(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
02 January 2020 (02.01.2020)

WIPO | PCT

(10) International Publication Number
WO 2020/003043 A1

(51) International Patent Classification:
G06N 3/02 (2006.01)

(21) International Application Number:
PCT/IB2019/054960

(22) International Filing Date:
13 June 2019 (13.06.2019)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
16/020,952 27 June 2018 (27.06.2018) US

(71) Applicant: **INTERNATIONAL BUSINESS MACHINES CORPORATION** [US/US]; New Orchard Road, Armonk, New York 10504 (US).

(71) Applicants *(for MG only)*: **IBM UNITED KINGDOM LIMITED** [GB/GB]; PO Box 41, North Harbour, Portsmouth Hampshire PO6 3AU (GB). **IBM (CHINA) INVESTMENT COMPANY LIMITED** [CN/CN]; 25/F, Pangu Plaza, No.27, Central North 4th Ring Road, Chaoyang District, Beijing 100101 (CN).

(72) Inventors: **VENKATARAMANI, Swagath**; IBM Corporation, PO Box 218, 1101 Kitchawan Road, Yorktown Heights, New York 10598 (US). **JAIN, Shubham**; IBM Corporation, PO Box 218, 1101 Kitchawan Road, Yorktown Heights, New York 10598 (US). **SRINIVASAN, Vijayalakshmi**; IBM Corporation, PO Box 218, 1101 Kitchawan Road, Yorktown Heights, New York 10598 (US). **CHOI, Jungwook**; IBM Corporation, PO Box 218, 1101 Kitchawan Road, Yorktown Heights, New York 10598 (US). **CHANG, Leland**; IBM Corporation, PO Box

(54) Title: LOW PRECISION DEEP NEURAL NETWORK ENABLED BY COMPENSATION INSTRUCTIONS

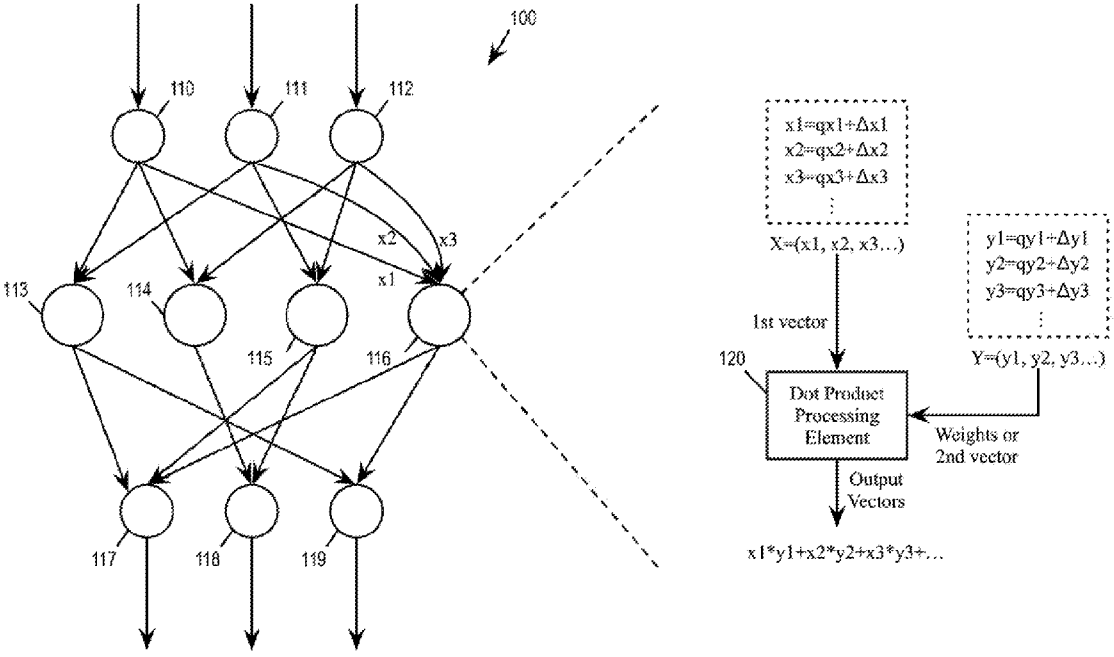

FIG. 1

(57) **Abstract:** A compensated deep neural network (compensated-DNN) is provided. A first vector having a set of components and a second vector having a set of corresponding components are received. A component of the first vector includes a first quantized value and a first compensation instruction,and a corresponding component of the second vector includes a second quantized value and a second compensation instruction. The first quantized value is multiplied with the second quantized value to compute a raw product value. The raw product value is compensated for a quantization error according to the first and second compensation instructions to produce a compensated product value. The compensated product value is added into an accumulated value for the dot product. The accumulated value is converted into an output vector of the dot product. The output vector includes an output quantized value and an output compensation instruction.

WO 2020/003043 A1

[Continued on next page]

218, 1101 Kitchawan Road, Yorktown Heights, New York 10598 (US).

(74) **Agent: WILLIAMS, Julian**; IBM United Kingdom Limited, Intellectual Property Law, Hursley Park, Winchester Hampshire SO21 2JN (GB).

(81) **Designated States** *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

Published:

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

LOW PRECISION DEEP NEURAL NETWORK ENABLED BY COMPENSATION INSTRUCTIONS

BACKGROUND

[0001] The present invention generally relates to computation by neural networks.

[0002] The field of Artificial Intelligence (AI) has witnessed a quintessential growth in recent years with the advent of Deep Neural Networks (DNNs) that have achieved state-of-the-art classification accuracies on many recognition tasks involving images, videos, text and natural language. DNNs are multi-layered networks of basic compute units called artificial neurons. Each layer in the network is associated with a set of weights. Each neuron in a layer evaluates a multi-input, single-output function that computes dot-product of its inputs and weights, followed by a non-linear activation function on the weighted sum. DNNs operate in two phases: (i) Training and (ii) Inference. Training is performed based on a labeled dataset, where the weights of the DNN are iteratively refined using the Stochastic Gradient Descent (SGD) algorithm. During inference, inputs hither to unseen are classified using the trained model.

SUMMARY

[0003] Some embodiments of the invention provide a processing element for an artificial neuron of a deep neural network (DNN). The processing element is configured to produce a dot product based on a first vector having a set of components and a second vector having a set of corresponding components. A component of the first vector includes a first quantized value and a first compensation instruction and a corresponding component of the second vector includes a second quantized value and a second compensation instruction. The processing element includes a computation module configured to multiply the first quantized value with the second quantized value to compute a raw product value. The processing element includes a compensation module configured to compensate the raw product value for a quantization error according to the first and second compensation instructions to produce a compensated product value. The processing element includes an accumulation module configured to add the compensated product value into an accumulated value for the dot product. The processing element includes a conversion module configured to convert the accumulated value into an output vector of the dot-product, the output vector including an output quantized value and an output compensation instruction.

[0004] A compensation instruction includes a direction bit and a magnitude bit for compensating a quantization error of the quantized value. In some embodiments, a compensation instruction includes no more than four bits that include a direction bit, a zero compensation bit, and two or less magnitude bits. The zero compensation bits indicates whether the quantization error is less than a threshold, such that the estimated quantization error is zero and the quantized value need not be compensated. In some embodiments, the compensation module is in a low

power mode when the zero compensation bits of both the first quantization instruction and the second quantization instruction indicate that the estimated quantization errors for both the first and second quantized values are zero.

[0005] The preceding Summary is intended to serve as a brief introduction to some embodiments of the invention. It is not meant to be an introduction or overview of all inventive subject matter disclosed in this document. The Detailed Description that follows and the Drawings that are referred to therein will further describe the embodiments of the invention described in the Summary as well as other embodiments of the invention. Accordingly, to understand all the embodiments of the invention described by this document, a Summary, Detailed Description and the Drawings are provided. Moreover, the claimed subject matter is not to be limited by the illustrative details in the Summary, Detailed Description, and the Drawings, but rather is to be defined by the appended claims, because the claimed subject matter can be embodied in other specific forms without departing from the scope of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

[0006] The drawings are of illustrative embodiments of the invention. They do not illustrate all embodiments of the invention. Other embodiments of the invention may be used in addition or instead. Details that may be apparent or unnecessary may be omitted to save space or for more effective illustration. Some embodiments of the invention may be practiced with additional components or steps and/or without all of the components or steps that are illustrated. When the same numeral appears in different drawings, it refers to the same or like components or steps.

