



- (51) International Patent Classification:  
Not classified
- (21) International Application Number:  
PCT/IB2018/000995
- (22) International Filing Date:  
30 July 2018 (30.07.2018)
- (25) Filing Language:  
English
- (26) Publication Language:  
English
- (30) Priority Data:  
62/538,722 30 July 2017 (30.07.2017) US  
62/538,724 30 July 2017 (30.07.2017) US  
62/548,990 23 August 2017 (23.08.2017) US
- (72) Inventors; and  
(71) Applicants: SITY, Elad [IL/IL]; Belinson 8, Kfar Saba (IL). HILLEL, Eliad [IL/IL]; Moran 10, Kfar Saba (IL).
- (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,

(54) Title: A MEMORY-BASED DISTRIBUTED PROCESSOR ARCHITECTURE

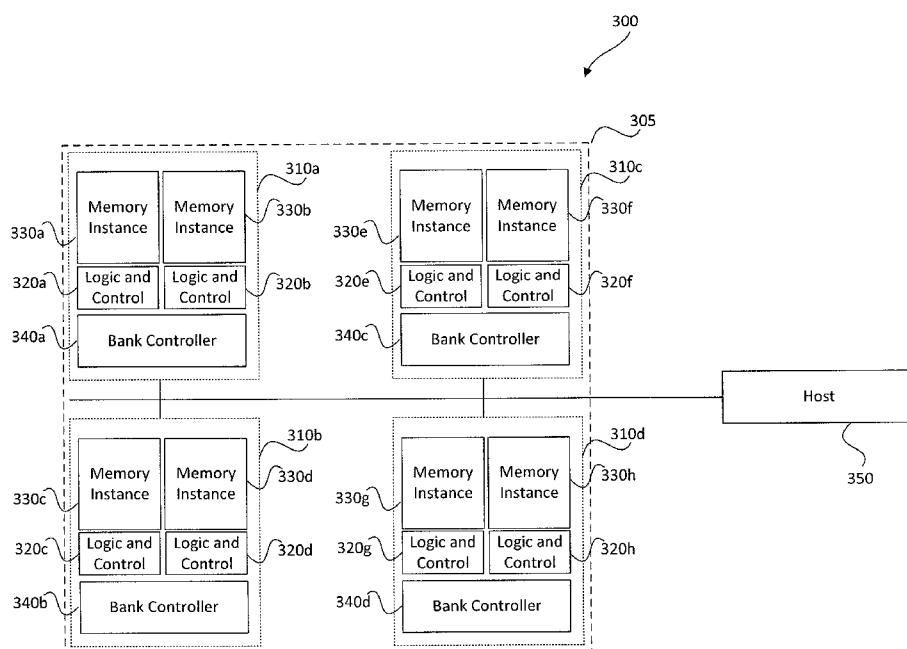


FIG. 3A

(57) Abstract: The present disclosure includes distributed processors and methods for compiling code for executed by the distributed processors. In one implementation, a distributed processor may include a substrate; a memory array disposed on the substrate; and a processing array disposed on the substrate. The memory array may include a plurality of discrete memory banks, and the processing array may include a plurality of processor subunits, each one of the processor subunits being associated with a corresponding, dedicated one of the plurality of discrete memory banks. The distributed processor may further include a first plurality of buses, each connecting one of the plurality of processor subunits to its corresponding, dedicated memory bank, and a second plurality of buses, each connecting one of the plurality of processor subunits to another of the plurality of processor subunits.



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:**

- *without international search report and to be republished upon receipt of that report (Rule 48.2(g))*

## A MEMORY-BASED DISTRIBUTED PROCESSOR ARCHITECTURE

### Cross References to Related Applications

[001] This application claims the benefit of priority of United States Provisional Patent Application No. 62/538,722, filed on July 30, 2017, United States Provisional Patent Application No. 5 62/538,724, filed on July 30, 2017, and United States Provisional Patent Application No. 62/548,990, filed on August 23, 2017. All of the foregoing applications are incorporated herein by reference in their entireties.

### BACKGROUND

#### Technical Field

10 [002] The present disclosure relates generally to apparatuses for facilitating memory-intensive operations. In particular, the present disclosure relates to hardware chips that include processing elements coupled to dedicated memory banks.

#### Background Information

15 [003] As processor speeds and memory sizes both continue to increase, a significant limitation on effective processing speeds is the von Neumann bottleneck. The von Neumann bottleneck results from throughput limitations resulting from conventional computer architecture. In particular, data transfer from memory to the processor is often bottlenecked compared to actual computations undertaken by the processor. Accordingly, the number of clock cycles to read and write from memory increases significantly with memory-intensive processes. These clock cycles result in lower effective processing 20 speeds because reading and writing from memory consumes clock cycles that cannot be used for performing operations on data. Moreover, the computational bandwidth of the processor is generally larger than the bandwidth of the buses that the processor uses to access the memory.

[004] These bottlenecks are particularly pronounced for memory-intensive processes, such as neural network and other machine learning algorithms; database construction, indexing searching, and 25 querying; and other tasks that include more reading and writing operation than data processing operations.

[005] Additionally, the rapid growth in volume and granularity of available digital data has created opportunities to develop machine learning algorithms and has enabled new technologies. However, it has also brought cumbersome challenges to the world of data bases and parallel computing. For example, the rise of social media and the Internet of Things (IoT) creates digital data at a record rate. 30 This new data can be used to create algorithms for a variety of purposes, ranging from new advertising techniques to more precise control methods of industrial processes. However, the new data has been difficult to store, process, analyze and handle.

[006] New data resources can be massive, sometimes in the order of peta- to zettabytes. Moreover, the growth rate of these data resources may exceed data processing capabilities. Therefore, 35 data scientists have turned to parallel data processing techniques, to tackle these challenges. In an effort to increase computation power and handle the massive amount of data, scientists have attempted to create systems and methods capable of parallel intensive computing. But these existing systems and methods have not kept up with the data processing requirements, often because the techniques employed are

limited by their demand of additional resources for data management, integration of segregated data, and analysis of the sectioned data.

[007] To facilitate the manipulation of large data sets, engineers and scientists now seek to improve the hardware used to analyze data. For example, new semiconductor processors or chips (such as those described herein) may be designed specifically for data intensive tasks by incorporating memory and processing functions in a single substrate fabricated in technologies more fitting for memory operations rather than arithmetic computation. With integrated circuits specifically designed for data-intensive tasks, it is possible to meet the new data processing requirements. Nonetheless, this new approach to tackle data processing of large data sets requires solving new issues in chip design and fabrication. For instance, if the new chips designed for data intensive tasks are manufactured with fabrication techniques and architectures used for common chips, they would have a poor performance and/or unacceptable yields. In addition, if the new chips are designed to operate with current data handling methods, they will have poor performance because current methods can limit the chip's ability to handle parallel operations.

[008] The present disclosure describes solutions for mitigating or overcoming one or more of the problems set forth above, among other problems in the prior art.

#### SUMMARY

[009] Embodiments consistent with the present disclosure provide apparatuses including hardware processing chips. The disclosed embodiments may use dedicated memory banks paired with processing elements to provide more efficient effective processing speeds than conventional processors. For example, consistent with the disclosed embodiments, the disclosed chips may include dedicated buses between each processing element and its corresponding memory banks. In addition, the disclosed chips may be free of arbiters and/or other hardware that controls timing of data transfers between the processing elements. Other disclosed non-transitory computer-readable media may store instructions for compiling higher-level instructions to lower-level instructions executed by hardware chips disclosed herein.

[010] Some embodiments of the present disclosure include a distributed processor, comprising: a substrate; a memory array disposed on the substrate, the memory array including a plurality of discrete memory banks; a processing array disposed on the substrate, the processing array including a plurality of processor subunits, each one of the processor subunits being associated with a corresponding, dedicated one of the plurality of discrete memory banks; a first plurality of buses, each connecting one of the plurality of processor subunits to its corresponding, dedicated memory bank; and a second plurality of buses, each connecting one of the plurality of processor subunits to another of the plurality of processor subunits.

[011] Other embodiments consistent with the present disclosure include a memory chip, comprising: a substrate; a memory array disposed on the substrate, the memory array including a plurality of discrete memory banks; a processing array disposed on the substrate, the processing array including a plurality of address generators, each one of the address generators being associated with a corresponding,

dedicated one of the plurality of discrete memory banks; and a plurality of buses, each connecting one of the plurality of address generators to its corresponding, dedicated memory bank.

[012] Another embodiment consistent with the present disclosure may include a distributed processor, comprising: a substrate; a memory array disposed on the substrate, the memory array including  
5 a plurality of discrete memory banks, wherein each of the discrete memory banks has a capacity greater than one megabyte; and a processing array disposed on the substrate, the processing array including a plurality of processor subunits, each one of the processor subunits being associated with a corresponding, dedicated one of the plurality of discrete memory banks.

[013] Still other embodiments consistent with the present disclosure may include a distributed  
10 processor, comprising: a substrate; a memory array disposed on the substrate, the memory array including a plurality of discrete memory banks; and a processing array disposed on the substrate, the processing array including a plurality of processor subunits, each one of the processor subunits being associated with a corresponding, dedicated one of the plurality of discrete memory banks; and a plurality of buses, each one of the plurality of buses connecting one of the plurality of processor subunits to at least another one  
15 of the plurality of processor subunits, wherein the plurality of buses are free of timing hardware logic components such that data transfers between processor subunits and across corresponding ones of the plurality of buses are uncontrolled by timing hardware logic components.

[014] Other embodiments may include a distributed processor on a memory chip, comprising: a  
20 substrate; a memory array disposed on the substrate, the memory array including a plurality of discrete memory banks; and a processing array disposed on the substrate, the processing array including a plurality of processor subunits, each one of the processor subunits being associated with a corresponding, dedicated one of the plurality of discrete memory banks; and a plurality of buses, each one of the plurality of buses connecting one of the plurality of processor subunits to a corresponding, dedicated one of the  
25 plurality of discrete memory banks, wherein the plurality of buses are free of timing hardware logic components such that data transfers between a processor subunit and a corresponding, dedicated one of the plurality of discrete memory banks and across a corresponding one of the plurality of buses are uncontrolled by timing hardware logic components.

[015] Other embodiments may include a distributed processor, comprising: a substrate; a  
30 memory array disposed on the substrate, the memory array including a plurality of discrete memory banks; and a processing array disposed on the substrate, the processing array including a plurality of processor subunits, each one of the processor subunits being associated with a corresponding, dedicated one of the plurality of discrete memory banks; and a plurality of buses, each one of the plurality of buses connecting one of the plurality of processor subunits to at least another one of the plurality of processor subunits, wherein the plurality of processor subunits are configured to execute software that controls  
35 timing of data transfers across the plurality of buses to avoid colliding data transfers on at least one of the plurality of buses.

[016] Other embodiments may include a distributed processor on a memory chip, comprising: a substrate; a plurality of processor subunits disposed on the substrate, each processor subunit being

configured to execute a series of instructions independent from other processor subunits, each series of instructions defining a series of tasks to be performed by a single processor subunit; a corresponding plurality of memory banks disposed on the substrate, each one of the plurality processor subunits being connected to at least one dedicated memory bank not shared by any others of the plurality of processor subunits; and a plurality of buses, each of the plurality of buses connecting one of the plurality of processor subunits to at least one other of the plurality of processor subunits, wherein data transfers across at least one of the plurality of buses are predefined by the series of instructions included in a processor subunit connected to the at least one of the plurality of buses.

[017] Other embodiments may include a distributed processor on a memory chip, comprising: a plurality of processor subunits disposed on the memory chip; a plurality of memory banks disposed on the memory chip, wherein each one of the plurality of memory banks is configured to store data independent from data stored in other ones of the plurality of memory banks, and wherein each one of the plurality of processor subunits is connected to at least one dedicated memory bank from among the plurality of memory banks; and a plurality of buses, wherein each one of the plurality of buses connects one of the plurality of processor subunits to one or more corresponding, dedicated memory banks from among the plurality of memory banks, wherein data transfers across a particular one of the plurality of buses are controlled by a corresponding processor subunit connected to the particular one of the plurality of buses.

[018] Other embodiments may include a distributed processor on a memory chip, comprising: a plurality of processor subunits disposed on the memory chip; a plurality of memory banks disposed on the memory chip, wherein each one of the plurality of processor subunits is connected to at least one dedicated memory bank from among the plurality of memory banks, and wherein each memory bank of the plurality of memory banks is configured to store data independent from data stored in other ones of the plurality of memory banks, and wherein at least some of the data stored in one particular memory bank from among the plurality of memory banks comprises a duplicate of data stored in at least another one of the plurality of memory banks; and a plurality of buses, wherein each one of the plurality of buses connects one of the plurality of processor subunits to one or more corresponding, dedicated memory banks from among the plurality of memory banks, wherein data transfers across a particular one of the plurality of buses are controlled by a corresponding processor subunit connected to the particular one of the plurality of buses.

[019] Other embodiments may include a distributed processor on a memory chip, comprising: a plurality of processor subunits disposed on the memory chip; a plurality of memory banks disposed on the memory chip, wherein each one of the plurality of processor subunits is connected to at least one dedicated memory bank from among the plurality of memory banks, and wherein each memory bank of the plurality of memory banks is configured to store data independent from data stored in other ones of the plurality of memory banks, and wherein at least some of the data stored in one particular memory bank from among the plurality of memory banks comprises a duplicate of data stored in at least another one of the plurality of memory banks; and a plurality of buses, wherein each one of the plurality of buses connects one of the plurality of processor subunits to one or more corresponding, dedicated memory

banks from among the plurality of memory banks, wherein data transfers across a particular one of the plurality of buses are controlled by a corresponding processor subunit connected to the particular one of the plurality of buses.

[020] Other embodiments may include a non-transitory computer-readable medium storing instructions for compiling a series of instructions for execution on a memory chip comprising a plurality of processor subunits and a plurality of memory banks, wherein each processor subunit from among the plurality of processor subunits is connected to at least one corresponding, dedicated memory bank from among the plurality of memory banks, the instructions causing at least one processor to: divide the series of instructions into a plurality of groups of sub-series instructions, the division comprising: assigning tasks associated with the series of instructions to different ones of the processor subunits, wherein the processor subunits are spatially distributed among the plurality of memory banks disposed on the memory chip; generating tasks to transfer data between pairs of the processor subunits of the memory chip, each pair of processor subunits being connected by a bus, and grouping the assigned and generated tasks into the plurality of groups of sub-series instructions, wherein each of the plurality of groups of sub-series instructions corresponds to a different one of the plurality of processor sub-units; generate machine code corresponding to each of the plurality of groups of subs-series instructions; and assign the generated machine code corresponding to each of the plurality of groups of subs-series instructions to a corresponding one of the plurality of processor subunits in accordance with the division.

[021] Other embodiments may include a memory chip, comprising: a plurality of memory banks, each memory bank having a bank row decoder, a bank column decoder, and a plurality of memory sub-banks, each memory sub-bank having a sub-bank row decoder and a sub-bank column decoder for allowing reads and writes to locations on the memory sub-bank, each memory sub-bank comprising: a plurality of memory mats, each memory mat having a plurality of memory cells, wherein the sub-bank row decoders and the sub-bank column decoders are connected to the bank row decoder and the bank column decoder.

[022] Other embodiments may include a memory chip, comprising: a plurality of memory banks, each memory bank having a bank controller and a plurality of memory sub-banks, each memory sub-bank having a sub-bank row decoder and a sub-bank column decoder for allowing reads and writes to locations on the memory sub-bank, each memory sub-bank comprising: a plurality of memory mats, each memory mat having a plurality of memory cells, wherein the sub-bank row decoders and the sub-bank column decoders process read and write requests from the bank controller.

[023] Other embodiments may include a memory chip, comprising: a plurality of memory banks, each memory bank having a having a bank controller for processing reads and writes to locations on the memory bank, each memory bank comprising: a plurality of memory mats, each memory mat having a plurality of memory cells and having a mat row decoder and a mat column decoder, wherein the mat row decoders and the mat column decoders process read and write requests from the sub-bank controller.

[024] Other embodiments may include a memory chip, comprising: a plurality of memory banks, each memory bank having a bank controller, a row decoder, and a column decoder for allowing reads and writes to locations on the memory bank; and a plurality of buses connecting each controller of the plurality of bank controllers to at least one other controller of the plurality of bank controllers.

5 [025] One aspect of the present disclosure is directed to a memory device including a substrate; a plurality of memory banks on the substrate; a plurality of primary logic blocks on the substrate, each of the plurality of primary logic blocks being connected to at least one of the plurality of memory banks; a plurality of redundant blocks on the substrate, each of the plurality of redundant blocks being connected to at least one of the memory banks, each of the plurality of redundant blocks replicating at least one of the plurality of primary logic blocks; and a plurality of configuration switches on the substrate, each one of the plurality of the configuration switches being connected to at least one of the plurality of primary logic blocks or to at least one of the plurality of redundant blocks. In the memory device, upon detection of a fault associated with one of the plurality of primary logic blocks: a first configuration switch of the plurality of configuration switches may be configured to disable the one of the plurality of primary logic blocks, and a second configuration switch of the plurality of configuration switches may be configured to enable one of the plurality of redundant blocks that replicates the one of the plurality of primary logic blocks.

[026] Another aspect of the present disclosure is directed to a distributed processor on a memory chip including a substrate; an address manager on the substrate; a plurality of primary logic blocks on the substrate, each of the plurality of primary logic blocks being connected to at least one of the plurality of memory banks; a plurality of redundant blocks on the substrate, each of the plurality of redundant blocks being connected to at least one of the plurality of memory banks, each of the plurality of redundant blocks replicating at least one of the plurality of primary logic blocks; and a bus on the substrate connected to each of the plurality of primary logic blocks, each of the plurality of redundant blocks, and the address manager. In the processor may assign running ID numbers to blocks in the plurality of primary logic blocks that pass a testing protocol; assign illegal ID numbers to blocks in the plurality of primary logic blocks that do not pass the testing protocol; and assign running ID numbers to blocks in the plurality of redundant blocks that pass the testing protocol.

[027] Yet another aspect of the present disclosure is directed to a method for configuring a distributed processor on a memory chip. The method may include: testing each one of a plurality of primary logic blocks on the substrate of the memory chip for at least one circuit functionality; identifying at least one faulty logic block in the plurality of primary logic blocks based on the testing results, the at least one faulty logic block being connected to at least one memory bank disposed on the substrate of the memory chip; testing at least one redundant block on the substrate of the memory chip for the at least one circuit functionality, the at least one redundant block replicating the at least one faulty logic block and being connected to the at least one memory bank; disabling the at least one faulty logic block by applying an external signal to a deactivation switch, the deactivation switch being connected with the at least one faulty logic block and being disposed on the substrate of the memory chip; and enabling the at least one



redundant block by applying the external signal to an activation switch, the activation switch being connected with the at least one redundant block and being disposed on the substrate of the memory chip.

[028] Another aspect of the present disclosure is directed to a method for configuring a distributed processor on a memory chip. The method may include enabling a plurality of primary logic blocks and a plurality of redundant blocks on the substrate of the memory; testing each one of the plurality of primary logic blocks on the substrate of the memory chip for at least one circuit functionality; identifying at least one faulty logic block in the plurality of primary logic blocks based on the testing results, the at least one faulty logic block being connected to at least one memory bank disposed on the substrate of the memory chip; testing at least one redundant block on the substrate of the memory chip for the at least one circuit functionality, the at least one redundant block replicating the at least one faulty logic block and being connected to the at least one memory bank; and disabling at least one redundant block by applying the external signal to an activation switch, the activation switch being connected with the at least one redundant block and being disposed on the substrate of the memory chip.

[029] One aspect of the present disclosure is directed to a processing device. The processing device may include a substrate; a plurality of memory banks on the substrate; a memory controller on the substrate connected to each one of the plurality of memory banks; and a plurality of processing units on the substrate, each one of the plurality of processing units being connected to the memory controller, the plurality of processing units including a configuration manager. In the processing device, the configuration manager is configured to receive a first indication of a task to be performed, the task requiring at least one computation; signal at least one selected processing unit from the plurality of processing units based upon a capability of the selected processing unit for performing the at least one computation; and transmitting a second indication to the at least one selected processing unit, and the memory controller is configured to route data from at least two memory banks to the at least one selected processing unit using at least one communication line, the at least one communication line being connected to the at least two memory banks and the at least one selected processing unit via the memory controller.

[030] Another aspect of the present disclosure is directed to a method performed for operating a distributed memory device. The method may include: compiling, by a compiler, a task for the distributed memory device, the task requiring at least one computation, the compiling may include determining a number of words that are required simultaneously to perform the task, and providing instructions for writing words that need to be accessed simultaneously in a plurality of memory banks disposed on the substrate when a number a number of words that can be accessed simultaneously from one of the plurality of memory banks is lower than the number of words that are required simultaneously; receiving, by a configuration manager disposed on the substrate, an indication to perform the task; and in response to receiving the indication, configuring a memory controller disposed in the substrate to: within a first line access cycle: access at least one first word from a first memory bank from the plurality of memory banks using a first memory line, send the at least one first word to at least one processing unit, and open a first memory line in the second memory bank to access a second address from the second memory bank from

the plurality of memory banks, and within a second line access cycle: access at least one second word from the second memory bank using the first memory line, send the at least one second word to at least one processing unit, and access a third address from the first memory bank using a second memory line in the first bank.

5 [031] Yet another aspect of the present disclosure is directed to a non-transitory computer-readable medium that stores instructions that, when executed by at least one processor, cause the at least one processor to determine a number of words that are required simultaneously to perform a task, the task requiring at least one computation; write words that need to be accessed simultaneously in a plurality of memory banks disposed on the substrate when a number a number of words that can be accessed  
10 simultaneously from one of the plurality of memory banks is lower than the number of words that are required simultaneously; transmit an indication to perform the task to a configuration manager disposed on the substrate; and transmit instructions to configure a memory controller disposed on the substrate to, within a first line access cycle: access at least one first word from a first memory bank from the plurality of memory banks using a first memory line, send the at least one first word to at least one processing unit,  
15 and open a first memory line in the second memory bank to access a second address from the second memory bank from the plurality of memory banks, and within a second line access cycle: access at least one second word from the second memory bank using the first memory line, send the at least one second word to at least one processing unit, and access a third address from the first memory bank using a second memory line in the first bank.

20 [032] Consistent with other disclosed embodiments, non-transitory computer-readable storage media may store program instructions, which are executed by at least one processing device and perform any of the methods described herein.

[033] The foregoing general description and the following detailed description are exemplary and explanatory only and are not restrictive of the claims.

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

[034] The accompanying drawings, which are incorporated in and constitute a part of this disclosure, illustrate various disclosed embodiments. In the drawings:

[035] FIG. 1 is a diagrammatic representation of a central processing unit (CPU).

[036] FIG. 2 is a diagrammatic representation of a graphics processing unit (GPU).

30 [037] FIG. 3A is a diagrammatic representation of an embodiment of an exemplary hardware chip consistent with the disclosed embodiments.

[038] FIG. 3B is a diagrammatic representation of another embodiment of an exemplary hardware chip consistent with the disclosed embodiments.

35 [039] FIG. 4 is a diagrammatic representation of a generic command executed by an exemplary hardware chip consistent with the disclosed embodiments.

[040] FIG. 5 is a diagrammatic representation of a specialized command executed by an exemplary hardware chip consistent with the disclosed embodiments.

[041] FIG. 6 is a diagrammatic representation of a processing group for use in an exemplary hardware chip consistent with the disclosed embodiments.

[042] FIG. 7A is a diagrammatic representation of a rectangular array of processing groups consistent with the disclosed embodiments.

5 [043] FIG. 7B is a diagrammatic representation of an elliptical array of processing groups consistent with the disclosed embodiments.

[044] FIG. 7C is a diagrammatic representation an array of hardware chips consistent with the disclosed embodiments.

10 [045] FIG. 7D is a diagrammatic representation another array of hardware chips consistent with the disclosed embodiments.

[046] FIG. 8 is a flowchart depicting an exemplary method for compiling a series of instructions for execution on an exemplary hardware chip consistent with the disclosed embodiments.

[047] FIG. 9 is a diagrammatic representation of a memory bank.

[048] FIG. 10 is a diagrammatic representation of a memory bank.

15 [049] FIG. 11 is a diagrammatic representation of an embodiment of an exemplary memory bank with sub-bank controls consistent with the disclosed embodiments.

[050] FIG. 12 is a diagrammatic representation of another embodiment of an exemplary memory bank with sub-bank controls consistent with the disclosed embodiments.

20 [051] FIG. 13 is a block diagram of an exemplary memory chip, consistent with disclosed embodiments.

[052] FIG. 14 is a block diagram of an exemplary redundant logic block set, consistent with disclosed embodiments.

[053] FIG. 15 is a block diagram for an exemplary logic block, consistent with disclosed embodiments.

25 [054] FIG. 16 are block diagrams of exemplary logic blocks connected with a bus, consistent with disclosed embodiments.

[055] FIG. 17 is a block diagram for exemplary logic blocks connected in series, consistent with disclosed embodiments.

30 [056] FIG. 18 is a block diagram of exemplary logic blocks connected in a two-dimension array, consistent with disclosed embodiments.

[057] FIG. 19 is a block diagram for exemplary logic blocks in a complex connection, consistent with disclosed embodiments.

[058] FIG. 20 is an exemplary flow chart illustrating a redundant block enabling process, consistent with disclosed embodiments.

35 [059] FIG. 21 is an exemplary flow chart illustrating an address assignment process, consistent with disclosed embodiments.

[060] FIG. 22 provides block diagrams for exemplary processing devices, consistent with disclosed embodiments.

[061] FIG. 23 is a block diagram of an exemplary processing device, consistent with disclosed embodiments.

[062] FIG. 24 includes exemplary memory configuration diagrams, consistent with disclosed embodiments.

5 [063] FIG. 25 is an exemplary flowchart illustrating a memory configuration process, consistent with disclosed embodiments.

[064] FIG. 26 is an exemplary flowchart illustrating a memory read process, consistent with disclosed embodiments.

10 [065] FIG. 27 is an exemplary flowchart illustrating a process execution, consistent with disclosed embodiments.

### DETAILED DESCRIPTION

[066] The following detailed description refers to the accompanying drawings. Wherever convenient, the same reference numbers are used in the drawings and the following description to refer to the same or similar parts. While several illustrative embodiments are described herein, modifications, adaptations and other implementations are possible. For example, substitutions, additions or  
15 modifications may be made to the components illustrated in the drawings, and the illustrative methods described herein may be modified by substituting, reordering, removing, or adding steps to the disclosed methods. Accordingly, the following detailed description is not limited to the disclosed embodiments and examples. Instead, the proper scope is defined by the appended claims.

#### 20 Processor Architecture

[067] As used throughout this disclosure, the term “hardware chip” refers to a semiconductor wafer (such as silicon or the like) on which one or more circuit elements (such as transistors, capacitors, resistors, and/or the like) are formed. The circuit elements may form processing elements or memory elements. A “processing element” refers to one or more circuit elements that, together, perform at least  
25 one logic function (such as an arithmetic function, a logic gate, other Boolean operations, or the like). A processing element may be a general-purpose processing element (such as a configurable plurality of transistors) or a special-purpose processing element (such as a particular logic gate or a plurality of circuit elements designed to perform a particular logic function). A “memory element” refers to one or more circuit elements that can be used to store data. A “memory element” may also be referred to as a  
30 “memory cell.” A memory element may be dynamic (such that electrical refreshes are required to maintain the data store), static (such that data persists for at least some time after power loss), or non-volatile memories.

[068] Processing elements may be joined to form processor subunits. A “processor subunit” may thus comprise a smallest grouping of processing elements that may execute at least one task or  
35 instructions (e.g., of a processor instruction set). For example, a subunit may comprise one or more general-purpose processing elements configured to execute instructions together, one or more general-purpose processing elements paired with one or more special-purpose processing elements configured to execute instructions in a complementary fashion, or the like. The processor subunits may be arranged on a

substrate (e.g., a wafer) in an array. Although the “array” may comprise a rectangular shape, any arrangement of the subunits in the array may be formed on the substrate.

[069] Memory elements may be joined to form memory banks. For example, a memory bank may comprise one or more lines of memory elements linked along at least one wire (or other conductive connection). Furthermore, the memory elements may be linked along at least one addition wire in another direction. For example, the memory elements may be arranged along wordlines and bitlines, as explained below. Although the memory bank may comprise lines, any arrangement of the elements in the bank may be used to form the bank on the substrate. Moreover, one or more banks may be electrically joined to at least one memory controller to form a memory array. Although the memory array may comprise a rectangular arrangement of the banks, any arrangement of the banks in the array may be formed on the substrate.

[070] As further used throughout this disclose, a “bus” refers to any communicative connection between elements of a substrate. For example, a wire or a line (forming an electrical connection), an optical fiber (forming an optical connection), or any other connection conducting communications between components may be referred to as a “bus.”

[071] Conventional processors pair general-purpose logic circuits with shared memories. The shared memories may store both instruction sets for execution by the logic circuits as well as data used for and resulting from execution of the instruction sets. As described below, some conventional processors use a caching system to reduce delays in performing pulls from the shared memory; however, conventional caching systems remain shared. Conventional processors include central processing units (CPUs), graphics processing units (GPUs), various application-specific integrated circuits (ASICs), or the like. FIG. 1 shows an example of a CPU, and FIG. 2 shows an example of a GPU.

[072] As shown in FIG. 1, a CPU 100 may comprise a processing unit 110 that includes one or more processor subunits, such as processor subunit 120a and processor subunit 120b. Although not depicted in FIG. 1, each processor subunit may comprise a plurality of processing elements. Moreover, the processing unit 110 may include one or more levels of on-chip cache. Such cache elements are generally formed on the same semiconductor die as processing unit 110 rather than being connected to processor subunits 120a and 120b via one or more buses formed in the substrate containing processor subunits 120a and 120b and the cache elements. An arrangement directly on the same die, rather than being connected via buses, is common for both first-level (L1) and second-level (L2) caches in conventional processors. Alternatively, in older processors, L2 caches were shared amongst processor subunits using back-side buses between the subunits and the L2 caches. Back-side buses are generally larger than front-side buses, described below. Accordingly, because cache is to be shared with all processor subunits on the die, cache 130 may be formed on the same die as processor subunits 120a and 120b or communicatively coupled to processor subunits 120a and 120b via one or more back-side buses. In both embodiments without buses (e.g., cache is formed directly on-die) as well as embodiments using back-side buses, the caches are shared between processor subunits of the CPU.

[073] Moreover, processing unit 110 communicates with shared memory 140a and memory 140b. For example, memories 140a and 140b may represent memory banks of shared dynamic random access memory (DRAM). Although depicted with two banks, most conventional memory chips include between eight and sixteen memory banks. Accordingly, processor subunits 120a and 120b may use shared memories 140a and 140b to store data that is then operated upon by processor subunits 120a and 120b. This arrangement, however, results in the buses between memories 140a and 140b and processing unit 110 acting as a bottleneck when the clock speeds of processing unit 110 exceed data transfer speeds of the buses. This is generally true for conventional processors, resulting in lower effective processing speeds than the stated processing speeds based on clock rate and number of transistors.

