensor descriptor is described with reference to FIG. 6G, other formats may be used. Further, different formats for input and output tensors may be used. Other variations are possible.

[00142] Further details regarding various functions supported by one embodiment of the Neural Network Processing Assist instruction are described below. Additional, fewer, and/or other functions may be supported.

[00143] Function Code 0: NNPA-QAF (Query Available Functions)

[00144] The Neural Network Processing Assist (NNPA) query function provides a mechanism to indicate selected information, such as, for instance, the availability of installed functions, installed parameter block formats, installed data types, installed data layout formats, maximum dimension index size and maximum tensor size. The information is obtained and placed in a selected location, such as a parameter block (e.g., parameter block 630). When the operation ends, reserved fields of the parameter block may be stored as zeros or may remain unchanged.

[00145] In execution of one embodiment of the query function, a processor, such as general-purpose processor 104, obtains information relating to a specific model of a selected processor, such as a specific model of a neural network processor, such as neural network processor 105. A specific model of a processor or machine has certain capabilities. Another model of the processor or machine may have additional, fewer and/or different capabilities and/or be of a different generation (e.g., a current or future generation) having additional, fewer and/or different capabilities. The obtained information is placed in a parameter block (e.g., parameter block 630) or other structure that is accessible to and/or for use with one or more applications that may use this information in further processing. In one example, the parameter block and/or information of the parameter block is maintained in memory. In other embodiments, the parameter block and/or information may be maintained in one or more hardware registers. As another example, the query function may be a privileged operation executed by the operating system, which makes available an application programming interface to make this information available to the application or non-privileged program. In yet a further example, the query function is performed by a special-purpose processor, such as neural network processor 105. Other variations are possible.

[00146] The information is obtained, e.g., by the firmware of the processor executing the query function. The firmware has knowledge of the attributes of the specific model of the specific processor (e.g., neural network processor). This information may be stored in, e.g., a control block, register and/or memory and/or otherwise be accessible to the processor executing the query function.

[00147] The obtained information includes, for instance, model-dependent detailed information regarding at least one or more data attributes of the specific processor, including, for instance, one or more installed or supported data types, one or more installed

or supported data layout formats and/or one or more installed or supported data sizes of the selected model of the specific processor. This information is model-dependent in that other models (e.g., previous models and/or future models) may not support the same data attributes, such as the same data types, data sizes, and/or data layout formats. When execution of the query function (e.g., NNPA-QAF function) completes, condition code 0, as an example, is set. Condition codes 1, 2 and 3 are not applicable to the query function, in one example. Further information relating to the obtained information is described below.

[00148] As indicated, in one example, the obtained information includes model-dependent information about one or more data attributes of, e.g., a particular model of a neural network processor. One example of a data attribute is installed data types of the neural network processor. For instance, a particular model of a neural network processor (or other processor) may support one or more data types, such as a NNP-data-type-1 data type (also referred to as a neural network processing-data-type-1 data type) and/or other data types, as examples. The NNP-data-type-1 data type is a 16-bit floating-point format that provides a number of advantages for deep learning training and inference computations, including, for instance: preserves the accuracy of deep learning networks; eliminates the subnormal format which simplifies rounding modes and handling of corner cases; automatic rounding to the nearest value for arithmetic operations; and special entities of infinity and not-a-number (NaN) are combined into one value (NINF), which is accepted and handled by arithmetic operations. NINF provides better defaults for exponent overflow and invalid operations (such as division by zero). This allows many programs to continue running without hiding such errors and without using specialized exception handlers. Other model-dependent data types are also possible.

[00149] One example of a format of the NNP-data-type-1 data type is depicted in FIG. 7. As depicted, in one example, NNP-data-type-1 data may be represented in a format 700, which includes, for instance, a sign 702 (e.g., bit 0), an exponent + 31 704 (e.g., bits 1-6) and a fraction 706 (e.g., bits 7-15).

[00150] Example properties of the NNP-data-type-1 format are depicted below:

[00151]	<u>Property</u>	<u>NNP-data-type-1</u>
[00152]	Format length (bits)	16 bits
[00153]	Biased-exponent length (bits)	6 bits

[00154]	Fraction length (bits)	9 bits
[00155]	Precision (p)	10 bits
[00156]	Maximum left-units-view exponent (Emax)	32
[00157]	Minimum left-units-view exponent (Emin)	-31
[00158]	Left-units-view (LUV) bias	31
[00159]	Nmax	$(1-2^{-9}) \times 2^{33} \approx 8.6 \times 10^9$
[00160]	Nmin	$(1+2^{-9}) \times 2^{-31} \approx 4.6 \times 10^{-10}$
[00161]	Dmin	---
[00162]	Where $\approx$ indicates that the value is approximate, Nmax is largest (in magnitude) representable finite number, and Nmin is smallest (in magnitude) representable number.	
[00163]	Further details relating to the NNP-data-type-1 data type are described below:	
[00164]	Biased Exponent: The bias that is used to allow exponents to be expressed as unsigned numbers is shown above. Biased exponents are similar to characteristics of the binary floating-point format, except that no special meanings are attached to biased exponents of all zeros and all ones, as described below with reference to the classes of the NNP-data-type-1 data type.	
[00165]	Significand: The binary point of a NNP-data-type-1 number is considered to be to the left of the leftmost fraction bit. To the left of the binary point there is an implied unit bit, which is considered to be one for normal numbers and zero for zeros. The fraction with the implied unit bit appended on the left is the significand of the number.	
[00166]	The value of a normal NNP-data-type-1 is the significand multiplied by the radix 2 raised to the power of the unbiased exponent.	
[00167]	Values of Non-Zero Numbers: The values of non-zero numbers are shown below:	
[00168]	<u>Number Class</u>	<u>Value</u>

[00169] Normal numbers $\pm 2^{e-31} \times (1.f)$

[00170] Where e is biased exponent shown in decimal, and f is fraction in binary.

[00171] In one embodiment, there are three classes of NNP-data-type-1 data, including numeric and related non-numeric entities. Each data item includes a sign, an exponent and a significand. The exponent is biased such that all biased exponents are non-negative unsigned numbers and the minimum biased exponent is zero. The significand includes an explicit fraction and an implicit unit bit to the left of the binary point. The sign bit is zero for plus and one for minus.

[00172] All non-zero finite numbers permitted have a unique NNP-data-type-1 representation. There are no subnormal numbers, which numbers might allow multiple representations for the same values, and there are no subnormal arithmetic operations. The three classes include, for instance:

[00173]	<u>Data Class</u>	<u>Sign</u>	<u>Biased Exponent</u>	<u>Unit Bit*</u>	<u>Fraction</u>
[00174]	Zero	$\pm$	0	0	0
[00175]	Normal numbers	$\pm$	0	1	Not 0
[00176]	Normal numbers	$\pm$	Not 0, not all ones	1	Any
[00177]	Normal numbers	$\pm$	All ones	-	Not all ones
[00178]	NINF	$\pm$	All ones	-	All ones

[00179] Where: – indicates does not apply, * indicates the unit bit is implied, NINF is not a number or infinity.

[00180] Further details regarding each of the classes are described below:

[00181] Zeros: Zeros have a biased exponent of zero and a zero fraction. The implied unit bit is zero.

[00182] Normal Numbers: Normal numbers may have a biased exponent of any value. When the biased exponent is 0, the fraction is to be non-zero. When the biased exponent is all ones, the fraction is not to be all ones. Other biased exponent values may have any fraction value. The implied unit bit is one for all normal numbers.

[00183] NINF: A NINF is represented by a biased exponent of all ones and a fraction of all ones. A NINF represents a value not in the range of representable values in NNP-data-type-1 (i.e., 16-bit floating point designed for deep learning that has 6 exponent bits and 9 fraction bits). Normally, NINFs are just propagated during computations so that they will remain visible at the end.

[00184] Although the NNP-data-type-1 data type is supported in one example, other specialized or non-standard data types may also be supported, as well as one or more standard data types including, but not limited to: IEEE 754 short precision, binary floating-point 16-bit, IEEE half precision floating point, 8-bit floating point, 4-bit integer format and/or 8-bit integer format, to name a few. These data formats have different qualities for neural network processing. As an example, smaller data types (e.g., less bits) can be processed faster and use less cache/memory, and larger data types provide greater result accuracy in the neural network. Each data type to be supported may have one or more assigned bits in the query parameter block (e.g., in installed data types field 636 of parameter block 630). For instance, specialized or non-standard data types supported by a particular processor are indicated in the installed data types field but standard data types are not indicated. In other embodiments, one or more standard data types are also indicated. Other variations are possible.

[00185] In one particular example, bit 0 of installed data types field 636 is reserved for the NNP-data-type-1 data type, and when it is set to, e.g., 1, it indicates that the processor supports NNP-data-type-1. In one example, the bit vector of installed data types is configured to represent up to 16 data types, in which a bit is assigned to each data type. However, a bit vector in other embodiments may support more or fewer data types. Further, a vector may be configured in which one or more bits are assigned to a data type. Many examples are possible and/or additional, fewer and/or other data types may be supported and/or indicated in the vector.

[00186] In one example, the query function obtains an indication of the data types installed on the model-dependent processor and places the indication in the parameter block by, e.g., setting one or more bits in installed data types field 636 of parameter block 630. Further, in one example, the query function obtains an indication of installed data layout formats (another data attribute) and places the information in the parameter block by, e.g., setting one or more bits in installed data layout formats field 638. Example data layout

formats include, for instance, a 4D-feature tensor layout and a 4D-kernel tensor layout. These data layout formats arrange data in storage for a tensor in a way that increases processing efficiency in execution of the functions of the Neural Network Processing Assist instruction. For instance, to operate efficiently, the Neural Network Processing Assist instruction uses input tensors provided in particular data layout formats. Although example layouts are provided, additional, fewer and/or other layouts may be provided for the functions described herein and/or other functions.

[00187] The use or availability of layouts for a particular processor model is provided by the vector of installed data layout formats (e.g., field 638 of parameter block 630). The vector is, for instance, a bit vector of installed data layout formats that allows the CPU to convey to applications which layouts are supported. For instance, bit 0 is reserved for the 4D-feature tensor layout, and when it is set to, e.g., 1, it indicates that the processor supports a 4D-feature tensor layout; and bit 1 is reserved for the 4D-kernel tensor layout, and when it is set to, e.g., 1, it indicates that the processor supports a 4D-kernel tensor layout. In one example, the bit vector of installed data layout formats is configured to represent up to 16 data layouts, in which a bit is assigned to each data layout. However, a bit vector in other embodiments may support more or fewer data layouts. Further, a vector may be configured in which one or more bits are assigned to data layouts. Many examples are possible. Further details regarding the 4D-feature tensor layout and the 4D-kernel tensor layout are described below. Again, other layouts may be used now or in the future to optimize performance.

[00188] In one example, the Neural Network Processing Assist instruction operates with 4D-tensors, i.e., tensors with 4 dimensions. These 4-D tensors are obtained from generic input tensors described herein in, e.g., row-major, i.e., when enumerating the tensor elements in increasing memory address order, the inner dimension called E1 will be stepped up first through the E1-index-size values starting with 0 through the E1-index-size -1, before the index of the E2 dimension will be increased and the stepping through the E1 dimension is repeated. The index of the outer dimension called the E4 dimension is increased last.

[00189] Tensors that have a lower number of dimensions (e.g., 3D- or 1D-tensors) will be represented as 4D-tensors with one or more dimensions of the 4D-tensor exceeding the original tensor dimensions set to 1.

[00190] The transformation of a row-major generic 4D-tensor with dimensions E4, E3, E2, E1 into a 4D-feature tensor layout (also referred to herein as NNPA data layout format 0 4D-feature tensor) is described herein:

[00191] A resulting tensor can be represented, for instance, as a 4D-tensor of, e.g., 64-element vectors or a 5D-tensor with dimensions:

[00192] $E4, \lceil E1/64 \rceil, E3, \lceil E2/32 \rceil * 32, 64$, where $\lceil \cdot \rceil$ refers to a ceil function. (Stated another way: $E4 * E3 * \text{ceil}(E2/32) * 32 * \text{ceil}(E1/64) * 64$ elements.)