[0007] FIG. 1 illustrates a compensated-DNN in which Multiply And Accumulate (MAC) operations are based on numerical values represented in Fixed Point with Error Compensation (FPEC) format to facilitate dynamic compensation of quantization errors, consistent with an embodiment of the invention.

[0008] FIG. 2 illustrates the FPEC format for representing floating-point numbers, consistent with an embodiment of the invention.

[0009] FIG. 3 illustrates the mapping between quantization errors and compensation bits of FPEC, consistent with an embodiment of the invention.

[0010] FIG. 4 illustrates an example design of a processing element of an artificial neuron.

[0011] FIG. 5 conceptually illustrates a process for computing a dot product between two vectors at an artificial neuron of a DNN, consistent with an embodiment of the invention.

[0012] FIG. 6 shows a block diagram of the components of a data processing system in accordance with an embodiment of the present invention.

DETAILED DESCRIPTION

[0013] In the following detailed description, numerous specific details are set forth by way of examples in order to provide a thorough understanding of the relevant teachings. However, it should be apparent that the present teachings may be practiced without such details. In other instances, well-known methods, procedures, components, and/or circuitry have been described at a relatively high-level, without detail, in order to avoid unnecessarily obscuring aspects of the present teachings.

[0014] DNNs are the state-of-the-art solutions in many recognition problems involving images, video, text, and natural language. However, the computational and storage demands imposed by these large-scale networks have been the primary bottleneck to their ubiquitous adoption, as the amount of data processed by DNNs impose significant computational challenges. A key scenario that exemplifies this extreme computational challenge is low-power inference, where DNN models are executed on deeply-embedded wearable and Internet of Things (IoT) devices that have stringent energy and area constraints. One approach to improve the efficiency of DNNs, specifically in the context of low-power platforms (such as mobile, wearables, and other IoT devices), is to explore low-precision implementations using low-precision fixed point (<16 bits) representation. However, low-precision implementations suffer from quantization errors that are inherent in any fixed-point implementation. The choice of bit-widths is therefore limited if the system is to maintain application-level accuracy. Increasing the network size and/or re-training the DNN has been proposed to minimize loss of accuracy due to quantization, albeit with limited success.

[0015] Some embodiments of the invention provide compensated-DNN, in which errors introduced by quantization are dynamically compensated during execution. Numbers in compensated-DNN are represented in FPEC format. The bits in FPEC are split between computation bits and compensation bits. The computation bits use conventional floating-point notation (FxP) to represent the number at low-precision. The compensation bits explicitly capture an estimate (direction and magnitude) of the quantization error in the representation. For a given word length, FPEC may use fewer computation bits compared to FxP representation. This enables a near-quadratic improvement in energy in the MAC operations in the DNN. In some embodiments, a low-overhead sparse compensation scheme based on the compensation bits is used to estimate the error accrued during MAC operations, which is then added to the MAC output to minimize the impact of quantization.

[0016] It is observed that MAC operations intrinsically exhibit the property of error compensation. Specifically, when MAC operations accumulate multiplication results, the quantization errors of the multiplication operations are also being accrued. If the quantization errors of different multiplication operations take opposite signs, the

quantization errors partially cancel each other out, reducing the quantization error in the eventual dot-product output. In practice, the quantization errors of the multiplication operations may be positive or negative. Whether the quantization error is positive or negative is determined based on several factors, including the rounding mode of the inputs, the sign of the inputs, among others. Regardless, the DNN dynamically estimates the accrued quantization error and explicitly offsets the error, thereby minimizing degradation in classification accuracy.

[0017] FIG. 1 illustrates a compensated-DNN 100 in which the MAC operations are based on numerical values represented in FPEC format to facilitate dynamic compensation of quantization errors, consistent with an embodiment of the invention. The DNN 100 includes artificial neurons 110-119. The artificial neurons 110-119 are organized into multiple interconnected layers. Each neuron receives a set of inputs and performs dot-product based on its received inputs to produce a single output. The inputs to an artificial neuron may include outputs from other artificial neurons and/or primary inputs of the network. Each neuron may be implemented as a module of software instructions or a module of electrical circuits. The inputs to the artificial neurons may be implemented as data structures in a computer memory or electrical wires in an electronic device.

[0018] Each artificial neuron 110-119 computes a dot-product between a first vector and a second vector. The components of the first vector may be a set of values received from the artificial neuron's interconnections. The components of the second vector may be a set of corresponding weights for the components of the first vector. The artificial neuron 116 receives inputs x1, x2, and x3 from outputs of neurons 110, 111, and 112. The artificial neuron 116 applies weights y1, y2, and y3 to the inputs x1, x2, and x3, respectively. The artificial neuron 116 includes a dot-product processing element 120 that is configured to compute a weighted sum of the first vector (corresponds to X vector in the figure) with components [x1, x2, x3] using weights in the second vector (corresponds to Y vector in the figure) with components [y1, y2, y3]. An example dot-product processing element 120 will be described by reference to FIG. 4 below.

[0019] The components of the X vector and the Y vector are in FPEC format. The bit-fields in FPEC are split into 2 groups: computation bits and compensation bits. For a value x (e.g., component x1 in the X vector or component y2 in the Y vector) that is represented in FPEC format, the computation bits provide the quantized value qx. The compensation bits specify an estimate of quantization error Δx, or an estimated error $\Delta x(est)$.

[0020] FIG. 2 illustrates the FPEC format for representing floating-point numbers, consistent with an embodiment of the invention. As illustrated, the bit-fields in FPEC are split into computation bits and compensation bits. The computation bits include a sign bit (Sign), integer bits (IB), and fractional bits (FB) to capture the range and resolution of numbers. The compensation bits specify the direction and magnitude of the error incurred during quantization. The compensation bits are sub-divided into 3 fields: (i) Error Direction Bit (EDB), (ii) Error Magnitude Bit (EMB) and (iii) Error Zero Bit (EZB). EDB indicates whether the number is rounded up or down, i.e., the direction of the quantization error. EMB denotes the magnitude of the quantization error. The compensation bits of a FPEC

number are used to compute an estimated quantization error or a compensation value for compensating the quantized value in the computation bits. The compensation bits of a FPEC are therefore also a compensation instruction.

[0021] If the EDB and EMB fields contain δ and θ respectively, then the estimated error is $-1^{\delta} * 2^{-FB-1-\theta}$. In some embodiments, the error estimates are constrained to powers of 2 to avoid multipliers in the compensation logic. In some embodiments, the EMB field is optional in a FPEC representation. For example, if only the direction of error is specified, θ is assumed to be 0 and the magnitude of the error is half of the resolution i.e., and the estimated error is $\pm 2^{-FB-1}$.

[0022] The EZB field is a single bit that indicates zero quantization error. EZB enables FPEC to limit compensation to only selected computations. When EZB is 1, EDB and EMB fields are don't cares (e.g., EDB and EMB are not processed by the processing element 120).

[0023] FIG. 3 illustrates the mapping between quantization errors and compensation bits of FPEC, consistent with an embodiment of the invention. For a value x represented in FPEC, the figure shows a continuous range 300 between qx and $qx+2^{-FB}$, which are two successive quantized values. The range is divided into a number of regions 301-306, whose estimated quantization errors are represented in FPEC. Quantization causes values above $qx+2^{-FB-1}$ to be rounded up to $qx+2^{-FB}$ and values below $qx+2^{-FB-1}$ to be rounded down to qx. The EDB indicating the direction of the rounding.

[0024] For actual values of x in regions 301 and 306 (close to qx or $qx + 2^{FB}$), the quantization error is sufficiently small. The EZB is therefore set to 1 and the estimated error $\Delta x(est)$ is 0. The size of the region 301 and 306 for setting EZB to 1 is controlled by a threshold that can be modulated to control the degree of sparsity in compensation.

[0025] For values in regions 302-305 where EZB = 0, the EMB field (optionally) splits the range in powers of 2. The values with larger quantization errors have EMB set to 1 (regions 303 and 304) and the estimated error $\Delta x(est)$ encoded in the compensation bits is 2^{-FB-1}. The values with small quantization errors have EMB set to 0 (regions 302 and 305) and the estimated error $\Delta x(est)$ encoded in the compensation bits is 2^{-FB-2}.

