



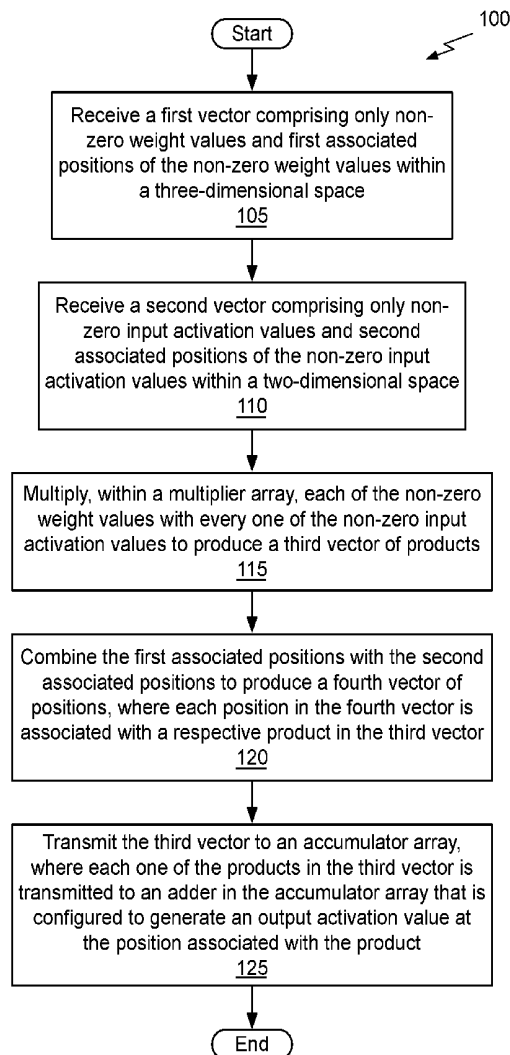
US 20180046916A1

(19) **United States**(12) **Patent Application Publication** (10) **Pub. No.: US 2018/0046916 A1**
(43) **Pub. Date: Feb. 15, 2018**
Dally et al.(54) **SPARSE CONVOLUTIONAL NEURAL NETWORK ACCELERATOR****Publication Classification**(71) Applicant: **NVIDIA Corporation**, Santa Clara, CA (US)(51) **Int. Cl.**
G06N 3/08 (2006.01)
G06F 7/523 (2006.01)
G06N 3/04 (2006.01)(72) Inventors: **William J. Dally**, Los Altos Hills, CA (US); **Angshuman Parashar**, Northborough, MA (US); **Joel Springer Emer**, Acton, MA (US); **Stephen William Keckler**, Austin, TX (US); **Larry Robert Dennison**, Medon, MA (US)(52) **U.S. Cl.**
CPC **G06N 3/08** (2013.01); **G06N 3/04** (2013.01); **G06F 7/523** (2013.01)(21) Appl. No.: **15/458,837**(22) Filed: **Mar. 14, 2017****Related U.S. Application Data**

(60) Provisional application No. 62/373,919, filed on Aug. 11, 2016.

(57) **ABSTRACT**

A method, computer program product, and system perform computations using a sparse convolutional neural network accelerator. Compressed-sparse data is received for input to a processing element, wherein the compressed-sparse data encodes non-zero elements and corresponding multi-dimensional positions. The non-zero elements are processed in parallel by the processing element to produce a plurality of result values. The corresponding multi-dimensional positions are processed in parallel by the processing element to produce destination addresses for each result value in the plurality of result values. Each result value is transmitted to a destination accumulator associated with the destination address for the result value.