[00193] An element $[e4][e3][e2][e1]$ of the generic tensor may be mapped to the following element of the resulting 5D-tensor:

[00194] $[e4][\lfloor e1/64 \rfloor][e3][e2][e1 \bmod 64]$, where $\lfloor \cdot \rfloor$ is a floor function and mod is modulo. (Stated another way: element $(E3 * e2_limit * e1_limit * e4x) + (e2_limit * e3x * 64) + (e2x * 64) + (\lfloor e1x/64 \rfloor * e2_limit * E3 * 64) + (e1x \bmod 64)$, where $e2_limit = \lceil E2/32 \rceil * 32$ and $e1_limit = \lceil E1/64 \rceil * 64$.)

[00195] The resulting tensor may be larger than the generic tensor. Elements of the resulting tensor with no corresponding elements in the generic tensor are called pad elements.

[00196] Consider the element $[fe4][fe1][fe3][fe2][fe0]$ of a NNPA data layout format 0 4D-feature tensor of a 64-element vector or its equivalent representation as a 5D-tensor of elements. This element is either a pad element or its corresponding element in the generic 4D tensor with dimensions E4, E3, E2, E1 can be determined with the following formula:

[00197] · if $fe2 \geq E2$ then this is an E2 (or page)-pad element

[00198] · else if $fe1 * 64 + fe0 \geq E1$ then this is an E1 (or row)-pad element

[00199] · else corresponding element in generic 4D tensor is:

[00200] $[fe4][fe3][fe2][fe1 * 64 + fe0]$

[00201] For convolutional neural network based artificial intelligence models, the meaning of the 4 dimensions of a feature tensor can generally be mapped to:

[00202] · E4: N – Size of mini-batch

- [00203] · E3: H – Height of the 3D-tensor/image
- [00204] · E2: W – Width of the 3D-tensor/image
- [00205] · E1: C – Channels or classes of the 3D-tensor
- [00206] For machine learning or recurrent neural network based artificial intelligence models, the meaning of the 4 dimensions of a 4D-feature tensor may generally be mapped to:
- [00207] · E4: T – Number of time-steps or models
- [00208] · E3: Reserved, generally set to 1
- [00209] · E2: N_{mb} – Minibatch size
- [00210] · E1: L – Features
- [00211] The NNPA data layout format 0 provides, e.g., two dimensional data locality with 4k-Bytes blocks of data (pages) as well as 4k-Byte block data alignment for the outer dimensions of the generated tensor.
- [00212] Pad element bytes are ignored for the input tensors and unpredictable for output tensors. PER storage-alteration on pad bytes is unpredictable.
- [00213] One example of an input data layout for a 4D-feature tensor layout, having dimensions E1, E2, E3 and E4, is shown in FIGS. 8A-8C, and an example output for the 4D-feature tensor layout is depicted in FIGS. 9A-9C. Referring to FIG. 8A, a 3D-tensor 800 is shown, which has dimensions E1, E2 and E3. In one example, each 3D-tensor includes a plurality of 2D-tensors 802. Thus, in the example shown, a plurality of 2D-tensors (e.g., 3 2D-tensors) creates a 3D-tensor, and a plurality of 3D-tensors (e.g., 3 3D-tensors) creates the 4D-tensor. The numbers in each 2D-tensor 802 describe memory offsets of where each of its elements would be in memory. The inputs are used to lay-out the data of the original tensor (e.g., original 4D-tensor of FIGS. 8A-8C) in memory, as shown in FIGS. 9A-9C, which correspond to FIGS. 8A-8C.
- [00214] In FIG. 9A, as an example, a unit of memory 900 (e.g., a memory page) includes a pre-selected number (e.g., 32) of rows 902, each of which is identified by, e.g., `e2_page_idx`; and each row has a pre-selected number (e.g., 64) of elements 904, each

identified by, e.g., `e1_page_idx`. If a row does not include the pre-selected number of elements, it is padded 906, referred to as row or E1 padding; and if the memory unit does not have a pre-selected number of rows, it is padded 908, referred to as page or E2 padding. As examples, the row padding is e.g., zeros or other values and the page padding is, e.g., existing values, zeros, or other values.

[00215] In one example, output elements of a row are provided in memory (e.g., in a page) based on element positions in the E1 direction of its corresponding input. For instance, referring to FIG. 8A, element positions 0, 1 and 2 of the three matrices shown (e.g., element positions at a same location in each matrix) are shown in row 0 of page 0 of FIG. 9A, etc. In this example, the 4D-tensor is small and all of the elements of each 2D-tensor representing the 4D-tensor fits in one page. However, this is only one example. A 2D-tensor may include one or more pages. As shown in FIG. 3A, the 2D-tensor in that example includes 12 pages. However, that is only one example. Again, a 2D-tensor may include one or more pages. If a 2D-tensor is created based on a reformatting of a 4D-tensor, then the number of pages of the 2D-tensor is based on the size of the 4D-tensor. In one example, one or more ceil functions are used to determine the number of rows in a 2D-tensor and the number of elements in each row, which will indicate how many pages are to be used. Other variations are possible.

[00216] The reformatted 2D-tensors (e.g., that are concatenated) are, in accordance with one or more aspects of the present invention, based on the 4D-feature tensor layout and are stored in memory, as described herein. The 2D-tensors input to the cell activations are, for instance, 4D-tensors in which E3 and E4 are set to one.

[00217] In addition to the 4D-feature tensor layout, in one example, a neural network processor may support a 4D-kernel tensor, which re-arranges the elements of a 4D-tensor to reduce the number of memory accesses and data gathering steps when executing certain artificial intelligence (e.g., neural network processing assist) operations, such as a convolution. As an example, a row-major generic 4D-tensor with dimensions E4, E3, E2, E1 is transformed into a NNPA data layout format 1 4D-kernel tensor (4D-kernel tensor), as described herein:

[00218] A resulting tensor can be represented as a 4D-tensor of, e.g., 64-element vectors or a 5D-tensor with dimensions:

[00219] $\lceil E1/64 \rceil, E4, E3, \lceil E2/32 \rceil * 32, 64$, where $\lceil \cdot \rceil$ refers to a ceil function. (Stated another way: $E4 * E3 * \text{ceil}(E2/32) * 32 * \text{ceil}(E1/64) * 64$ elements.)

[00220] An element $[e4][e3][e2][e1]$ of the generic tensor may be mapped to the following element of the resulting 5D-tensor:

[00221] $\lfloor e1/64 \rfloor [e4][e3][e2][e1 \text{ MOD } 64]$, where $\lfloor \cdot \rfloor$ refers to a floor function and mod is modulo. Stated another way: element $(\lfloor e1x/64 \rfloor * E4 * E3 * e2_limit * 64) + (e4x * E3 * e2_limit * 64) + (e3x * e2_limit * 64) + (e2x * 64) + (e1x \text{ mod } 64)$, where $e2_limit = \lceil E2/32 \rceil * 32$ and $e1_limit = \lceil E1/64 \rceil * 64$.

[00222] The resulting tensor may be larger than the generic tensor. Elements of the resulting tensor with no corresponding elements in the generic tensor are called pad elements.

[00223] Consider the element $[fe1][fe4][fe3][fe2][fe0]$ of a NNPA data layout format 1 4D-feature tensor of 64-element vectors or its equivalent representation as a 5D-tensor of elements. This element is either a pad element or its corresponding element in the generic 4D tensor with dimensions $E4, E3, E2, E1$ can be determined with the following formula:

[00224] · if $fe2 \geq E2$ then this is an $E2$ (or page)-pad element

[00225] · else if $fe1 * 64 + fe0 \geq E1$ then this is an $E1$ (or row)-pad element

[00226] · else corresponding element in generic 4D tensor is

[00227] $[fe4][fe3][fe2][fe1 * 64 + fe0]$

[00228] For convolutional neural network based artificial intelligence models, the meaning of the 4 dimensions of a kernel tensor can generally be mapped to:

[00229] · $E4$: H – Height of the 3D-tensor/image

[00230] · $E3$: W – Width of the 3D-tensor/image

[00231] · $E2$: C – Number of Channels of the 3D-tensor

[00232] · $E1$: K – Number of Kernels

[00233] The NNPA data layout format 1 provides, e.g., two dimensional kernel parallelism within 4k-Byte blocks of data (pages) as well as 4k-Byte block data alignment for the outer dimensions of the generate tensor for efficient processing.

[00234] Pad bytes are ignored for the input tensors. PER storage-alteration on pad bytes is unpredictable.

[00235] Again, although example data layout formats include a 4D-feature tensor layout and a 4D-kernel tensor layout, other data layout formats may be supported by the processor (e.g., neural network processor 105). An indication of supported data layouts is obtained and placed in the query parameter block by setting one or more bits in, e.g., field 638.

[00236] The query parameter block also includes, in accordance with one or more aspects of the present invention, other data attribute information, which includes, e.g., supported size information for the data. A processor, such as a neural network processor, typically has limitations based on internal buffer sizes, processing units, data bus structures, firmware limitations, etc. that can limit the maximum size of tensor dimensions and/or the overall size of a tensor. Therefore, the query function provides fields to convey these limits to applications. For instance, the processor, based on executing the query function, obtains various data sizes, such as a maximum dimension index size (e.g., 65,536 elements) and a maximum tensor size (e.g., 8 GB), and includes this information in fields 640 and 642, respectively, of the parameter block (e.g., parameter block 630). Additional, fewer and/or other size information may also be supported by the processor (e.g., neural network processor 105), and thus, obtained and placed in the parameter block, e.g., fields 640, 642 and/or other fields. In other embodiments, the limitations could be smaller or larger, and/or the sizes may be in other units, such as bytes instead of elements, elements instead of bytes, etc. Further, other embodiments allow for different maximum sizes of each dimension, rather than the same maximum for all dimensions. Many variations are possible.

[00237] In accordance with one or more aspects of the present invention, a query function is provided to determine model-dependent information relating to a specific processor. (A processor may also support standard data attributes, such as standard data types, standard data layouts, etc., which are implied and not necessarily presented by the query function; although, in other embodiments, the query function may indicate all or various selected subsets of data attributes, etc.) Although example information is provided,

other information may be provided in other embodiments. The obtained information, which may be different for different models of a processor and/or of different processors, is used to perform artificial intelligence and/or other processing. The artificial intelligence and/or other processing may employ one or more non-query functions of, e.g., the Neural Network Processing Assist instruction. A specific non-query function employed in the processing is performed by executing the Neural Network Processing Assist instruction one or more times and specifying the non-query specific function.

[00238] Further details of example non-query functions supported by the Neural Network Processing Assist instruction are described below (additional, fewer and/or other functions may be supported in other embodiments):

[00239] Function Code 16: NNPA-ADD (Addition)

[00240] When the NNPA-ADD function is specified, each element of the input tensor 1 described by tensor descriptor 1 is added to the corresponding element of the input tensor 2 described by tensor descriptor 2, and the resulting sum is placed in the corresponding element of the output tensor described by the output tensor descriptor.

[00241] In one example, if the specified data layout in any of the specified tensor descriptors does not specify a 4D-feature tensor (e.g., data layout = 0) or if the data type in any specified tensor descriptor does not specify NNP-data-type-1 (e.g., data type = 0), response code, e.g., 0010 hex or 0011 hex, respectively, is set in general register 0 and the instruction completes with condition code, e.g., 1.

[00242] The shape, the data layout and the data type of input tensor 1, input tensor 2 and the output tensor are to be the same, in one example; otherwise, a general operand data exception is recognized.

[00243] The output tensor descriptor 2, input tensor descriptor 3, function-specific-parameters 1-5 and function-specific-save-area-address fields are ignored, in one example.

[00244] Function Code 17: NNPA-SUB (Subtraction)

[00245] When the NNPA-SUB function is specified, each element of the input tensor 2 described by tensor descriptor 2 is subtracted from the corresponding element of the input tensor 1 described by tensor descriptor 1, and the resulting difference is placed in the corresponding element of the output tensor.

[00246] In one example, if the specified data layout in any of the specified tensor descriptors does not specify a 4D-feature tensor (e.g., data layout = 0) or if the data type in any specified tensor descriptor does not specify NNP-data-type-1 (e.g., data type = 0), response code, e.g., 0010 hex or 0011 hex, respectively, is set in general register 0 and the instruction completes with condition code, e.g., 1.

[00247] The shape, the data layout and the data type of input tensor 1, input tensor 2 and the output tensor are to be the same, in one example; otherwise, a general operand data exception is recognized.