[0026] Table 1 below illustrates example values that are represented by FPEC format in which the bit-widths of EDB, EMB, and EZB are all 1. For each actual value represented by FPEC, the table shows its corresponding quantized value, quantization error, FPEC estimated error, and FPEC representation. Note that a 'd' represents a "don't care" bit. The bit-width of IB is 3 and the bit-width of FB is 0.

x (actual value)	qx (quantized value)	Δx (quantization error)	Δx(est) (estimated error by FPEC)	FPEC representation [SB,IB,FB] [EDB,EMB,EZB]
2.4	2	0.4	0.5	0 010 010
2.1	2	0.1	0	0 010 dd1
2.8	3	-0.2	-0.25	0 011 100
3.05	3	0.05	0	0 011 dd1

Table 1: Example values represented by FPEC format

[0027] As mentioned, each artificial neuron computes the dot product of its inputs and weights. The X vector [x_1, x_2, x_3,…] represents inputs and a Y vector [y_1, y_2, y_3,…] represents the weights. The X vector and the Y vector are quantized to QX = [qx_1, qx_2, qx_3,…] and QY =[qy_1, qy_2, qy_3,…], with quantization errors ΔX =[Δx_1, Δx_2, Δx_3,…] and ΔY =[Δy_1, Δy_2, Δy_3,…]. In other words, $X=QX+\Delta X$ where $x_i=qx_i+\Delta x_i$, and $Y=QY+\Delta Y$ where $y_i=qy_i+\Delta y_i$. The dot product is therefore

[0028] $X \bullet Y = QX \bullet QY+\Delta Y \bullet QX+\Delta X \bullet QY+\Delta X \bullet \Delta Y$ Eq. (1)

[0029] The quantization error for the dot product is therefore

[0030] $\Delta(X \bullet Y) = \Delta Y \bullet QX+\Delta X \bullet QY+\Delta X \bullet \Delta Y$ Eq. (2)

[0031] The computation bits of the components of the X and Y vectors provide values for quantized values QX and QY. The compensation bits of the components of X and Y provide estimates for quantization errors ΔX and ΔY, or estimated errors $\Delta X(est)$ and $\Delta Y(est)$. The estimated quantization errors $\Delta X(est)$ and $\Delta Y(est)$ can then be used to compute the estimated quantization error for the dot product:

[0032] $\Delta(X \bullet Y)(est) = \Delta Y(est) \bullet QX+\Delta X(est) \bullet QY+\Delta X(est) \bullet \Delta Y(est)$ Eq. (3)

[0033] In some embodiments of the invention, approximations are used during the computation of the estimated quantization error of the dot product. First, since the magnitude of quantization error is typically smaller than the actual value (i.e., $QX,QY >> \Delta X,\Delta Y$), the higher order term $\Delta X(est) \bullet \Delta Y(est)$ is ignored. Second, input quantization errors are approximated to the nearest power of 2, thereby converting multipliers to shifters when computing $\Delta Y(est) \bullet QX$ and $\Delta X(est) \bullet QY$. Third, at most 1 or 2 bits are used to capture the magnitude of quantization errors (so only 1 or 2 EMB bits and 3 or 4 compensation bits in total), which makes the shifters less expensive. Fourth,

$\Delta X(est)$ and $\Delta Y(est)$ are made sparse by ignoring smaller quantization errors in the input. Therefore, compensation is performed only for certain selected computations, and the compensation logic is active for only a fraction of the overall execution cycles. These approximations significantly reduce the energy expended for estimating the quantization error.

[0034] FIG. 4 illustrates an example design of a processing element 400 of an artificial neuron. The processing element (PE) performs dot product computation between an X vector and a Y vector. The processing element accepts values in FPEC format. The processing element 400 is therefore also referred to as an FPEC-PE. The FPEC-PE performs compensation of quantization errors by computing estimated errors based on the FPEC compensation bits. Approximations are used during the computation of the estimated errors.

[0035] As illustrated, the FPEC-PE 400 includes (i) a Computation unit 410, (ii) a Compensation unit 420, and (iii) a Conversion unit 430. In some embodiments, the Computation unit 410, the Compensation unit 420, and the Conversion unit 430 are modules of software instructions being executed by one or more processing units (e.g., a processor) of a computing device. In some embodiments, the Computation unit 410, the Compensation unit 420, and the Conversion unit 430 are modules of hardware circuits implemented by one or more integrated circuits (ICs) of an electronic apparatus. An example computing device that may implement the FPEC-PE will be described by reference to FIG. 6 below.

[0036] The computation unit 410 includes a multiply-and-accumulate engine 412 that evaluates the dot-product using the values in [IB,FB] fields (computation bits) of FPEC. Specifically, the multiply-and-accumulate engine 412 multiplies a component from the X vector with a corresponding component from the Y vector to produce a raw product value to be added by an adder 414 to an accumulator register 416.

[0037] The compensation unit 420 simultaneously evaluates the quantization error at the dot-product output using the [EDB,EMB,EZB] fields (compensation bits). This involves shifting the X vector component using EMB bits of xi (at shifter 421) and the Y vector component using EMB bits of yi (at shifter 422) and appropriately adding/subtracting (at adder 424) them from the compensation sum (at adder 426 and ErrComp register 428) based on the respective EDB bits. The EZB bits of x_i and y_i combine (logic AND) to render the compensation unit 420 inactive or in low power mode (by e.g., clock gating, signal gating, etc.). After all inputs are processed, the compensation sum (i.e., estimated error for the dot product $\Delta(X \bullet Y)(est)$) is added to the accumulator register 412 in the computation unit 410 to produce a compensated dot-product output.

[0038] The conversion unit 430 quantizes the output to its desired FPEC format. The conversion unit 430 includes a shifter to scale the output based on [IB,FB]. The conversion unit 430 includes a bit-wise logic to infer output compensation bits [EDB,EMB,EZB] based on the scaled bits. The output can be used as a vector to another

processing element that uses the quantized values in the output to compute a raw product value and compensates the computed raw product according to the compensation bits in the output.

[0039] It is worth noting that the PE-FPEC 400 is energy efficient because FPEC representation allows quantization error to be expressed by fewer computation bits. This yields a near-quadratic improvement in the computation unit 410, as the multiplier, adder, and accumulator bit-widths are correspondingly reduced. On the other hand, the adders and registers in the compensation unit 420 do incur a linear increase in logic. However, they are used sporadically, such as when the EZBs of x_i and y_i are asserted (once every 3 to 5 computations in experiments). The overhead in the conversion unit 430 is negligible, as it is active only once per dot-product, which takes thousands of computation cycles in large-scale DNNs.

[0040] By representing values in the FPEC format, an artificial neuron is able to perform dot product operations with minimal quantization errors by using fewer computation bits and very few compensation bits. An architecture based on FPEC such as the processing element 400 can rely on approximations to compute estimated quantization errors with reduced energy expenditure and computation resources. This improves DNNs in terms of size and power, which is critical for lower power applications.

[0041] FIG. 5 conceptually illustrates a process 500 for computing a dot product between two vectors at an artificial neuron of a DNN, consistent with an embodiment of the invention. In some embodiments, a processing element (e.g., PE-FPEC 400) of an artificial neuron performs the process 500. In some embodiments of the invention, one or more processing units (e.g., processor) of a computing device implementing the DNN 100 perform the process 500 by executing instructions stored in a computer readable medium.

[0042] The process 500 starts when the processing element receives (at 510) a first vector having a set of components (e.g., X vector). Each component of the first vector having a first quantized value (e.g., qx) and a first compensation instruction (e.g., estimated quantization error $\Delta X(est)$). The process element also receives (at 520) a second vector having a corresponding set of components (e.g., Y vector). Each component of the second vector has a second quantized value (e.g., qy) and a second compensation instruction (e.g., estimated quantization error $\Delta Y(est)$).

[0043] For some embodiments of the invention, the components of the first vector correspond to input connections to the artificial neuron, while the components of the second vector correspond to weights that are to be applied to the input connections. For a vector component that is in FPEC format, its quantized value is represented by the computation bits [IB,FB] and its compensation instruction includes the compensation bits [EDB, EMB, EZB]. The compensation bits also represent an estimated quantization error for the vector component.

[0044] The processing element multiplies (at 530) a component of the first vector with a corresponding component of the second vector by multiplying the first quantized value with the second quantized value to produce a raw product value as part of the dot product computation.

[0045] The processing element determines (at 535) whether to compensate for quantization errors based on EZB of the first and second compensation instructions. If the EZB of both the first and second compensation instructions are asserted (indicating both $\Delta X(est)$ and $\Delta Y(est)$ are zero), the process proceeds to 560 as there is no need to perform compensation. Otherwise the process proceeds to 540 to perform quantization error compensation.