[074] As shown in FIG. 2, similar deficiencies persist in GPUs. A GPU 200 may comprise a processing unit 210 that includes one or more processor subunits (e.g., subunits 220a, 220b, 220c, 220d, 220e, 220f, 220g, 220h, 220i, 220j, 220k, 220l, 220m, 220n, 220o, and 220p). Moreover, the processing unit 210 may include one or more levels of on-chip cache and/or register files. Such cache elements are generally formed on the same semiconductor die as processing unit 210. Indeed, in the example of FIG. 2, cache 210 is formed on the same die as processing unit 210 and shared amongst all of the processor subunits, while caches 230a, 230b, 230c, and 230d are formed on a subset of the processor subunits, respectively, and dedicated thereto.

[075] Moreover, processing unit 210 communicates with shared memories 250a, 250b, 250c, and 250d. For example, memories 250a, 250b, 250c, and 250d may represent memory banks of shared DRAM. Accordingly, the processor subunits of processing unit 210 may use shared memories 250a, 250b, 250c, and 250d to store data that is then operated upon by the processor subunits. This arrangement, however, results in the buses between memories 250a, 250b, 250c, and 250d and processing unit 210 acting as a bottleneck, similar to the bottleneck described above for CPUs.

#### Overview of Disclosed Hardware Chips

[076] FIG. 3A is a diagrammatic representation of an embodiment depicting an exemplary hardware chip 300. Hardware chip 300 may comprise a distributed processor designed to mitigate the bottlenecks described above for CPUs, GPUs, and other conventional processors. A distributed processor may include a plurality of processor subunits distributed spatially on a single substrate. Moreover, as explained above, in distributed processors of the present disclosure, corresponding memory banks are also spatially distributed on the substrate. In some embodiments, a distributed processor may be associated with a set of instructions, and each one of the processor subunits of the distributed processor may be responsible for performing one or more tasks included in the set of instructions.

[077] As depicted in FIG. 3A, hardware chip 300 may comprise a plurality of processor subunits, e.g., logic and control subunits 320a, 320b, 320c, 320d, 320e, 320f, 320g, and 320h. As further depicted in FIG. 3A, each processor subunit may have a dedicated memory instance. For example, logic and control subunit 320a is operably connected to dedicated memory instance 330a, logic and control subunit 320b is operably connected to dedicated memory instance 330b, logic and control subunit 320c is operably connected to dedicated memory instance 330c, logic and control subunit 320d is operably

connected to dedicated memory instance 330d, logic and control subunit 320e is operably connected to dedicated memory instance 330e, logic and control subunit 320f is operably connected to dedicated memory instance 330f, logic and control subunit 320g is operably connected to dedicated memory instance 330g, and logic and control subunit 320h is operably connected to dedicated memory instance 330h.

[078] Although FIG. 3A depicts each memory instance as a single memory bank, hardware chip 300 may include two or more memory banks as a dedicated memory instance for a processor subunit on hardware chip 300. Furthermore, although FIG. 3A depicts each processor subunit as comprising both a logic component and a control for the dedicated memory bank(s), hardware chip 300 may use controls for the memory banks that are separate, at least in part, from the logic components. Moreover, as depicted in FIG. 3A, two or more processor subunits and their corresponding memory banks may be grouped, e.g., into processing groups 310a, 310b, 310c, and 310d. A “processing group” may represent a spatial distinction on a substrate on which hardware chip 300 is formed. Accordingly, a processing group may include further controls for the memory banks in the group, e.g., controls 340a, 340b, 340c, and 340d. Additionally or alternatively, a “processing group” may represent a logical grouping for the purposes of compiling code for execution on hardware chip 300. Accordingly, a compiler for hardware chip 300 (further described below) may divide an overall set of instructions between the processing groups on hardware chip 300.

[079] Furthermore, host 350 may provide instructions, data, and other input to hardware chip 300 and read output from the same. Accordingly, a set of instructions may be executed entirely on a single die, e.g., the die hosting hardware chip 300. Indeed, the only communications off-die may include the loading of instructions to hardware chip 300, any input sent to hardware chip 300, and any output read from hardware chip 300. Accordingly, all calculations and memory operations may be performed on-die (on hardware chip 300) because the processor subunits of hardware chip 300 communicate with dedicated memory banks of hardware chip 300.

[080] FIG. 3B is a diagrammatic representation of an embodiment depicting another exemplary hardware chip 300'. Although depicted as an alternative to hardware chip 300, the architecture depicted in FIG. 3B may be combined, at least in part, with the architecture depicted in FIG. 3A.

[081] As depicted in FIG. 3B, hardware chip 300' may comprise a plurality of processor subunits, e.g., processor subunits 350a, 350b, 350c, and 350d. As further depicted in FIG. 3B, each processor subunit may have a plurality of dedicated memory instances. For example, processor subunit 350a is operably connected to dedicated memory instances 330a and 330b, processor subunit 350b is operably connected to dedicated memory instances 330c and 330d, processor subunit 350c is operably connected to dedicated memory instances 330e and 330f, and processor subunit 350d is operably connected to dedicated memory instances 330g and 330h. Moreover, as depicted in FIG. 3B, the processor subunits and their corresponding memory banks may be grouped, e.g., into processing groups 310a, 310b, 310c, and 310d. As explained above, a “processing group” may represent a spatial distinction

on a substrate on which hardware chip 300' is formed and/or a logical grouping for the purposes of compiling code for execution on hardware chip 300'.

[082] As further depicted in FIG. 3B, the processor subunits may communicate with each other via buses. For example, as shown in FIG. 3B, processor subunit 350a may communicate with processor subunit 350b via bus 360a, with processor subunit 350c via bus 360c, and with processor subunit 350d via bus 360f. Similarly, processor subunit 350b may communicate with processor subunit 350a via bus 360a (as described above), with processor subunit 350c via bus 360e, and with processor subunit 350d via bus 360d. In addition, processor subunit 350c may communicate with processor subunit 350a via bus 360c (as described above), with processor subunit 350b via bus 360e (as described above), and with processor subunit 350d via bus 360b. Accordingly, processor subunit 350d may communicate with processor subunit 350a via bus 360f (as described above), with processor subunit 350b via bus 360d (as described above), and with processor subunit 350c via bus 360b (as described above). One of ordinary skill will understand that fewer buses than depicted in FIG. 3B may be used. For example, bus 360e may be eliminated such that communications between processor subunit 350b and 350c pass through processor subunit 350a and/or 350d. Similarly, bus 360f may be eliminated such that communications between processor subunit 350a and processor subunit 350d pass through processor subunit 350b or 350c.

[083] Moreover, one of ordinary skill will understand that architectures other than those depicted in FIGS. 3A and 3B may be used. For example, an array of processing groups, each with a single processor subunit and memory instance, may be arranged on a substrate. Processor subunits may additionally or alternatively form part of controllers for corresponding dedicated memory banks, part of controllers for memory mats of corresponding dedicated memory, or the like.

[084] In view of the architecture described above, hardware chips 300 and 300' may provide significant increases in efficiency for memory-intensive tasks as compared with traditional architectures. For example, database operations and artificial intelligence algorithms (such as neural networks) are examples of memory-intensive tasks for which traditional architectures are less efficient than hardware chips 300 and 300'. Accordingly, hardware chips 300 and 300' may be referred to as database accelerator processors and/or artificial intelligence accelerator processors.

#### Configuring the Disclosed Hardware Chips

[085] The hardware chip architecture described above may be configured for execution of code. For example, each processor subunit may individually execute code (defining a set of instructions) apart from other processor subunits in the hardware chip. Accordingly, rather than relying on an operating system to manage multithreading or using multitasking (which is concurrency rather than parallelism), hardware chips of the present disclosure may allow for processor subunits to operate fully in parallel.

[086] In addition to a fully parallel implementation described above, at least some of the instructions assigned to each processor subunit may be overlapping. For example, a plurality of processor subunits on a distributed processor may execute overlapping instructions as, for example, an implementation of an operating system or other management software, while executing non-overlapping



instructions in order to perform parallel tasks within the context of the operating system or other management software.

[087] FIG. 4 depicts an exemplary process 400 for executing a generic command with processing group 410. For example, processing group 410 may comprise a portion of a hardware chip of the present disclosure, e.g., hardware chip 300, hardware chip 300', or the like.

[088] As depicted in FIG. 4, a command may be sent to processor subunit 430, which is paired with dedicated memory instance 420. An external host (e.g., host 350) may send the command to processing group 410 for execution. Alternatively, host 350 may have sent an instruction set including the command for storage in memory instance 420 such that processor subunit 430 may retrieve the command from memory instance 420 and execute the retrieved command. Accordingly, the command may be executed by processing element 440, which is a generic processing element configurable to execute the received command. Moreover, processing group 410 may include a control 460 for memory instance 420. As depicted in FIG. 4, control 460 may perform any reads and/or writes to memory instance 420 required by processing element 440 when executing the received command. After execution of the command, processing group 410 may output the result of the command, e.g., to the external host or to a different processing group on the same hardware chip.

[089] In some embodiments, as depicted in FIG. 4, processor subunit 430 may further include an address generator 450. An "address generator" may comprise a plurality of processing elements that are configured to determine addresses in one or more memory banks for performing reads and writes and may also perform operations on the data located at the determined addresses (e.g., addition, subtraction, multiplication, or the like). For example, address generator 450 may determine addresses for any reads or writes to memory. In one example, address generator 450 may increase efficiency by overwriting a read value with a new value determined based on the command when the read value is no longer needed. Additionally or alternatively, address generator 450 may select available addresses for storage of results from execution of the command. This may allow for scheduling of result read-off for a later clock cycle, when it is more convenient for the external host. In another example, address generator 450 may determine addresses to read from and write to during a multi-cycle calculation, such as a vector or matrix multiply-accumulate calculation. Accordingly, address generator 450 may maintain or calculate memory addresses for reading data and writing intermediate results of the multi-cycle calculation such that processor subunit 430 may continue processing without having to store these memory addresses.

[090] FIG. 5 depicts an exemplary process 500 for executing a specialized command with processing group 510. For example, processing group 510 may comprise a portion of a hardware chip of the present disclosure, e.g., hardware chip 300, hardware chip 300', or the like.

[091] As depicted in FIG. 5, a specialized command (e.g., a multiply-accumulate command) may be sent to processing element 530, which is paired with dedicated memory instance 520. An external host (e.g., host 350) may send the command to processing element 530 for execution. Accordingly, the command may be executed at a given signal from the host by processing element 530, a specialized processing element configurable to execute particular commands (including the received command).

Alternatively, processing element 530 may retrieve the command from memory instance 520 for execution. Thus, in the example of FIG. 5, processing element 530 is a multiply-accumulate (MAC) circuit configured to execute MAC commands received from the external host or retrieved from memory instance 520. After execution of the command, processing group 410 may output the result of the command, e.g., to the external host or to a different processing group on the same hardware chip. Although depicted with a single command and a single result, a plurality of commands may be received or retrieved and executed, and a plurality of results may be combined on processing group 510 before output.

[092] Although depicted as a MAC circuit in FIG. 5, additional or alternative specialized circuits may be included in processing group 510. For example, a MAX-read command (which returns the max value of a vector) a MAX0-read command (a common function also termed a rectifier, which returns the entire vector but also does MAX with 0), or the like may be implemented.

[093] Although depicted separately, the generalized processing group 410 of FIG. 4 and the specialized processing group 510 of FIG. 5 may be combined. For example, a generic processor subunit may be coupled to one or more specialized processor subunits to form a processor subunit. Accordingly, the generic processor subunit may be used for all instructions not executable by the one or more specialized processor subunits.

[094] One of ordinary skill will understand that neural network implementation and other memory-intensive tasks may be handled with specialized logic circuits. For example, database queries, packet inspection, string comparison, and other functions may increase in efficiency if executed by the hardware chips described herein.

#### A Memory-Based Architecture for Distributed Processing

[095] On hardware chips consistent with the present disclosure, dedicated buses may transfer data between processor subunits on the chip and/or between the processor subunits and their corresponding dedicated memory banks. The use of dedicated buses may reduce arbitration costs because competing requests are either not possible or easily avoided using software rather than hardware.

[096] FIG. 6 schematically depicts a diagrammatic representation of a processing group 600. Processing group 600 may be for use in a hardware chip, e.g., hardware chip 300, hardware chip 300', or the like. Processor subunit 610 may be connected via buses 630 to memory 620. Memory 620 may comprise a Randomly Accessible Memory (RAM) element that stores data and code for execution by processor subunit 610. In some embodiments, memory 620 may be an N-Way memory (wherein N is a number equal to or larger than 1 that implies the number of segments in an interleaved memory 620). Because processor subunit 610 is coupled to memory 620 dedicated to processor subunit 610 via bus 630, N may be kept relatively small without compromising the execution performance. This represents an improvement over conventional multiway register files or caches where a lower N generally results in lower execution performance, and a higher N generally results in large area and power loss.

[097] The size of memory 620, the number of ways, and the width of bus 630 may be adjusted to meet the requirements of tasks and application implementations of a system using processing group

600 according to, for instance, the size of data involved in the task or tasks. Memory element 620 may comprise one or more types of memory known in the art, e.g., volatile memory (such as RAM, DRAM, SRAM, phase-change RAM (PRAM), magnetoresistive RAM (MRAM), resistive RAM (ReRAM), or the like) or non-volatile memory (such as flash or ROM). According to some embodiments, a portion of  
5 memory element 620 may comprise a first memory type, while another portion may comprise another memory type. For instance, the code region of a memory element 620 may comprise a ROM element, while a data region of the memory element 620 may comprise a DRAM element. Another example for such partitioning is storing the weights of a neural network in flash while storing the data for calculation in DRAM.

10 [098] Processor subunit 610 comprises a processing element 640 that may comprise a processor. The processor can be pipelined or not pipelined, a customized Reduced Instruction Set Computing (RISC) element or other processing scheme, implemented on any commercial Integrated Circuit (IC) known in the art (such as ARM, ARC, RISC-V, etc.), as appreciated by one of ordinary skill. Processing element 640 may comprise a controller that, in some embodiments, includes an Arithmetic  
15 Logic Unit (ALU) or other controller.

[099] According to some embodiments, processing element 640, which executes received or stored code, may comprise a generic processing element and, therefore, be flexible and capable of performing a wide variety of processing operations. Non-dedicated circuitry typically consumes more power than specific-operation-dedicated circuitry when comparing the power consumed during  
20 performance for a specific operation. Therefore, when performing specific complex arithmetic calculations, processing element 640 may consume more power and perform less efficiently than dedicated hardware. Therefore, according to some embodiments, a controller of processing element 640 may be designed to perform specific operations (e.g., addition or “move” operations).

[0100] In one example, the specific operations may be performed by one or more accelerators  
25 650. Each accelerator may be dedicated and programmed to perform a specific calculation (such as multiplication, floating point vector operations, or the like). By using accelerator(s), the average power consumed per calculation per processor subunit may be lowered, and the calculation throughput henceforth increases. Accelerator(s) 650 may be chosen according to an application that the system is designed to implement (e.g., execution of neural networks, execution of database queries, or the like).  
30 Accelerator(s) 650 may be configured by processing element 640 and may operate in tandem therewith for lowering power consumption and accelerating calculations and computations. The accelerators may additionally or alternatively be used to transfer data between memory and MUXs/DEMUXs/input/output ports (e.g., MUX 650 and DEMUX 660) of processing group 600, such as a smart DMA (direct memory access) peripheral.

35 [0101] Accelerator(s) 650 may be configured to perform a variety of functions. For instance, one accelerator may be configured to perform 16-bit floating point calculation or 8-bit integer calculations, which are often used in neural networks. Another example of an accelerator function is a 32-bit floating point calculation, which is often used during a training stage of a neural network. Yet another example of

an accelerator function is query processing, such as that used in databases. In some embodiments, accelerator(s) 650 may comprise specialized processing elements to perform these functions and/or may be configured according to configuration data, stored on the memory element 620, such that it may be modified.

5 [0102] Accelerator(s) 650 may additionally or alternatively implement a configurable scripted list of memory movements to time movements of data to/from memory 620 or to/from other accelerators and/or inputs/outputs. Accordingly, as explained further below, all the data movement inside the hardware chip using processing group 600 may use software synchronization rather than hardware synchronization. For example, an accelerator in one processing group (e.g., group 600) may transfer data  
10 from its input to its accelerator every tenth cycle and then output data at the next cycle, thereby letting the information flow from the memory of the processing group to another one.

[0103] As further depicted in FIG. 6, in some embodiments, processing group 600 may further comprise at least one input multiplexer (MUX) 660 connected to its input port and at least one output DEMUX 670 connected to its output port. These MUXs/DEMUXs may be controlled by control signals  
15 (not shown) from processing element 640 and/or from one of accelerator(s) 650, determined according to a current instruction being carried out by processing element 640 and/or the operation executed by an accelerator of accelerator(s) 650. In some scenarios, processing group 600 may be required (according to a predefined instruction from its code memory) to transfer data from its input port to its output port. Accordingly, one or more of the input MUXs (e.g., MUX 660) may be directly connected via one or more  
20 buses to an output DEMUX (e.g., DEMUX 670), in addition to each of the DEMUXs/MUXs being connected to processing element 640 and accelerator(s) 650.

[0104] The processing group 600 of FIG. 6 may be arrayed to form a distributed processor, for example, as depicted in FIG. 7A. The processing groups may be disposed on substrate 710 to form an array. In some embodiments, substrate 710 may comprise a semiconductor substrate, such as silicon.  
25 Additionally or alternatively, substrate 710 may comprise a circuit board, such as a flexible circuit board.

[0105] As depicted in FIG. 7A, substrate 710 may include, disposed thereon, a plurality of processing groups, such as processing group 600. Accordingly, substrate 710 includes a memory array that includes a plurality of banks, such as banks 720a, 720b, 720c, 720d, 720e, 720f, 720g, and 720h. Furthermore, substrate 710 includes a processing array that may include a plurality of processor subunits,  
30 such as subunits 730a, 730b, 730c, 730d, 730e, 730f, 730g, and 730h.

[0106] Furthermore, as explained above, each processing group may include a processor subunit and one or more corresponding memory banks dedicated to the processor subunit. Accordingly, as depicted in FIG. 7A, each subunit is associated with a corresponding, dedicated memory bank, e.g.: Processor subunit 730a is associated with memory bank 720a, processor subunit 730b is associated with  
35 memory bank 720b, processor subunit 730c is associated with memory bank 720c, processor subunit 730d is associated with memory bank 720d, processor subunit 730e is associated with memory bank 720e, processor subunit 730f is associated with memory bank 720f, processor subunit 730g is associated with memory bank 720g, processor subunit 730h is associated with memory bank 720h.

[0107] To allow each processor subunit to communicate with its corresponding, dedicated memory bank(s), substrate 710 may include a first plurality of buses connecting one of the processor subunits to its corresponding, dedicated memory bank(s). Accordingly, bus 740a connects processor subunit 730a to memory bank 720a, bus 740b connects processor subunit 730b to memory bank 720b, bus 740c connects processor subunit 730c to memory bank 720c, bus 740d connects processor subunit 730d to memory bank 720d, bus 740e connects processor subunit 730e to memory bank 720e, bus 740f connects processor subunit 730f to memory bank 720f, bus 740g connects processor subunit 730g to memory bank 720g, and bus 740h connects processor subunit 730h to memory bank 720h. Moreover, to allow each processor subunit to communicate with other processor subunits, substrate 710 may include a second plurality of buses connecting one of the processor subunits to another of the processor subunits. In the example of FIG. 7A, bus 750a connects processor subunit 730a to processor subunit 750e, bus 750b connects processor subunit 730a to processor subunit 750b, bus 750c connects processor subunit 730b to processor subunit 750f, bus 750d connects processor subunit 730b to processor subunit 750c, bus 750e connects processor subunit 730c to processor subunit 750g, bus 750f connects processor subunit 730c to processor subunit 750d, bus 750g connects processor subunit 730d to processor subunit 750h, bus 750h connects processor subunit 730h to processor subunit 750g, bus 750i connects processor subunit 730g to processor subunit 750g, and bus 750j connects processor subunit 730f to processor subunit 750e.

[0108] Accordingly, in the example arrangement shown in FIG. 7A, the plurality of logic processor subunits is arranged in at least one row and at least one column. The second plurality of buses connect each processor subunit to at least one adjacent processor subunit in the same row and to at least one adjacent processor subunit in the same column. FIG. 7A may be referred to as a “partial tile connection.”

[0109] The arrangement shown in FIG. 7A may be modified to form a “full tile connection.” A full tile connection includes additional buses connecting diagonal processor subunits. For example, the second plurality of buses may include additional buses between processor subunit 730a and processor subunit 730f, between processor subunit 730b and processor subunit 730e, between processor subunit 730b and processor subunit 730g, between processor subunit 730c and processor subunit 730f, between processor subunit 730c and processor subunit 730h, and between processor subunit 730d and processor subunit 730g.

[0110] A full tile connection may be used for convolution calculations, in which data and results stored in a near processor subunit are used. For example, during convolutional image processing, each processor subunit may receive a tile of the image (such as a pixel or a group of pixels). In order to calculate the convolution results, each processor subunit may acquire data from all eight adjacent processor subunits, each of which have received a corresponding tile. In a partial tile connection, the data from the diagonal adjacents may be passed through other adjacent processor subunits connected to the processor subunit. Accordingly, the distributed processor on a chip may be an artificial intelligence accelerator processor.

[0111] In a specific example of a convolutional calculation, an  $N \times M$  image may be divided across a plurality of processor subunits. Each processor subunit may perform a convolution with an  $A \times B$  filter on its corresponding tile. To perform the filtering on one or more pixels on a boundary between tiles, each processor subunit may require data from neighboring processor subunits having tiles including pixels on the same boundary. Accordingly, the code generated for each processor subunit configures the subunit to calculate the convolutions and pull from one of the second plurality of buses whenever data is needed from an adjacent subunit. Corresponding commands to output data to the second plurality of buses are provided to the subunits to ensure proper timing of needed data transfers.

[0112] The partial tile connection of FIG. 7A may be modified to be an  $N$ -partial tile connection. In this modification, the second plurality of buses may further connect each processor subunit to processor subunits within a threshold distance of the processor subunit (e.g., within  $n$  processor subunits) in the four directions along which the buses of FIG. 7A run (i.e., up, down, left, and right). A similar modification may be made to the full-tile connection (to result in an  $N$ -full tile connection) such that the second plurality of buses further connects each processor subunit to processor subunits within a threshold distance of the processor subunit (e.g., within  $n$  processor subunits) in the four directions along which the buses of FIG. 7A run in addition to the two diagonal directions.

[0113] Other arrangements are possible. For example, in the arrangement shown in FIG. 7B, bus 750a connects processor subunit 730a to processor subunit 730d, bus 750b connects processor subunit 730a to processor subunit 730b, bus 750c connects processor subunit 730b to processor subunit 730c, and bus 750d connects processor subunit 730c to processor subunit 730d. Accordingly, in the example arrangement shown in FIG. 7B, the plurality of processor subunits is arranged in a star pattern. The second plurality of buses connect each processor subunit to at least one adjacent processor subunit within the star pattern.

[0114] Further arrangements (not shown) are possible. For example, a neighbor connection arrangement may be used such that the plurality of processor subunits is arranged in one or more lines (e.g., similar to that depicted in FIG. 7A). In a neighbor connection arrangement, the second plurality of buses connect each processor subunit to a processor subunit to the left in the same line, to a processor subunit to the right in the same line, to the processor subunits both to the left and to the right in the same line, etc.

[0115] In another example, an  $N$ -linear connection arrangement may be used. In an  $N$ -linear connection arrangement, the second plurality of buses connect each processor subunit to processor subunits within a threshold distance of the processor subunit (e.g., within  $n$  processor subunits). The  $N$ -linear connection arrangement may be used with the line array (described above), the rectangular array (depicted in FIG. 7A), the elliptical array (depicted in FIG. 7B), or any other geometrical array.

[0116] In yet another example, an  $N$ -log connection arrangement may be used. In an  $N$ -log connection arrangement, the second plurality of buses connect each processor subunit to processor subunits within a threshold power of two distance of the processor subunit (e.g., within  $2^n$  processor subunits). The  $N$ -log connection arrangement may be used with the line array (described above), the

rectangular array (depicted in FIG. 7A), the elliptical array (depicted in FIG. 7B), or any other geometrical array.

[0117] Any of the connection schemes described above may be combined for use in the same hardware chip. For example, a full tile connection may be used in one region while a partial tile  
5 connection is used in another region. In another example, an N-linear connection arrangement may be used in one region while an N-full tile connection is used in another region.

[0118] Alternatively to or in addition with dedicated buses between processor subunits of the memory chip, one or more shared buses may be used to interconnect all (or a subset of) the processor subunits of a distributed processor. Collisions on the shared buses may still be avoided by timing data  
10 transfers on the shared buses using code executed by the processor subunits, as explained further below. Additionally with or alternatively to shared buses, configurable buses may be used to dynamically connect processor subunits to form groups of processors units connected to separated buses. For example, the configurable buses may include transistors or other mechanisms that may be controlled by processor subunit to direct data transfers to a selected processor subunit.

[0119] In both FIGS. 7A and 7B, the plurality of processor subunits of the processing array is  
15 spatially distributed among the plurality of discrete memory banks of the memory array. In other alternative embodiments (not shown), the plurality of processor subunits may be clustered in one or more regions of the substrate, and the plurality of memory banks may be clustered in one or more other regions of the substrate. In some embodiments, a combination of spatial distribution and clustering may be used  
20 (not shown). For example, one region of the substrate may include a cluster of processor subunits, another region of the substrate may include a cluster of memory banks, and yet another region of the substrate may include processing arrays distributed amongst memory banks.

[0120] One of ordinary skill will recognize that arraying processor groups 600 on a substrate is not an exclusive embodiment. For example, each processor subunit may be associated with at least two  
25 dedicated memory banks. Accordingly, processing groups 310a, 310b, 310c, and 310d of FIG. 3B may be used in lieu of or in combination with processing group 600 to form the processing array and the memory array. Other processing groups including, for example, three, four, or more dedicated memory banks (not shown) may be used.

[0121] Each of the plurality of processor subunits may be configured to execute software code  
30 associated with a particular application independently, relative to other processor subunits included in the plurality of processor subunits. For example, as explained below, a plurality of sub-series of instructions may be grouped as machine code and provided to each processor subunit for execution.

[0122] In some embodiments, each dedicated memory bank comprises at least one dynamic random access memory (DRAM). Alternatively, the memory banks may comprise a mix of memory  
35 types, such as static random access memory (SRAM), DRAM, Flash or the like.

[0123] In conventional processors, data sharing between processor subunits is usually performed with shared memory. Shared memory typically requires a large portion of chip area and/or performed a bus that is managed by additional hardware (such as arbiters). The bus results in bottlenecks, as described

above. In addition, the shared memory, which may be external to the chip, typically includes cache coherency mechanisms and more complex caches (e.g., L1 cache, L2 cache, and shared DRAM) in order to provide accurate and up-to-date data to the processor subunits. As explained further below, the dedicated buses depicted in FIGS. 7A and 7B allow for hardware chips that are free of hardware management (such as arbiters). Moreover, the use of dedicated memories as depicted in FIGS. 7A and 7B allow for the elimination of complex caching layers and coherency mechanism.

[0124] Instead, in order to allow each processor subunit to access data calculated by other processor subunits and/or stored in memory banks dedicated to the other processor subunits, buses are provided whose timing is performed dynamically using code individually executed by each processor subunit. This allows for elimination of most, if not all, bus management hardware as conventionally used. Moreover, complex caching mechanisms are replaced with direct transfers over these buses, resulting in lower latency times during memory reads and writes.

#### Memory-Based Processing Arrays

[0125] As depicted in FIGS. 7A and 7B, a memory chip of the present disclosure may operate independently. Alternatively, memory chips of the present disclosure may be operably connected with one or more additional integrated circuits, such as a memory device (e.g., one or more DRAM banks), a system-on-a-chip, a field-programmable gate array (FPGA), or other processing and/or memory chip. In such embodiments, tasks in a series of instructions executed by the architecture may be divided (e.g., by a compiler, as described below) between processor subunits of the memory chip and any processor subunits of the additional integrated circuit(s). For example, the other integrated circuits may comprise a host (e.g., host 350 of FIG. 3A) that inputs instructions and/or data to the memory chip and receives output therefrom.

[0126] In order to interconnect memory chips of the present disclosure with one or more additional integrated circuits, the memory chip may include a memory interface, such as a memory interface complying with a Joint Electron Device Engineering Council (JEDEC) standard or any of its variants. The one or more additional integrated circuits may then connect to the memory interface. Accordingly, if the one or more additional integrated circuits are connected to a plurality of memory chips of the present disclosure, data may be shared between the memory chips through the one or more additional integrated circuits. Additionally or alternatively, the one or more additional integrated circuits may include buses to connect to buses on the memory chips of the present disclosure such that the one or more additional integrated circuits may execute code in tandem with the memory chips of the present disclosure. In such embodiments, the one or more additional integrated circuits further assist with distributed processing even though they may be on different substrates than the memory chips of the present disclosure.

[0127] Furthermore, memory chips of the present disclosure may be arrayed in order to form an array of distributed processors. For example, one or more buses may connect a memory chip 770a to an additional memory chip 770b, as depicted in FIG. 7C. In the example of FIG. 7C, memory chip 770a includes processor subunits with one or more corresponding memory banks dedicated to each processor



subunit, e.g.: Processor subunit 730a is associated with memory bank 720a, processor subunit 730b is associated with memory bank 720b, processor subunit 730e is associated with memory bank 720c, and processor subunit 730f is associated with memory bank 720d. Buses connect each processor subunit to its corresponding memory bank. Accordingly, bus 740a connects processor subunit 730a to memory bank 720a, bus 740b connects processor subunit 730b to memory bank 720b, bus 740c connects processor subunit 730e to memory bank 720c, and bus 740d connects processor subunit 730f to memory bank 720d. Moreover, bus 750a connects processor subunit 730a to processor subunit 750e, bus 750b connects processor subunit 730a to processor subunit 750b, bus 750c connects processor subunit 730b to processor subunit 750f, and bus 750d connects processor subunit 730e to processor subunit 750f. Other arrangements of memory chip 770a may be used, for example, as described above.

