



(43) International Publication Date  
20 September 2018 (20.09.2018)

- (51) International Patent Classification:  
G06F 15/00 (2006.01) G06F 15/76 (2006.01)
- (21) International Application Number:  
PCT/US2018/022106
- (22) International Filing Date:  
13 March 2018 (13.03.2018)
- (25) Filing Language:  
English
- (26) Publication Language:  
English
- (30) Priority Data:  
62/471,340 14 March 2017 (14.03.2017) US  
62/471,367 15 March 2017 (15.03.2017) US  
62/471,368 15 March 2017 (15.03.2017) US  
62/471,372 15 March 2017 (15.03.2017) US  
62/472,579 17 March 2017 (17.03.2017) US

Jianbin [CN/US]; 11772 Sorrento Valley Road, Suite 123, San Diego, California 92121 (US).

(74) Agent: XU, Na; IPro PLLC, 21800 Town Center Plaza, #266A-816, Sterling, Virginia 20164 (US).

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

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ,

- (72) Inventors; and
- (71) Applicants: LI, Yuan [CN/US]; 11772 Sorrento Valley Road, Suite 123, San Diego, California 92121 (US). ZHU,

(54) Title: RECONFIGURABLE PARALLEL PROCESSING




FIG. 2

(57) Abstract: Processors, systems and methods are provided for thread level parallel processing. A processor may comprise a plurality of processing elements (PEs) that each may comprise a configuration buffer, a sequencer coupled to the configuration buffer of each of the plurality of PEs and configured to distribute one or more PE configurations to the plurality of PEs, and a gasket memory coupled to the plurality of PEs and being configured to store at least one PE execution result to be used by at least one of the plurality of PEs during a next PE configuration.



TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

**Published:**

- *with international search report (Art. 21(3))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

## RECONFIGURABLE PARALLEL PROCESSING

### Related Application Field

**[0001]** The application claims priority to U.S. Provisional Application 62/471,340, filed March 14, 2017, entitled "Reconfigurable Parallel Processing," U.S. Provisional Application 62/471,367, filed March 15, 2017, entitled "Circular Reconfiguration for Reconfigurable Parallel Processor," U.S. Provisional Application 62/471,368, filed March 15, 2017, entitled "Private Memory Structure for Reconfigurable Parallel Processor," U.S. Provisional Application 62/471,372, filed March 15, 2017, entitled "Shared Memory Structure for Reconfigurable Parallel Processor," U.S. Provisional Application 62/472,579, filed March 17, 2017, entitled "Static Shared Memory Access for Reconfigurable Parallel Processor," the contents of these applications are hereby incorporated by reference in their entities.

### Technical Field

**[0002]** The disclosure herein relates to computer architectures, particularly relates to reconfigurable processors.

### Background

**[0003]** Reconfigurable computing architecture with large amount of processing array can meet demand of computation power while keeping the power and silicon area efficient. Unlike field-programmable gate array (FPGA), Coarse-Grained Reconfigurable Architecture (CGRA) utilizes larger processing elements like arithmetic logic units (ALU) as its building blocks. It provides features of reconfigurability using high level language to quickly program the processing element (PE) array. One typical design of CGRA is shown in FIG. 1. It consists of a PE array, a

configuration memory, a memory unit as a frame buffer and interconnections between the PEs and PEs to the frame buffer.

**[0004]** In general, CGRA is an approach to explore loop level parallelism. It is not specifically targeted to handle thread level parallelism. With any data dependency from one iteration to the next, the parallelism is largely limited. Therefore, the 2D array size is intended to be limited to an 8x8 PE array in most of the designs.

**[0005]** Graphics processing unit (GPU) architecture has provided a way to execute parallel threads in a Same Instruction Multiple Thread (SIMT) fashion. It is especially suitable for massive parallel computing applications. In these applications, typically no dependency is assumed between threads. This type of parallelism is beyond loop level parallelism within a software task, which CGRA is designed for. The thread level parallelism can be easily scalable beyond single core execution to multicore execution. The thread parallelism provides optimization opportunities and makes the PE array more efficient and more capable and it is easily made larger than 8x8. GPU, however, is not reconfigurable. Therefore, there is a need in the art to develop a next generation processor that can harness the processing capability of both CGRA and GPU.

### Summary

**[0006]** The present disclosure describes apparatus, methods and systems for massive parallel data processing. A processor according to various embodiments of the present disclosure may be designed to take advantage of massive thread level parallelism similar to GPU using programable processor array similar to CGRA. In one embodiment, a processor may efficiently process threads which are identical to each other but with different data, similar to SIMT

architecture. A software program's data dependency graph may be mapped to a virtual data path with infinite length. Then the virtual data path may be chopped into segments that can be fit into multiple physical data paths, each physical data path may have its configuration context. A sequencer may distribute configurations of each PE into its configuration FIFO and similarly to switch boxes. A gasket memory may be used to temporarily store outputs of one physical data path configuration and give it back to the processing elements for the next configuration.

Memory ports may be used to calculate addresses for read and write. FIFOs may be used to allow each PE for independent operation. Data stored in a memory unit may be accessed through either private or shared memory access method. The same data can be access through different access methods in different part of the software program to reduce data movement between memories.

**[0007]** In an exemplary embodiment, there is provided a processor comprising: a plurality of processing elements (PEs) each comprising a configuration buffer; a sequencer coupled to the configuration buffer of each of the plurality of PEs and configured to distribute one or more PE configurations to the plurality of PEs; and a gasket memory coupled to the plurality of PEs and being configured to store at least one PE execution result to be used by at least one of the plurality of PEs during a next PE configuration.

**[0008]** According to an embodiment, the processor may further comprise a plurality of switch boxes coupled to the sequencer to receive switch box configurations from the sequencer, each of the plurality of switch boxes may be associated with a respective PE of the plurality of PEs and configured to provide input data switching for the respective PE according to the switch box configurations.

**[0009]** According to an embodiment, the plurality of switch boxes and their associated PEs may be arranged in a plurality of columns, a first switch box in a first column of the plurality of columns may be coupled between the gasket memory and a first PE in the first column of the plurality of columns, and a second PE in a last column of the plurality of columns may be coupled to the gasket memory.

**[0010]** According to an embodiment, the processor may further comprise a memory unit for providing data storage for the plurality of PEs; and a plurality of memory ports each arranged in a separate column of the plurality of columns for the plurality of PEs to access the memory unit.

**[0011]** According to an embodiment, the processor may further comprise a plurality of inter-column switch boxes (ICSBs) coupled to the sequencer to receive ICSB configurations from the sequencer, the plurality of ICSBs may be configured to provide data switching between neighboring columns of the plurality of columns according to the ICSB configurations.

**[0012]** According to an embodiment, the plurality of memory ports (MPs) may be coupled to the sequencer to receive MP configurations from the sequencer and configured to operate in a private access mode or a shared access mode during one MP configuration.

**[0013]** According to an embodiment, a piece of data stored in the memory unit may be accessed through the private access mode and the shared access mode in different part of a program without the piece of data being moved in the memory unit.

**[0014]** According to an embodiment, each of the plurality of columns comprises one PE, the plurality of PEs may be identical and form one row of repetitive identical PEs.

**[0015]** According to an embodiment, each of the plurality of columns may comprise two or more PEs and the plurality of PEs form two or more rows.

**[0016]** According to an embodiment, a first row of PEs may be configured to implement a first set of instructions and a second row of PEs may be configured to implement a second set of instructions, at least one instruction of the second set of instructions is not in the first set of instructions, and the of plurality of columns may be identical and form repetitive columns.

**[0017]** According to an embodiment, each of the plurality of PEs may comprise a plurality of arithmetic logic units (ALUs) that may be configured to execute same instruction in parallel threads.

**[0018]** According to an embodiment, each of the plurality of PEs may comprise a plurality of data buffers for the plurality of ALUs and may be configured to operate independently.

**[0019]** According to an embodiment, each of the plurality of memory ports may be configured to access the memory unit using a vector address, and in the private access mode, one address in the vector address may be routed to one memory bank of the memory unit according to a thread index and all private data for one thread may be located in a same memory bank.

**[0020]** According to an embodiment, each of the plurality of memory ports may be configured to access the memory unit using a vector address, and in the shared access mode, one address in the vector address may be routed in a defined region across memory banks regardless of the thread index and data shared to all threads may be spread in all memory banks.

**[0021]** In another exemplary embodiment, there is provided a method comprising: mapping an execution kernel into a virtual data path at a processor, wherein the execution kernel includes a sequence of instructions to be executed by the processor, and the processor comprises various reconfigurable units that include a gasket memory; chopping the virtual data path into one or more physical data paths; delivering configurations to various reconfigurable units of the

processor for the various reconfigurable units to form the one or more physical data paths to execute the sequence of instructions; and executing the processor to complete the one or more physical data paths by operating the various reconfigurable units according to the configurations, including routing data from one physical data path to the gasket memory to be used in a future physical data path as input.

**[0022]** According to an embodiment, the various reconfigurable units may further comprise a plurality of processing elements, a plurality of switch boxes each associated with a separate processing element, a plurality of memory ports that provide access to a memory unit for the plurality of processing elements, and a plurality of inter-column switch boxes, where each of the various reconfigurable units are reconfigured by applying a next configuration independently from other reconfigurable units.

**[0023]** According to an embodiment, each of the plurality of PEs may comprise a plurality of arithmetic logic units (ALUs) that may be configured to execute same instruction in parallel threads.

**[0024]** According to an embodiment, each of the plurality of memory ports may be configured to operate in a private access mode or a shared access mode during one configuration.

**[0025]** According to an embodiment, the method may further comprise accessing a piece of data stored in the memory unit through the private access mode and the shared access mode in different physical data paths without the piece of data being moved in the memory unit.

**[0026]** According to an embodiment, each of the memory ports may be configured to access the memory unit using a vector address, and in the private access mode, one address in the vector address may be routed to one memory bank of the memory unit according to a thread

index and all private data for one thread may be located in a same memory bank, and in the shared access mode, one address in the vector address may be routed in a defined region across memory banks regardless of the thread index and data shared to all threads may be spread in all memory banks.

**[0027]** According to an embodiment, each of the plurality of PEs may comprise a plurality of data buffers for the plurality of ALUs and may be configured to operate independently during one physical data path.

**[0028]** According to an embodiment, the plurality of PEs may form a PE array and the execution kernel may be mapped into one or more physical data paths on the processor based on a size of the PE array, connections between the plurality of PEs, and memory access capability.

**[0029]** According to an embodiment, the various reconfigurable units may form multiple repetitive columns, and each of the one or more physical data paths may be fitted into the multiple repetitive columns, and data flows between the repetitive columns may be in one direction.

**[0030]** In yet another exemplary embodiment, there is provided a system comprising: a processor, comprising: a sequencer configured to map an execution kernel to be executed by the processor into a virtual data path and chop the virtual data path into one or more physical data paths; a plurality of processing elements (PEs) coupled to the sequencer, each of the plurality of PEs comprising a configuration buffer configured to receive PE configurations for the one or more physical data paths from the sequencer; and a gasket memory coupled to the plurality of PEs and being configured to store data from one of the one or more physical data paths to be used by another physical data path of the one or more physical data paths as input.

**[0031]** According to an embodiment, the processor may further comprise a plurality of switch boxes (SBs) coupled to the sequencer to receive SB configurations for the one or more physical data paths from the sequencer, each of the plurality of SBs being associated with a respective PE of the plurality of PEs and configured to provide input data switching for the respective PE according to the SB configurations.

**[0032]** According to an embodiment, the plurality of switch boxes and their associated PEs may be arranged in a plurality of columns, a first switch box in a first column of the plurality of columns may be coupled between the gasket memory and a first PE in the first column of the plurality of columns, and a second PE in a last column of the plurality of columns may be coupled to the gasket memory.

**[0033]** According to an embodiment, the processor may further comprise a memory unit for providing data storage for the plurality of PEs and a plurality of memory ports each arranged in a separate column of the plurality of columns for the plurality of PEs to access the memory unit.

**[0034]** According to an embodiment, the processor may further comprise a plurality of inter-column switch boxes (ICSBs) coupled to the sequencer to receive ICSB configurations for the one or more physical data paths from the sequencer, the plurality of ICSBs being configured to provide data switching between neighboring columns of the plurality of columns according to the ICSB configurations.

**[0035]** According to an embodiment, the plurality of memory ports (MPs) may be coupled to the sequencer to receive MP configurations for the one or more physical data paths from the sequencer and configured to operate in a private access mode or a shared access mode during one MP configuration.

