



(19) **United States**

(12) **Patent Application Publication**

Ross et al.

(10) **Pub. No.: US 2020/0057942 A1**

(43) **Pub. Date: Feb. 20, 2020**

(54) **NEURAL NETWORK PROCESSOR**

Publication Classification

(71) Applicant: **Google LLC**, Mountain View, CA (US)

(51) **Int. Cl.**
G06N 3/08 (2006.01)
G06N 5/04 (2006.01)
G06N 3/063 (2006.01)
G06F 15/80 (2006.01)

(72) Inventors: **Jonathan Ross**, Mountain View, CA (US); **Norman Paul Joppi**, Palo Alto, CA (US); **Andrew Everett Phelps**, Middleton, WI (US); **Reginald Clifford Young**, Palo Alto, CA (US); **Thomas Norrie**, Mountain View, CA (US); **Gregory Michael Thorson**, Waunakee, WI (US); **Dan Luu**, Madison, WI (US)

(52) **U.S. Cl.**
CPC **G06N 3/08** (2013.01); **G06F 15/8046** (2013.01); **G06N 3/063** (2013.01); **G06N 5/04** (2013.01)

(21) Appl. No.: **16/663,876**

(57) **ABSTRACT**

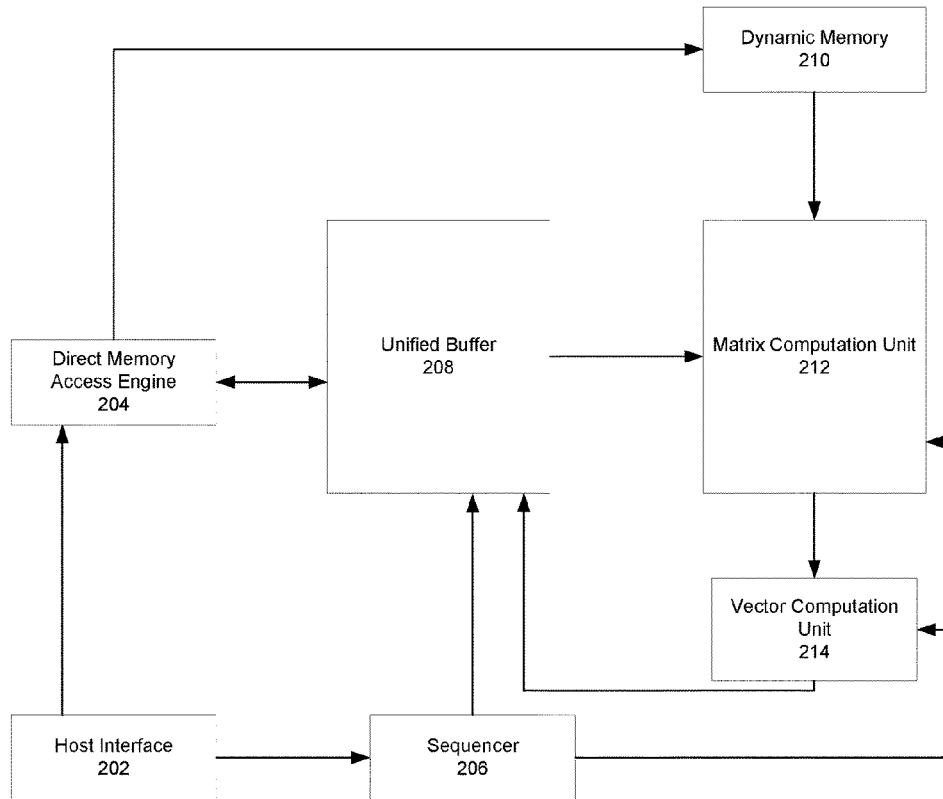
(22) Filed: **Oct. 25, 2019**

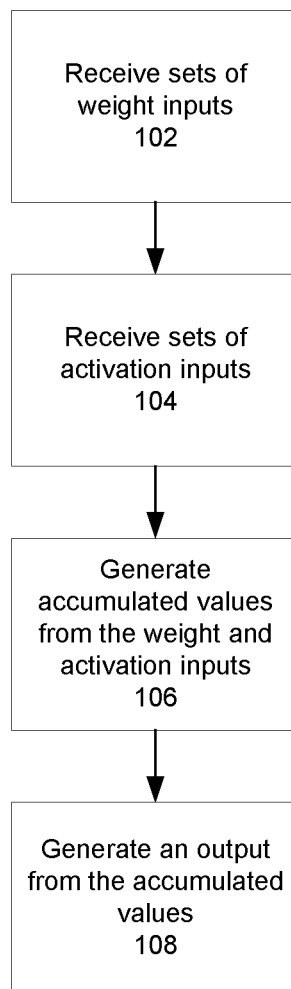
A circuit for performing neural network computations for a neural network comprising a plurality of neural network layers, the circuit comprising: a matrix computation unit configured to, for each of the plurality of neural network layers: receive a plurality of weight inputs and a plurality of activation inputs for the neural network layer, and generate a plurality of accumulated values based on the plurality of weight inputs and the plurality of activation inputs; and a vector computation unit communicatively coupled to the matrix computation unit and configured to, for each of the plurality of neural network layers: apply an activation function to each accumulated value generated by the matrix computation unit to generate a plurality of activated values for the neural network layer.

Related U.S. Application Data

(63) Continuation of application No. 16/529,782, filed on Aug. 1, 2019, which is a continuation of application No. 15/686,615, filed on Aug. 25, 2017, which is a continuation of application No. 14/844,524, filed on Sep. 3, 2015, now Pat. No. 9,747,546.