[0128] Similarly, memory chip 770b includes processor subunits with one or more corresponding memory banks dedicated to each processor subunit, e.g.: Processor subunit 730c is associated with memory bank 720e, processor subunit 730d is associated with memory bank 720f, processor subunit 730g is associated with memory bank 720g, and processor subunit 730h is associated with memory bank 720h. Buses connect each processor subunit to its corresponding memory bank. Accordingly, bus 740e connects processor subunit 730c to memory bank 720e, bus 740f connects processor subunit 730d to memory bank 720f, bus 740g connects processor subunit 730g to memory bank 720g, and bus 740h connects processor subunit 730h to memory bank 720h. Moreover, bus 750g connects processor subunit 730c to processor subunit 750g, bus 750h connects processor subunit 730d to processor subunit 750h, bus 750i connects processor subunit 730c to processor subunit 750d, and bus 750j connects processor subunit 730g to processor subunit 750h. Other arrangements of memory chip 770b may be used, for example, as described above.

[0129] The processor subunits of memory chip 770a and 770b may be connected using one or more buses. Accordingly, in the example of FIG. 7C, bus 750e may connect processor subunit 730b of memory chip 770a and processor subunit 730c of memory chip 770b, and bus 750f may connect processor subunit 730f of memory chip 770a and processor subunit 730c of memory 770b. For example, bus 750e may serve as an input bus to memory chip 770b (and thus an output bus for memory chip 770a) while bus 750f may serve as an input bus to memory chip 770a (and thus an output bus for memory chip 770b) or vice versa. Alternatively, buses 750e and 750f may both server as two-way buses between memory chips 770a and 770b.

[0130] Buses 750e and 750f may include direct wires or may be interleaved on a high-speed connection in order to reduce the pins used for the inter-chip interface between memory chip 770a and integrated circuit 770b. Moreover, any of the connection arrangements described above used in the memory chip itself may be used to connect the memory chip to one or more additional integrated circuits. For example, memory chip 770a and 770b may be connected using a full-tile or partial-tile connection rather than only two buses as shown in FIG. 7C.

[0131] Accordingly, although depicted using buses 750e and 750f, architecture 760 may include fewer buses or additional buses. For example, a single bus between processor subunits 730b and 730c or

between processor subunits 730f and 730c may be used. Alternatively, additional buses, e.g., between processor subunits 730b and 730d, between processor subunits 730f and 730d, or the like, may be used.

[0132] Furthermore, although depicted as using a single memory chip and an additional integrated circuit, a plurality of memory chips may be connected using buses as explained above. For example, as depicted in the example of FIG. 7C, memory chips 770a, 770b, 770c, and 770d are connected in an array. Each memory chip includes processor subunits and dedicated memory banks similar to the memory chips described above. Accordingly, a description of these components is not repeated here.

[0133] In the example of FIG. 7C, memory chips 770a, 770b, 770c, and 770d are connected in a loop. Accordingly, bus 750a connects memory chips 770a and 770d, bus 750c connects memory chips 770a and 770b, bus 750e connects memory chips 770b and 770c, and bus 750g connects memory chips 770c and 770d. Although memory chips 770a, 770b, 770c, and 770d may be connected with full-tile connections, partial-tile connections, or other connection arrangements, the example of FIG. 7C allows for fewer pin connections between memory chips 770a, 770b, 770c, and 770d.

#### Relatively Large Memories

[0134] Embodiments of the present disclosure may use dedicated memories of relatively large size as compared with shared memories of conventional processors. The use of dedicated memories rather than shared memories allows for gains in efficiency to continue without tapering off with memory increases. This allows for memory-intensive tasks such as neural network processing and database queries to be performed more efficiently than in conventional processors, where the efficiency gains of increasing shared memory taper off due to the von Neumann bottleneck.

[0135] For example, in distributed processors of the present disclosure, a memory array disposed on the substrate of the distributed processor may include a plurality of discrete memory banks. Each of the discrete memory banks may have a capacity greater than one megabyte, as well as a processing array disposed on the substrate, including a plurality of processor subunits. As explained above, each one of the processor subunits may be associated with a corresponding, dedicated one of the plurality of discrete memory banks. In some embodiments, the plurality of processor subunits may be spatially distributed among the plurality of discrete memory banks within the memory array. By using dedicated memories of at least one megabyte, rather than shared caches of a few megabytes for a large CPU or GPU, the distributed processors of the present disclosure gain efficiencies that are not possible in conventional systems due to the von Neumann bottleneck in CPUs and GPUs.

[0136] Different memories may be used as the dedicated memories. For example, each dedicated memory bank may comprise at least one DRAM bank. Alternatively, each dedicated memory bank may comprise at least one static random access memory bank. In other embodiments, different types of memories may be combined on a single hardware chip.

[0137] As explained above, each dedicated memory may be at least one megabyte. Accordingly, each dedicated memory bank may be the same size or at least two of the plurality of memory banks may have different sizes.

[0138] Moreover, as described above, the distributed processor may include a first plurality of buses, each connecting one of the plurality of processor subunits to a corresponding, dedicated memory bank and a second plurality of buses, each connecting one of the plurality of processor subunits to another one of the plurality of processor subunits.

5           Synchronization Using Software

[0139] As explained above, hardware chips of the present disclosure may manage data transfers using software rather than hardware. In particular, because the timings of transfers on the buses, reads and writes to the memories, and calculations of the processor subunits are set by the sub-series of instructions executed by the processor subunits, hardware chips of the present disclosure may execute code to prevent  
10 collisions on the buses. Accordingly, hardware chips of the present disclosure may avoid hardware mechanisms conventionally used to manage data transfers (such as network controllers within in a chip, packet parsers and packets transferors between processor subunits, bus arbitrators, a plurality of buses to avoid arbitration, or the like).

[0140] If hardware chips of the present disclosure transferred data conventionally, connecting N  
15 processor subunits with buses would require bus arbitration or wide MUXs controlled by an arbiter. Instead, as described above, embodiments of the present disclosure may use a bus that is only a wire, an optical cable, or the like between processor subunits, where the processor subunits individually execute code to avoid collision on the buses. Accordingly, embodiments of the present disclosure may preserve space on the substrate as well as materials cost and efficiency losses (e.g., due to power and time  
20 consumption by arbitration). The efficiency and space gains are even greater when compared to other architectures using first-in-first-out (FIFO) controllers and/or mailboxes.

[0141] Furthermore, as explained above, each processor subunit may include one or more accelerators in addition to one or more processing elements. In some embodiments, the accelerator(s) may read and write from the buses rather than the processing element(s). In such embodiments, additional  
25 efficiency may be obtained by allowing the accelerator(s) to transmit data during the same cycle in which the processing element(s) perform one or more calculations. Such embodiments, however, require additional materials for the accelerator(s). For example, additional transistors may be required for fabrication of the accelerator(s).

[0142] The code also may account for the internal behavior, including timing and latencies, of  
30 the processor subunits (e.g., including the processing elements and/or accelerators forming part of the processor subunit). For example, a compiler (as described below) may perform pre-processing that accounts for the timing and latencies when generating the sub-series of instructions that control the data transfers.

[0143] In one example, a plurality of processor subunits may be assigned a task of calculating a  
35 neural network layer containing a plurality of neurons fully-connected to a previous layer of a larger plurality of neurons. Assuming data of the previous layer is evenly spread between the plurality of processor subunits, one way to perform the calculation may be to configure each processor subunit to transmit the data of the previous layer to the main bus in turn and then each processor subunit will

multiply this data by the weight of the corresponding neuron that the subunit implements. Because each processor subunit calculates more than one neuron, each processor subunit will transmit the data of the previous layer a number of times equal to the number of neurons. Thus, the code of each processor subunit is not the same as the code for other processor subunits because the subunits will transmit at  
5 different times.

[0144] In some embodiments, a distributed processor may comprise a substrate (e.g., a semiconductor substrate, such as silicon and/or a circuit board, such as a flexible circuit board) with a memory array disposed on the substrate, the memory array including a plurality of discrete memory banks, and a processing array disposed on the substrate, the processing array including a plurality of  
10 processor subunits, as depicted, e.g., in FIGS. 7A and 7B. As explained above, each one of the processor subunits may be associated with a corresponding, dedicated one of the plurality of discrete memory banks. Moreover, as depicted, e.g., in FIGS. 7A and 7B, the distributed processor may further comprise a plurality of buses, each one of the plurality of buses connecting one of the plurality of processor subunits to at least another one of the plurality of processor subunits.

[0145] As explained above, the plurality of buses may be controlled in software. Accordingly, the plurality of buses may be free of timing hardware logic components such that data transfers between processor subunits and across corresponding ones of the plurality of buses are uncontrolled by timing hardware logic components. In one example, the plurality of buses may be free of bus arbiters such that data transfers between processor subunits and across corresponding ones of the plurality of buses are  
15 uncontrolled by bus arbiters.

[0146] In some embodiments, as depicted, e.g., in FIGS. 7A and 7B, the distributed processor may further comprise a second plurality of buses connecting one of the plurality of processor subunits to a corresponding, dedicated memory bank. Similar to the plurality of buses described above, the second plurality of buses may be free of timing hardware logic components such that data transfers between  
20 processor subunits and corresponding, dedicated memory banks are uncontrolled by timing hardware logic components. In one example, the second plurality of buses may be free of bus arbiters such that data transfers between processor subunits and corresponding, dedicated memory banks are uncontrolled by bus arbiters.

[0147] As used herein, the phrase “free of” does not necessarily imply the absolute absence of components, such as timing hardware logic components (e.g., bus arbiters, arbitration trees, FIFO controllers, mailboxes, or the like). Such components may still be included in a hardware chip described as “free of” those components. Instead, the phrase “free of” refers to the function of the hardware chip; that is, a hardware chip “free of” timing hardware logic components controls the timing of its data transfers without use of the timing hardware logic components, if any, included therein. For example, a  
25 hardware chip that executes code including sub-series of instructions that control data transfers between processor subunits of the hardware chip, even if the hardware chip includes timing hardware logic components as a secondary precaution to protect against collisions due to errors in the executed code.

[0148] As explained above, the plurality of buses may comprise at least one of wires or optical fibers between corresponding ones of the plurality of processor subunits. Accordingly, in one example, a distributed processor free of timing hardware logic components may include only wires or optical fibers without bus arbiters, arbitration trees, FIFO controllers, mailboxes, or the like.

5 [0149] In some embodiments, the plurality of processor subunits is configured to transfer data across at least one of the plurality of buses in accordance with code executed by the plurality of processor subunits. Accordingly, as explained below, a compiler may organize sub-series of instructions, each sub-series comprising code executed by a single processor subunit. The sub-series instructions may instruct  
10 the processor subunit when to transfer data onto one of the buses and when to retrieve data from the buses. When the sub-series are executed in tandem across the distributed processor, the timing of transfers between the processor subunits may be governed by the instructions to transfer and retrieve included in the sub-series. Thus, the code dictates timing of data transfers across at least one of the plurality of buses. The compiler may generate code to be executed by a single processor subunit. Additionally, the compiler may generate code to be executed by groups of processor subunits. In some cases, the compiler may treat  
15 all the processor subunits together as if they were one super-processor (e.g., a distributed processor), and the compiler may generate code for execution by that defined super-processor/distributed processor.

[0150] As explained above and depicted in FIGS. 7A and 7B, the plurality of processor subunits may be spatially distributed among the plurality of discrete memory banks within the memory array. Alternatively, the plurality of processor subunits may be clustered in one or more regions of the substrate,  
20 and the plurality of memory banks may be clustered in one or more other regions of the substrate. In some embodiments, a combination of spatial distribution and clustering may be used, as explained above.

[0151] In some embodiments, a distributed processor may comprise a substrate (e.g., a semiconductor substrate, including silicon and/or a circuit board, such as a flexible circuit board) with a memory array disposed on the substrate, the memory array including a plurality of discrete memory  
25 banks. A processing array may also be disposed on the substrate, the processing array including a plurality of processor subunits, as depicted, e.g., in FIGS. 7A and 7B. As explained above, each one of the processor subunits may be associated with a corresponding, dedicated one of the plurality of discrete memory banks. Moreover, as depicted, e.g., in FIGS. 7A and 7B, the distributed processor may further comprise a plurality of buses, each one of the plurality of buses connecting one of the plurality of  
30 processor subunits to a corresponding, dedicated one of the plurality of discrete memory banks.

[0152] As explained above, the plurality of buses may be controlled in software. Accordingly, the plurality of buses may be free of timing hardware logic components such that data transfers between a processor subunit and a corresponding, dedicated one of the plurality of discrete memory banks and across a corresponding one of the plurality of buses are not controlled by timing hardware logic  
35 components. In one example, the plurality of buses may be free of bus arbiters such that data transfers between processor subunits and across corresponding ones of the plurality of buses are uncontrolled by bus arbiters.

[0153] In some embodiments, as depicted, e.g., in FIGS. 7A and 7B, the distributed processor may further comprise a second plurality of buses connecting one of the plurality of processor subunits to at least another one of the plurality of processor subunits. Similar to the plurality of buses described above, the second plurality of buses may be free of timing hardware logic components such that data transfers between processor subunits and corresponding, dedicated memory banks are uncontrolled by timing hardware logic components. In one example, the second plurality of buses may be free of bus arbiters such that data transfers between processor subunits and corresponding, dedicated memory banks are uncontrolled by bus arbiters.

[0154] In some embodiments, the distributed processor may use a combination of software timing with hardware timing components. For example, a distributed processor may comprise a substrate (e.g., a semiconductor substrate, including silicon and/or a circuit board, such as a flexible circuit board) with a memory array disposed on the substrate, the memory array including a plurality of discrete memory banks. A processing array may also be disposed on the substrate, the processing array including a plurality of processor subunits, as depicted, e.g., in FIGS. 7A and 7B. As explained above, each one of the processor subunits may be associated with a corresponding, dedicated one of the plurality of discrete memory banks. Moreover, as depicted, e.g., in FIGS. 7A and 7B, the distributed processor may further comprise a plurality of buses, each one of the plurality of buses connecting one of the plurality of processor subunits to at least another one of the plurality of processor subunits. Moreover, as explained above, the plurality of processor subunits may be configured to execute software that controls timing of data transfers across the plurality of buses to avoid colliding data transfers on at least one of the plurality of buses. In such an example, the software may control the timing of the data transfers, but the transfers themselves may be controlled, at least in part, by one or more hardware components.

[0155] In such embodiments, the distributed processor may further comprise a second plurality of buses connecting one of the plurality of processor subunits to a corresponding, dedicated memory bank. Similar to the plurality of buses described above, the plurality of processor subunits may be configured to execute software that controls timing of data transfers across the second plurality of buses to avoid colliding data transfers on at least one of the second plurality of buses. In such an example, as explained above, the software may control the timing of the data transfers, but the transfers themselves may be controlled, at least in part, by one or more hardware components.

#### Division of Code

[0156] As explained above, hardware chips of the present disclosure may execute code in parallel across processor subunits included on a substrate forming the hardware chip. Additionally, hardware chips of the present disclosure may perform multitasking. For example, hardware chips of the present disclosure may perform area multitasking, in which one group of processor subunits of the hardware chip execute one task (e.g., audio processing) while another group of processor subunits of the hardware chip execute another task (e.g., image processing). In another example, hardware chips of the present disclosure may perform timing multitasking, in which one or more processor subunits of the hardware chip execute one task during a first period of time and another task during a second period of

time. A combination of area and timing multitasking may also be used such that one task may be assigned to a first group of processor subunits during a first period of time while another task may be assigned to a second group of processor subunits during the first period of time, after which a third task may be assigned to processor subunits included in the first group and the second group during a second period of time.

5 [0157] In order to organize machine code for execution on memory chips of the present disclosure, machine code may be divided between processor subunits of the memory chip. For example, a processor on a memory chip may comprise a substrate and a plurality of processor subunits disposed on the substrate. The memory chip may further comprise a corresponding plurality of memory banks  
10 disposed on the substrate, each one of the plurality processor subunits being connected to at least one dedicated memory bank not shared by any other processor subunit of the plurality of processor subunits. Each processor subunit on the memory chip may be configured to execute a series of instructions independent from other processor subunits. Each series of instructions may be executed by configuring one or more general processing elements of the processor subunit in accordance with code defining the series of instructions and/or by activating one or more special processing elements (e.g., one or more  
15 accelerators) of the processor subunit in accordance with a sequence provided in the code defining the series of instructions.

[0158] Accordingly, each series of instructions may define a series of tasks to be performed by a single processor subunit. A single task may comprise an instruction within an instruction set defined by  
20 the architecture of one or more processing elements in the processor subunit. For example, the processor subunit may include particular registers, and a single task may push data onto a register, pull data from a register, perform an arithmetic function on data within a register, perform a logic operation on data within a register, or the like. Moreover, the processor subunit may be configured for any number of operands, such as a 0-operand processor subunit (also called a “stack machine”), a 1-operand processor subunit  
25 (also called an accumulator machine), a 2-operand processor subunit (such as a RISC), a 3-operand processor subunit (such as a complex instruction set computer (CISC)), or the like. In another example, the processor subunit may include one or more accelerators, and a single task may activate an accelerator to perform a specific function, such as a MAC function, a MAX function, a MAX-0 function, or the like.

[0159] The series of instructions may further include tasks for reading and writing from the  
30 dedicated memory banks of the memory chip. For example, a task may include writing a piece of data to a memory bank dedicated to the processor subunit executing the task, reading a piece of data from a memory bank dedicated to the processor subunit executing the task, or the like. In some embodiments, the reading and writing may be performed by the processor subunit in tandem with a controller of the memory bank. For example, the processor subunit may execute a read or write task by sending a control  
35 signal to the controller to perform the read or write. In some embodiments, the control signal may include a particular address to use for reads and writes. Alternatively, the processor subunit may defer to the memory controller to select an available address for the reads and writes.

[0160] Additionally or alternatively, the reading and writing may be performed by one or more accelerators in tandem with a controller of the memory bank. For example, the accelerators may generate the control signals for the memory controller, similar to how the processor subunit generates control signals, as described above.

5 [0161] In any of the embodiments described above, an address generator may also be used to direct the reads and writes to specific addresses of a memory bank. For example, the address generator may comprise a processing element configured to generate memory addresses for reads and writes. The address generator may be configured to generate addresses in order to increase efficiency, e.g., by writing  
10 results of a later calculation to the same address as the results of a former calculation that are no longer needed. Accordingly, the address generator may generate the controls signals for the memory controller, either in response to a command from the processor subunit (e.g., from a processing element included therein or from one or more accelerator(s) therein) or in tandem with the processor subunit. Additionally or alternatively, the address generator may generate the addresses based on some configuration or registers for example generating a nested loop structure to iterate on certain addresses in the memory at a  
15 certain pattern.

[0162] In some embodiments, each series of instructions may comprise a set of machine code defining a corresponding series of tasks. Accordingly, the series of tasks described above may be encapsulated within machine code comprising the series of instructions. In some embodiments, as explained below with respect to FIG. 8, the series of tasks may be defined by a compiler configured to  
20 distribute a higher-level series of tasks amongst the plurality of logic circuits as a plurality of series of tasks. For example, the compiler may generate the plurality of series of tasks based on the higher-level series of tasks such that the processor subunits, executing each corresponding series of tasks in tandem, perform the same function as outlined by the higher-level series of tasks.

[0163] As explained further below, the higher-level series of tasks may comprise a set of  
25 instructions in a human-readable programming language. Correspondingly, the series of tasks for each processor subunit may comprise lower-level series of tasks, each of which comprises a set of instructions in a machine code.

[0164] As explained above with respect to FIGS. 7A and 7B, the memory chip may further  
30 comprise a plurality of buses, each bus connecting one of the plurality of processor subunits to at least one other of the plurality of processor subunits. Moreover, as explained above, data transfers on the plurality of buses may be controlled using software. Accordingly, data transfers across at least one of the plurality of buses may be predefined by the series of instructions included in a processor subunit connected to the at least one of the plurality of buses. Therefore, one of the tasks included in the series of  
35 instructions may include outputting data to one of the buses or pulling data from one of the buses. Such tasks may be executed by a processing element of the processor subunit or by one or more accelerators included in the processor subunit. In the latter embodiment, the processor subunit may perform a calculation or send a control signal to a corresponding memory bank in the same cycle during which accelerator(s) pull data from or place data on one of the buses.



[0165] In one example, the series of instructions included in the processor subunit connected to the at least one of the plurality of buses may include a sending task that comprises a command for the processor subunit connected to the at least one of the plurality of buses to write data to the at least one of the plurality of buses. Additionally or alternatively, the series of instructions included in the processor subunit connected to the at least one of the plurality of buses may include a receiving task that comprises a command for the processor subunit connected to the at least one of the plurality of buses to read data from the at least one of the plurality of buses.

[0166] Additionally or alternatively to distribution of code amongst processor subunits, data may be divided between memory banks of the memory chip. For example, as explained above, a distributed processor on a memory chip may comprise a plurality of processor subunits disposed on the memory chip and a plurality of memory banks disposed on the memory chip. Each one of the plurality of memory banks may be configured to store data independent from data stored in other ones of the plurality of memory banks, and each one of the plurality of processor subunits may be connected to at least one dedicated memory bank from among the plurality of memory banks. For example, each processor subunit may have access to one or more memory controllers of one or more corresponding memory banks dedicated to the processor subunit, and no other processor subunit may have access to these corresponding one or more memory controllers. Accordingly, the data stored in each memory bank may be unique to the dedicated processor subunit. Moreover, the data stored in each memory bank may be independent of the memory stored in other memory banks because no memory controllers may be shared between memory banks.

[0167] In some embodiments, as described below with respect to FIG. 8, the data stored in each of the plurality of memory banks may be defined by a compiler configured to distribute data amongst the plurality of memory banks. Moreover, the compiler may be configured to distribute data defined in a higher-level series of tasks amongst the plurality of memory banks using a plurality of lower-level tasks distributed amongst corresponding processor subunits.

[0168] As explained further below, the higher-level series of tasks may comprise a set of instructions in a human-readable programming language. Correspondingly, the series of tasks for each processor subunit may comprise lower-level series of tasks, each of which comprises a set of instructions in a machine code.

[0169] As explained above with respect to FIGS. 7A and 7B, the memory chip may further comprise a plurality of buses, each bus connecting one of the plurality of processor subunits to one or more corresponding, dedicated memory banks from among the plurality of memory banks. Moreover, as explained above, data transfers on the plurality of buses may be controlled using software. Accordingly, data transfers across a particular one of the plurality of buses may be controlled by a corresponding processor subunit connected to the particular one of the plurality of buses. Therefore, one of the tasks included in the series of instructions may include outputting data to one of the buses or pulling data from one of the buses. As explained above, such tasks may be executed by (i) a processing element of the processor subunit or (ii) one or more accelerators included in the processor subunit. In the latter

embodiment, the processor subunit may perform a calculation or use buses connecting the processor subunit to other processor subunits in the same cycle during which accelerator(s) pull data from or place data on one of the buses connected to the one or more corresponding, dedicated memory banks.

[0170] Therefore, in one example, the series of instructions included in the processor subunit  
5 connected to the at least one of the plurality of buses may include a sending task. The sending task may comprise a command for the processor subunit connected to the at least one of the plurality of buses to write data to the at least one of the plurality of buses for storage in the one or more corresponding, dedicated memory banks. Additionally or alternatively, the series of instructions included in the processor subunit connected to the at least one of the plurality of buses may include a receiving task. The receiving  
10 task may comprise a command for the processor subunit connected to the at least one of the plurality of buses to read data from the at least one of the plurality of buses for storage in the one or more corresponding, dedicated memory banks. Accordingly, the sending and receiving tasks in such embodiments may comprise control signals that are sent, along the at least one of the plurality of buses, to one or more memory controllers of the one or more corresponding, dedicated memory banks. Moreover,  
15 the sending and receiving tasks may be executed by one portion of the processing subunit (e.g., by one or more accelerators thereof) concurrently with a calculation or other task executed by another portion of the processing subunit (e.g., by one or more different accelerators thereof). An example of such a concurrent execution may include a MAC-relay command, in which receiving, multiplying, and sending are executed in tandem.

[0171] In addition to distributing data amongst the memory banks, particular portions of data  
20 may be duplicated across different memory banks. For example, as explained above, a distributed processor on a memory chip may comprise a plurality of processor subunits disposed on the memory chip and a plurality of memory banks disposed on the memory chip. Each one of the plurality of processor subunits may be connected to at least one dedicated memory bank from among the plurality of memory  
25 banks, and each memory bank of the plurality of memory banks may be configured to store data independent from data stored in other ones of the plurality of memory banks. Moreover, at least some of the data stored in one particular memory bank from among the plurality of memory banks may comprise a duplicate of data stored in at least another one of the plurality of memory banks. For example, a number, string, or other type of data used in the series of instructions may be stored in a plurality of memory banks  
30 dedicated to different processor subunits rather than being transferred from one memory bank to other processor subunits in the memory chip.

[0172] In one example, parallel string matching may use data duplication described above. For example, a plurality of strings may be compared to the same string. A conventional processor would compare each string in the plurality to the same string in sequence. On a hardware chip of the present  
35 disclosure, the same string may be duplicated across the memory banks such that the processor subunits may compare a separate string in the plurality to the duplicated string in parallel.

[0173] In some embodiments, as described below with respect to FIG. 8, the at least some data duplicated across the one particular memory bank from among the plurality of memory banks and the at

least another one of the plurality of memory banks is defined by a compiler configured to duplicate data across memory banks. Moreover, the compiler may be configured to duplicate the at least some data using a plurality of lower-level tasks distributed amongst corresponding processor subunits.

[0174] Duplication of data may be useful for certain tasks that re-use the same portions of data across different calculations. By duplicating these portions of data, the different calculations may be distributed amongst processor subunits of the memory chip for parallel execution while each processor subunit may store the portions of data in, and access the stored portions from, a dedicated memory bank (rather than pushing and pulling the portions of data across buses connecting the processor subunits). In one example, the at least some data duplicated across the one particular memory bank from among the plurality of memory banks and the at least another one of the plurality of memory banks may comprise weights of a neural network. In this example, each node in the neural network may be defined by at least one processor subunit from among the plurality of processor subunits. For example, each node may comprise machine code executed by the at least one processor subunit defining the node. In this example, duplication of the weights may allow each processor subunit to execute machine code to effect, at least in part, a corresponding node while only accessing one or more dedicated memory banks (rather than performing data transfers with other processor subunits). Because the timing of reads and writes to the dedicated memory bank(s) are independent of other processor subunits while the timing of data transfers between processor subunits requires timing synchronization (e.g., using software, as explained above), duplication of memory to avoid data transfers between processor subunits may produce further efficiencies in overall execution.

[0175] As explained above with respect to FIGS. 7A and 7B, the memory chip may further comprise a plurality of buses, each bus connecting one of the plurality of processor subunits to one or more corresponding, dedicated memory banks from among the plurality of memory banks. Moreover, as explained above, data transfers on the plurality of buses may be controlled using software. Accordingly, data transfers across a particular one of the plurality of buses may be controlled by a corresponding processor subunit connected to the particular one of the plurality of buses. Therefore, one of the tasks included in the series of instructions may include outputting data to one of the buses or pulling data from one of the buses. As explained above, such tasks may be executed by (i) a processing element of the processor subunit or (ii) one or more accelerators included in the processor subunit. As further explained above, such tasks may include a sending task and/or a receiving tasks that comprise control signals that are sent, along the at least one of the plurality of buses, to one or more memory controllers of the one or more corresponding, dedicated memory banks.

[0176] FIG. 8 depicts a flowchart of a method 800 for compiling a series of instructions for execution on an exemplary memory chip of the present disclosure, e.g., as depicted in FIGS. 7A and 7B. Method 800 may be implemented by any conventional processor, whether generic or special-purpose.

[0177] Method 800 may be executed as a portion of a computer program forming a compiler. As used herein, a “compiler” refers to any computer program that converts a higher-level language (e.g., a procedural language, such as C, FORTRAN, BASIC, or the like; an object-oriented language, such as

Java, C++, Pascal, Python, or the like; etc.) to a lower-level language (e.g., assembly code, object code, machine code, or the like). The compiler may allow a human to program a series of instructions in a human-readable language, which is then converted to a machine-executable language.

[0178] At step 810, the processor may assign tasks associated with the series of instructions to  
5 different ones of the processor subunits. For example, the series of instructions may be divided into subgroups, the subgroups to be executed in parallel across the processor subunits. In one example, a neural network may be divided into its nodes, and one or more nodes may be assigned to separate processor subunits. In this example, each subgroup may comprise a plurality of nodes connected across  
10 different layers. Thus, a processor subunit may implement a node from a first layer of the neural network, a node from a second layer connected to the node from the first layer implemented by the same processor subunit, and the like. By assigning nodes based on their connections, data transfers between the processor subunits may be lessened, which may result in greater efficiency, as explained above.

[0179] As explained above depicted in FIGS. 7A and 7B, the processor subunits may be spatially distributed among the plurality of memory banks disposed on the memory chip. Accordingly, the  
15 assignment of tasks may be, at least in part, a spatial divisional as well as a logical division.

[0180] At step 820, the processor may generate tasks to transfer data between pairs of the processor subunits of the memory chip, each pair of processor subunits being connected by a bus. For example, as explained above, the data transfers may be controlled using software. Accordingly, processor subunits may be configured to push and pull data on buses at synchronized times. The generated tasks  
20 may thus include tasks for performing this synchronized pushing and pulling of data.

[0181] As explained above, step 820 may include pre-processing to account for the internal behavior, including timing and latencies, of the processor subunits. For example, the processor may use known times and latencies of the processor subunits (e.g., the time to push data to a bus, the time to pull data from a bus, the latency between a calculation and a push or pull, or the like) to ensure that the  
25 generated tasks synchronize. Therefore, the data transfers comprising at least one push by one or more processor subunits and at least one pull by one or more processor subunits may occur simultaneously rather than incurring a delay due to timing differences between the processor subunits, latencies of the processor subunits, or the like.

[0182] At step 830, the processor may group the assigned and generated tasks into the plurality  
30 of groups of sub-series instructions. For example, the sub-series instructions may each comprise a series of tasks for execution by a single processor subunit. Therefore, each of the plurality of groups of sub-series instructions may correspond to a different one of the plurality of processor sub-units. Accordingly, steps 810, 820, and 830 may result in dividing the series of instructions into a plurality of groups of sub-series instructions. As explained above, step 820 may ensure that any data transfers between the different  
35 groups are synchronized.

[0183] At step 840, the processor may generate machine code corresponding to each of the plurality of groups of subs-series instructions. For example, the higher-level code representing sub-series

instructions may be converted to lower-level code, such as machine code, executable by corresponding processor subunits.

[0184] At step 850, the processor may assign the generated machine code corresponding to each of the plurality of groups of sub-series instructions to a corresponding one of the plurality of processor subunits in accordance with the division. For example, the processor may label each sub-series instructions with an identifier of the corresponding processor subunit. Thus, when the sub-series instructions are uploaded to a memory chip for execution (e.g., by host 350 of FIG. 3A), each sub-series may configure a correct processor subunit.