**[0036]** According to an embodiment, a piece of data stored in the memory unit may be accessed through the private access mode and the shared access mode in different physical data paths of the one or more physical data paths without the piece of data being moved in the memory unit.

**[0037]** According to an embodiment, each of the plurality of columns may comprise one PE and the plurality of PEs may be identical and may form one row of repetitive identical PEs.

**[0038]** According to an embodiment, each of the plurality of columns may comprise two or more PEs and the plurality of PEs form two or more rows.

**[0039]** According to an embodiment, a first row of PEs may be configured to implement a first set of instructions and a second row of PEs may be configured to implement a second set of instructions, at least one instruction of the second set of instructions may be not in the first set of instructions, and the of plurality of columns may be identical and may form repetitive columns.

**[0040]** According to an embodiment, each of the plurality of PEs may comprise a plurality of arithmetic logic units (ALUs) that may be configured to execute same instruction in parallel threads.

**[0041]** According to an embodiment, each of the plurality of PEs may comprise a plurality of data buffers for the plurality of ALUs and may be configured to operate independently.

**[0042]** According to an embodiment, each of the plurality of memory ports may be configured to access the memory unit using a vector address, and in the private access mode, one address in the vector address may be routed to one memory bank of the memory unit according to a thread index and all private data for one thread may be located in a same memory bank.

**[0043]** According to an embodiment, each of the plurality of memory ports may be configured to access the memory unit using a vector address, and in the shared access mode, one address in the vector address may be routed in a defined region across memory banks regardless of the thread index and data shared to all threads may be spread in all memory banks.

**[0044]** According to an embodiment, the plurality of PEs may form a PE array and the execution kernel may be mapped into one or more physical data paths on the processor based on a size of the PE array, connections between the plurality of PEs, and memory access capability.

**[0045]** In yet another exemplary embodiment, there is provided a processor comprising: a plurality of processing elements (PEs); a plurality of switch boxes arranged in a plurality of columns, each of the plurality of switch boxes being associated with a respective PE and configured to provide input data switching for the respective PE; a plurality of memory ports arranged in the plurality of columns and being coupled to a memory unit and a top switch box in each column of the plurality of columns, each of the plurality of memory port being configured to provide data access to the memory unit for one or more switch boxes in a respective column; a plurality of inter-column switch boxes (ICSBs) each coupled to a bottom switch box in each column of the plurality of columns; and a gasket memory with its input coupled to a memory port, a PE, one or more switch boxes and an ICSB in a last column of the plurality of columns, and its output coupled to a memory port, one or more switch boxes and an ICSB in a first column of the plurality of columns.

**[0046]** According to an embodiment, the processor may further comprise a sequencer coupled to the plurality of PEs, the plurality of switch boxes, the plurality of ICSBs, the plurality of memory ports and the gasket memory to deliver configurations to these components.

**[0047]** According to an embodiment, the processor may further comprise a configuration memory coupled to the sequencer to store compiled configurations for the sequencer to decode and deliver.

**[0048]** According to an embodiment, the processor may further comprise a memory unit for providing data storage for the processor.

**[0049]** In another exemplary embodiment, there is provided a processor comprising: a plurality of processing elements (PEs) each comprising a configuration buffer and a plurality of arithmetic logic units (ALUs), and each configured to operate independently according to respective PE configurations stored in the configuration buffer; and a gasket memory coupled to the plurality of PEs and being configured to store at least one PE execution result to be used by at least one of the plurality of PEs during a next PE configuration.

**[0050]** According to an embodiment, the processor may further comprise a plurality of switch boxes each comprising a configuration buffer configured to store switch box configurations, each of the plurality of switch boxes being associated with a respective PE of the plurality of PEs and configured to provide input data switching for the respective PE according to the switch box configurations.

**[0051]** According to an embodiment, the plurality of switch boxes and their associated PEs may be arranged in a plurality of columns, a first switch box in a first column of the plurality of columns may be coupled between the gasket memory and a first PE in the first column of the plurality of columns, and a second PE in a last column of the plurality of columns may be coupled to the gasket memory.

**[0052]** According to an embodiment, the processor may further comprise: a memory unit for providing data storage for the plurality of PEs; and a plurality of memory ports each arranged in a separate column of the plurality of columns for the plurality of PEs to access the memory unit.

**[0053]** According to an embodiment, the processor may further comprise a plurality of inter-column switch boxes (ICSBs) each comprising a configuration buffer configured to store ICSB configurations, the plurality of ICSBs may be configured to provide data switching between neighboring columns of the plurality of columns according to the ICSB configurations.

**[0054]** According to an embodiment, each of the plurality of memory ports (MPs) may comprise a configuration buffer to store MP configurations and may be configured to operate in a private access mode or a shared access mode during one MP configuration.

**[0055]** According to an embodiment, a piece of data stored in the memory unit may be accessed through the private access mode and the shared access mode in different part of a program without the piece of data being moved in the memory unit.

**[0056]** According to an embodiment, each of the plurality of columns may comprise one PE, the plurality of PEs are identical and form one row of repetitive identical PEs.

**[0057]** According to an embodiment, each of the plurality of columns may comprise two or more PEs and the plurality of PEs form two or more rows.

**[0058]** According to an embodiment, a first row of PEs may be configured to implement a first set of instructions and a second row of PEs may be configured to implement a second set of instructions, at least one instruction of the second set of instructions is not in the first set of instructions, and the of plurality of columns may be identical and form repetitive columns.

**[0059]** According to an embodiment, each of the plurality of memory ports may be configured to access the memory unit using a vector address, and in the private access mode, one address in the vector address may be routed to one memory bank of the memory unit according to a thread index and all private data for one thread may be located in a same memory bank.

**[0060]** According to an embodiment, each of the plurality of memory ports may be configured to access the memory unit using a vector address, and in the shared access mode, one address in the vector address may be routed in a defined region across memory banks regardless of the thread index and data shared to all threads may be spread in all memory banks.

**[0061]** According to an embodiment, each of the plurality of PEs may comprise a plurality of data buffers for the plurality of ALUs and may be configured to operate independently.

**[0062]** In yet another exemplary embodiment, there is provided a processor comprising: a plurality of processing elements (PEs) arranged in a plurality of columns; a plurality of switch boxes (SBs) each associated with a separate PE of the plurality of PEs to providing data switching; and a gasket memory coupled to the plurality of PEs and being configured to store at least one PE execution result to be delivered to at least one of the plurality of PEs via a switch box for the PE execute result to be used as input data during a next PE configuration.

**[0063]** In another exemplary embodiment, there is provided a method comprising mapping an execution kernel into a virtual data path. The execution kernel may include a sequence of instructions to be executed by a processor, and the processor may comprise various reconfigurable units that form repetitive columns. The method may further comprises: chopping the virtual data path into one or more physical data paths to fit each physical data path into the repetitive columns respectively; and delivering configurations to various

reconfigurable units of the processor for the various reconfigurable units to form the one or more physical data paths to execute the sequence of instructions.

**[0064]** In yet another exemplary embodiment, there is provided a method comprising mapping an execution kernel into a virtual data path for a processor to execute, the processor comprising various reconfigurable units that form repetitive columns; chopping the virtual data path into a plurality of physical data paths including a first physical data path that fits into the repetitive columns and a second physical data path that fits into the repetitive columns; and delivering configurations to various reconfigurable units for the repetitive columns to form the first physical data path to execute a first part of the execution kernel and form the second physical data path to execute a second part of the execution kernel.

**[0065]** In yet an exemplary embodiment, there is provided a processor comprising: a plurality of reconfigurable units including a plurality of processing elements (PEs) and a plurality of memory ports (MPs) for the plurality of PEs to access a memory unit, each of the plurality of reconfigurable units comprising a configuration buffer and a reconfiguration counter; and a sequencer coupled to the configuration buffer of each of the plurality of reconfigurable units and configured to distribute a plurality of configurations to the plurality of reconfigurable units for the plurality of PEs and the plurality of memory ports to execute a sequence of instructions.

**[0066]** According to an embodiment, each of the plurality of configurations may include a specified number, and the reconfigurable counter of each of the plurality of PEs and the plurality of memory ports may be configured to count for a respective PE or MP to repeat an instruction of the sequence of instructions the specified number of times.

**[0067]** According to an embodiment, the plurality of reconfigurable units may further include a plurality of data switching units, each of the plurality of data switching units may be configured to a data switching setting according to a current data switching configuration the specified number of times.

**[0068]** According to an embodiment, the plurality of reconfigurable units may further include a gasket memory, the gasket memory may comprise a plurality of data buffers, an input configuration buffer, an output configuration buffer, a plurality of input reconfigurable counters and a plurality of output reconfigurable counters, and the gasket memory may be configured to perform reconfiguration for input and output independently.

**[0069]** According to an embodiment, the plurality of configurations may include a first set of configurations for the plurality of reconfigurable units to form a first physical data path and a second set of configurations for the plurality of reconfigurable units to form a second physical data path, the gasket memory may be configured to store data from the first physical data path to be used as input to the second physical data path.

**[0070]** According to an embodiment, each of the plurality of reconfigurable units may be configured to switch to a next configuration independently after its reconfigurable counter reaches the specified number.

**[0071]** According to an embodiment, each of the plurality of memory ports may be configured to operate in a private memory access mode or a shared memory access mode during one configuration.

**[0072]** According to an embodiment, a piece of data stored in the memory unit may be accessed through the private memory access mode and the shared memory access mode in

configurations for different physical data paths without the piece of data being moved in the memory unit.

**[0073]** According to an embodiment, each of the plurality of memory ports may be configured to access the memory unit using a vector address, and in the private memory access mode, one address in the vector address may be routed to one memory bank of the memory unit according to a thread index and all private data for one thread may be located in a same memory bank.

**[0074]** According to an embodiment, each of the plurality of memory ports may be configured to access the memory unit using a vector address, and in the shared memory access mode, one address in the vector address may be routed in a defined region across memory banks regardless of thread index and data shared to all threads may be spread in all memory banks.

**[0075]** According to an embodiment, each of the plurality of PEs may comprise a plurality of arithmetic logic units (ALUs) that may be configured to execute a same instruction in parallel threads.

**[0076]** According to an embodiment, each of the plurality of PEs may comprise a plurality of data buffers and may be configured to operate independently.

**[0077]** According to an embodiment, the plurality of PEs may form two or more rows.

According to an embodiment, a first row of PEs may be configured to implement a first set of instructions and a second row of PEs may be configured to implement a second set of instructions, at least one instruction of the second set of instructions is not in the first set of instructions

**[0078]** According to an embodiment, the plurality of PEs and the plurality of memory ports (MPs) may be arranged in repetitive columns.

**[0079]** According to an embodiment, each of the sequence of instructions may be executed by one of the plurality of PEs or one of the plurality of memory ports as a stage of pipeline according to a respective configuration.

**[0080]** In yet another exemplary embodiment, there is provided a method comprising: delivering a plurality of configurations to a plurality of reconfigurable units of a processor for the plurality of reconfigurable units to form a plurality of physical data paths to execute a sequence of instructions, each of the plurality of configurations including a specified number; repeating a respective operation at each of the plurality of reconfigurable units for the specified number of times, including executing a first instruction of the sequence of instructions at a first reconfigurable processing element (PE) the specified number of times according to a first configuration in a first physical data path; and reconfiguring each of the plurality of reconfiguration units to a new configuration after repeating the respective operation the specified number of times, including executing a second instruction of the sequence of instructions at the first reconfigurable PE the specified number of times according to a second configuration in a second physical data path.

**[0081]** According to an embodiment, the plurality of reconfiguration units may include a plurality of PEs and a plurality of memory ports, and at least one instruction of the sequence of instructions may be a memory access instruction and executed by a memory port the specified number of times before the memory port is reconfigured by applying a next memory port configuration.

**[0082]** According to an embodiment, the plurality of reconfigurable units may further include a plurality of data switching units, and each of the plurality of data switching units may be configured to repeat a respective operation by applying a data switching setting according to a current data switching configuration the specified number of times.

**[0083]** According to an embodiment, the method may further comprise a gasket memory. The gasket memory may comprise a plurality of data buffers, an input configuration buffer, an output configuration buffer, a plurality of input reconfigurable counters and a plurality of output reconfigurable counters. And the gasket memory may be configured to perform reconfiguration for input and output independently.

**[0084]** According to an embodiment, the method may further comprise storing data from the first physical data path in the gasket memory to be used as input to the second physical data path.

**[0085]** According to an embodiment, each of the plurality of PEs may comprise a plurality of arithmetic logic units (ALUs) that may be configured to execute a same instruction in parallel threads.