[0046] At 540, the processing element computes a compensation value (or estimated quantization error) for the raw product value based on EMB and EDB in the first and second compensation instructions. The process unit compensates (at 550) the raw product value by using the computed compensation value to produce a compensated product value. The processing element then adds (at 555) the compensated product value to an accumulated value for the dot product. The process then proceeds to 580.

[0047] At 560, the processing element disables the error compensation operation by placing the compensation unit in low power mode, disabling the compensation unit, etc., since the estimated error of both X component and Y component are zeros. The processing element adds (at 570) the raw product value to the accumulated value for the dot product without computing the compensation value. The process then proceeds to 580.

[0048] At 580, the processing element determines whether there are more components in the first and second vectors that have yet to be multiplied and accumulated for the dot product. If so, the process returns to 530 to process the next pair of components from the first and second vectors. If all of the components have been multiplied and accumulated for the dot product, the process proceeds to 590.

[0049] At 590, the processing element converts the accumulated value into an output vector of the dot product that includes an output quantized value and an output compensation instruction (with EZB, EDB, and EMB). This operation corresponds to the conversion unit 430. The process 500 then ends.

[0050] The present invention may be embodied in a system, a method, and/or a computer program product at any possible technical detail level of integration. The computer program product may include a computer readable storage medium (or media) having computer readable program instructions thereon for causing a processor to carry out aspects of the present invention.

[0051] The computer readable storage medium can be a tangible device that can retain and store instructions for use by an instruction execution device. The computer readable storage medium may be, for example, but is not limited to, an electronic storage device, a magnetic storage device, an optical storage device, an electromagnetic

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

[0052] Computer readable program instructions described herein can be downloaded to respective computing/processing devices from a computer readable storage medium or to an external computer or external storage device via a network, for example, the Internet, a local area network, a wide area network and/or a wireless network. The network may comprise copper transmission cables, optical transmission fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers. A network adapter card or network interface in each computing/processing device receives computer readable program instructions from the network and forwards the computer readable program instructions for storage in a computer readable storage medium within the respective computing/processing device. Computer readable program instructions for carrying out operations of the present disclosure may be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, configuration data for integrated circuitry, or either source code or object code written in any combination of one or more programming languages, including an object oriented programming language such as Smalltalk, C++, or the like, and procedural programming languages, such as the "C" programming language or similar programming languages. The computer readable program instructions may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider). In some embodiments of the invention, electronic circuitry including, for example, programmable logic circuitry, field-programmable gate arrays (FPGA), or programmable logic arrays (PLA) may execute the computer readable program instructions by utilizing state information of the computer readable program instructions to personalize the electronic circuitry, in order to perform aspects of the present invention.

[0053] Aspects of the present invention are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems), and computer program products according to embodiments of the disclosure. It will be understood that each block of the flowchart illustrations and/or block diagrams, and

combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions. These computer readable program instructions may be provided to a processor of a computer or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks. These computer readable program instructions may also be stored in a computer readable storage medium that can direct a computer, a programmable data processing apparatus, and/or other devices to function in a particular manner, such that the computer readable storage medium having instructions stored therein comprises an article of manufacture including instructions which implement aspects of the function/act specified in the flowchart and/or block diagram block or blocks.

[0054] The computer readable program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other device to cause a series of operational steps to be performed on the computer, other programmable apparatus or other device to produce a computer implemented process, such that the instructions which execute on the computer, other programmable apparatus, or other device implement the functions/acts specified in the flowchart and/or block diagram block or blocks. The flowchart and block diagrams in the Figures (e.g., FIG. 5) illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present disclosure. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of instructions, which comprises one or more executable instructions for implementing the specified logical function(s). In some alternative implementations, the functions noted in the blocks may occur out of the order noted in the Figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts or carry out combinations of special purpose hardware and computer instructions.

[0055] FIG. 6 shows a block diagram of the components of data processing systems 600 and 650 that may be used to implement a DNN, an artificial neuron, or a FPEC processing element in accordance with an embodiment of the present invention. It should be appreciated that FIG. 6 provides only an illustration of one implementation and does not imply any limitations with regard to the environments in which different embodiments of the invention may be implemented. Many modifications to the depicted environments may be made based on design and implementation requirements.

[0056] Data processing systems 600 and 650 are representative of any electronic device capable of executing machine-readable program instructions. Data processing systems 600 and 650 may be representative of a smart

phone, a computer system, PDA, or other electronic devices. Examples of computing systems, environments, and/or configurations that may represented by data processing systems 600 and 650 include, but are not limited to, personal computer systems, server computer systems, thin clients, thick clients, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, network PCs, minicomputer systems, and distributed cloud computing environments that include any of the above systems or devices.

[0057] The data processing systems 600 and 650 may include a set of internal components 600 and a set of external components 650 illustrated in FIG. 6. The set of internal components 600 includes one or more processors 620, one or more computer-readable RAMs 622 and one or more computer-readable ROMs 624 on one or more buses 626, and one or more operating systems 628 and one or more computer-readable tangible storage devices 630. The one or more operating systems 628 and programs such as the programs for executing the process 500 are stored on one or more computer-readable tangible storage devices 630 for execution by one or more processors 620 via one or more RAMs 622 (which typically include cache memory). In the embodiment illustrated in FIG. 6, each of the computer-readable tangible storage devices 630 is a magnetic disk storage device of an internal hard drive. Alternatively, each of the computer-readable tangible storage devices 630 is a semiconductor storage device such as ROM 624, EPROM, flash memory or any other computer-readable tangible storage device that can store a computer program and digital information.

[0058] The set of internal components 600 also includes a R/W drive or interface 632 to read from and write to one or more portable computer-readable tangible storage devices 686 such as a CD-ROM, DVD, memory stick, magnetic tape, magnetic disk, optical disk or semiconductor storage device. The instructions for executing the processes 500 can be stored on one or more of the respective portable computer-readable tangible storage devices 686, read via the respective R/W drive or interface 632 and loaded into the respective hard drive 630.

[0059] The set of internal components 600 may also include network adapters (or switch port cards) or interfaces 636 such as a TCP/IP adapter cards, wireless Wi-Fi interface cards, or 3G or 4G wireless interface cards or other wired or wireless communication links. Instructions of processes or programs described above can be downloaded from an external computer (e.g., server) via a network (for example, the Internet, a local area network or other, wide area network) and respective network adapters or interfaces 636. From the network adapters (or switch port adaptors) or interfaces 636, the instructions and data of the described programs or processes are loaded into the respective hard drive 630. The network may comprise copper wires, optical fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers.

[0060] The set of external components 650 can include a computer display monitor 670, a keyboard 680, and a computer mouse 684. The set of external components 650 can also include touch screens, virtual keyboards, touch pads, pointing devices, and other human interface devices. The set of internal components 600 also includes device drivers 640 to interface to computer display monitor 670, keyboard 680 and computer mouse 684. The

device drivers 640, R/W drive or interface 632 and network adapter or interface 636 comprise hardware and software (stored in storage device 630 and/or ROM 624).

[0061] The descriptions of the various embodiments of the present disclosure have been presented for purposes of illustration, but are not intended to be exhaustive or limited to the embodiments disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the described embodiments. The terminology used herein was chosen to best explain the principles of the embodiments, the practical application or technical improvement over technologies found in the marketplace, or to enable others of ordinary skill in the art to understand the embodiments disclosed herein.

CLAIMS

1. A computer-implemented method for computing a dot-product, comprising:

receiving a first vector having a set of components and a second vector having a set of corresponding components, wherein a component of the first vector comprises a first quantized value and a first compensation instruction and a corresponding component of the second vector comprises a second quantized value and a second compensation instruction;

multiplying the first quantized value with the second quantized value to compute a raw product value;

compensating the raw product value for a quantization error according to the first and second compensation instructions to produce a compensated product value;

adding the compensated product value into an accumulated value for the dot product; and

converting the accumulated value into an output vector of the dot product, the output vector comprising an output quantized value and an output compensation instruction.

2. The computer-implemented method of claim 1, wherein the dot-product is a first dot-product, the method further comprising using the output vector as a vector to compute a second raw product value for a second dot-product and to compensate the second raw product according to the output compensation instruction.

3. The computer-implemented method of claim 1, wherein:

the first compensation instruction comprises a first direction bit and a first magnitude bit for compensating a quantization error of the first quantized value; and

the second compensation instruction comprises a second direction bit and a second magnitude bit for compensating a quantization error of the second quantized value.