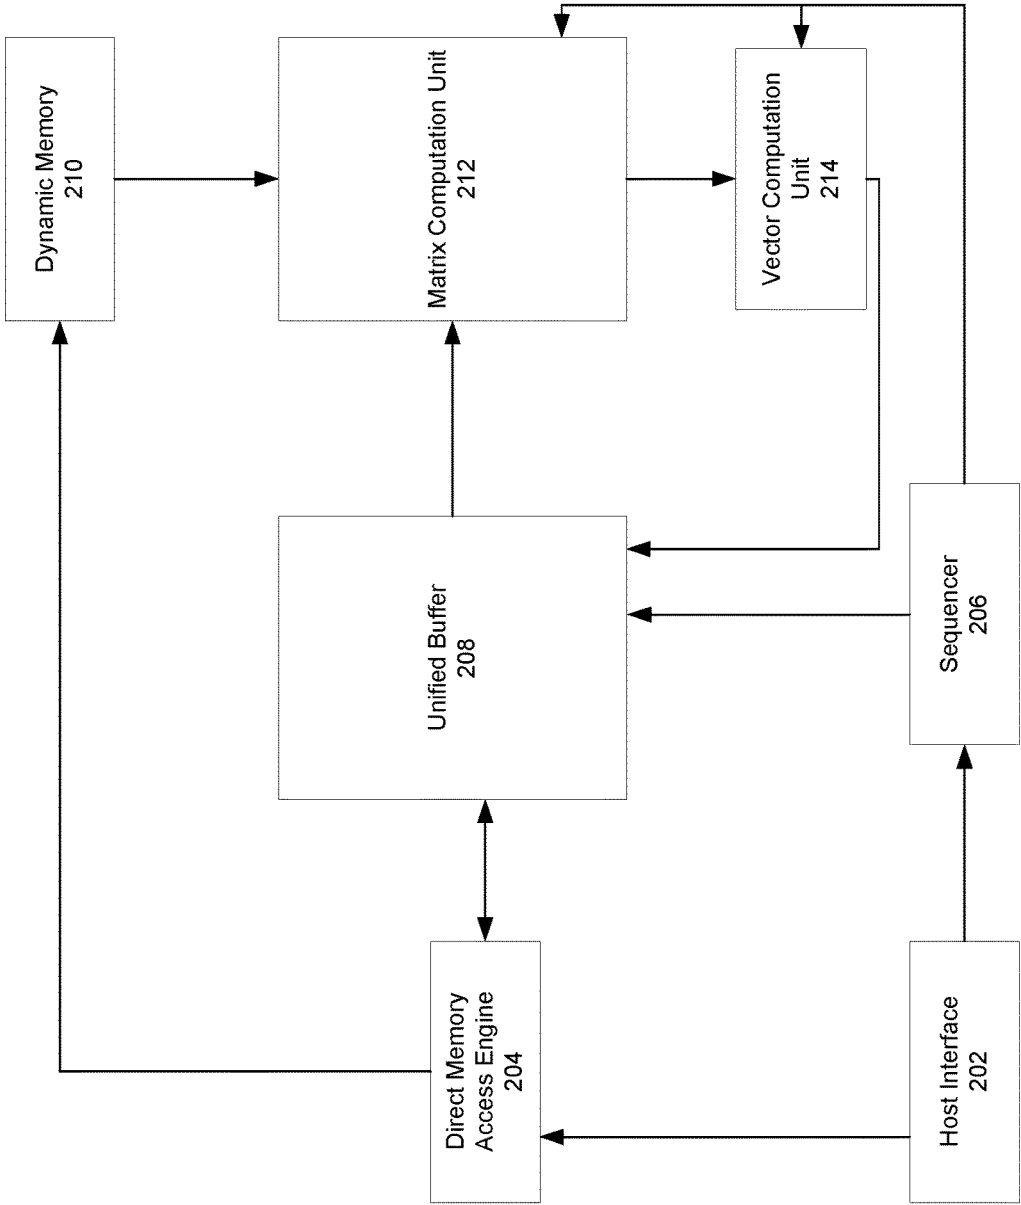
(60) Provisional application No. 62/164,931, filed on May 21, 2015.