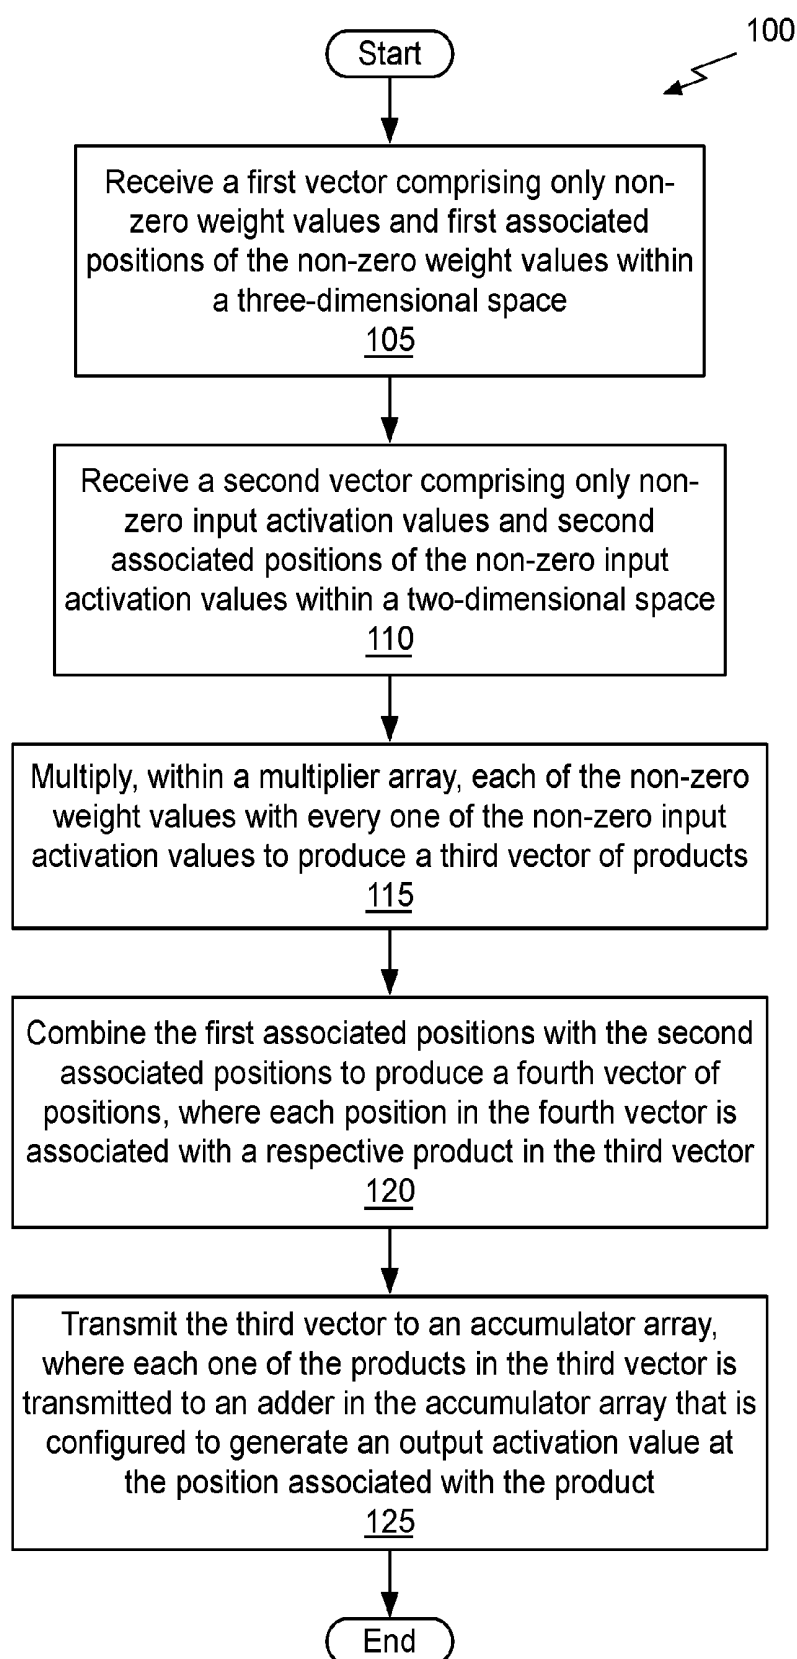


Fig. 1

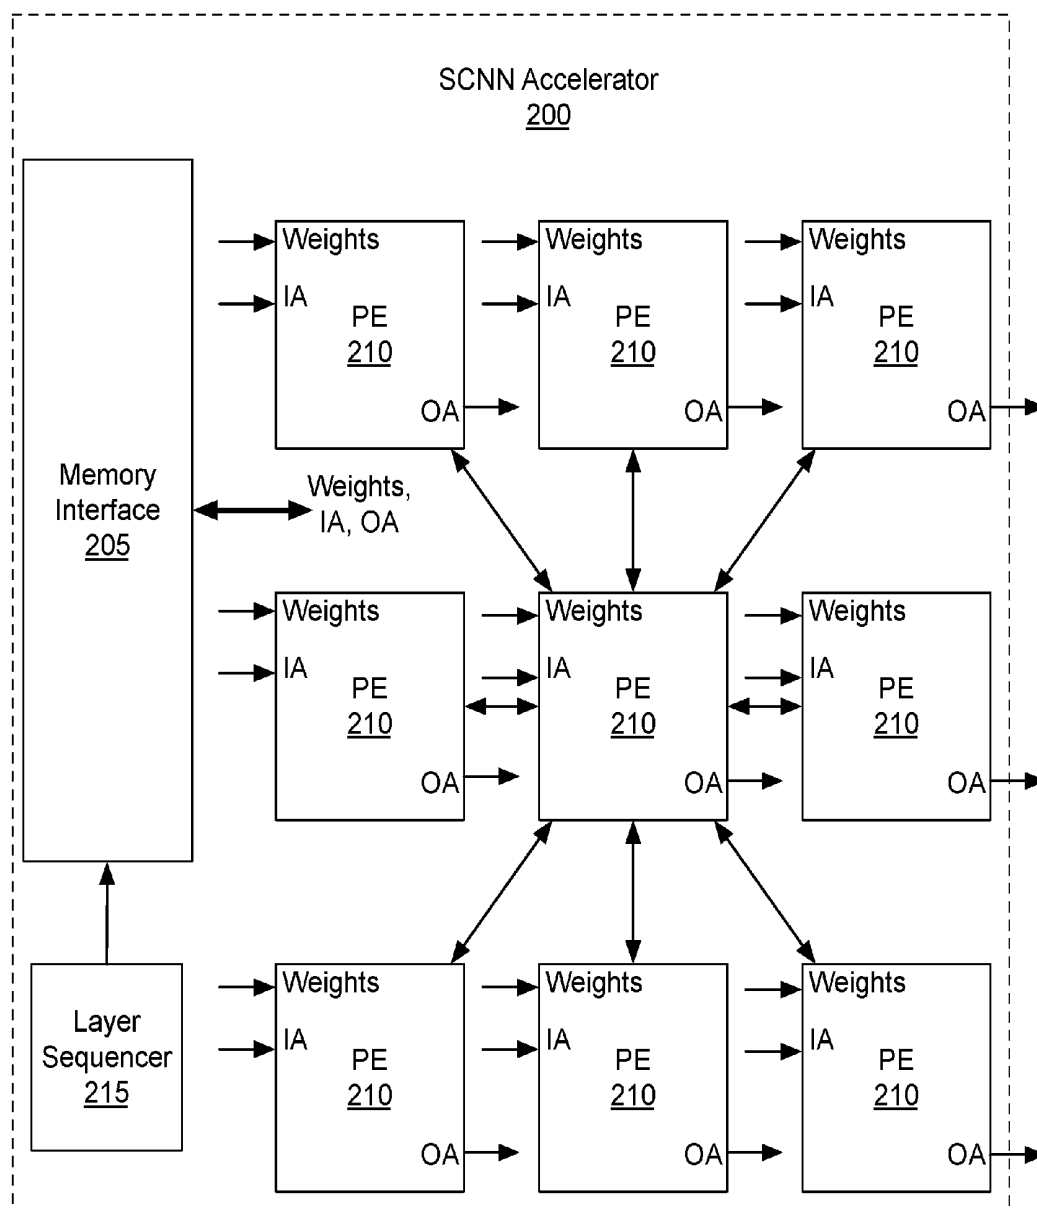


Fig. 2A

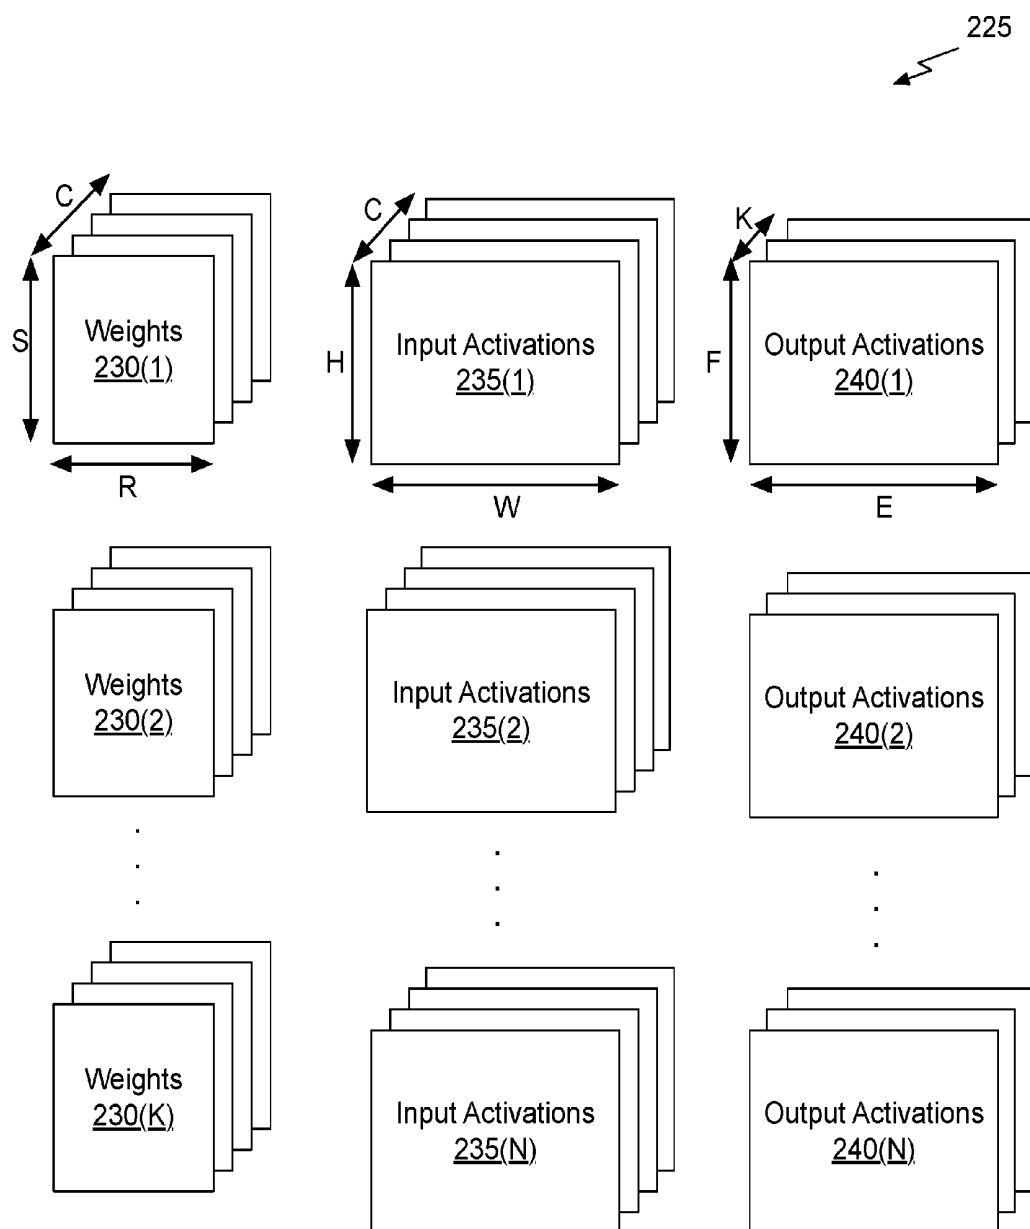


Fig. 2B

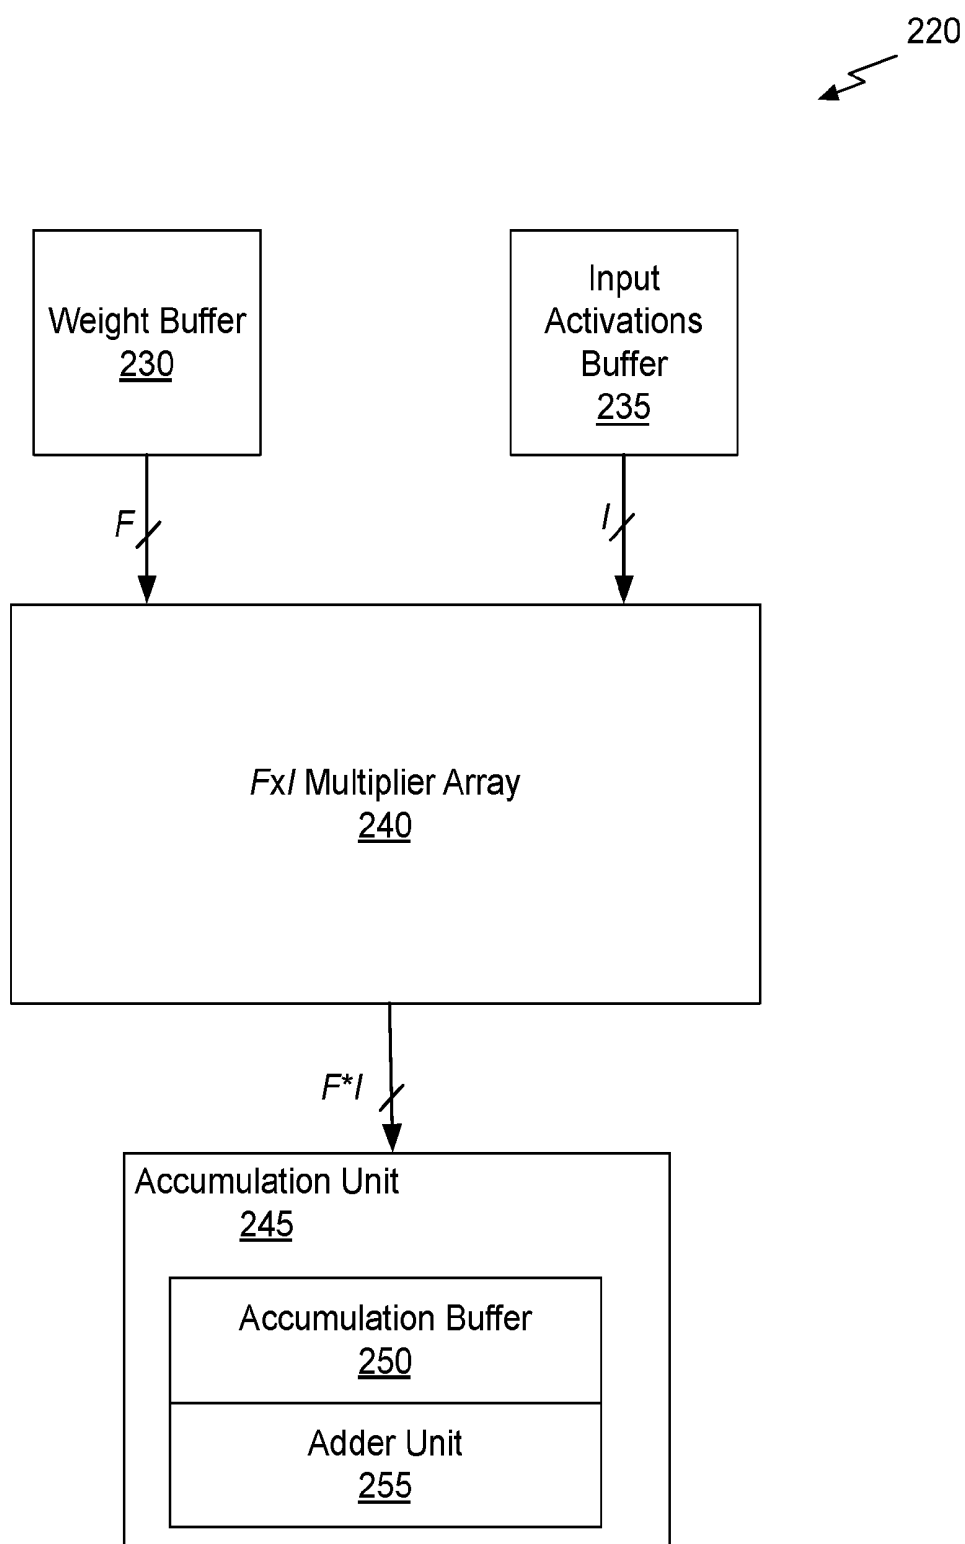


Fig. 2C

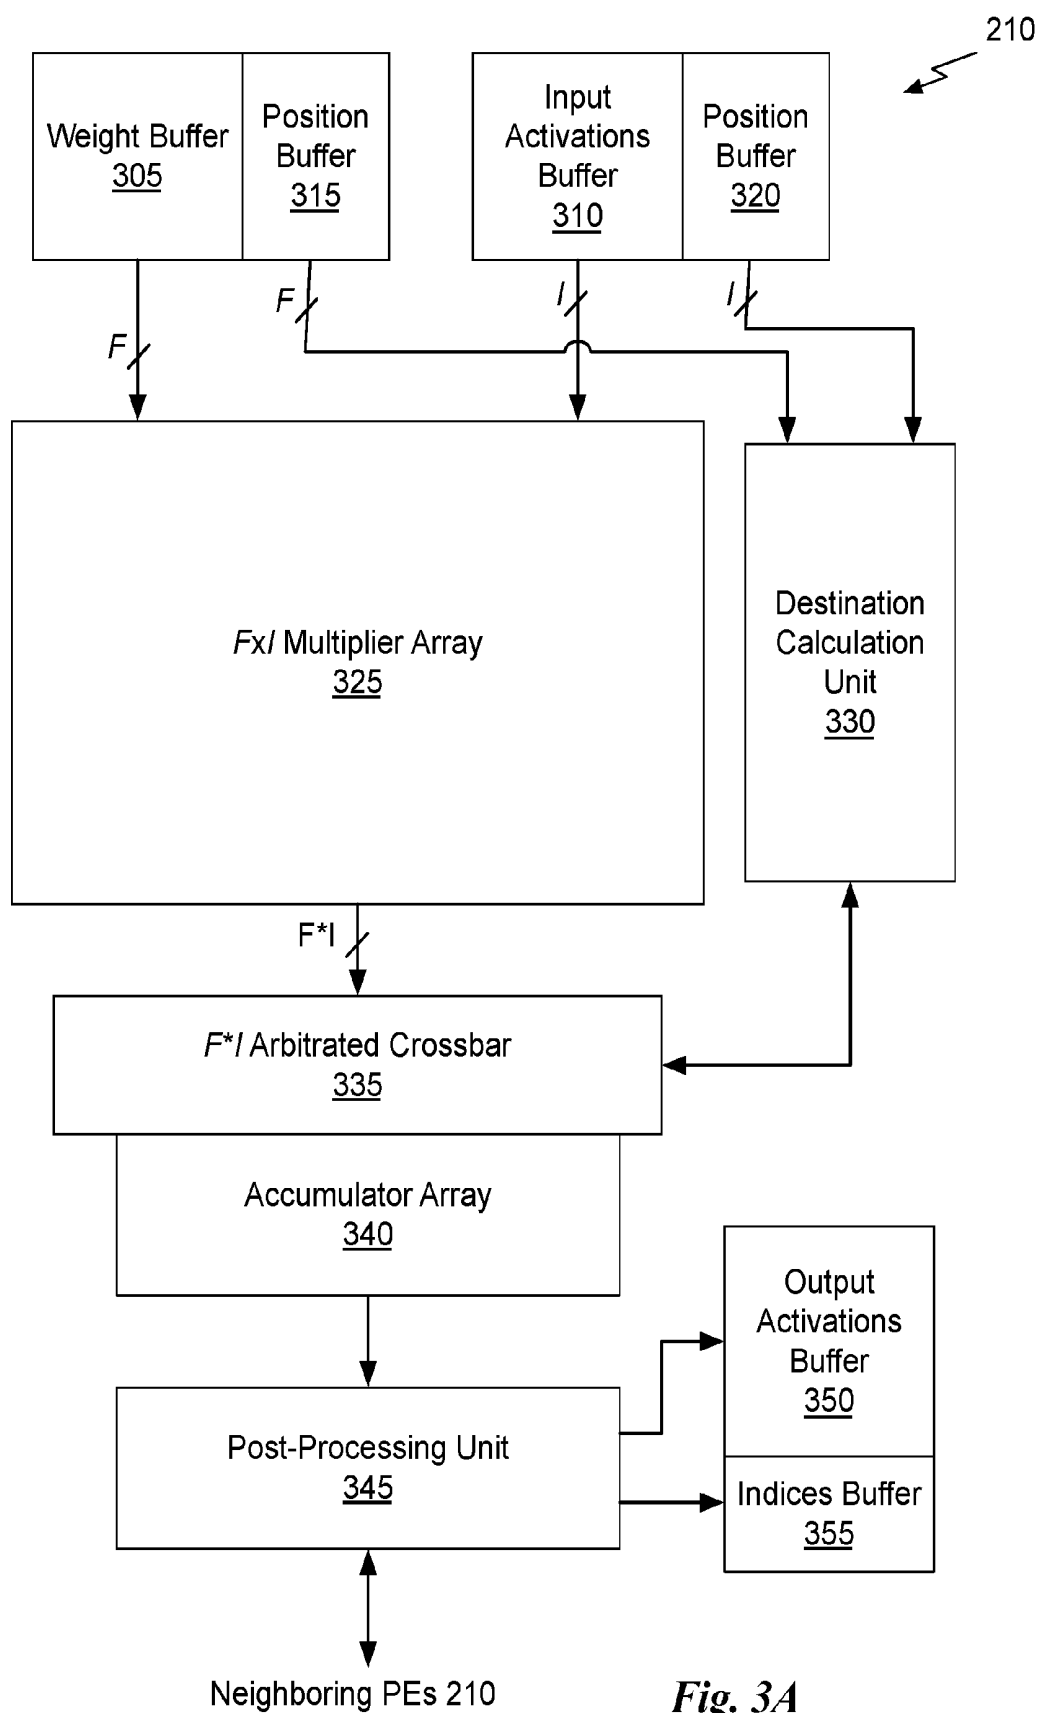


Fig. 3A

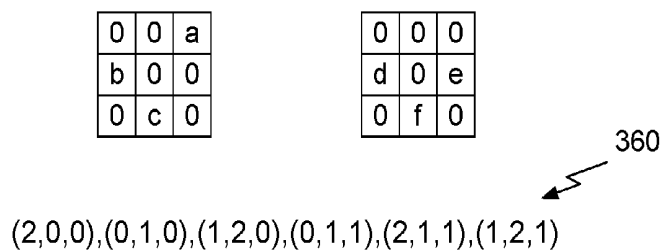


Fig. 3B

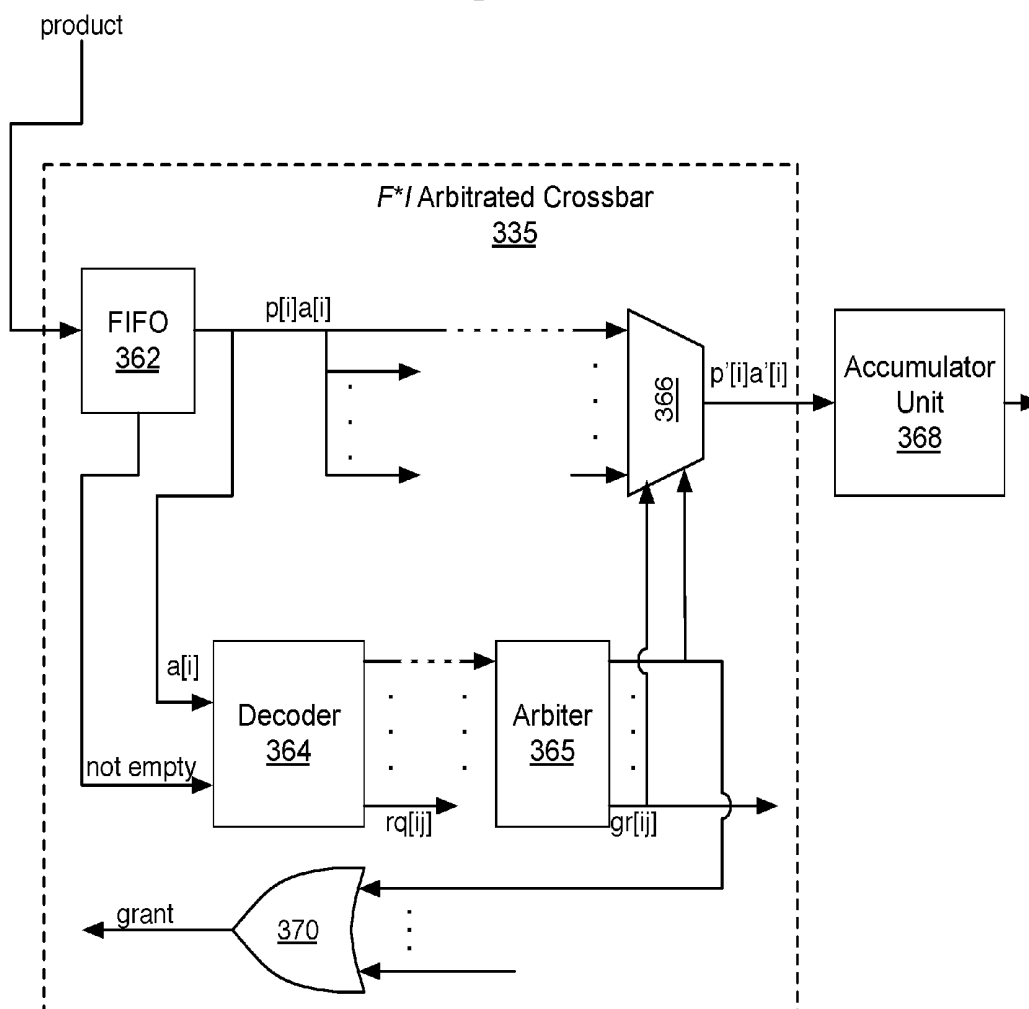


Fig. 3C

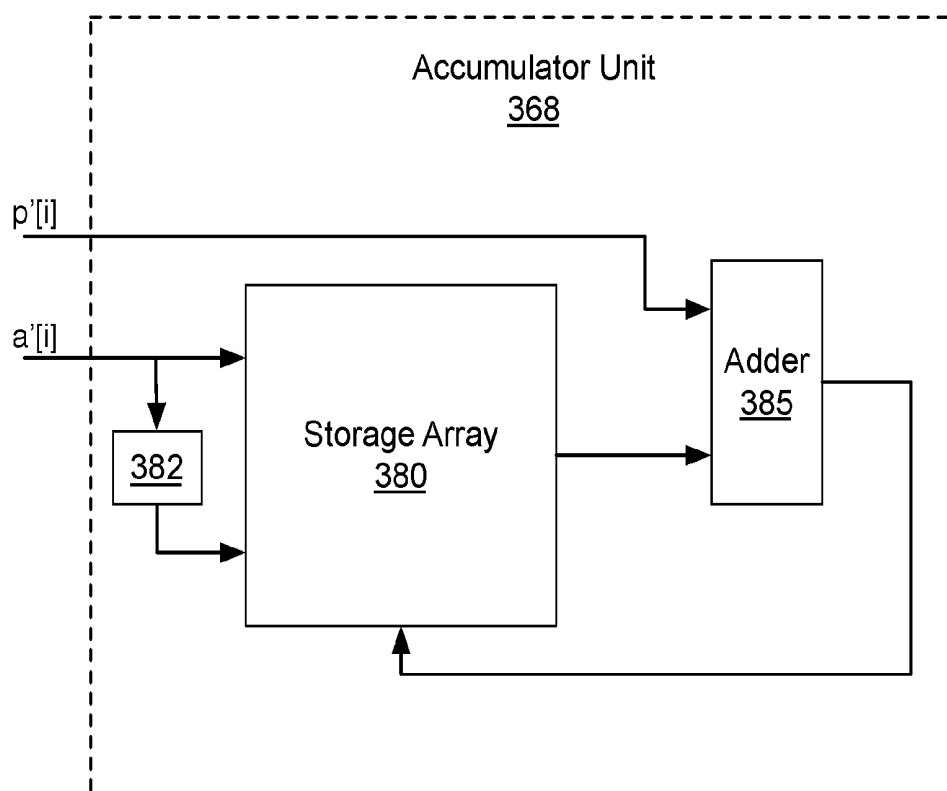
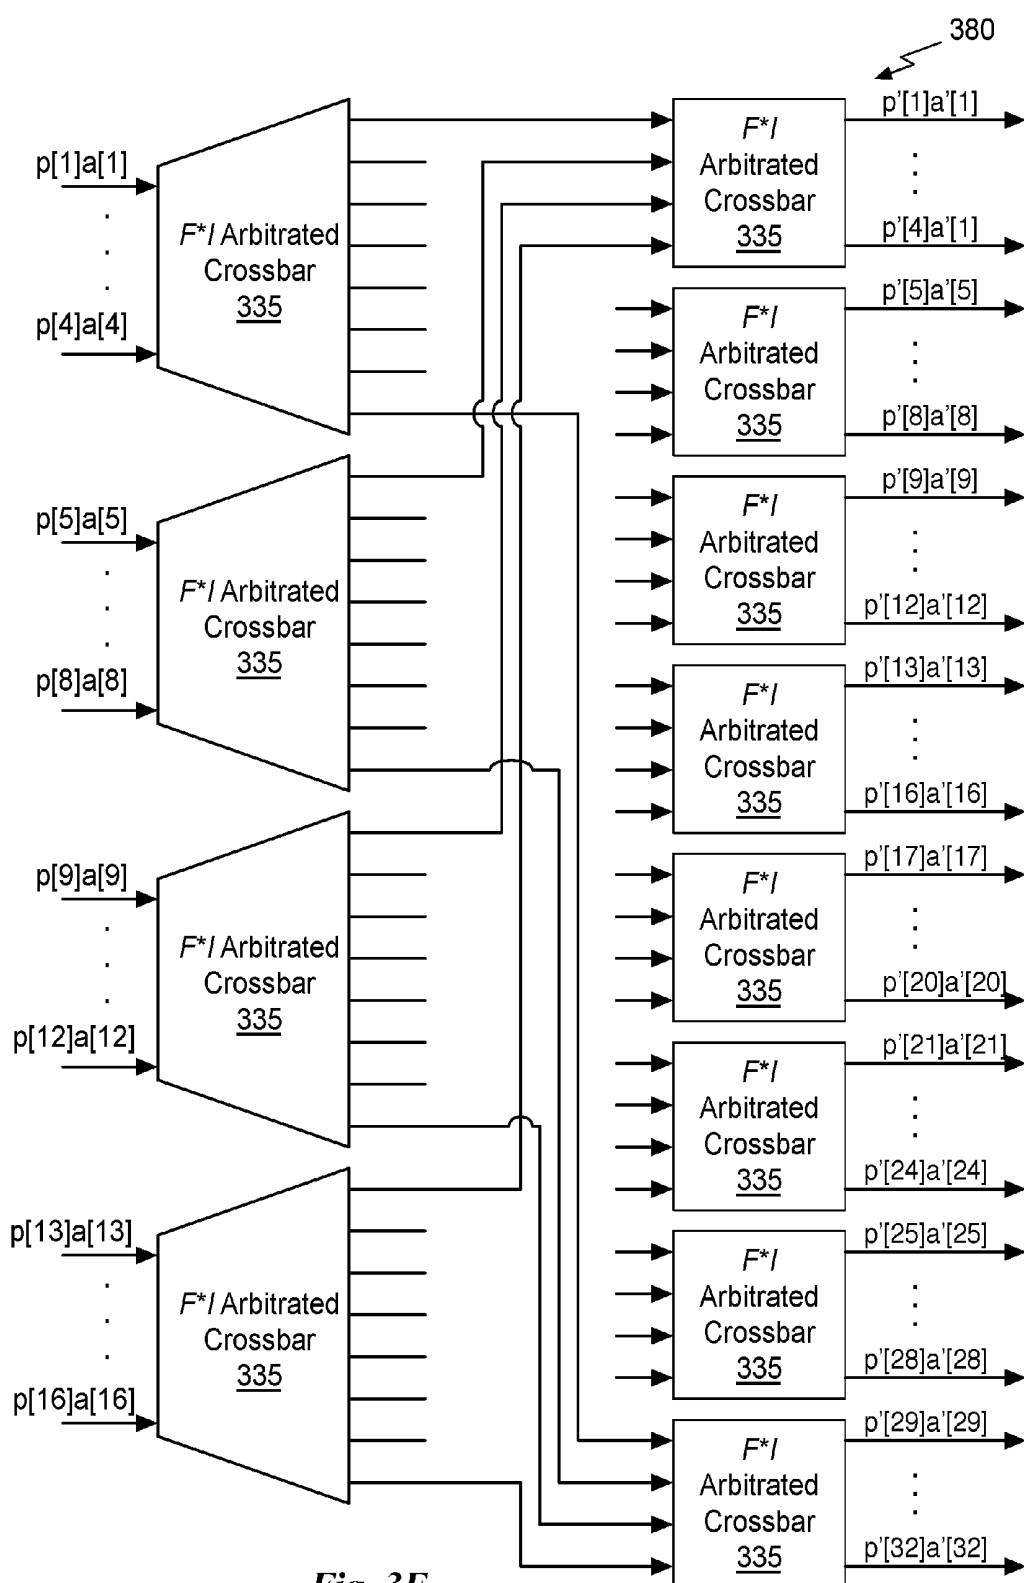
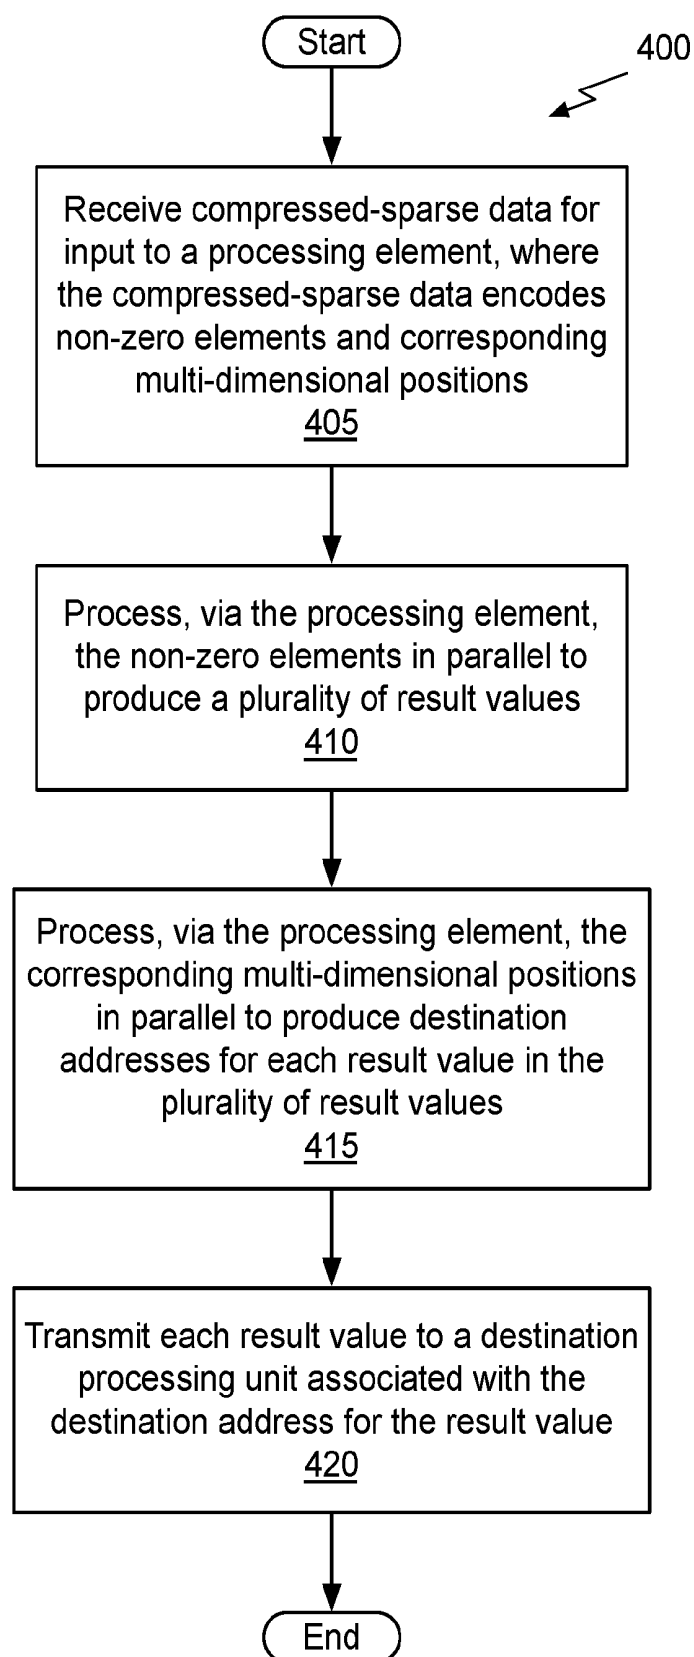