[0185] In some embodiments, assigning tasks associated with the series of instructions to the different ones of the processor subunits may depend, at least in part, on a spatial proximity between two or more of the processor subunits on the memory chip. For example, as explained above, efficiency may be increased by lessening the number of data transfers between processor subunits. Accordingly, the processor may minimize data transfers that move data across more than two of processor subunits. Therefore, the processor may use a known layout of the memory chip in combination with one or more optimization algorithms (such as a greedy algorithm) in order to assign sub-series to processor subunits in a way that maximizes (at least locally) adjacent transfers and minimizes (at least locally) transfers to non-neighboring processor subunits.

[0186] Method 800 may include further optimizations for the memory chips of the present disclosure. For example, the processor may group data associated with the series of instructions based on the division and assign the data to the memory banks in accordance with the grouping. Accordingly, the memory banks may hold data used for the sub-series instructions assigned to each processor subunit to which each memory bank is dedicated.

[0187] In some embodiments, grouping the data may include determining at least a portion of the data to duplicate in two or more of the memory banks. For example, as explained above, some data may be used across more than one sub-series instructions. Such data may be duplicated across the memory banks dedicated to the plurality of processor subunits to which the different sub-series instructions are assigned. This optimization may further reduce data transfers across processor subunits.

[0188] The output of method 800 may be input to a memory chip of the present disclosure for execution. For example, a memory chip may comprise a plurality of processor subunits and a corresponding plurality of memory banks, each processor subunit being connected to at least one memory bank dedicated to the processor subunit, and the processor subunits of the memory chip may be configured to execute the machine code generated by method 800. As explained above with respect to FIG. 3A, host 350 may input the machine code generated by method 800 to the processor subunits for execution.

#### Sub-banks and Sub-controllers

[0189] In conventional memory banks, controllers are provided at the bank level. Each bank includes a plurality of mats, which are typically arranged in a rectangular manner but may be arranged in any geometrical shape. Each mat includes a plurality of memory cells, which are also typically arranged

in a rectangular manner but may be arranged in any geometrical shape. Each cell may store a single bit of data (e.g., depending on whether the cell is retained at a high voltage or a low voltage).

[0190] An example of this conventional architecture is depicted in FIGS. 9 and 10. As shown in FIG. 9, at the bank level, a plurality of mats (e.g., mats 930-1, 930-2, 940-1, and 940-2) may form bank 5 900. In a conventional rectangular organization, bank 900 may be controlled across global wordlines (e.g., wordline 950) and global bitlines (e.g., bitline 960). Accordingly, row decoder 910 may select the correct wordline based on an incoming control signal (e.g., a request for a read from an address, a request for a write to an address, or the like) and global sense amplifier 920 (and/or a global column decoder, not shown in FIG. 9) may select the correct bitline based on the control signal. Amplifier 920 may also 10 amplify any voltage levels from a selected bank during a read operation. Although depicted as using a row decoder for initial selecting and performing amplification along columns, a bank may additionally or alternatively use a column decoder for initial selecting and perform amplification along rows.

[0191] FIG. 10 depicts an example of a mat 1000. For example, mat 1000 may form a portion of a memory bank, such as bank 900 of FIG. 9. As depicted in FIG. 10, a plurality of cells (e.g., cells 1030- 15 1, 1030-2, and 1030-3) may form mat 1000. Each cell may comprise a capacitor, a transistor, or other circuitry that stores at least one bit of data. For example, a cell may comprise a capacitor that is charged to represent a '1' and discharged to represent a '0' or may comprise a flip-flop having a first state representing a '1' and a second state representing a '0.' A conventional mat may comprise, for example, 512 bits by 512 bits. In embodiments where mat 1000 forms a portion of MRAM, ReRAM, or the like, a cell may comprise a transistor, resistor, capacitor or other mechanism for isolating an ion or portion of a 20 material that stores at least one bit of data. For example, a cell may comprise a electrolyte ion, a portion of chalcogenide glass, or the like, having a first state representing a '1' and a second state representing a '0.'

[0192] As further depicted in FIG. 10, in a conventional rectangular organization, mat 1000 may 25 be controlled across local wordlines (e.g., wordline 1040) and local bitlines (e.g., bitline 1050). Accordingly, wordline drivers (e.g., wordline driver 1020-1, 1020-2, . . . , 1020-x) may control the selected wordline to perform a read, write, or refresh based on a control signal from a controller associated with the memory bank of which mat 1000 forms a part (e.g., a request for a read from an address, a request for a write to an address, a refresh signal). Moreover, local sense amplifiers (e.g., local 30 amplifiers 1010-1, 1010-2, . . . , 1010-x) and/or local column decoders (not shown in FIG. 10) may control the selected bitline to perform a read, write, or refresh. The local sense amplifiers may also amplify any voltage levels from a selected cell during a read operation. Although depicted as using a wordline driver for initial selecting and performing amplification along columns, a mat may instead use a bitline driver for initial selecting and perform amplification along rows.

[0193] As explained above, a large number of mats are duplicated to form a memory bank. 35 Memory banks may be grouped to form a memory chip. For example, a memory chip may comprise eight to thirty-two memory banks. Accordingly, pairing processor subunits with memory banks on a conventional memory chip may result in only eight to thirty two processor subunits. Accordingly,

embodiments of the present disclosure may include memory chips with additional sub-bank hierarchy. These memory chips of the present disclosure may then include processor subunits with memory sub-banks used as the dedicated memory banks paired with the processor subunits allowing for a larger number of sub processors, which may then achieve higher parallelism and performance of in-memory computing.

5 [0194] In some embodiments of the present disclosure, the global row decoder and global sense amplifier of bank 900 may be replaced with sub-bank controllers. Accordingly, rather than sending control signals to a global row decoder and a global sense amplifier of the memory bank, a controller of the memory bank may direct the control signal to the appropriate sub-bank controller. The direction may  
10 be controlled dynamically or may be hard-wired (e.g., via one or more logic gates). In some embodiments, fuses may be used to indicate the controller of each sub bank or mat whether to block or pass the control signal to the appropriate sub-bank or mat. In such embodiments, faulty sub-banks may thus be deactivated using the fuses.

[0195] In one example of such embodiments, a memory chip may include a plurality of memory  
15 banks, each memory bank having a bank controller and a plurality of memory sub-banks, each memory sub-bank having a sub-bank row decoder and a sub-bank column decoder for allowing reads and writes to locations on the memory sub-bank. Each sub-bank may comprise a plurality of memory mats, each memory mat having a plurality of memory cells and may have internally local row decoders, column decoders, and/or local sense amplifiers. The sub-bank row decoders and the sub-bank column decoders  
20 may process read and write requests from the bank controller or from a sub-bank processor subunit used for in memory computations on the sub-bank memory, as described below. Additionally, each memory sub-bank may further have a controller configured to determine whether to process read requests and write requests from the bank controller and/or to forward them to the next level (e.g., of row and column decoders on a mat) or to block the requests, e.g., to allow an internal processing element or processor  
25 subunit to access the memory. In some embodiments, the bank controller may be synchronized to a system clock. However, the sub-bank controllers may be not synchronized to the system clock.

[0196] As explained above, the use of sub-banks may allow for the inclusion of a larger number processor subunits in the memory chip than if processor subunits were paired with memory banks of conventional chips. Accordingly, each sub-bank may further have a processor subunit using the sub-bank  
30 as a dedicated memory. As explained above, the processor subunit may comprise a RISC, a CISC, or other general-purpose processing subunit and/or may comprise one or more accelerators. Additionally, the processor subunit may include an address generator, as explained above. In any of the embodiments described above, each processor subunit may be configured to access a sub-bank dedicated to the processor subunit using the row decoder and the column decoder of the sub-bank without using the bank  
35 controller. The processor sub-unit associated with the sub-bank may also handle the memory mats (including the decoder and memory redundancy mechanisms, described below) and/or determine whether a read or write request from an upper level (e.g., the bank level or the memory level) is forwarded and handled accordingly.

[0197] In some embodiments, the sub-bank controller may further include a register that stores a state of the sub-bank. Accordingly, the sub-bank controller may return an error if the sub-bank controller receives a control signal from the memory controller while the register indicates that the sub-bank is in use. In embodiments where each sub-bank further includes a processor subunit, the register may indicate an error if the processor subunit in the sub-bank is accessing the memory in conflict with an external request from the memory controller.

[0198] FIG. 11 shows an example of another embodiment of a memory bank using sub-bank controllers. In the example of FIG. 11, bank 1100 has a row decoder 1110, a column decoder 1120, and a plurality of memory sub-banks (e.g., sub-banks 1170a, 1170b, and 1170c) with sub-bank controllers (e.g., controllers 1130a, 1130b, and 1130c). The sub-bank controllers may include address resolvers (e.g., resolvers 1140a, 1140b, and 1140c), which may determine whether to pass a request to one or more sub-banks controlled by the sub-bank controller.

[0199] The sub-bank controllers may further include one or more logic circuits (e.g., logic 1150a, 1150b, and 1150c). For example, a logic circuit comprising one or more processing elements may allow for one or more operations, such as refreshing of cells in the sub-bank, clearing of cells in the sub-bank, or the like, to be performed without processing requests externally from bank 1100. Alternatively, the logic circuit may comprise a processor subunit, as explained above, such that the processor sub-unit has any sub-banks controlled by the sub-bank controller as corresponding, dedicated memory. In the example of FIG. 11, logic 1150a may have sub-bank 1170a as a corresponding, dedicated memory, logic 1150b may have sub-bank 1170b as a corresponding, dedicated memory, and logic 1150c may have sub-bank 1170c as a corresponding, dedicated memory. In any of the embodiments described above, the logic circuits may have buses to the sub-banks, e.g., buses 1131a, 1131b, or 1131c. As further depicted in FIG. 11, the sub-bank controllers may each include a plurality of decoders, such as a sub-bank row decoder and a sub-bank column decoder for allowing reads and writes, either by a processing element or processor subunit or by a higher-level memory controller issuing commands, to locations on the memory sub-bank(s). For example, sub-bank controller 1130a includes decoders 1160a, 1160b, and 1160c, sub-bank controller 1130b includes decoders 1160d, 1160e, and 1160f, and sub-bank controller 1130c includes decoders 1160g, 1160h, and 1160i. The sub-bank controllers may, based on a request from bank row decoder 1110, select a wordline using the decoders included in the sub-bank controllers. The described system may allow a processing element or processor subunit of the sub-bank to access the memory without interrupting other banks and even other sub-banks, thereby allowing each sub-bank processor subunit to perform memory computations in parallel with the other sub-bank processor subunits.

[0200] Furthermore, each sub-bank may comprise a plurality of memory mats, each memory mat having a plurality of memory cells. For example, sub-bank 1170a includes mats 1190a-1, 1190a-2, . . . , 1190a-x; sub-bank 1170b includes mats 1190b-1, 1190b-2, . . . , 1190b-x; and sub-bank 1170c includes mats 1190c-1, 1190c-2, . . . , 1190c-3. As further depicted in FIG. 11, each sub-bank may include at least one decoder. For example, sub-bank 1170a includes decoder 1180a, sub-bank 1170b includes decoder



1180b, and sub-bank 1170c includes decoder 1180c. Accordingly, bank column decoder 1120 may select a global bitline (e.g., bitline 1121a or 1121b) based on external requests while the sub-bank selected by bank row decoder 1110 may use its column decoder to select a local bitline (e.g., bitline 1181a or 1181b) based on local requests from the logic circuit to which the sub-bank is dedicated. Accordingly, each processor subunit may be configured to access a sub-bank dedicated to the processor subunit using the row decoder and the column decoder of the sub-bank without using the bank row decoder and the bank column decoder. Thus, each processor subunit may access a corresponding sub-bank without interrupting other sub-banks. Moreover, sub-bank decoders may reflect accessed data to the bank decoders when the request to the sub-bank is external to the processor subunit. Alternatively, in embodiments where each sub-bank has only one row of memory mats, the local bitlines may be the bitlines of the mat rather than bitlines of the sub-bank.

[0201] A combination of embodiments using sub-bank row decoders and sub-bank column decoders with the embodiment depicted in FIG. 11 may be used. For example, the bank row decoder may be eliminated but the bank column decoder retained and local bitlines used.

[0202] FIG. 12 shows an example of an embodiment of a memory sub-bank 1200 having a plurality of mats. For example, sub-bank 1200 may represent a portion of sub-bank 1100 of FIG. 11 or may represent an alternative implementation of a memory bank. In the example of FIG. 12, sub-bank 1200 includes a plurality of mats (e.g., mats 1240a and 1240b). Moreover, each mat may include a plurality of cells. For example, mat 1240a includes cells 1260a-1, 1260a-2, . . . , 1260a-x, and mat 1240b includes cells 1260b-1, 1260b-2, . . . , 1260b-x.

[0203] Each mat may be assigned a range of addresses that will be assigned to the memory cells of the mat. These addresses may be configured at production such that mats may be shuffled around and such that faulted mats may be deactivated and left unused (e.g., using one or more fuses, as explained further below).

[0204] Sub-bank 1200 receives read and write requests from memory controller 1210. Although not depicted in FIG. 12, requests from memory controller 1210 may be filtered through a controller of sub-bank 1200 and directed to an appropriate mat of sub-bank 1200 for address resolution. Alternatively, at least a portion (e.g., higher bits) of an address of a request from memory controller 1210 may be transmitted to all mats of sub-bank 1200 (e.g., mats 1240a and 1240b) such that each mat may process the full address and the request associated with the address only if the mat's assigned address range includes the address specified in the command. Similar to the sub-bank direction described above, the mat determination may be dynamically controlled or may be hardwired. In some embodiments, fuses may be used to determine the address range for each mat, also allowing for disabling of faulty mats by assigning an illegal address range. Mats may additionally or alternatively be disabled by other common methods or connection of fuses.

[0205] In any of the embodiments described above, each mat of the sub-bank may include a row decoder (e.g., row decoder 1230a or 1230b) for selection of a wordline in the mat. In some embodiments, each mat may further include fuses and comparators (e.g., 1220a and 1220b). As described above, the

comparators may allow each mat to determine whether to process an incoming request, and the fuses may allow each mat to deactivate if faulty. Alternatively, row decoders for the bank and/or sub-bank may be used rather than a row decoder in each mat.

[0206] Furthermore, in any of the embodiments described above, a column decoder included in the appropriate mat (e.g., column decoder 1250a or 1250b) may select a local bitline (e.g., bitline 1251 or 1253). The local bitline may be connected to a global bitline of the memory bank. In embodiments where the sub-bank has local bitlines of its own, the local bitline of the cell may be further connected to the local bitline of the sub-bank. Accordingly, data in the selected cell may be read through the column decoder (and/or sense amplifier) of the cell, then through the column decoder (and/or sense amplifier) of the sub-bank (in embodiments including a sub-bank column decoder and/or sense amplifier), and then through the column decoder (and/or sense amplifier) of the bank.

[0207] Mat 1200 may be duplicated and arrayed to form a memory bank (or a memory sub-bank). For example, a memory chip of the present disclosure may comprise a plurality of memory banks, each memory bank having a plurality of memory sub-banks, and each memory sub-bank having a sub-bank controller for processing reads and writes to locations on the memory sub-bank. Furthermore, each memory sub-bank may comprise a plurality of memory mats, each memory mat having a plurality of memory cells and having a mat row decoder and a mat column decoder (e.g., as depicted in FIG. 12). The mat row decoders and the mat column decoders may process read and write requests from the sub-bank controller. For example, the mat decoders may receive all requests and determine (e.g., using a comparator) whether to process the request based on a known address range of each mat, or the mat decoders may only receive requests within the known address range based on selection of a mat by the sub-bank (or bank) controller.

#### Controller Data Transfers

[0208] Any of the memory chips of the present disclosure may also share data using memory controllers (or sub-bank controllers or mat controllers) in addition to sharing data using processing subunits. For example, a memory chip of the present disclosure may comprise a plurality of memory banks (e.g., an SRAM bank, a DRAM bank, or the like), each memory bank having a bank controller, a row decoder, and a column decoder for allowing reads and writes to locations on the memory bank, as well as a plurality of buses connecting each controller of the plurality of bank controllers to at least one other controller of the plurality of bank controllers. The plurality of buses may be similar to the buses connecting the processing subunits, as described above, but connecting the bank controllers directly rather than through the processing subunits. Furthermore, although described as connecting the bank controllers, buses may additionally or alternatively connect sub-bank controllers and/or mat controllers.

[0209] In some embodiments, the plurality of buses may be accessed without interruption of data transfers on main buses of the memory banks connected to one or more processor subunits. Accordingly, a memory bank (or sub-bank) may transmit data to or from a corresponding processor subunit in the same clock cycle as transmitting data to or from a different memory bank (or sub-bank). In embodiments where each controller is connected to a plurality of other controllers, the controllers may be configurable for

selection of one other of the other controllers for sending or receiving of data. In some embodiments, each controller may be connected to at least one neighboring controller (e.g., pairs of spatially adjacent controllers may be connected to one another).

#### Redundant Logic in Memory Circuits

5 [0210] The disclosure is generally directed to a memory chip with primary logic portions for on-chip data processing. The memory chip may include redundant logic portions, which may replace defective primary logic portions to increase the fabrication yield of the chip. Thus, the chip may include on-chip components that allow a configuration of logic blocks in the memory chip based on individual testing of the logic portions. This feature of the chip may increase yields because a memory chip with  
10 larger areas dedicated to logic portions is more susceptible to fabrication failures. For example, DRAM memory chips with large redundant logic portions may be susceptible to fabrication issues that reduce yield. However, implementing redundant logic portions may result in increased yield and reliability because it provides a manufacturer or user of DRAM memory chips to turn on or off full logic portions while maintaining the ability of high parallelism. It should be noted that here and throughout the  
15 disclosure, example of certain memory types (such as DRAM) may be identified in order to facilitate the explanation of disclosed embodiments. It is to be understood, however, that in such instances the identified memory types are not intended to be limiting. Rather, memory types such as DRAM, Flash, SRAM, ReRAM, PRAM, MRAM, ROM, or any other memory may be used together with the disclosed embodiments even if fewer examples are specifically identified in a certain section of the disclosure.

20 [0211] Fig. 13 is a block diagram of an exemplary memory chip 1300, consistent with disclosed embodiments. Memory chip 1300 may be implemented as a DRAM memory chip. Memory chip 1300 may also be implemented as any type of memory volatile or non-volatile, such as Flash, SRAM, ReRAM, PRAM, and/or MRAM, etc. Memory chip 1300 may include a substrate 1301 in which an address manager 1302, a memory array 1304 including a plurality of memory banks, 1304(a,a) to 1304(z,z), a  
25 memory logic 1306, a business logic 1308, and a redundant business logic 1310 are disposed. Memory logic 1306 and business logic 1308 may constitute primary logic blocks, while redundant business logic 1310 may constitute redundant blocks. In addition, memory chip 1300 may include configuration switches, which may include deactivation switches 1312, and an activation switches 1314. Deactivation switches 1312 and activation switches 1314 may also be disposed in the substrate 1301. In this  
30 Application, memory logic 1306, business logic 1308, and redundant business logic 1310 may also be collectively referred to as the “logic blocks.”

[0212] Address manager 1302 may include row and column decoders or other type of memory auxiliaries. Alternatively, or additionally, address manager 1302 may include a microcontroller or processing unit.

35 [0213] In some embodiments, as shown in Fig. 13, memory chip 1300 may include a single memory array 1304 that may arrange the plurality of memory blocks in a two-dimensional array on substrate 1301. In other embodiments, however, memory chip 1300 may include multiple memory arrays 1304 and each of the memory arrays 1304 may arrange memory blocks in different configurations. For

example, memory blocks in at least one of the memory arrays (also known as memory banks) may be arranged in a radial distribution to facilitate routing between address manager 1302 or memory logic 1306 to the memory blocks.

[0214] Business logic 1308 may be used to do the in-memory computation of an application that is not related to the logic used to manage the memory itself. For example, business logic 1308 may implement functions related to AI such as floating, integer, or MAC operations used as activation functions. In addition, business logic 1308 may implement data base related functions like min, max, sort, count, among others. Memory logic 1306 may perform tasks related to memory management, including (but not limited to) read, write, and refresh operations. Therefore, business logic may be added in one or more of the bank level, mats level, or a group of mats level. Business logic 1308 may have one or more address outputs and one or more data inputs/outputs. For instance, business logic 1308 can address by row\column lines to address manager 1302. In certain embodiments, however, the logic blocks may be additionally or alternatively addressed via data inputs\outputs.

[0215] Redundant business logic 1310 may be a replicate of business logic 1308. In addition, redundant business logic 1310 may be connected to deactivation switches 1312 and/or activation switches 1314, which may include small fuse\anti-fuse, and used for logic disabling or enabling one of the instances (e.g., an instance which is connected by default) and enable one of the other logic blocks (e.g., an instance which is disconnected by default). In some embodiments, as further described in connection to Fig. 15, the redundancy of blocks may be local within a logic block, such as business logic 1308.

[0216] In some embodiments, the logic blocks in memory chip 1300 may be connected to subsets of memory array 1304 with dedicated buses. For example, a set of memory logic 1306, business logic 1308, and redundant business logic 1310 may be connected to the first row of memory blocks in memory array 1304 (i.e., memory blocks 1304 (a,a) to 1304 (a,z)). The dedicated buses may allow associated logic blocks to quickly access data from the memory blocks without requirements of opening communication lines through, for example, address manager 1302.

[0217] Each of the plurality of primary logic blocks may be connected to at least one of the plurality of memory banks 1304. Also, redundant blocks, such as redundant business block 1310, may be connected to at least one of the memory instances 1304(a,a)-(z,z). Redundant blocks may replicate at least one of the plurality of primary logic blocks, such as memory logic 1306 or business logic 1308. Deactivation switches 1312 may be connected to at least one of the plurality of primary logic blocks and activation switches 1314 may be connected to at least one of the plurality of redundant blocks.

[0218] In these embodiments, upon detecting of a fault associated with one of the plurality of primary logic blocks (memory logic 1306 and/or business logic 1308), deactivation switches 1312 may be configured to disable the one of the plurality of primary logic blocks. Simultaneously, activation switches 1314 may be configured to enable one of the plurality of redundant blocks, such as redundant logic block 1310, that replicates the one of the plurality of primary logic blocks.

[0219] In addition, activation switches 1314 and deactivation switches 1312, which may collectively be referred to as "configuration switches," may include an external input to configure the

status of the switch. For instance, activation switches 1314 may be configured so an activation signal in the external input causes a closed switch condition, while deactivation switches 1312 may be configured so a deactivation signal in the external input causes an open switch condition. In some embodiments, all configuration switches in 1300 may be deactivated by default and become activated or enabled after a test  
5 indicates an associated logic block is functional and a signal is applied in the external input.

Alternatively, in some cases, all configuration switches in 1300 may be enabled by default and may be deactivated or disabled after a test indicates an associated logic block is not functional and a deactivation signal is applied in the external input.

[0220] Regardless of whether a configuration switch is initially enabled or disabled, upon  
10 detection of a fault associated with an associated logic block, the configuration switch may disable the associated logic block. In cases where the configuration switch is initially enabled, the state of the configuration switch may be changed to disabled in order to disable the associated logic block. In cases where the configuration switch is initially disabled, the state of the configuration switch may be left in its disabled state in order to disable the associated logic block. For example, the result of an operability test  
15 may indicate that a certain logic block is nonoperational or that it fails to operate within certain specifications. In such cases, the logic block may be disabled by not enabling its corresponding configuration switch.

[0221] In some embodiments, configuration switches may be connected to two or more logic blocks and may be configured to choose between different logic blocks. For example, a configuration  
20 switch may be connected to both business logic 1308 and redundant logic block 1310. Configuration switch may enable redundant logic block 1310 while disabling business logic 1308.

[0222] Alternatively, or additionally, at least one of the plurality of primary logic blocks (memory logic 1306 and/or business logic 1308) may be connected to a subset of the plurality of memory banks or memory instances 1304 with a first dedicated connection. Then, at least one of the plurality of  
25 redundant blocks (such as redundant business logic 1310), which replicates the at least one of the plurality of primary logic blocks, may be connected to the subset of the same plurality of memory banks or instances 1304 with a second dedicated connection.

[0223] Moreover, memory logic 1306 may have different functions and capabilities than business logic 1308. For example, while memory logic 1306 may be designed to enable read and write  
30 operations in the memory bank 1304, business logic 1308 may be designed to perform in-memory computations. Therefore, if the business logic 1308 includes a first business logic block, and the business logic 1308 includes a second business logic block (like redundant business logic 1310), it is possible to disconnect defective business logic 1308 and reconnect redundant business logic 1310 without missing any capability.

[0224] In some embodiments, configuration switches (including deactivation switches 1312 and  
35 activation switches 1314) may be implemented with a fuse, an anti-fuse, or a programmable device (including a one-time programmable device), or other form of non-volatile memory.

[0225] Fig. 14 is a block diagram of an exemplary redundant logic block set 1400, consistent with disclosed embodiments. In some embodiments, redundant logic block set 1400 may be disposed in substrate 1301. Redundant logic block set 1400 may include at least one of business logic 1308, and redundant business logic 1310, connected to switches 1312 and 1314, respectively. In addition, business logic 1308 and redundant business logic 1310 may be connected to an address bus 1402 and a data bus 1404.

[0226] In some embodiments, as shown in Fig. 14, the switches 1312 and 1314 may connect logic blocks to a clock node. In this way, the configuration switches may engage or disengage the logic blocks from the clock signal, effectively activating or deactivating the logic blocks. In other embodiments, however, switches 1312 and 1314 may connect logic blocks to other nodes for activation or deactivation. For instance, configuration switches may connect logic blocks to a voltage supply node (e.g., VCC) or to the ground node (e.g., GND) or clock signal. In this way, the logic blocks may be enable of disable by the configuration switches because they would create an open circuit or cut-off the logic block power supply.

[0227] In some embodiments, as shown in Fig. 14, address bus 1402 and data bus 1404 may be in opposite sides of the logic blocks, which are connected in parallel to each one of the buses. In this way, routing of the different on-chip components may be facilitated by the logic block set 1400.

[0228] In some embodiments, each one of the plurality of deactivation switches 1312 couple at least one of the plurality of primary logic blocks with a clock node, and each one of the plurality of activation switches 1314 may be couple at least one of the plurality of redundant blocks with the clock node allowing to connect\disconnect the clock as a simple activation\deactivation mechanism.

[0229] Redundant business logic 1310 of redundant logic block set 1400 allows the designer to choose, based on area and routing, the blocks that are worth duplication. For example, a chip designer may select larger blocks for duplication because larger blocks may be more error prone. Thus, a chip designer may decide to duplicate large logic blocks. On the other hand, a designer may prefer to duplicate smaller logic blocks because they are easily duplicated without a significant loss of space. Moreover, using the configuration in Fig. 14, a designer may easily choose to duplicate logic blocks depending on the statistics of errors per area.

[0230] Fig. 15 is a block diagram for an exemplary logic block 1500, consistent with disclosed embodiments. The logic block may be business logic 1308 and/or redundant business logic 1310. In other embodiments, however, the exemplary logic block may describe memory logic 1306 or other component of memory chip 1300.

[0231] Logic block 1500 presents yet another embodiment where the logic redundancy is used within a small processor pipeline. The logic block 1500 may include a register 1508, a fetch circuit 1504, decoder 1506, and a write-back circuit 1518. In addition, logic block 1500 may include a computation unit 1510 and a duplicated computing unit 1512. However, in other embodiments, logic block 1500 may include other units that do not comprise a controller pipeline but include sporadic processing elements that comprise a required business logic.

[0232] Computation unit 1510 and duplicated computation unit 1512 may include a digital circuit capable of performing digital calculations. For example, computation unit 1510 and duplicated computation unit 1512 may include an arithmetic logic unit (ALU) to perform arithmetic and bitwise operations on binary numbers. Alternatively, computation unit 1510 and duplicated computation unit 5 1512 may include a floating-point unit (FPU), which operates on floating point numbers. In addition, in some embodiments computation unit 1510 and duplicated computation unit 1512 may implement data base related functions like min, max, count, and compare operations, among others.

[0233] In some embodiments, as shown in Fig. 15, computation unit 1510 and duplicated computation unit 1512 may be connected to switching circuits 1514 and 1516. When activated the 10 switching circuits may enable or disable the computing units.

[0234] In logic block 1500, the duplicated computation unit 1512 may replicate the computation unit 1510. Moreover, in some embodiments, register 1508, fetch circuit 1504, decoder 1506, and write-back circuit 1518 (collectively referred to as the local logic units) may be smaller in size than the computation unit 1510. Because larger elements are more prone to issues during fabrication, a designer 15 may decide to replicate larger units (such as computation unit 1510) instead of smaller units (such as the local logic units). Depending on historic yields and error rates, however, a designer may elect to duplicate local logic units additionally or alternatively to large units (or the entire block). For example, computation unit 1510 may be larger, and thus more error prone, than register 1508, fetch circuit 1504, decoder 1506, and write-back circuit 1518. A designer may choose to duplicate computation unit 1510 instead of the 20 other elements in logic block 1500 or the whole block.

[0235] Logic block 1500 may include a plurality of local configuration switches, each one of the plurality of local configuration switches being connected to at least one of the at least one of computation unit 1510 or duplicated computation unit 1512. Local configuration switches may be configured to 25 disable computation unit 1510 and enable duplicated computation unit 1512 when a fault is detected in the computation unit 1510.

[0236] Fig. 16 shows block diagrams of exemplary logic blocks connected with a bus, consistent with disclosed embodiments. In some embodiments, logic blocks 1602 (which may represent memory logic 1306, business logic 1308, or redundant business logic 1310) may be independent of each other, may be connected via a bus, and may be activated externally by addressing them specifically. For 30 example, memory chip 1300 may include many logic blocks, each logic block having an ID number. In other embodiments, however, logic blocks 1602 may represent larger units comprised of a plurality one or more of memory logic 1306, business logic 1308, or redundant business logic 1310.

[0237] In some embodiments, each one of logic blocks 1602 may be redundant with the other logic blocks 1602. This complete redundancy, in which all blocks may operate as primary or redundant 35 blocks, may improve fabrication yields because a designer may disconnect faulty units while maintaining functionality of the overall chip. For example, a designer may have the ability to disable logic areas that are prone to errors but maintain similar computation capabilities because the all duplicate blocks may be connected to the same address and data buses. For example, the initial number of logic blocks 1602 may

greater than a target capability. Then, disabling some logic blocks 1602 would not affect the target capability.

[0238] A bus connected to the logic blocks may include address bus 1614, command lines 1616, and data lines 1618. As shown in Fig. 16, each one of the logic blocks may be connected independently from each line in the bus. In certain embodiments, however, logic blocks 1602 may be connected in a hierarchical structure to facilitate routing. For instance, each line in the bus may be connected to a multiplexer that routes the line to different logic blocks 1602.

[0239] In some embodiments, to allow external access without knowing the internal chip structure, which may change due to enable and disabled units, each one of the logic blocks may include Fused IDs such as fused identification 1604. Fused identification 1604 may include an array of switches (like fuses) that determine an ID and may be connected to a managing circuit. For example, fused identification 1604 may be connected to address manager 1302. Alternatively, fused identification 1604 may be connected to higher memory address units. In these embodiments, fused identification 1604 may be configurable to for a specific address. For example, fused identification 1604 may include a programmable, non-volatile device that determines a final ID based on instructions received from a managing circuit.