**[0086]** According to an embodiment, each of the plurality of memory ports may be configured to operate in a private memory access mode or a shared memory access mode during one configuration.

**[0087]** According to an embodiment, the method may further comprise accessing a piece of data stored in a memory unit through the private memory access mode and the shared memory access mode in different physical data paths without the piece of data being moved in the memory unit.

**[0088]** According to an embodiment, each of the plurality of memory ports may be configured to access the memory unit using a vector address, and in the private memory access mode, one address in the vector address is routed to one memory bank of a memory unit according to a thread index and all private data for one thread are located in a same memory bank.

**[0089]** According to an embodiment, each of the plurality of memory ports may be configured to access the memory unit using a vector address, and in the shared memory access mode, one address in the vector address may be routed in a defined region across memory banks regardless of thread index and data shared to all threads are spread in all memory banks.

**[0090]** According to an embodiment, each of the plurality of PEs may comprise a plurality of data buffers and may be configured to operate independently during one physical data path configuration.

**[0091]** According to an embodiment, the plurality of PEs may form a PE array and the sequence of instructions may be mapped into one or more physical data paths on the processor based on a size of the PE array, connections between the plurality of PEs, and memory access capability.

**[0092]** In yet another exemplary embodiment, there is provided a method comprising:  
delivering a first set of configurations to a plurality of reconfigurable units of a processor for the plurality of reconfigurable units to form a first physical data path to execute a first part of a sequence of instructions, each of the first set of configurations including a specified number;  
delivering a second set of configurations to the plurality of reconfigurable units for the plurality of reconfigurable units to form a second physical data path to execute a second part of a sequence of instructions, each of the second set of configurations including the specified number;  
applying the first set of configurations at the plurality of reconfigurable units for each

of the plurality of reconfiguration units to repeat a respective operation the specified number of times to execute the first physical data path; storing data from the first physical data path to a gasket memory; and applying the second set of configurations at the plurality of reconfigurable units for each of the plurality of reconfiguration units to repeat a respective operation the specified number of times to execute the second physical data path, with the data stored in the gasket memory as input to the second physical data path.

**[0093]** According to an embodiment, the gasket memory may comprise a plurality of data buffers, an input configuration buffer, an output configuration buffer, a plurality of input reconfigurable counters and a plurality of output reconfigurable counters, and wherein the gasket memory may be configured to perform reconfiguration for input and output independently.

**[0094]** According to an embodiment, the plurality of reconfiguration units may include a plurality of PEs and a plurality of memory ports, and at least one instruction of the sequence of instructions may be a memory access instruction and executed by a memory port the specified number of times before the memory port may be reconfigured by applying a next memory port configuration.

**[0095]** According to an embodiment, the plurality of reconfigurable units may further include a plurality of data switching units, and each of the plurality of data switching units may be configured to repeat a respective operation by applying a data switching setting according to a current data switching configuration the specified number of times.

**[0096]** According to an embodiment, each of the plurality of PEs may comprise a plurality of arithmetic logic units (ALUs) that may be configured to execute a same instruction in parallel threads.

**[0097]** According to an embodiment, each of the plurality of memory ports may be configured to operate in a private memory access mode or a shared memory access mode during one configuration.

**[0098]** According to an embodiment, the method may further comprise accessing a piece of data stored in a memory unit through the private memory access mode and the shared memory access mode in different physical data paths without the piece of data being moved in the memory unit.

**[0099]** According to an embodiment, each of the plurality of memory ports may be configured to access the memory unit using a vector address, wherein in the private memory access mode, one address in the vector address may be routed to one memory bank of a memory unit according to a thread index and all private data for one thread are located in a same memory bank.

**[00100]** According to an embodiment, each of the plurality of memory ports may be configured to access the memory unit using a vector address, and in the shared memory access mode, one address in the vector address may be routed in a defined region across memory banks regardless of thread index and data shared to all threads may be spread in all memory banks.

**[00101]** According to an embodiment, each of the plurality of PEs may comprise a plurality of data buffers and may be configured to operate independently during one physical data path configuration.

**[00102]** According to an embodiment, the plurality of PEs may form a PE array and the sequence of instructions may be mapped into one or more physical data paths on the processor based on a size of the PE array, connections between the plurality of PEs, and memory access capability.

**[00103]** In another exemplary embodiment, there is provided a processor comprising: a plurality of processing elements (PEs) each having a plurality of arithmetic logic units (ALUs) that are configured to execute a same instruction in parallel threads; and a plurality of memory ports (MPs) for the plurality of PEs to access a memory unit, each of the plurality of MPs comprising an address calculation unit configured to generate respective memory addresses for each thread to access a different memory bank in the memory unit.

**[00104]** According to an embodiment, the address calculation unit may have a first input coupled to a base address input that provides a base address common to all threads, a second input coupled to a vector address that provides address offsets for each thread individually, and a third input coupled to a counter that is configured to provide thread indexes.

**[00105]** According to an embodiment, one address in the vector address may be routed to one memory bank according to a thread index.

**[00106]** According to an embodiment, the memory unit may comprise a plurality of memory caches each associated with a different memory bank.

**[00107]** According to an embodiment, each of the plurality of memory ports may be coupled to the plurality of memory caches.

**[00108]** According to an embodiment, each memory bank may comprise a plurality of memory words and a cache miss in a memory cache causes a word to be fetched from a memory bank associated with the memory cache.

**[00109]** According to an embodiment, each of the plurality of PEs may comprise a plurality of data buffers to store data for each thread separately.

**[00110]** According to an embodiment, the processor may further comprise a sequencer coupled to the plurality of memory ports, and each of the plurality of memory ports may comprise a configuration buffer to receive one or more configurations from the sequencer, and each memory port may be configured to provide a same memory access pattern during one configuration.

**[00111]** According to an embodiment, consecutive data pieces for one thread may be located in one word of a memory bank and continue in a next word of the memory bank.

**[00112]** According to an embodiment, consecutive data pieces for one thread may be located in a same position of consecutive words of a memory bank.

**[00113]** According to an embodiment, the plurality of MPs may be configured in a one column mode, in which one MP may be configured to access the memory unit for all concurrent threads in one PE and the address offsets may be independent for each thread.

**[00114]** According to an embodiment, the plurality of MPs may be configured in a linear mode, in which multiple MPs may be configured to access the memory unit, wherein a first MP may be configured to access the memory unit for all concurrent threads in a first PE, and a

second MP may be configured to access the memory unit for all concurrent threads in a second PE, the address offsets in the second MP may be linear to the address offsets in the first MP.

**[00115]** According to an embodiment, the plurality of MPs may be configured in a reverse linear mode, in which multiple MPs may be configured to access the memory unit, wherein a first MP may be configured to access the memory unit for all concurrent threads in a first PE, and a second MP may be configured to access the memory unit for all concurrent threads in a second PE, the address offsets in the second MP may be reverse linear to the address offsets in the first MP.

**[00116]** According to an embodiment, the plurality of MPs may be configured in an overlap mode, in which multiple MPs are configured to access the memory unit, wherein a first MP is configured to access the memory unit for all concurrent threads in a first PE, and a second MP is configured to access the memory unit for all concurrent threads in a second PE, the address offsets in the second MP have overlap with the address offsets in the first MP.

**[00117]** According to an embodiment, the plurality of MPs may be configured in a non-unity stride mode, in which multiple MPs are configured to access the memory unit, wherein a first MP is configured to access the memory unit for all concurrent threads in a first PE, and a second MP is configured to access the memory unit for all concurrent threads in a second PE, the address offsets in the second MP and the address offsets in the first MP may be spaced by a stride.

**[00118]** According to an embodiment, the plurality of MPs may be configured in a random mode, in which multiple MPs may be configured to access the memory unit, and address offsets in different MPs may be random numbers.

**[00119]** According to an embodiment, the memory unit may comprise a plurality of memory caches each associated with a different memory bank and the random numbers may be within a range depending on a size of a memory cache.

**[00120]** According to an embodiment, the memory unit may be configured to be used as registers to store spilled variables for register spilling.

**[00121]** In another exemplary embodiment, there is provided a method comprising: generating a plurality of memory addresses by an address calculation unit in a memory port of a plurality of memory ports, wherein the plurality of memory ports provide access to a memory unit for a plurality of processing elements (PEs) each having a plurality of arithmetic logic units (ALUs) that are configured to execute a same instruction in parallel threads; and accessing a plurality of memory banks in the memory unit using the plurality of memory addresses with each thread accessing a different memory bank in the memory unit.

**[00122]** According to an embodiment, the address calculation unit may have a first input coupled to a base address input that provides a base address common to all threads, a second input coupled to a vector address that provides address offsets for each thread individually, and a third input coupled to a counter that is configured to provide thread indexes, and the address calculation unit may be configured to generate the plurality of memory addresses using the first input, the second input and the third input.

**[00123]** According to an embodiment, one address in the vector address may be routed to one memory bank according to a thread index.

**[00124]** According to an embodiment, the memory unit may comprise a plurality of memory caches each associated with a different memory bank, and accessing the plurality of memory banks in the memory unit may comprise accessing the plurality of memory caches.

**[00125]** According to an embodiment, each of the plurality of memory ports may be coupled to the plurality of memory caches.

**[00126]** According to an embodiment, the method may further comprise fetching a word from a plurality of words of a memory bank when there is a cache miss in a memory cache associated with the memory bank.

**[00127]** According to an embodiment, the method may further comprise storing data for each thread in a separate data buffer in each of the plurality of PEs.

**[00128]** According to an embodiment, the method may further comprise receiving one or more configurations by the memory port from a sequencer, and the memory port may be configured to provide a same memory access pattern during one configuration.

**[00129]** According to an embodiment, consecutive data pieces for one thread may be located in one word of a memory bank and continue in a next word of the memory bank.

**[00130]** According to an embodiment, consecutive data pieces for one thread may be located in a same position of consecutive words of a memory bank.

**[00131]** According to an embodiment, accessing the plurality of memory banks in the memory unit may be in a one column mode, in which one MP may be configured to access the memory unit for all concurrent threads in one PE and the address offsets are independent for each thread.

**[00132]** According to an embodiment, accessing the plurality of memory banks in the memory unit may be in a linear mode, in which multiple MPs may be configured to access the memory unit, and a first MP may be configured to access the memory unit for all concurrent threads in a first PE, and a second MP may be configured to access the memory unit for all concurrent threads in a second PE, address offsets in the second MP may be linear to address offsets in the first MP.

**[00133]** According to an embodiment, accessing the plurality of memory banks in the memory unit may be in a reverse linear mode, in which multiple MPs may be configured to access the memory unit, and a first MP may be configured to access the memory unit for all concurrent threads in a first PE, and a second MP may be configured to access the memory unit for all concurrent threads in a second PE, address offsets in the second MP may be reverse linear to address offsets in the first MP.

**[00134]** According to an embodiment, accessing the plurality of memory banks in the memory unit may be in an overlap mode, in which multiple MPs may be configured to access the memory unit, and a first MP may be configured to access the memory unit for all concurrent threads in a first PE, and a second MP may be configured to access the memory unit for all concurrent threads in a second PE, address offsets in the second MP may have overlap with address offsets in the first MP.

**[00135]** According to an embodiment, accessing the plurality of memory banks in the memory unit may be in a non-unity stride mode, in which multiple MPs may be configured to access the memory unit, and a first MP may be configured to access the memory unit for all concurrent threads in a first PE, and a second MP may be configured to access the memory unit

for all concurrent threads in a second PE, address offsets in the second MP and address offsets in the first MP may be spaced by a stride.

**[00136]** According to an embodiment, accessing the plurality of memory banks in the memory unit may be in a random mode, in which multiple MPs may be configured to access the memory unit, and address offsets in different MPs may be random numbers.

**[00137]** According to an embodiment, the memory unit may comprise a plurality of memory caches each associated with a different memory bank and the random numbers may be within a range depending on a size of a memory cache.

**[00138]** According to an embodiment, the method may further comprise storing variables for register spilling.

**[00139]** In an exemplary embodiment, there is provided a processor comprising: a memory unit comprising a plurality of memory banks; a plurality of processing elements (PEs) each having a plurality of arithmetic logic units (ALUs) that are configured to execute a same instruction in parallel threads; and a plurality of memory ports (MPs) for the plurality of PEs to access the memory unit, each of the plurality of MPs comprising an address calculation unit configured to generate respective memory addresses for each thread to access a different memory bank in the memory unit.

**[00140]** According to an embodiment, the address calculation unit may have a first input coupled to a base address input that provides a base address common to all threads, a second input coupled to a vector address that provides address offsets for each thread individually, and a third input coupled to a counter that is configured to provide thread indexes.