Fig. 3D



*Fig. 4A*

430

0	3	0	6	0	0
0	0	0	0	7	0
4	0	5	0	0	8

Fig. 4B

Code	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Weight	R0	R1	R2	R3	R4	R5	R6	R7	R8	R10	S0	S1	S2	S3	K0	K1
IA	X0	X1	X2	X3	X4	X5	X6	X7	X8	X10	Y16	X24	Y0	Y1	Y2	Y3

Fig. 4C

435

0	3	0	6	7	0	0	9	0	0	0	0	0	0
0	0	0	0	0	0	10	0	11	13	0	14	0	0
4	0	5	0	8	0	0	12	0	0	0	0	0	0

Fig. 4D

440

S1	S4	S1	S0	S0	S5	S2	S1	S1	S1	S4	S1
1,0,0	0,2,0	2,2,0	0,0,1	1,0,1	1,2,1	1,0,2	0,1,2	2,1,2	1,2,2	0,1,3	2,1,3

Fig. 4E

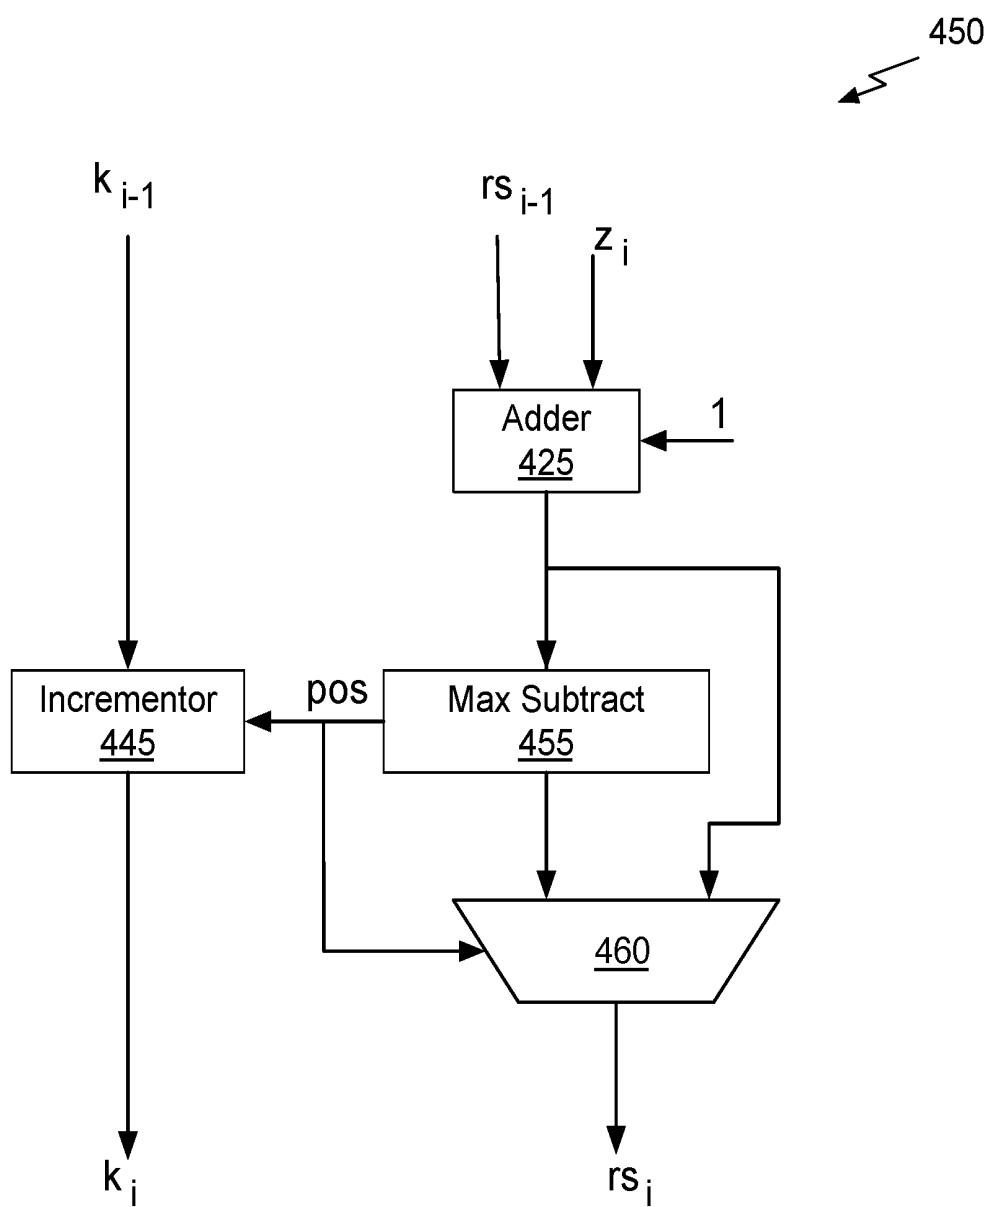


Fig. 4F

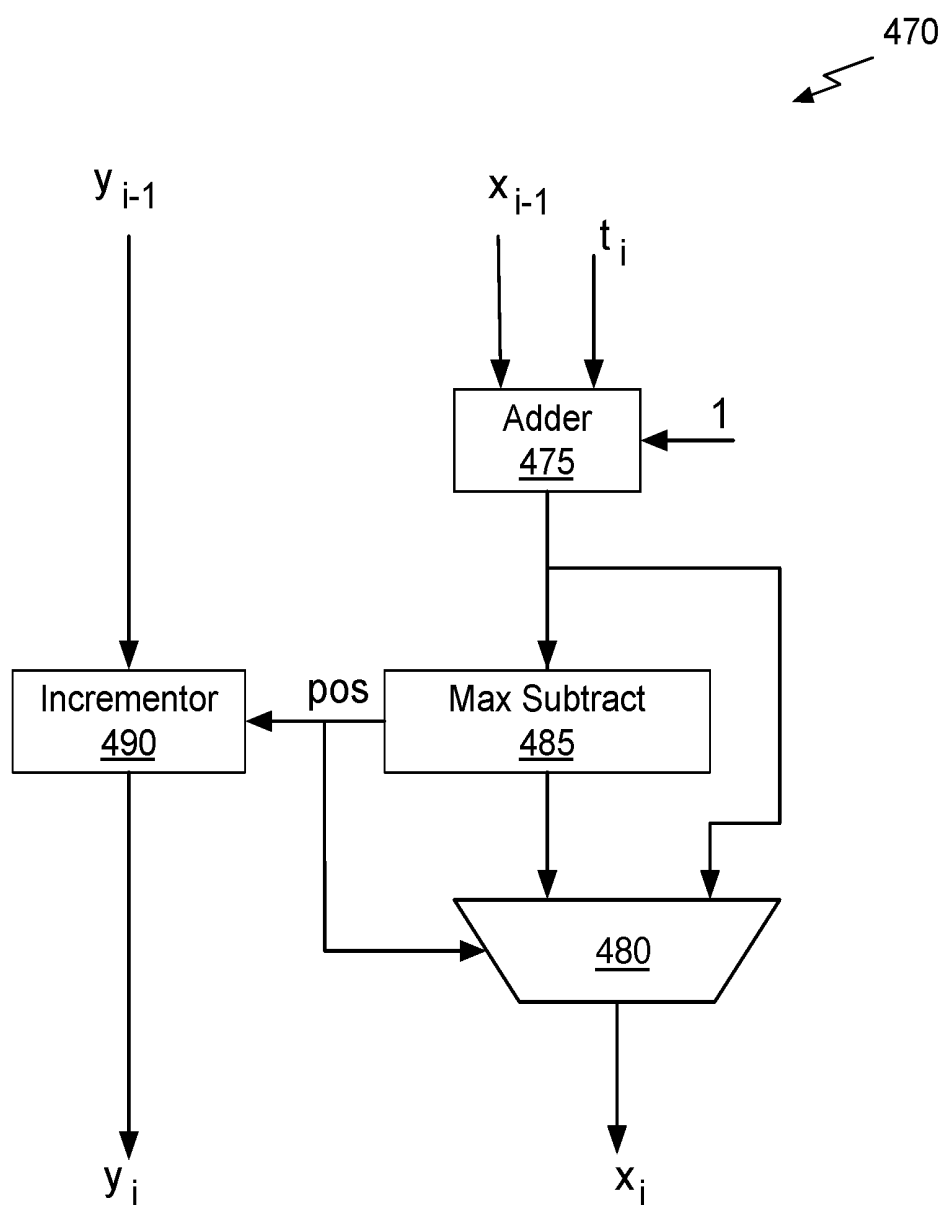


Fig. 4G

505

code	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
value	0	1	2	3	4	5	6	7	12	16	24	32	48	64	96	128

Fig. 5A

510

code	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
value	0	1	2	3	4	5	6	7	Y0	Y1	Y2	Y3	Y4	Y5	Y6	Y7

Fig. 5B

515

code	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
value	0	1	2	3	4	5	6	7	8	9	10	Y0	Y1	Y2	Y3	Y4

Fig. 5C

520

code	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
value	0	1	2	3	4	5	6	7	8	9	10	11	12	13	K0	K1

Fig. 5D

525

code	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
value	0	1	2	3	4	5	6	7	8	9	S0	S1	S2	S3	K0	K1

Fig. 5E

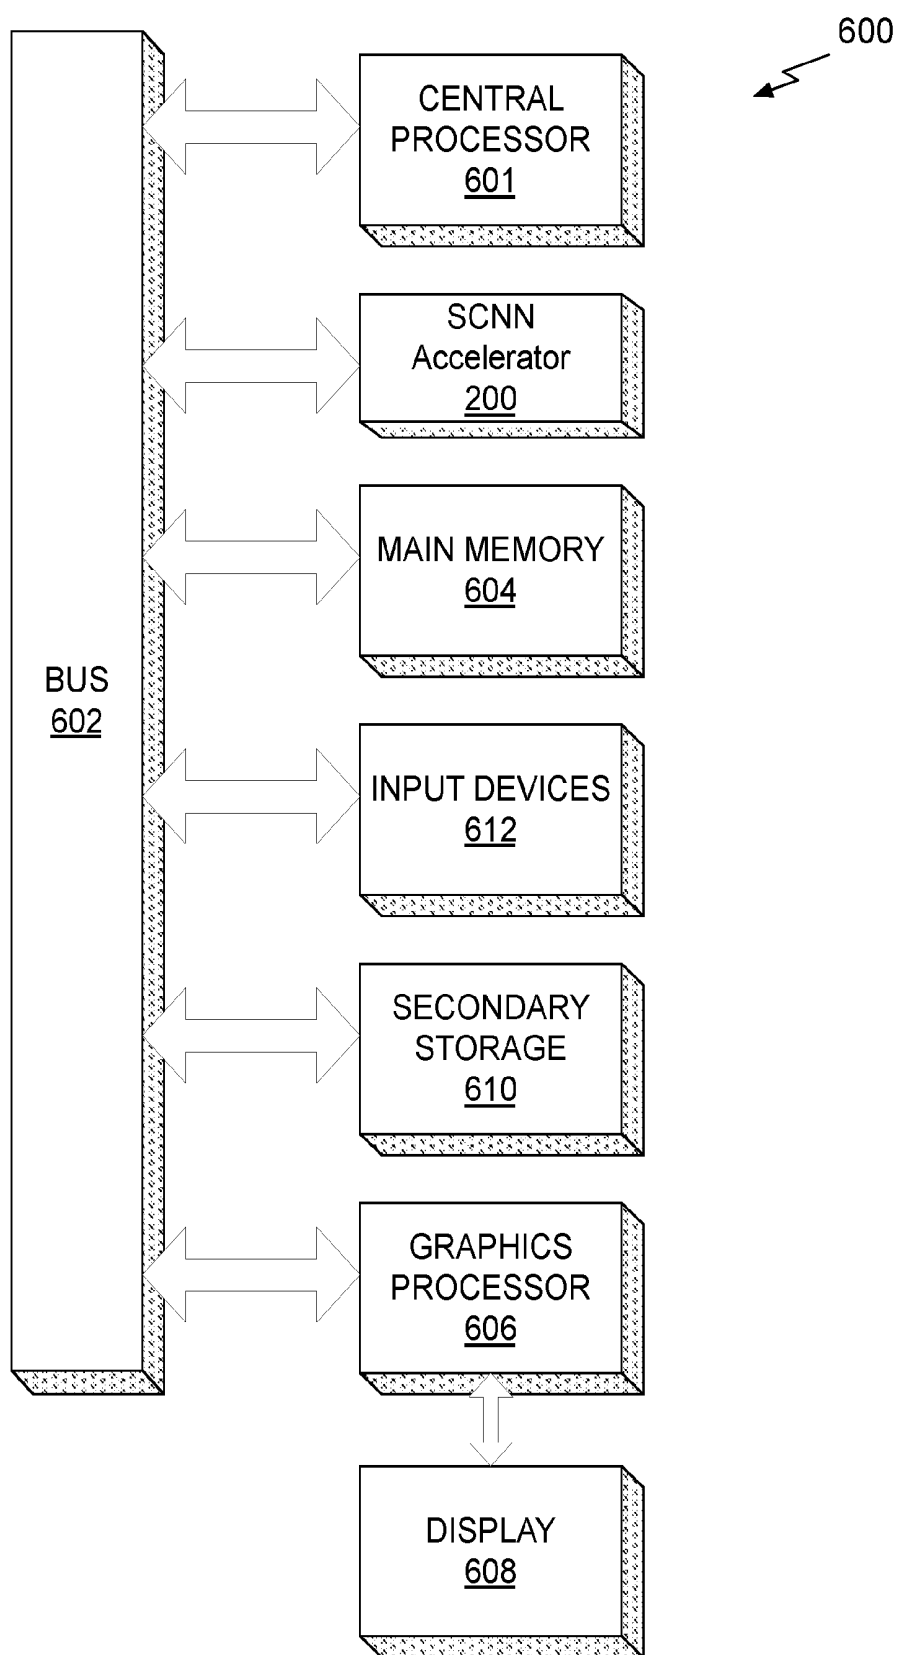


Fig. 6

SPARSE CONVOLUTIONAL NEURAL NETWORK ACCELERATOR

CLAIM OF PRIORITY

[0001] This application claims the benefit of U.S. Provisional Application No. 62/373,919 (Attorney Docket No. NVIDP11374/16-SC-0139-US01) titled "Sparse Convolutional Neural Network Accelerator," filed Aug. 11, 2016, the entire contents of which is incorporated herein by reference.