100 ↗

FIG. 1



200

FIG. 2




FIG. 3

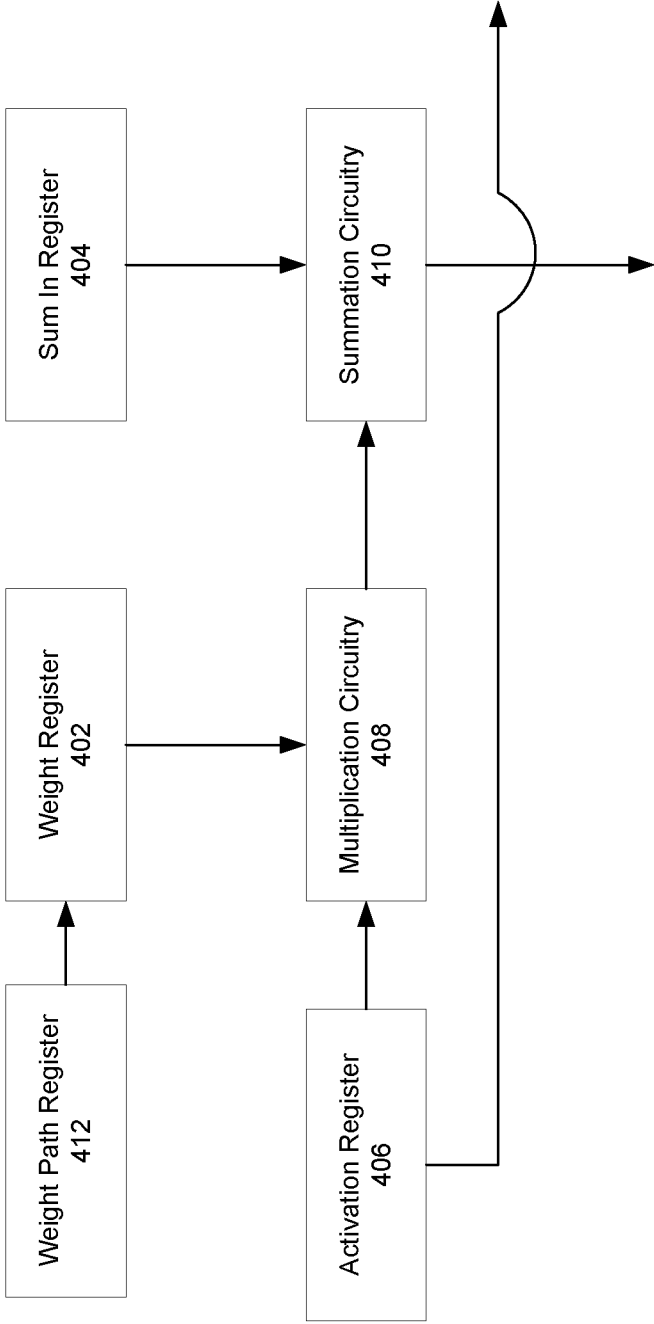



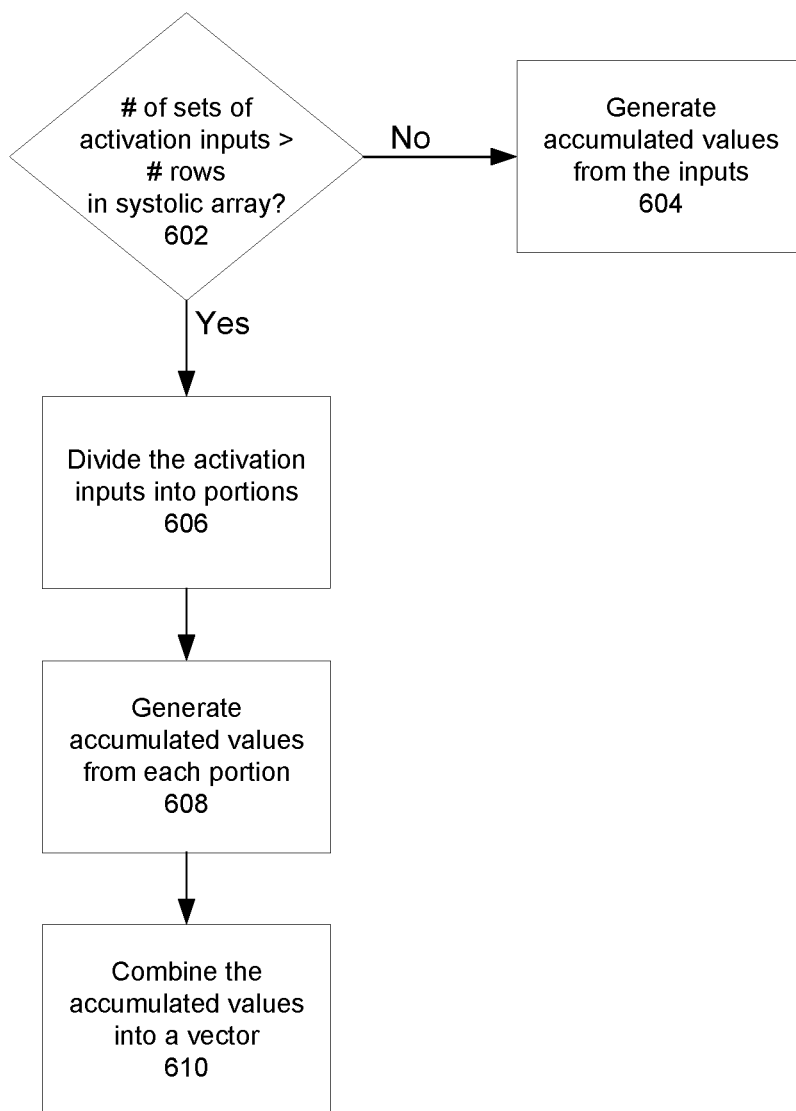
FIG. 4





500

FIG. 5



600 ↗

FIG. 6

NEURAL NETWORK PROCESSOR

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This application is a continuation of U.S. application Ser. No. 16/529,782, filed Aug. 1, 2019, which is a continuation of U.S. application Ser. No. 15/686,615, filed on Aug. 25, 2017, which is a continuation of U.S. application Ser. No. 14/844,524, filed on Sep. 3, 2015, which is a non-provisional of and claims priority to U.S. Provisional Patent Application No. 62/164,931, filed on May 21, 2015, the entire contents of which are hereby incorporated by reference.

BACKGROUND

[0002] This specification relates to computing neural network inferences in hardware.

[0003] Neural networks are machine learning models that employ one or more layers of models to generate an output, e.g., a classification, for a received input. Some neural networks include one or more hidden layers in addition to an output layer. The output of each hidden layer is used as input to the next layer in the network, i.e., the next hidden layer or the output layer of the network. Each layer of the network generates an output from a received input in accordance with current values of a respective set of parameters.

SUMMARY

[0004] In general, this specification describes a special-purpose hardware circuit that computes neural network inferences.

[0005] In general, one innovative aspect of the subject matter described in this specification can be embodied in a circuit for performing neural network computations for a neural network comprising a plurality of neural network layers, the circuit comprising: a matrix computation unit configured to, for each of the plurality of neural network layers: receive a plurality of weight inputs and a plurality of activation inputs for the neural network layer, and generate a plurality of accumulated values based on the plurality of weight inputs and the plurality of activation inputs; and a vector computation unit communicatively coupled to the matrix computation unit and configured to, for each of the plurality of neural network layers: apply an activation function to each accumulated value generated by the matrix computation unit to generate a plurality of activated values for the neural network layer.

[0006] Implementations can include one or more of the following features. A unified buffer communicatively coupled to the matrix computation unit and the vector computation unit, where the unified buffer is configured to receive and store output from the vector computation unit, and the unified buffer is configured to send the received output as input to the matrix computation unit. A sequencer configured to receive instructions from a host device and generate a plurality of control signals from the instructions, where the plurality of control signals control dataflow through the circuit; and a direct memory access engine communicatively coupled to the unified buffer and the sequencer, where the direct memory access engine is configured to send the plurality of activation inputs to the unified buffer, where the unified buffer is configured to send the plurality of activation inputs to the matrix computation