[00248] The output tensor descriptor 2, input tensor descriptor 3, function-specific-parameters 1-5 and function-specific-save-area-address fields are ignored, in one example.

[00249] Function Code 18: NNPA-MUL (Multiplication)

[00250] When the NNPA-MUL function is specified, the product of each element of the input tensor 1 (the multiplier) described by tensor descriptor 1 and the corresponding element of the input tensor 2 (the multiplicand) described by tensor descriptor 2 is placed in the corresponding element of the output tensor.

[00251] In one example, if the specified data layout in any of the specified tensor descriptors does not specify a 4D-feature tensor (e.g., data layout = 0) or if the data type in any specified tensor descriptor does not specify NNP-data-type-1 (e.g., data type = 0), response code, e.g., 0010 hex or 0011 hex, respectively, is set in general register 0 and the instruction completes with condition code, e.g., 1.

[00252] The shape, the data layout and the data type of input tensor 1, input tensor 2 and the output tensor are to be the same, in one example; otherwise, a general operand data exception is recognized.

[00253] The output tensor descriptor 2, input tensor descriptor 3, function-specific-parameters 1-5 and function-specific-save-area-address fields are ignored, in one example.

[00254] Function Code 19: NNPA-DIV (Division)

[00255] When the NNPA-DIV function is specified, each element of the input tensor 1 described by tensor descriptor 1 (the dividend) is divided by the corresponding element of

the input tensor 2 (the divisor) described by tensor descriptor 2, and the quotient is placed in the corresponding element of the output tensor.

[00256] In one example, if the specified data layout in any of the specified tensor descriptors does not specify a 4D-feature tensor (e.g., data layout = 0) or if the data type in any specified tensor descriptor does not specify NNP-data-type-1 (e.g., data type = 0), response code, e.g., 0010 hex or 0011 hex, respectively, is set in general register 0 and the instruction completes with condition code, e.g., 1.

[00257] The shape, the data layout and the data type of input tensor 1, input tensor 2 and the output tensor are to be the same, in one example; otherwise, a general operand data exception is recognized.

[00258] The output tensor descriptor 2, input tensor descriptor 3, function-specific-parameters 1-5 and function-specific-save-area-address fields are ignored, in one example.

[00259] Function Code 20: NNPA-MIN (Minimum)

[00260] When the NNPA-MIN function is specified, each element of the input tensor 1 described by tensor descriptor 1 is compared to the corresponding element of the input tensor 2 described by tensor descriptor 2. The smaller of the two values is placed into the corresponding element of the output tensor descriptor. If both values are equal, then the value is placed in the corresponding element of the output tensor.

[00261] In one example, if the specified data layout in any of the specified tensor descriptors does not specify a 4D-feature tensor (e.g., data layout = 0) or if the data type in any specified tensor descriptor does not specify NNP-data-type-1 (e.g., data type = 0), response code, e.g., 0010 hex or 0011 hex, respectively, is set in general register 0 and the instruction completes with condition code, e.g., 1.

[00262] The shape, the data layout and the data type of input tensor 1, input tensor 2 and the output tensor are to be the same, in one example; otherwise, a general operand data exception is recognized.

[00263] The output tensor descriptor 2, input tensor descriptor 3, function-specific-parameters 1-5 and function-specific-save-area-address fields are ignored, in one example.

[00264] Function Code 21: NNPA-MAX (Maximum)

[00265] When the NNPA-MAX function is specified, each element of the input tensor 1 described by tensor descriptor 1 is compared to the corresponding element of the input tensor 2 described by tensor descriptor 2. The greater of the two values is placed in the corresponding element of the output tensor descriptor. If both values are the same, then the value is placed in the corresponding element of the output tensor.

[00266] In one example, if the specified data layout in any of the specified tensor descriptors does not specify a 4D-feature tensor (e.g., data layout = 0) or if the data type in any specified tensor descriptor does not specify NNP-data-type-1 (e.g., data type = 0), response code, e.g., 0010 hex or 0011 hex, respectively, is set in general register 0 and the instruction completes with condition code, e.g., 1.

[00267] The shape, the data layout and the data type of input tensor 1, input tensor 2 and the output tensor are to be the same, in one example; otherwise, a general operand data exception is recognized.

[00268] The output tensor descriptor 2, input tensor descriptor 3, function-specific-parameters 1-5 and function-specific-save-area-address fields are ignored, in one example.

[00269] Function Code 32: NNPA-LOG (Natural Logarithm)

[00270] When the NNPA-LOG function is specified, for each element of the input tensor described by tensor descriptor 1, if that element is greater than zero, the corresponding element in the output tensor, described by the output tensor descriptor, is the natural logarithm of that element. Otherwise, the corresponding element in the output tensor is not numerically representable and the value associated with negative infinity in the target data type is stored.

[00271] In one example, if the specified data layout in any of the specified tensor descriptors does not specify a 4-D feature tensor (e.g., data layout = 0) or if the data type in any specified tensor descriptor does not specify NNP-data-type-1 (e.g., data type = 0), response code, e.g., 0010 hex or 0011 hex, respectively, is set in general register 0 and the instruction completes with condition code, e.g., 1.

[00272] The shape, the data layout and the data type of input tensor 1 and the output tensor are to be the same, in one example; otherwise, a general operand data exception is recognized.

[00273] The output tensor descriptor 2, input tensor descriptor 2, input tensor descriptor 3, function-specific-parameters 1-5 and function-specific-save-area-address fields are ignored, in one example.

[00274] Function Code 33: NNPA-EXP (Exponential)

[00275] When the NNPA-EXP function is specified, for each element of the input tensor described by tensor descriptor 1, the corresponding element in the output tensor described by the output tensor descriptor is the exponential of that element.

[00276] In one example, if the specified data layout in any of the specified tensor descriptors does not specify a 4D-feature tensor (e.g., data layout = 0) or if the data type in any specified tensor descriptor does not specify NNP-data-type-1 (e.g., data type = 0), response code, e.g., 0010 hex or 0011 hex, respectively, is set in general register 0 and the instruction completes with condition code, e.g., 1.

[00277] The shape, the data layout and the data type of input tensor 1 and the output tensor are to be the same, in one example; otherwise, a general operand data exception is recognized.

[00278] The output tensor descriptor 2, input tensor descriptor 2, input tensor descriptor 3, function-specific-parameters 1-5 and function-specific-save-area-address fields are ignored, in one example.

[00279] Function Code 49: NNPA-RELU (Rectified Linear Unit)

[00280] When the NNPA-RELU function is specified, for each element of the input tensor described by tensor descriptor 1, if that element is less than or equal to zero, the corresponding element in the output tensor described by the output tensor descriptor is zero. Otherwise, the corresponding element in the output tensor is the minimum of the element in the input tensor and the clipping value specified in function-specific-parameter 1.

[00281] As an example, function-specific-parameter 1 defines the clipping value for the RELU operation. For instance, the clipping value is in bits 16-31 of function-specific-parameter 1. The clipping value is specified in, e.g., the NNPA-data-type-1 format. A clipping value of zero indicates to use the maximum positive value; in other words, no clipping is performed. If a negative value is specified, a general operand data exception is recognized.

[00282] In one example, if the specified data layout in any of the specified tensor descriptors does not specify a 4D-feature tensor (e.g., data layout = 0) or if the data type in any specified tensor descriptor does not specify NNP-data-type-1 (e.g., data type = 0), response code, e.g., 0010 hex or 0011 hex, respectively, is set in general register 0 and the instruction completes with condition code, e.g., 1.

[00283] The shape, the data layout and the data type of input tensor 1 and the output tensor are to be the same, in one example; otherwise, a general operand data exception is recognized.

[00284] The output tensor descriptor 2, input tensor descriptor 2, input tensor descriptor 3, and function-specific-save-area-address fields are ignored, in one example. Function-specific-parameters 2-5 are to contain zeros, in one example.

[00285] Function Code 50: NNPA-TANH (Tangent)

[00286] When the NNPA-TANH function is specified, for each element of the input tensor described by tensor descriptor 1, the corresponding element value in the output tensor described by the output tensor descriptor is the hyperbolic tangent of that element.

[00287] In one example, if the specified data layout in any of the specified tensor descriptors does not specify a 4D-feature tensor (e.g., data layout = 0) or if the data type in any specified tensor descriptor does not specify NNP-data-type-1 (e.g., data type = 0), response code, e.g., 0010 hex or 0011 hex, respectively, is set in general register 0 and the instruction completes with condition code, e.g., 1.

[00288] The shape, the data layout and the data type of input tensor 1 and the output tensor are to be the same, in one example; otherwise, a general operand data exception is recognized.

[00289] The output tensor descriptor 2, input tensor descriptor 2, input tensor descriptor 3, function-specific-parameters 1-5 and function-specific-save-area-address fields are ignored, in one example.

[00290] Function Code 51: NNPA-SIGMOID

[00291] When the NNPA-SIGMOID function is specified, for each element of the input tensor described by tensor descriptor 1, the corresponding element in the output tensor described by the output tensor descriptor is the sigmoidal of that element.

[00292] In one example, if the specified data layout in any of the specified tensor descriptors does not specify a 4D-feature tensor (e.g., data layout = 0) or if the data type in any specified tensor descriptor does not specify NNP-data-type-1 (e.g., data type = 0), response code, e.g., 0010 hex or 0011 hex, respectively, is set in general register 0 and the instruction completes with condition code, e.g., 1.

[00293] The shape, the data layout and the data type of input tensor 1 and the output tensor are to be the same, in one example; otherwise, a general operand data exception is recognized.

[00294] The output tensor descriptor 2, input tensor descriptor 2, input tensor descriptor 3, function-specific-parameters 1-5 and function-specific-save-area-address fields are ignored, in one example.

[00295] Function Code 52: NNPA-SOFTMAX

[00296] When the NNPA-SOFTMAX function is specified, for each vector in dimension-1 of the input tensor 1, the corresponding vector in the output tensor is computed, as described below:

[00297] * The maximum value of the vector is computed.

[00298] * The summation of the exponentials of the difference between each element in dimension-1 of the vector and the maximum value computed above is computed. If both the element in dimension-1 of the input vector and the maximum value computed above are numeric values, and the difference is non-numeric, the result of the exponential for that element is forced to zero.

[00299] * For each element in the vector, an intermediate quotient is formed of the exponential of the difference between the element and the maximum value computed above divided by the summation computed above. An optional activation function is applied to this intermediate quotient to form the corresponding element in the output vector.

[00300] This process is repeated for, e.g., all dimension-4-index-size x dimension-3-index-size x dimension-2-index-size vectors in dimension-1.

[00301] In one example, a NNPA-SOFTMAX function-specific-parameter 1 controls the activation function. As an example, an ACT field (e.g., bits 28-31) of function-specific-parameter 1 specifies the activation function. Example activation functions include:

[00302]	<u>ACT</u>	<u>Activation Function</u>
---------	------------	----------------------------

[00303]	0	No activation function performed
---------	---	----------------------------------

[00304]	1	LOG
---------	---	-----

[00305]	2-15	Reserved
---------	------	----------

[00306] If a reserved value is specified for the ACT field, a response code of, e.g., F001 hex, is reported and the operation completes with condition code, e.g., 1.

[00307] In one example, if the specified data layout in any of the specified tensor descriptors does not specify a 4D-feature tensor (e.g., data layout = 0) or if the data type in any specified tensor descriptor does not specify NNP-data-type-1 (e.g., data type = 0), response code, e.g., 0010 hex or 0011 hex, respectively, is set in general register 0 and the instruction completes with condition code, e.g., 1.

[00308] In one example, if the dimension-3-index-size of the input tensor is not equal to one, a response code of, e.g., F000 hex is stored and the instruction completes with condition codes, e.g., 1.

[00309] The shape, the data layout and the data type of input tensor 1 and the output tensor are to be the same, in one example; otherwise, a general operand data exception is recognized.

[00310] The output tensor descriptor 2, input tensor descriptor 2 and input tensor descriptor 3 are ignored, in one example. Function-specific parameters 2-5 are to contain zeros, in one example.

[00311] An 8 K-byte function specific save area may be used by this function.

[00312] In one embodiment, when obtaining the vector in dimension-1, the elements may not be contiguous in memory depending on the specified data layout format. If all elements of a dimension-1 vector of the input tensor 1 contain the largest magnitude negative number representable in the specified data type, results may be less accurate.