[0240] A distributed processor on a memory chip may be designed with the configuration depicted in Fig. 16. A testing procedure executed as BIST at chip wakeup or at factory testing may assign running ID numbers to blocks in the plurality of primary logic blocks (memory logic 1306 and business logic 1308) that pass a testing protocol. A testing procedure may also assign illegal ID numbers to blocks in the plurality of primary logic blocks that do not pass the testing protocol. The test procedure may also assign running ID numbers to blocks in the plurality of redundant blocks (redundant logic block 1310) that pass the testing protocol. Because redundant blocks replace failing primary logic blocks, the blocks in the plurality of redundant blocks assigned running ID numbers may be equal to, or greater than, the blocks in the plurality of primary logic blocks assigned illegal ID numbers, thereby disabling the block. In addition, each one of the plurality of primary logic blocks and each one of the plurality of redundant blocks may include at least one fused identification 1604. Also, as shown in Fig. 16, the bus connecting logic blocks 1602 may include a command line, a data line, and an address line.

[0241] In other embodiments, however, all logic blocks 1602 that are connected to the bus will start disabled and with no ID number. Tested one by one, each good logic block will get a running ID number, and those logic blocks not working will remain with illegal ID, which would disable these blocks. In this manner, redundant logic blocks may improve the fabrication yields by replacing blocks that are known to be defective during the testing process.

[0242] Address bus 1614 may couple a managing circuit to each one of the plurality of memory banks, each one of the plurality of primary logic blocks, and each one of the plurality of redundant blocks. These connections allow the managing circuit to, upon detection of the fault associated with a primary logic blocks (such as business logic 1308), assign an invalid address to the one of the plurality of primary logic blocks and assign a valid address to the one of the plurality of redundant blocks.



[0243] For example, as shown in Fig. 16A, illegal IDs are configured to all logic blocks 1602(a)-(c) (e.g., address 0xFFFF). After testing logic blocks 1602(a) and 1602(c) are verified to be functional while logic block 1602(b) is not functional. In Fig. 16A unshaded logic blocks may represent logic blocks that passed the functionality test successfully, while shaded logic blocks may represent logic blocks that failed the test for functionality. Then, the test procedure changes the illegal IDs to legal IDs for logic blocks that are functional while leaving the illegal IDs for logic blocks that are not functional. As an example, in Fig. 16A, the address for logic blocks 1602(a) and 1602(c) is changed from 0xFFFF to 0x001 and 0x002, respectively. In contrast, the address for logic block 1602(b) remains the illegal address 0xFFFF. In some embodiments, the ID is changed by programming a corresponding fused identification

5  
10 1604.

[0244] Different results from the testing of logic blocks 1602 may result in a different configuration. For example, as shown in Fig. 16B, address manager 1302 may initially assign illegal IDs to all logic blocks 1602 (i.e., 0xFFFF). The testing results, however, may indicate that both logic blocks 1602(a) and 1602(b) are functional. In these cases, testing of logic block 1602(c) may not be necessary because memory chip 1300 may require only two logic blocks. Therefore, to minimize testing resources, logic blocks may be tested only according to the minimum number of functional logic blocks needed by the product definition of 1300, leaving other logic blocks untested. Fig. 16B also shows unshaded logic blocks, which represent tested logic blocks that passed the test for functionality, and shaded logic blocks, which represent untested logic blocks.

15  
20

[0245] In these embodiments, a production tester (external or internal, automatic or manual) or a controller executing a BIST at startup, may change illegal IDs to running IDs for tested logic blocks that are functional while leaving the illegal IDs to untested logic blocks. As an example, in Fig. 16B, the address for logic blocks 1602(a) and 1602(b) is changed from 0xFFFF to 0x001 and 0x002, respectively. In contrast, the address for untested logic block 1602(c) remains with the illegal address 0xFFFF.

25

[0246] Fig. 17 is a block diagram for exemplary units 1702 and 1712 connected in series, consistent with disclosed embodiments. Fig. 17 may represent an entire system or chip. Alternatively, Fig. 17 may represent a block in a chip containing other functional blocks.

[0247] Units 1702 and 1712 may represent complete units that include a plurality of logic blocks such as memory logic 1306 and/or business logic 1308. In these embodiments units 1702 and 1712 may also include elements required to perform operations such as address manager 1302. In other embodiments, however, units 1702 and 1712 may represent logic units such as business logic 1308 or redundant business logic 1310.

30

[0248] Fig. 17 presents embodiments in which units 1702 and 1712 may need to communicate between themselves. In such cases, units 1702 and 1712 may be connected in series. However, a non-working unit may break the continuity between the logic blocks. Therefore, the connection between units may include a bypass option when a unit needs to be disabled due to a defect. The bypass option can also be a part of the bypassed unit itself.

35

[0249] In Fig. 17 units may be connected in series (e.g., 1702(a)-(c)), and a failing unit (e.g., 1702(b)) may be bypassed when it is defective. The units may further be connected in parallel with switching circuits. For example, in some embodiments units 1702 and 1712 may be connected with switching circuits 1722 and 1728, as depicted in Fig. 17. In the example depicted in Fig. 17, unit 1702(b) is defective. For example, unit 1702(b) does not pass a test for a circuit functionality. Therefore, unit 1702(b) may be disabled using, for example, activation switches 1314 (not shown in Figure 17) and/or switching circuit 1722(b) may be activated to bypass unit 1702(b) and sustain the connectivity between logic blocks.

[0250] Accordingly, when a plurality of primary units are connected in series, each one of the plurality of units may be connected in parallel with a parallel switch. Upon detection of a fault associated with the one of the plurality of units, the parallel switch connected to the one of the plurality of units may be activated to connect two of the plurality of units.

[0251] In other embodiments, as shown in Fig. 17, switching circuits 1728 may include a sampling point or more that would cause a cycle or cycles delay maintaining synchronization between different lines of units. When a unit is disabled, shorting the connection between adjacent logic blocks may generate synchronization errors with other calculations. For example, if a task requires data from both A and B lines, and each of A and B is carried by an independent series of units, disabling a unit would cause a desynchronization between the lines that would require further data management. To prevent desynchronizations, sample circuits 1730 may simulate the delay caused by the disabled unit 1712(b). Nonetheless, in some embodiments, the parallel switch may include an anti-fuse instead of a sampling circuit 1730.

[0252] Figure 18 is a block diagram of exemplary units connected in a two-dimension array, consistent with disclosed embodiments. Fig. 18 may represent an entire system or chip. Alternatively, Fig. 18 may represent a block in a chip containing other functional blocks.

[0253] Units 1806 may represent autonomous units that include a plurality of logic blocks such as memory logic 1306 and/or business logic 1308. However, in other embodiments units 1806 may represent logic units such as business logic 1308. Where convenient, discussion of Fig. 18 may refer to elements identified in Fig. 13 (e.g., memory chip 1300) and discussed above.

[0254] As shown in Fig. 18, units may be arranged in a two-dimensional array in which units 1806 (which may include or represent one or more of memory logic 1306, business logic 1308, or redundant business logic 1310) are interconnected via switching boxes 1808 and connection boxes 1810. In addition, in order to control the configuration of the two-dimensional array, the two-dimensional array may include I/O blocks 1804 in the periphery of the two-dimensional array.

[0255] Connection boxes 1810 may be programmable and reconfigurable devices that may respond to signals inputted from the I/O blocks 1804. For example, connection boxes may include a plurality of input pins from units 1806 and may also be connected to switching boxes 1808. Alternatively, connection boxes 1810 may include a group of switches connecting pins of programmable

logic cells with routing tracks, while switching boxes 1808 may include a group of switches connecting different tracks.

[0256] In certain embodiments, connection boxes 1810 and switching boxes 1808 may be implemented with configuration switches such as switches 1312 and 1314. In such embodiments, connection boxes 1810 and switching boxes 1808 may be configured by a production tester or a BIST executed at chip startup.

[0257] In some embodiments, connection boxes 1810 and switching boxes 1808 may be configured after units 1806 are tested for a circuit functionality. In such embodiments, I/O blocks 1804 may be used to send testing signals to units 1806. Depending on the test results, I/O blocks 1804 may send programming signals that configure connection boxes 1810 and switching boxes 1808 in a manner disabling the units 1806 that fail the testing protocol and enabling units 1806 that pass the testing protocol.

[0258] In such embodiments, the plurality of primary logic blocks and the plurality of redundant blocks may be disposed on the substrate in a two-dimensional grid. Therefore, each one of the plurality of primary units 1806 and each one of the plurality of redundant blocks, such as redundant business logic 1310, may be interconnected with switching boxes 1808, and an input block may be disposed in the periphery of each line and each column of the two-dimensional grid.

[0259] Fig. 19 is a block diagram for exemplary units in a complex connection, consistent with disclosed embodiments. Fig. 19 may represent an entire system. Alternatively, Fig. 19 may represent a block in a chip containing other functional blocks.

[0260] The complex connection of Fig. 19 includes units 1902(a)-(f) and configuration switches 1904(a)-(h). Units 1902 may represent autonomous units that include a plurality of logic blocks such as memory logic 1306 and/or business logic 1308. However, in other embodiments units 1902 may represent logic units such as memory logic 1306, business logic 1308, or redundant business logic 1310. Configuration switches 1904 may include any of deactivation switches 1312 and activation switches 1314.

[0261] As shown in Fig. 19, the complex connection may include units 1902 in two planes. For example, the complex connection may include two independent substrates separated in the z-axis. Alternatively, or additionally, units 1902 may be arranged in two surfaces of a substrate. For example, with the objective to reduce the area of memory chip 1300, substrate 1301 may be arranged in two overlapping surfaces and connected with configuration switches 1904 arranged in three dimensions. Configuration switches may include deactivation switches 1312 and/or activation switches 1314.

[0262] A first plane of the substrate may include “main” unit 1902. These blocks may be enabled by default. In such embodiments, a second plain may include “redundant” unit 1902. These units may be disabled by default.

[0263] In some embodiments, configuration switches 1904 may include anti-fuses. Thus, after testing of units 1902, the blocks may be connected in a tile of functional units by switching certain anti-fuses to “always-on” and disable selected units 1902, even if they are in a different plane. In the example

presented in Fig. 19, one of the 'main' units (unit 1902(e)) is not working. Fig. 19 may represent nonfunctional or untested blocks as shaded blocks while tested or functional blocks may be unshaded. Therefore, configuration switches 1904 are configured so one of the logic blocks in a different plane (e.g., unit 1902(f)) becomes active. In this way even though one of the main logic blocks was defective, the memory chip is still working by replacing a spare logic unit.

5 [0264] Fig. 19 additionally shows that one of the units 1902 (i.e., 1902(c)) in the second plane is not tested or enabled because the main logic blocks are functional. For example, in Fig. 19, both main units 1902(a) and 1902(d) passed a test for functionality. Thus, units 1902(c) was not tested or enabled. Therefore, Fig. 19 shows the ability to specifically select the logic blocks that become active depending  
10 on testing results.

[0265] In some embodiments, as shown in Fig. 19, not all units 1902 in a first plain may have a corresponding spare or redundant blocks. However, in other embodiments, all units may be redundant with each other for complete redundancy where all units are both primary or redundant. In addition, while some implementations may follow the star network topology depicted in Fig. 19, other implementation  
15 may use parallel connections, serial connections, and/or couple the different elements with configuration switches in parallel or in series.

[0266] Fig. 20 is an exemplary flowchart illustrating a redundant block enabling process 2000, consistent with disclosed embodiments. The enabling process 2000 may be implemented for memory chip 1300 and specially for DRAM memory chips. In some embodiments, process 2000 may include steps of  
20 testing each one of a plurality of logic blocks on the substrate of the memory chip for at least one circuit functionality, identifying faulty logic blocks in the plurality of primary logic blocks based on the testing results, testing at least one redundant or additional logic block on the substrate of the memory chip for the at least one circuit functionality, disabling the at least one faulty logic block by applying an external signal to a deactivation switch, and enabling the at least one redundant block by applying the external  
25 signal to an activation switch, the activation switch being connected with the at least one redundant block and being disposed on the substrate of the memory chip. The description of Fig. 20 below further elaborates on each step of process 2000.

[0267] Process 2000 may include testing a plurality of logic blocks (step 2002), such as business block 1308 and a plurality of redundant blocks (e.g., redundant business block 1310). The testing may be  
30 before packaging using, for example, probing stations for on-wafer testing. Step 2000, however, may also be performed after packaging.

[0268] The testing in step 2002 may include applying a finite sequence of testing signals to every logic block in memory chip 1300 or a subset of logic blocks in memory chip 1300. The testing signals may include requesting a computation that is expected to yield a 0 or a 1. In other embodiments,  
35 the testing signal may request reading a specific address in a memory bank or writing in a specific memory bank.

[0269] Testing techniques may be implemented to test the response of the logic blocks under iterative processes in step 2002. For example, the test may involve testing logic blocks by transmitting

instructions to write data in a memory bank and then verifying the integrity of the written data. In some embodiments, the testing may include repeating the algorithm with data inversed.

[0270] In alternative embodiments, the testing of step 2002 may include running a model of the logic blocks to generate a target memory image based on a set of testing instructions. Then, the same sequence of instructions may be executed to the logic blocks in the memory chip, and the results may be recorded. The residual memory image of the simulation may also be compared to the image taken from the test, and any mismatch may be flagged as a failure.

[0271] Alternatively, in step 2002, testing may include shadow modeling, where a diagnostic is generated but the results are not necessarily predicted. Instead, the test using shadow modeling may be run in parallel on both the memory chip and a simulation. For example, when the logic blocks in the memory chip complete an instruction or task, the simulation may be signaled to execute the same instruction. Once the logic blocks in the memory chip finalize the instructions, the two models' architectural states may be compared. If there is a mismatch, then a failure is flagged.

[0272] In some embodiments, all logic blocks (including, e.g., each one of memory logic 1306, business logic 1308, or redundant business logic 1310) may be tested in step 2002. In other embodiments, however, only subsets of the logic blocks may be tested in different testing rounds. For example, in a first round of testing only memory logic 1306 and associated blocks may be tested. In a second round, only business logic 1308 and associated blocks may be tested. In a third round, depending on the results of the first two rounds, logic blocks associated with redundant business logic 1310 may be tested.

[0273] Process 2000 may continue to step 2004. In step 2004, faulty logic blocks may be identified, and faulty redundant blocks may also be identified. For example, logic blocks that do not pass the testing of step 2002 may be identified as faulty blocks in step 2004. In other embodiments, however, only certain faulty logic blocks may be initially identified. For example, in some embodiments, only logic blocks associated with business logic 1308 may be identified, and faulty redundant blocks are only identified if they are required for substituting a faulty logic block. In addition, identifying faulty blocks may include writing on a memory bank or a nonvolatile memory the identification information of the identified faulty blocks.

[0274] In step 2006, faulty logic blocks may be disabled. For example, using a configuration circuit, the faulty logic blocks may be disabled by disconnecting them from clock, ground, and/or power nodes. Alternatively, faulty logic blocks may be disabled by configuring connection boxes in an arrangement that avoids the logic blocks. Yet, in other embodiments, faulty logic blocks may be disabled by receiving an illegal address from address manager 1302.

[0275] In step 2008, redundant blocks that duplicate the faulty logic blocks may be identified. To support the same capabilities of the memory chips even though some logic blocks have failed, in step 2008, redundant blocks that are available and can duplicate faulty logic blocks may be identified. For example, if a logic block that performs multiplications of vectors is determined to be faulty, in step 2008, an address manager 1302 or an on-chip controller may identify an available redundant logic block that also performs multiplication of vectors.

[0276] In step 2010, the redundant blocks identified in step 2008 may be enabled. In contrast to the disable operation of step 2006, in step 2010, the identified redundant blocks may be enabled by connecting them to clock, ground, and/or power nodes. Alternatively, identified redundant blocks may be enabled by configuring connection boxes in an arrangement that connects the identified redundant blocks. Yet, in other embodiments, identified redundant blocks may be enabled by receiving a running address at the test procedure execution time.

[0277] Fig. 21 is an exemplary flow chart illustrating an address assignment process 2100, consistent with disclosed embodiments. The address assignment process 2100 may be implemented for memory chip 1300 and specially for a DRAM memory chips. As described in relation to Fig. 16, in some embodiments, logic blocks in memory chip 1300 may be connected to a data bus and have an address identification. Process 2100 describes an address assignment method that disables faulty logic blocks and enables logic blocks that pass a test. The steps described in process 2100 will be described as being performed by a production tester or a BIST executed at chip startup; however, other components of memory chip 1300 and/or external devices may also perform one or more steps of process 2100.

[0278] In step 2102, the tester may disable all logic and redundant blocks by assigning an illegal identification to each logic block at a chip level.

[0279] In step 2104, the tester may execute a testing protocol of a logic block. For example, the tester may run testing methods described in step 2002 for one or more of the logic blocks in memory chip 1300.

[0280] In step 2106, depending on the results of the test in step 2104, the tester may determine whether the logic block is defective. If the logic block is not defective (step 2106: no), address manager may assign a running ID to the tested logic block in step 2108. If the logic block is defective (step 2106: yes), address manager 1302 may leave the illegal ID for the defective logic block in step 2110.

[0281] In step 2112, address manager 1302 may select a redundant logic block that replicates the defective logic block. In some embodiments, the redundant logic block that replicates the defective logic block may have the same components and connections to the defective logic blocks. In other embodiments, however, the redundant logic block may have different components and/or connections to the defective logic blocks but be able to perform an equivalent operation. For example, if the defective logic block is designed to perform multiplication of vectors, the selected redundant logic block would also be capable of performing multiplication of vectors, even if it does not have the same architecture as the defective unit.

[0282] In step 2114, address manager 1302 may test the redundant block. For instance, the tester may apply the testing techniques applied in step 2104 to the identified redundant block.

[0283] In step 2116, based on the results of testing in step 2114, the tester may determine whether the redundant block is defective. In step 2118, if the redundant block is not defective (step 2116: no), the tester may assign a running ID to the identified redundant block. In some embodiments, process 2100 may return to step 2104 after step 2118, creating an iteration loop to test all logic blocks in the memory chip.

[0284] If the tester determines the redundant block is defective (step 2116: yes), in step 2120, the tester may determine if additional redundant blocks are available. For example, the tester may query a memory bank with information regarding available redundant logic blocks. If redundant logic blocks are available (step 2120: yes), the tester may return to step 2112 and identify a new redundant logic block replicating the defective logic block. If redundant logic blocks are not available (step 2120: no), in step 2122, the tester may generate an error signal. The error signal may include information of the defective logic block and the defective redundant block.

#### Coupled memory banks

[0285] The presently disclosed embodiments also include a distributed high-performance processor. The processor may include a memory controller that interfaces memory banks and processing units. The processor may be configurable to expedite delivery of data to the processing units for calculations. For example, if a processing unit requires two data instances to perform a task, the memory controller may be configured so communication lines independently provide access to the information from two data instances. The disclosed memory architecture seeks to minimize hardware requirements that are associated with complex cache memory and complex register files schemes. Normally, processor chips include cache hierarchies that allow cores to work directly with registers. However, the cache operations require significant die area and consume additional power. The disclosed memory architecture avoids the use of a cache hierarchy by adding logic components in the memory.

[0286] The disclosed architecture also enables strategic (or even optimized) placement of data in memory banks. Even if the memory banks have a single port and high latency, the disclosed memory architecture may enable high performance and avoid memory accessing bottlenecks by strategically positioning data in different blocks of memory banks. With the goal of providing a continuous stream of data to the processing units, a compilation optimization step may determine how data should be stored in memory banks for specific or generic tasks. Then, the memory controller, which interfaces processing units and memory banks, may be configured to grant access to specific processing units when they require data to perform operations.

[0287] The configuration of the memory chip may be performed by a processing unit (e.g., a configuration manager) or an external interface. The configuration may be also written by a compiler or other SW tool. In addition, the configuration of the memory controller may be based on the available ports in the memory banks and the organization of data in the memory banks. Accordingly, the disclosed architecture may provide processing units with a constant flow of data or simultaneous information from different memory blocks. In this way, computation tasks within the memory may be quickly processed by avoiding latency bottlenecks or cache memory requirements.

[0288] Moreover, data stored in the memory chip may be arranged based on compilation optimization steps. The compilation may allow for building of processing routines in which the processor efficiently assigns tasks to processing units without memory latency associated delays. The compilation may be performed by a compiler and transmitted to a host connected to an external interface in the substrate. Normally, high latency for certain access patterns and/or low numbers of ports would result in

data bottlenecks for processing units requiring the data. The disclosed compilation, however, may position data in memory banks in a way that enables processing units to continuously receive data even with disadvantageous memory types.

[0289] Furthermore, in some embodiments, a configuration manager may signal required  
5 processing units based on computations that are required by a task. Different processing units or logic blocks in the chip may have specialized hardware or architectures for different tasks. Therefore, depending on the task that will be performed, a processing unit, or a group of processing units, may be selected to perform the task. The memory controller on the substrate may be configurable to route data, or grant access, according to the selection of processing subunits to improve data transfer rates. For  
10 example, based on the compilation optimization and the memory architecture, processing units may be granted access to memory banks when they are required to perform a task.

[0290] Moreover, the chip architecture may include on-chip components that facilitate transfer of data by reducing the time required to access data in the memory banks. Therefore, the present disclosure describes chip architecture(s), along with a compilation optimization step, for a high-  
15 performance processor capable of performing specific or generic tasks using simple memory instances. The memory instances may have high latency in random access and/or low number of ports, such as those used in a DRAM device or other memory-oriented technologies, but the disclosed architecture may overcome these shortcomings by enabling a continuous (or nearly continuous) flow of data from memory banks to processing units.

[0291] In this application, simultaneous communication may refer to communication within a  
20 clock cycle. Alternatively, simultaneous communication may refer to sending information within a predetermine amount of time. For example, simultaneous communication may refer to communication within a few nanoseconds.

[0292] Fig. 22 provides block diagrams for exemplary processing devices, consistent with  
25 disclosed embodiments. Fig. 22A shows a first embodiment of a processing device 2200 in which a memory controller 2210 connects a first memory block 2202 and a second memory block 2204 using multiplexers. Memory controller 2210 may also connect at least a configuration manager 2212, a logic block 2214, and multiple accelerators 2216(a)-(n). Fig. 22B shows a second embodiment of processing device 2200 in which memory controller 2210 connects memory blocks 2202 and 2204 using a bus that  
30 connects memory controller 2210 with at least a configuration manager 2212, a logic block 2214, and multiple accelerators 2216(a)-(n). In addition, host 2230 may be external and connected to processing device 2200 through, for example, an external interface.

[0293] Memory blocks 2202 and 2204 may include a DRAM mats or group of mats, DRAM  
banks, MRAM\ PRAM\ RERAM\ SRAM units, Flash mats, or other memory technologies. Memory  
35 blocks 2202 and 2204 may alternatively include non-volatile memories, a flash memory device, a Resistive Random Access Memory (ReRAM) device, or a Magnetoresistive Random Access Memory (MRAM) device.



[0294] Memory blocks 2202 and 2204 may additionally include a plurality of memory cells arranged in rows and columns between a plurality of word lines (not shown) and a plurality of bit lines (not shown). The gates of each row of memory cells may be connected to a respective one of the plurality of word lines. Each column of memory cells may be connected to a respective one of the plurality of bit lines.

[0295] In other embodiments, a memory area (including memory blocks 2202 and 2204) is built from simple memory instances. In this application, the term “memory instance” may be used interchangeably with the term “memory block.” The memory instances (or blocks) may have poor characteristics. For example, the memories may be only one port memories and may have high random-access latency. Alternatively, or additionally, the memories may be inaccessible during column and line changes and face data access problems related to, for example, capacity charging and/or circuitry setups. Nonetheless, the architecture presented in Fig. 22 still facilitates parallel processing in the memory device by allowing dedicated connections between memory instances and processing units and arranging the data in a certain manner that takes the characteristics of the blocks into account.

[0296] In some device architectures, memory instances may include several ports, facilitating the parallel operations. Nonetheless, in such embodiments, the chip may still achieve an improved performance when data is compiled and organized based on the chip architecture. For example, a compiler may improve the efficiency of access in the memory area by providing instructions and organizing data placement, so it can be readily access even using one-port memories.

[0297] Furthermore, memory blocks 2202 and 2204 may be multiple types for memory in a single chip. For example, memory blocks 2202 and 2204 may be eFlash and eDRAM. Also, memory blocks may include DRAM with instances of ROM.

[0298] Memory controller 2210 may include a logic circuit to handle the memory access and return the results to the rest of the modules. For example, memory controller 2210 may include an address manager and selection devices, such as multiplexers, to route data between the memory blocks and processing units or grant access to the memory blocks. Alternatively, Memory controller 2210 may include double data rate (DDR) memory controllers used to drive DDR SDRAM, where data is transferred on both rising and falling edges of the system's memory clock.

[0299] In addition, memory controller 2210 may constitute Dual Channel memory controllers. The incorporation of dual channel memory may facilitate control of parallel access lines by memory controller 2210. The parallel access lines may be configured to have identical lengths to facilitate synchronization of data when multiple lines are used in conjunction. Alternatively, or additionally, the parallel access lines may allow access of multiple memory ports of the memory banks.

[0300] In some embodiments processing device 2200 may include one or more muxes that may be connected to processing units. The processing units may include configuration manager 2212, logic block 2214, and accelerators 2216, which may be connected directly to the mux. Also, memory controller 2210 may include at least one data input from a plurality of memory banks or blocks 2202 and 2204 and at least one data output connected to each one of the plurality of processing units. With this configuration,

memory controller 2210 may simultaneously receive data from memory banks or memory blocks 2202 and 2204 via the two data inputs, and simultaneously transmits data received via to the at least one selected processing unit via the two data outputs. In some embodiments, however, the at least one data input and at least one data output may be implemented in a single port allowing only read or write operations. In such embodiments, the single port may be implemented as a data bus including data, address, and command lines.

[0301] Memory controller 2210 may be connected to each one of the plurality of memory blocks 2202 and 2204, and may also connect to processing units via, for example, a selection switch. Also processing units on the substrate, including configuration manager 2212, logic block 2214, and accelerators 2216, may be independently connected to memory controller 2210. In some embodiments, configuration manager 2212 may receive an indication of a task to be performed and, in response, configure memory controller 2210, accelerators 2216, and/or logic blocks 2214 according to a configuration stored in memory or supplied externally. Alternatively, memory controller 2210 may be configured by an external interface. The task may require at least one computation that may be used to select at least one selected processing unit from the plurality of processing units. Alternatively, or additionally, the selection may be based at least in part upon a capability of the selected processing unit for performing the at least one computation. In response, memory controller 2210 may grant access to the memory banks, or route data between the at least one selected processing unit and at least two memory banks, using dedicated buses and/or in a pipelined memory access.

[0302] In some embodiments, first memory block 2202 of at least two memory blocks may be arranged on a first side of the plurality of processing units; and second memory bank 2204 of the at least two memory banks may be arranged on a second side of the plurality of processing units opposite to the first side. Further, a selected processing unit to perform the task, for instance accelerator 2216(n), may be configured to access second memory bank 2204 during a clock cycle in which a communication line is opened to the first memory bank or first memory block 2202. Alternatively, the selected processing unit may be configured to transfer data to second memory block 2204 during a clock cycle in which a communication line is opened to first memory block 2202.

[0303] In some embodiments, memory controller 2210 may be implemented as an independent element, as shown in Figure 22. In other embodiments, however, memory controller 2210 may be embedded in the memory area or may be disposed along accelerators 2216(a)-(n).

[0304] A processing area in processing device 2200 may include configuration manager 2212, logic block 2214, and accelerators 2216(a)-(n). Accelerators 2216 may include multiple processing circuits with pre-defined functions and may be defined by a specific application. For example, an accelerator may be a vector multiply accumulate (MAC) unit or a Direct Memory Access (DMA) unit handling memory moving between modules. Accelerators 2216 may also be able to calculate their own address and request the data from memory controller 2210 or write data to it. For example, configuration manager 2212 may signal at least one of accelerators 2216 that he can access the memory bank. Then accelerators 2216 may configure memory controller 2210 to route data or grant access to themselves. In

addition, accelerators 2216 may include at least one arithmetic logic unit, at least one vector handling logic unit, at least one string compare logic unit, at least one register, and at least one direct memory access.

[0305] Configuration manager 2212 may include digital processing circuits to configure  
5 accelerators 2216 and instructs execution of tasks. For example, configuration manager 2212 may be  
connected to memory controller 2210 and each one of the plurality of accelerators 2216. Configuration  
manager 2212 may have its own dedicated memory to hold the configurations of accelerators 2216.  
Configuration manager 2212 may use the memory banks to fetch commands and configurations via  
10 memory controller 2210. Alternatively, configuration manager 2212 may be programmed through an  
external interface. In certain embodiments, configuration manager 2212 may be implemented with an on-  
chip reduced instruction set computer (RISC) or an on-chip complex CPU with its own cache hierarchy.  
In some embodiments, configuration manager 2212 may also be omitted and the accelerators can be  
configured through an external interface.

[0306] Processing device 2200 may also include an external interface (not shown). The external  
15 interface allows access to the memory from an upper level, such a memory bank controller which receives  
the command from external host 2230 or on-chip main processor or access to the memory from external  
host 2230 or on-chip main processor. The external interface may allow programming of the configuration  
manager 2212 and the accelerators 2216 by writing configurations or code to the memory via memory  
controller 2210 to be used later by configuration manager 2212 or the units 2214 and 2216 themselves.  
20 The external interface, however, may also directly program processing units without being routed through  
memory controller 2210. In case configuration manager 2212 is a microcontroller, configuration manager  
2212 may allow loading of code from a main memory to the controller local memory via the external  
interface. Memory controller 2210 may be configured to interrupt the task in response to receiving a  
request from the external interface.

[0307] The external interface may include multiple connectors associated with logic circuits that  
25 provide a glue-less interface to a variety of elements on the processing device. The external interface may  
include: Data I/O Inputs for data reads and output for data writes; External address outputs; External CE0  
chip select pins; Active-low chip selectors; Byte enable pins; a pin for wait states on the memory cycle; a  
Write enable pin; an Output enable-active pin; and read-write enable pin. Therefore, the external interface  
30 has the required inputs and outputs to control processes and obtain information from the processing  
device. For example, the external interface may conform to JEDEC DDR standards. Alternatively, or  
additionally, external interface may conform to other standards such as SPI/OSPI or UART.

[0308] In some embodiments, the external interface may be disposed on the chip substrate and  
may be connected external host 2230. The external host may gain access to memory blocks 2202 and  
35 2204, memory controller 2210, and processing units via the external interface. Alternatively, or  
additionally, external host 2230 may read and write to the memory or may signal configuration manager  
2212, through read and write commands, to perform operations such as starting a process and/or stopping  
a process. In addition, external host 2230 may configure the accelerators 2216 directly. In some

embodiments, external host 2230 be able to perform read/write operations directly on memory blocks 2202 and 2204.