**[00141]** In an exemplary embodiment, there is provided a processor comprising: a processing element (PE) having a plurality of arithmetic logic units (ALUs) that are configured to execute a same instruction in parallel threads; and a memory port (MP) for the PE to access a memory unit, the MP comprising an address calculation unit configured to generate respective memory addresses for each thread to access a different memory bank in the memory unit.

**[00142]** According to an embodiment, the PE may be one of a plurality of PEs that each has a plurality of ALUs configured to execute a same instruction in parallel threads.

**[00143]** According to an embodiment, the MP may be one of a plurality of MPs that each has an address calculation unit configured to generate respective memory addresses for each thread in one of plurality of PEs to access a different memory bank in the memory unit.

**[00144]** In yet another exemplary embodiment, there is provided a method comprising: generating a plurality of memory addresses by an address calculation unit in a memory port, wherein the memory port provides access to a memory unit for a processing element (PE) having a plurality of arithmetic logic units (ALUs) configured to execute a same instruction in parallel threads; and accessing a plurality of memory banks in the memory unit using the plurality of memory addresses with each thread accessing a different memory bank in the memory unit.

**[00145]** According to an embodiment, the PE may be one of a plurality of PEs that each may have a plurality of ALUs configured to execute a same instruction in parallel threads.

**[00146]** According to an embodiment, the MP may be one of a plurality of MPs that each may have an address calculation unit configured to generate respective memory addresses for each thread in one of plurality of PEs to access a different memory bank in the memory unit.

**[00147]** In an exemplary embodiment, there is provided a processor comprising: a plurality of processing elements (PEs) each having a plurality of arithmetic logic units (ALUs) that are configured to execute a same instruction in parallel threads; and a plurality of memory ports (MPs) for the plurality of PEs to access a memory unit, each of the plurality of MPs comprising an address calculation unit configured to generate respective memory addresses for each thread to access a common area in the memory unit.

**[00148]** According to an embodiment, the address calculation unit may have a first input coupled to a base address input that provides a base address common to all threads, and a second input coupled to a vector address that provides address offsets for each thread individually.

**[00149]** According to an embodiment, the address calculation unit may be configured to generate a number of memory addresses that match a number of threads in a PE.

**[00150]** According to an embodiment, each of the plurality of MPs may further comprise a plurality of selection units coupled to the number of memory addresses, each of the plurality of selection units may be configured to select zero or more memory addresses to be routed to one memory bank of the memory unit.

**[00151]** According to an embodiment, each selection unit may be configured with a mask for a different memory bank of the memory unit.

**[00152]** According to an embodiment, one MP may be configured to access the memory unit for all threads in one PE and the address offsets may be same for all threads.

**[00153]** According to an embodiment, multiple MPs may be configured to access the memory units for threads in different PEs, the address offsets may be the same within one MP but different for different MPs.

**[00154]** According to an embodiment, one MP may be configured to access the memory unit for all threads in one PE and the address offsets may be sequential in the MP.

**[00155]** According to an embodiment, multiple MPs may be configured to access the memory unit for threads in different PEs, the address offsets may be the sequential within each MP respectively.

**[00156]** According to an embodiment, one MP may be configured to access the memory unit for all threads in one PE, the address offsets may be sequential with discontinuity.

**[00157]** According to an embodiment, the plurality of MPs may be configured to access the memory unit for different threads in different PEs, the address offsets may be sequential with discontinuity in each of the MPs respectively.

**[00158]** According to an embodiment, one MP may be configured to access the memory unit for all threads in one PE, the address offsets may be linear with non-unity stride.

**[00159]** According to an embodiment, multiple MPs may be configured to access the memory unit for all threads in one PE, the address offsets may be random but within a small range  $C$  to  $C+R$  dependent on a size of the memory cache.

**[00160]** According to an embodiment, multiple MPs may be configured to access the memory unit for threads in different PEs, the address offsets may be random but with a small range  $C$  to  $C+R$  dependent on a size of the memory cache.

**[00161]** According to an embodiment, the common area may include all memory banks of the memory unit.

**[00162]** According to an embodiment, the memory unit may comprise a plurality of memory caches each associated with a different memory bank.

**[00163]** According to an embodiment, each of the plurality of memory ports may be coupled to the plurality of memory caches.

**[00164]** According to an embodiment, each memory bank may comprise a plurality of memory words and a cache miss in a memory cache may cause a memory word to be fetched from a memory bank associated with the memory cache.

**[00165]** According to an embodiment, each of the plurality of PEs may comprise a plurality of data buffers to store data for each thread separately.

**[00166]** According to an embodiment, the processor may comprise a sequencer coupled to the plurality of memory ports, and each of the plurality of memory ports may comprise a configuration buffer to receive one or more configurations from the sequencer, and each memory port may be configured to provide a same memory access pattern during one configuration.

**[00167]** In yet another exemplary embodiment, there is provided a method comprising: generating a plurality of memory addresses by an address calculation unit in a memory port of a plurality of memory ports, wherein the plurality of memory ports provide access to a memory unit for a plurality of processing elements (PEs) each having a plurality of arithmetic logic units (ALUs) that are configured to execute a same instruction in parallel threads; and accessing a

plurality of memory banks in the memory unit using the plurality of memory addresses with all threads accessing a common area in the memory unit.

**[00168]** According to an embodiment, the address calculation unit may take a base address common to all threads as a first input and a vector address that provides address offsets for each thread individually as a second input to generate the plurality of memory address.

**[00169]** According to an embodiment, the address calculation unit may be configured to generate a number of memory addresses that match a number of threads in a PE.

**[00170]** According to an embodiment, accessing the plurality of memory banks may comprise selecting zero or more memory addresses to be routed to one memory bank of the memory unit using a plurality of selection units respectively.

**[00171]** According to an embodiment, each selection unit may be configured with a mask for a different memory bank of the memory unit.

**[00172]** According to an embodiment, one MP may be configured to access the memory unit for all threads in one PE and the address offsets may be same for all threads.

**[00173]** According to an embodiment, multiple MPs may be configured to access the memory units for threads in different PEs, the address offsets may be the same within one MP but different for different MPs.

**[00174]** According to an embodiment, one MP may be configured to access the memory unit for all threads in one PE and the address offsets may be sequential in the MP.

**[00175]** According to an embodiment, multiple MPs may be configured to access the memory unit for threads in different PEs, the address offsets may be the sequential within each MP respectively.

**[00176]** According to an embodiment, one MP may be configured to access the memory unit for all threads in one PE, the address offsets may be sequential with discontinuity.

**[00177]** According to an embodiment, the plurality of MPs may be configured to access the memory unit for different threads in different PEs, the address offsets may be sequential with discontinuity in each of the MPs respectively.

**[00178]** According to an embodiment, one MP may be configured to access the memory unit for all threads in one PE, the address offsets may be linear with non-unity stride.

**[00179]** According to an embodiment, multiple MPs may be configured to access the memory unit for all threads in one PE, the address offsets may be random but within a small range  $C$  to  $C+R$  dependent on a size of the memory cache.

**[00180]** According to an embodiment, multiple MPs may be configured to access the memory unit for threads in different PEs, the address offsets may be random but with a small range  $C$  to  $C+R$  dependent on a size of the memory cache.

**[00181]** According to an embodiment, the common area may include all memory banks of the memory unit.

**[00182]** According to an embodiment, the memory unit may comprise a plurality of memory caches each associated with a different memory bank.

**[00183]** According to an embodiment, each of the plurality of memory ports may be coupled to the plurality of memory caches.

**[00184]** According to an embodiment, each memory bank may comprise a plurality of memory words and a cache miss in a memory cache may cause a memory word to be fetched from a memory bank associated with the memory cache.

**[00185]** According to an embodiment, each of the plurality of PEs may comprise a plurality of data buffers to store data for each thread separately.

**[00186]** According to an embodiment, the method may further comprise receiving one or more configurations from a sequencer for each of the plurality of memory ports, wherein each memory port is configured to provide a same memory access pattern during one configuration.

**[00187]** In yet an exemplary embodiment, there is provided a processor comprising: a memory unit comprising: a plurality of memory banks; a plurality of processing elements (PEs) each having a plurality of arithmetic logic units (ALUs) that are configured to execute a same instruction in parallel threads; and a plurality of memory ports (MPs) for the plurality of PEs to access the memory unit, each of the plurality of MPs comprising an address calculation unit configured to generate respective memory addresses for each thread to access a common area across the plurality of memory banks in the memory unit.

**[00188]** In an exemplary embodiment, there is provided a processor comprising: a processing element (PE) having a plurality of arithmetic logic units (ALUs) that are configured to execute a same instruction in parallel threads; and a memory port (MP) for the PE to access a memory unit, the MP comprising an address calculation unit configured to generate respective memory addresses for each thread to access a common area in the memory unit.

**[00189]** According to an embodiment, the PE may be one of a plurality of PEs that each may have a plurality of ALUs configured to execute a same instruction in parallel threads.

**[00190]** According to an embodiment, the MP may be one of a plurality of MPs that each may have an address calculation unit configured to generate respective memory addresses for each thread in one of plurality of PEs to access the common area in the memory unit.

**[00191]** In yet another exemplary embodiment, there is provided a method comprising: generating a plurality of memory addresses by an address calculation unit in a memory port, wherein the memory port provides access to a memory unit for a processing element (PE) having a plurality of arithmetic logic units (ALUs) configured to execute a same instruction in parallel threads; and accessing a plurality of memory banks in the memory unit using the plurality of memory addresses with each thread accessing a common area in the memory unit.

**[00192]** According to an embodiment, the PE may be one of a plurality of PEs that each may have a plurality of ALUs configured to execute a same instruction in parallel threads.

**[00193]** According to an embodiment, the MP may be one of a plurality of MPs that each may have an address calculation unit configured to generate respective memory addresses for each thread in one of plurality of PEs to access the common area in the memory unit.

**[00194]** In yet an exemplary embodiment, there is provided a processor comprising: a plurality of processing elements (PEs) each comprising: an arithmetic logic unit (ALU); a data buffer associated with the ALU; and an indicator associated with the data buffer to indicate whether a piece of data inside the data buffer is to be reused for repeated execution of a same instruction as a pipeline stage.

**[00195]** According to an embodiment, the processor further comprise a plurality of memory ports (MPs) for the plurality of PEs to access a memory unit, each of the plurality of MPs may comprise an address calculation unit configured to generate respective memory

addresses for each thread to access a common area in the memory unit, and one MP of the plurality of MPs responsible for loading a piece of data from the memory unit for the piece of data to be reused at a PE may be configured to load the piece of data only once.

**[00196]** According to an embodiment, the MP responsible for loading the piece of data to be reused may be configured to determine that the piece of data is to be reused by determining that multiple threads to be executed at the PE are loading the piece of data using a same memory address.

**[00197]** According to an embodiment, at least one piece of data to be reused may be an execution result generated by one PE of the plurality of PEs.

**[00198]** According to an embodiment, each PE of the plurality of PEs may further comprise a configuration buffer to store configurations for each PE and a reconfiguration counter to count a number of repeated execution, each configuration may specify an instruction to be executed by a respective PE and the number for the instruction to be repeated during a respective configuration.

**[00199]** According to an embodiment, the ALU may be a vector ALU and the data buffer may be a vector data buffer, each data buffer of the vector data buffer may be associated with one ALU of the vector ALU.

**[00200]** According to an embodiment, the processor may further comprise a plurality of memory ports (MPs) for the plurality of PEs to access a memory unit, each of the plurality of MPs may comprise an address calculation unit configured to generate respective memory addresses for each thread to access a common area in the memory unit, and each MP of the plurality of MPs may comprise at least one data buffer for temporarily storing data loaded from

the memory unit and each of the at least one data buffer may have an indicator associated therewith to indicate whether a piece of data stored therein is to be reused for other load operation.

**[00201]** According to an embodiment, each PE may comprise a plurality of data buffers associated with the ALU, each of the plurality data buffers may be configured to store a separate input for the ALU and may have an associated indicator to indicate whether a respective input is to be reused for repeated execution.

**[00202]** In yet another exemplary embodiment, there is provided a method comprising: determining that a piece of data is to be shared and reused by all threads at a processing element (PE) of a processor during one configuration applied at the PE; loading the piece of data once into a data buffer of the PE; setting an indicator associated with the data buffer to indicate that the piece of data is to be reused; and executing a same instruction with the piece of data as an input at the PE repeatedly a number of times as a pipeline stage, the same instruction and the number being specified by the configuration.