FIELD OF THE INVENTION

[0002] The present invention relates to convolutional neural networks, and more particularly to a sparse convolutional neural network accelerator.

BACKGROUND

[0003] Driven by the availability of massive data and the computational capability to process it, deep learning has recently emerged as a critical tool for solving complex problems across a wide range of domains, including image recognition, speech processing, natural language processing, language translation, and autonomous vehicles. Convolutional neural networks (CNNs) have become the most popular algorithmic approach for deep learning for many of these domains. High performance and extreme energy efficiency are critical for deployments of CNNs in a wide range of situations, especially mobile platforms such as autonomous vehicles, cameras, and electronic personal assistants.

[0004] Employing CNNs can be decomposed into two tasks: (1) training—in which the parameters of a neural network are learned by observing massive numbers of training examples, and (2) classification—in which a trained neural network is deployed in the field and classifies the observed data. Today, training is often done on graphics processing units (GPUs) or farms of GPUs, while classification depends on the application and can employ central processing units (CPUs), GPUs, field-programmable gate arrays (FPGAs), or application-specific integrated circuits (ASICs).

[0005] During the training process, a deep learning expert will typically architect the network, establishing the number of layers in the neural network, the operation performed by each layer, and the connectivity between layers. Many layers have parameters, typically filter weights, that determine exact computation performed by the layer. The objective of the training process is to learn the filter weights, usually via a stochastic gradient descent-based excursion through the space of weights. The training process typically employs a forward-propagation calculation for each training example, a measurement of the error between the computed and desired output, and then back-propagation through the network to update the weights. Inference has similarities, but only includes the forward-propagation calculation. Nonetheless, the computation requirements for inference can be prohibitively large, particularly with the emergence of deeper networks (hundreds of layers) and larger inputs sets, such as high-definition video. Furthermore, the energy efficiency of this computation is important, especially for mobile platforms, such as autonomous vehicles, cameras, and electronic personal assistants. The computation requirements and energy consumption of a neural network for machine learning presents challenges for mobile platforms.

Thus, there is a need for addressing these issues and/or other issues associated with the prior art.

SUMMARY

[0006] A method, computer program product, and system perform computations using a sparse convolutional neural network accelerator. Compressed-sparse data is received for input to a processing element, wherein the compressed-sparse data encodes non-zero elements and corresponding multi-dimensional positions. The non-zero elements are processed in parallel by the processing element to produce a plurality of result values. The corresponding multi-dimensional positions are processed in parallel by the processing element to produce destination addresses for each result value in the plurality of result values. Each result value is transmitted to a destination accumulator associated with the destination address for the result value.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] FIG. 1 illustrates a flowchart of a method for performing computations using a Sparse Convolutional Neural Network (SCNN) Accelerator, in accordance with one embodiment;

[0008] FIG. 2A illustrates a block diagram of a SCNN accelerator, in accordance with one embodiment;

[0009] FIG. 2B illustrates a conceptual diagram of the organization of input activations and filter weights for processing by the SCNN accelerator shown in FIG. 2A, in accordance with one embodiment;

[0010] FIG. 2C illustrates a block diagram of a processing element, in accordance with one embodiment;

[0011] FIG. 3A illustrates a block diagram of another processing element, in accordance with one embodiment;

[0012] FIG. 3B illustrates two 3×3 weight kernels and positions, in accordance with one embodiment;

[0013] FIG. 3C illustrates a single-stage F*I arbitrated crossbar, in accordance with one embodiment;

[0014] FIG. 3D illustrates an accumulator unit, in accordance with one embodiment;

[0015] FIG. 3E illustrates a two-stage F*I arbitrated crossbar, in accordance with one embodiment;

[0016] FIG. 4A illustrates a flowchart of a method for compressing weight and input activation values, in accordance with one embodiment;

[0017] FIG. 4B illustrates a tile of weight values for two output channels, in accordance with one embodiment;

[0018] FIG. 4C illustrates a coding scheme for weights and input activations (IA), in accordance with one embodiment;

[0019] FIG. 4D illustrates weight values for four 3×3 convolution kernels, in accordance with one embodiment;

[0020] FIG. 4E illustrates an encoding of the positions for the weight values in the four 3×3 convolution kernels shown in FIG. 4D, in accordance with one embodiment;

[0021] FIG. 4F shows a block diagram for determining the (r,s) weight coordinates, in accordance with one embodiment;

[0022] FIG. 4G shows a block diagram for determining the (x,y) input activation coordinates, in accordance with one embodiment;

[0023] FIG. 5A illustrates a non-linear coding scheme for input activation zero-count values, in accordance with one embodiment;

[0024] FIG. 5B illustrates another coding scheme for input activation zero-count values, in accordance with one embodiment;

[0025] FIG. 5C illustrates another coding scheme for input activation zero-count values, in accordance with one embodiment;

[0026] FIG. 5D illustrates another coding scheme for weight zero-count values, in accordance with one embodiment;

[0027] FIG. 5E illustrates another coding scheme for weight zero-count values, in accordance with one embodiment; and

[0028] FIG. 6 illustrates an exemplary system in which the various architecture and/or functionality of the various previous embodiments may be implemented.

DETAILED DESCRIPTION

[0029] Neural networks typically have significant redundancy and can be pruned dramatically during training without substantively affecting accuracy of the neural network. The number of weights that can be eliminated varies widely across the layers of the neural network, but typically ranges from 20% to 80%. Eliminating weights results in a neural network with a substantial number of zero values, which can potentially reduce the computational requirements of inference.

[0030] The inference computation also offers a further optimization opportunity. In particular, many neural networks employ the rectified linear unit (ReLU) function that clamps all negative activation values to zero as a non-linear operator. The activations are the output values of an individual layer that are passed as inputs to the next layer. For typical data sets, 50-70% of the activations are clamped to zero. Since the multiplication of weights and activations is the key computation for inference, the combination of activations that are zero and weights that are zero can reduce the amount of computation required by over an order of magnitude. A sparse CNN (SCNN) accelerator architecture described herein, exploits weight and/or activation sparsity to reduce energy consumption and improve processing throughput. The SCNN accelerator architecture couples an algorithmic dataflow that eliminates all multiplications with a zero operand while employing a compressed representation of both weights and activations through almost the entire computation. In one embodiment, each non-zero weight and activation value is represented by a (value, position) pair.

[0031] Additional benefits can be achieved by a compressed or compact encoding for sparse weights and/or activations that include several zeros, thus allowing more weight and/or activation values to fit in on-chip random access memory (RAM) and reducing the number of energy-costly dynamic random access memory (DRAM) accesses to read activations and weights. Furthermore, transmitting the compact encoding may reduce the number of transitions on buses, further reducing energy consumption. Finally, only the non-zero elements of weights and input activations are provided as operands to the multipliers, ensuring that each multiplier within a processing element (PE) generates a product that affects an output activation value. In the context of the following description, activation refers to an input activation and/or an output activation. In the context of the following description, the weight and activation values are multi-bit values representing zero, positive values, or nega-

tive values. In the context of the following description, the positions are coordinates in an N dimensional space.

[0032] FIG. 1 illustrates a flowchart of a method 100 for performing computations using a SCNN accelerator, in accordance with one embodiment. Although method 100 is described in the context of a processing element within a SCNN accelerator, the method 100 may also be performed by a program, custom circuitry, or by a combination of custom circuitry and a program. Furthermore, persons of ordinary skill in the art will understand that any system that performs method 100 is within the scope and spirit of embodiments of the present invention.

[0033] At step 105, a first vector comprising only non-zero weight values and first associated positions of the non-zero weight values within a three-dimensional (3D) space are received. In one embodiment, the first vector is received from a memory. In one embodiment, the first vector is received by a processing element (PE) within a SCNN accelerator, such as the SCNN accelerator 200 described in conjunction with FIG. 2A.

[0034] At step 110, a second vector comprising only non-zero input activation values and second associated positions of the non-zero input activation values within a two-dimensional (2D) space are received. In one embodiment, the second vector is received from a memory. In one embodiment, the second vector is received by a PE within a SCNN accelerator, such as the SCNN accelerator 200 described in conjunction with FIG. 2A. In one embodiment, the second vector is generated by the SCNN accelerator 200 during processing of a previous layer of a neural network.

[0035] At step 115, each one of the non-zero weight values is multiplied with every one of the non-zero input activation values, within a multiplier array, to produce a third vector of products. At step 120, the first associated positions are combined with the second associated positions to produce a fourth vector of positions, where each position in the fourth vector is associated with a respective product in the third vector. In one embodiment, the combining comprises performing a vector addition to sum coordinates in the first associated positions with the second associated positions to produce the fourth vector of positions, where each position in the fourth vector is associated with a respective product in the third vector.

[0036] At step 125, the third vector is transmitted to an accumulator array, where each one of the products in the third vector is transmitted to an adder in the accumulator array that is configured to generate an output activation value at the position associated with the product. In one embodiment, the third vector is transmitted through an array of buffers in the accumulator array, where each one of the buffers is coupled to an input of one of the adders in the accumulator array.

[0037] More illustrative information will now be set forth regarding various optional architectures and features with which the foregoing framework may or may not be implemented, per the desires of the user. It should be strongly noted that the following information is set forth for illustrative purposes and should not be construed as limiting in any manner. Any of the following features may be optionally incorporated with or without the exclusion of other features described.

Sparse Convolutional Neural Network Accelerator

[0038] FIG. 2A illustrates a block diagram of the SCNN 200, in accordance with one embodiment. SCNN 200 couples an algorithmic dataflow that eliminates all multiplications with a zero operand while transmitting a compact representation of weights and/or input activations between memory and logic blocks within the SCNN 200. The SCNN 200 includes a memory interface 205, layer sequencer 215, and an array of processing elements (PEs) 210.

[0039] The memory interface 205 reads weight and activation data from a memory coupled to the SCNN 200 the memory interface 205 may also write weight and/or activation data from the SCNN 200 to the memory. In one embodiment, all of the activation data is stored within the PEs 210, so that only weight data is access through the memory interface 205. The weight and/or activation data may be stored in the memory in a compact format or an expanded format. The compact format may comprise vectors including only non-zero elements (weights or activations) and positions associated with the non-zero elements.

[0040] The memory may be implemented using dynamic random access memory (DRAM), or the like. In one embodiment, the memory interface 205 or the PEs 210 are configured to compact multi-bit data, such as the weights, input activations, and output activations. The layer sequencer 215 controls the reading of the memory to obtain the compact input activations and compact weights. The compact input activations and compact weights may be stored within the memory interface 205 before being transmitted to the PEs 210.

[0041] In one embodiment, the compact activations and the compact weights are data sequences encoded as non-zero elements and positions. In one embodiment, the non-zero elements and positions are each encoded as a (value, position) pair. If needed, the compact activations and compact weights may be expanded to generate data sequences of weights and activations that include multi-bit zero and non-zero elements. Importantly, when the weights and input activations are in compact form, only non-zero weights and input activations are transferred from the memory interface 205 to the PEs 210. In one embodiment, the non-zero elements are 8 bits and the positions are 4 bits. However, the non-zero elements may be more than 8 bits or fewer bits than 8 bits and the positions may be more than 4 bits or fewer than 4 bits.

[0042] The layer sequencer 215 reads the weights and outputs weight vectors to be multiplied by the PEs 210. In one embodiment, the weights are in compact form and are read from off-chip DRAM only once and stored within the SCNN accelerator 200. In one embodiment, the layer sequencer 215 broadcasts a weight vector to each PE 210 and sequences through multiple activation vectors before broadcasting another weight vector. In one embodiment, the layer sequencer 215 broadcasts an input activation vector to each PE 210 and sequences through multiple weight vectors before broadcasting another input activation vector. Products generated by the multipliers within each PE 210 are accumulated to produce intermediate values (e.g., partial sums) that become the output activations after one or more iterations. When the output activations for a neural network layer have been computed and stored in an output activation buffer, the layer sequencer 215 may proceed to process a next layer by applying the output activations as input activations.

[0043] Each PE 210 includes a multiplier array that accepts a vector of weights (weight vector) and a vector of input activations (activation vector), where each multiplier within the array is configured to generate a product from one input activation value in the activation vector and one weight in the weight vector. The weights and input activations in the vectors can all be multiplied by one another in the manner of a Cartesian product. For example, if the input vectors are a,b,c,d and p,q,r,s respectively, the output is a 16-vector with values $a*p$, $a*q$, $a*r$, $a*s$, $b*p$, $b*q$, $b*r$, $b*s$, $c*p$, $c*q$, $c*r$, $c*s$, $d*p$, $d*q$, $d*r$, and $d*s$.

[0044] Importantly, only non-zero weights and input activations are transmitted to the multiplier array within each PE 210. Additionally, the input activation vectors may be reused within each PE 210 in an input stationary fashion against a number of weight vectors to reduce data accesses. The products generated by the multipliers are then summed together to generate the partial sums and the output activations. However, because the zero values have been removed, the mapping of the products to accumulators may vary for each product that is generated within the multiplier array. For example, in a conventional implementation where the zero values are maintained, the products generated during a clock cycle may be summed together to produce a partial sum. In contrast, the products generated during a clock cycle within a PE 210 may not necessarily be summed together to produce a partial sum. Therefore, output coordinates associated with each multiplication are tracked within the PE 210 and an output position (defined by the output coordinates) and product are provided to a scatter accumulator array for summing. The scatter accumulator allows for transmission of any product to any adder, based on the output position associated with the product. In one embodiment, the PEs 210 are configured to perform convolution operations on the weights and input activations. Summing the products in the adders completes the convolution operation and generates the output activations.

[0045] The SCNN 200 may be configured to implement CNN algorithms that are a cascaded set of pattern recognition filters trained with supervision. A CNN consists of a series of layers, which include convolutional layers, non-linear scalar operator layers, and layers that downsample the intermediate data, for example by pooling. The convolutional layers represent the core of the CNN computation and are characterized by a set of filters that are usually 1×1 or 3×3 , and occasionally 5×5 or larger. The values of these filters are the weights that are trained using a training set for the network. Some deep neural networks (DNNs) also include fully-connected layers, typically toward the end of the DNN. During classification, a new image (in the case of image recognition) is presented to the neural network, which classifies images into the training categories by computing in succession each of the layers in the neural network. The SCNN 200 accelerates the convolutional layers, receiving weights and input activations and generating output activations.

[0046] Sparsity in a layer of a CNN is defined as the fraction of zeros in the layer's weight and input activation matrices. The primary technique for creating weight sparsity is to prune the network during training. In one embodiment, any weight with an absolute value that is close to zero (e.g. below a defined threshold) is set to zero. The pruning process has the effect of removing weights from the filters, and sometimes even forcing an output activation to always

equal zero. The remaining network may be retrained, to regain the accuracy lost through naïve pruning. The result is a smaller network with accuracy extremely close to the original network. The process can be iteratively repeated to reduce network size while maintaining accuracy.

[0047] Activation sparsity occurs dynamically during inference and is highly dependent on the data being processed. In one embodiment, activations having negative values are clamped to zero. In one embodiment, input activations having an absolute value below a defined threshold are set to zero.

[0048] In one embodiment, a compaction engine within the PE 210 sets output activations having an absolute value below a defined threshold to zero. If the activations are in a compacted format, the compaction engine reformats the activations as needed after setting one or more activations to zero to produce compacted activations. After completing computation of the output activations for a layer of a CNN, each element in the output activation matrices that is below a threshold value may be set to zero before the output activation data is passed to the next layer.

SCNN Calculation Dataflow

[0049] The core operation in a CNN layer is a two-dimensional sliding-window convolution of an $R \times S$ element filter over a $W \times H$ element input activation plane to produce a $W \times H$ element output activation plane. There can be multiple (C) input activation planes, which are referred to as input channels. A distinct filter is applied to each input activation channel, and the filter output for each of the C channels are accumulated together element-wise into a single output activation plane. Multiple filters (K) can be applied to the same body of input activations to produce K output channels of output activations. Finally, a batch of length N of groups of C channels of input activation planes can be applied to the same volume of filter weights.

[0050] FIG. 2B illustrates input activations, weights, and output activations for a single CNN layer, in accordance with one embodiment. The set of computations for the complete layer can be formulated as a loop nest over the seven variables (N, K, C, W, H, R, and S). Because multiply-add operations are associative (modulo rounding errors, which are ignored in the context of the following description), all permutations of the seven loop variables are legal. TABLE 1 shows an example loop nest based on one such permutation. The nest may be concisely described as $N \rightarrow K \rightarrow C \rightarrow W \rightarrow H \rightarrow R \rightarrow S$. Each point in the seven-dimensional space formed from the variables represents a single multiply-accumulate operation. Note that for the remainder of the description, a batch size of 1 is assumed, which is a common batch size for inferencing tasks.