[00313] Function Code 64: NNPA-BATCHNORM (Batch Normalization)

[00314] When the NNPA-BATCHNORM function is specified, for each vector in dimension-1 of the input 1 tensor, the corresponding vector in dimension-1 of the output tensor is computed by multiplying each element in the vector by the corresponding element in the dimension-1 vector that makes up the input 2 tensor. The full precision product is then added to the corresponding element in the dimension-1 vector that makes up the input 3 tensor and then rounding to the precision of the specified data type of the output tensor. This process is repeated for, e.g., all dimension-4-index-size x dimension-3-index-size x dimension-2-index-size vectors in dimension-1.

[00315] In one example, if the specified data layout in any of the specified tensor descriptors does not specify a 4D-feature tensor (e.g., data layout = 0) or if the data type in any specified tensor descriptor does not specify NNP-data-type-1 (e.g., data type = 0), response code, e.g., 0010 hex or 0011 hex, respectively, is set in general register 0 and the instruction completes with condition code, e.g., 1.

[00316] In one example, the following conditions are to be true, otherwise a general operand data exception is recognized:

[00317] * The shape and the data layout of input tensor 1 and the output tensor are to be the same.

[00318] * The data type of the input tensors and the output tensor are to be the same.

[00319] * The dimension-1-index-size of input tensors 1, 2, 3 and the output tensor are to be the same.

[00320] * The dimension 2, 3 and 4 index sizes of input tensors 2 and 3 are to be one.

[00321] The output tensor descriptor 2 function-specific-save-area-address fields are ignored, in one example. Function-specific-parameters 2-5 are to contain zeros, in one example.

[00322] Function Code 80: NNPA-MAXPOOL2D

Function Code 81: NNPA-AVGPOOL2D

[00323] When either the NNPA-MAXPOOL2D or the NNPA-AVGPOOL2D function is specified, input tensor 1, described by the input tensor 1 descriptor, is reduced by the specified operation to summarize windows of the input. The windows of the input are selected by moving a 2D sliding window over dimension indices 2 and 3. The summary of the window is an element in the output tensor. The sliding window dimensions are described by, e.g., function-specific-parameter 4 and function-specific-parameter 5. The amount that the sliding window moves over the input 1 tensor when computing adjacent output tensor elements is called the stride. The sliding window stride is specified by, e.g., function-specific-parameter 2 and function-specific-parameter 3. When the NNPA-MAXPOOL2D operation is specified, the Max operation defined below is performed on the window. When the NNPA-AVGPOOL2D operation is specified, the AVG operation defined below is performed on the window. If the specified padding type is Valid, all elements in the window are added to the collection used to compute the resulting output element. If the specified padding type is Same, depending on the location of the window, only a subset of elements from the window may be added to the collection used to compute the resulting output element.

[00324] In one example, a CollectElements operation adds an element to the collection of elements and increments the number of elements in the collection. Each time the window start position moves, the collection is emptied. It is unpredictable whether elements not required to perform the operations are accessed.

[00325] Max Operation: In one example, the maximum value of the collection of elements in the window is computed by comparing all elements in the collection to each other and returning the largest value.

[00326] Avg (Average) Operation: In one example, the average value of the collection of elements in the window is computed as the summation of all elements in the collection divided by the number of elements in the collection.

[00327] In one example, fields are allocated as follows:

[00328] * A pooling function-specific-parameter 1 controls the padding type. For instance, bits 29-31 of function-specific-parameter 1 include a PAD field that specifies the padding type. Example types include, for instance:

	<u>PAD</u>	<u>Padding Type</u>
[00330]	0	Valid
[00331]	1	Same
[00332]	2-7	Reserved

[00333] If a reserved value is specified for the PAD field, a response code of, e.g., F000 hex is reported and the operation completes with condition code, e.g., 1.

[00334] In one example, bit positions 0-28 of function-specific-parameter 1 are reserved and are to contain zeros.

[00335] * Function-specific-parameter 2 contains, e.g., a 32-bit unsigned binary integer that specifies the dimension-2-stride (D2S) which specifies the number of elements the sliding window moves in dimension 2.

[00336] * Function-specific-parameter 3 contains, e.g., a 32-bit unsigned binary integer that specifies the dimension-3-stride (D3S) which specifies the number of elements the sliding window moves in dimension 3.

[00337] * Function-specific-parameter 4 contains, e.g., a 32-bit unsigned binary integer that specifies the dimension-2-window-size (D2WS) which specifies the number of elements in dimension 2 the sliding window contains.

[00338] * Function-specific-parameter 5 contains, e.g., a 32-bit unsigned binary integer that specifies the dimension-3-window-size (D3WS) which specifies the number of elements in dimension 3 the sliding window contains.

[00339] In one example, the specified values in function-specific-parameters 2-5 are to be less than or equal to the maximum dimension index size, and the specified values in

function-specific-parameters 4-5 are to be greater than zero; otherwise, response code, e.g., 0012 hex is reported and the operation completes with condition code, e.g., 1.

[00340] If the dimension-2-stride and the dimension-3-stride are both zero and either the dimension-2-window size or the dimension-3-window size is greater than, e.g., 1024, response code, e.g., F001 hex is stored. If the dimension-2-stride and the dimension-3-stride are both greater than, e.g., zero and either the dimension-2-window-size or the dimension-3-window-size is greater than, e.g., 64, response code, e.g., F002 hex is stored. If the dimension-2-stride and the dimension-3-stride are both greater than, e.g., zero and either the dimension-2 stride or the dimension-3 stride is greater than, e.g., 30, response code, e.g., F003 hex is stored. If the dimension-2-stride and the dimension-3-stride are both greater than, e.g., zero and either the input tensor dimension-2-index-size or the input tensor dimension-3-index-size is greater than, e.g., 1024, response code, e.g., F004 hex is stored. For all of the above conditions, the instruction completes with condition code, e.g., 1.

[00341] In one example, if the specified data layout in any of the specified tensor descriptors does not specify a 4D-feature tensor (e.g., data-layout = 0) or if the data-type in any specified tensor descriptor does not specify NNP-data-type-1 (e.g., data-type = 0), response code, e.g., 0010 hex or 0011 hex, respectively, is set in general register 0 and the instruction completes with condition code, e.g., 1.

[00342] In one example, the following conditions are to be true, otherwise, a general operand data exception is recognized:

- [00343] * The dimension-4-index-sizes and dimension-1-index-sizes of the input tensor and the output tensor are to be the same.
- [00344] * The data layout and the data type of the input tensor and the output tensor are to be the same.
- [00345] * If the dimension-2-stride and the dimension-3-stride are both zero, the following additional conditions are to be true, in one example:
 - [00346] * The input tensor dimension-2-index-size is to be equal to the dimension-2-window size.
 - [00347] * The input tensor dimension-3-index-size of the input tensor is to be equal to the dimension-3-window-size.

- [00348] * The dimension-2-index-size and the dimension-3-index-size of the output tensor are to be one.
- [00349] * The specified padding is to be valid.
- [00350] * If either the dimension-2-stride or the dimension-3-stride is non-zero, then both strides are to be non-zero, in one example.
- [00351] * If the dimension-2-stride and the dimension-3-stride are both greater than zero, the following additional conditions are to be true, in one example:
- [00352] * When the specified padding is Valid, the dimension-2-window-size is to be less than or equal to the dimension-2-index-size of the input tensor.
- [00353] * When the specified padding is Valid, the dimension-3-window-size is to be less than or equal to the dimension-3-index-size of the input tensor.
- [00354] * When the specified padding is Same, the following relationships between the dimension-2-index-size and dimension-3-index size of the input and output tensors are to be satisfied (Pooling Same Padding):
- [00355]
$$O1D2IS = \left\lceil \frac{I1D2IS}{D2S} \right\rceil$$
- [00356]
$$O1D3IS = \left\lceil \frac{I1D3IS}{D3S} \right\rceil$$
- [00357] where:
- [00358] IxDyIS Dimension-y-index-size of the input tensor x defined in tensor descriptor x.
- [00359] OxDyIS Dimension-y-index-size of the output tensor x defined in tensor descriptor x.
- [00360] D2S Dimension-2-stride.
- [00361] D3S Dimension-3-stride.

[00362] * When the specified padding is Valid, the following relationships between the dimension-2-index-size and dimension-3-index-size of the input and output tensors are to be satisfied (Pooling Valid Padding):

$$[00363] \quad O1D2IS = \left\lceil \frac{(I1D2IS - D2WS + 1)}{D2S} \right\rceil$$

$$[00364] \quad O1D3IS = \left\lceil \frac{(I1D3IS - D3WS + 1)}{D3S} \right\rceil$$

[00365] where D2WS is dimension-2-window size and D3WS is dimension-3-window size.

[00366] The output tensor descriptor 2, input tensor descriptors 2 and 3, and function-specific-save-area-address field are ignored.

[00367] Function Code 96: NNPA-LSTMACT (Long Short-Term Memory Activation)

[00368] When the NNPA-LSTMACT function is specified, input tensor 1 (e.g., a reformatted, concatenated input tensor), described by the input tensor 1 descriptor, split into four sub-tensors for each dimension 4 index value, along with input tensor 2 (e.g., a reformatted, concatenated input tensor) described by the input tensor 2 descriptor, split into four sub-tensors for each dimension 4 index value, and input tensor 3 described by the input tensor 3 descriptor are the inputs to a LSTMACT operation. At the end of the LSTMACT operation, results are written to output tensor 1 described by the output tensor 1 (e.g., a reformatted, concatenated output tensor) descriptor and output tensor 2 (e.g., a reformatted, concatenated output tensor) described by the output tensor 2 descriptor.

[00369] In one example, if the specified data layout in any of the specified tensor descriptors does not specify a 4D-feature tensor (e.g., data-layout = 0) or if the data-type in any specified tensor descriptor does not specify NNP-data-type-1 (e.g., data-type = 0), response code 0010 hex or 0011 hex, respectively, is set in general register 0 and the instruction completes with condition code, e.g., 1.

[00370] In one embodiment, the following conditions are to be true, otherwise, a general operand data exception is recognized:

[00371] * The dimension-4-index-size for input tensor 3, and output tensors 1 and 2 are to be equal to, e.g., one.

- [00372] * The dimension-4-index-size for input tensor 1 and input tensor 2 are to be equal to, e.g., four.
- [00373] * The dimension-3-index-size for, e.g., all input tensors and the two output tensors are to be equal to, e.g., one.
- [00374] * The data layout and data type of, e.g., all input tensors and the two output tensors are to be the same.
- [00375] * The dimension-1-index-size of, e.g., all input tensors and the two output tensors are to be the same.
- [00376] * The dimension-2-index-size of, e.g., all input tensors and the two output tensors are to be the same.
- [00377] The function-specific-save-area address fields are ignored, in one example. Function-specific-parameters 1-5 are to contain zeros, in one example.
- [00378] Further details regarding one embodiment of a long short-term cell activation are described herein with reference, e.g., to FIG. 4A and FIGS. 5A-5B.
- [00379] Function Code 97: NNPA-GRUACT (Gated Recurrent Unit Activation)
- [00380] When the NNPA-GRUACT function is specified, input tensor 1 (e.g., a reformatted, concatenated input tensor), described by the input tensor 1 descriptor, split into three sub-tensors for each dimension 4 index value, along with input tensor 2 (e.g., a reformatted, concatenated input tensor) described by the input tensor 2 descriptor, split into three sub-tensors for each dimension 4 index value, and input tensor 3 described by the input tensor 3 descriptor are the inputs to a GRUACT operation. At the end of the GRUACT operation, the output tensor (e.g., a reformatted, concatenated output tensor) described by the output tensor descriptor is stored.
- [00381] In one example, if the specified data layout in any of the specified tensor descriptors does not specify a 4D-feature tensor (e.g., data-layout = 0) or if the data-type in any specified tensor descriptor does not specify NNP-data-type-1 (e.g., data-type = 0), response code, e.g., 0010 hex or 0011 hex, respectively, is set in general register 0 and the instruction completes with condition code, e.g., 1.