unit, and where the direct memory access engine is configured to read result data from the unified buffer. A memory unit configured to send the plurality of weight inputs to the matrix computation unit, and where the direct memory access engine is configured to send the plurality of weight inputs to the memory unit. The matrix computation unit is configured as a two dimensional systolic array comprising a plurality of cells. The plurality of weight inputs is shifted through a first plurality of cells along a first dimension of the systolic array, and where the plurality of activation inputs is shifted through a second plurality of cells along a second dimension of the systolic array. For a given layer in the plurality of layers, a count of the plurality of activation inputs is greater than a size of the second dimension of the systolic array, and where the systolic array is configured to: divide the plurality of activation inputs into portions, where each portion has a size less than or equal to the size of the second dimension; generating, for each portion, a respective portion of accumulated values; and combining each portion of accumulated values to generate a vector of accumulated values for the given layer. For a given layer in the plurality of layers, a count of the plurality of weight inputs is greater than a size of the first dimension of the systolic array, and where the systolic array is configured to: divide the plurality of weight inputs into portions, where each portion has a size less than or equal to the size of the first dimension; generating, for each portion, a respective portion of accumulated values; and combining each portion of accumulated values to generate a vector of accumulated values for the given layer. Each cell in the plurality of cells comprises: a weight register configured to store a weight input; an activation register configured to store an activation input and configured to send the activation input to another activation register in a first adjacent cell along the second dimension; a sum-in register configured to store a previously summed value; multiplication circuitry communicatively coupled to the weight register and the activation register, where the multiplication circuitry is configured to output a product of the weight input and the activation input; and summation circuitry communicatively coupled to the multiplication circuitry and the sum-in register, where the summation circuitry is configured to output a sum of the product and the previously summed value, and where the summation circuitry is configured to send the sum to another sum-in register in a second adjacent cell along the first dimension. One or more cells in the plurality of cells are each configured to store the respective sum in a respective accumulator unit, where the respective sum is an accumulated value. The first dimension of the systolic array corresponds to columns of the systolic array, and where the second dimension of the systolic array corresponds to rows of the systolic array. The vector computation unit normalizes each activated value to generate a plurality of normalized values. The vector computation unit pools one or more activated values to generate a plurality of pooled values.

[0007] Particular embodiments of the subject matter described in this specification can be implemented so as to realize one or more of the following advantages. Implementing a neural network processor in hardware improves efficiency, e.g., increase speed and throughput and reduce power and cost, over implementations in software. This can be useful for inference applications. Integrating components of the neural network processor into one circuit allows inferences to be computed without incurring penalties of

off-chip communication. Additionally, the circuit can process neural network layers that have a number of inputs, e.g., a number of weight inputs or a number of activation inputs, larger than a size of a dimension of a matrix computation unit within the circuit. For example, the circuit can process a large number of weight inputs per neuron of the neural network.

[0008] The details of one or more embodiments of the subject matter of this specification are set forth in the accompanying drawings and the description below. Other features, aspects, and advantages of the subject matter will become apparent from the description, the drawings, and the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] FIG. 1 is a flow diagram of an example method for performing a computation for a given layer of a neural network.

[0010] FIG. 2 shows an example neural network processing system.

[0011] FIG. 3 shows an example architecture including a matrix computation unit.

[0012] FIG. 4 shows an example architecture of a cell inside a systolic array.

[0013] FIG. 5 shows an example architecture of a vector computation unit.

[0014] FIG. 6 is a flow diagram of another example process for performing, using a systolic array, the computation for a given neural network layer having more activation inputs than rows in the systolic array.

[0015] Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION

[0016] A neural network having multiple layers can be used to compute inferences. For example, given an input, the neural network can compute an inference for the input. The neural network computes this inference by processing the input through each of the layers of the neural network. In particular, the layers of the neural network are arranged in a sequence, each with a respective set of weights. Each layer receives an input and processes the input in accordance with the set of weights for the layer to generate an output.

[0017] Therefore, in order to compute an inference from a received input, the neural network receives the input and processes it through each of the neural network layers in the sequence to generate the inference, with the output from one neural network layer being provided as input to the next neural network layer. Data inputs to a neural network layer, e.g., either the input to the neural network or the outputs of the layer below the layer in the sequence, to a neural network layer can be referred to as activation inputs to the layer.

[0018] In some implementations, the layers of the neural network are arranged in a directed graph. That is, any particular layer can receive multiple inputs, multiple outputs, or both. The layers of the neural network can also be arranged such that an output of a layer can be sent back as an input to a previous layer.

[0019] FIG. 1 is a flow diagram of an example process 100 for performing a computation for a given layer of a neural network using a special-purpose hardware circuit. For convenience, the method 100 will be described with respect to a system having one or more circuits that performs the

method 100. The method 100 can be performed for each layer of the neural network in order to compute an inference from a received input.

[0020] The system receives sets of weight inputs (step 102) and sets of activation inputs (step 104) for the given layer. The sets of weight inputs and the sets of activation inputs can be received from dynamic memory and a unified buffer, respectively, of the special-purpose hardware circuit. In some implementations, both the sets of weight inputs and the sets of activation inputs can be received from the unified buffer.

[0021] The system generates accumulated values from the weight inputs and the activation inputs using a matrix multiplication unit of the special-purpose hardware circuit (step 106). In some implementations, the accumulated values are dot products of the sets of weight inputs and the sets of activation inputs. That is, for one set of weights, the system can multiply each weight input with each activation input and sum the products together to form an accumulated value. The system can then compute dot products of other set of weights with other sets of activation inputs.

[0022] The system can generate a layer output from the accumulation values (step 108) using a vector computation unit of the special-purpose hardware circuit. In some implementations, the vector computation unit applies an activation function to the accumulated values, which will be described further below in reference to FIG. 5. The output of the layer can be stored in the unified buffer for use as an input to a subsequent layer in the neural network or can be used to determine the inference. The system finishes processing the neural network when a received input has been processed through each layer of the neural network to generate the inference for the received input.

[0023] FIG. 2 shows an example special-purpose integrated circuit 200 for performing neural network computations. The system 200 includes a host interface 202. The host interface 202 can receive instructions that include parameters for a neural network computation. The parameters can include at least one or more of the following: how many layers should be processed, corresponding sets of weight inputs for each layer of the layer, an initial set of activation inputs, i.e., the input to the neural network from which the inference is to be computed, corresponding input and output sizes of each layer, a stride value for the neural network computation, and a type of layer to be processed, e.g., a convolutional layer or a fully connected layer.

[0024] The host interface 202 can send the instructions to a sequencer 206, which converts the instructions into low level control signals that control the circuit to perform the neural network computations. In some implementations, the control signals regulate dataflow in the circuit, e.g., how the sets of weight inputs and the sets of activation inputs flow through the circuit. The sequencer 206 can send the control signals to a unified buffer 208, a matrix computation unit 212, and a vector computation unit 214. In some implementations, the sequencer 206 also sends control signals to a direct memory access engine 204 and dynamic memory 210. In some implementations, the sequencer 206 is a processor that generates clock signals. The sequencer 206 can use timing of the clock signals to, at appropriate times, send the control signals to each component of the circuit 200. In some other implementations, the host interface 202 passes in a clock signal from an external processor.