**[00203]** According to an embodiment, the method may further comprise loading the piece of data from a memory unit for the piece of data to be loaded into the data buffer of the PE, wherein the processor comprises a plurality of PEs and a plurality of memory ports (MPs) for the plurality of PEs to access the memory unit, wherein one MP of the plurality of MPs responsible for loading the piece of data from the memory unit is configured to load the piece of data only once.

**[00204]** According to an embodiment, the method may further comprise generating the piece of data as an execution result by one of a plurality of PEs of the processor.

**[00205]** According to an embodiment, the method may further comprise receiving the configuration and storing the configuration in a configuration buffer of the PE, wherein the configuration may specify an instruction to be executed by the PE and the number for the instruction to be repeated.

**[00206]** According to an embodiment, determining that the piece of data may be to be shared and reused by all threads at the PE comprises determining that the all threads may be using a same memory address to access the piece of data.

**[00207]** According to an embodiment, the method may further comprise loading the piece of data once into a data buffer of a memory port, the memory port providing access to a memory unit for the PE; setting an indicator associated with the data buffer of the memory port to indicate that the piece of data is to be reused for other load operations accessing a same memory address.

**[00208]** In yet an exemplary embodiment, there is provided a processor comprising: a plurality of processing elements (PEs) each comprising: a vector arithmetic logic unit (ALU) including a plurality of ALUs; a plurality of data buffer associated with each of the plurality of ALU; and a plurality of indicators each associated with a separate data buffer to indicate whether a piece of data inside a respective data buffer is to be reused for repeated execution of a same instruction as a pipeline stage at a respective PE.

**[00209]** According to an embodiment, the processor may further comprise a plurality of memory ports (MPs) for the plurality of PEs to access a memory unit, each of the plurality of MPs may comprise an address calculation unit configured to generate respective memory addresses for each thread to access a common area in the memory unit, wherein one MP of the

plurality of MPs responsible for loading a piece of data from the memory unit for the piece of data to be reused at a PE may be configured to load the piece of data only once.

**[00210]** According to an embodiment, the MP responsible for loading the piece of data to be reused may be configured to determine that the piece of data may be to be reused by determining that multiple threads to be executed at the PE may be loading the piece of data using a same memory address.

**[00211]** According to an embodiment, at least one piece of data to be reused may be an execution result generated by one PE of the plurality of PEs.

**[00212]** According to an embodiment, the processor may further comprise a plurality of memory ports (MPs) for the plurality of PEs to access a memory unit, each of the plurality of MPs may comprise an address calculation unit configured to generate respective memory addresses for each thread to access a common area in the memory unit, and each MP of the plurality of MPs may comprise at least one data buffer for temporarily storing data loaded from the memory unit and each of the at least one data buffer may have an indicator associated therewith to indicate whether a piece of data stored therein may be to be reused for other load operation.

**[00213]** According to an embodiment, each of the plurality of data buffers may be a vector data buffer having a plurality of data buffer units and a piece of data to be reused for repeated execution may be duplicated in all data buffer units of one vector data buffer.

**[00214]** According to an embodiment, each of the plurality of data buffers may be a vector data buffer having a plurality of data buffer units and a piece of data to be reused for repeated execution may be stored only in one data buffer unit of one vector data buffer.

**[00215]** In yet another exemplary embodiment, there is provided a method comprising: receiving a first configuration and a second configuration at a reconfigurable unit of a processor, the reconfigurable unit having a configuration buffer to store the first configuration and the second configuration; executing a first operation a first number of times according to the first configuration, the first configuration being part of a first physical data path for executing a first part of a sequence of instructions; and reconfiguring the reconfigurable unit to execute a second operation a second number of times according to the second configuration, the second configuration being part of a second physical data path for executing a second part of a sequence of instructions.

**[00216]** In another exemplary embodiment, there is provided a method comprising: executing a first instruction at a reconfigurable processing element a number of times according to a first configuration, the reconfigurable processing element being configured to be part of a first physical data path during the first configuration; delivering an execution result from the reconfigurable processing element to a gasket memory to temporarily store the execution result after each execution of the first instruction; and feeding the execution result stored in the gasket memory to a second physical data path.

### **Brief Description of Figures**

**[00217]** FIG. 1 schematically shows a prior art CGRA with a PE array.

**[00218]** FIG. 2 schematically shows a processor in accordance with an embodiment of the present disclosure.

**[00219]** FIG. 3A schematically shows a memory system for a processor in accordance with an embodiment of the present disclosure.

**[00220]** FIG. 3B schematically shows a first memory mapping for a private memory access mode in accordance with an embodiment of the present disclosure.

**[00221]** FIG. 3C schematically shows a second memory mapping for a private memory access mode in accordance with an embodiment of the present disclosure.

**[00222]** FIG. 3D schematically shows a memory mapping for a shared memory access in accordance with an embodiment of the present disclosure.

**[00223]** FIG. 4A schematically shows a first memory access configuration for a memory port in accordance with an embodiment of the present disclosure.

**[00224]** FIG. 4B schematically shows a second memory access configuration for a memory port in accordance with an embodiment of the present disclosure.

**[00225]** FIG. 5 schematically shows a switch box for a processor in accordance with an embodiment of the present disclosure.

**[00226]** FIG. 6A schematically shows a processing element for a processor in accordance with an embodiment of the present disclosure.

**[00227]** FIG. 6B schematically shows another processing element for a processor in accordance with an embodiment of the present disclosure.

**[00228]** FIG. 7 schematically shows an inter-column switch box for a processor in accordance with an embodiment of the present disclosure.

**[00229]** FIG. 8 schematically shows a gasket memory for a processor in accordance with an embodiment of the present disclosure.

**[00230]** FIG. 9A schematically shows a dependency graph of an execution kernel in accordance with an embodiment of the present disclosure.

**[00231]** FIG. 9B schematically shows the dependency graph of the execution kernel of FIG. 9A mapped into a virtual data path for a processor in accordance with an embodiment of the present disclosure.

**[00232]** FIG. 9C schematically shows the virtual data path of FIG. 9B divided into physical data paths for a processor in accordance with an embodiment of the present disclosure.

**[00233]** FIG. 10 schematically shows a pipeline operation for a processor in accordance with an embodiment of the present disclosure.

**[00234]** FIGs. 11A-11I schematically shows a configuration process for a processor in accordance with an embodiment of the present disclosure.

**[00235]** FIG. 12A schematically shows a pipeline operation for executing a stream of instructions using a processor in accordance with an embodiment of the present disclosure.

**[00236]** FIG. 12B schematically shows a reduced pipeline operation for executing a stream of instructions using a processor in accordance with an embodiment of the present disclosure.

**[00237]** FIG. 13 is a flowchart of a method for executing an execution kernel in accordance with an embodiment of the present disclosure.

**[00238]** FIG. 14 is a flowchart of a method for reconfiguration in accordance with an embodiment of the present disclosure.

**[00239]** FIG. 15 is a flowchart of another method for reconfiguration in accordance with an embodiment of the present disclosure.

**[00240]** FIG. 16 is a flowchart of another method for reconfiguration in accordance with an embodiment of the present disclosure.

**[00241]** FIG. 17 is a flowchart of a method for accessing memory in accordance with an embodiment of the present disclosure.

**[00242]** FIG. 18 is a flowchart of another method for accessing memory in accordance with an embodiment of the present disclosure.

**[00243]** FIG. 19 is a flowchart of a method for reusing a piece of data in accordance with an embodiment of the present disclosure.

### **Detailed Description**

**[00244]** Reference will now be made in detail to the embodiments of the present teaching, examples of which are illustrated in the accompanying drawings. Like elements in the various figures are denoted by like reference numerals for consistency. While the present teaching will be described in conjunction with the embodiments, it will be understood that they are not intended to limit the present teaching to these embodiments. On the contrary, the present teaching is intended to cover alternatives, modifications, and equivalents, which may be included within the spirit and scope of the present teaching as defined by the appended claims.

**[00245]** In addition, in the following detailed description of embodiments of the present teaching, numerous specific details are set forth in order to provide a thorough understanding of the present teaching. However, it will be recognized by one of ordinary skill in the art that the present teaching may be practiced without these specific details. In other instances, well-known methods, procedures, components, and circuits have not been described in detail as not to unnecessarily obscure aspects of the embodiments of the present teaching.

**[00246]** FIG. 2 schematically shows a processor 200 according to an embodiment of a processor. The processor 200 may comprise a direct memory access (DMA) module 202, a configuration memory 204, a sequencer 206, a control processor 208, a program cache 210, a memory unit 212, a PE array 214 and a gasket memory 216. The DMA module 202 may be coupled to an external bus 230 and may be controlled by the control processor 208. The DMA module 202 may be responsible to move executable instructions and non-executable data in and out from the external bus 230. The program cache 210 may store instructions and data to be used by the control processor 208 to control the operation of the DMA module 202. In one embodiment, the instructions and data stored in the program cache 210 may be used by the control processor 208 to handle sequencer programs.

**[00247]** It should be noted that as used herein, a “coupling” between two components, such as one component being “coupled” to another component may refer to an electronic connection between two components, which may include but not limited to, by electronic wiring, through an electronic element (e.g., a resistor, a transistor), etc. Moreover, in some embodiments, the processor 200 may be configured for massive thread level parallel processing. For example, one processing element (PE) in the PE array 214 may comprise multiple arithmetic logic units (ALUs) that may be configured to perform same operation but on different data (e.g., each in a separate thread). That is, in these embodiments with multiple ALUs, each PE may be configured to operate in a Single Instruction Multiple Threads (SIMT) fashion. In one embodiment, a PE with a vector address and a vector data input may generate vector data output. In some embodiments, a thread may also be referred to as a stream.

**[00248]** To provide data for multiple threads to be executed concurrently, in some embodiments, some relevant electronic connections between components of the processor 200 may be in vectors. For example, a vector address of  $H \times G$  may have  $H$  number of  $G$  bits addresses, a vector data connection of  $K \times W$  may have  $K$  number of  $W$  bits data. It should also be noted that although not shown in any of the figures, data or address connections between different components may be accompanied by one or more signal lines. For example, a busy signal line may exist between a first component and a second component, and may be used by the first component to send a busy signal to the second component indicating that the first component is not ready to accept valid data or address signals. Moreover, a valid signal line may also exist between the first and second components, and may be used by the second component to send a valid signal to the first component indicating that valid data or address signals have been put on the connection wires.

**[00249]** The configuration memory 204 may store data path programs consists of executable instructions and/or data loading instructions for one or more data paths. In one embodiment, the data path programs stored in the configuration memory 204 may be sequence(s) of compiled instructions. For example, a data path program may include instructions to be executed by the PE array 214, which represents configuration information to be executed by which PEs when conditions are met, and how each data path components may hold or transmit data

**[00250]** The sequencer 206 may decode the instruction stored in the configuration memory 204 and move a decoded instruction into the memory unit 212 and a physical data path. The physical data path may include various components of the PE array 214 (e.g.,

components of the PE array 214 that will be involved in the executing on, staging and/or movement of data) and the gasket memory 216. The decoded instruction may be delivered to various components in a package, which may be referred to as a configuration package or simply a configuration. In addition to the decoded instruction, a configuration package for one component may include some other parameters (e.g., a number specifying how many times an instruction is to be repeatedly executed or how many times data passes through a data switching unit in one configuration setting). In one embodiment, a physical data path configuration may be referred to as a physical data path program, which may comprise individual configurations for various components included in a physical data path. Although not shown, there may be a configuration bus connecting the sequencer 206 to the components of a data path for individual configurations to be delivered to these components via the bus respectively.

**[00251]** The memory unit 212 may be a data staging area to store data received from the external bus 230 and store execution result data generated by the PE array 214 (before these results may be transmitted away via the external bus 230). In some embodiments, the memory unit 212 may be an in processor cache for a large memory system external of the processor 200. The PE array 214 may comprise a plurality of memory ports (MPs) 220.1-220.N, a plurality of switch boxes (SBs) 222.1-222.N, a plurality of processing elements (PEs) 218.1-218.N and a plurality of inter-column switch boxes (ICSBs) 224.1-224.N. In the example shown in FIG. 2, the PE array 214 may comprise one row of processing elements 218.1-218.N (e.g., one dimensional array), which may include N PEs with N being an integer number. In the embodiments described herein, the number N may be 32. But this is an example and it may be another

integer number in other embodiments, such as, but not limited to, 16, 32, 64, etc. Although these example numbers given may be power of 2, but the number of PEs in one row need not be a power of 2 in one embodiment. In some embodiments, the PE array 214 may be a two-dimensional array that may comprise multiple rows of processing elements (e.g., one or more rows of PEs may be positioned underneath the PEs 218.1-218.N). It should be noted that the PE array 214 may be a composite of MPs, SBs, ICSBs and PEs for illustration purpose and used to refer to these components collectively. In some embodiments, the memory unit 212 and the MPs 220.1-220.N may be collectively referred to collectively as a multi-port memory unit.