TABLE 1

seven dimensional CNN loop nest	
for n = 1 to N	
for k = 1 to K	
for c = 1 to C	
for x = 1 to W	
for y = 1 to H	
for r = 1 to R	
for s = 1 to S	
out[n][k][x][y] +=	

TABLE 1-continued

seven dimensional CNN loop nest	
in[n][c][x+r-1][y+s-1] *	
filter[k][c][r][s];	

[0051] The simple loop nest shown in TABLE 1 can be transformed in numerous ways to capture different reuse patterns of the activations and weights and to map the computation to a hardware accelerator implementation, such as the SCNN accelerator 200. A CNN's dataflow defines how the loops are ordered, partitioned, and parallelized and the choice of dataflow may have a significant impact on the area and energy-efficiency of an architecture.

[0052] While the concept of dataflow has been studied for dense architectures, sparse architectures can also employ various alternative dataflows, each with its own set of trade-offs. One such specific dataflow, described herein, is a sparse planar-tiled input-stationary (PTIS-sparse). PTIS-sparse enables reuse patterns that exploit the characteristics of sparse weights and activations. First, an equivalent dense dataflow (PTIS-dense) is described to explain the decomposition of the computations. Then, the specific features for PTIS-sparse are described.

[0053] FIG. 2C illustrates a PE 220, in accordance with one embodiment. To understand the temporal component of the PTIS-dense dataflow, the operation of the PE 220 is described. PTIS employs an input-stationary computation order in which an input activation is held stationary at the computation units as the input activation is multiplied by all the filter weights needed to make all contributions of the input activation to each of the K output channels (a $K \times R \times S$ subvolume). Thus, each input activation will contribute to a volume of $K \times R \times S$ output activations. The input-stationary computation order maximizes the reuse of the input activations, while paying a cost to stream the weights to the PEs 220. Accommodating multiple input channels (C) adds an additional outer loop and results in the loop nest $C \rightarrow W \rightarrow H \rightarrow R \rightarrow S$.

[0054] The PTIS -dense dataflow relies on input buffers, weight buffer 230 and input activations buffer 235, for storing weights and input activations, respectively. An accumulator buffer 250 stores the partial sums of the output activations. A read-add-write operation is performed for every access to a previously-written partial sum in the accumulator buffer 250. The accumulator buffer 250 combined with an attached adder unit 255 forms an accumulation unit 245.

[0055] Parameters of contemporary networks cause the weight buffer 230 and input activations buffer 235 to be large and energy-expensive to access. The input-stationary temporal loop nest amortizes the energy cost of accessing the weight buffer 230 and input activations buffer 235 over multiple weight buffer 235 and accumulation buffer 250 accesses. More precisely, the register in which the stationary input is held over $K \times R \times S$ iterations serves as an inner buffer, filtering accesses to the larger input buffer (e.g., weight buffer 230 or input activations buffer 235).

[0056] Unfortunately, the stationary feature of input activations comes at the cost of more accesses to the weights in the weight buffer 230 (or the memory) and partial sums in the accumulator buffer 250. Blocking the weights and partial sums in the output channel (K) dimension can increase reuse of the weight buffer 230 and accumulator buffer 250,

improving energy efficiency. The output channel variable (K) may be factored into K_c (called a output-channel group) and K/K_c is the number of output-channel groups. In one embodiment, only the weights and outputs for a single output-channel group are stored at a time inside the weight buffer 230 and the accumulation buffer 250. Thus, the sub-volumes that are housed in buffers at the computation unit are:

[0057] Weights: $K_c \times R \times S$

[0058] Input Activations: $C \times W \times H$

[0059] Partial Sums: $K_c \times W \times H$

An outer loop over all the K/K_c output-channel tiles results in the complete loop nest $K/K_c \rightarrow C \rightarrow W \rightarrow H \rightarrow K_c \rightarrow R \rightarrow S$. Note that each iteration of the outer loop will require the weight buffer 230 to be refilled and the accumulation buffer 250 must be drained and cleared, while the contents of the input activations buffer 235 will be fully reused because the same input activations are used across all output channels. [0060] To exploit the parallelism of many multipliers within a PE 220, a vector of F filter-weights may be fetched from the weight buffer 230 and a vector of I inputs may be fetched from the input activations buffer 235. The vectors are delivered to an array of $F \times I$ multipliers 240 to compute a full Cartesian product of output partial sums. Each product yields a useful partial sum such that no extraneous fetches or computations are performed. PTIS-sparse exploits this same property to perform efficient computations on compressed-sparse weights and input activations.

[0061] The multiplier outputs (e.g., products) are sent to the accumulation unit 245, which updates the partial sums stored in the accumulation buffer 250. Each product is accumulated with a partial sum at the output coordinates in the output activation space that matches (i.e., equals) a position associated with the product. The output positions for the products are computed in parallel with the products (not shown in FIG. 2C). In one embodiment, coordinates defining the output positions are computed by a state machine in the accumulation unit 245. The number of adders in the adder unit 255 does not necessarily equal the number of multipliers in the $F \times I$ multiplier array 240. However, the accumulation unit 245 must employ at least $F \times I$ adders in the adder unit 255 to match the throughput of the $F \times I$ multiplier array 240.

[0062] TABLE 2 shows pseudo-code for the PTIS-dense dataflow, including blocking in the K dimension (A,C), fetching vectors of input activations and weights (B,D), and computing the Cartesian product in parallel (E,F). Note that this PTIS-dense dataflow is simply a reordered, partitioned and parallelized version of the pseudo-code shown in TABLE 1.

TABLE 2

pseudo-code for the PTIS-dense dataflow	
	BUFFER wt_buf[C][Kc][R][S/F][F]; BUFFER in_buf[C][Wt][Ht/I][I]; BUFFER acc_buf[Kc][Wt][Ht]; BUFFER out_buf[K/Kc][Kc][Wt][Ht];
(A)	for $k' = 1$ to K/Kc { for $c = 1$ to C for $x = 1$ to Wt for $y = 1$ to Ht/I { in[1:I] = in_buf[c][x][y][1:I];
(B)	in[1:I] = in_buf[c][x][y][1:I];

TABLE 2-continued

pseudo-code for the PTIS-dense dataflow	
(C)	for $k = 1$ to Kc for $r = 1$ to R for $s = 1$ to S/F {
(D)	wt[1:F] = wt_buf[c][k][r][s][1:F];
(E)	parallel_for ($i = 1$ to I) * ($f = 1$ to F)
(F)	acc_buf[k][x-r+1][y-s+1] += in[i]*wt[f];
	} out_buf[k][1:K][1:Wt][1:Ht] = acc_buf[1:K][1:Wt][1:Ht];
	}

Note that the output positions associated with an output buffer (out_buf) may be computed using the loop indices as shown in section (F) of TABLE 2.

[0063] To scale beyond the practical limits of multiplier count and buffer sizes within a PE 220, a tiling strategy may be employed to spread the work across an array of PEs 210 so that each PE 210 can operate independently. In one embodiment of the PTIS-dense technique, the $W \times H$ element activation plane is partitioned into smaller $W_t \times H_t$ element tiles that are distributed across the PEs 210 within the SCNN accelerator 200. Each tile extends fully into the input-channel dimension C , resulting in an input-activation volume of $C \times W_t \times H_t$ assigned to each PE 210. Weights are broadcast to the PEs 210 and each PE 210 operates on an exclusive subset of the input and output activation space. In other words, there is no duplication of input activations or output activations between the PEs 210.

[0064] Unfortunately, strictly partitioning both input and output activations into $W_t \times H_t$ tiles does not work because the sliding-window nature of the convolution operation introduces cross-tile dependencies at tile edges. These dependencies are called halos. Halos can be resolved in two ways. The first technique for handling halos is to size the input activations buffer 235 in each PE 210 to be slightly larger than $C \times W_t \times H_t$ to accommodate the halos. The halo input activation values are replicated across adjacent PEs 210, but the computed products are strictly private to each PE 210. Replicated input activation values can be multicast when the input activation values are stored into the input activations buffer 235. The second technique for handling halos is to size the accumulation buffer in each PE 210 to be slightly larger than $K_c \times W \times H$ to accommodate the halos. The halos now contain incomplete partial sums that must be communicated to neighbor PEs 210 for accumulation. In one embodiment, the communication between neighboring PEs 210 occurs at the end of computing each output-channel group.

[0065] The PTIS-sparse technique is a natural extension of the PTIS-dense technique, with the PTIS-sparse technique exploiting sparsity in the weights and activations. The PTIS-sparse dataflow is specifically designed to operate on compressed-sparse (i.e., compacted) encodings of the weights and input activations and to produce a compressed-sparse encoding of the output activations. At a CNN layer boundary, the output activations of the previous layer become the input activations of the next layer. The specific format used to generate the compressed-sparse encoded data is orthogonal to the sparse architecture itself. What is key is that decoding a sparse format ultimately yields a non-zero data value and a position indicating the coordinates of the

value in the weight or input activation matrices. In one embodiment, the position is defined by an index or an address, such as an address corresponding to one of the accumulation buffers 250 or adder units 255.

[0066] FIG. 3A illustrates a block diagram of a PE 210, in accordance with one embodiment. The PE 210 is configured to support the PTIS-sparse dataflow. Like the PE 220 shown in FIG. 2C, the PE 210 includes a weight buffer 305, an input activations buffer 310, and an FxI multiplier array 325. Parallelism within a PE 210 is accomplished by processing a vector of F non-zero filter weights a vector of I non-zero input activations in within the FxI multiplier array 325. FxI products are generated each processing cycle by each PE 210 in the SCNN accelerator 200. In one embodiment $F=I=4$. In other embodiments, F and I may be any positive integer and the value of F may be greater than or less than I. The values of F and I may each be tuned to balance overall performance and circuit area. With typical density values of 30% for both weights and activations, 16 multiplies of the compressed sparse weight and input activation values is equivalent to 178 multiplies in a dense accelerator that processes weight and input activation values including zeros.

[0067] The accumulator array 340 may include one or more accumulation buffers and adders to store the products generated in the multiplier array 325 and sum the products into the partial sums. The PE 210 also includes position buffers 315 and 320, indices buffer 355, destination calculation unit 330, F*I arbitrated crossbar 335, and a post-processing unit 345.

[0068] To facilitate easier decoding of the compressed-sparse data, weights are grouped into compressed-sparse blocks at the granularity of an output-channel group, with $K_c \times R \times S$ weights encoded into one compressed-sparse block. Likewise, input activations are encoded at the granularity of input channels, with a block of $W_c \times H_c$ encoded into one compressed-sparse block. At each access, the weight buffer 305 and the position buffer 315 deliver a vector of F non-zero filter weights along with the associated positions (e.g. coordinates) within the $K_c \times R \times S$ region, respectively. Similarly, the input activations buffer 310 and position buffer 320 deliver a vector of I non-zero input activations and the associated positions (e.g., coordinates) within the $W_c \times H_c$ region, respectively. Similar to the PTIS-dense dataflow, the FxI multiplier array 325 computes the full cross-product of FxI partial sum outputs, with no extraneous computations. Unlike a dense architecture that includes zero values, output coordinates defining the output positions are not derived from loop indices in a state machine, but are instead derived from the positions (e.g. coordinates) of the non-zero elements embedded in the compressed format.

[0069] Even though calculating the output positions of the products is not difficult, unlike the PTIS-dense technique, the products are not typically contiguous using the PTIS-sparse technique. Thus, the products generated by the FxI multiplier array 325 need to be scattered to discontinuous positions within the $K_c \times W_c \times H_c$ output space. Because any partial sum in the output space can be zero, the accumulator array 340 stores data in a dense format that may include both non-zero and zero values. In fact, output activations will probabilistically have high density even with a very low density (i.e., high sparsity) of weights and input activations, until the output activations pass through a ReLU operation.

[0070] To accommodate the accumulation of sparse partial sums, the monolithic $K_c \times W_c \times H_c$ accumulation buffer 250 used in the PTIS-dense dataflow is modified to be a distributed array of smaller accumulation buffers accessed via a scatter network which can be implemented as a crossbar switch, such as the FxI arbitrated crossbar 335. The FxI arbitrated crossbar 335 routes FxI products to an array of A accumulator units based on the output positions associated with each product. The positions may be translated to form an address. A particular product is transmitted to the one accumulator unit in the accumulator array 340 that is configured to compute the output activation for the position associated with the product. Taken together, a scatter accumulator array comprising the FxI arbitrated crossbar 335 and accumulator array 340 is associated with a $K_c \times W_c \times H_c$ address range. The address space is distributed across the A accumulator units and each accumulator unit includes a bank of addressable storage and an adder to accumulate a partial sum (when processing of a tile is complete, the partial sum is an output activation) for the output position.

[0071] The PTIS-sparse technique can be implemented via small adjustments to the pseudo-code shown in TABLE 2. Instead of a dense vector fetches, (B) and (D) are modified fetch the compressed sparse input activations and weights, respectively. In addition, the positions of the non-zero elements in the compressed-sparse form of the data structures are fetched from the respective buffers (not shown in TABLE 2). After the weights, input activations, and positions are fetched, the accumulator buffer (F) is indexed with the output positions computed from the sparse weight and sparse input activations.

[0072] In one embodiment, the accumulation unit 245 shown in FIG. 2C and the scatter accumulator array is doubled buffered so that products generated for one tile of weights are accumulated into one set of adders within the accumulator array 340 while registers in the accumulator array 340 storing partial products for the previous tile are accessed to resolve halos and encode the resulting output activations into the compressed format. Finally, when the computation for the output-channel group has been completed, the accumulator array 340 is drained and the compressed output activations are stored into the output activations buffer 350 and the output coordinates are stored into the indices buffer 355.

[0073] TABLE 3 shows pseudo-code for the PTIS-sparse dataflow. Referring to FIG. 2A, the layer sequencer 215 controls the memory interface 205 to read the weights once from off-chip DRAM in fully compressed form and transmit the weights to the PEs 210. Within each PE 210, the weights are ordered by tile (i.e., output channel group) (g), then by input channel (c), then by output channel within the tile (k). The per-PE computation using the tile/input channel/output channel ordering is shown in TABLE 3.

TABLE 3

pseudo-code for the PTIS-sparse dataflow	
For each tile g in G // Each tile is a set of contiguous channels $g=\{k_i, k_{i+1}, \dots, k_j\}$	
For input channel c in C //	
For each input activation a in c // read each input activation exactly once per tile	
For each weight w in g with input channel c //	
p = a * w ; // compute product and the associated	

TABLE 3-continued

pseudo-code for the PTIS-sparse dataflow	
	output position
	A[p.position] += p.value ; // sum into appropriate accumulator
For position j in Halo	
	Send A[j] to home PE and sum into home accumulator
	Apply non-linear function to A[*]
	Encode A[*] in sparse format and transfer to output activations buffer

Processing Element

[0074] Referring to FIG. 3A, as the weights are read by the memory interface 205 from DRAM, the weights are broadcast to the PEs 210 and held locally in a per-PE weight buffer 305. The input activations may be read by the memory interface 205 from DRAM or transmitted from the output activations buffer 350 and stored locally in a per-PE input activations buffer 310.

[0075] A state machine within the destination calculation unit 330 operates on the weight and input activations in the order defined by the PTIS-sparse dataflow to produce an output-channel group of $K_c \times W_c \times H_c$ partial sums inside the accumulator array 340. First, a vector F of compressed weights and a vector I of compressed input activations are fetched from the weight buffer 305 and the input activations buffer 310, respectively. The vectors are distributed into the FxI multiplier array 325 that computes a form of the cartesian product of the vectors.

[0076] While the vectors are processed by the FxI multiplier array 325 to compute products, the positions from the sparse-compressed weights and activations are processed by the destination calculation unit 330 to compute the output positions associated with the products. The FxI products are delivered to an array of A accumulator units within the accumulator array 340, that are addressed by the output positions. Each accumulator unit within the accumulator array 340 includes an addressable bank of storage, adder, and a register for storing partial sums associated with the output-channel group being processed. When processing of an output-channel group is complete, the partial sum stored in each register is the output activation value for one of the output positions. In one embodiment, the accumulator units are double-buffered so that one set of registers can store new partial sums while the second set of registers are drained out by the post-processing unit 345. When the output-channel group is complete, the post-processing unit 345 performs the following tasks: (1) exchange partial sums with neighboring PEs 210 for the halo regions at the boundary of the PE's 210 output activations, (2) apply the non-linear activation (e.g. ReLU), pooling, and dropout functions, and (3) compress the output activations into the compressed-sparse form and write the compressed-sparse output activations into the output activations buffer 350 and write the output positions associated with the compressed-sparse output activations into the indices buffer 355. In one embodiment, the post-processing unit 345 includes a compaction engine that is configured to encode the output activations and output positions into the compressed-sparse form.

[0077] In one embodiment, the weight buffer 305 is a first-in first-out FIFO buffer (WFIFO). The weight buffer 305 should have enough storage capacity to hold all of the non-zero weights for one input channel within one tile (i.e.,

for the inner most nested “For” in TABLE 3). When possible, the weights and input activations are held in the weight buffer 305 and input activations buffer 310, respectively, and are never swapped out to DRAM. If the output activation volume of a neural network layer can serve as the input activation volume for the next neural network layer, then the output activations buffer 350 is logically swapped with the input activations buffer 310 between processing of the different neural network layers. Similarly, the indices buffer 355 is logically swapped with the buffer 320 between processing the different neural network layers.

[0078] In one embodiment, when the weight buffer 305 within any PE 210 becomes full, broadcasting of the weight values into the weight buffer 305 is stalled. If the weight buffer 305 is big enough to hold a few input channels of a tile, some PEs 210 can move ahead to the next input channel while one or more other PEs 210 are a few channels behind—smoothing out load imbalance between PEs 210. In one embodiment, the weight buffer 305 has enough storage capacity to hold more than all of the weights in a tile (i.e., output-channel group) to smooth over some load imbalance between PEs 210.

[0079] The different logic blocks within the PE 210 may be pipelined as needed to achieve a target clock rate. However, the pipeline registers between pipeline stages need to freeze if the logic block receiving data output by the pipeline registers is stalled. Alternatively, elastic buffers can be used between the pipeline stages to simplify the distribution of a ready signal that indicates data can be accepted.