[0025] The host interface 202 can send the sets of weight inputs and the initial set of activation inputs to the direct memory access engine 204. The direct memory access engine 204 can store the sets of activation inputs at the unified buffer 208. In some implementations, the direct memory access stores the sets of weights to dynamic memory 210, which can be a memory unit. In some implementations, the dynamic memory is located off of the circuit.

[0026] The unified buffer 208 is a memory buffer. It can be used to store the set of activation inputs from the direct memory access engine 204 and outputs of the vector computation unit 214. The vector computation unit will be described in more detail below with reference to FIG. 5. The direct memory access engine 204 can also read the outputs of the vector computation unit 214 from the unified buffer 208.

[0027] The dynamic memory 210 and the unified buffer 208 can send the sets of weight inputs and the sets of activation inputs, respectively, to the matrix computation unit 212. In some implementations, the matrix computation unit 212 is a two-dimensional systolic array. The matrix computation unit 212 can also be a one-dimensional systolic array or other circuitry that can perform mathematical operations, e.g., multiplication and addition. In some implementations, the matrix computation unit 212 is a general purpose matrix processor.

[0028] The matrix computation unit 212 can process the weight inputs and the activation inputs and provide a vector of outputs to the vector computation unit 214. In some implementations, the matrix computation unit sends the vector of outputs to the unified buffer 208, which sends the vector of outputs to the vector computation unit 214. The vector computation unit can process the vector of outputs and store a vector of processed outputs to the unified buffer 208. The vector of processed outputs can be used as activation inputs to the matrix computation unit 212, e.g., for use in a subsequent layer in the neural network. The matrix computation unit 212 and the vector computation unit 214 will be described in more detail below with reference to FIG. 3 and FIG. 5, respectively.

[0029] FIG. 3 shows an example architecture 300 including a matrix computation unit. The matrix computation unit is a two-dimensional systolic array 306. The two-dimensional systolic array 306 can be a square array. The array 306 includes multiple cells 304. In some implementations, a first dimension 320 of the systolic array 306 corresponds to columns of cells and a second dimension 322 of the systolic array 306 corresponds to rows of cells. The systolic array can have more rows than columns, more columns than rows, or an equal number of columns and rows.

[0030] In the illustrated example, value loaders 302 send activation inputs to rows of the array 306 and a weight fetcher interface 308 sends weight inputs to columns of the array 306. In some other implementations, however, activation inputs are transferred to the columns and weight inputs are transferred to the rows of the array 306.

[0031] The value loaders 302 can receive the activation inputs from a unified buffer, e.g., the unified buffer 208 of FIG. 2. Each value loader can send a corresponding activation input to a distinct left-most cell of the array 306. The left-most cell can be a cell along a left-most column of the array 306. For example, value loader 312 can send an activation input to cell 314. The value loader can also send the activation input to an adjacent value loader, and the

activation input can be used at another left-most cell of the array 306. This allows activation inputs to be shifted for use in another particular cell of the array 306.

[0032] The weight fetcher interface 308 can receive the weight input from a memory unit, e.g., the dynamic memory 210 of FIG. 2. The weight fetcher interface 308 can send a corresponding weight input to a distinct top-most cell of the array 306. The top-most cell can be a cell along a top-most row of the array 306. For example, the weight fetcher interface 308 can send weight inputs to cells 314 and 316.

[0033] In some implementations, a host interface, e.g., the host interface 202 of FIG. 2, shifts activation inputs throughout the array 306 along one dimension, e.g., to the right, while shifting weight inputs throughout the array 306 along another dimension, e.g., to the bottom. For example, over one clock cycle, the activation input at cell 314 can shift to an activation register in cell 316, which is to the right of cell 314. Similarly, the weight input at cell 316 can shift to a weight register at cell 318, which is below cell 314.

[0034] On each clock cycle, each cell can process a given weight input and a given activation input to generate an accumulated output. The accumulated output can also be passed to an adjacent cell along the same dimension as the given weight input. An individual cell is described further below with reference FIG. 4.

[0035] The accumulated output can be passed along the same column as the weight input, e.g., towards the bottom of the column in the array 306. In some implementations, at the bottom of each column, the array 306 can include accumulator units 310 that store and accumulate each accumulated output from each column when performing calculations with layers having more weight inputs than columns or layers having more activation inputs than rows. In some implementations, each accumulator unit stores multiple parallel accumulations. This will be described further below with reference to FIG. 6. The accumulator units 310 can accumulate each accumulated output to generate a final accumulated value. The final accumulated value can be transferred to a vector computation unit, e.g., the vector computation unit 502 of FIG. 5. In some other implementations, the accumulator units 310 passes the accumulated values to the vector computation unit without performing any accumulations when processing layers with fewer weight inputs than columns or layers having fewer activation inputs than rows.

[0036] FIG. 4 shows an example architecture 400 of a cell inside a systolic array, e.g., the systolic array 306 of FIG. 3.

[0037] The cell can include an activation register 406 that stores an activation input. The activation register can receive the activation input from a left adjacent cell, i.e., an adjacent cell located to the left of the given cell, or from a unified buffer, depending on the position of the cell within the systolic array. The cell can include a weight register 402 that stores a weight input. The weight input can be transferred from a top adjacent cell or from a weight fetcher interface, depending on the position of the cell within the systolic array. The cell can also include a sum in register 404. The sum in register 404 can store an accumulated value from the top adjacent cell. Multiplication circuitry 408 can be used to multiply the weight input from the weight register 402 with the activation input from the activation register 406. The multiplication circuitry 408 can output the product to summation circuitry 410.

[0038] The summation circuitry can sum the product and the accumulated value from the sum in register 404 to generate a new accumulated value. The summation circuitry 410 can then send the new accumulated value to another sum in register located in a bottom adjacent cell. The new accumulated value can be used as an operand for a summation in the bottom adjacent cell.

[0039] The cell can also shift the weight input and the activation input to adjacent cells for processing. For example, the weight register 402 can send the weight input to another weight register in the bottom adjacent cell. The activation register 406 can send the activation input to another activation register in the right adjacent cell. Both the weight input and the activation input can therefore be reused by other cells in the array at a subsequent clock cycle.