[0309] In some embodiments, configuration manager 2212 and accelerators 2216 may be configured to connect the device area with the memory area using direct buses depending on the target task. For example, a subset of accelerators 2216 may connect with memory instances 2204 when the subset of accelerators has the capability to perform computations required to execute the task. By doing such a separation, it is possible to assure that dedicated accelerators get the bandwidth (BW) needed to memory blocks 2202 and 2204. Moreover, this configuration with dedicated buses may allow splitting a large memory to smaller instances or blocks because connecting memory instances to memory controller 2210 allows quick access to data in different memories even with high row latency time. To achieve the parallelization of connection, memory controller 2210 may be connected to each of the memory instances with data, address, and/or control buses.

[0310] The above-discussed inclusion of memory controller 2210 may eliminate the requirement of a cache hierarchy or complex register file in the processing device. Although the cache hierarchy can be added to give added capabilities, the architecture in processing device processing device 2200 may allow a designer to add enough memory blocks or instances based on the processing operations and manage the instances accordingly without a cache hierarchy. For example, the architecture in processing device processing device 2200 may eliminate requirements of a cache hierarchy by implementing a pipelined memory access. In the pipelined memory access, processing units may receive a sustaining flow of data in every cycle certain data lines may be opened (or activated) while other data lines receive or transmit data. The sustained flow of data using independent communication lines may allow an improved execution speed and minimum latency due to line changes.

[0311] Moreover, the disclosed architecture in Fig. 22 enables a pipelined memory access it may be possible to organize data in a low number of memory blocks and save power losses caused by line switching. For example, a In some embodiments, a compiler may communicate host 2230 the organization of, or a method to organize, data in memory banks to facilitate access to data during a given task. Then, configuration manager 2212 may define which memory banks, and in some cases which ports of the memory banks, may be accessed by the accelerators. This synchronization between the location of data in memory banks and the access method to data, improves computing tasks by feeding data to the accelerators with minimum latency. For example, in embodiments in which configuration manager 2212 includes a RISC\CPU, the method may be implemented in offline software (SW) and then the configuration manager 2212 may be programmed to execute the method. The method may be developed in any language executable by RISC/CPU computers and may be executed on any platform. The inputs of the method may include configuration of the memories behind memory controller and the data itself along with the pattern of memory accesses. In addition, the method may be implemented in a language or machine language specific to the embodiment and may also be just a series of configuration values in binary or text.

[0312] As discussed above, in some embodiments, a compiler may provide instructions to host 2230 for organizing data in memory blocks 2202 and 2204 in preparation of a pipelined memory access. The pipelined memory access may generally include steps of: receiving a plurality of addresses of a plurality of memory banks or memory blocks 2202 and 2204; accessing the plurality of memory banks according to the received addresses using independent data lines; supplying data from a first address through a first communication line to at least one of the plurality of processing units and opening a second communication line to a second address, the first address being in a first memory bank of the plurality of memory banks, the second address being in second memory bank 2204 of the plurality of memory banks; and supplying data from the second address through the second communication line to the at least one of the plurality of processing units and opening a third communication line to a third address in the first memory bank in the first line within a second clock cycle. In some embodiments, the pipelined memory access may be executed with two memory blocks being connected to a single port. In such embodiments, memory controller 2210 may hide the two memory blocks behind a single port but transmit data to the processing units with the pipelined memory access approach.

[0313] In some embodiments, a compiler can run on host 2230 before executing a task. In such embodiments, the compiler may be able to determine a configuration of data flow based on the architecture of the memory device since the configuration would be known to the compiler.

[0314] In other embodiments, if the configuration of memory blocks 2204 and 2202 is unknown at offline time, the pipelined method can run on host 2230 which may arrange data in memory blocks before starting calculations. For example, host 2230 may directly write data in memory blocks 2204 and 2202. In such embodiments, processing units, such as configuration manager 2212 and memory controller 2210 may not have information regarding required hardware until run time. Then, it may be necessary to delay the selection of an accelerator 2216 until a task starts running. In these situations, the processing units or memory controller 2210 may randomly select an accelerator 2216 and create a test data access pattern, which may be modified as the task is executed.

[0315] Nonetheless, when the task is known in advance, a compiler may organize data and instructions in memory banks for host 2230 to provide to a processing unit, such as configuration manager 2212, to set signal connections that minimize access latency. For example, in some cases  $n$  words may be needed at the same time by accelerators 2216. However, each memory instance supports retrieving only  $m$  words at a time, where “ $m$ ” and “ $n$ ” are integers and  $m < n$ . Thus, the compiler may place the needed data across different memory instances or blocks facilitating data access. Also, to avoid line miss latencies, a host may split data in different lines of different memory instances if processing device 2200 includes multiple memory instances. The division of data may allow accessing the next line of data in the next instance while still using data from the current instance.

[0316] For example, accelerator 2216(a) may be configured to multiply two vectors. Each one of the vectors may be stored in independent memory blocks, such as memory blocks 2202 and 2204, and each vector may include multiple words. Therefore, to complete a task requiring a multiplication by accelerator 2216(a), it may be necessary to access the two memory blocks and retrieve multiple words.

However, in some embodiments, memory blocks only allow access of one word per clock cycle. For instance, memory blocks may have a single port. In these cases, to expedite data transmittal during an operation, a compiler may organize the words composing vectors in different memory blocks allowing parallel and/or simultaneous reading of the words. In these situations, a compiler may store words in memory blocks that have a dedicated line. For instance, if each vector includes two words and memory controller has direct access to four memory blocks, a compiler may arrange data in four memory blocks, each one transmitting a word and expediting data delivery. Moreover, in embodiments when memory controller 2210 may have more than a single connection to each memory block, the compiler may instruct configuration manager 2212 (or other processing unit) to access ports specific ports. In this way, processing device 2200 may perform a pipelined memory access, continuously providing data to processing units by simultaneously loading words in some lines and transmitting data in other lines. Thus, this pipelined memory access avoid may avoid latency issues.

[0317] Fig. 23 is a block diagram of an exemplary processing device 2300, consistent with disclosed embodiments. The block diagram shows a simplified processing device 2300 displaying a single accelerator in the form of MAC Unit 2302, configuration manager 2304 (equivalent or similar to configuration manager 2212), memory controller 2306 (equivalent or similar to memory controller 2210), and a plurality of memory blocks 2308(a)-(d).

[0318] In some embodiments, MAC unit 2302 may be a specific accelerator for processing a particular task. By way of example, the processing device 2300 may be tasked with 2D-convolutions. Then, configuration manager 2304 can signal an accelerator that has the appropriate hardware to perform calculations associated with the task. For instance, MAC unit 2302 may have four internal incrementing counters (logical adders and registers to manage the four loops needed by a convulsion calculation) and a multiply accumulate unit. Configuration manager 2304 may signal MAC unit 2302 to process incoming data and execute the task. Configuration manager 2304 may transmit an indication to MAC unit 2302 to execute the task. In these situations, MAC unit 2302 may iterate over calculated addresses, multiply the numbers, and accumulate them to an internal register.

[0319] In some embodiments, configuration manager 2304 may configure the accelerators while memory controller 2306 grants access to blocks 2308 and MAC unit 2302 using dedicated buses. In other embodiments, however, memory controller 2306 can directly configure the accelerators based on instructions received from configuration manger 2304 or an external interface. Alternatively, or additionally, configuration manager 2304 can pre-load a few configurations and allow the accelerator to iteratively run on different addresses with different sizes. In such embodiments, configuration manager 2304 may include a cache memory that stores a command before it is transmitted to at least one of the plurality of processing units, such as accelerators 2216. However, in other embodiments configuration manager 2304 may not include a cache.

[0320] In some embodiments, configuration manager 2304 or memory controller 2306 may receive addresses that need to be accessed for a task. Configuration manager 2304 or memory controller 2306 may check a register to determine whether the address is already in a loaded line to one of memory

blocks 2308. If so, memory controller 2306 may read the word from memory block 2308 and pass it to the MAC unit 2302. If the address is not in a loaded line, configuration manager 2304 may request memory controller 2306 may load the line and signal MAC unit 2302 to delay until it is retrieved.

[0321] In some embodiments, as shown in Fig. 23, memory controller 2306 may include two  
5 inputs form two independent addresses. But if more than two addresses should be accessed simultaneously, and these addresses are in a single memory block (for example it is only in of memory blocks 2308(a)), memory controller 2306 or configuration manager 2304 may raise an exception. Alternatively, configuration manager 2304 may return invalid data signal when the two addresses can only be access through a single line. In other embodiments, the unit may delay the process execution until  
10 it is possible to retrieve all needed data. This may diminish the overall performance. Nonetheless, a compiler may be able to find a configuration and data placement that would prevent delays.

[0322] In some embodiments, a compiler may create a configuration or instruction set for processing device 2300 that may configure configuration manager 2304 and memory controller 2306 and accelerator 2302 to handle situations in which multiple addresses need to be accessed from a single  
15 memory block but the memory block has one port. For instance, a compiler may re-arrange data in memory blocks 2308 such that processing units may access multiple lines in memory blocks 2308.

[0323] In addition, memory controller 2306 may also work simultaneously on more than one input at the same time. For example, memory controller 2306 may allow accessing one of memory blocks 2308 through one port and supplying the data while receiving a request from a different memory block in  
20 another input. Therefore, this operation may result in and accelerator 2216 tasked with the exemplary 2D-convolutions receiving data from dedicated lines of communication with the pertinent memory blocks.

[0324] Additionally, or alternatively, memory controller 2306 or a logic block may hold refresh counters for every memory block 2308 and handle the refresh of all lines. Having such a counter allows memory controller 2306 to slip in the refresh cycles between dead access times from the devices.

[0325] Furthermore, memory controller 2306 may be configurable to perform the pipelined memory access, receiving addresses and opening lines in memory blocks before supplying the data. The pipelined memory access may provide data to processing units without interruption or delayed clock cycles. For example, while memory controller 2306 or one of the logic blocks access data with the right line in Fig. 23, it may be transmitting data in the left line. These methods will be explained in greater  
30 detail in connection to Figure 26.

[0326] In response to the required data, processing device 2300 may use multiplexors and/or other switching devices to choose which device gets serviced to perform a given task. For example, configuration manager 2304 may configure multiplexers so at least two data lines reach the MAC unit 2302. In this way, a task requiring data from multiple addresses, such as 2D-convolutions, may be  
35 performed faster because the vectors or words requiring multiplication during convolution can reach the processing unit simultaneously, in a single clock. This data transferring method may allow the processing units, such as accelerators 2216, to quickly output a result.

[0327] In some embodiments, configuration manager 2304 may be configurable to execute processes based on priority of tasks. For example, configuration manager 2304 can be configured to let a running process finish without any interruptions. In that case, configuration manager 2304 may provide an instruction or configurations of a task to accelerators 2216, let them run uninterrupted, and switch multiplexers only when the task is finished. However, in other embodiments, configuration manager 2304 may interrupt a task and reconfigure data routing when it receives a priority task, such a request from an external interface. Nevertheless, with enough memory blocks 2308, memory controller 2306 may be configurable to route data, or grant access, to processing units with dedicated lines that do not have to be changed until a task is completed. Moreover, in some embodiments, all devices may be connected by buses to the entries of configuration manager 2304, and the devices may manage access between themselves and the buses (e.g., using the same logic as a multiplexer). Therefore, memory controller 2306 may be directly connected to a number of memory instances or memory blocks.

[0328] Alternatively, memory controller 2306 may be connected directly to memory sub-instances. In some embodiments, each memory instance or block can be built from sub-instances (for example, DRAM may be built from mats with independent data lines arranged in multiple sub-blocks). Further, the instances may include at least one of DRAM mats, DRAM, banks, flash mats, or SRAM mats or any other type of memory. Then, memory controller 2306 may include dedicated lines to address sub-instances directly to minimize latency during a pipelined memory access.

[0329] In some embodiments, memory controller 2306 may also hold the logic needed for a specific memory instance (such as row\col decoders, refresh logic, etc.) and memory blocks 2308 may handle its own logic. Therefore, memory blocks 2308 may get an address and generate commands for return\write data.

[0330] Fig. 24 depicts exemplary memory configuration diagrams, consistent with disclosed embodiments. In some embodiments, a compiler generating code or configuration for processing device 2200 may perform a method to configure loading from memory blocks 2202 and 2204 by pre-arranging data in each block. For example, a compiler may prearrange data so each word required for a task is correlated to a line of memory instance or memory block(s). But for tasks that require more memory blocks than the one available in processing device 2200, a compiler may implement methods of fitting data in more than one memory location of each memory block. The compiler may also store data in sequence and evaluate the latency of each memory block to avoid line miss latency. In some embodiments, the host may be part of a processing unit, such as configuration manager 2212, but in other embodiments the compiler host may be connected to processing device 2200 via an external interface. In such embodiments, the host may run compiling functions, such as the ones described for the compiler.

[0331] In some embodiments, configuration manager 2212 may be a CPU or a micro controller (uC). In such embodiments, configuration manager 2212 may have to access the memory to fetch commands or instructions placed in the memory. A specific compiler may generate the code and place it in the memory in a manner that allows for consecutive commands to be stored in the same memory line and across a number of memory banks to allow for the pipelined memory access also on the fetched



command. In these embodiments, configuration manager 2212 and memory controller 2210 may be capable of avoiding row latency in linear execution by facilitating the pipelined memory access.

[0332] The previous case of linear execution of a program described a method for a compiler to recognize and place the instructions to allow for pipelined memory execution. However other software structures may be more complex and would require the compiler to recognize them and act accordingly. For example, in case a task requires loops and branches, a compiler may place all the loop code inside a single line so that the single line can be looped without line opening latency. Then, memory controller 2210 may not need to change lines during an execution.

[0333] In some embodiments, configuration manager 2212 may include internal caching or small memory. The internal caching may store commands that are executed by configuration manager 2212 to handle branches and loops. For example, commands in internal caching memory may include instructions to configure accelerators for accessing memory blocks.

[0334] Fig. 25 is an exemplary flowchart illustrating a possible memory configuration process 2500, consistent with disclosed embodiments. Where convenient in describing memory configuration process 2500, reference may be made to the identifiers of elements depicted in Fig. 22 and described above. In some embodiments, process 2500 may be executed by a compiler that provides instructions to a host connected through an external interface. In other embodiments, process 2500 may be executed by components of processing device 2200, such as configuration manager 2212.

[0335] In general, process 2500 may include determining a number of words required simultaneously to perform the task; determining a number of words that can be accessed simultaneously from each one of the plurality of memory banks; and dividing the number of words required simultaneously between multiple memory banks when the number of words required simultaneously is greater than the number of words that can be accessed simultaneously. Moreover, dividing the number of words required simultaneously may include executing a cyclic organization of words and sequentially assigning one word per memory bank.

[0336] More specifically, process 2500 may begin with step 2502, in which a compiler may receive a task specification. The specification include required computations and/or a priority level.

[0337] In step 2504, a compiler may identify an accelerator, or group of accelerators, that may perform the task. Alternatively, the compiler may generate instructions so the processing units, such as configuration manager 2212, may identify an accelerator to perform the task. For example, using the required computation configuration manger 2212 may identify accelerators in the group of accelerators 2216 that may process the task.

[0338] In step 2506, the compiler may determine a number of words that needs to be simultaneously accessed to execute the task. For example, the multiplication of two vectors requires access to at least two vectors, and the compiler may therefore determine that vector words must be simultaneously accessed to perform the operation.

[0339] In step 2508, the compiler may determine a number of cycles necessary to execute the task. For example, if the task requires a convolution operation of four by-products, the compiler may determine that at least 4 cycles will be necessary to perform the task.

5 [0340] In step 2510, the compiler may place words that are needed to be accessed simultaneously in different memory banks. In that way, memory controller 2210 may be configured to open lines to different memory instances and access the required memory blocks within a clock cycle, without any required cached data.

10 [0341] In step 2512, the compiler place words that are accessed sequentially in the same memory banks. For example, in the case that four cycles of operations are required, the compiler may generate instructions to write needed words in sequential cycles in a single memory block to avoid changing lines between different memory blocks during execution.

15 [0342] In step 2514, compiler generate instructions for programing processing units, such as configuration manager 2212. The instructions may specify conditions to operate a switching device (such as a multiplexor) or configure a data bus. With such instructions, configuration manager 2212 may configure memory controller 2210 to route data from, or grant access to, memory blocks to processing units using dedicated lines of communication according to a task.

20 [0343] Fig. 26 is an exemplary flowchart illustrating a memory read process 2600, consistent with disclosed embodiments. Where convenient in describing memory read process 2600, reference may be made to the identifiers of elements depicted in Fig. 22 and described above. In some embodiments, as described below, process 2600 may be implemented by memory controller 2210. In other embodiments, however, process 2600 may be implemented by other elements in the processing device 2200, such as configuration manager 2212.

25 [0344] In step 2602, memory controller 2210, configuration manager 2212, or other processing units may receive an indication to route data from, or grant access to, a memory bank. The request may specify an address and a memory block.

[0345] In some embodiments, the request may be received via a data bus specifying a read command in line 2218 and address in line 2220. In other embodiments, the request may be received via demultiplexers connected to memory controller 2210.

30 [0346] In step 2604, configuration manager 2212, a host, or other processing units, may query an internal register. The internal register may include information regarding opened lines to memory banks, opened addresses, opened memory blocks, and/or upcoming tasks. Based on the information in the internal register, it may be determined whether there are lines opened to the memory bank and/or whether the memory block received the request in step 2602. Alternatively, or additionally, memory controller 2210 may directly query the internal register.

35 [0347] If the internal register indicates that the memory bank is not loaded in an opened line (step 2606: no), process 2600 may continue to step 2616 and a line may be loaded to a memory bank associated with the received address. In addition, memory controller 2210 or a processing unit, such as configuration manager 2212, may signal a delay to the element requesting information from the memory

address in step 2616. For example, if accelerator 2216 is requesting the memory information that is located an already occupied memory block, memory controller 2210 may send a delay signal to the accelerator in step 2618. In step 2620, configuration manager 2212 or memory controller 2210 may update the internal register to indicate a line has opened to a new memory bank or a new memory block.

5 [0348] If the internal register indicates that the memory bank is loaded in an opened line (step 2606: yes), process 2600 may continue to step 2608. In step 2608, it may be determined whether the line loaded the memory bank is being used for a different address. If the line is being used for a different address (step 2608: yes), it would indicate that there are two instances in a single block and, therefore, they cannot be accessed simultaneously. Thus, an error or exemption signal may be send to the element  
10 requesting information from the memory address in step 2616. But, if the line is not being used for a different address (step 2608: no), a line may be opened for the address and retrieve data from the target memory bank and continue to step 2614 to transmit data to the to the element requesting information from the memory address.

[0349] With process 2600, processing device 2200 has the ability to establish direct connections  
15 between processing units and the memory blocks or memory instances that contain the required information to perform a task. This organization of data would enable reading information from organized vectors in different memory instances, as well as allow the retrieval of information simultaneously from different memory blocks when a device requests a plurality of these addresses.

[0350] Fig. 27 is an exemplary flowchart illustrating an execution process 2700, consistent with  
20 disclosed embodiments. Where convenient in describing execution process 2700, reference may be made to the identifiers of elements depicted in Fig. 22 and described above.

[0351] In step 2702, a compiler or a local unit, such as configuration manager 2212, may receive  
an indication of a task that needs to be performed. The task may include a single operation (e.g., multiplication) or a more complex operation (e.g., convolution between matrixes). The task may also  
25 indicate a required computation.

[0352] In step 2704, the compiler or configuration manager 2212 may determine a number of  
words that is required simultaneously to perform the task. For example, configuration a compiler may determine two words are required simultaneously to perform a multiplication between vectors. In another  
example, a 2D convolution task, configuration manager 2212 may determine that “n” times “m” words  
30 are required for a convolution between matrices, where “n” and “m” are the matrices dimensions. Moreover, in step 2704, configuration manager 2212 may also determine a number of cycles necessary to perform the task.

[0353] In step 2706, depending on the determinations in step 2704, a compiler may write words  
that need to be accessed simultaneously in a plurality of memory banks disposed on the substrate. For  
35 instance, when a number a number of words that can be accessed simultaneously from one of the plurality of memory banks is lower than the number of words that are required simultaneously, a compiler may organize data in multiple memory banks to facilitate access to the different required words within a clock. Moreover, when configuration manager 2212 or the compiler determine a number of cycles is necessary

to perform the task, the compiler may write words that are needed in sequential cycles in a single memory bank of the plurality of memory banks to prevent switching of lines between memory banks.

[0354] In step 2708, memory controller 2210 may be configured to read or grant access to at least one first word from a first memory bank from the plurality of memory banks or blocks using a first memory line..

[0355] In step 2170, a processing unit, for example one of accelerators 2216, may process the task using the at least one first word.

[0356] In step 2712, memory controller 2210 may be configured to open a second memory line in a second memory bank. For example, based on the tasks and using the pipelined memory access approach, memory controller 2210 may be configured to open a second memory line in a second memory block where information required for the tasks was written in step 2706. In some embodiments, the second memory line may be opened when the task in step 2170 is about to be completed. For example, if a task requires 100 clocks, the second memory line may be opened in the 90th clock.

[0357] In some embodiments, steps 2708-2712 may be executed within one line access cycle.

[0358] In step 2714, memory controller 2210 may be configured to grant access to data from at least one second word from the second memory bank using the second memory line opened in step 2710.

[0359] In step 2176, a processing unit, for example one of accelerators 2216, may process the task using the at least second word.

[0360] In step 2718, memory controller 2210 may be configured to open a second memory line in the first memory bank. For example, based on the tasks and using the pipelined memory access approach, memory controller 2210 may be configured to open a second memory line to the first memory block. In some embodiments, the second memory line to the first block may be opened when the task in step 2176 is about to be completed.

[0361] In some embodiments, steps 2714-2718 may be executed within one line access cycle.

[0362] In step 2720, memory controller 2210 may read or grant access to at least one third word from the first memory bank from the plurality of memory banks or blocks using a second memory line in the first bank or a first line in a third bank and continuing in different memory banks.

[0363] The foregoing description has been presented for purposes of illustration. It is not exhaustive and is not limited to the precise forms or embodiments disclosed. Modifications and adaptations will be apparent to those skilled in the art from consideration of the specification and practice of the disclosed embodiments. Additionally, although aspects of the disclosed embodiments are described as being stored in memory, one skilled in the art will appreciate that these aspects can also be stored on other types of computer readable media, such as secondary storage devices, for example, hard disks or CD ROM, or other forms of RAM or ROM, USB media, DVD, Blu-ray, 4K Ultra HD Blu-ray, or other optical drive media.

[0364] Computer programs based on the written description and disclosed methods are within the skill of an experienced developer. The various programs or program modules can be created using any of the techniques known to one skilled in the art or can be designed in connection with existing software.

For example, program sections or program modules can be designed in or by means of .Net Framework, .Net Compact Framework (and related languages, such as Visual Basic, C, etc.), Java, C++, Objective-C, HTML, HTML/AJAX combinations, XML, or HTML with included Java applets.

- [0365] Moreover, while illustrative embodiments have been described herein, the scope of any  
5 and all embodiments having equivalent elements, modifications, omissions, combinations (e.g., of aspects  
across various embodiments), adaptations and/or alterations as would be appreciated by those skilled in  
the art based on the present disclosure. The limitations in the claims are to be interpreted broadly based on  
the language employed in the claims and not limited to examples described in the present specification or  
during the prosecution of the application. The examples are to be construed as non-exclusive.  
10 Furthermore, the steps of the disclosed methods may be modified in any manner, including by reordering  
steps and/or inserting or deleting steps. It is intended, therefore, that the specification and examples be  
considered as illustrative only, with a true scope and spirit being indicated by the following claims and  
their full scope of equivalents.

**WHAT IS CLAIMED IS:**

1. A distributed processor, comprising:
  - a substrate;
  - a memory array disposed on the substrate, the memory array including a plurality of discrete
- 5 memory banks;
  - a processing array disposed on the substrate, the processing array including a plurality of processor subunits, each one of the processor subunits being associated with a corresponding, dedicated one of the plurality of discrete memory banks;
  - a first plurality of buses, each connecting one of the plurality of processor subunits to its
- 10 corresponding, dedicated memory bank; and
  - a second plurality of buses, each connecting one of the plurality of processor subunits to another of the plurality of processor subunits.
2. The distributed processor of claim 1, wherein the substrate is a semiconductor substrate.
3. The distributed processor of claim 1, wherein the substrate is a circuit board.
- 15 4. The distributed processor of claim 1, wherein the plurality of processor subunits of the processing array are spatially distributed among the plurality of discrete memory banks of the memory array.
5. The distributed processor of claim 1, wherein the distributed processor on a chip is an artificial intelligence accelerator processor.
- 20 6. The distributed processor of claim 1, wherein each of the plurality of processor subunits is configured to execute software code associated with a particular application independently relative to other processor subunits included in the plurality of processor subunits.
7. The distributed processor of claim 1, wherein the plurality of logic processor subunits are arranged in at least one row and at least one column, and the second plurality of buses connect each
- 25 processor subunit to at least one adjacent processor subunit in the same row and to at least one adjacent processor subunit in the same column.
8. The distributed processor of claim 1, wherein the plurality of processor subunits are arranged in a star pattern, and the second plurality of buses connect each processor subunit to at least one adjacent processor subunit within the star pattern.
- 30 9. The distributed processor of claim 1, wherein each processor subunit is associated with at least two dedicated memory banks.
10. The distributed processor of claim 1, wherein each dedicated memory bank comprises at least one dynamic random access memory.
11. A memory chip, comprising:
  - a substrate;
  - a memory array disposed on the substrate, the memory array including a plurality of discrete
- 35 memory banks;

a processing array disposed on the substrate, the processing array including a plurality of logic portions each including an address generator, wherein each one of the address generators is associated with a corresponding, dedicated one of the plurality of discrete memory banks; and

5 a plurality of buses, each connecting one of the plurality of address generators to its corresponding, dedicated memory bank.

12. The memory chip of claim 11, wherein each address generator is configured to determine addresses to be accessed for processing in a memory bank corresponding to the address generator.

13. The memory chip of claim 11, wherein the processing array further includes a plurality of accelerators, each accelerator being associated with a corresponding address generator and configured to  
10 perform a specific function.

14. The memory chip of claim 13, wherein the specific function comprises a multiply-accumulate function, a minimum function, a maximum function, a compare function, or a count function.

15. The memory chip of claim 113, wherein the memory chip is at least one of a DRAM, Flash, SRAM, ReRAM, PRAM, MRAM, or ROM memory chip.

16. The memory chip of claim 11, wherein each processor subunit comprises a reduced instruction set computer (RISC) processor or a complex instruction set computer (CISC) processor.

17. The memory chip of claim 11, further comprising a memory interface connected to an external host.

18. The memory chip of 17, wherein the memory interface comprises an interface compliant  
20 with at least one Joint Electron Device Engineering Council (JEDEC) standard or its variants.

19. The memory chip of claim 11, wherein each of the plurality of logic portions corresponds to at least one memory bank of the plurality of discrete memory banks, a plurality of memory mats included in a single memory bank of the plurality of discrete memory banks, or a single memory mat included in a single memory bank of the plurality of discrete memory banks.

20. A distributed processor, comprising:  
25 a substrate;

a memory array disposed on the substrate, the memory array including a plurality of discrete memory banks, wherein each of the discrete memory banks has a capacity greater than one megabyte; and

30 a processing array disposed on the substrate, the processing array including a plurality of processor subunits, each one of the processor subunits being associated with a corresponding, dedicated one of the plurality of discrete memory banks.

21. The distributed processor of claim 20, further including:

a first plurality of buses, each connecting one of the plurality of processor subunits to a corresponding, dedicated memory bank; and

35 a second plurality of buses, each connecting one of the plurality of processor subunits to another one of the plurality of processor subunits.

22. The distributed processor of claim 20, wherein each dedicated memory bank comprises at least one dynamic random access memory bank.

23. The distributed processor of claim 20, wherein each dedicated memory bank comprises at least one static random access memory bank.
24. The distributed processor of claim 20, wherein each dedicated memory bank is the same size.
- 5 25. The distributed processor of claim 20, wherein at least two of the plurality of memory banks have different sizes.
26. The distributed processor of claim 20, wherein the plurality of processor subunits are spatially distributed among the plurality of discrete memory banks within the memory array.
27. The distributed processor of claim 20, wherein the substrate includes a semiconductor  
10 substrate.
28. A distributed processor, comprising:  
a substrate;  
a memory array disposed on the substrate, the memory array including a plurality of discrete memory banks; and  
15 a processing array disposed on the substrate, the processing array including a plurality of processor subunits, each one of the processor subunits being associated with a corresponding, dedicated one of the plurality of discrete memory banks; and  
a plurality of buses, each one of the plurality of buses connecting one of the plurality of processor subunits to at least another one of the plurality of processor subunits,  
20 wherein the plurality of buses are free of timing hardware logic components such that data transfers between processor subunits and across corresponding ones of the plurality of buses are uncontrolled by timing hardware logic components.
29. The distributed processor of claim 28, wherein the plurality of buses are free of bus arbiters such that data transfers between processor subunits and across corresponding ones of the plurality  
25 of buses are uncontrolled by bus arbiters.
30. The distributed processor of claim 28, wherein the plurality of buses comprise at least one of wires or optical fibers between corresponding ones of the plurality of processor subunits.
31. The distributed processor of claim 28, wherein the plurality of processor subunits are configured to transfer data across at least one of the plurality of buses in accordance with code executed  
30 by the plurality of processor subunits.
32. The distributed processor of claim 31, wherein the code dictates timing of data transfers across at least one of the plurality of buses.
33. The distributed processor of claim 28, further comprising a second plurality of buses, wherein each of the second plurality of buses connects one of the plurality of processor subunits to a  
35 corresponding, dedicated memory bank.
34. The distributed processor of claim 33, wherein the second plurality of buses are free of timing hardware logic components such that data transfers between processor subunits and corresponding, dedicated memory banks are uncontrolled by timing hardware logic components.



35. The distributed processor of claim 33, wherein the second plurality of buses are free of bus arbiters such that data transfers between processor subunits and corresponding, dedicated memory banks are uncontrolled by bus arbiters.

36. The distributed processor of claim 28, wherein the plurality of processor subunits are spatially distributed among the plurality of discrete memory banks within the memory array.