[0080] In one embodiment, the weight buffer 305 is a FIFO buffer that includes a tail pointer, a channel pointer, and a head pointer. The layer sequencer 215 controls the “input” side of the weight buffer 305, pushing weight vectors into the weight buffer 305. The tail pointer is not allowed to advance over the channel pointer. A full condition is signaled when the tail pointer will advance past the channel pointer when another write vector is stored. The buffer 315 may be implemented in the same manner as weight buffer 305 and is configured to store the positions associated with each weight vector. In one embodiment, the weight buffer 305 outputs a weight vector of F weights $\{w[0] \dots w[F-1]\}$ and the buffer 315 outputs the associated positions $\{x[0] \dots x[F-1]\}$. Each position specifies r, s, and k for a weight. The output channel k is encoded relative to the tile. For example, if the tile contains channels 40-47, channel 42 is encoded as k=2—an offset of 2 from 40, the base of the tile.

[0081] The destination calculation unit 330 controls the head and channel pointers (HeadPtr and ChannelPtr) of the weight buffer 305 and the buffer 315, to sequence the computation of a tile. The input activations buffer 310 and buffer 320 may be a set of registers or SRAM that are configured to store the input activations and the positions associated with each input activation value. The destination calculation unit 330 also controls a pointer (IA ptr) into the input activations buffer 310 and the buffer 320, to sequence the computation of a tile. The sequence that is implemented by the destination calculation unit 330 corresponds to the three inner loops of the pseudo-code shown in TABLE 3. Pseudo-code for operation of the destination calculation unit 330 is shown in TABLE 4. ScatterAdd is a function that transmits the products to the A accumulator units within the accumulator array 340.

TABLE 4

pseudo-code for sequencing computations for a tile
<pre> For c = 1 to C ChannelPtr = start of channel c in the weight buffer IAPtr = start of channel c in the input activations buffer IACnt = Number of non-zero IA entries in channel c // each entry is / wide Do HeadPtr = ChannelPtr // Rewind the weight buffer to start of channel c WCnt = Number of non-zero weight entries in c // each entry is F wide A[1:I] = IABuffer[IAPtr] // read next I activations, value and position for each Do W[1:F] = Wbuffer[HeadPtr] // read next F weights, value and position for each P[1:P] = OuterProduct(W [1: F], A[1: I]) // multiply all pairs of weights and IAs ScatterAdd(P[1:P]) HeadPtr = HeadPtr+1 WCnt=WCnt-1 While(WCnt > 0) IAPtr = IAPtr + 1 IACnt = IACnt-1 While(IACnt > 0) </pre>

While the pseudo-code shown in TABLE 4 is several lines long, each iteration of the inner loop takes a single cycle and the overhead of incrementing the counters and testing the loop bounds takes place in parallel. Therefore, the F*I multiplier array 335 performs F*I multiplies (of values and positions) each processing cycle unless the weight buffer 305 goes empty or the F*I arbitrated crossbar 335 signals that it cannot accept inputs. When the processing is not stalled, the destination calculation unit 330 increments the head pointers each processing cycle, outputting another vector of F weights (and associated positions) each processing cycle. The destination calculation unit 330 continues to increment the head pointer, each processing cycle that the processing is not stalled, until the next increment would pass the end of the current channel (i.e. pass the channel pointer). When the end of the current channel is reached, the destination calculation unit 330, advances the IAPtr and the head pointer is rolled back (i.e., rewound) to the start of the current channel. The IAPtr is then used to read the next vector of I input activations and the rolled back head pointer is used to read the first vector of F weights. The destination calculation unit 330 then sequences all of the weights for another vector of input activations to produce another vector of products. When the last vector of input activations for channel c is processed, the destination calculation unit 330 advances to channel c+1 by setting the channel pointer to point to the first weight vector of the channel c+1.

[0082] At the end of an input channel not all F weights or I activations may be valid. Invalid activations are indicated by a value of zero and will not result in a request to the ScatterAdd function. The end of an input channel c is identified by count. The weights and activations for each input channel are preceded by a count of the non-zero elements for the channel. At the start of the channel IACnt and WCnt are initialized to the number of I-wide or F-wide entries for the channel. IACnt and WCnt are decremented after each vector is consumed and checked for zero to determine the end of the channel. In one embodiment, to avoid losing a processing cycle reading IACnt and WCnt for

a channel, the counts are kept in a pair of separate small RAMs—one for weight counts and one for IA counts (not shown in FIG. 3A).

Position Conversion to an Accumulator Address

[0083] FIG. 3B illustrates two 3x3 weight kernels and positions, in accordance with one embodiment. A first set of weights for k=1 includes the non-zero elements a, b, and c and a second set of weights for k=2 includes the non-zero elements d, e, and f. The (r,s,k) format encodes positions for the non-zero weights as the following position vector: (2,0,1),(0,1,1),(1,2,1),(0,1,2),(2,1,2),(1,2,2).

[0084] As multiplication is performed on the “value” component of each (value, position) pair, the destination calculation unit 330 performs a vector addition on the positions—giving a (x, y, k) position (e.g., output coordinates) for the resulting product. Specifically, for each product, the x coordinates associated with the weight and input activation positions are summed and the y coordinates associated with the weight and input activation positions are summed to produce the (x, y, k) position for the resulting product. For example, summing the first position in the weight position vector with a set of four positions of non-zero input activations, (7,3), (12,3), (20,3), and (24,3) produces a product position vector (9,3,1),(14,3,1),(22,3,1), and (26,3,1).

[0085] The destination calculation unit 330 then linearizes the coordinates of the output position to produce an accumulator address that is output to the F*I arbitrated crossbar 335. TABLE 5 is pseudo-code for the operations performed in the F*I multiplier array 325 and destination calculation unit 330.

TABLE 5

pseudo-code for product and position computations
<pre> forall f = 0: F -1 // in parallel forall i = 0:-1 // in parallel t = f + i*F ; // index of product p[t].value = w[f].value * a[i].value ; // compute value p[t].x = w[f].x + a[i].x ; // x-coordinate of position p[t].y = w[f].y + a[i].y ; // y-coordinate of position p[t].k = w[f].k ; // k-coordinate of position endfor endfor </pre>

The “forall” in TABLE 5 implies that all P iterations of the inner loop are performed in parallel—in a single cycle. In one embodiment, after computing the output position of each product p[t] in (x,y,k) form, the output position is linearized into an accumulator address p[t].a according to the formula:

$$p[t].a = p[t].x + p[t].y * \text{max_x_oa} + p[t].k * \text{max_x_oa} * \text{max_y_oa} \quad \text{Equation (1)}$$

[0086] Note that max_x_oa is typically greater than max_x_ia by one less than R, the width of the convolution kernel, max_x_weight. Similarly, max_y_oa is typically greater than max_y_ia by one less than S, the height of the convolution kernel, max_y_weight. max_x_oa and max_y_oa refer to the dimensions of the halo. Continuing with the previous example, the output position vector (9,3,0),(14,3,0),(22,3,0), and (26,3,0) is converted into 105, 110, 118, and 122, assuming an output tile has a max_x_oa=32.

[0087] The F*I arbitrated crossbar 335 transmits the products to the associated accumulator in the accumulator array 340 based on the product positions. The low bits of the linearized accumulator address are used by the F*I arbitrated crossbar 335 to route each product to an accumulator unit in the accumulator array 340 and the product is added to a partial sum by the adder within the accumulator array 340 that is selected by the high bits of the address. The operation of the F*I arbitrated crossbar 335 is described in detail in conjunction with FIG. 3C.

[0088] When arbitration is used and two products are associated with the same output position (e.g., address), one of the two products is transmitted by the F*I arbitrated crossbar 335 and stored in an accumulator unit within the accumulator array 340 while the other product destined for the same accumulator unit is stalled by the F*I arbitrated crossbar 335. Each accumulator unit may be considered to be a bank of addressable storage combined with an adder so that products associated with the same address can be accumulated. In one embodiment, when one product is stalled, output registers in the F*I multiplier array 325 are stalled and computation of new products stalls. In one embodiment, a FIFO buffer at the output of each multiplier in the F*I multiplier array 325 is used to smooth load imbalance between accumulator units. Performance improvement may be provided when the number of banks A is larger than the number of products F*I. In one embodiment, $A=2 \cdot F \cdot I$, where $F \cdot I=16$ and $A=32$.

[0089] After all partial sums for one tile have been computed, the doubled-buffered accumulator array 340 is switched. The PE 210 may begin processing of the next tile using the “primary” of accumulator array 340 while the post-processing unit 345 begins post-processing of the last tile in parallel using the “secondary” accumulator array 340. The post-processing unit 345 performs the following steps: halo resolution, non-linear function evaluation, and encoding. The adders and registers in the “secondary” accumulator array 340 are also cleared to force the partial sums for a subsequent tile to values of zero as the encoding process is completed.

Scatter-Add

[0090] A scatter-add function is performed by a combination of the F*I arbitrated crossbar 335 and the accumulator array 340. The F*I arbitrated crossbar 335 receives $F \cdot I=P$ products and output positions from the $F \times I$ multiplier array 325. In one embodiment, the output positions are represented as linear addresses. The products are routed to adders within the accumulator array 340, where each product is routed to a particular adder selected by the linear address associated with the product. In one embodiment, the products are routed to the adders through a buffer (e.g., an accumulator unit). The product is then added to the value stored in the register paired with the adder to generate a partial sum. TABLE 6 is pseudo-code for the scatter-add function performed by the F*I arbitrated crossbar 335 and the accumulator array 340.

TABLE 6

pseudo-code for the scatter-add function
<pre> Forall (t = 0:(I*F-1)) do Acc[p[t].a] = Acc[p[t].a] + p[t].value ; endfor </pre>

[0091] To simplify addressing of the adders, the number of accumulator units should be a power of 2, $A=2^b$, where b

is an integer. The low order bits of address a select the accumulator unit containing $Acc[a]$, and the high order bits of address a specify an offset within the storage within the accumulator unit. The F*I arbitrated crossbar 335 includes a network to route values to the appropriate accumulator unit. A should be larger than F*I to reduce contention for the accumulator units and provide adequate processing throughput. For small values of A, the network can be a single stage of arbitrated multiplexers. For larger values of A, a multi-stage network may be used to reduce wiring complexity. In one embodiment a FIFO is provided on each input to the F*I arbitrated crossbar 335 to smooth out load imbalance between the accumulator units.

[0092] FIG. 3C illustrates a single-stage F*I arbitrated crossbar 335, in accordance with one embodiment. The single-stage F*I arbitrated crossbar 335 includes a FIFO 362, decoder 364, arbiter 365, multiplexer 366, and an OR-gate 370. An accumulator unit 368 within the accumulator array 340 is coupled to the output of the mux 366. Head-of-line blocking at the accumulator array 340 can be avoided by using multiple input FIFOs at the accumulator units 368 and having each FIFO hold (p,a) pairs for a subset of the accumulator units 368. Only one input and one accumulator unit 368 are shown in FIG. 3C. A complete F*I arbitrated crossbar 335 includes P FIFOs 362, P decoders 364, P OR-gates 370, A arbiters 365, and A multiplexers 366 coupled to A accumulator units 368.

[0093] Products $p[i]$ are pushed into the FIFO 362. In one embodiment, the FIFO 362 has a depth of 2 or 3. When any one of the FIFOs 362 fills, the F*I arbitrated crossbar 335 becomes not-ready and stalls the $F \times I$ multiplier array 325. The output of the FIFO 362 consists of a product $p[i]$ and an address $a[i]$. Product $p[i]$ from input i is connected to the ith input of the multiplexer 366 at the input to each accumulator unit 368. The low bits of address $a[i]$ are decoded by the decoder 364 to a one-hot request vector $r[i][j]$. Across all inputs, if $r[i][j]$ is true, it implies that input i is making a request for the jth accumulator unit 368. When the FIFO 362 is empty the decoder 364 is disabled, so that no requests are asserted. In one embodiment, selecting the low-bits of $a[i]$ is replaced with a hash to spread the addresses in the accumulator array 340 across the accumulator units 368 to reduce bank conflicts.

[0094] Each accumulator unit 368 in the accumulator array 340 functions as a bank of storage (e.g., a latch or register array) associated with an adder. The requests $r[i][j]$ from the decoder 364 to the accumulator unit 368 are input to the arbiter 365. The arbiter 365 generates a grant vector $gr[i][j]$ (selecting the winning i for the j accumulator unit 368). Across all the accumulator units 368, if bit $gr[i][j]$ of the $P \times I$ grant matrix is true, it implies that input i has been granted access to the j accumulator unit 368 for the next cycle. The grant signals are used both to control the multiplexer 366, to select the winning product and address from the multiplexer inputs and providing an indication back to the FIFO 362—so the winning product is dequeued from the FIFO 362 at the end of the processing cycle.

[0095] FIG. 3D illustrates the accumulator unit 368, in accordance with one embodiment. The accumulator unit 368 includes a flip-flop 382, storage array 380, and an adder 385. The address output of the multiplexer 366 is used to select one latch or register from the storage array 382 for output to the adder 385. The storage array 380 stores partial sums and is read using the address $a'[i]$. The product $p'[i]$ received by

the accumulator unit **368** is summed with the partial sum stored in the storage array **380** at the location associated with the address $a'[i]$. As shown in FIG. 3D, the address $a'[i]$ is pipeline by a flip-flop **382** and thereby delayed by one clock cycle to be used as a write address for storing the sum output by the adder **385**. In other embodiments, $a'[i]$ may be delayed by more than one clock cycle to write the sum generated by the adder **385** to accumulate the product into the partial sum.

[0096] Head-of-line blocking at the accumulator array **340** can be avoided by using multiple input FIFOs at the accumulator unit **368** and having each FIFO hold (p,a) pairs for a subset of the accumulator units **368**. In one embodiment, a separate FIFO is provided at each the input to each adder **385** for each one of the accumulator units **368** (i.e., virtual output queueing is employed at outputs of the accumulator units **368**). A drawback of the single-stage F*I arbitrated crossbar **335** shown in FIG. 3C is complex wiring because there is a direct path from every product input to every accumulator unit **368** resulting in PxA paths. For example, with P=16 and A=32, there are 612 paths, each carrying a product, an address, a request and returning a grant. The wiring complexity can be reduced by factoring the scatter-add function.

[0097] FIG. 3E illustrates a two-stage F*I arbitrated crossbar **380**, in accordance with one embodiment. Although the two-stage F*I arbitrated crossbar **335** is described for P=16 and A=32, other values of P and A may be used in two or more stages. A first stage is 4 instances of the single-stage F*I arbitrated crossbar **335** with P=4 and A=8. A second stage is 8 instances of the single-stage F*I arbitrated crossbar **335** with P=4 and A=4. Each of the stages requires 128 direct paths. The number of stages can be increased to reduce the number of direct paths. In one embodiment, FIFOs are included at the intermediate stages of a multi-stage arbitrated crossbar. However, if all arbitration can be completed in one processing cycle, FIFOs at the intermediate stages do not necessarily provide any benefit in terms of processing throughput.

[0098] The energy of accessing the accumulator array **340** may be reduced by combining products associated with the same output position. In one embodiment, to maximize the probability of combining, products are buffered at the accumulator units **368** in a combining buffer (e.g., a FIFO with 8 entries) and the products are only accumulated into the partial sum when the combining buffer becomes full. Addresses of arriving products are compared to entries in the combining buffer and when an address of an arriving product matches the address of a stored product, the arriving product is summed with the stored product. In one embodiment, the combining buffers have multiple write ports, allowing two or more arriving products to be inserted into the combining buffer simultaneously.

Post-Processing

[0099] The post processing unit **345** performs three functions: halo resolution, non-linear function evaluation, and sparse output activation encoding. In one embodiment, the accumulator array **340** is double-buffered. The three functions are performed on completed tile of output activations in the secondary accumulator array **340** while a current tile of output activations is being computed in the primary accumulator array **340**.

[0100] The operation count of the post-processing unit **345** is relatively low compared to the FxI multiplier array **325**. The FxI multiplier array **325** performs a 6-nested loop (over x,y,r,s,c,k) while the post-processing unit **345** only performs a 3-nested loop (over x,y,k). Thus, a post-processing unit **345** performing one operation per cycle should keep pace with a FxI multiplier array **325** that performs 16 operations per cycle. In one embodiment, the post-processing unit **345** is implemented using a microcontroller or a state machine. The pseudo-code for halo resolution is shown in TABLE 7.

TABLE 7

Pseudo-code for halo resolution	
For halo region q = 0:7	
x1, x2, y1, y2, xo, yo, pe = haloDescriptor[q] ; // 7-tuple descriptor	
for each halo region	
for x = x1:x2	
for y = y1:y2	
for k = 0: hl	
src = linearAddress(x,y,k) ; // compute	
accumulator address in this PE	
dst = linearAddress(x+xo,y+yo,k) ; // compute	
address in destination PE	
send(pe, dst, A[src]) ;	
endfor	
endfor	
endfor	
endfor	

[0101] The Pseudo-code shown in TABLE 7 iterates over the eight halo regions. Each region is described by a 7-tuple that is loaded from a region descriptor table. The 7-tuple includes the x-and y-ranges of the halo region in the source PE **210** (x1:x2, y1:y2). The 7-tuple includes the x- and y-offset (xo, yo) to translate a position in this PE **210** to a position in the destination PE **210**. (The offsets are signed values). Finally, the 7-tuple includes the neighbor number of the destination PE **210**. The linearAddress function converts (x,y,k) to a linear accumulator address as:

$$\text{linearAddress}(x,y,k) = x + y * \text{max_x_oa} + k * \text{max_x_oa} * \text{max_y_oa} \quad \text{Equation (2)}$$

[0102] Consider an example where R×S=3×3 convolutions are performed on input activations with dimensions 50×50×c, and the result is a set of output activations with dimensions 52×52×|hl|. The halo consists of eight regions—four edges and four corners. The eight region descriptors for this case are shown in TABLE 8.