**[00252]** The plurality of MPs 220.1-220.N may be gateways for data flow between the PE array 214 and the memory unit 212. Each MP 220.1-220.N may be coupled to the memory unit 212 respectively to read from and write to the memory unit 212. With the exception of MP 220.1 and MP 220.N, all MPs may be coupled to two adjacent MPs such that each MP may be configured to receive data from a first MP and/or transmit data to a second MP. The electronic coupling between MPs may provide a one-way flow of data (e.g., if one computation configuration specifies that data may flow from one MP to a next MP). For example, as shown in FIG.2, MP 220.1 may be coupled to MP 220.2 for one-way flow of data, MP 220.2 may be coupled to MP 220.3 for one-way flow of data. The last MP 220.N may be an exception and coupled to the gasket memory 216, which may provide a temporary storage for data. The first MP 220.1 may be another exception in that it may receive one-way flow of data from the gasket memory 216. In some embodiments, the MPs 220.1-220.N may form a data routing bus along a PE row direction. That is, data may be routed between MPs in a direction that is parallel to the direction that data may be routed between PEs. In embodiments with two-

dimensional PE array 214, each MP 220.1-220.N may be shared by one column of PEs. In one embodiment, the gasket memory 216 may be used as a data buffer, for example, first-in-first-out (FIFO), to collect data from the PE array and feed it back to the PE array for a new configuration.

**[00253]** In some embodiments, the PEs and MPs may be statically programmed with instructions for one configuration. For example, the instructions may be programmed to the PEs and MPs as a stage of pipeline and no instructions are changed during one configuration. The address computation instructions and memory access instructions like read or store may be mapped to the memory ports (MP) and other instructions may be mapped to PEs.

**[00254]** As illustrated in FIG. 2, each of the ICSB 224.1-224.N may be coupled to a respective SB 222.1-222.N and two adjacent ICSBs, with the exception of ICSB 224.1 and ICSB 224.N. ICSB 224.1 may be coupled to the MP 222.1, the gasket memory 216 and the ICSB 224.2. And ICSB 224.N may be coupled to the MP 222.N, the gasket memory 216 and the ICSB 224.N-1. In some embodiments, the ICSBs 220.1-220.N may form another data routing bus along a PE row direction. That is, data may be routed between ICSBs in a direction that is parallel to the direction that data may be routed between PEs. An ICSB may also be referred to as a data switching unit. In one embodiment, one or more ICSBs may be used to route data to bypass one or more PEs.

**[00255]** The SBs 222.1-222.N may be configured to provide data switching for neighboring PEs, PEs to data routing buses, and the data routing bus formed by the MPs 220.1-220.N and the data routing bus formed by the ICSBs 224.1-224.N. For example, the switch box 222.1 may be configured to provide data switching for data to be delivered to the processing

element 218.1 from the gasket memory 216, the MP 220.1 and the ICSB 224.1. Moreover, the switch box 222.1 may be configured to route data between the gasket memory 216, the MP 220.1 and the ICSB 224.1. As another example, the switch box 222.2 may be configured to provide data switching for data to be delivered to the processing element 218.2 from the processing element 218.1, the MP 220.2 and the ICSB 224.2. Moreover, the switch box 222.2 may be configured to route data between the processing element 218.2, the MP 220.2 and the ICSB 224.2. In yet another example, the switch box 222.N may be configured to provide data switching for data to be delivered to the processing element 218.N from the PE 218.N-1, the MP 220.N and the ICSB 224.N. Moreover, the switch box 222.N may be configured to route data between PE 218.N-1, MP 220.N and ICSB 224.N. A SB may also be referred to as a data switching unit.

**[00256]** An exemplary data path may be illustrated by the exemplary internal connections of the MPs 222.1 to 222.N. For example, as shown in FIG. 2, MP 222.1 may show that two inputs of PE 218.1 may be coupled to two outputs from MP 220.1, MP 222.2 may show that two inputs of PE 218.2 may be coupled to two outputs from MP 220.2 and two inputs of PE 218.2 may be coupled to two outputs from PE 218.1, MP 222.3 may show that two inputs of PE 218.3 may be coupled to two outputs from MP 220.3 and two inputs of PE 218.3 may be coupled to two outputs from PE 218.2,, and so forth, until MP 222.N may show that two inputs of PE 218.N may be coupled to two outputs from MP 220.N and two inputs of PE 218.N may be coupled to two outputs from PE 218.N-1.

**[00257]** To simplify wording, a MP 220 may refer to one of the MPs 220.1-220.N, a SB 222 may refer to one of the SBs 222.1-222.N, a PE 218 may refer to one of the PEs 218.1-218.N and an ICSB 224 may refer to one of the ICSB 224.1-224.N.

**[00258]** FIG. 3A schematically shows a memory system for a processor in accordance with an embodiment of the present disclosure. The memory system may comprise a memory unit 300 and the plurality of memory ports 220.1 through 220.N. The memory unit 300 may be an embodiment of the memory unit 212 in FIG.2, and may comprise a plurality of memory banks (e.g., memory bank 0 denoted as 302.1, memory bank 1 denoted as 302.N, memory bank N-1 denoted as 302.N, etc.) and a plurality of memory caches 304.1 through 304.N. Each of the memory bank 302 may be coupled to a respective cache 304. For example, the memory bank 302.1 may be coupled to the cache 304.1, the memory bank 302.2 may be coupled to the cache 304.2, the memory bank 302.N may be coupled to the cache 304.N, and so on. Each cache 304 may be individually coupled to all of the plurality of memory ports 220.1-220.N. For example, the cache 304.1 may be coupled to the MPs 220.1-220.N, the cache 304.2 may be coupled to the MPs 220.1-220.N, the cache 304.N may be coupled to the MPs 220.1-220.N, and so on.

**[00259]** In addition to being individually coupled to all caches 304.1-304.N of the memory unit 300, the MPs 220.1-220.N may be chained to form the row direction data routing bus, with the MP 220.1 and the MP 220.N being coupled at one end to the gasket memory 216 respectively (shown in FIG. 2). Each of MPs 220.1-220.N may further comprise a write data (WData) input 306 and a read data (RData) output 308. For example, the MP 220.1 may comprise a write data input 306.1 and a read data output 308.1, the MP 220.2 may comprise a write data input 306.2 and a read data output 308.2, the MP 220.N may comprise a write data

input 306.N and a read data output 308.N, and so on. The write data inputs 306 and read data outputs 308 may be coupled to respective outputs and inputs of the SBs 222.1-222.N. In one embodiment, each of the write data inputs 306 and each of the read data outputs 308 may be configured for vector data connections. For example, the write data input 306.1 may be a 32x32 input or two 32x16 inputs, and the read data output 308.1 may be a 32x32 output or two 32x32 outputs. As used herein, a data input or data output may also be referred to a data port.

**[00260]** The memory unit 300 and MPs 220.1-220.N may support two accessing modes: a private memory access mode and a shared memory access mode, which may also be referred to as the private memory access method and shared memory access method. In one MP, multiple data units may be read or written using a vector address. These addresses of one vector may be different from each other. In the private memory access mode, one address in a vector address may be routed to one memory bank according to the thread index. All private data for one thread may be located in the same memory bank. In shared memory access mode, each MP may access anywhere in the defined region regardless of thread index. Data shared to all threads may be spread in all memory banks.

**[00261]** The memory unit structure is illustrated in **Error! Reference source not found.A** as one example. For each column of a PE array, it may have one MP with multiple buses going through. The memory port may be configured as shared (e.g., shared memory access mode) or private (e.g., private memory access mode). Each memory port may be further coupled to a data cache network.

**[00262]** FIG. 3B schematically shows a first memory mapping for a private memory access mode in accordance with an embodiment of the present disclosure. Each memory bank

302.1 through 302.N may comprise a plurality of “words.” In the embodiment shown in FIG. 3B, each word of the memory bank may be 512-bit wide and may contain 32 data units that each may be 16 bits. In **Error! Reference source not found.**consecutive data units of a thread “i” may be referred to as  $S_i(0), S_i(1), \dots$ , and stored in a memory bank-i. For example, data units  $S_0(0), S_0(1), \dots, S_0(31)$  for thread zero (“0”) may be stored in the first word in the memory bank 302.1, and data units  $S_0(32), S_0(33), \dots, S_0(63)$  for thread 0 may be stored in the second word in the memory bank 302.1, and so on. Similarly, data units  $S_1(0), S_1(1), \dots, S_1(31)$  for thread one (“1”) may be stored in the first word in the memory bank 302.2, and data units  $S_1(32), S_1(33), \dots, S_1(63)$  for thread 1 may be stored in the second word in the memory bank 302.2, and so on. And data units  $S_{31}(0), S_{31}(1), \dots, S_{31}(31)$  for thread 31 may be stored in the first word in the memory bank 302.N, and data units  $S_{31}(32), S_{31}(33), \dots, S_{31}(63)$  for thread 31 may be stored in the second word in the memory bank N-1, and so on.

**[00263]** In one embodiment of this first memory mapping, data units for different threads may be intended to be stored in different memory banks and wrap back to the first bank for thread N. For example, for N equal to 32, data units for the 32<sup>nd</sup> thread may be stored to memory bank 0 (e.g., data units  $S_{32}(0)$  through  $S_{32}(31)$  in memory bank 0), data units for the 33<sup>rd</sup> thread may be stored to memory bank 1 (e.g., data units  $S_{33}(0)$  through  $S_{33}(31)$  in memory bank 1), data units for the 63<sup>rd</sup> thread may be stored to memory bank N-1 (e.g., data units  $S_{63}(0)$  through  $S_{63}(31)$  in memory bank 0), and so on.

**[00264]** For the same memory structure of FIG. 3B, data could be mapped in a different way. FIG. 3C schematically shows a second memory mapping for a private memory access

mode in accordance with an embodiment of the present disclosure. The memory unit 300 shown in FIG. 3C may comprise the same plurality of memory banks as in FIG. 3A, and each word in the memory banks 302.1-302.N of FIG. 3C may also be 512 bits wide and each data unit be 16 bits wide. In **Error! Reference source not found.** consecutive data units of a thread  $i$  may still be stored in a memory bank- $i$  but in different words. For example, data units  $S_0(0)$ ,  $S_0(1)$  and so on for thread 0 may be stored in different words in a column direction in memory bank 302.1; data units  $S_1(0)$ ,  $S_1(1)$  and so on for thread 1 may be stored in different words in a column direction in memory bank 302.2; data units  $S_{31}(0)$ ,  $S_{31}(1)$  and so on for thread 32<sup>nd</sup> may be stored in different words in a column direction in memory bank 302.N; and so on.

**[00265]** In one embodiment of this second memory mapping, data units for different threads may be intended to be stored in different memory banks and wrap back to the first bank for thread  $N$  and integer multiple of  $N$  (e.g.,  $2N$ ,  $3N$ , etc.). Moreover, data units of a group of different threads with same index may be mapped to the same word of a memory bank. For example, for  $N$  equal to 32, data units for the 32<sup>nd</sup> thread may be stored to memory bank 302.1 in different words (e.g., data units  $S_{32}(0)$  through  $S_{32}(99)$  in memory bank 302.1 in a second column, with data units  $S_0(m)$  and  $S_{32}(m)$  in the same word,  $m$  being the index of the data unit in a thread), data units for the 33<sup>rd</sup> thread may be stored to memory bank 302.2 in different words (e.g., data units  $S_{33}(0)$  through  $S_{33}(99)$  in memory bank 302.2 in a second column, with data units  $S_1(m)$  and  $S_{33}(m)$  in the same word,  $m$  being the index of the data unit in a thread), data units for the 63<sup>rd</sup> thread may be stored to memory bank 302.N (e.g., data units  $S_{63}(0)$  through  $S_{63}(99)$  in memory bank 0, with data units  $S_{31}(m)$  and  $S_{63}(m)$  in the same word,  $m$  being the index of the data unit in a thread), and so on. Because each word has 32 data units,

the last data unit in the first row of memory bank 302.1 may be the first data unit S992(0) of the thread 992, the last data unit in the first row of memory bank 302.2 may be the first data unit S993(0) of the thread 993, and so on until the last data unit in the first row of memory bank 302.N may be the first data unit S1023(0) of the thread 1023. It should be noted that a thread may have more than 99 data units and  $S_i(99)$  (e.g., S0(99), etc.) may not be the last data units for a thread and dotted lines may represent that more data units may exist and stored in a memory bank.