[0040] In some implementations, the cell also includes a control register. The control register can store a control signal that determines whether the cell should shift either the weight input or the activation input to adjacent cells. In some implementations, shifting the weight input or the activation input takes one or more clock cycles. The control signal can also determine whether the activation input or weight inputs are transferred to the multiplication circuitry 408, or can determine whether the multiplication circuitry 408 operates on the activation and weight inputs. The control signal can also be passed to one or more adjacent cells, e.g., using a wire.

[0041] In some implementations, weights are pre-shifted into a weight path register 412. The weight path register 412 can receive the weight input, e.g., from a top adjacent cell, and transfer the weight input to the weight register 402 based on the control signal. The weight register 402 can statically store the weight input such that as activation inputs are transferred to the cell, e.g., through the activation register 406, over multiple clock cycles, the weight input remains within the cell and is not transferred to an adjacent cell. Therefore, the weight input can be applied to multiple activation inputs, e.g., using the multiplication circuitry 408, and respective accumulated values can be transferred to an adjacent cell.

[0042] FIG. 5 shows an example architecture 500 of a vector computation unit 502. The vector computation unit 502 can receive a vector of accumulated values from a matrix computation unit, e.g., the matrix computation unit described in reference to FIG. 2.

[0043] The vector computation unit 502 can process the vector of accumulated values at the activation unit 504. In some implementations, the activation unit includes circuitry that applies a non-linear function to each accumulated value to generate activation values. For example, the non-linear function can be $\tanh(x)$, where x is an accumulated value.

[0044] Optionally, the vector computation unit 502 can normalize the activation values in a normalization unit 506 that generates normalized values from the activation values.

[0045] Also optionally, the vector computation unit 502 can pool values, either activation values or normalization values, using a pooling unit 508. The pooling unit 508 can apply an aggregation function to one or more of the normalized values to generate pooled values. In some implementations, the aggregation functions are functions that return a maximum, minimum, or average of the normalized values or of a subset of the normalized values.

[0046] Control signals 510 can be transferred, e.g., by the sequencer 206 of FIG. 2, and can regulate how the vector

computation unit 502 processes the vector of accumulated values. That is, the control signals 510 can regulate whether the activation values are pooled, normalized, or both. The control signals 510 can also specify the activation, normalization, or pooling functions, as well as other parameters for normalization and pooling, e.g., a stride value.

[0047] The vector computation unit 502 can send values, e.g., activation values, normalized values, or pooled values, to a unified buffer, e.g., the unified buffer 208 of FIG. 2.

[0048] In some implementations, the pooling unit 508 receives the activation values instead of the normalization unit 506, and the pooling unit 508 sends the pooled values to the normalization unit 506, which generates normalized values to be stored in the unified buffer.

[0049] FIG. 6 is a flow diagram of example process for performing, using a systolic array, the computation for a given neural network layer having more activation inputs than rows in the systolic array. For convenience, the process 600 will be described with respect to a system that performs the process 600. In some implementations, a host interface or a sequencer performs the process 600, e.g., the host interface 202 or the sequencer 206, respectively, of FIG. 2. In some other implementations, the host interface receives instructions from an external processor that performs the process 600.

[0050] As described above, each layer can have multiple sets of activation inputs and each set of weight inputs can be transferred to cells at distinct rows of the array. In some implementations, some layers of the neural network have more sets of activation inputs than there are rows of the array.

[0051] The system can determine, e.g., using a comparator, whether there are more sets of activation inputs for the given neural network layer than there are rows in the systolic array. In some implementations, the system makes the determination at compile time. A set of activation inputs can correspond to the activation inputs provided to a single row of the array.

[0052] If there are more rows than sets of activation inputs (step 602), the system can generate accumulated values as described above in the systolic array 306 of FIG. 3 (step 604).

[0053] If there are more sets of activation inputs to be processed than there are rows in the array (step 602), the system can divide the sets of activation inputs into portions so that each portion has a size less than or equal to a number of rows in the array (step 606).

[0054] The system then can generate, for each portion of activation inputs, a portion of accumulated values (step 608). An accumulated value can be a sum of products of activation and weight inputs to cells along a given column, e.g., as described in systolic array 306 of FIG. 3. Each portion of accumulated values can be stored in a buffer until all portions of activation inputs have been processed. The buffer can be a buffer in accumulator units 310 of FIG. 3, a buffer in the systolic array, or the unified buffer 208 of FIG. 2.

[0055] The system can then combine all portions of accumulated values into a vector of accumulated values (step 610). In particular, the system can access the buffer of previously stored portions of accumulated values and accumulate, e.g., using accumulator units 310 of FIG. 3, the accumulated values to generate a vector of the accumulated

values. The system can send the vector of the accumulated values to a vector computation unit, e.g., the vector computation unit 214 of FIG. 2.

[0056] For example, if there are 256 rows in the array and there are 300 sets of activation inputs to process at a given layer, the system can generate 256 final accumulated values from 256 sets of activation inputs for complete utilization of the systolic array and store the 256 final accumulated values in a buffer. The system can then generate 44 final accumulated values from the 44 remainder sets of activation inputs. Finally, the system can combine all 300 final accumulated values to form a vector and send the vector to the vector computation unit.

[0057] If there are more sets of weight inputs than columns to the array, the system can perform similar operations. That is, the system can divide the sets of weight inputs into portions having fewer sets of weight inputs than a number of columns in the array, generate accumulated values for each portion, and combine the accumulated values into a vector for use in the vector computation unit. In some implementations, instead of comparing the number of sets of weight inputs with the number of columns in the array, the system can compare the number of accumulated values with the number of columns in the array.

[0058] Although the system has been described with weight inputs being transferred to columns of the array and activation inputs being transferred to rows of the array, in some implementations, the weight inputs are transferred to rows of the array and the activation inputs are transferred to columns of the array.

[0059] Although the hardware is described to be for computing inferences, the hardware can be used for one or more of the following: convolutional or fully-connected neural network training, linear or logistic regression, clustering, e.g., k-means clustering, video-encoding, and image processing.