37. The distributed processor of claim 28, wherein the substrate includes a semiconductor substrate.

38. A distributed processor on a memory chip, comprising:

a substrate;

a memory array disposed on the substrate, the memory array including a plurality of discrete memory banks; and

a processing array disposed on the substrate, the processing array including a plurality of processor subunits, each one of the processor subunits being associated with a corresponding, dedicated one of the plurality of discrete memory banks; and

a plurality of buses, each one of the plurality of buses connecting one of the plurality of processor subunits to a corresponding, dedicated one of the plurality of discrete memory banks,

wherein the plurality of buses are free of timing hardware logic components such that data transfers between a processor subunit and a corresponding, dedicated one of the plurality of discrete memory banks and across a corresponding one of the plurality of buses are uncontrolled by timing hardware logic components.

39. A distributed processor, comprising:

a substrate;

a memory array disposed on the substrate, the memory array including a plurality of discrete memory banks; and

a processing array disposed on the substrate, the processing array including a plurality of processor subunits, each one of the processor subunits being associated with a corresponding, dedicated one of the plurality of discrete memory banks; and

a plurality of buses, each one of the plurality of buses connecting one of the plurality of processor subunits to at least another one of the plurality of processor subunits, wherein the plurality of processor subunits are configured to execute software that controls timing of data transfers across the plurality of buses to avoid colliding data transfers on at least one of the plurality of buses.

40. A distributed processor on a memory chip, comprising:

a substrate;

a plurality of processor subunits disposed on the substrate, each processor subunit being configured to execute a series of instructions independent from other processor subunits, each series of instructions defining a series of tasks to be performed by a single processor subunit;

a corresponding plurality of memory banks disposed on the substrate, each one of the plurality of processor subunits being connected to at least one dedicated memory bank not shared by any others of the plurality of processor subunits; and

5 a plurality of buses, each of the plurality of buses connecting one of the plurality of processor subunits to at least one other of the plurality of processor subunits,

wherein data transfers across at least one of the plurality of buses are predefined by the series of instructions included in a processor subunit connected to the at least one of the plurality of buses.

41. The distributed processor on a memory chip of claim 40, wherein each series of instructions comprises a set of machine code defining a corresponding series of tasks.

10 42. The distributed processor on a memory chip of claim 41, wherein the series of tasks are defined by a compiler configured to distribute a higher-level series of tasks amongst the plurality of logic circuits as a plurality of series of tasks.

43. The distributed processor on a memory chip of claim 42, wherein the higher-level series of tasks comprises a set of instructions in a human-readable programming language.

15 44. The distributed processor on a memory chip of claim 40, wherein the series of instructions included in the processor subunit connected to the at least one of the plurality of buses includes a sending task that comprises a command for the processor subunit connected to the at least one of the plurality of buses to write data to the at least one of the plurality of buses.

20 45. The distributed processor on a memory chip of claim 40, wherein the series of instructions included in the processor subunit connected to the at least one of the plurality of buses includes a receiving task that comprises a command for the processor subunit connected to the at least one of the plurality of buses to read data from the at least one of the plurality of buses.

46. A distributed processor on a memory chip, comprising:

a plurality of processor subunits disposed on the memory chip;

25 a plurality of memory banks disposed on the memory chip, wherein each one of the plurality of memory banks is configured to store data independent from data stored in other ones of the plurality of memory banks, and wherein each one of the plurality of processor subunits is connected to at least one dedicated memory bank from among the plurality of memory banks; and

30 a plurality of buses, wherein each one of the plurality of buses connects one of the plurality of processor subunits to one or more corresponding, dedicated memory banks from among the plurality of memory banks,

wherein data transfers across a particular one of the plurality of buses are controlled by a corresponding processor subunit connected to the particular one of the plurality of buses.

35 47. The distributed processor on a memory chip of claim 46, wherein the data stored in each of the plurality of memory banks are defined by a compiler configured to distribute data amongst the plurality of memory banks.

48. The distributed processor on a memory chip of claim 47, wherein the compiler is configured to distribute data defined in a higher-level series of tasks amongst the plurality of memory banks using a plurality of lower-level tasks distributed amongst corresponding processor subunits.

49. The distributed processor on a memory chip of claim 48, wherein the higher-level series of tasks comprises a set of instructions in a human-readable programming language.

50. The distributed processor on a memory chip of claim 48, wherein the lower-level series of tasks comprises a set of instructions in a machine code.

51. A distributed processor on a memory chip, comprising:  
a plurality of processor subunits disposed on the memory chip;  
a plurality of memory banks disposed on the memory chip, wherein each one of the plurality of processor subunits is connected to at least one dedicated memory bank from among the plurality of memory banks, and wherein each memory bank of the plurality of memory banks is configured to store data independent from data stored in other ones of the plurality of memory banks, and wherein at least some of the data stored in one particular memory bank from among the plurality of memory banks comprises a duplicate of data stored in at least another one of the plurality of memory banks; and  
a plurality of buses, wherein each one of the plurality of buses connects one of the plurality of processor subunits to one or more corresponding, dedicated memory banks from among the plurality of memory banks,

wherein data transfers across a particular one of the plurality of buses are controlled by a corresponding processor subunit connected to the particular one of the plurality of buses.

52. The distributed processor on a memory chip of claim 51, wherein the at least some data duplicated across the one particular memory bank from among the plurality of memory banks and the at least another one of the plurality of memory banks is defined by a compiler configured to duplicate data across memory banks.

53. The distributed processor on a memory chip of claim 51, wherein the at least some data duplicated across the one particular memory bank from among the plurality of memory banks and the at least another one of the plurality of memory banks comprises weights of a neural network.

54. The distributed processor on a memory chip of claim 53, wherein each node in the neural network is defined by at least one processor subunit from among the plurality of processor subunits.

55. The distributed processor on a memory chip of claim 54, wherein each node comprises machine code executed by the at least one processor subunit defining the node.

56. A non-transitory computer-readable medium storing instructions for compiling a series of instructions for execution on a memory chip comprising a plurality of processor subunits and a plurality of memory banks, wherein each processor subunit from among the plurality of processor subunits is connected to at least one corresponding, dedicated memory bank from among the plurality of memory banks, the instructions causing at least one processor to:  
divide the series of instructions into a plurality of groups of sub-series instructions, the division comprising:

assigning tasks associated with the series of instructions to different ones of the processor subunits, wherein the processor subunits are spatially distributed among the plurality of memory banks disposed on the memory chip;

5 generating tasks to transfer data between pairs of the processor subunits of the memory chip, each pair of processor subunits being connected by a bus, and

grouping the assigned and generated tasks into the plurality of groups of sub-series instructions, wherein each of the plurality of groups of sub-series instructions corresponds to a different one of the plurality of processor sub-units;

10 generate machine code corresponding to each of the plurality of groups of subs-series instructions; and

assign the generated machine code corresponding to each of the plurality of groups of subs-series instructions to a corresponding one of the plurality of processor subunits in accordance with the division.

57. The non-transitory computer-readable medium of claim 56, wherein the assigning of tasks associated with the series of instructions to the different ones of the processor subunits depends on a spatial proximity between two or more of the processor subunits on the memory chip.

58. The non-transitory computer-readable medium of claim 56, wherein the instructions further cause the at least one processor to:

group data associated with the series of instructions based on the division; and

assign the data to the memory banks in accordance with the grouping.

20 59. The non-transitory computer-readable medium of claim 58, wherein grouping the data comprises determining at least a portion of the data to duplicate in two or more of the memory banks.

60. A memory chip comprising a plurality of processor subunits and a corresponding plurality of memory banks, each processor subunit being connected to at least one memory bank dedicated to the processor subunit, the processor subunits of the memory chip being configured to execute the machine code generated according to claim 56.

61. A memory chip, comprising:

30 a plurality of memory banks, each memory bank having a bank row decoder, a bank column decoder, and a plurality of memory sub-banks, each memory sub-bank having a sub-bank row decoder and a sub-bank column decoder for allowing reads and writes to locations on the memory sub-bank, each memory sub-bank comprising:

a plurality of memory mats, each memory mat having a plurality of memory cells,

wherein the sub-bank row decoders and the sub-bank column decoders are connected to the bank row decoder and the bank column decoder.

62. The memory chip of claim 61, wherein each memory sub-bank further has a sub-bank controller configured to determine whether to processor read requests and write requests from a controller of the memory bank.

63. The memory chip of claim 62, wherein the controller of the memory is synchronized to a system clock.

64. The memory chip of claim 62, wherein the controllers of the memory sub-banks are not synchronized to a system clock.

65. The memory chip of claim 61, wherein each sub-bank further has a processor subunit using the sub-bank as a dedicated memory.

5 66. The memory chip of claim 61, wherein the processor subunit comprises a configurable processor subunit or an accelerator.

67. The memory chip of claim 65, wherein each processor subunit is configured to access a sub-bank dedicated to the processor subunit using the row decoder and the column decoder of the sub-bank without using the bank row decoder and the bank column decoder.

10 68. A memory chip, comprising:

a plurality of memory banks, each memory bank having a bank controller and a plurality of memory sub-banks, each memory sub-bank having a sub-bank row decoder and a sub-bank column decoder for allowing reads and writes to locations on the memory sub-bank, each memory sub-bank comprising:

15 a plurality of memory mats, each memory mat having a plurality of memory cells, wherein the sub-bank row decoders and the sub-bank column decoders process read and write requests from the bank controller.

69. The memory chip of claim 68, wherein each memory sub-bank further has a sub-bank controller configured to determine whether to processor read requests and write requests from the bank controller.

20 70. The memory chip of claim 69, wherein the bank controller is synchronized to a system clock.

71. The memory chip of claim 69, wherein the controllers of the memory sub-banks are not synchronized to a system clock.

25 72. The memory chip of claim 68, wherein each sub-bank further has a processor subunit using the sub-bank as a dedicated memory.

73. The memory chip of claim 72, wherein each processor subunit is configured to access a sub-bank dedicated to the processor subunit using the row decoder and the column decoder of the sub-bank without using the bank controller.

30 74. A memory chip, comprising:

a plurality of memory banks, each memory bank having a having a bank controller for processing reads and writes to locations on the memory bank, each memory bank comprising:

a plurality of memory mats, each memory mat having a plurality of memory cells and having a mat row decoder and a mat column decoder,

35 wherein the mat row decoders and the mat column decoders process read and write requests from the sub-bank controller.

75. The memory chip of claim 74, wherein each memory mat further has a comparator configured to determine whether to process read requests and write requests from the sub-bank controller based on the command address from the bank controller.

76. The memory chip of claim 74, wherein each memory mat has an assigned address range that is determined by one or more fuses.

77. The memory chip of claim 76, wherein the one or more fuses are configured to disable a memory mat that is faulty.

78. The memory chip of claim 74, wherein each sub-bank further has a processor subunit using the sub-bank as a dedicated memory.

79. A memory chip, comprising:  
a plurality of memory banks, each memory bank having a bank controller, a row decoder, and a column decoder for allowing reads and writes to locations on the memory bank; and  
a plurality of buses connecting each controller of the plurality of bank controllers to at least one other controller of the plurality of bank controllers.

80. The memory chip of claim 79, wherein the plurality of buses may be accessed without interruption of data transfers on main buses of the memory banks.

81. The memory chip of claim 79, wherein each controller is connected to a plurality of other controllers and is configurable for selection of one other of the other controllers for sending or receiving of data.

82. The memory chip of claim 79, wherein each memory bank comprises a dynamic random access memory bank.

83. The memory chip of claim 79, wherein each controller is configurable and is configured to determine addresses for reading and writing in the memory bank having the controller.

84. The memory chip of claim 79, wherein each controller is configured to process data incoming from other controllers before passing the data to the memory bank having the controller.

85. The memory chip of claim 79, wherein each controller is connected to another spatially adjacent controller.

86. A memory device, comprising:  
a substrate;  
a plurality of memory banks on the substrate;  
a plurality of primary logic blocks on the substrate, each of the plurality of primary logic blocks being connected to at least one of the plurality of memory banks;  
a plurality of redundant blocks on the substrate, each of the plurality of redundant blocks being connected to at least one of the memory banks, each of the plurality of redundant blocks replicating at least one of the plurality of primary logic blocks; and  
a plurality of configuration switches on the substrate, each one of the plurality of the configuration switches being connected to at least one of the plurality of primary logic blocks or to at least one of the plurality of redundant blocks;

wherein upon detection of a fault associated with one of the plurality of primary logic blocks:

a first configuration switch of the plurality of configuration switches is configured to disable the one of the plurality of primary logic blocks, and

5 a second configuration switch of the plurality of configuration switches is configured to enable one of the plurality of redundant blocks that replicates the one of the plurality of primary logic blocks.

87. The memory device of claim 86, wherein

the plurality of configuration switches comprise a plurality of activation switches and a plurality of deactivation switches,

10 each one of the plurality of activation switches and each one of the plurality of deactivation switches comprise an external input;

each one of the plurality of activation switches is configured so an activation signal in the external input causes a closed switch condition; and

15 each one of the plurality of deactivation switches is configured so a deactivation signal in the external input causes an open switch condition.

88. The memory device of claim 86, wherein

at least one of the plurality of primary logic blocks is connected to a subset of the plurality of memory banks with a first dedicated connection, and

20 at least one of the plurality of redundant blocks, which replicates the at least one of the plurality of primary logic blocks, is connected to the subset of the plurality of memory banks with a second dedicated connection.

89. The memory device of claim 87, wherein

each one of the plurality of configuration switches couple at least one of the plurality of primary or redundant logic blocks with at least one of a clock node or a power node.

25 90. The memory device of claim 86, wherein the plurality of primary logic blocks comprise: at least one memory logic block configured to enable read and write operations in the memory bank; and

at least one business logic block configured to perform in-memory computations.

91. The memory device of claim 90, wherein

30 the at least one business logic block comprises a first business logic block; and

the plurality of redundant blocks comprise a second business logic block replicating the first business logic block.

92. The memory device of claim 86, wherein the plurality of configuration switches comprise at least one of a fuse, an anti-fuse, a non-volatile memory device, or a one-time programmable device.

35 93. The memory device of claim 86, wherein each one of the plurality of primary logic blocks and each one of the plurality of redundant blocks are connected to an address bus and a data bus.

94. The memory device of claim 86, wherein at least one of the plurality of primary logic blocks comprise:

at least one local logic unit;  
at least one computation unit; and  
at least one duplicated unit,  
wherein,

5                   the at least one duplicated unit replicates the at least one computation unit, and  
                  the at least one local logic unit has a smaller size than the at least one computation unit.

95.     The memory device of claim 94, wherein the at least one of the plurality of primary logic  
blocks comprise:

                  a plurality of local configuration switches, each one of the plurality of local configuration  
10 switches being connected to at least one of the at least one computation unit or the at least one duplicated  
unit,

                  wherein the local configuration switches are configured to disable the at least one local  
computation unit and enable the at least one duplicated unit when a fault is detected in the at least one  
computation unit.

15           96.     The memory device of claim 86 further comprising  
                  an address manager; and  
                  an address bus coupling the address manager to each one of the plurality of memory banks, each  
one of the plurality of primary logic blocks, and each one of the plurality of redundant blocks,

                  wherein upon detection of the fault associated with the one of the plurality of primary logic  
20 blocks, an invalid address is assigned to the one of the plurality of primary logic blocks and a valid  
address is assigned to the one of the plurality of redundant blocks.

                  97.     The memory device of claim 86, wherein  
                  the plurality of primary logic blocks are connected in series;  
                  each one of the plurality of primary logic blocks is connected in parallel with a parallel switch;  
25 and

                  upon detection of the fault associated with the one of the plurality of primary logic blocks, the  
parallel switch connected to the one of the plurality of primary logic blocks is activated to couple two of  
the plurality of primary logic blocks.

98.     The memory device of claim 97, wherein the parallel switch comprises an anti-fuse.

30           99.     The memory device of claim 97, wherein the parallel switch comprises a sampling circuit  
with a selected cycle delay.

                  100.    The memory device of claim 86, wherein  
                  the plurality of primary logic blocks and the plurality of redundant blocks are disposed on the  
substrate in a two-dimensional grid;  
35           each one of the plurality of primary logic blocks and each one of the plurality of redundant blocks  
are interconnected with connection boxes; and  
                  an input block is disposed in a periphery of each line and each column of the two-dimensional  
grid.



101. The memory device of claim 86, wherein the memory chip comprises at least one of a DRAM, a Flash, an SRAM, a ReRAM, a PRAM, or a MRAM memory.

102. The memory device of claim 86, wherein each one of the plurality of primary logic blocks has equivalent computing capabilities to at least one of the plurality of redundant blocks.

5 103. The memory device of claim 86, wherein the plurality of primary logic blocks comprise at least one local computing unit and at least one redundant computing unit, the at least one redundant computing unit replicating the at least one local computing unit.

104. The memory device of claim 86, wherein the first configuration switch is further configured to enable the one of the plurality of redundant blocks that replicates the one of the plurality of  
10 primary logic blocks.

105. A distributed processor on a memory chip, comprising:

a substrate;

an address manager on the substrate;

a plurality of primary logic blocks on the substrate, each of the plurality of primary logic blocks  
15 being connected to at least one of the plurality of memory banks;

a plurality of redundant blocks on the substrate, each of the plurality of redundant blocks being connected to at least one of the plurality of memory banks, each of the plurality of redundant blocks replicating at least one of the plurality of primary logic blocks; and

a bus on the substrate connected to each of the plurality of primary logic blocks, each of the  
20 plurality of redundant blocks, and the address manager,

wherein the processor is configured to:

assign running ID numbers to blocks in the plurality of primary logic blocks that pass a testing protocol;

25 assign illegal ID numbers to blocks in the plurality of primary logic blocks that do not pass the testing protocol; and

assign running ID numbers to blocks in the plurality of redundant blocks that pass the testing protocol.

106. The distributed memory chip of claim 105, wherein the blocks in the plurality of redundant blocks assigned running ID numbers is equal to, or greater than, the blocks in the plurality of  
30 primary logic blocks assigned illegal ID numbers.

107. The distributed memory chip of claim 106, wherein each one of the plurality of primary logic blocks and each one of the plurality of redundant blocks comprise a fused identification circuit.

108. The distributed memory chip of claim 107, wherein the bus comprises a command line, a data line, and an address line.

35 109. A method for configuring a distributed processor on a memory chip, comprising:

testing each one of a plurality of primary logic blocks on the substrate of the memory chip for at least one circuit functionality;

identifying at least one faulty logic block in the plurality of primary logic blocks based on the testing results, the at least one faulty logic block being connected to at least one memory bank disposed on the substrate of the memory chip;

5 testing at least one redundant block on the substrate of the memory chip for the at least one circuit functionality, the at least one redundant block replicating the at least one faulty logic block and being connected to the at least one memory bank;

disabling the at least one faulty logic block by applying an external signal to a deactivation switch, the deactivation switch being connected with the at least one faulty logic block and being disposed on the substrate of the memory chip; and

10 enabling the at least one redundant block by applying the external signal to an activation switch, the activation switch being connected with the at least one redundant block and being disposed on the substrate of the memory chip.

110. A method for configuring a distributed processor on a memory chip, comprising:

15 enabling a plurality of primary logic blocks and a plurality of redundant blocks on the substrate of the memory;

testing each one of the plurality of primary logic blocks on the substrate of the memory chip for at least one circuit functionality;

20 identifying at least one faulty logic block in the plurality of primary logic blocks based on the testing results, the at least one faulty logic block being connected to at least one memory bank disposed on the substrate of the memory chip;

testing at least one redundant block on the substrate of the memory chip for the at least one circuit functionality, the at least one redundant block replicating the at least one faulty logic block and being connected to the at least one memory bank;

25 disabling at least one redundant block by applying the external signal to an activation switch, the activation switch being connected with the at least one redundant block and being disposed on the substrate of the memory chip.

111. A processing device, comprising:

a substrate;

a plurality of memory banks on the substrate;

30 a memory controller on the substrate connected to each one of the plurality of memory banks; and a plurality of processing units on the substrate, each one of the plurality of processing units being connected to the memory controller, the plurality of processing units comprising a configuration manager;

wherein the configuration manager is configured to:

35 receive a first indication of a task to be performed, the task requiring at least one computation;

signal at least one selected processing unit from the plurality of processing units based upon a capability of the selected processing unit for performing the at least one computation; and

transmit a second indication to the at least one selected processing unit, and wherein the memory controller is configured to:

route data from at least two memory banks to the at least one selected processing unit using at least one communication line, the at least one communication line being connected to the at least two memory banks and the at least one selected processing unit via the memory controller.

5

10

112. The processing device of claim 111, wherein transmitting the second indication comprises communicating an instruction to execute the task by the at least one selected processing unit, and wherein the at least one selected processing unit is configured to transmit or receive data to and from the memory controller.

113. The processing device of claim 111, wherein transmitting the second indication comprises communicating the instruction to route data to the memory controller.

15

114. The processing device of claim 111, wherein the at least one selected processing unit is configured to open a memory line in a second memory bank from the at least two memory banks during a series of accesses to an opened memory line in a first memory bank from the at least two memory banks.

115. The processing device of claim 114, wherein the selected processing unit is configured to transfer data to the second memory bank during a line access period in which a communication line is opened to the first memory bank.

20

116. The processing device of claim 111, wherein the memory controller comprises at least two data inputs from the plurality of memory banks and at least two data outputs connected to each one of the plurality of processing units;

the memory controller is configured to simultaneously receive data from two memory banks via the two data inputs; and

25

the memory controller is configured to simultaneously transmit data received via the two data inputs to the at least one selected processing unit via the two data outputs.

117. The processing device of claim 111, wherein the plurality of processing units comprise a plurality of accelerators configured for pre-defined tasks.

118. The processing device of claim 117, wherein the plurality of accelerators comprise at least one of a vector multiply accumulate unit or a direct memory access.

30

119. The processing device of claim 117, wherein the configuration manager comprises at least one of a RISC processor or a micro-controller.

120. The processing device of claim 111, further comprising an external interface connected to the memory banks.

35

121. The processing device of claim 111, wherein the processing device is further configured to:

supply data from a first address through a first memory line to at least one of the plurality of processing units and open a second address in a second memory line within a line access period, the first

address being in a first memory bank of the plurality of memory banks, the second address being in a second memory bank of the plurality of memory banks; and

supply data from the second address through the second memory line to the at least one of the plurality of processing units and open a third address in the first memory bank in the first line within a second line access period.

5

122. The processing device of claim 111, wherein a compiler is configured to:

determine a number of words required simultaneously to perform the task;

determine a number of words that can be accessed simultaneously from each one of the plurality

10

of memory banks; and

divide the number of words required simultaneously between multiple memory banks when the number of words required simultaneously is greater than the number of words that can be accessed simultaneously.

123. The processing device of claim 122, wherein the words comprise machine instructions.

15

124. The processing device of claim 111, wherein the configuration manager comprises a local memory that stores a command to be transmitted to at least one of the plurality of processing units.

125. The processing device of claim 111, wherein the memory controller is configured to interrupt the task in response to receiving a request from an external interface.

20

126. The processing device of claim 111, wherein the plurality of memory banks includes at least one of DRAM mats, DRAM, banks, flash mats, or SRAM mats.

127. The processing device of claim 111, wherein the plurality of processing units comprise at least one arithmetic logic unit, at least one vector handling logic unit, at least one register, and at least one direct memory access.

25

128. The processing device of claim 111, wherein the configuration manager and the plurality of processing units are configured to hand over access to the memory controller between each other after finalizing a task.

129. A method performed for operating a distributed memory device comprising:

compiling, by a compiler, a task for the distributed memory device, the task requiring at least one computation, the compiling comprising:

30

determining a number of words that are required simultaneously to perform the task, and

providing instructions for writing words that need to be accessed simultaneously in a

plurality of memory banks disposed on the substrate when a number of words that can be accessed simultaneously from one of the plurality of memory banks is lower than the number of words that are required simultaneously;

35

receiving, by a configuration manager disposed on the substrate, an indication to perform the task; and

in response to receiving the indication, configuring a memory controller disposed in the substrate

to:

within a first line access cycle:

access at least one first word from a first memory bank from the plurality of memory banks using a first memory line,

send the at least one first word to at least one processing unit, and

5 open a first memory line in the second memory bank to access a second address from the second memory bank from the plurality of memory banks, and

within a second line access cycle:

access at least one second word from the second memory bank using the first memory line,

10 send the at least one second word to at least one processing unit, and

access a third address from the first memory bank using a second memory line in the first bank.

130. The method of claim 129 further wherein the compiling the task further comprises: determining a number of cycles necessary to perform the task; and

15 writing words that are needed in sequential cycles in a single memory bank of the plurality of memory banks.

131. A non-transitory computer-readable medium that stores instructions that, when executed by at least one processor, cause the at least one processor to:

20 determine a number of words that are required simultaneously to perform a task, the task requiring at least one computation;

write words that need to be accessed simultaneously in a plurality of memory banks disposed on the substrate when a number a number of words that can be accessed simultaneously from one of the plurality of memory banks is lower than the number of words that are required simultaneously;

transmit an indication to perform the task to a configuration manager disposed on the substrate;

25 and

transmit instructions to configure a memory controller disposed on the substrate to,

within a first line access cycle:

access at least one first word from a first memory bank from the plurality of memory banks using a first memory line,

30 send the at least one first word to at least one processing unit, and


open a first memory line in the second memory bank to access a second address from the second memory bank from the plurality of memory banks, and

within a second line access cycle:

35 access at least one second word from the second memory bank using the first memory line,

send the at least one second word to at least one processing unit, and

access a third address from the first memory bank using a second memory line in the first bank.



**FIG. 1**




FIG. 2




FIG. 3A






FIG. 3B




FIG. 4




FIG. 5




FIG. 6




FIG. 7A

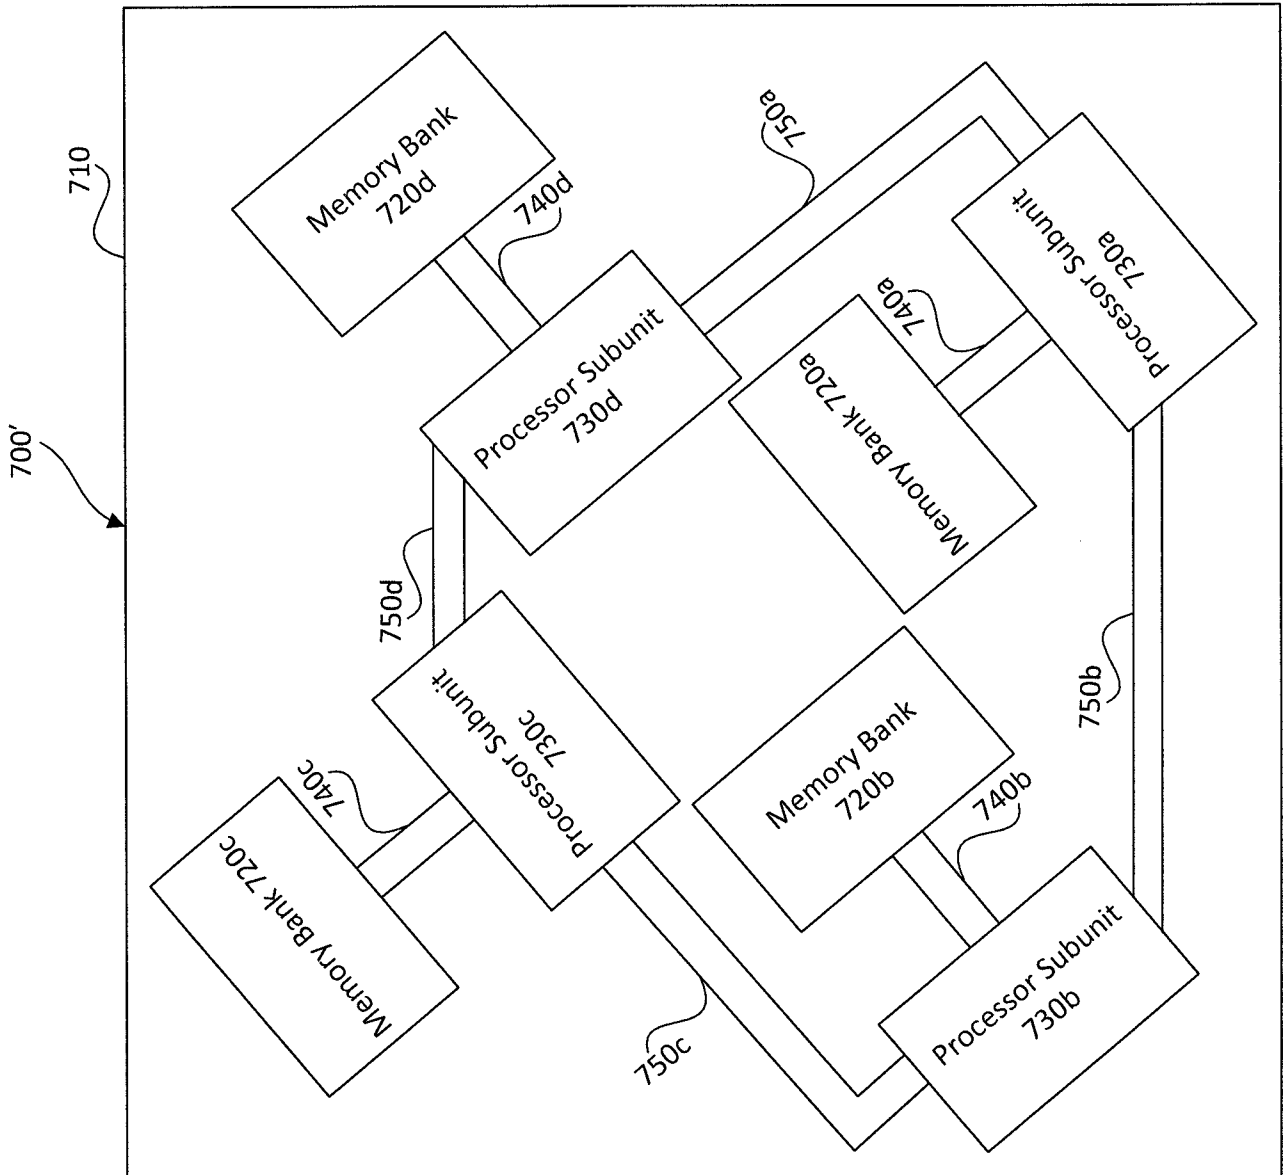


FIG. 7B




FIG. 7C


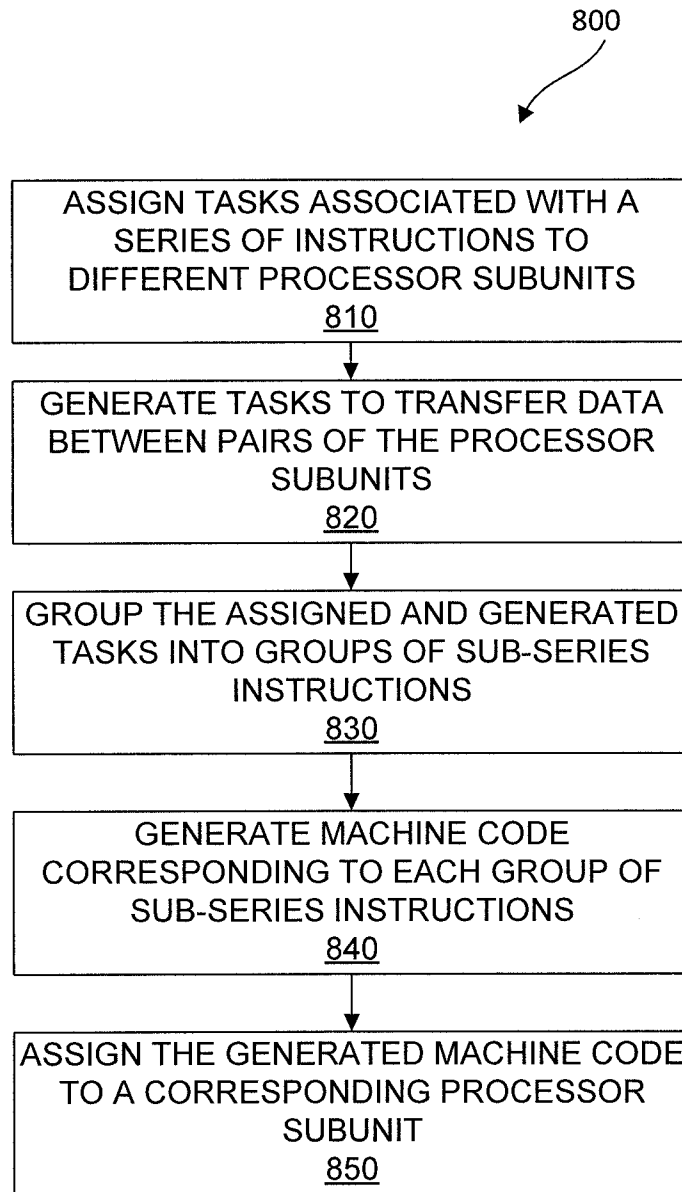
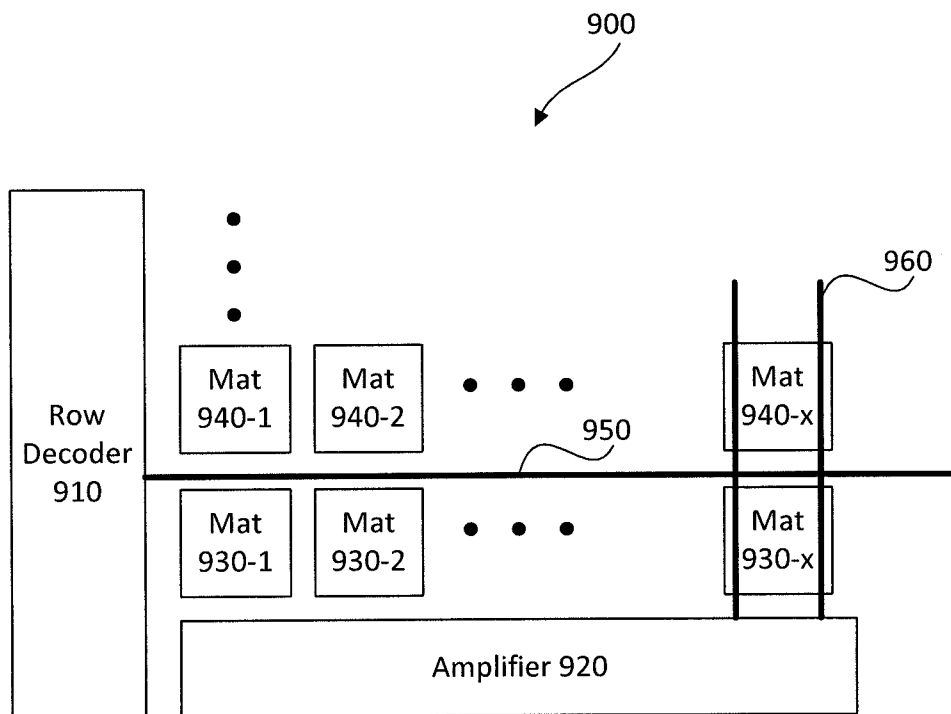