[00382] In one embodiment, the following conditions are to be true, otherwise, a general operand data exception is recognized:

- [00383] * The dimension-4-index-size of the output tensor and input tensor 3 are to be equal to, e.g., one.
- [00384] * The dimension-4-index-size for the input tensor 1 and input tensor 2 are to be equal to, e.g., three.
- [00385] * The dimension-3-index-size for, e.g., all input tensors and the output tensor are to be equal to, e.g., one.
- [00386] * The dimension-1-index-size of, e.g., all input tensors and the output tensor are to be the same.
- [00387] * The dimension-2-index-size of, e.g., all input tensors and the output tensor are to be the same.
- [00388] * The data layout and data type of, e.g., all input tensors and the output tensor are to be the same.

[00389] The output tensor descriptor 2 and function-specific save area address fields are ignored, in one example. Function-specific-parameters 2-5 are to contain zeros, in one example.

[00390] Further details regarding one embodiment of a gated recurrent unit cell activation are described herein with reference, e.g., to FIG. 4B.

[00391] Function Code 112: NNPA-CONVOLUTION

[00392] When the NNPA-CONVOLUTION function is specified, for each output element in the output tensor described by the output tensor 1 descriptor, a 3-dimensional input-1 window consisting of dimension indices 3, 2, and 1 is selected from input tensor 1, described by the input tensor 1 descriptor. A 3-dimensional input-2 window of the same size consisting of dimension indices 4, 3, and 2 is selected from tensor 2, described by the input tensor 2 descriptor. The elements in the input-1 window are multiplied by the corresponding elements in the input-2 window and all of the products are added together to create an initial summation. This initial summation is added to the corresponding element

of input tensor 3 to compute an intermediate summation value. The element of the output tensor is the result of the specified activation function performed on the intermediate summation. If no activation function is specified, the output element is equal to the intermediate summation.

[00393] If the specified padding type is Valid, all elements in the window are used to compute the resulting initial summation. If the specified padding type is Same, depending on the location of the window, some elements of the input-1 window may be implied zero, when computing the resulting initial summation.

[00394] It is unpredictable whether elements not required to perform the operation are accessed.

[00395] In one example, fields of a function-specific-parameter used by the convolution function are allocated, as follows:

[00396] * A NNPA-CONVOLUTION function-specific-parameter 1 controls the padding type and the activation function. In one example, bits 29-31 of function-specific-parameter 1 include a PAD field that specifies the padding type. Example types are below:

[00397]	<u>PAD</u>	<u>Padding Type</u>
[00398]	0	Valid
[00399]	1	Same
[00400]	2-7	Reserved

[00401] If a reserved value is specified for the PAD field, a response code of, e.g., F000 hex is reported and the operation completes with condition code, e.g., 1.

[00402] Further, in one example, bits 24-27 of the NNPA-CONVOLUTION function-specific-parameter 1 include an activation field that specifies activation functions. Example functions are below:

[00403]	<u>ACT</u>	<u>Activation Function</u>
---------	------------	----------------------------

- | | | |
|---------|------|----------------------------------|
| [00404] | 0 | No activation function performed |
| [00405] | 1 | RELU |
| [00406] | 2-15 | Reserved |
- [00407] If an activation function of RELU is specified, the resulting output element value is determined, as follows: if the intermediate summation value is less than or equal to zero, the corresponding element in the output tensor is zero; otherwise, the corresponding element in the output tensor is the minimum of the intermediate summation value and the clipping value specified in function-specific-parameter 4.
- [00408] If a reserved value is specified for the ACT field, a response code of, e.g., F001 hex is reported and the operation completes with condition code, e.g., 1.
- [00409] * Function-specific-parameter 2 contains, e.g., a 32-bit unsigned binary integer that specifies the dimension-2 (D2S) stride which specifies the number of elements the sliding window moves in dimension 2.
- [00410] * Function-specific-parameter 3 contains, e.g., a 32-bit unsigned binary integer that specifies the dimension-3 (D3S) stride which specifies the number of elements the sliding window moves in dimension 3.
- [00411] The specified values in function-specific-parameters 2-3 are to be less than the maximum dimension index size; otherwise a response code, e.g., 0012 hex is reported and the operation completes with condition code, e.g., 1.
- [00412] * Function-specific-parameter 4 defines the clipping value for the optional RELU operation. In one example, the clipping value is in bits 16-31 of function-specific-parameter 4.
- [00413] In one example, if the ACT field is zero, this field is ignored. If the ACT field specifies RELU, the clipping value is specified in NNP-data-type-1 format. A clipping value of zero indicates to use the maximum positive value; in other

words, no clipping is performed. If a non-zero is specified, a general operand data exception is recognized.

[00414] In one example, if the specified data layout in any of the specified tensor descriptors except input tensor 2 does not specify a 4D-feature tensor (e.g., data-layout = 0) or if the specified data layout in input tensor 2 does not specify a 4D-kernel tensor (e.g., data-layout = 1), response code, e.g., 0010 hex is set in general register 0 and the instruction completes with condition code, e.g., 1. In one example, if the data-type in any specified tensor descriptor does not specify NNP-data-type-1 (e.g., data-type = 0), response code, e.g., 0011 hex, is set in general register 0 and the instruction completes with condition code, e.g., 1.

[00415] If the dimension-2-stride and dimension-3-stride are both zero and the dimension-3-index size or the dimension-4-index size of input tensor 2 is greater than, e.g., 448, response code, e.g., F002 hex is stored. If the dimension-2-stride and the dimension-3-stride are both greater than zero and either the dimension-3-index size or the dimension-4-index-size of the input tensor 2 is greater than, e.g., 64, response code, e.g., F003 hex is stored and the operation completes with condition code, e.g., 1. If either the dimension 2 stride or the dimension 3 stride is greater than, e.g., 13, response code, e.g., F004 hex is stored and the operation completes with condition code, e.g., 1.

[00416] In one example, the following conditions are to be true, otherwise, a general operand data exception is recognized:

- [00417] * The data layout of input tensor 1, input tensor 3, and the output tensor are to be the same.
- [00418] * The data type of all input tensors and the output tensor are to be the same.
- [00419] * The dimension-2, dimension-3 and dimension-4 index sizes of the input 3 tensor are to be 1.
- [00420] * The dimension-4-index-size of the output tensor is to be equal to the dimension-4-index-size of the input 1 tensor.
- [00421] * The dimension-1-index-size of the output tensor is to be equal to the dimension-1 index size of the input 2 tensor and the dimension-1-index size of the input 3 tensor.

- [00422] * The dimension-1-index-size of the input 1 tensor is to be equal to the dimension-2 index size of the input 2 tensor.
- [00423] * If the dimension-2-stride and the dimension-3-stride are both zero, the following additional conditions are to be true, in one example:
- [00424] * The input 1 tensor dimension-2-index-size is to be equal to the dimension-3-index size of input 2 tensor.
- [00425] * The input 1 tensor dimension-3-index-size of the input tensor is to be equal to the dimension-4-index-size of input 2 tensor.
- [00426] * The dimension-2-index-size and the dimension-3-index-size of the output tensor are to be one.
- [00427] * The specified padding is to be Valid.
- [00428] * If either the dimension-2-stride or the dimension-3-stride is non-zero, then both strides are to be non-zero.
- [00429] * If the dimension-2-stride and the dimension-3-stride are both greater than zero, the following additional conditions are to be true, in one example:
- [00430] * When the specified padding is Valid, the dimension-2-index-size of the input 1 tensor is to be greater than or equal to the dimension-3-index-size of input tensor 2.
- [00431] * When the specified padding is Valid, the dimension-3-index-size of the input 1 tensor is to be greater than or equal to the dimension-4-index-size of the input 2 tensor.
- [00432] * When the specified padding is Same, the following relationships between the dimension-2-index-size and dimension-3-index-size of the input 1 tensor and output tensor is to be satisfied, in one example (Convolution Same Padding):

$$[00433] \quad O1D2IS = \left\lceil \frac{I1D2IS}{D2S} \right\rceil$$

$$[00434] \quad O1D3IS = \left\lceil \frac{I1D3IS}{D3S} \right\rceil$$

- [00435] where:
- [00436] O1D2IS Dimension-2-index-size of the output tensor.
- [00437] O1D3IS Dimension-3-index-size of the output tensor.
- [00438] I1D2IS Dimension-2-index-size of the input 1 tensor.
- [00439] I1D3IS Dimension-3-index-size of the input 1 tensor.
- [00440] D2S Dimension-2-Stride.
- [00441] D3S Dimension-3-Stride.
- [00442] * When the specified padding is Valid, the following relationships between the dimension-2-index-size and dimension-3-index-size of the input 1 tensor, dimension-3-index-size and dimension-4-index-size of the input 2 tensor and output tensor are to be satisfied, in one example (Convolution Valid Padding):

$$[00443] \quad O1D2IS = \left\lceil \frac{(I1D2IS - I2D3IS + 1)}{D2S} \right\rceil$$

$$[00444] \quad O1D3IS = \left\lceil \frac{(I1D3IS - I2D4IS + 1)}{D3S} \right\rceil$$

- [00445] where:
- [00446] O1D2IS Dimension-2-index-size of the output tensor.
- [00447] O1D3IS Dimension-3-index-size of the output tensor.
- [00448] I1D2IS Dimension-2-index-size of the input 1 tensor.
- [00449] I1D3IS Dimension-3-index-size of the input 1 tensor.
- [00450] I2D3IS Dimension-3-index-size of the input 2 tensor.
- [00451] I2D4IS Dimension-4-index-size of the input 2 tensor.
- [00452] D2S Dimension-2-Stride.
- [00453] D3S Dimension-3-Stride.

[00454] The output tensor descriptor 2 and the function-specific-save-area address fields are ignored, in one example. Function-specific-parameter 5 is to contain zeros, in one example.

[00455] Function Code 113: NNPA-MATMUL-OP (Matrix Multiplication Operation)

[00456] When the NNPA-MATMUL-OP function is specified, each element in the output tensor described by the output tensor descriptor is computed as described below, in one example:

[00457] * A dimension-1-vector is selected from the input tensor 1, described by the input tensor 1 descriptor, using the get-dimension-1-vector operation described below.

[00458] * A dimension-2-vector is selected from the input tensor 2, described by the input tensor 2 descriptor, using the get-dimension-2-vector operation described below.

[00459] * An intermediate dot product of the dimension-1-vector and the dimension-2-vector is computed using the dot product operation described below.

[00460] * An operation is performed on the intermediate dot product and the element of the input tensor 3, described by the input tensor 3 descriptor, with the same dimension index 4 and dimension index 1 values as the output tensor element. The resulting element is stored in the output tensor. A fused operation is determined by function-specific-parameter 1 and described below.

[00461] Get-dimension-1-vector Operation: For a specified output element, a dimension-1 vector is selected from the input-1 tensor where the input dimension-4-index is the output dimension-4-index, the input-dimension-3-index is the output dimension-3-index, and the input dimension-2-index is the output dimension-2-index.

[00462] Get-dimension-2-vector Operation: For a specified output element, a dimension-2 vector is selected from the input-2 tensor where the input dimension-4-index is the output dimension-4-index, the input-dimension-3-index is the output dimension-3-index, and the input dimension-1-index is the output dimension-1-index.

[00463] Dot Product Operation: The intermediate dot product of two vectors of the same size and data type is computed as the summation of products of each element in the input vector 1 and the corresponding element of the input vector 2.

[00464] Fused Operation: Function-specific-parameter 1 controls the operation performed on the intermediate dot product and the corresponding element from input tensor 3. In one example, a NNPA-MATMUL-OP function-specific-parameter 1 includes an operation field in, e.g., bits 24-31. The operation field specifies the operation performed. Example operations are indicated below:

[00465]	<u>OPERATION</u>	<u>Operation Type</u>
[00466]	0	Addition
[00467]	1	Compare if dot product is high
[00468]	2	Compare if dot product is not low
[00469]	3	Compare if dot product and element are equal
[00470]	4	Compare if dot product and element are not equal
[00471]	5	Compare if dot product is not high
[00472]	6	Compare if dot product is low

[00473] In one example, for an operation type of addition, the input tensor 3 element is added to the intermediate dot product. For operation types of comparison, the intermediate dot product is compared to the input tensor 3 element and if the comparison is true, the result is set to a value of, e.g., +1; otherwise, it is set to a value of, e.g., +0, in the data type specified for the output tensor.

[00474] In one example, all other values of the OPERATION field are reserved. If a reserved value is specified for the OPERATION field, a response code of, e.g., F000 hex, is reported and the operation completes with condition code, e.g., 1.