[0060] Embodiments of the subject matter and the functional operations described in this specification can be implemented in digital electronic circuitry, in tangibly-embodied computer software or firmware, in computer hardware, including the structures disclosed in this specification and their structural equivalents, or in combinations of one or more of them. Embodiments of the subject matter described in this specification can be implemented as one or more computer programs, i.e., one or more modules of computer program instructions encoded on a tangible non transitory program carrier for execution by, or to control the operation of, data processing apparatus. Alternatively or in addition, the program instructions can be encoded on an artificially generated propagated signal, e.g., a machine-generated electrical, optical, or electromagnetic signal, that is generated to encode information for transmission to suitable receiver apparatus for execution by a data processing apparatus. The computer storage medium can be a machine-readable storage device, a machine-readable storage substrate, a random or serial access memory device, or a combination of one or more of them.

[0061] The term “data processing apparatus” encompasses all kinds of apparatus, devices, and machines for processing data, including by way of example a programmable processor, a computer, or multiple processors or computers. The apparatus can include special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application specific integrated circuit). The apparatus can also

include, in addition to hardware, code that creates an execution environment for the computer program in question, e.g., code that constitutes processor firmware, a protocol stack, a database management system, an operating system, or a combination of one or more of them.

[0062] A computer program (which may also be referred to or described as a program, software, a software application, a module, a software module, a script, or code) can be written in any form of programming language, including compiled or interpreted languages, or declarative or procedural languages, and it can be deployed in any form, including as a standalone program or as a module, component, subroutine, or other unit suitable for use in a computing environment. A computer program may, but need not, correspond to a file in a file system. A program can be stored in a portion of a file that holds other programs or data, e.g., one or more scripts stored in a markup language document, in a single file dedicated to the program in question, or in multiple coordinated files, e.g., files that store one or more modules, sub programs, or portions of code. A computer program can be deployed to be executed on one computer or on multiple computers that are located at one site or distributed across multiple sites and interconnected by a communication network.

[0063] The processes and logic flows described in this specification can be performed by one or more programmable computers executing one or more computer programs to perform functions by operating on input data and generating output. The processes and logic flows can also be performed by, and apparatus can also be implemented as, special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application specific integrated circuit).

[0064] Computers suitable for the execution of a computer program include, by way of example, can be based on general or special purpose microprocessors or both, or any other kind of central processing unit. Generally, a central processing unit will receive instructions and data from a read only memory or a random access memory or both. The essential elements of a computer are a central processing unit for performing or executing instructions and one or more memory devices for storing instructions and data. Generally, a computer will also include, or be operatively coupled to receive data from or transfer data to, or both, one or more mass storage devices for storing data, e.g., magnetic, magneto optical disks, or optical disks. However, a computer need not have such devices. Moreover, a computer can be embedded in another device, e.g., a mobile telephone, a personal digital assistant (PDA), a mobile audio or video player, a game console, a Global Positioning System (GPS) receiver, or a portable storage device, e.g., a universal serial bus (USB) flash drive, to name just a few.

[0065] Computer readable media suitable for storing computer program instructions and data include all forms of nonvolatile memory, media and memory devices, including by way of example semiconductor memory devices, e.g., EPROM, EEPROM, and flash memory devices; magnetic disks, e.g., internal hard disks or removable disks; magneto optical disks; and CD ROM and DVD-ROM disks. The processor and the memory can be supplemented by, or incorporated in, special purpose logic circuitry.

[0066] To send for interaction with a user, embodiments of the subject matter described in this specification can be implemented on a computer having a display device, e.g., a

CRT (cathode ray tube) or LCD (liquid crystal display) monitor, for displaying information to the user and a keyboard and a pointing device, e.g., a mouse or a trackball, by which the user can send input to the computer. Other kinds of devices can be used to send for interaction with a user as well; for example, feedback provided to the user can be any form of sensory feedback, e.g., visual feedback, auditory feedback, or tactile feedback; and input from the user can be received in any form, including acoustic, speech, or tactile input. In addition, a computer can interact with a user by sending documents to and receiving documents from a device that is used by the user; for example, by sending web pages to a web browser on a user's client device in response to requests received from the web browser.

[0067] Embodiments of the subject matter described in this specification can be implemented in a computing system that includes a back end component, e.g., as a data server, or that includes a middleware component, e.g., an application server, or that includes a front end component, e.g., a client computer having a graphical user interface or a Web browser through which a user can interact with an implementation of the subject matter described in this specification, or any combination of one or more such back end, middleware, or front end components. The components of the system can be interconnected by any form or medium of digital data communication, e.g., a communication network. Examples of communication networks include a local area network ("LAN") and a wide area network ("WAN"), e.g., the Internet.

[0068] The computing system can include clients and servers. A client and server are generally remote from each other and typically interact through a communication network. The relationship of client and server arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other.

[0069] While this specification contains many specific implementation details, these should not be construed as limitations on the scope of any invention or of what may be claimed, but rather as descriptions of features that may be specific to particular embodiments of particular inventions. Certain features that are described in this specification in the context of separate embodiments can also be implemented in combination in a single embodiment. Conversely, various features that are described in the context of a single embodiment can also be implemented in multiple embodiments separately or in any suitable subcombination. Moreover, although features may be described above as acting in certain combinations and even initially claimed as such, one or more features from a claimed combination can in some cases be excised from the combination, and the claimed combination may be directed to a subcombination or variation of a subcombination.

[0070] Similarly, while operations are depicted in the drawings in a particular order, this should not be understood as requiring that such operations be performed in the particular order shown or in sequential order, or that all illustrated operations be performed, to achieve desirable results. In certain circumstances, multitasking and parallel processing may be advantageous. Moreover, the separation of various system modules and components in the embodiments described above should not be understood as requiring such separation in all embodiments, and it should be understood that the described program components and

systems can generally be integrated together in a single software product or packaged into multiple software products.

[0071] Particular embodiments of the subject matter have been described. Other embodiments are within the scope of the following claims. For example, the actions recited in the claims can be performed in a different order and still achieve desirable results. As one example, the processes depicted in the accompanying figures do not necessarily require the particular order shown, or sequential order, to achieve desirable results. In certain implementations, multitasking and parallel processing may be advantageous.