4. The computer-implemented method of claim 1, wherein compensating the raw product value for a quantization error according to the first and second compensation instructions to produce a compensated product value comprises:

shifting the first quantization value according to the first magnitude bit;

shifting the second quantization value according to the second magnitude bit;

adding the shifted first quantization value and the shifted second quantization value to produce a compensation value according to the first and second direction bits; and

adding the compensation value with the raw product value to produce the compensated product value.

5. The computer-implemented method of claim 1, further comprising:

placing a compensation circuit in a low power mode upon determining that the first compensation instruction includes a bit that indicates that a quantization error of the first quantized value is less than a threshold

and the second compensation instruction includes a bit that indicates that a quantization error of the second quantized value is less than a threshold.

6. The computer-implemented method of claim 1, wherein the first compensation instruction comprises no more than four bits, comprising a direction bit, a zero compensation bit, and two or less magnitude bits.

7. A computing device comprising:

a processor; and

a storage device storing a set of instructions, wherein an execution of the set of instructions by the processor configures the computing device to perform the steps of a method as claimed in any of claims 1 to 6.

8. The computing device of claim 1, wherein the output compensation instruction comprises a bit to indicate whether the output quantized value is to be compensated by an estimated quantization error.

9. An integrated circuit (IC) comprising:

a processing element circuit configured to produce a dot product based on a first vector having a set of components and a second vector having a set of corresponding components, wherein a component of the first vector comprises a first quantized value and a first compensation instruction and a corresponding component of the second vector comprises a second quantized value and a second compensation instruction,

the processing element circuit comprising:

a computation circuit configured to multiply the first quantized value with the second quantized value to compute a raw product value;

a compensation circuit configured to compensate the raw product value for a quantization error according to the first and second compensation instructions to produce a compensated product value; and

an accumulation circuit configured to add the compensated product value into an accumulated value for the dot product.

10. The IC of claim 9, wherein the processing element further comprises a conversion circuit configured to convert the accumulated value into an output vector of the dot-product, the output vector comprising an output quantized value and an output compensation instruction.

11. The IC of claim 10, wherein the processing element is a first processing element, wherein the output vector is used as a vector to a second processing element that is configured to use the output quantized value to compute a second raw product value and compensate the second raw product according to the output compensation instruction.

12. The IC of claim 10, wherein the output compensation instruction comprises a bit to indicate whether an estimated quantization error for the output quantized value is zero.

13. The IC of claim 9, wherein:

the first compensation instruction comprises a first direction bit and a first magnitude bit for compensating a quantization error of the first quantized value; and

the second compensation instruction comprises a second direction bit and a second magnitude bit for compensating a quantization error of the second quantized value.

14. The IC of claim 13, wherein compensating the raw product value for a quantization error according to the first and second compensation instructions to produce a compensated product value comprises:

shifting the first quantization value according to the first magnitude bit;

shifting the second quantization value according to the second magnitude bit;

adding the shifted first quantization value and the shifted second quantization value to produce a compensation value according to the first and second direction bits; and

adding the compensation value with the raw product value to produce the compensated product value.

15. The IC of claim 9, wherein the compensation circuit is in a low power mode upon determining that the first compensation instruction includes a bit that indicates that an estimated quantization error of the first quantized value is zero and the second compensation instruction includes a bit that indicates that an estimated quantization error of the second quantized value is zero.

16. The IC of claim 9, wherein the first compensation instruction comprises no more than four bits, comprising a direction bit, a zero compensation bit, and two or less magnitude bits.

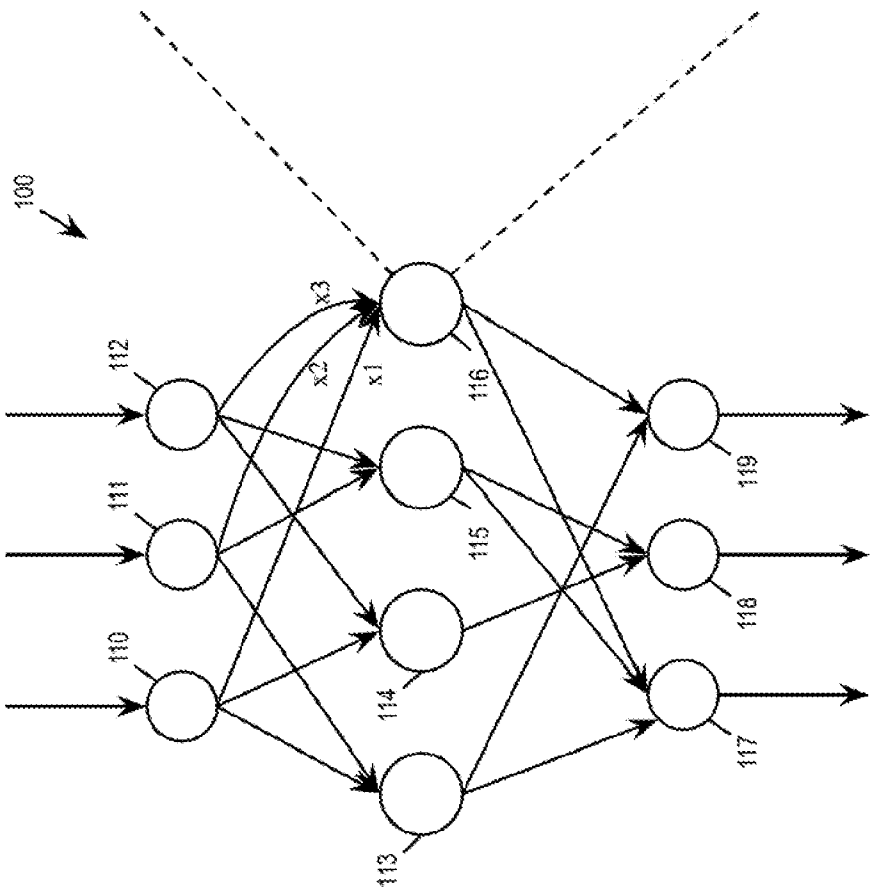
100
110
111
112
113
114
115
116
117
118
119
x1
x2
x3

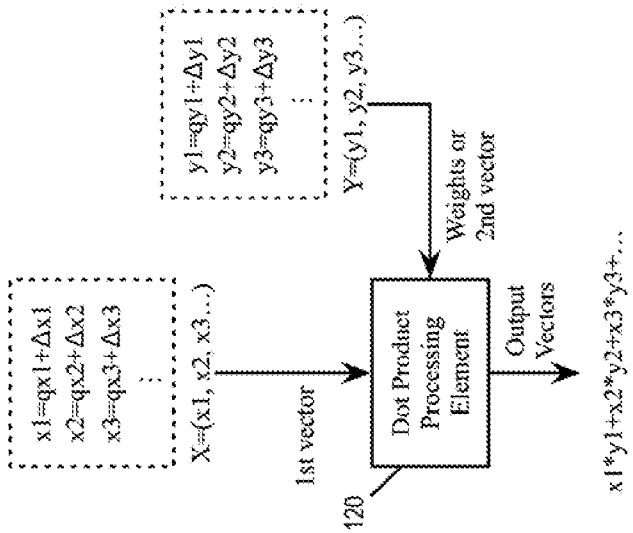
x1=qx1+Δx1
x2=qx2+Δx2
x3=qx3+Δx3
X=(x1, x2, x3...)
1st vector
y1=qy1+Δy1
y2=qy2+Δy2
y3=qy3+Δy3
Y=(y1, y2, y3...)
Weights or 2nd vector
Dot Product Processing Element
120
Output Vectors
x1*y1+x2*y2+x3*y3+...