[00475] In one example, if the specified data layout in any of the specified tensor descriptors does not specify a 4D-feature tensor (e.g., data layout = 0) or if the data-type in any specified tensor descriptor does not specify NNP-data-type-1 (e.g., data-type = 0),

response code, e.g., 0010 hex or 0011 hex, respectively, is set in general register 0 and the instruction completes with condition code, e.g., 1.

[00476] In one embodiment, the following conditions are to be true, otherwise, a general operand data exception is recognized:

- [00477] * The dimension-4-index-size of all input tensors and the output tensor are to be the same.
- [00478] * The dimension-3-index-size of all input tensors and the output tensor are to be equal to one.
- [00479] * The dimension-2-index-size of the input tensor 3 is to be equal to one.
- [00480] * The dimension-2-index-size of input tensor 1 and the output tensor are to be the same.
- [00481] * The dimension-1-index-size of input tensor 1 and the dimension-2-index-size of input tensor 2 are to be the same.
- [00482] * The dimension-1-index-size of input tensor 2, input tensor 3, and the output tensor are to be the same.
- [00483] * The data layout and data type of all input tensors and the output tensor are to be the same.

[00484] In one embodiment, the output tensor descriptor 2 and function-specific-save-area-address fields are ignored. Function-specific-parameters 2-5 are to contain zeros, in an example.

[00485] Function Code 114: NNPA-MATMUL-OP-BCAST23 (Matrix Multiplication Operation – Broadcast 23)

[00486] When the NNPA-MATMUL-OP-BCAST23 function is specified, each element in the output tensor described by the output tensor descriptor is computed, as described below, in one example:

- [00487] * A dimension-1-vector is selected from the input tensor 1, described by the input tensor 1 descriptor, using the get-dimension-1-vector operation described below.
- [00488] * A dimension-2-vector is selected from the input tensor 2, described by the input tensor 2 descriptor, using the get-dimension-2-vector operation described below.
- [00489] * The dot product of a dimension-1-vector and a dimension-2-vector is computed using the dot product operation described below.
- [00490] * The element of the input tensor 3, described by the input tensor 3 descriptor, with the same dimension index 1 value as the output tensor element is added to the previously computed dot product and stored in the output tensor.

[00491] Get-dimension-1-vector Operation: For a specified output element, a dimension-1 vector is selected from the input-1 tensor where the input dimension-4-index is the output dimension-4-index, the input-dimension-3-index is the output dimension-3-index, and the input dimension-2-index is the output dimension-2-index.

[00492] Get-dimension-2-vector Operation: For a specified output element, a dimension-2 vector is selected from the input-2 tensor where the input dimension-4-index is one, the input-dimension-3-index is the output dimension-3-index, and the input dimension-1-index is the output dimension-1-index.

[00493] Dot Product Operation: The intermediate product of two vectors of the same size and data type is computed as the summation of products of each element in the input vector 1 and the corresponding element of the input vector 2.

[00494] In one example, if the specified data layout in any of the specified tensor descriptors does not specify a 4D-feature tensor (e.g., data layout = 0) or if the data-type in any specified tensor descriptor does not specify NNP-data-type-1 (e.g., data-type = 0), a response code, e.g., 0010 hex or 0011 hex, respectively, is set in general register 0 and the instruction completes with condition code, e.g., 1.

[00495] In one embodiment, the following conditions are to be true, otherwise, a general operand data exception is recognized:

- [00496] * The dimension-4-index-size of the input tensor 1 and the output tensor are to be the same.
- [00497] * The dimension-4-index-size of input tensor 2 and the input tensor 3 is to be equal to one.
- [00498] * The dimension-3-index-size of all input tensors and the output tensor is to be equal to one.
- [00499] * The dimension-2-index-size of input tensor 3 is to be equal to one.
- [00500] * The dimension-2-index-size of input tensor 1 and the output tensor is to be the same.
- [00501] * The dimension-1-index-size of input tensor 1 and the dimension-2-index-size of input tensor 2 are to be the same.
- [00502] * The dimension-1-index-size of input tensor 2, input tensor 3, and the output tensor are to be the same.
- [00503] * The data layout and data type of all input tensors and the output tensor are to be the same.
- [00504] In one embodiment, the output tensor descriptor 2 and function-specific-save-area-address fields are ignored. Function-specific-parameters 1-5 are to contain zeros, in one example.
- [00505] For the Neural Network Processing Assist instruction, in one embodiment, if the output tensor overlaps with any input tensor or the parameter block, results are unpredictable.
- [00506] A specification exception is recognized when execution of the Neural Network Processing Assist instruction is attempted and the parameter block is not designated on, e.g., a doubleword boundary, as an example.
- [00507] A general operand data exception is recognized when execution of the Neural Network Processing Assist instruction is attempted and there are, for instance, tensor descriptor inconsistencies.

[00508] Resulting Condition Codes for the Neural Network Processing Assist instruction include, for instance: 0 – Normal completion; 1 – Response code is set; 2 --; 3 – CPU-determined amount of data processed.

[00509] In one embodiment, the priority of execution for the Neural Network Processing Assist instruction includes, for instance:

[00510] 1.-7. Exceptions with the same priority as the priority of program interruption conditions for the general case.

[00511] 8.A Condition code 1 due to an unassigned or uninstalled function code specified.

[00512] 8.B Specification exception due to parameter block not designated on doubleword boundary.

[00513] 9. Access exceptions for an access to the parameter block.

[00514] 10. Condition code 1 due to specified format of the parameter block not supported by the model.

[00515] 11.A Condition code 1 due to the specified tensor data layouts are not supported.

[00516] 11.B General operand data exception due to differing data layouts between tensor descriptors.

[00517] 12.A Condition code 1 due to conditions other than those included in items 8.A, 10 and 11.A above and 12.B.1 below.

[00518] 12.B.1 Condition code 1 due to invalid output tensor data type for NNPA-RELU (Rectified Linear Unit) and NNPA-CONVOLUTION (other available functions not described herein).

[00519] 12.B.2 General operand data exception for invalid value for NNPA-RELU function-specific-parameter 1 and NNPA-CONVOLUTION function-specific-parameter 4.

[00520] 13.A Access exceptions for an access to the output tensor.

[00521] 13.B Access exceptions for an access to the input tensors.

[00522] 13.C Access exceptions for an access to the function-specific-save-area.

[00523] 14. Condition code 0.

[00524] As described herein, a single instruction (e.g., the Neural Network Processing Assist instruction) is configured to perform a plurality of functions, including a query function and a plurality of non-query functions. Each non-query function may operate on tensors, such as 4D-tensors (or tensors of other sizes). To facilitate processing that uses tensors, in accordance with one or more aspects of the present invention, tensors are reformatted into a plurality of, e.g., 2D-tensors having certain characteristics to improve processing. For instance, a reformatted tensor has easily calculatable addresses and may be loaded/stored in one operation, increasing bandwidth and improving system performance. This is a result of, for instance, starting a tensor on a memory boundary and having fixed dimensions (made possible using padding).

[00525] In one example, the reformatting of the tensors is performed based on a processor (e.g., general processor 104) obtaining the Neural Network Processing Assist instruction that specifies a non-query function. The tensor(s) that are specified are reformatted using, e.g., the tensor descriptor information provided in the parameter block (e.g., tensor descriptor 660, 665 of FIG. 6G). Address information relating to the reformatted tensor(s) are provided to the special-purpose processor (e.g., neural network processor 105) for use in performing the function specified by the instruction.

[00526] In one example, an instruction (e.g., the Neural Network Processing Assist instruction) implements a recurrent neural network cell activation (e.g., a long short-term memory cell activation, a gated recurrent unit cell activation and/or other cell activations), in which the input and/or output data uses a concatenated data layout in memory of tensors to prevent reformatting of data between operations. As an example, for the concatenation of input data, weight tensors are independently 2D-transformed and concatenated within timesteps prior to a multiplication operation. A single invocation of the instruction computes all multiplications of the input feature across timesteps at once to intermediate results. The intermediate results are provided in memory address contiguous tensors to compute the activation.

[00527] For concatenation of the output data, the result tensor includes a concatenation of 2D-reformatted results of a timestep. Each timestep result tensor includes memory address contiguous tensors of the complete results of the recurrent neural network computation. The result tensor of a timestep can be directly used in the computation of the next timestep without data manipulation or copy operations.

[00528] In one or more further aspects, individual activations and operations are combined in one instruction executed at once in the accelerator. Recurrent neural networks rely, in one example, on long short-term memory networks or gated recurrent unit networks. For each timestep (one operation after another), a number of activations (e.g., sigmoid, tanh) and operations (e.g., addition, subtraction and/or multiplication) are applied to a hidden state (e.g., previously learned), input state and cell state. Calling an accelerator (e.g., neural network processor 105) for each of these steps is detrimental for the overall performance of the processor and/or system due to, at least, the start-up time of the accelerator. Significant speed-up is achieved based on the individual activations and operations being combined in one instruction executed at once in the accelerator, in accordance with an aspect of the present invention. A single instruction is implemented, in accordance with an aspect of the present invention, combining the individual activations and combination functions. Thus, there is only one invocation; intermediate computation data is stored in the accelerator instead of being written back to memory; the SIMD width and the pipelined nature of the accelerator can be used to do more computation in parallel with less cycle per computer; and higher precision is used for intermediate results leading to enhanced accuracy and higher stability for long short-term memory and/or gated recurrent unit operations. The combining of, e.g., multiply and add operations provides a higher accuracy without a loss of precision on intermediate results. Further, higher numeric accuracy may be achieved by saving intermediate computations in the accelerator in higher precisions.

[00529] Further, in accordance with one or more aspects of the present invention, the matrix multiplication operations used to provide a concatenated result tensor that is input to a cell activation are separate from the cell activations, reducing complexity of single operations and allowing reuse of the basic blocks for other recurrent neural networks. The architected instruction provides spatially close input and output data sources to reduce address translations.

[00530] In accordance with one or more aspects, activations of inputs in an internal format are computed, and the computations are combined, producing one or more outputs in an input numerical format. As an example, the internal format is a model-dependent format for, e.g., the neural network processor. In one example, the internal format used may have a different numerical precision than the input/output numerical format to increase accuracy or reduce compute time and power.

[00531] Further, in accordance with one or more aspects, multiple activations are encapsulated in one instruction. The instruction provides modularity without breaking the activation in very small chunks. Further, the instruction uses concatenated input and output formats for the activations, providing savings in processing time and increasing processing speed.

[00532] One or more aspects of the present invention are inextricably tied to computer technology and facilitate processing within a computer, improving performance thereof. The reformatted concatenated tensors and/or instruction that defines and/or uses such tensors may be used in many technical fields, such as in computer processing, artificial intelligence, recurrent neural networks, medical processing, engineering, automotive technologies, manufacturing, etc. By using reformatted concatenated tensors, as described herein, certain optimizations are provided including optimizations in performing complex calculations used in various technical fields, improving those fields by increasing bandwidth, providing efficiencies and/or reducing execution time.

[00533] Further details of one embodiment of facilitating processing within a computing environment, as it relates to one or more aspects of the present invention, are described with reference to FIGS. 10A and 10B.

[00534] Referring to FIG. 10A, an instruction to perform a recurrent neural network cell activation is executed 1000. The executing includes, for instance, performing a plurality of operations of the recurrent neural network cell activation to provide a result of the recurrent neural network cell activation 1002. As an example, the plurality of operations is performed in a single invocation of the instruction 1004.

[00535] Using a single invocation of the instruction to perform a plurality of operations reduces complexity, reduces use of system resources and increases system performance.

[00536] In one example, the plurality of operations includes one or more sigmoid functions and one or more tangent functions 1006. In one example, the plurality of operations includes tensor element-wise add and tensor element-wise multiplication operations 1008.

[00537] As an example, the plurality of operations includes one or more sigmoid functions, one or more tangent functions, one or more tensor element-wise add operations and one or more tensor element-wise multiplication operations 1010.

[00538] In one example, one or more inputs to the instruction include one or more concatenated tensors 1012. A concatenated tensor may be directly used by an instruction executing on, e.g., an accelerator performing a cell activation of a recurrent neural network. The concatenated tensor may be accessed in one operation saving processing time and increasing processing speed. Further, there are fewer tensor pointers to be managed and there is a reduction in the copying or reorganizing of tensor data between invocations of the accelerator, improving processing speed.