**[00266]** The data units for thread 1024 and higher number of threads may be wrapped from the first column of memory bank 0 and so on. For example, with  $m$  being the index, data units for threads 1024, 1056 and so on until 2016 (e.g., S1024( $m$ ), S1056( $m$ ) and so on until S2016( $m$ )) may be in one word of the memory bank 0; data units for threads 1025, 1057 and so on until 1057 (e.g., S1025( $m$ ), S1057( $m$ ) and so on until S2017( $m$ )) may be in one word of the memory bank 1; and data units for threads 105, 1087 and so on until 2047 (e.g., S1055( $m$ ), S1087( $m$ ) and so on until S2047( $m$ )) may be in one word of the memory bank N-1.

**[00267]** FIG. 3D schematically shows one example of memory mapping for a shared memory access in accordance with an embodiment of the present disclosure. The memory unit 300 shown in FIG. 3D may comprise the same plurality of memory banks as in FIG. 3A, each word in the memory banks 302.1-302.N of FIG. 3D may also be 512 bits wide and each data unit be 16 bits wide. In this example memory mapping for shared memory access, consecutive data units  $A(0)$ ,  $A(1)$ , ..., may be stored in different memory banks in interleaved fashion. For example, for  $N$  equal to 32,  $A(0)$ ,  $A(1)$ ,  $A(2)$ , and so on may be spread out in  $N$  memory banks, with  $A(1)$  in memory bank 0,  $A(2)$  in memory bank 1, and so on until  $A(31)$  in memory bank N-1;

and wrap back with A(N) in memory bank 0 in the same word as A(0), A(N+1) in memory bank 1 in the same word as A(1), and so on, until A(N+31) in memory bank N-1 in the same word as A(31); and wrap back and so on, until A(992) (e.g. A(31N)) in memory bank 0 in the same word as A(0), A(993) (e.g., A(31N+1)) in memory bank 1 in the same word as A(1), and so on until A(1023) (e.g., A(31N+31)) memory bank N-1 in the same word as A(31). After one word of a memory bank may be filled, more consecutive data units may be spread in another word in the memory banks. For example, A(1024), A(1056), through A(2016) may be in another word in memory bank 0; A(1025), A(1057), through A(2017) may be in another word in memory bank 1; and so on, until A(1055), A(1087), through A(2047) may be in another word in memory bank N-1.

**[00268]** Regardless of private or shared memory access modes, each of the caches 304.1-304.N of a memory unit 300 may comprise multiple cache lines that each may temporarily store a memory word from a corresponding memory bank. For example, cache 304.1 may comprise multiple cache lines that each may be configured to temporarily store one word retrieved from the memory bank 302.1 (e.g., memory bank 0), cache 304.2 may comprise multiple cache lines each configured to temporarily store one word retrieved from the memory bank 302.2 (e.g., memory bank 1), cache 304.N may comprise multiple cache lines each configured to temporarily store one word retrieved from the memory bank 302.N (e.g., memory bank N-1), and so on. A cache miss may be generated when one or more data pieces (e.g., one or more data units) requested are not in the cache. In one embodiment, one memory word of a memory bank of the memory unit 300 (e.g., in either FIG. 3B, FIG. 3C or FIG. 3D) may be fetched as one cache line into a cache when there is a cache miss. Generally, the larger the cache size,

more cache lines may be used to store memory words and the lower cache miss rate may be expected. In some embodiments, storage units in a cache may be implemented as registers.

**[00269]** Data storage in the memory unit 212 may be accessed by the MPs 220.1-220.N via the caches 304.1-304.N. The memory ports (MP) at each column may be configured with same components to carry out the memory operations, for example, calculating addresses and issuing read and/or store operations. In some embodiments, one cache 304 may be accessed by multiple MPs at the same time. Each of the MPs may be configured to provide the two accessing modes: the private memory access mode and the shared memory access mode. Due to the nature of SIMT, memory read or write instructions mapped to a MP for different threads belong to the same type, either shared or private. Moreover, a MP may be configured for private or shared memory access mode for a duration of a configuration.

**[00270]** FIG. 4A schematically shows a first memory access configuration for a memory port (MP) 400 in accordance with an embodiment of the present disclosure. The MP 400 may be an embodiment of a memory port 220. The first memory access configuration shown in FIG. 4A may be one example configuration for the private memory access mode. During operation, the MP 400 may receive a memory port (MP) configuration for a physical data path from the sequencer and the MP configuration may specify that for this physical data path the MP 400 may be configured for the private memory access mode. The MP 400 may comprise an address calculation unit 402 and a counter 404. The address calculation unit 402 may take a base address as a first input, an offset as a second input, and a third input from the counter 404. The base address may be a common address for all threads. The offset may be coupled to an address input port of the MP 400, which may be configured to accept a vector address.

The vector address may comprise individual addresses for each concurrent thread and these address may be referred to as address offsets. In one embodiment, the base address in the first memory access configuration may contain the starting address of the thread-0 data unit 0, e.g.,  $S_0(0)$ . The number of concurrent threads may be limited by the number of ALUs in a PE, and the width of the vector address and vector data bus. For example, if the number of ALUs in a PE is  $N$ , and the vector address may comprise  $N$  addresses and vector data bus may comprise  $N$  data buses, there may be  $N$  concurrent threads. The offset input may be the vector address for the  $N$ -threads. Each address offset (e.g., each address of the vector address) may be independently programmed/calculated.

**[00271]** The third input from the counter 404 may provide thread numbers (e.g., indexes) for the address calculation unit 402 and therefore, the counter 404 may be referred to as a thread counter. In one embodiment, the address vector, read data vector and write data vector may be simply split into each memory bank with a one-to-one mapping so that the data of different threads may be mapped into different memory banks. For example, the  $i$ -th address in the vector address may be for thread  $i$  (lower case letter “ $i$ ” to denote a thread number, which may start from zero for the first thread), and the counter 404 may provide a thread number vector to the address calculation unit 402 so the address calculation unit 402 may generate  $N$  addresses as  $A_0, A_1 \dots, A_{N-1}$  in this example corresponding to the vector size of ALU. Each address in the vector address may be mapped to an address  $A_i$  and a corresponding address output for a corresponding memory bank (e.g.,  $A_0$  coupled to the address port 410.1 for the memory bank 0 cache 304.1,  $A_{N-1}$  coupled to the address port 410.N for memory bank N-1 cache 304.N, etc.). The  $i$ -th data lines in the vector write data port WData 406 may be

mapped to  $WD_i$  (e.g.,  $WD_0$  coupled to the write data port 412.1 for memory bank 0 cache 304.1,  $WD_{N-1}$  coupled to the write data port 412.N for memory bank N-1 cache 304.N, etc.). The  $i$ -th data lines in the vector read data port RData 408 may be mapped to  $RD_i$  (e.g.,  $RD_0$  coupled to the read data port 414.1 for memory bank 0 cache 304.1,  $RD_{N-1}$  coupled to the read data port 414.N for memory bank N-1 cache 304.N, etc.). No bus switch may be needed for this configuration and there may be no memory contention at this level.

**[00272]** It should be noted that the number of memory bank does not need to be identical to the vector size. For example, a vector (e.g., vector ALU, vector address, vector data ports) may have a vector size =  $V$ , a PE array may have a number of columns =  $N$ , and a memory unit may have a number of memory banks =  $M$ , and  $V$ ,  $N$  and  $M$  may be all different. For convenience, the capital letter  $N$  may be used herein to denote the vector size, the number of columns of PEs, and the number of memory banks, but the number represented by  $N$  may be equal or different in different components.

**[00273]** For thread numbers larger than the number  $N$ , the address calculation unit 402 and the counter 404 may generate a memory mapping that wraps around to  $N$  memory banks. For example, thread 32 may be mapped to memory bank 0 cache 304.1 (e.g.,  $S32(0)$  to memory bank 302.1 in FIGs. 3B and 3C), thread 63 may be mapped to memory bank N-1 cache 304.N (e.g.,  $S63(0)$  to memory bank 302.N in FIGs. 3B and 3C).

**[00274]** FIG. 4B schematically shows a second memory access configuration for the MP 400 in accordance with an embodiment of the present disclosure. The second memory access configuration shown in FIG. 4B may be one example configuration for the shared memory access mode. During operation, the MP 400 may receive a memory port (MP) configuration for

a physical data path from the sequencer and the MP configuration may specify that for this physical data path the MP 400 may be configured for the shared memory access mode. The address calculation unit 402 may take the base address as a first input and the offset as a second input, same as the first memory access configuration in FIG. 4A. But the counter 404 is not used in the shared memory access mode and the input from the counter 404 may be ignored. The base address of the shared memory is common to all the threads but the offsets could be different from each thread. The address calculation unit 402 may generate N addresses as  $A_0, A_1 \dots, A_N - 1$  in this example corresponding to the vector size of ALU. In contrast to the first memory access configuration in FIG. 4A, in which each address  $A_i$  may be mapped to one memory bank, in the second memory access configuration, the N addresses from the address calculation unit 402 may be delivered a plurality of address selection units (e.g., "Select 2" units 416.1 through 416.N). Each address selection unit 416.1 through 416.N may also take a mask as an input, for example, "Bank 0", ..., and "Bank N-1", and addresses targeted to a particular memory bank may be arbitrated by a respective selection unit such that a few addresses may be selected. The upper number of addresses that may be selected may depend on design considerations, for example, 2, 3 or more. And because the vector address may have a fixed number of addresses, if more than one address is directed to one memory bank, there may be one or more memory banks that have no addresses directed to. The selected addresses may be mapped to a memory port for a memory bank cache (e.g., address port 426.1 for memory bank 0 cache 304.1, address port 426.N for memory bank N-1 cache 304.N, etc.). In one embodiment, for example, up to two addresses may be selected from N

addresses and each address port 426.1 through 426.N may be configured to deliver up to two memory addresses for a corresponding memory bank.

**[00275]** Because more than one address may be selected for one memory bank, write data selection units (e.g., "Select 2" units 418.1 through 418.N) and read data selection units (e.g., "Select" units 420.1 through 420.N) may be provided to map multiple data ports from the vector data ports WData 406 and RData 408 to one memory bank. Each of the write data selection unit 418.1 through 418.N may take an input from a corresponding data selection unit 416.1 through 416.N, and map multiple write data lines from the write data lines WD<sub>0</sub> through WD<sub>N-1</sub> to a corresponding write data port for a selected memory bank (e.g., write data port 422.1 for memory bank 0 cache 304.1, write data port 422.N for memory bank N-1 cache 304.N). Each of the read data selection unit 420.1 through 420.N may take an input from a corresponding data selection unit 416.1 through 416.N passed over by a corresponding selection unit 418.1 through 418.N, and map multiple read data lines from the read data lines RD<sub>0</sub> through RD<sub>N-1</sub> to a corresponding read data port for a selected memory bank (e.g., read data port 424.1 for memory bank 0 cache 304.1, read data port 424.N for memory bank N-1 cache 304.N). In an embodiment in which up to two addresses may be selected from N addresses, the width of the address ports 426.1 through 426.N, the write data ports 422.1 through 422.N and the read data ports 424.1 through 424.N may be doubled of that of the address ports 410.1 through 410.N, the write data ports 412.1 through 412.N and the read data ports 414.N.

**[00276]** Embodiments of a processor may comprise a large amount of ALUs and support massive amount of parallel threads. The memory access could be very busy. It may be

extremely expensive to use multiport memory to meet the requirement. The complexity may also become very high if large amount of memory banks is used. The example private memory access may reduce the complexity of memory structure and support many typical memory access patterns for parallel processing. Some typical private memory access patterns are listed below.