TABLE 8

Halo region descriptors for R = S = 3 and W = H = 50							
Region	x1	x2	y1	y2	xo	yo	PE
Left	0	0	1	51	51	0	(-1, 0)
Top	1	51	0	0	0	51	(0, -1)
Right	52	52	1	51	-51	0	(1, 0)
Bottom	1	51	52	52	0	-51	(0, 1)
Upper-Left	0	0	0	0	51	51	(-1, -1)
Upper-Right	52	52	0	0	-51	51	(1, -1)
Lower-Right	52	52	52	52	-51	-51	(1, 1)
Lower-Left	0	0	52	52	51	-51	(-1, 1)

In the example, the Left region specifies a source range of (0,1:51), an offset of (51,0), and a PE **210** having coordinates of (-1,0) relative to the current PE **210**. The offset makes the destination range (51,1:51). The post-processing unit **345**

uses the descriptor to read the accumulator array **340** within the PE **210**, walking the left edge and sending value, position pairs to a neighboring PE **210** to the left (-1, 0). The neighboring PE **210** handles the value, position pairs in the same manner as the value, position pairs coming from the FxI multiplier array **325** except that the value, position pairs are input to the secondary accumulator units **368**. Additional input ports are provided by the FxI arbitrated crossbar **335** to route the value, position pairs from each of the neighboring PEs **210** to the secondary accumulator units **368**. The PEs **210** on the edges and corners of the PE array within the SCNN accelerator **200** will be missing 3 (edge) or 5 (corner) neighbors. The descriptors for the missing neighbors are marked invalid causing the post-processing unit **345** to skip halo resolution for non-existent neighbors.

[0103] After halo resolution for a PE **210** and all of its immediate neighbors is complete, the post-processing unit **345** scans the accumulator array **340** and performs a non-linear function for each output activation in the tile. The pseudo-code for the non-linear function is shown in TABLE 9.

TABLE 9

The non-linear function	
<pre> for x = x1: x2 // non-halo region of x for y = y1: y2 // non-halo region of y for k = 0: h-1 // entire tile in k addr = linearAddress(x,y,k) ; A[addr] = function(A[addr]) ; endfor endfor endfor </pre>	

[0104] The pseudo-code shown in TABLE 9 iterates over the non-halo region of the accumulator array **340**. The non-halo region includes all of the accumulator units in the accumulator array **340** that were not part of an edge or corner region. For the previous example, the non-halo region is (1:51,1:51). The most common non-linear function to be applied is the rectified non-linear function (ReLU) that converts negative values to zero, but other functions (such as sigmoid) may also be used. Some functions may be approximated as piecewise linear functions. In one embodiment, positive values below a predetermined threshold are forced to zero and negative values above a predetermined threshold are forced to zero.

[0105] After the non-linear function is applied to the tile in the secondary registers within the accumulator array **340**, tile is encoded to compress out non-zero elements. The pseudo-code for the compression operation is shown in TABLE 10.

TABLE 10

Pseudo-code for the compression operation	
<pre> OAptr = startOAptr ; for k = 0: h-1 // for each channel in the tile lastNZPos = (0,0) ; // initialize last non-zero position to start of channel nzCount = 0 ; // count non-zeros in this channel for x = x1: x2 // non-halo region of x for y = y1: y2 // non-halo region of y addr = linearAddress(x,y,k) ; if (A[addr] != 0) then // process non-zero pos = encode(lastNZPos, (x,y)) ; </pre>	

TABLE 10-continued

Pseudo-code for the compression operation	
<pre> OARAM [OAptr] = (A[addr], pos) ; OAptr = OAptr + 1 ; lastNZPos = (x,y) ; endif endfor endfor emitCount(nzCount) ; endfor </pre>	

[0106] The pseudo-code shown in TABLE 10 walks the accumulator array **340** one channel of output activations at a time and writes a (value, position) entry to the output activations buffer **350** and indices buffer **355** for every non-zero output activation value. The function “encode” encodes the position relative to the last position using one of the methods described below. Note that “encode” may need to emit one or more “dummy” values (a non-zero having a value of zero) if the difference between the current position (x,y) and “lastNZPos” cannot be directly encoded. After each channel is processed, the number of non-zeros in that channel (nzCount) is stored in a separate table. When encoding outputs, the OAptr addresses single (value, position) entries in the output activations buffer **350** and indices buffer **355**. After all tiles in a layer of the neural network have been completed, the output activations buffer **350** and indices buffer **355** and input activations buffer **310** and buffer **320**, respectively, switch functions and the next layer of the neural network is processed. When the output activations buffer **350** and indices buffer **355** is switched, the IAptr reads four vectors of (value, position) at a time.

[0107] To increase parallelism beyond a single PE **210**, multiple PEs **210** can be operated in parallel with each working on a disjoint three-dimensional tile of input activations. Because of the end-to-end compression of activations, both the input and output activations of each tile may be stored local to the PE **210** that processes the tile, further reducing energy-hungry data transmission. Overall, the SCNN accelerator **200** provides efficient compressed storage and delivery of input operands to the FxI multiplier array **325**, high reuse of the input operands in the FxI multiplier array **325**, and that spends no processing cycles on multiplications with zero operands.

Compressed-Sparse Weights and Activations

[0108] In addition to reducing energy consumption for weights and input activations of zero, the SCNN **200** architecture exploits sparse weights and activations. A dense encoding of sparse weights and activations is used to reduce the bandwidth needed to transmit the weight and activation values from the memory to the SCNN **200**, between different levels of the memory hierarchy, and between the different logic circuits within the SCNN **200**. Input data, such as weights and activations with zeros can be represented in a compact form referred to as compressed-sparse format. The amount by which the input data may be compressed increases as the number of zeros increases. However, even when only 10% of the multi-bit elements equal zero, it may be worthwhile to encode the input data in the compressed-sparse format. Encoding the sparse weights and/or activations reduces the data footprint, which allows larger matrices to be stored within a given size storage structure, such as the

input activation buffer **235** and the weight buffer **230**. In one embodiment the weight buffer **230** and the input activation buffer **235** each carry a 10-bit overhead for each 16-bit value to encode multi-dimensional positions of non-zero elements in the compressed-sparse format.

[0109] FIG. 4A illustrates a flowchart of a method **400** for processing compressed-sparse data in the SCNN **200**, in accordance with one embodiment. Although method **400** is described in the context of a processing element within the SCNN **200**, the method **400** may also be performed by a program, custom circuitry, or by a combination of custom circuitry and a program. Furthermore, persons of ordinary skill in the art will understand that any system that performs method **400** is within the scope and spirit of embodiments of the present invention.

[0110] At step **405**, compressed-sparse data is received for input to the PE **210**, where the compressed-sparse data encodes non-zero elements and corresponding multi-dimensional positions. In one embodiment, the compressed-sparse data represents weight values. In another embodiment, the compressed-sparse data represents input activation values.

[0111] At step **410**, the non-zero elements are processed in parallel by the PE **210** to produce a plurality of result values. In one embodiment, the non-zero elements are multiplied within the FxI multiplier array **325** to produce result values that are products. At step **415**, the corresponding multi-dimensional positions are processed in parallel to produce destination addresses for each result value in the plurality of result values. In one embodiment, the multi-dimensional positions are processed in the destination calculation unit **330** to produce a destination accumulator address associated with a location in the accumulator array **340** for each one of the result values. Specifically, the destination accumulator address may indicate a location in the storage array **380** (i.e., bank) within an accumulator unit **368**. At step **420**, each result value is transmitted to an accumulator unit **368** that is associated with the destination address for the result value. In one embodiment, each result value is a product that is transmitted through the F*I arbitrated crossbar **335** to one of the accumulator units **368** based on the corresponding destination address.

[0112] In one embodiment, the SCNN **200** uses a simple compressed-sparse encoding approach based on a run-length encoding scheme. A data vector may be extracted from the compressed-sparse encoded data, where the data vector is a sequence non-zero values. An index vector may be extracted from the compressed-sparse encoded data, where the index vector is a sequence of zero-counts (the number of zeros between each non-zero element). For example, a compressed-space encoding of the data shown in FIG. 3B is (a, b, c, d, e, f) and (2, 0, 3, 4, 1, 1) representing a data vector and a corresponding index vector, where each element in the index vector is a number of zeros preceding the corresponding non-zero element.

[0113] Determining the coordinates of a location in the accumulator array **340** for each product output by a multiplier in the FxI multiplier array **325** requires reading the index vectors for F and I and combining the index vectors with the coordinates of a portion of the output activation space currently being processed. Four bits per index allows for up to 15 zeros to appear between any two non-zero elements. When more than 15 zeros occur between two non-zero elements, a zero-value placeholder (i.e., zero pad) is inserted as an intervening non-zero element without

incurring any noticeable degradation in compression efficiency. With an expected non-zero element density of 30% there will be approximately 2 zeros between non-zero elements on average.

[0114] While the SCNN **200** will operate most efficiently when the activations fit in the input activation buffers **235**, large networks may require activations to be saved to and restored from DRAM through the memory interface **205**. Therefore, the SCNN **200** may employ a tiling approach that operates on a 2D subset of the activation space at a time. The DRAM accesses needed to read one tile of input activations can be hidden by pipelining the read operations in tandem with the computation of the previous tile of output activations. Similarly, reading of the weights from DRAM can be performed on a tile granularity.

[0115] In one embodiment, the weights are encoded in a compressed-sparse format of tiles that include at most K output channels, and the tiles are ordered by input channel. The goal is to maximize re-use of input activations under the constraint of a fixed number of accumulators (and hence a limit on the number of output channels). The compressed-sparse format allows reading of W weights and corresponding positions (r,s,k) in parallel for an input channel c. Therefore, a format where weights and positions have fixed locations is desired—so incremental decoding is not needed. Weight values are a four-dimensional matrix, where x, y, c, and k are the four dimensions. A tile is a slice of a weight data vector $k \in \{k_1, k_2, \dots, k_K\}$ —i.e., arbitrary r,s values but with k restricted to a set of K values. A tile may be encoded in a compressed-sparse format including K (the number of output channels), k_1, k_2, \dots, k_K (the actual numbers of the K output channels), and C (the number of input channels in the tile). For each input channel, the compressed-sparse format includes a delta coded index c for the input channel (i.e., a difference from the previous input channel) and a count of the non-zero weights in the input channel. For each output channel k, the compressed-sparse format includes three parameters for each non-zero weight in kernel ck. A first parameter is the number of zeros between the previous non-zero weight and the current weight. Note that the zeros at the end of one kernel and the beginning of the next kernel are coded together. A second parameter is a coded weight value W_{xyck} represented as either a binary weight or an index into a codebook.

[0116] FIG. 4B illustrates a tile **340** of weight values for two output channels, in accordance with one embodiment. In one embodiment, 3x3 convolutions may be performed using the tile **340** of weight values over two input channels to generate results for two output channels. The tile **340** of weight values is sparse and may be represented in a compressed-sparse format.

[0117] In one embodiment, the tile **340** of weight values is encoded as {2, 1, 2, 4, 0, 6, 1, 3, 4, 4, 1, 5, 0, 6, 3, 7, 3, 8, 0, ...}. The first four symbols indicate the “shape” of the tile K=2 with $k_1=1$ and $k_2=2$, and C=4. The first 0 indicates the first input channel at offset of 0 from the starting position, c=0. The following 6 indicates there are six non-zero weights in the first input channel. The next six symbols are zero-count, weight pairs that encode the c=0, k=1 kernel. The 1 implies there is 1 zero before the 3, and the first 4 indicates 4 zeros between the 3 and the 4. Because the 5 is in the last position for c=0, k=1 we know that the zero after the 5 starts encoding the next channel. The next six symbols encode the c=0, k=2 kernel. The final 0 indicates that there

are no empty channels before the next input channel, so the next symbols will encode channel $c=1$. The sequence of zero-counts indicates the number of zeros before the first non-zero weight value and between adjacent pairs of non-zero weight values.

[0118] After stripping off the tile header (2, 1, 2, 4) and the channel header (0, 6) the next 12 symbols can be read in parallel as the data vector and index vector, giving six weights along with the corresponding r,s,k positions. A running sum is needed to compute a linear index for each weight, and the linear indices are then converted to r,s,k position coordinates. To make it easier to decode the linear index into r,s coordinates, r_{max} may be rounded up to the next power of 2. For example a 3×3 kernel becomes a 3×4 ($s_{max} \times r_{max}$) kernel, with the last column of weights set to zeros. In one embodiment, when a running sum is used to compute the linear index, the low two bits are r and the remaining bits are s .

[0119] Each r,s,k position for a weight or (x,y) position for an input activation may be calculated using the position coordinates of the previous weight or input activation, respectively. The weight position calculation is shown in TABLE 11, where “value” is the zero-count.

TABLE 11

Pseudo-code for position calculations
<pre> if (type == 'R') then position.r = last_position.r + value + 1 ; position.s = last_position.s ; position.k = last_position.k ; // weights only wrap() ; // optional else if (type == 'S') then position.r = value ; position.s = last_position.s + 1 ; position.k = last_position.k ; // weights only wrap() ; // optional else if (type == 'K') then position.r = value ; position.s = 0 ; position.k = last_position.k + 1 ; // weights only endif </pre>

An r coordinate type (R) encodes the zero-count, i.e., the number of zeros between the last non-zero element and the current non-zero element. When a running sum in each dimension (e.g., position. r and position. s) exceeds the maximum r dimension value (r_{max}), the position can be optionally wrapped—incrementing y and reducing r by the r_{max} . The y coordinate type (S) increments the s coordinate of the position by one and sets the r position to the value. The k coordinate type (K) increments the k coordinate of the position, resets s to zero, and sets r to the value. The wrap procedure is shown in TABLE 12, where max_r is r_{max} and max_s is s_{max} .

TABLE 12

Pseudo-code for the position calculations with wrapping
<pre> if (position.r > max_r) then // max_r is max_r_weight or max_x_ia position.s = position.s + Truncate(position.r/max_r) ; position.r = position.r % max_r ; endif </pre>

TABLE 12-continued

Pseudo-code for the position calculations with wrapping
<pre> if (position.s > max_s) then // only applies to weights, max_s is max_s_weight position.k = position.k + Truncate(position.s/max_s) ; position.s = position.s % max_s ; endif </pre>

[0120] Wrapping can potentially lead to denser encodings—by providing more encoding options. However, supporting wrapping requires more complex decoding circuitry to perform divides and modulo operations. An intermediate option is to perform wrapping but to restrict r_{max} and s_{max} to be powers of two—simplifying the divide and modulo to shift and mask operations, respectively. Alternatively the wrap can be omitted and the appropriate coordinate type is required to advance the s or k coordinate. The (r,s,k) coordinates may be replaced with (x,y) coordinates, omitting k , to perform the position calculations for input activations.

[0121] In one embodiment, the weights may be represented as direct 16-bit or 8-bit values paired with a variable bit-width “code” value that is used to index a “codebook” to read the associated zero-count. Different codebooks may be used by different tiles. The encoding of the coordinate type and zero-count value should be done in a manner that maximizes encoding efficiency by providing more encodings for more common coordinate types and zero-count values.

[0122] FIG. 4C illustrates a coding scheme for weights and input activations (IA), in accordance with one embodiment. A 4-bit code indicates the coordinate type and zero-count value. Other coding schemes are possible and a coding scheme may use more or less than four bits. The weight codes have R, S, and K coordinate types while the activation codes have only X and Y coordinate types. For the weights, a larger number of codes (10) are devoted to the R coordinate type because it is most commonly used. The increment between values need not be one. For example, the zero-count values of R9 and X9 are not included to allow more “reach” between non-zero elements. Nine zeros between two non-zero elements can be encoded as an R4 (or X4) appended to R4 (or X4) with a zero weight value padded between the two runs of zeros. For the activation coding, a more aggressive coding of long runs of zeros is allowed with large increment gaps between codes.

[0123] As groups of F weights and I input activations are read from the weight buffer 305 and input activations buffer 310 each cycle, the position portion of the weights and input activations read from the buffers 315 and 320, respectively, are decoded from the 4-bit values shown in the table shown in FIG. 4C to full positions (x, y) for activations and (r,s,k) for weights. As previously explained, the $F \times I$ multiplier array 325 takes F weights and I input activations and produces $P = F \times I$ products. Each product is associated with a position that is computed by the destination calculation unit 330. For all product-position pairs, the non-zero weight and input activation values in compressed-sparse format are multiplied without expansion. The position portion of the compressed-sparse format includes zero-counts that are decoded into (r,s,k) for each weight and (x,y) for each input activation and then added to produce an (x,y,k) position for the corresponding product. The product-position computation is shown earlier in TABLE 5.

[0124] FIG. 4D illustrates weight values for four 3×3 convolution kernels 435, in accordance with one embodiment. FIG. 4E illustrates an encoding 440 of the positions for the weight values in the four 3×3 convolution kernels 435, in accordance with one embodiment. The first row of the encoding 440 comprises a stream of 12 codes, one for each non-zero weight value in the four 3×3 convolution kernels 435. With wrapping and $r_{max}=s_{max}=3$, the positions are encoded in the first row of the encoding 440. The first S1 corresponds to the zero in the upper left position that is followed by a 3. Following the first S1, S4 corresponds to the zero in the first row following the 3 and the three zeros in the second row of the first convolution kernel. Following the S4, a second S1 corresponds to the one zero in the third row of the first convolution kernel between the 4 and 5. Following the second S1 are two S0s corresponding to the lack of zeros between 5 and 6 and between 6 and 7 in the first row of the second convolution kernel. Following the two S0s is an S5 corresponding to the five zeros before the 8 in the third row of the second convolution kernel. The remaining codes can be derived in a similar fashion.

[0125] The second row of the encoding 440 shown in FIG. 4E illustrates the positions for the non-zero weight values in the four 3×3 convolution kernels 435. The positions may be determined based on the codes in the first row. Starting from an initial position of (0,0,0), the first S1 is decoded into the position (r,s,k)=(1,0,0) shown in the second row that corresponds to the position of the weight value of 3 in the first row of the first convolution kernel. The first S4 is decoded into the position (r,s,k)=(0,2,0) shown in the second row that corresponds to the position of the weight value of 4 in the third row of the first convolution kernel. The remaining positions can be derived in a similar fashion.