[00539] In one example, referring to FIG. 10B, the result is an output tensor 1014, and the output tensor is an input to another invocation of the instruction 1016, as an example.

[00540] As examples, the recurrent neural network cell activation includes a long short-term memory cell activation 1020 or the recurrent neural network cell activation includes a gated recurrent unit cell activation 1022.

[00541] In one example, the performing the plurality of operations of the recurrent neural network cell activation is performed by an accelerator and produces intermediate computation data 1024. The intermediate computation data is stored in the accelerator, as an example 1026.

[00542] In one example, the performing the plurality of operations includes performing the plurality of operations on spatially close input data 1028.

[00543] Other variations and embodiments are possible.

[00544] Aspects of the present invention may be used by many types of computing environments. Another example of a computing environment to incorporate and use one or more aspects of the present invention is described with reference to FIG. 11A. As an example, the computing environment of FIG. 11A is based on the z/Architecture®

instruction set architecture offered by International Business Machines Corporation, Armonk, New York. The z/Architecture instruction set architecture, however, is only one example architecture. Again, the computing environment may be based on other architectures, including, but not limited to, the Intel® x86 architectures, other architectures of International Business Machines Corporation, and/or architectures of other companies. Intel is a trademark or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

[00545] In one example, a computing environment 10 includes a central electronics complex (CEC) 11. Central electronics complex 11 includes a plurality of components, such as, for instance, a memory 12 (a.k.a., system memory, main memory, main storage, central storage, storage) coupled to one or more processors, such as one or more general-purpose processors (a.k.a., central processing units (CPUs) 13) and one or more special-purpose processors (e.g., neural network processor 31), and to an input/output (I/O) subsystem 14.

[00546] As examples, the one or more special-purpose processors may be separate from the one or more general-purpose processors and/or at least one special-purpose processor may be embedded within at least one general-purpose processor. Other variations are also possible.

[00547] I/O subsystem 14 can be a part of the central electronics complex or separate therefrom. It directs the flow of information between main storage 12 and input/output control units 15 and input/output (I/O) devices 16 coupled to the central electronics complex.

[00548] Many types of I/O devices may be used. One particular type is a data storage device 17. Data storage device 17 can store one or more programs 18, one or more computer readable program instructions 19, and/or data, etc. The computer readable program instructions can be configured to carry out functions of embodiments of aspects of the invention.

[00549] Central electronics complex 11 can include and/or be coupled to removable/non-removable, volatile/non-volatile computer system storage media. For example, it can include and/or be coupled to a non-removable, non-volatile magnetic media (typically called a "hard drive"), a magnetic disk drive for reading from and writing to a removable, non-

volatile magnetic disk (e.g., a "floppy disk"), and/or an optical disk drive for reading from or writing to a removable, non-volatile optical disk, such as a CD-ROM, DVD-ROM or other optical media. It should be understood that other hardware and/or software components could be used in conjunction with central electronics complex 11. Examples include, but are not limited to: microcode or millicode, device drivers, redundant processing units, external disk drive arrays, RAID systems, tape drives, and data archival storage systems, etc.

[00550] Further, central electronics complex 11 can be operational with numerous other general-purpose or special-purpose computing system environments or configurations. Examples of well-known computing systems, environments, and/or configurations that may be suitable for use with central electronics complex 11 include, but are not limited to, personal computer (PC) systems, server computer systems, thin clients, thick clients, handheld or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputer systems, mainframe computer systems, and distributed cloud computing environments that include any of the above systems or devices, and the like.

[00551] Central electronics complex 11 provides in one or more embodiments logical partitioning and/or virtualization support. In one embodiment, as shown in FIG. 11B, memory 12 includes, for example, one or more logical partitions 20, a hypervisor 21 that manages the logical partitions, and processor firmware 22. One example of hypervisor 21 is the Processor Resource/System Manager (PR/SMTM), offered by International Business Machines Corporation, Armonk, New York. PR/SM is a trademark or registered trademark of International Business Machines Corporation in at least one jurisdiction.

[00552] Each logical partition 20 is capable of functioning as a separate system. That is, each logical partition can be independently reset, run a guest operating system 23 such as the z/OS[®] operating system, offered by International Business Machines Corporation, Armonk, New York, or other control code 24, such as coupling facility control code (CFCC), and operate with different programs 25. An operating system or application program running in a logical partition appears to have access to a full and complete system, but in reality, only a portion of it is available. Although the z/OS operating system is offered as an example, other operating systems offered by International Business Machines

Corporation and/or other companies may be used in accordance with one or more aspects of the present invention.

[00553] Memory 12 is coupled to, e.g., CPUs 13 (FIG. 11A), which are physical processor resources that can be allocated to the logical partitions. For instance, a logical partition 20 may include one or more logical processors, each of which represents all or a share of a physical processor resource 13 that can be dynamically allocated to the logical partition.

[00554] In yet a further embodiment, the central electronics complex provides virtual machine support (either with or without logical partitioning support). As shown in FIG. 11C, memory 12 of central electronics complex 11 includes, for example, one or more virtual machines 26, a virtual machine manager, such as a hypervisor 27, that manages the virtual machines, and processor firmware 28. One example of hypervisor 27 is the z/VM[®] hypervisor, offered by International Business Machines Corporation, Armonk, New York. The hypervisor is sometimes referred to as a host. z/VM is a trademark or registered trademark of International Business Machines Corporation in at least one jurisdiction.

[00555] The virtual machine support of the central electronics complex provides the ability to operate large numbers of virtual machines 26, each capable of operating with different programs 29 and running a guest operating system 30, such as the Linux[®] operating system. Each virtual machine 26 is capable of functioning as a separate system. That is, each virtual machine can be independently reset, run a guest operating system, and operate with different programs. An operating system or application program running in a virtual machine appears to have access to a full and complete system, but in reality, only a portion of it is available. Although z/VM and Linux are offered as examples, other virtual machine managers and/or operating systems may be used in accordance with one or more aspects of the present invention. The registered trademark Linux[®] is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

[00556] Another embodiment of a computing environment to incorporate and use one or more aspects of the present invention is described with reference to FIG. 12A. In this example, a computing environment 36 includes, for instance, a native central processing unit (CPU) 37, a memory 38, and one or more input/output devices and/or interfaces 39 coupled to one another via, for example, one or more buses 40 and/or other connections. As

examples, computing environment 36 may include a PowerPC[®] processor offered by International Business Machines Corporation, Armonk, New York; an HP Superdome with Intel[®] Itanium[®] II processors offered by Hewlett Packard Co., Palo Alto, California; and/or other machines based on architectures offered by International Business Machines Corporation, Hewlett Packard, Intel Corporation, Oracle, and/or others. PowerPC is a trademark or registered trademark of International Business Machines Corporation in at least one jurisdiction. Itanium is a trademark or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

[00557] Native central processing unit 37 includes one or more native registers 41, such as one or more general purpose registers and/or one or more special purpose registers used during processing within the environment. These registers include information that represents the state of the environment at any particular point in time.

[00558] Moreover, native central processing unit 37 executes instructions and code that are stored in memory 38. In one particular example, the central processing unit executes emulator code 42 stored in memory 38. This code enables the computing environment configured in one architecture to emulate another architecture. For instance, emulator code 42 allows machines based on architectures other than the z/Architecture instruction set architecture, such as PowerPC processors, HP Superdome servers or others, to emulate the z/Architecture instruction set architecture and to execute software and instructions developed based on the z/Architecture instruction set architecture.

[00559] Further details relating to emulator code 42 are described with reference to FIG. 12B. Guest instructions 43 stored in memory 38 comprise software instructions (e.g., correlating to machine instructions) that were developed to be executed in an architecture other than that of native CPU 37. For example, guest instructions 43 may have been designed to execute on a processor based on the z/Architecture instruction set architecture, but instead, are being emulated on native CPU 37, which may be, for example, an Intel Itanium II processor. In one example, emulator code 42 includes an instruction fetching routine 44 to obtain one or more guest instructions 43 from memory 38, and to optionally provide local buffering for the instructions obtained. It also includes an instruction translation routine 45 to determine the type of guest instruction that has been obtained and to translate the guest instruction into one or more corresponding native instructions 46. This

translation includes, for instance, identifying the function to be performed by the guest instruction and choosing the native instruction(s) to perform that function.

[00560] Further, emulator code 42 includes an emulation control routine 47 to cause the native instructions to be executed. Emulation control routine 47 may cause native CPU 37 to execute a routine of native instructions that emulate one or more previously obtained guest instructions and, at the conclusion of such execution, return control to the instruction fetch routine to emulate the obtaining of the next guest instruction or a group of guest instructions. Execution of the native instructions 46 may include loading data into a register from memory 38; storing data back to memory from a register; or performing some type of arithmetic or logic operation, as determined by the translation routine.

[00561] Each routine is, for instance, implemented in software, which is stored in memory and executed by native central processing unit 37. In other examples, one or more of the routines or operations are implemented in firmware, hardware, software or some combination thereof. The registers of the emulated processor may be emulated using registers 41 of the native CPU or by using locations in memory 38. In embodiments, guest instructions 43, native instructions 46 and emulator code 42 may reside in the same memory or may be disbursed among different memory devices.

[00562] An instruction that may be emulated includes the Neural Network Assist Processing instruction described herein, in accordance with one or more aspects of the present invention. Further, other instructions and/or one or more aspects of tensor processing (including, but not limited to, defining, generating, reformatting and/or concatenating tensors) may be emulated, in accordance with one or more aspects of the present invention.

[00563] The computing environments described above are only examples of computing environments that can be used. Other environments, including but not limited to, non-partitioned environments, partitioned environments, cloud environments and/or emulated environments, may be used; embodiments are not limited to any one environment. Although various examples of computing environments are described herein, one or more aspects of the present invention may be used with many types of environments. The computing environments provided herein are only examples.

[00564] Each computing environment is capable of being configured to include one or more aspects of the present invention.

[00565] One or more aspects may relate to cloud computing.

[00566] It is to be understood that although this disclosure includes a detailed description on cloud computing, implementation of the teachings recited herein are not limited to a cloud computing environment. Rather, embodiments of the present invention are capable of being implemented in conjunction with any other type of computing environment now known or later developed.

[00567] Cloud computing is a model of service delivery for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, network bandwidth, servers, processing, memory, storage, applications, virtual machines, and services) that can be rapidly provisioned and released with minimal management effort or interaction with a provider of the service. This cloud model may include at least five characteristics, at least three service models, and at least four deployment models.

[00568] Characteristics are as follows:

[00569] On-demand self-service: a cloud consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with the service's provider.

[00570] Broad network access: capabilities are available over a network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops, and PDAs).

[00571] Resource pooling: the provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to demand. There is a sense of location independence in that the consumer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter).

[00572] Rapid elasticity: capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and rapidly released to quickly scale in. To the

consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.

[00573] Measured service: cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.

[00574] Service Models are as follows:

[00575] Software as a Service (SaaS): the capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a web browser (e.g., web-based e-mail). The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

[00576] Platform as a Service (PaaS): the capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including networks, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations.

[00577] Infrastructure as a Service (IaaS): the capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls).

[00578] Deployment Models are as follows:

[00579] Private cloud: the cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on-premises or off-premises.

[00580] Community cloud: the cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be managed by the organizations or a third party and may exist on-premises or off-premises.

[00581] Public cloud: the cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.

[00582] Hybrid cloud: the cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds).

[00583] A cloud computing environment is service oriented with a focus on statelessness, low coupling, modularity, and semantic interoperability. At the heart of cloud computing is an infrastructure that includes a network of interconnected nodes.

[00584] Referring now to FIG. 13, illustrative cloud computing environment 50 is depicted. As shown, cloud computing environment 50 includes one or more cloud computing nodes 52 with which local computing devices used by cloud consumers, such as, for example, personal digital assistant (PDA) or cellular telephone 54A, desktop computer 54B, laptop computer 54C, and/or automobile computer system 54N may communicate. Nodes 52 may communicate with one another. They may be grouped (not shown) physically or virtually, in one or more networks, such as Private, Community, Public, or Hybrid clouds as described hereinabove, or a combination thereof. This allows cloud computing environment 50 to offer infrastructure, platforms and/or software as services for which a cloud consumer does not need to maintain resources on a local computing device. It is understood that the types of computing devices 54A-N shown in FIG. 13 are intended to be illustrative only and that computing nodes 52 and cloud computing environment 50 can communicate with any type of computerized device over any type of network and/or network addressable connection (e.g., using a web browser).