FIG. 1

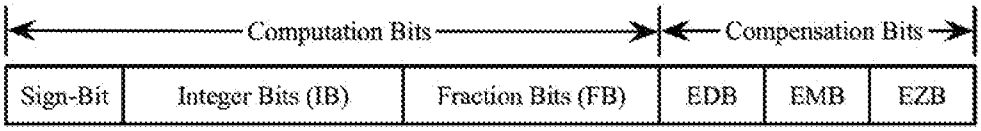
Computation Bits
Compensation Bits
Sign-Bit
Integer Bits (IB)
Fraction Bits (FB)
EDB
EMB
EZB

FIG. 2

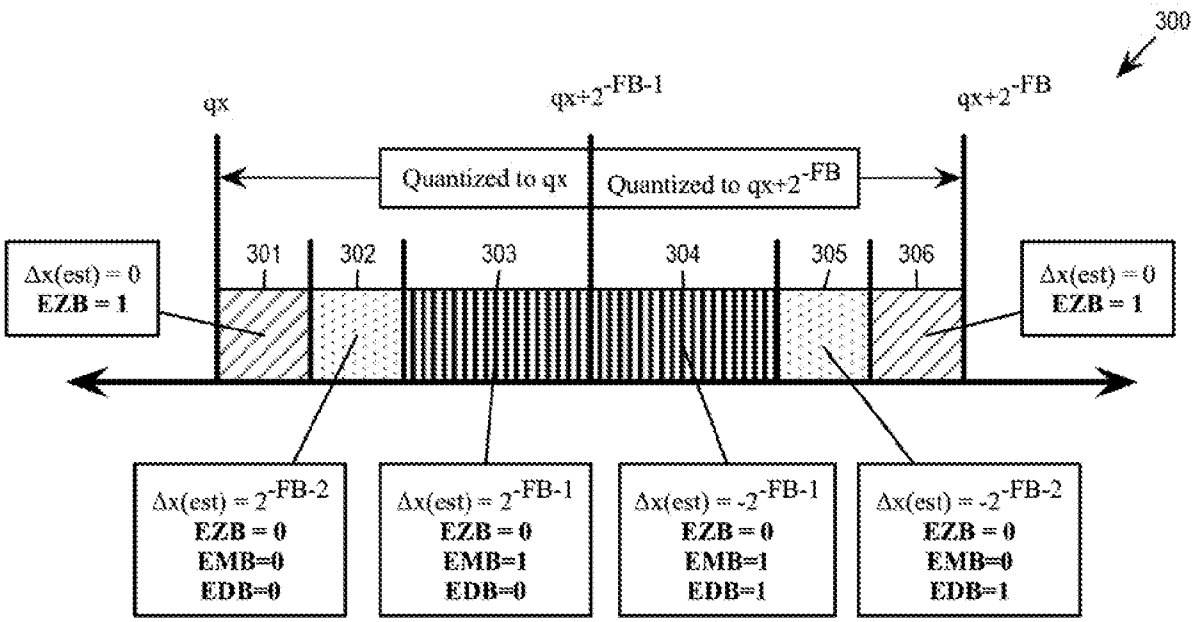
300
qx
qx+2-FB-1
qx+2-FB
Quantized to qx
Quantized to qx+2-FB
301
302
303
304
305
306
Δx(est) = 0
EZB = 1
Δx(est) = 0
EZB = 1
Δx(est) = 2-FB-2
EZB = 0
EMB=0
EDB=0
Δx(est) = 2-FB-1
EZB = 0
EMB=1
EDB=0
Δx(est) = -2-FB-1
EZB = 0
EMB=1
EDB=1
Δx(est) = -2-FB-2
EZB = 0
EMB=0
EDB=1

FIG. 3

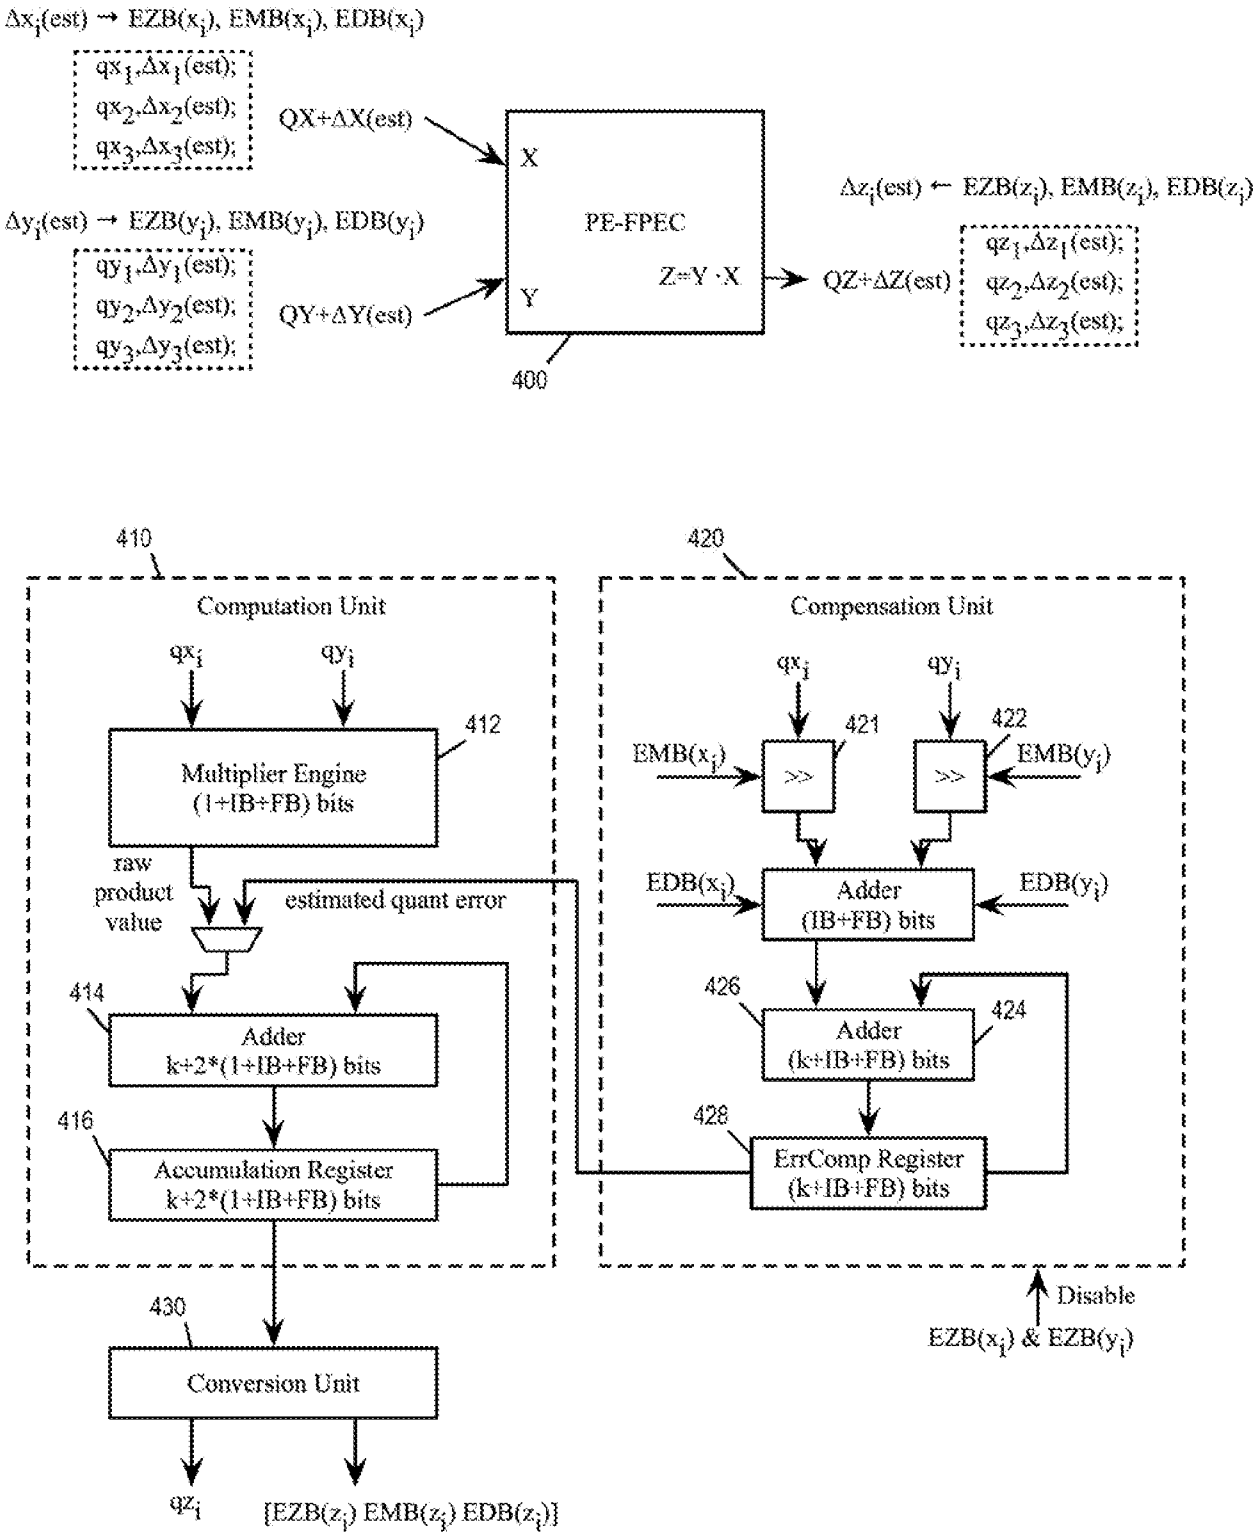
Δx_i(est) → EZB(x_i), EMB(x_i), EDB(x_i)
qx_1,Δx_1(est);
qx_2,Δx_2(est);
qx_3,Δx_3(est);
QX+ΔX(est)
X
PE-FPEC
Z=Y ·X
Y
QY+ΔY(est)
Δy_i(est) → EZB(y_i), EMB(y_i), EDB(y_i)
qy_1,Δy_1(est);
qy_2,Δy_2(est);
qy_3,Δy_3(est);
400
Δz_i(est) ← EZB(z_i), EMB(z_i), EDB(z_i)
QZ+ΔZ(est)
qz_1,Δz_1(est);
qz_2,Δz_2(est);
qz_3,Δz_3(est);
410
Computation Unit
qx_i
qy_i
412
Multiplier Engine
(1+IB+FB) bits
raw
product
value
estimated quant error
414
Adder
k+2*(1+IB+FB) bits
416
Accumulation Register
k+2*(1+IB+FB) bits
430
Conversion Unit
qz_i
[EZB(z_i) EMB(z_i) EDB(z_i)]
420
Compensation Unit
qx_i
qy_i
421
422
EMB(x_i)
EMB(y_i)
>>
>>
EDB(x_i)
Adder
(IB+FB) bits
EDB(y_i)
426
Adder
(k+IB+FB) bits
424
428
ErrComp Register
(k+IB+FB) bits
Disable
EZB(x_i) & EZB(y_i)