FIG. 7D







**FIG. 8**



**FIG. 9**



**FIG. 10**



**FIG. 11**




FIG. 12






FIG. 13

1400



**FIG. 14**



**FIG. 15**



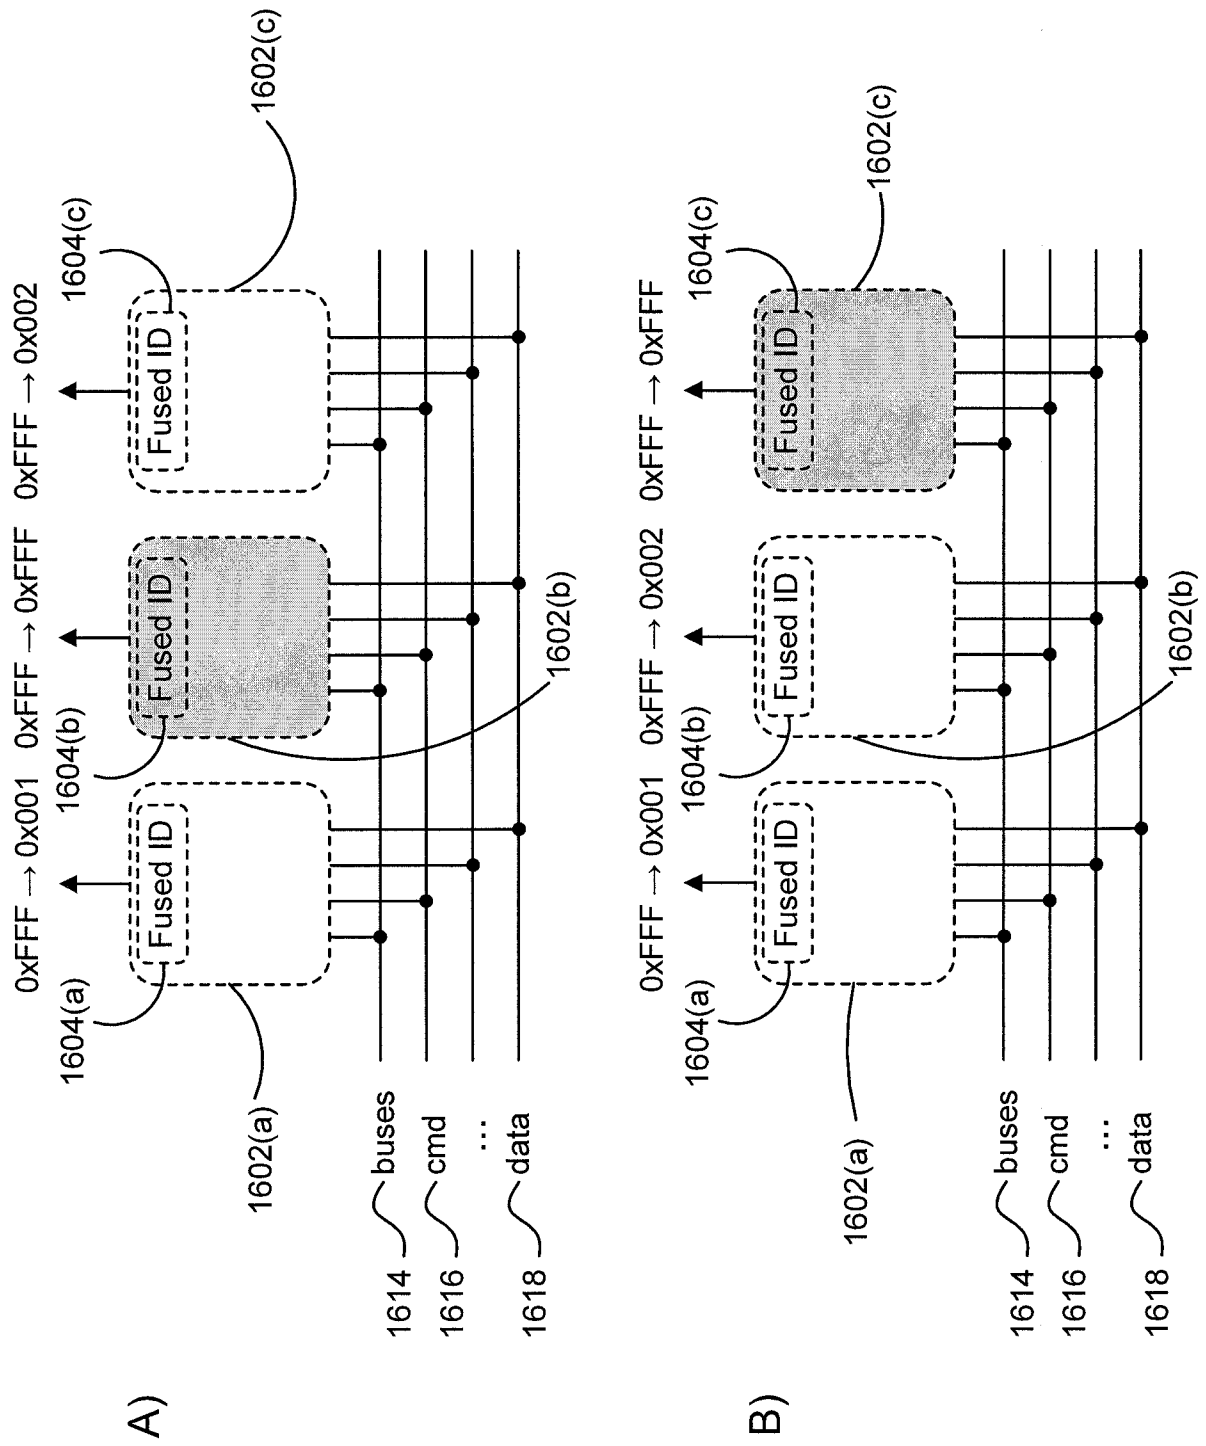


FIG. 16




FIG. 17




FIG. 18




FIG. 19

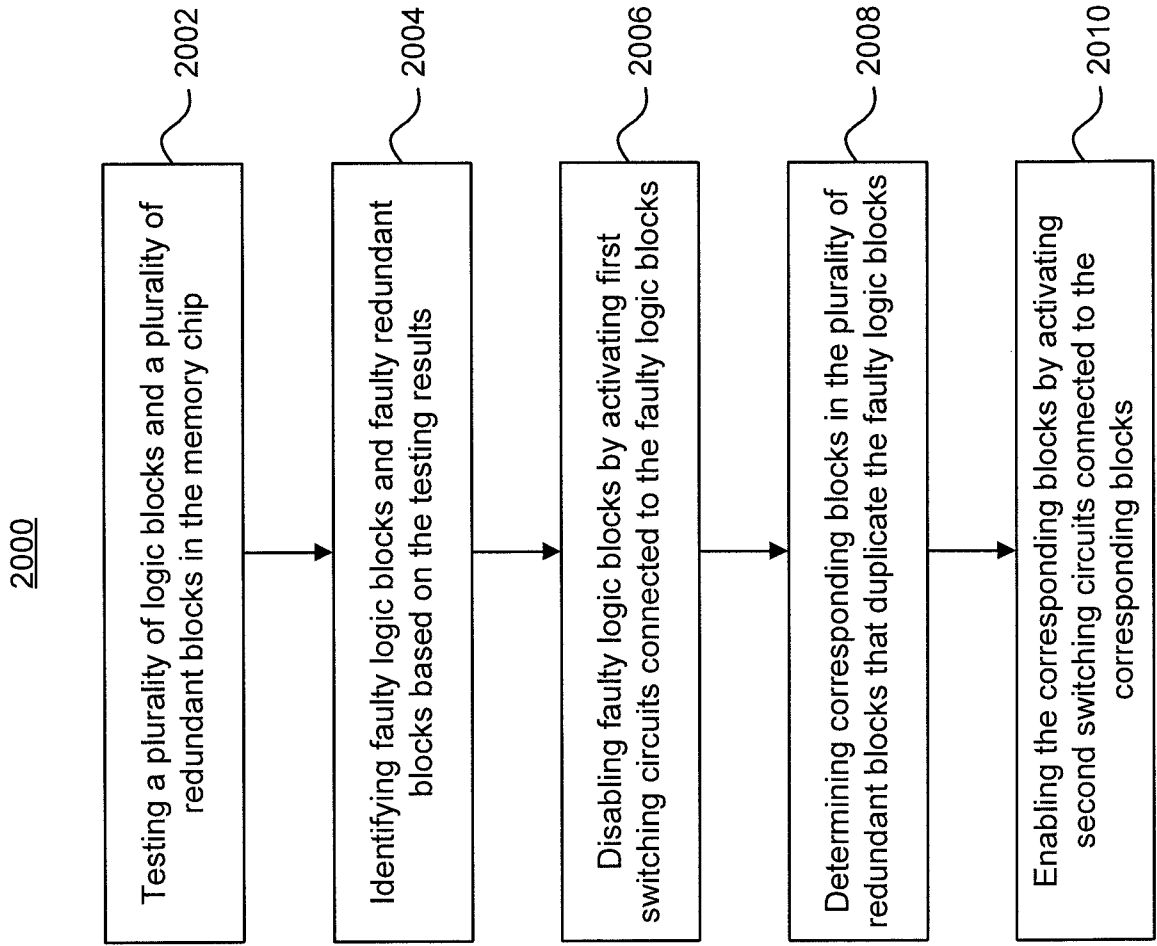


FIG. 20

2100

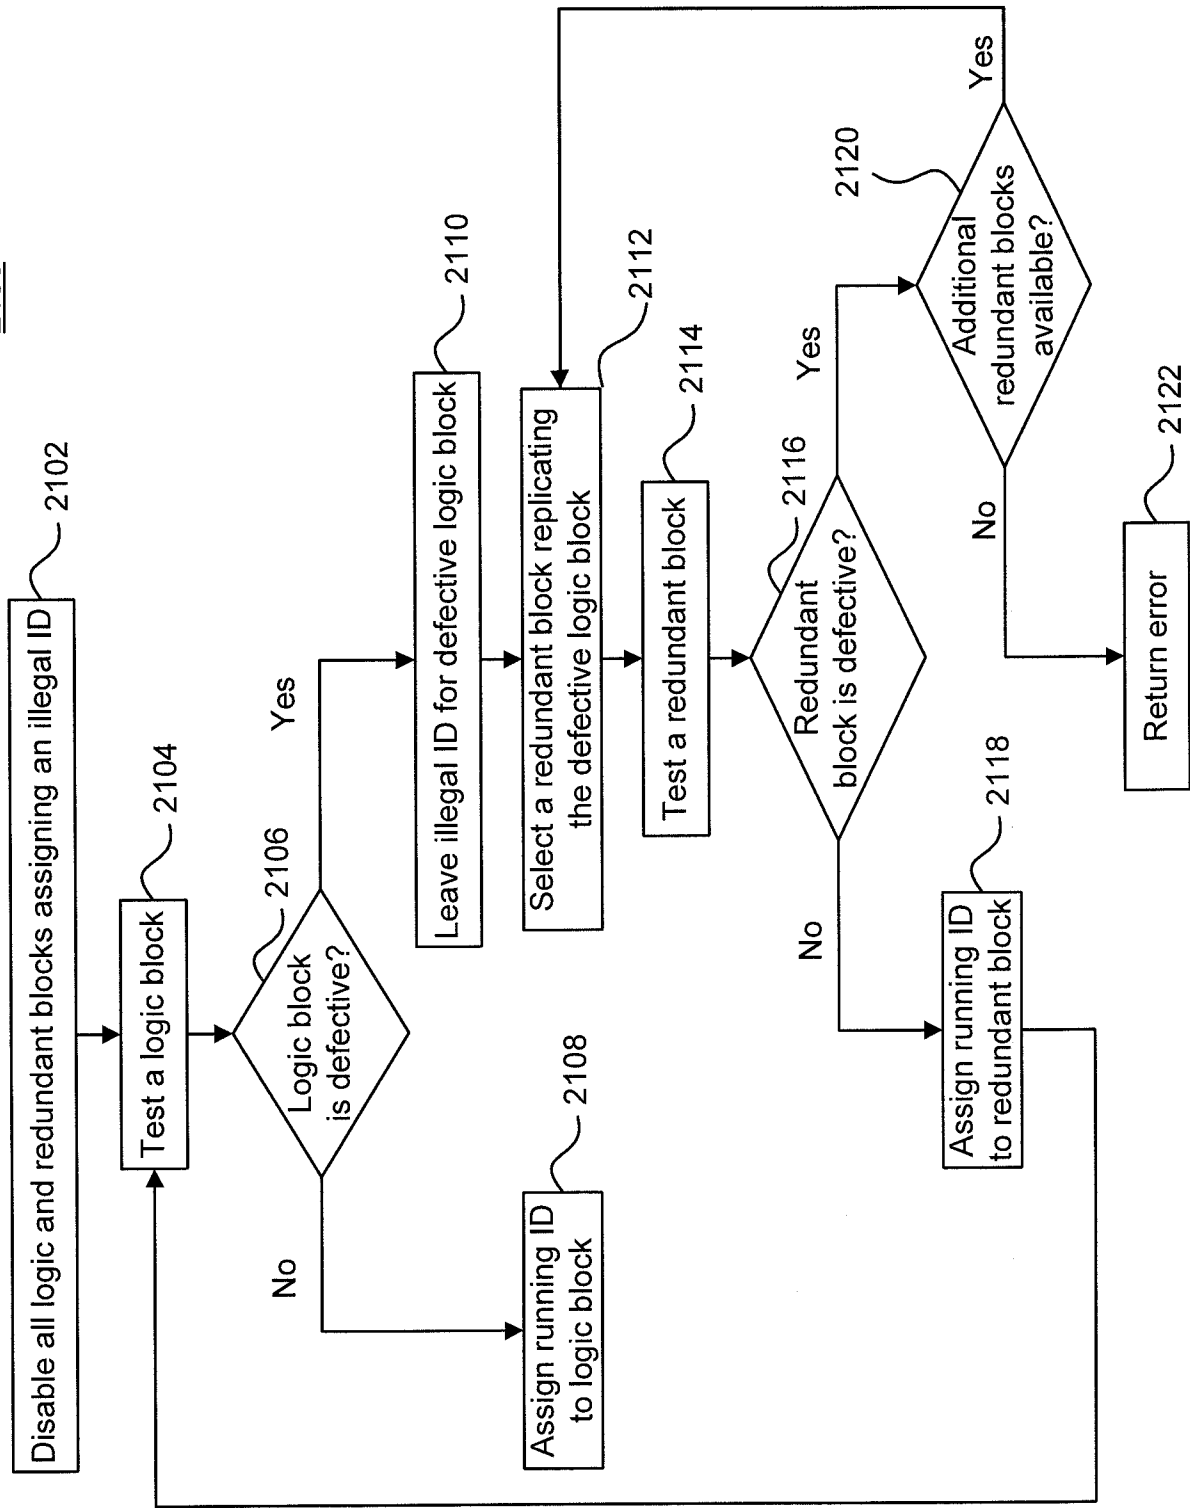


FIG. 21




FIG. 22

2200




FIG. 23



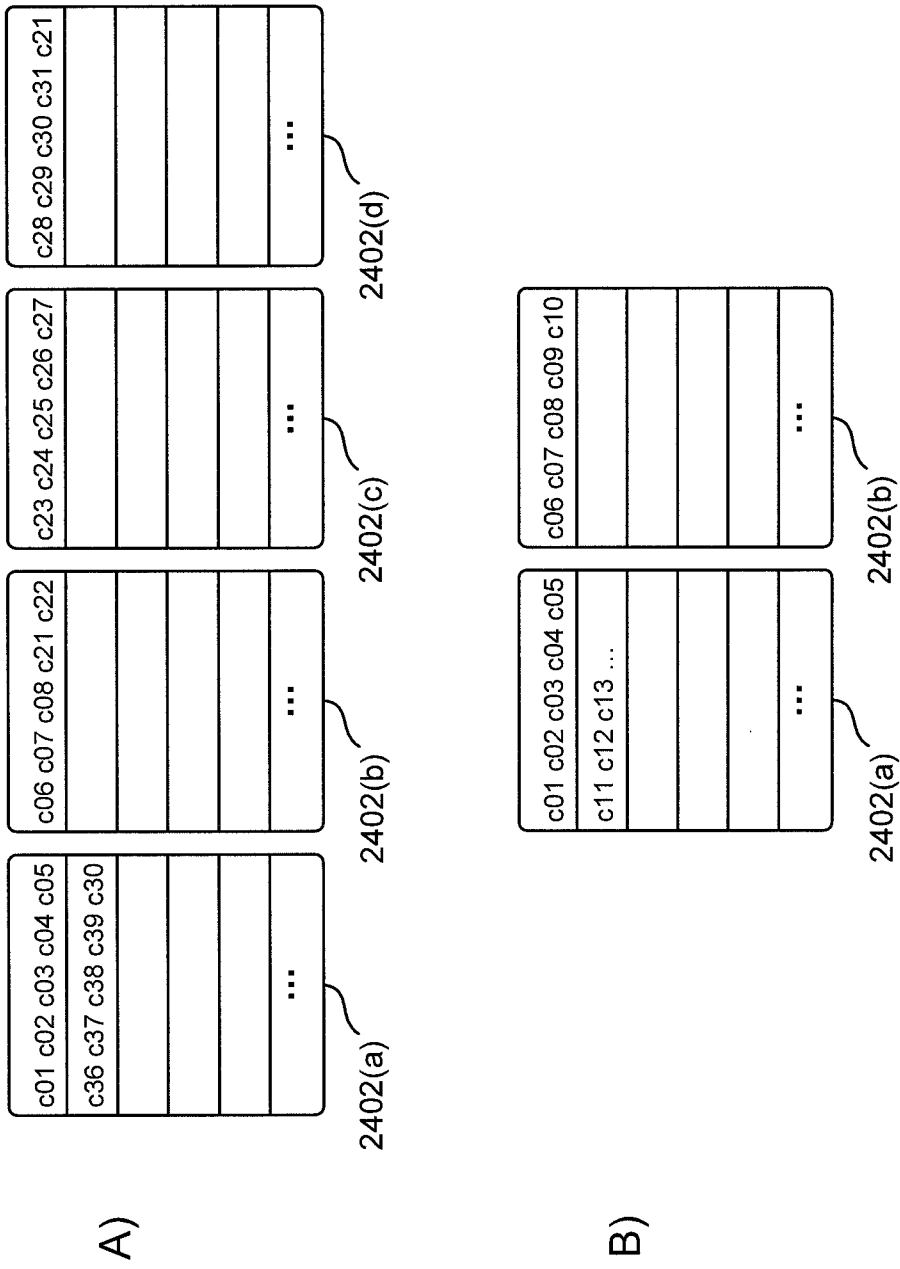
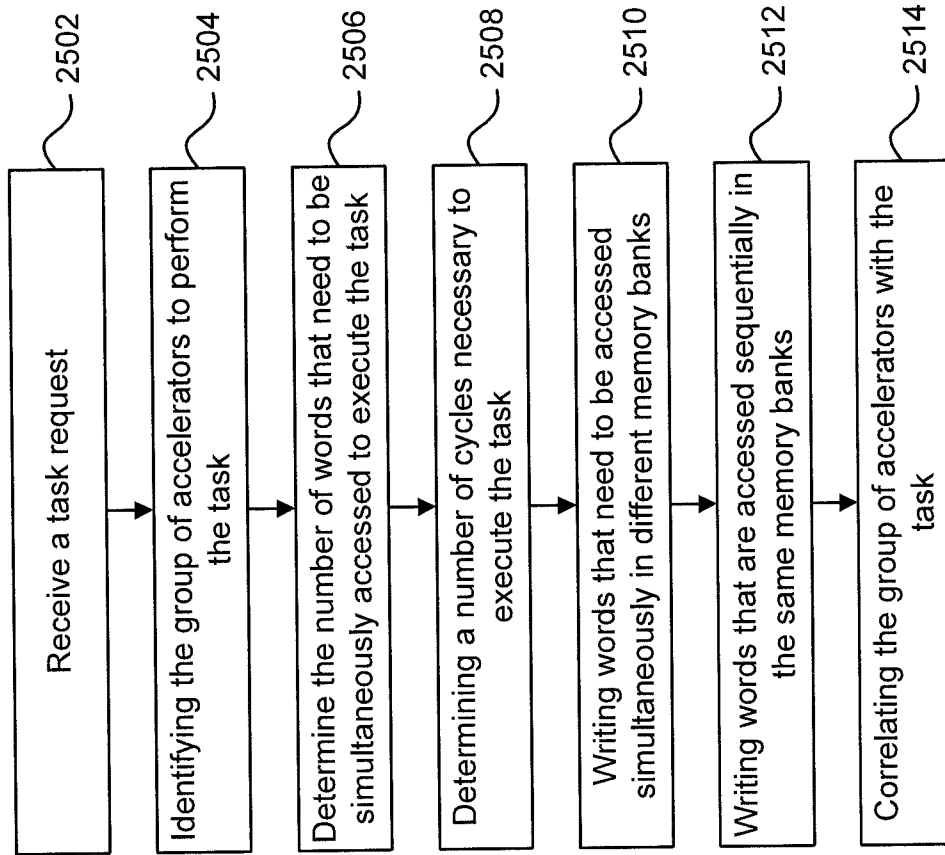


FIG. 24

2500



**FIG. 25**

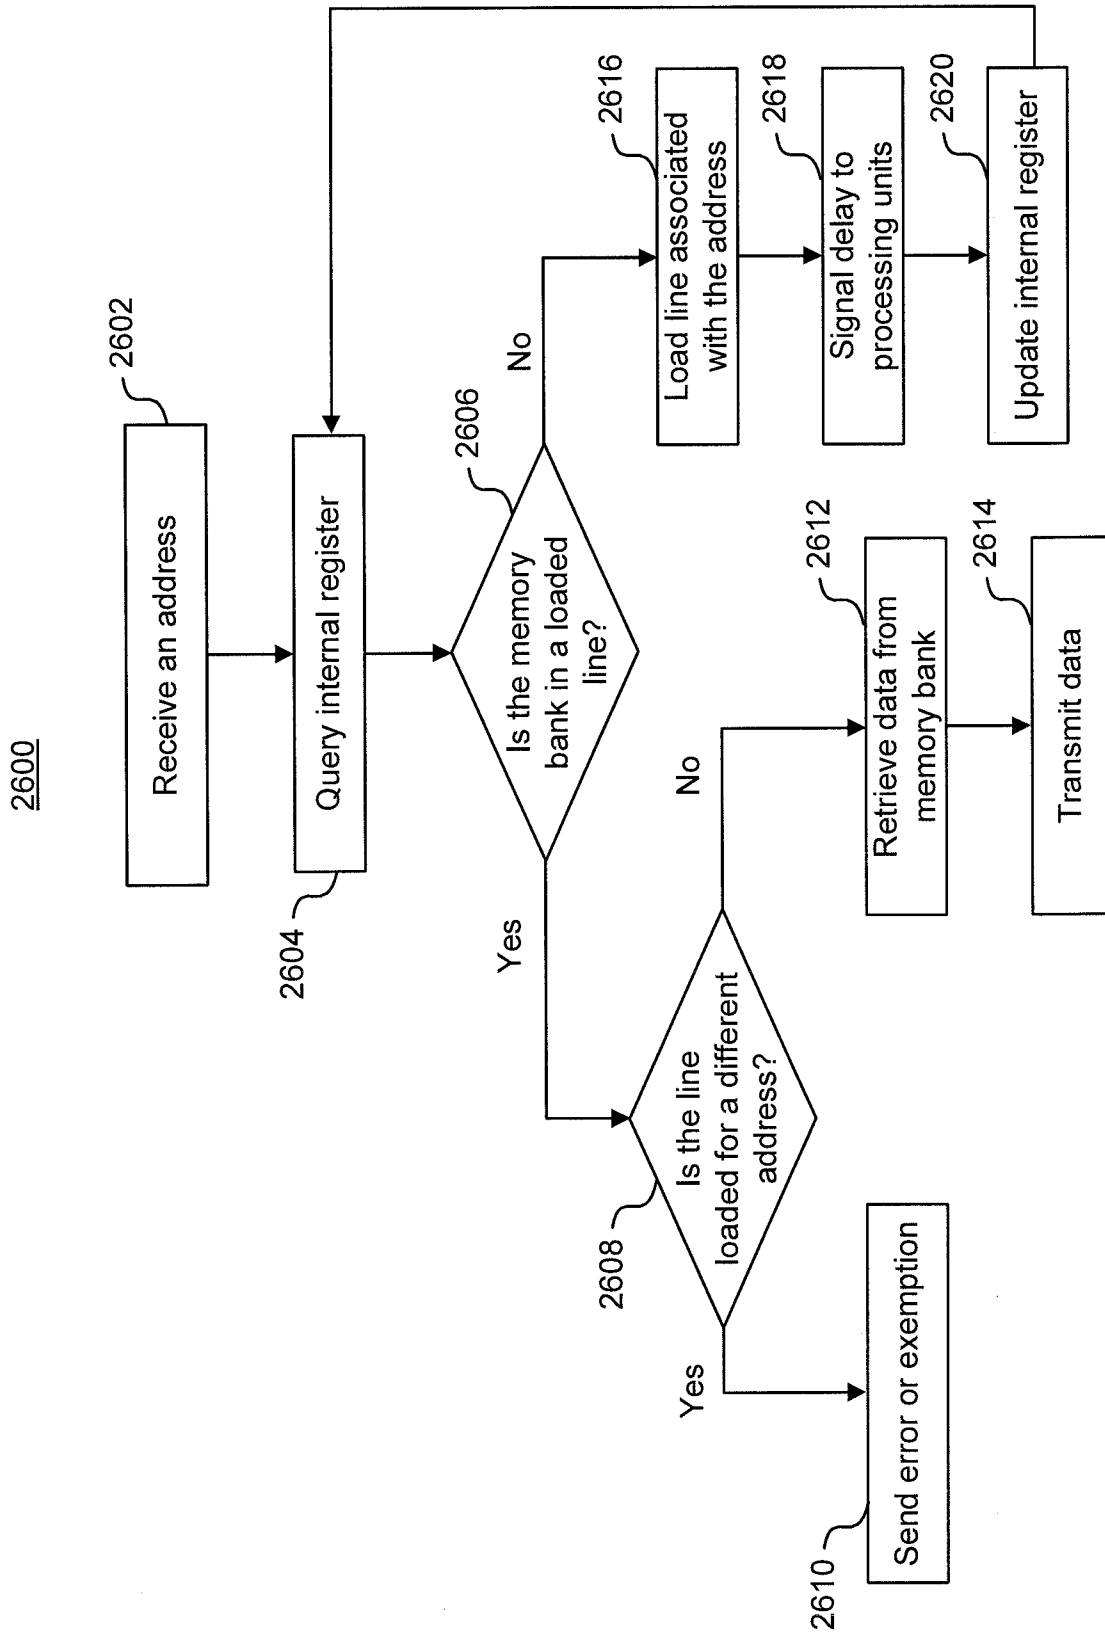


FIG. 26




FIG. 27