[00585] Referring now to FIG. 14, a set of functional abstraction layers provided by cloud computing environment 50 (FIG. 13) is shown. It should be understood in advance that the components, layers, and functions shown in FIG. 14 are intended to be illustrative only and embodiments of the invention are not limited thereto. As depicted, the following layers and corresponding functions are provided:

[00586] Hardware and software layer 60 includes hardware and software components. Examples of hardware components include: mainframes 61; RISC (Reduced Instruction Set Computer) architecture based servers 62; servers 63; blade servers 64; storage devices 65; and networks and networking components 66. In some embodiments, software components include network application server software 67 and database software 68.

[00587] Virtualization layer 70 provides an abstraction layer from which the following examples of virtual entities may be provided: virtual servers 71; virtual storage 72; virtual networks 73, including virtual private networks; virtual applications and operating systems 74; and virtual clients 75.

[00588] In one example, management layer 80 may provide the functions described below. Resource provisioning 81 provides dynamic procurement of computing resources and other resources that are utilized to perform tasks within the cloud computing environment. Metering and Pricing 82 provide cost tracking as resources are utilized within the cloud computing environment, and billing or invoicing for consumption of these resources. In one example, these resources may include application software licenses. Security provides identity verification for cloud consumers and tasks, as well as protection for data and other resources. User portal 83 provides access to the cloud computing environment for consumers and system administrators. Service level management 84 provides cloud computing resource allocation and management such that required service levels are met. Service Level Agreement (SLA) planning and fulfillment 85 provide pre-arrangement for, and procurement of, cloud computing resources for which a future requirement is anticipated in accordance with an SLA.

[00589] Workloads layer 90 provides examples of functionality for which the cloud computing environment may be utilized. Examples of workloads and functions which may be provided from this layer include: mapping and navigation 91; software development and lifecycle management 92; virtual classroom education delivery 93; data analytics processing 94; transaction processing 95; and tensor and/or neural network assist processing 96.

[00590] Aspects of the present invention may be a system, a method, and/or a computer program product at any possible technical detail level of integration. The computer program product may include a computer readable storage medium (or media) having computer readable program instructions thereon for causing a processor to carry out aspects of the present invention.

[00591] The computer readable storage medium can be a tangible device that can retain and store instructions for use by an instruction execution device. The computer readable storage medium may be, for example, but is not limited to, an electronic storage device, a magnetic storage device, an optical storage device, an electromagnetic storage device, a semiconductor storage device, or any suitable combination of the foregoing. A non-exhaustive list of more specific examples of the computer readable storage medium includes the following: a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), a static random access memory (SRAM), a portable compact disc read-only memory (CD-ROM), a digital versatile disk (DVD), a memory stick, a floppy disk, a mechanically encoded device such as punch-cards or raised structures in a groove having instructions recorded thereon, and any suitable combination of the foregoing. A computer readable storage medium, as used herein, is not to be construed as being transitory signals per se, such as radio waves or other freely propagating electromagnetic waves, electromagnetic waves propagating through a waveguide or other transmission media (e.g., light pulses passing through a fiber-optic cable), or electrical signals transmitted through a wire.

[00592] Computer readable program instructions described herein can be downloaded to respective computing/processing devices from a computer readable storage medium or to an external computer or external storage device via a network, for example, the Internet, a local area network, a wide area network and/or a wireless network. The network may comprise copper transmission cables, optical transmission fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers. A network adapter card or network interface in each computing/processing device receives computer readable program instructions from the network and forwards the computer readable program instructions for storage in a computer readable storage medium within the respective computing/processing device.

[00593] Computer readable program instructions for carrying out operations of the present invention may be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, configuration data for integrated circuitry, or either source code or object code written in any combination of one or more programming languages, including an object oriented programming language such as Smalltalk, C++, or the like, and procedural programming languages, such as the "C" programming language or similar programming languages. The computer readable program instructions may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider). In some embodiments, electronic circuitry including, for example, programmable logic circuitry, field-programmable gate arrays (FPGA), or programmable logic arrays (PLA) may execute the computer readable program instructions by utilizing state information of the computer readable program instructions to personalize the electronic circuitry, in order to perform aspects of the present invention.

[00594] Aspects of the present invention are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems), and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions.

[00595] These computer readable program instructions may be provided to a processor of a computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks. These computer readable program instructions may also be stored in a computer readable storage medium that can direct a computer, a programmable data processing apparatus, and/or other devices to function in a particular manner, such that the computer readable storage medium having

instructions stored therein comprises an article of manufacture including instructions which implement aspects of the function/act specified in the flowchart and/or block diagram block or blocks.

[00596] The computer readable program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other device to cause a series of operational steps to be performed on the computer, other programmable apparatus or other device to produce a computer implemented process, such that the instructions which execute on the computer, other programmable apparatus, or other device implement the functions/acts specified in the flowchart and/or block diagram block or blocks.

[00597] The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of instructions, which comprises one or more executable instructions for implementing the specified logical function(s). In some alternative implementations, the functions noted in the blocks may occur out of the order noted in the Figures. For example, two blocks shown in succession may, in fact, be accomplished as one step, executed concurrently, substantially concurrently, in a partially or wholly temporally overlapping manner, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts or carry out combinations of special purpose hardware and computer instructions.

[00598] In addition to the above, one or more aspects may be provided, offered, deployed, managed, serviced, etc. by a service provider who offers management of customer environments. For instance, the service provider can create, maintain, support, etc. computer code and/or a computer infrastructure that performs one or more aspects for one or more customers. In return, the service provider may receive payment from the customer under a subscription and/or fee agreement, as examples. Additionally or alternatively, the service provider may receive payment from the sale of advertising content to one or more third parties.

[00599] In one aspect, an application may be deployed for performing one or more embodiments. As one example, the deploying of an application comprises providing computer infrastructure operable to perform one or more embodiments.

[00600] As a further aspect, a computing infrastructure may be deployed comprising integrating computer readable code into a computing system, in which the code in combination with the computing system is capable of performing one or more embodiments.

[00601] As yet a further aspect, a process for integrating computing infrastructure comprising integrating computer readable code into a computer system may be provided. The computer system comprises a computer readable medium, in which the computer medium comprises one or more embodiments. The code in combination with the computer system is capable of performing one or more embodiments.

[00602] Although various embodiments are described above, these are only examples. For instance, computing environments of other architectures can be used to incorporate and/or use one or more aspects. Further, different instructions or operations may be used. Additionally, different types of registers and/or different registers may be used. Further, other data formats, data layouts and/or data sizes may be supported. In one or more embodiments, one or more general-purpose processors, one or more special-purpose processor or a combination of general-purpose and special-purpose processors may be used. Many variations are possible.

[00603] Various aspects are described herein. Further, many variations are possible without departing from a spirit of aspects of the present invention. It should be noted that, unless otherwise inconsistent, each aspect or feature described herein, and variants thereof, may be combinable with any other aspect or feature.

[00604] Further, other types of computing environments can benefit and be used. As an example, a data processing system suitable for storing and/or executing program code is usable that includes at least two processors coupled directly or indirectly to memory elements through a system bus. The memory elements include, for instance, local memory employed during actual execution of the program code, bulk storage, and cache memory which provide temporary storage of at least some program code in order to reduce the number of times code must be retrieved from bulk storage during execution.

[00605] Input/Output or I/O devices (including, but not limited to, keyboards, displays, pointing devices, DASD, tape, CDs, DVDs, thumb drives and other memory media, etc.) can be coupled to the system either directly or through intervening I/O controllers. Network adapters may also be coupled to the system to enable the data processing system to become coupled to other data processing systems or remote printers or storage devices through intervening private or public networks. Modems, cable modems, and Ethernet cards are just a few of the available types of network adapters.

[00606] The terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting. As used herein, the singular forms "a", "an" and "the" are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will be further understood that the terms "comprises" and/or "comprising", when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components and/or groups thereof.

[00607] The corresponding structures, materials, acts, and equivalents of all means or step plus function elements in the claims below, if any, are intended to include any structure, material, or act for performing the function in combination with other claimed elements as specifically claimed. The description of one or more embodiments has been presented for purposes of illustration and description, but is not intended to be exhaustive or limited to in the form disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art. The embodiment was chosen and described in order to best explain various aspects and the practical application, and to enable others of ordinary skill in the art to understand various embodiments with various modifications as are suited to the particular use contemplated.

CLAIMS

1. A computer program product for facilitating processing within a computing environment, the computer program product comprising:

one or more computer readable storage media and program instructions collectively stored on the one or more computer readable storage media to perform a method comprising:

executing an instruction to perform a recurrent neural network cell activation, the executing comprising:

performing a plurality of operations of the recurrent neural network cell activation to provide a result of the recurrent neural network cell activation, the plurality of operations performed in a single invocation of the instruction.

2. The computer program product according to the preceding claim, wherein the plurality of operations includes one or more sigmoid functions and one or more tangent functions.

3. The computer program product according to any of the preceding claims, wherein the plurality of operations includes tensor element-wise add and tensor element-wise multiplication operations.

4. The computer program product according to any of the preceding claims, wherein the plurality of operations includes one or more sigmoid functions, one or more tangent functions, one or more tensor element-wise add operations and one or more tensor element-wise multiplication operations.

5. The computer program product according to any of the preceding claims, wherein one or more inputs to the instruction include one or more concatenated tensors.

6. The computer program product according to any of the preceding claims, wherein the result is an output tensor, the output tensor being an input to another invocation of the instruction.

7. The computer program product according to any of the preceding claims, wherein the recurrent neural network cell activation comprises a long short-term memory cell activation.

8. The computer program product according to any of the preceding claims, wherein the recurrent neural network cell activation comprises a gated recurrent unit cell activation.

9. The computer program product according to any of the preceding claims, wherein the performing the plurality of operations of the recurrent neural network cell activation is performed by an accelerator and produces intermediate computation data, and wherein the method further comprises storing the intermediate computation data in the accelerator.

10. The computer program product according to any of the preceding claims, wherein the performing the plurality of operations includes performing the plurality of operations on spatially close input data.

11. A computer system for facilitating processing within a computing environment, the computer system comprising:

a memory; and

at least one processor in communication with the memory, wherein the computer system is configured to perform a method, said method comprising:

executing an instruction to perform a recurrent neural network cell activation, the executing comprising:

performing a plurality of operations of the recurrent neural network cell activation to provide a result of the recurrent neural network cell activation, the plurality of operations performed in a single invocation of the instruction.

12. The computer system according to the preceding claim, wherein the plurality of operations includes one or more sigmoid functions, one or more tangent functions, one or more tensor element-wise add operations and one or more tensor element-wise multiplication operations.

13. The computer system according to any of the two preceding claims, wherein one or more inputs to the instruction include one or more concatenated tensors.

14. The computer system according to any of the three preceding claims, wherein the recurrent neural network cell activation comprises a long short-term memory cell activation or a gated recurrent unit cell activation.

15. The computer system according to any of the four preceding claims, wherein the performing the plurality of operations of the recurrent neural network cell activation is performed by an accelerator and produces intermediate computation data, and wherein the method further comprises storing the intermediate computation data in the accelerator.

16. A computer-implemented method of facilitating processing within a computing environment, the computer-implemented method comprising:

executing an instruction to perform a recurrent neural network cell activation, the executing comprising:

performing a plurality of operations of the recurrent neural network cell activation to provide a result of the recurrent neural network cell activation, the plurality of operations performed in a single invocation of the instruction.

17. The computer-implemented method according to the preceding claim, wherein the plurality of operations includes one or more sigmoid functions, one or more tangent functions, one or more tensor element-wise add operations and one or more tensor element-wise multiplication operations.

18. The computer-implemented method according to any of the two preceding claims, wherein one or more inputs to the instruction include one or more concatenated tensors.

19. The computer-implemented method according to any of the three preceding claims, wherein the recurrent neural network cell activation comprises a long short-term memory cell activation or a gated recurrent unit cell activation.

20. The computer-implemented method according to any of the four preceding claims, wherein the performing the plurality of operations of the recurrent neural network cell activation is performed by an accelerator and produces intermediate computation data, and further comprising storing the intermediate computation data in the accelerator.