[0126] In one embodiment, linear indices are derived for the codes in the top row of the encoding 440 by computing a running sum starting at -1 and adding one for each weight value along with the zero-count value. Extracting the zero-counts from the top row produces {1,4,1,0,0,5,2,1,1,1,4,1}. If r_{max} is set to 4 instead of 3 (for a convolution kernel that is 4×3 instead of 3×3), the zero-counts become {1,6,1,0,0,7,3,2,1,2,6,1}. The zero-counts are then converted to a running sum, starting at -1 and adding one to each position for each of the corresponding weights. The running sum that is a linear index L_i of the zero-counts C_i is {1,8,10,11,12,20,24,27,29,32,39,41}, where $L_i=L_{i-1}+C_i+1$ and $L_0=-1$. The linear index is then converted to position coordinates (r,s,k).

[0127] When r_{max} is set to 4 (or another power of 2 value), r can be extracted by just removing the low two bits. A division by 3 is required to separate k and s from the remaining bits. The division may be avoided by rounding the kernel dimensions to 4×4 (or another power of 2 value in each dimension), the compressed-sparse coding may not be as dense due to the extra zeros. Note that the k coordinate in the (r,s,k) position is not the absolute address of the output channel, but rather the temporary address of the accumulator currently holding the output channel. The positions extracted from the linear index are shown in the second row of the encoding 440.

[0128] FIG. 4F shows a block diagram 450 for determining the (r,s,k) weight coordinates, in accordance with one embodiment. r_{max} is assumed to be small (smaller than the maximum zero-count) and is therefore rounded up to a power of two, so that the r and s coordinates can be handled

as a single field rs with the low bits being r and the high bits being s. An adder 425 sums a zero-count, z_i and 1 with rs_{i-1} to produce a tentative rs_i . The division operation to separate k and s does not actually require a divide, but may instead be performed using a running divide technique. At each step, computing the running sum, the tentative rs_i may be compared to $rs_{max}=r_{max}*s_{max}$. If the sum is greater than or equal to rs_{max} , rs_{max} is subtracted from the tentative rs_i and k is incremented. The running divide technique may be used to separate r and s when r_{max} is not rounded up to the next power of two.

[0129] A max subtract 455 subtracts rs_{max} from the tentative rs_i value output by the adder 425 and determines if the result is positive as indicated by the pos signal output by the max subtract 455. If the result is positive, the result of the subtraction is retained and selected for output as rs_i by a multiplexer 460. If the result is not positive, the multiplexer 460 selects the tentative rs_i for output as rs_i . An incrementor 455 receives k_{i-1} and increments k_{i-1} to update the output k_i when the result is positive. Note that if rs_{max} is smaller than the maximum zero-count, it may be necessary to compare against $2*r_{max}$ and other multiples. However, when rs_{max} is that small, in one embodiment, rs_{max} is rounded up to the next power of 2 and a running sum should be computed on a combined krs field.

[0130] In one embodiment, the encoding for input activations is the same as for weights except that the (r,s) coordinates are replaced with (x,y) coordinates and the k coordinate is omitted. However the size of an input activation slice may be considerably larger. For a high definition (HD) image size of 1920×1080 pixels distributed into an 8×8 array of PEs 210, each PE 210 will hold a 240×135 slice. At the other extreme, a deep convolutional layer may be only 14×14 having an x_{max} of just 1 or 2. When large sizes are too large to be rounded up to powers of 2, input activations may employ the running divide technique to separate x, y, and k.

[0131] FIG. 4G shows a block diagram 470 for determining the (x,y) input activation coordinates, in accordance with one embodiment. The calculation used for the input activation coordinates is similar to calculating the weight coordinates except that: (1) there is no k field and the positions are all from the same input channel c; and (2) the x coordinate is compared to x_{max} at each step and, if needed, x_{max} is subtracted. For input activations, x_{max} can become large, making it costly to round up to the next power of two.

[0132] An adder 475 sums a zero-count, t_i and 1 with x_{i-1} to produce a tentative x_i . A max subtract 485 subtracts x_{max} from the tentative x_i value output by the adder 475 and determines if the result is positive as indicated by the pos signal output by the max subtract 485. If the result is positive, the result of the subtraction is retained and selected for output as x_i by a multiplexer 480. If the result is not positive, the multiplexer 480 selects the tentative x_i for output as x_i . An incrementor 490 receives y_{i-1} and increments y_{i-1} to update the output y_i when the result is positive.

[0133] Note that the input activation coordinate system is tied to the halo such that, for a 3×3 convolution kernel, the current input activations start at (1,1). Once the (r,s,k) positions of the weights are computed and the (x,y) positions of the input activations are computed by the destination calculation unit 330, the r and x coordinates are summed and the s and y coordinates are summed by the destination calculation unit 330 to compute the output activation positions in (x,y,k) form. The destination calculation unit 330

then converts the output activation positions to a linear accumulator address according to:

$$\text{address}_i = x + y * x_{\text{max_halo}} + k * x_{\text{max_halo}} * y_{\text{max_halo}}$$

[0134] Note that $x_{\text{max_halo}}$ and $y_{\text{max_halo}}$ refer to the dimensions of the halo and (x, y, k) is the output activation position. The values being multiplied by y and k can be rounded up, if needed, to reduce the cost of the multiply. However the rounding may increase the cost of the accumulators in terms of extra operations that are not necessary.

[0135] FIG. 5A illustrates a non-linear coding scheme 505 for input activation zero-count values, in accordance with one embodiment. One technique for reducing the cost of rounding up x_{max} (or xy_{max}) to the next power of two is to space the zero-count symbols non-linearly. The first row of the coding scheme 505 is the four-bit code and the second row is the corresponding zero-count value. The first eight codes encode linear zero-counts from 0 to 7 as was described in conjunction with FIG. 4C. The next eight codes, however encode larger, non-linear zero-count values (e.g., 12, 18, 16, 24, 32, 48, 64, 96, and 128) to “jump around” the large empty region produced by rounding up x_{max} . For example, if x_{max} is 129 and is rounded up to 256, it may be necessary to jump by 128.

[0136] FIG. 5B illustrates another coding scheme 510 for input activation zero-count values, in accordance with one embodiment. The coding scheme 510 allows the zero-count to specify that the x-coordinate should be set to the specified value and that the y-coordinate should be incremented. As in the coding scheme 505, the first eight codes specify zero-counts of 0 to 7. The next eight codes, of the form Yn , instruct the destination calculation unit 330 to increment the y coordinate and set the x coordinate to $x=n$. With this form of encoding, there is no need to first convert to a linear index. The zero-count codes can be directly converted to (x, y) .

[0137] FIG. 5C illustrates another coding scheme 515 for input activation zero-count values, in accordance with one embodiment. Because most cases do not increment Y, it makes sense to have more “normal” codes than “Y-increment codes”. Therefore, the coding scheme 510 includes 11 codes that allow the zero-count to specify that the x-coordinate should be set to the specified value and includes 5 codes for incrementing the y-coordinate.

[0138] FIG. 5D illustrates another coding scheme 520 for weight zero-count values, in accordance with one embodiment. For weights the non-zero values are encoded in a three-dimensional r, s, k space, so r_{max} may be rounded up to the next power of two and jumps to the next channel k are encoded with distinguished zero-count values. The coding scheme 520 allows the zero-count to specify that the r-coordinate should be set to the specified value and that the k-coordinate should be incremented. The first fourteen codes specify zero-counts of 0 to 13. The last two codes, of the form Kn , instruct the destination calculation unit 330 to zero r and skip to the next output channel k .

[0139] FIG. 5E illustrates another coding scheme 525 for weight zero-count values, in accordance with one embodiment. The coding scheme 525 allows the zero-count to specify that the r-coordinate should be set to the specified value and that either the s-coordinate or the k-coordinate should be incremented. The first ten codes specify zero-counts of 0 to 9. The next three codes, of the form Sn , instruct the destination calculation unit 330 to zero r and

increment s . The last two codes, of the form Kn , instruct the destination calculation unit 330 to zero r and s and increment k .

[0140] As shown in FIG. 5A-5E, the compressed-sparse format may encode the non-zero weights and input activations as dense vectors of values along with a skip-encoded (i.e., nonlinear) vector of codes representing the position, where z is the value. In one embodiment, one or more codes may specify one of (i) add $z+1$ to the last coordinate, wrapping in r, s , and/or k as appropriate ($r=r+z+1$, wrap), (ii) skip to the next row ($s=s+1$, $x=z$, wrap), or (iii) skip to the next channel ($k=k+1$, $s=0$, $r=z$, wrap).

[0141] Previous efforts to exploit sparsity in CNN accelerators have focused on reducing energy or saving time, which will invariably also save energy. Eliminating the multiplication when an input operand is zero by gating an operand input to a multiplier is a natural way to save energy. Gating an operand will save energy, but the number of processing cycles will not be reduced. The SCNN accelerator 200 also saves energy by eliminating all the unnecessary multiplications, and when any input operand is zero the circuitry is not even prepared to perform a multiplication operation, thus saving time as well.

[0142] An additional approach to reducing energy is to reduce data transfer costs when the data is sparse. The input activations may be compressed for transfer to and from DRAM to save energy (and time) by reducing the number of DRAM accesses. However, conventional systems expand the compressed input activations before the input activations are loaded into an on-chip buffer, so the input activations are stored in expanded form. Thus, there is no savings on transfers from one internal buffer to another internal buffer or to the multipliers. In contrast, the SCNN accelerator 200 uses a compressed representation for all data coming from DRAM, and maintains the compressed representation in the on-die buffers. In contrast, SCNN accelerator 200 keeps both weights and activations in a compressed form in both DRAM and internal buffers. This saves data transfer time and energy on all data transfers and allows the chip to hold larger models for a given amount of internal storage.

[0143] The SCNN accelerator 200 exploits sparsity in both weights and activations using the sparse planar-tiled input-stationary (PTIS-sparse) dataflow. The PTIS-sparse dataflow enables the SCNN accelerator 200 to use a novel Cartesian product-based computation architecture that maximizes reuse of weights and input activations within a set of distributed PEs 210. In addition, the PTIS-sparse dataflow allows the use of a dense compressed-sparse encoding for both weights and activations to be used through almost the entire processing flow. The amount of data that is transferred within the SCNN accelerator 200 is reduced and the amount of on-die storage capacity is effectively increased. Results show that for an equivalent area, the SCNN accelerator 200 architecture achieves higher energy efficiency compared with an energy-optimized dense architecture when the weights and activations are each less than 85% dense. On three contemporary networks the SCNN accelerator 200 architecture achieves performance improvements over the dense architecture by a factor of 2.6 times while still being energy-efficient by a factor of 2.5 times.

Exemplary System

[0144] FIG. 6 illustrates an exemplary system 600 in which the various architecture and/or functionality of the

various previous embodiments may be implemented. As shown, a system 600 is provided including at least one SCNN Accelerator 200 that is connected to a communication bus 602. The communication bus 602 may be implemented using any suitable protocol, such as PCI (Peripheral Component Interconnect), PCI-Express, AGP (Accelerated Graphics Port), HyperTransport, or any other bus or point-to-point communication protocol(s). The system 600 also includes a main memory 604. Control logic (software) and data are stored in the main memory 604 which may take the form of random access memory (RAM).

[0145] The system 600 also includes a central processor 601 (e.g., CPU), input devices 612, a graphics processor 606, and a display 608, i.e. a conventional CRT (cathode ray tube), LCD (liquid crystal display), LED (light emitting diode), plasma display or the like. User input may be received from the input devices 612, e.g., keyboard, mouse, touchpad, microphone, and the like. In one embodiment, the graphics processor 606 may include a plurality of shader modules, a rasterization module, etc. Each of the foregoing modules may even be situated on a single semiconductor platform to form a graphics processing unit (GPU).

[0146] In the present description, a single semiconductor platform may refer to a sole unitary semiconductor-based integrated circuit or chip. It should be noted that the term single semiconductor platform may also refer to multi-chip modules with increased connectivity which simulate on-chip operation, and make substantial improvements over utilizing a conventional CPU and bus implementation. Of course, the various modules may also be situated separately or in various combinations of semiconductor platforms per the desires of the user.

[0147] The system 600 may also include a secondary storage 610. The secondary storage 610 includes, for example, a hard disk drive and/or a removable storage drive, representing a floppy disk drive, a magnetic tape drive, a compact disk drive, digital versatile disk (DVD) drive, recording device, universal serial bus (USB) flash memory. The removable storage drive reads from and/or writes to a removable storage unit in a well-known manner.

[0148] Computer programs, or computer control logic algorithms, input data for the SCNN Accelerator 200, output data generated by the SCNN Accelerator 200, and the like may be stored in the main memory 604 and/or the secondary storage 610. Such computer programs, when executed, enable the system 600 to perform various functions. The memory 604, the storage 610, and/or any other storage are possible examples of computer-readable media.

[0149] In one embodiment, the architecture and/or functionality of the various previous figures may be implemented in the context of the SCNN Accelerator 200, central processor 601, the graphics processor 606, an integrated circuit (not shown) that is capable of at least a portion of the capabilities of one or more of the SCNN Accelerator 200, the central processor 601, and the graphics processor 606, a chipset (i.e., a group of integrated circuits designed to work and sold as a unit for performing related functions, etc.), and/or any other integrated circuit for that matter.

[0150] Still yet, the architecture and/or functionality of the various previous figures may be implemented in the context of a general computer system, a circuit board system, a game console system dedicated for entertainment purposes, an application-specific system, and/or any other desired system. For example, the system 600 may take the form of a

desktop computer, laptop computer, server, workstation, game consoles, embedded system, and/or any other type of logic. Still yet, the system 600 may take the form of various other devices including, but not limited to a personal digital assistant (PDA) device, a mobile phone device, a television, etc.

[0151] Further, while not shown, the system 600 may be coupled to a network (e.g., a telecommunications network, local area network (LAN), wireless network, wide area network (WAN) such as the Internet, peer-to-peer network, cable network, or the like) for communication purposes.

[0152] While various embodiments have been described above, it should be understood that they have been presented by way of example only, and not limitation. Thus, the breadth and scope of a preferred embodiment should not be limited by any of the above-described exemplary embodiments, but should be defined only in accordance with the following claims and their equivalents.

What is claimed is:

1. A method, comprising:

receiving compressed-sparse data for input to a processing element, wherein the compressed-sparse data encodes non-zero elements and corresponding multi-dimensional positions;

processing, via the processing element, the non-zero elements in parallel to produce a plurality of result values;

processing, via the processing element, the corresponding multi-dimensional positions in parallel to produce destination addresses for each result value in the plurality of result values; and

transmitting each result value to a destination accumulator associated with the destination address for the result value.

2. The method of claim 1, wherein the non-zero elements are weight values.

3. The method of claim 1, wherein the non-zero elements are input activation values.

4. The method of claim 1, wherein the multi-dimensional positions are encoded as a sequence of zero-counts and each zero-count is a number of zeros between pairs of the non-zero elements.

5. The method of claim 4, wherein the processing the corresponding multi-dimensional positions comprises:

converting the zero-counts into coordinates of the multi-dimensional positions;

for each result value, summing the coordinates in each dimension of the multi-dimensional positions with additional coordinates associated with additional non-zero elements to produce the corresponding destination addresses.

6. The method of claim 5, further comprising computing a running sum for each zero-count in the sequence of zero-counts to produce the coordinates of the multi-dimensional positions.

7. The method of claim 1, wherein the multi-dimensional positions are encoded a sequence of codes and each code defines a zero-count that is a number of zeros between pairs of the non-zero elements.

8. The method of claim 7, wherein a first code of the codes defines a number of rows to skip in a first coordinate dimension associated with the multi-dimensional positions.

9. The method of claim 7, wherein the codes define a linear progression of the zero-counts.

10. The method of claim 7, wherein the codes define a non-linear progression of the zero-counts.

11. The method of claim 7, wherein the processing the corresponding multi-dimensional positions comprises:

converting the zero-counts into coordinates of the multi-dimensional positions;

for each result value, summing the coordinates in each dimension of the multi-dimensional positions with additional coordinates associated with additional non-zero elements to produce the corresponding destination addresses.

12. The method of claim 7, further comprising computing a running sum for each zero-count defined by the sequence of codes to produce the coordinates of the multi-dimensional positions.

13. The method of claim 1, wherein the multi-dimensional positions are three-dimensional positions.

14. The method of claim 1, wherein each result value is a product that is transmitted through an arbitrated crossbar to the destination accumulator.

15. A convolutional neural network accelerator, comprising:

an array of processing elements, each processing element configured to:

receive compressed-sparse data that encodes non-zero elements and corresponding multi-dimensional positions;

process the non-zero elements in parallel to produce a plurality of result values;

process the corresponding multi-dimensional positions in parallel to produce destination addresses for each result value in the plurality of result values; and

transmit each result value to a destination accumulator associated with the destination address for the result value.

16. The convolutional neural network accelerator of claim 14, wherein each processing element comprises a multiplier

array that is configured to produce products that are a portion of the plurality of result values.

17. The convolutional neural network accelerator of claim 14, wherein the multi-dimensional positions are encoded as a sequence of zero-counts and each zero-count is a number of zeros between pairs of the non-zero elements.

18. The convolutional neural network accelerator of claim 17, wherein the processing the corresponding multi-dimensional positions comprises:

converting the zero-counts into coordinates of the multi-dimensional positions; for each result value, summing the coordinates in each dimension of the multi-dimensional positions with additional coordinates associated with additional non-zero elements to produce the corresponding destination addresses.

19. The convolutional neural network accelerator of claim 14, wherein the multi-dimensional positions are encoded as a sequence of codes and each code defines a zero-count that is a number of zeros between pairs of the non-zero elements.

20. A system, comprising:

a memory storing compressed-sparse data that encodes non-zero elements and corresponding multi-dimensional positions; and

a convolutional neural network accelerator, comprising: an array of processing elements, each processing element configured to:

receive the compressed-sparse data;

process the non-zero elements in parallel to produce a plurality of result values;

process the corresponding multi-dimensional positions in parallel to produce destination addresses for each result value in the plurality of result values; and

transmit each result value to a destination accumulator associated with the destination address for the result value.

* * * * *