FIG. 4

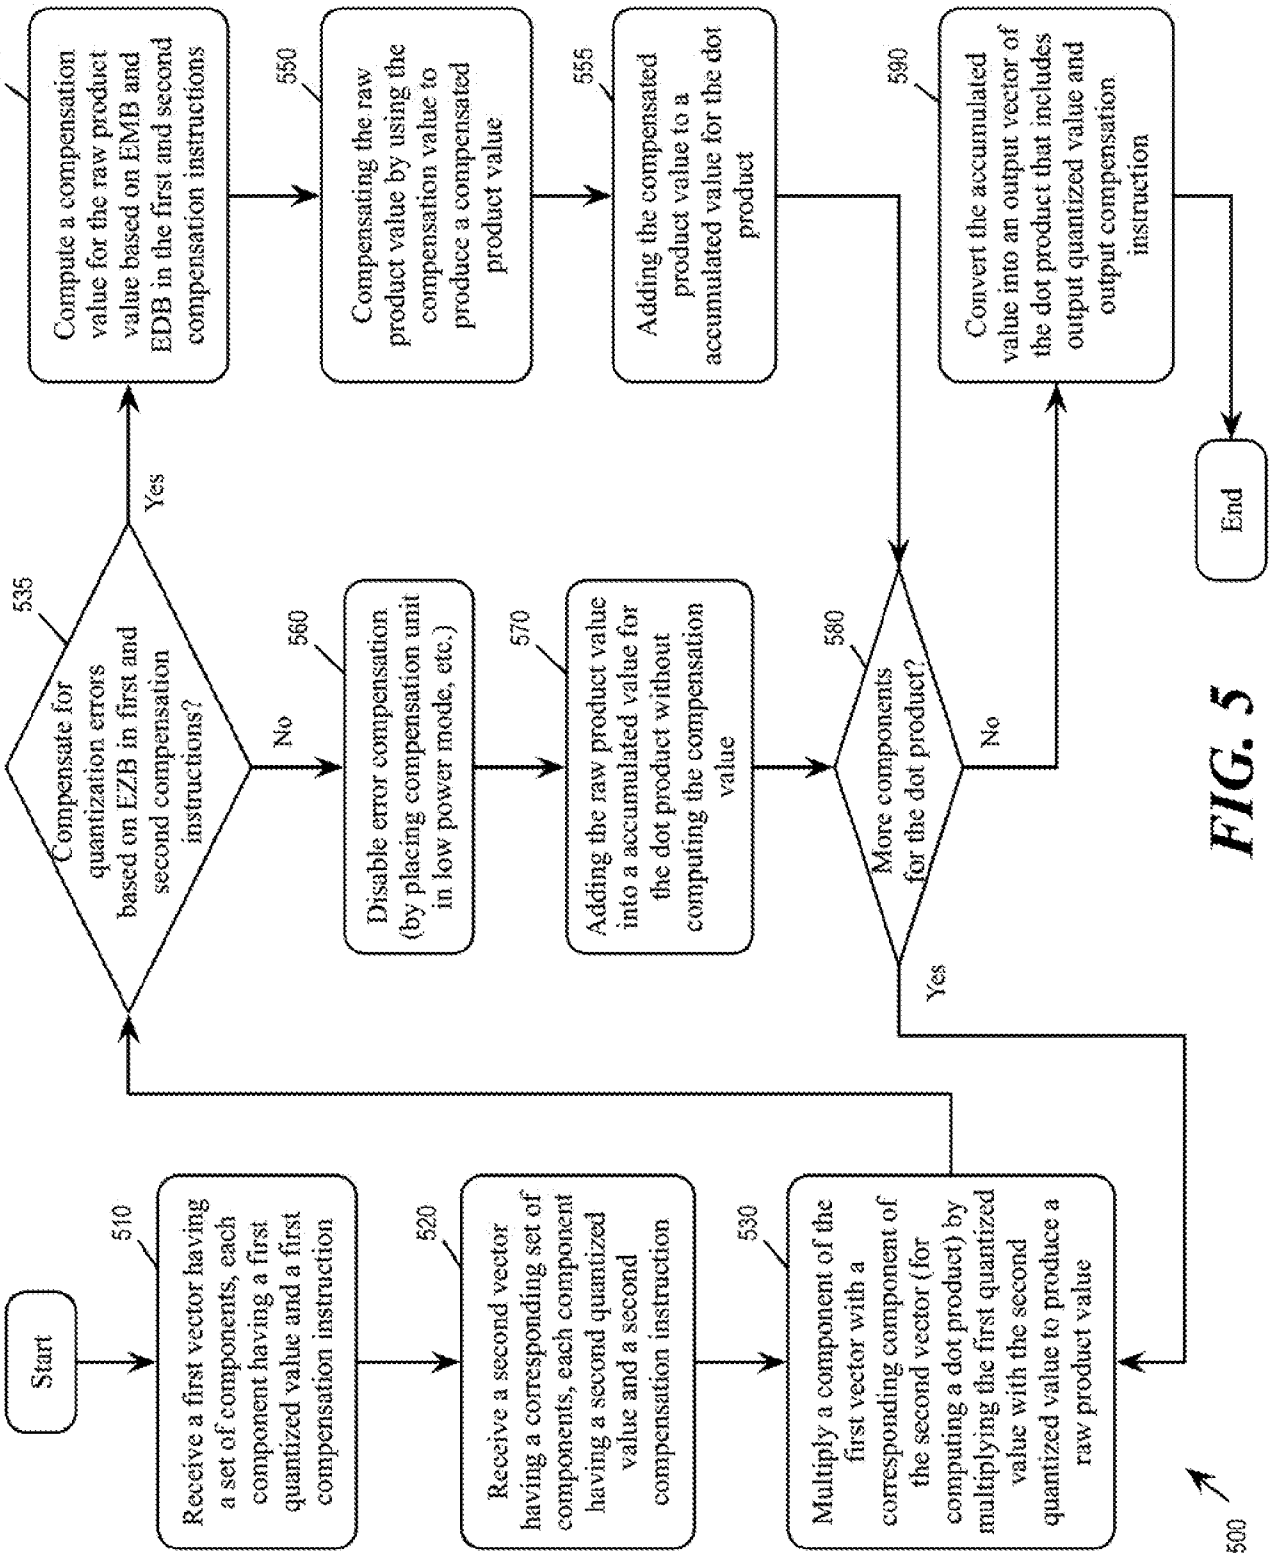
Start
510
Receive a first vector having a set of components, each component having a first quantized value and a first compensation instruction
520
Receive a second vector having a corresponding set of components, each component having a second quantized value and a second compensation instruction
530
Multiply a component of the first vector with a corresponding component of the second vector (for computing a dot product) by multiplying the first quantized value with the second quantized value to produce a raw product value
535
Compensate for quantization errors based on EZB in first and second compensation instructions?
Yes
No
540
Compute a compensation value for the raw product value based on EMB and EDB in the first and second compensation instructions
550
Compensating the raw product value by using the compensation value to produce a compensated product value
555
Adding the compensated product value to a accumulated value for the dot product
560
Disable error compensation (by placing compensation unit in low power mode, etc.)
570
Adding the raw product value into a accumulated value for the dot product without computing the compensation value
580
More components for the dot product?
Yes
No
590
Convert the accumulated value into an output vector of the dot product that includes output quantized value and output compensation instruction
End
500