What is claimed is:

1. (canceled)
2. A circuit for performing neural network computations for a neural network comprising a plurality of neural network layers, the circuit comprising:
 - a matrix computation unit configured to compute multiplications using a plurality of cells of the matrix computation unit to generate accumulated values, wherein each of the multiplications is between a weight for a neural network layer and an input to the neural network layer; and
 - a vector computation unit configured to (i) generate a plurality of activated values for the neural network layer based on the accumulated values generated by the matrix computation unit and (ii) generate an output for the neural network layer based on the plurality of activated values.
3. The circuit of claim 2, wherein:
 - the matrix computation unit is configured as a multi-dimensional systolic array; and
 - the plurality of cells are arranged along at least a first dimension and a second dimension of the systolic array, the first dimension being different than the second dimension.
4. The circuit of claim 3, wherein the matrix computation unit is configured to:
 - load one or more weights for the neural network layer into the plurality of cells of the matrix computation unit; and
 - shift one or more weights through the plurality of cells of the matrix computation unit.
5. The circuit of claim 4, wherein the matrix computation unit is configured to:
 - shift one or more inputs for the neural network layer into the plurality of cells of the matrix computation unit to perform the multiplications between the weight for the neural network layer and the input to the neural network layer.
6. The circuit of claim 4, wherein the vector computation unit is configured to:
 - apply an activation function to each of the accumulated values provided by the matrix computation unit; and
 - generate the plurality of activated values for the neural network layer based on the activation function that is applied to each of the accumulated values.
7. The circuit of claim 4, wherein each activated value of the plurality of activated values represents an activation input to a second neural network layer and the matrix computation unit is configured to:

- shift one or more weights for the second neural network layer through a first plurality of cells of the systolic array along a first, column dimension of the systolic array; and
- shift a plurality of activation inputs to the second neural network layer through a second plurality of cells of the systolic array along a second, row dimension of the systolic array.
- 8.** The circuit of claim **7**, wherein the matrix computation unit is configured to:
- generate a vector of accumulated values for the second neural network layer based on dot products of multiplications between the one or more weights for the second neural network layer and different activation inputs to the second neural network layer that are shifted along the second, row dimension of the of the systolic array.
- 9.** The circuit of claim **8**, wherein the vector computation unit is configured to:
- receive the vector of accumulated values generated by the matrix computation unit; and
 - generate a vector of activation values in response to applying an activation function to each accumulated value in the vector of accumulated values.
- 10.** The circuit of claim **2**, wherein:
- the matrix computation unit is configured as a multi-dimensional systolic array;
 - the systolic array includes the plurality of cells being arranged along at least a first dimension of the systolic array and a second dimension of the systolic array; and
 - the first dimension and the second dimension are the same dimension.
- 11.** A method for performing neural network computations using a circuit configured to implement a neural network comprising a plurality of neural network layers, the method comprising:
- receiving, by a matrix computation unit in the circuit, a plurality of weights for a neural network layer and inputs to the neural network layer;
 - computing, using a plurality of cells of a matrix computation unit, multiplications between a weight for the neural network layer and one or more of the inputs to the neural network layer; and
 - generating, using activation circuitry of a vector computation unit in the circuit, an output for the neural network layer based on the multiplications.
- 12.** The method of claim **11**, wherein generating the output for the neural network layer comprises:
- generating, by the matrix computation unit, accumulated values for the neural network layer based on the multiplications; and
 - generating, by the vector computation unit, a plurality of activated values for the neural network layer based on the accumulated values generated by the matrix computation unit.
- 13.** The method of claim **12**, wherein computing the multiplications between the weight for the neural network layer and one or more of the inputs comprises:
- computing the multiplications using a multi-dimensional systolic array of the matrix computation unit, and
 - wherein the plurality of cells are arranged along at least a first dimension and a second dimension of the systolic array, the first dimension being different than the second dimension.
- 14.** The method of claim **13**, wherein receiving the plurality of weights for the neural network layer comprises:
- loading one or more weights for the neural network layer into a distinct cell of the plurality of cells of the matrix computation unit; and
 - shifting one or more weights through one or more cells of the plurality of cells of the matrix computation unit.
- 15.** The method of claim **14**, wherein receiving the inputs to the neural network layer comprises:
- shifting one or more inputs for the neural network layer into the distinct cell of the plurality of cells of the matrix computation unit to perform the multiplications between the weight for the neural network layer and one or more of the inputs to the neural network layer.
- 16.** The method of claim **12**, wherein generating the plurality of activated values for the neural network layer comprises:
- applying, by the activation circuitry of the vector computation unit, an activation function to each of the accumulated values generated by the matrix computation unit; and
 - generating the plurality of activated values for the neural network layer based on the activation function applied to each of the accumulated values.
- 17.** The method of claim **16**, comprising:
- receiving, by the vector computation unit, a vector of accumulated values generated by the matrix computation unit; and
 - generating, by the vector computation unit, a vector of activation values in response to applying the activation function to each accumulated value in the vector of accumulated values.
- 18.** The method of claim **13**, wherein each activated value of the plurality of activated values represents an activation input to a second neural network layer and the method comprises:
- shifting one or more weights for the second neural network layer through a first plurality of cells of the systolic array along a first, column dimension of the systolic array; and
 - shifting a plurality of activation inputs to the second neural network layer through a second plurality of cells of the systolic array along a second, row dimension of the systolic array.
- 19.** The method of claim **18**, comprising:
- generating a vector of accumulated values for the second neural network layer based on dot products of multiplications between the one or more weights for the second neural network layer and different activation inputs provided to the second neural network layer that are shifted along the second, row dimension of the of the systolic array.
- 20.** The method of claim **12**, wherein:
- the matrix computation unit is configured as a multi-dimensional systolic array;
 - the systolic array includes the plurality of cells being arranged along at least a first dimension of the systolic array and a second dimension of the systolic array; and
 - the first dimension and the second dimension are the same dimension.
- 21.** One or more non-transitory machine-readable storage devices for storing instructions that are executable by one or more processing devices to cause performance of operations for performing neural network computations using a circuit

configured to implement a neural network comprising a plurality of neural network layers, the operations comprising:

receiving, by a matrix computation unit in the circuit, a plurality of weights for a neural network layer and inputs to the neural network layer;

computing, using a plurality of cells of the matrix computation unit, multiplications between a weight for the neural network layer and one or more of the inputs to the neural network layer; and

generating, using activation circuitry of a vector computation unit in the circuit, an output for the neural network layer based on the multiplications.

* * * * *