FIG. 5

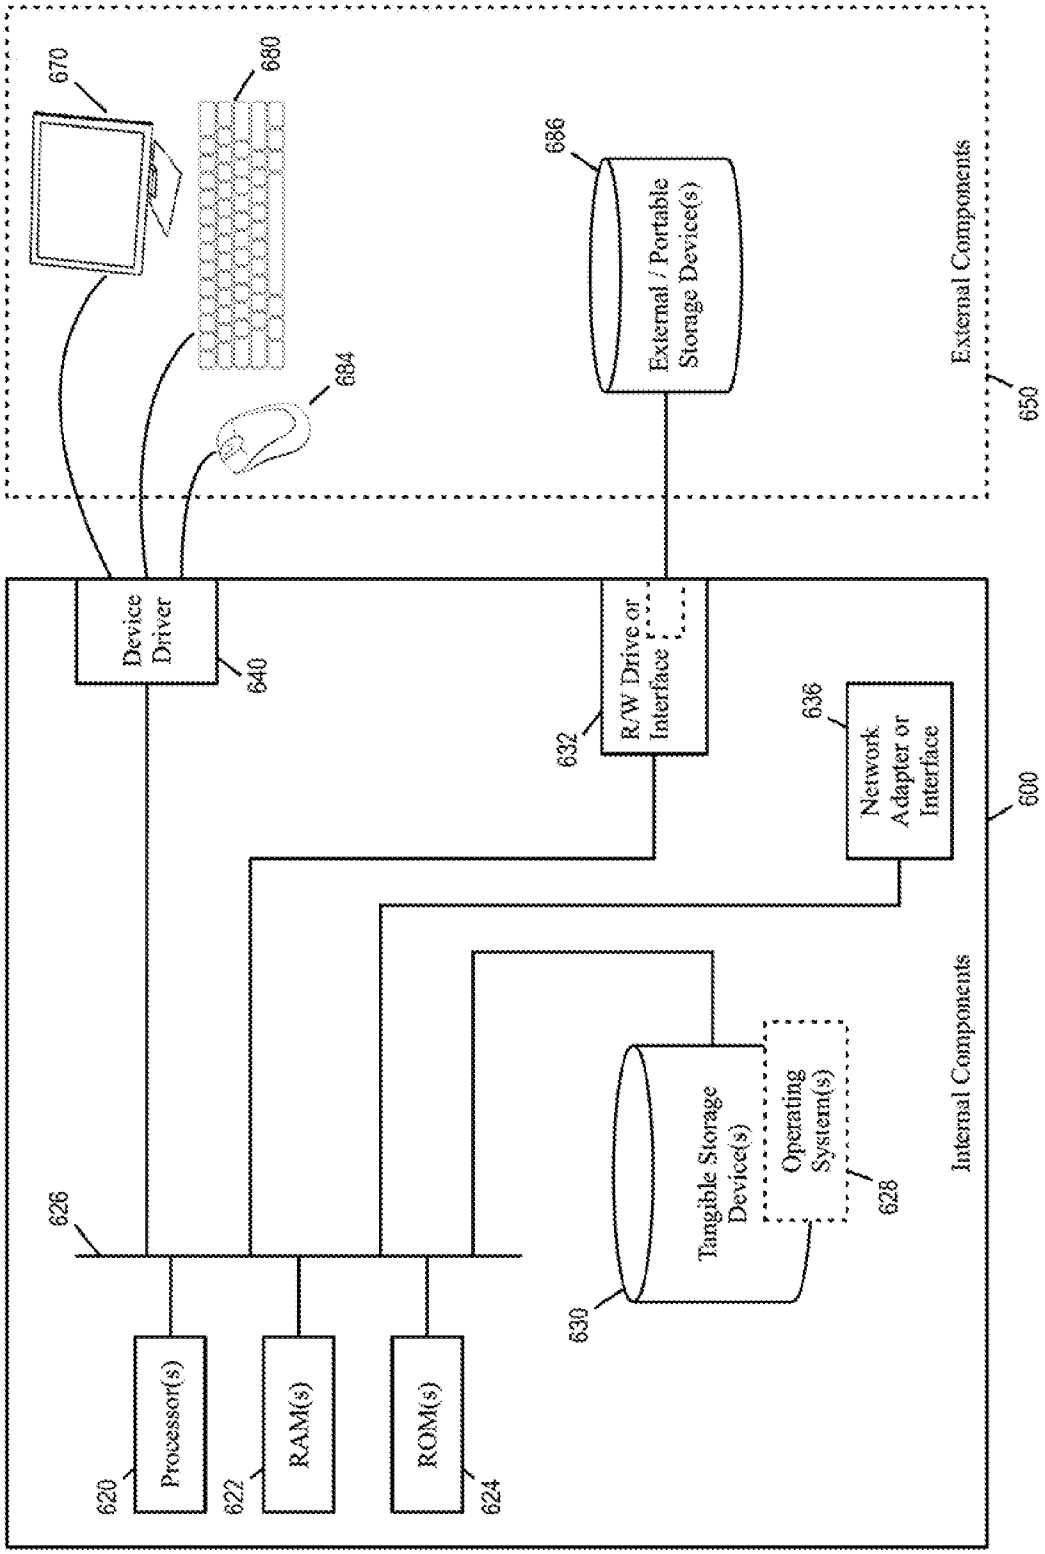
670
680
684
686
External / Portable Storage Device(s)
External Components
650
Device Driver
640
R/W Drive or Interface
632
636
Network Adapter or Interface
600
626
Processor(s)
620
RAM(s)
622
ROM(s)
624
Tangible Storage Device(s)
630
Operating System(s)
628
Internal Components

FIG. 6

INTERNATIONAL SEARCH REPORT

International application No.
PCT/IB2019/054960

A. CLASSIFICATION OF SUBJECT MATTER

G06N 3/02(2006.01)i

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06N, G06F

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

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

CNABS, DWPI, SIPOABS, CNKI: low, precision, deep, neural network, compensation, dot product, vector, quantized, multiply, add, convert, accumulate, error, raw, instruction, value

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2018129937 A1 (SALEFORCE COM INC) 10 May 2018 (2018-05-10) the whole document	1-16
A	CN 106203622 A (HANGZHOU HUAWEI DIGITAL TECHNOLOGY CO LTD) 07 December 2016 (2016-12-07) the whole document	1-16
A	CN 107862378 A (VERISILICON MICROELECTRONICS SHANGHAI CO LTD ET AL.) 30 March 2018 (2018-03-30) the whole document	1-16

☐ Further documents are listed in the continuation of Box C. ☑ See patent family annex.

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
17 September 2019	**30 September 2019**

Name and mailing address of the ISA/CN	Authorized officer
National Intellectual Property Administration, PRC **6, Xitucheng Rd., Jimen Bridge, Haidian District, Beijing 100088** **China**	**ZHAN,Qianqian**
Facsimile No. **(86-10)62019451**	Telephone No. **86-(10)-62411115**

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/IB2019/054960

Patent document cited in search report	Publication date (day/month/year)	Patent family member(s)	Publication date (day/month/year)
US 2018129937 A1	10 May 2018	CA 3040153 A1	11 May 2018
		WO 2018085724 A1	11 May 2018
		CA 3040188 A1	11 May 2018
		AU 2017355537 A1	02 May 2019
		CN 109923559 A	21 June 2019
		AU 2017355535 A1	02 May 2019
		US 2018129931 A1	10 May 2018
		WO 2018085722 A1	11 May 2018
		CN 109952580 A	28 June 2019
CN 106203622 A	07 December 2016	CN 106203622 B	25 December 2018
CN 107862378 A	30 March 2018	None	