Memory Access Pattern for Private Memory Access	
1)	One column mode: One MP may access the memory unit for all threads (all executing same instruction) in one PE, arbitrary and independent offset may be used for each ALU thread. I.e., $\{a_0 = C_0, a_1 = C_1, \dots, a_{31} = C_{31}\}$ , where $C_0$ may be the address offset for thread-0, $C_1$ may be the address offset for thread-1. There may be not memory contention in this case due to explicit memory partitioning for different threads accessing data located in different memory banks.
2)	Linear mode: Multiple MPs may access the memory unit for concurrent threads in different PEs, address offsets for each MP may be linear. For example, for concurrent threads in a first PE, the address offsets may be $\{a_0=C, a_1=C, \dots, a_{31}=C\}$ for a first MP, for concurrent threads in a second PE, the address offsets may be $\{b_0=C+1, b_1=C+1, \dots, b_{31}=C+1\}$ for a second MP, and so on for other concurrent threads in other PEs.
3)	Reverse linear mode: Multiple MPs may access the memory unit for concurrent threads in different PEs, address offsets for each MP may be reverse linear. For example, for concurrent threads in a first PE, the address offsets may be $\{a_0=C, a_1=C, \dots, a_{31}=C\}$ for a first MP, for concurrent threads in a second PE, the

	address offsets may be $\{b_0=C-1, b_1=C-1, \dots, b_{31}=C-1\}$ for a second MP, and so on for other concurrent threads in other PEs
4)	Overlap mode: Multiple MPs may access the memory unit for concurrent threads in different PEs, address offsets for each MP may be overlapped. For example, for concurrent threads in a first PE, the address offsets may be $\{a_0=C, a_1=C+1, \dots, a_{31}=C+31\}$ for a first MP, for concurrent threads in a second PE, the address offsets may be $\{b_0=C+1, b_1=C+2, \dots, b_{31}=C+32\}$ for a second MP, and so on for other concurrent threads in other PEs
5)	Non-unity stride mode: Multiple MPs may access the memory unit for concurrent threads in different PEs, address offsets from each MP may be spaced by a stride D. For example, for concurrent threads in a first PE, the address offsets may be $\{a_0=C, a_1=C, \dots, a_{31}=C\}$ for a first MP, for concurrent threads in a second PE, the address offsets may be $\{b_0=C+D, b_1=C+D, \dots, b_{31}=C+D\}$ for a second MP, for concurrent threads in a third PE, the address offsets may be $\{c_0=C+2D, c_1=C+3D, \dots, c_{31}=C+3D\}$ for a third MP, and so on for other concurrent threads in other PEs
6)	Random mode: Multiple MPs may access the memory unit for concurrent threads in different PEs, address offsets from each MP may be random numbers. For example, for concurrent threads in a first PE, the address offsets may be may be $\{a_0=n_{0,0}, a_1=n_{0,1}, \dots, a_{31}=n_{0,31}\}$ for a first MP, for concurrent threads in a second PE, the address offsets may be $\{b_0=n_{1,0}, b_1=n_{2,0}, \dots, b_{31}=n_{31,0}\}$ for a second MP,

	and so on for other concurrent threads in other PEs, where $n_{(i,j)}$ may be different random numbers.
--	---

TABLE 1. PRIVATE MEMORY ACCESS PATTERNS

**[00277]** In some embodiments, the private memory access may allow random data access from all threads at the same time but to different memory area for each thread. This enables programmers to write software in conventional style, without complicated data vectorization and detailed knowledge of underlying processor hardware architecture. This may enable same-instruction-multiple-thread (SIMT) programming to be applicable to an embodiment of a PE array. That is, one instruction may be concurrently executed by multiple threads in one PE.

**[00278]** Due to the non-overlapping nature, the total throughput may be the sum of throughputs of all threads. Embodiments of the private memory access mode may support large throughput from simultaneous access from each thread. The first and second memory data mapping may allow minimum memory contention in typical private data access patterns. Embodiments of private memory access may also reduce the complexity of memory system. The number of memory banks may be significantly reduced. The parallel cache structure may also reduce the total cache size since each content in the cache may be unique. Moreover, embodiments of private memory access may significantly reduce access to the memory banks by allowing simultaneous cache access from multiple memory ports.

**[00279]** In one embodiment, for a PE array size with 32x32 ALUs, only 32 memory banks may be required using the private memory access configuration (e.g., as shown in FIG. 4A). This

may greatly reduce the number of memory banks from 1024 memory banks for a conventional design to support different addresses used by each thread.

**[00280]** Different memory access patterns may use different mapping methods, both mappings in FIG. 3B and FIG. 3C may be supported by an embodiment of the memory access configuration shown in FIG. 4A using an address generation instruction.

**[00281]** The first memory mapping in FIG. 3B may be good to handle case-1, case-2, case-3 and case-4 in Table 1. In case-6 in Table 1, if the range is within the cache size, it also can be handled well. The second memory mapping in FIG. 3C may be good to handle case-1, case-2, case-3, and case-5 in Table 1.

**[00282]** In some embodiments, register spilling may occur. Register spilling may refer to scenarios that when a compiler is generating machine code, there are more live variables than the number of registers the machine may have and thus some variables may be transferred or spilled to memory. Memory for register spilling may be private to each thread, these spilled variables may need to be stored in private memory. Due to the fact that all address offset for register spilling may be identical to each thread, they are similar to Non-unity stride mode in case-5 of Table 1 and the spilled variables may be stored using the second memory mapping as shown in FIG. 3C, and there may be no memory contention.

**[00283]** The example shared memory access mode may also reduce the complexity of memory structure and support many typical memory access patterns for parallel processing. Some typical shared memory access patterns are listed below.

	Memory Access Pattern for Shared Memory Access
--	--

<p>1)</p>	<p>One MP may access the memory unit for all threads in one PE, same addresses for all these concurrent threads. For example, the address offsets may be <math>\{a_0=C, a_1=C, \dots, a_{31}=C\}</math> for the concurrent threads</p>
<p>2)</p>	<p>Multiple MPs may access the memory unit for concurrent threads in different PEs, address offsets may be the same within one MP but different for different MPs. For example, for concurrent threads in a first PE, the address offsets may be <math>\{a_0=C_1, a_1=C_1, \dots, a_{31}=C_1\}</math> for a first MP, for concurrent threads in a second PE, the address offsets may be <math>\{b_0=C_2, b_1=C_2, \dots, b_{31}=C_2\}</math> for a second MP, and so on for other concurrent threads in other PEs</p>
<p>3)</p>	<p>One MP may access the memory unit for all threads in one PE, the addresses may be sequential, for example, the address offsets may be <math>\{a_0=C, a_1=C+1, \dots, a_{31}=C+31\}</math> for the concurrent threads</p>
<p>4)</p>	<p>Multiple MPs may access the memory unit for concurrent threads in different PEs, for example, for concurrent threads in a first PE, the address offsets may be <math>\{a_0=C_1, a_1=C_1+1, \dots, a_{31}=C_1+31\}</math> for a first MP, for concurrent threads in a second PE, the address offsets may be <math>\{b_0=C_2, b_1=C_2+1, \dots, b_{31}=C_2+31\}</math> for a second MP, and so on for other concurrent threads in other PEs</p>
<p>5)</p>	<p>One MP may access the memory unit for all threads in one PE, address may be sequential with discontinuity, for example, the address offsets may be <math>\{a_0=C_1, a_1=C_1+1, a_2=C_1+K_1, a_3=C_2+K_1+1, \dots\}</math> for the concurrent threads</p>

<p>6)</p>	<p>Multiple MPs may access the memory unit for different concurrent threads in different PEs, address may be sequential with discontinuity, for example, for concurrent threads in a first PE, the address offsets may be <math>\{a_0=C_1, a_1=C_1+1, a_2=C_1+K_1, a_3=C_2+K_1+1, \dots\}</math> for a first MP, for concurrent threads in a second PE, the address offsets may be <math>\{b_0=C_2, b_1=C_2+1, b_2=C_2+K_2, b_3=C_2+K_2+1, \dots\}</math> for a second MP, etc.</p>
<p>7)</p>	<p>One MP may access the memory unit for all threads in one PE, address may be linear with non-unity stride, for example, the address offsets may be <math>\{a_0=C_1, a_1=C_1+K, a_2=C_1+2K, a_3=C_2+3K, \dots\}</math> for the concurrent threads</p>
<p>8)</p>	<p>One MP may access the memory unit for all threads in one PE, addresses may be random but within a small range, for example, the address offsets may be <math>\{a_0 = n_0, a_1 = n_1, a_2 = n_2, a_3 = n_3, \dots\}</math>, where <math>n_0, n_1, n_2, \dots, \in (C, C + R)</math>, which may mean <math>n_0, n_1, n_2, \dots</math>, are in the range of C to C+R, the range C to C+R may be what can be stored in the overall memory cache (e.g., R may be the capacity of all memory caches combined, and data stored in memory addresses C to C+R may be loaded and distributed across all memory caches at the same time)</p>
<p>9)</p>	<p>Multiple MPs may access the memory unit for concurrent threads in different PEs, addresses may be random but with a small range, for example, for concurrent threads in a first PE, the address offsets may be <math>\{a_0=n_{00}, a_1=n_{01}, a_2=n_{02}, \dots\}</math> for a first MP, for concurrent threads in a second PE, the address offsets may be <math>\{b_0=n_{10}, b_1=n_{11}, a_2=n_{12}, \dots\}</math> for a second MP, etc., where <math>n_{00}, \dots, n_{3131} \in (C, C+R)</math></p>

TABLE 2. SHARED MEMORY ACCESS PATTERNS

**[00284]** In some embodiments, the shared memory access may allow random data accesses from each parallel thread at the same time. All threads may access anywhere in a common area in the memory unit. In one embodiment, the common area may be a shared memory space that includes all memory banks. In another embodiment, the common area may be a shared memory space across a plurality of memory banks. This may enable programmers to write software in conventional style, without complicated data vectorization and detailed knowledge of underlying processor hardware architecture. This may also enable SIMT programming to be applicable to an embodiment of a PE array.

**[00285]** Embodiments of shared memory access may reduce the complexity of memory system. The number of memory banks may be significantly reduced. The parallel cache structure may also reduce the total cache size since each content in the cache may be unique. Moreover, embodiments of shared memory access may significantly reduce access to the memory banks by allowing simultaneous cache access from multiple memory ports.

**[00286]** It should be noted that as shown in FIG. 2, each MP may provide two read data ports (two arrows from the memory unit 212 to a MP and two arrows to a corresponding SB, e.g., for MP 220.1, two arrows from the memory unit 212 at the top and two arrows to SB 222.1 at the bottom) and two write data ports (two arrows to the memory unit 212 and two arrows from a corresponding SB, e.g., for MP 220.1, two arrows from the memory unit 212 at the top and two arrows to SB 222.1 at the bottom), in addition to two inputs and outputs from and to a next MP or the gasket memory 216 (e.g., two horizontal arrows from left to right between the MPs, and from MP 220.N to the gasket memory 216 and from the gasket memory 216 to the MP 220.1). That is, in some embodiments, a MP 400 may have a duplicate set of

components shown in FIGs. 4A and 4B to provide two concurrent memory access for two concurrent data access operations, for example, two reads, two writes, or one read and one write. In some embodiments, the two concurrent data access operations may be vector data operations. That is, the two read operations may be two vector data reads (e.g., reading  $2N$  pieces of data, with  $N$  being the number of concurrent threads), the two write operations may be two vector data writes (e.g., writing  $2N$  pieces of data, with  $N$  being the number of concurrent threads), and the one read and one write operations may be one vector data read and one vector data write (e.g., reading  $N$  pieces of data and writing  $N$  pieces of data with  $N$  being the number of concurrent threads). In one embodiment, for example, to perform concurrent one read and one write operations, a MP may read  $N$  pieces of data for a PE in the same column and at the same time write  $N$  pieces of data generated from a PE in a preceding column.

**[00287]** In one embodiment, for a PE array size with  $32 \times 32$  ALUs, only 32 memory banks may be needed using the shared memory access configuration (e.g., as shown in FIG. 4B). This may greatly reduce the number of memory banks from 1024 memory banks for a conventional design to support different addresses used by each thread.

**[00288]** FIG. 5 schematically shows a switch box (SB) 500 in accordance with an embodiment of the present disclosure. The SB 500 may be an embodiment of a SB 222 and may comprise a plurality of data inputs and data outputs, and inter-connections that couple the data outputs to data inputs for data switching. The data inputs of SB 500 may comprise data inputs 502.1, 502.2, 512.1, 512.2, 514.1 and 514.2. The data outputs of SB 500 may comprise data outputs 504.1, 504.2, 506.1, 506.2, 508.1, 508.2, 510.1 and 510.2. The SB 500 may further

comprise a configuration buffer 518 and a corresponding configuration input 516. The configuration buffer 518 may be implemented as a First-In-First-Out buffer and referred to as C-FIFO 518. The configuration input 516 may be coupled externally to the configuration bus that's coupled to the sequencer 206 for the SB 500 to receive configurations from the sequencer 206. The configurations for the SB 500 may be referred to as SB configurations. Moreover, the SB 500 may further comprise a plurality of counters 520.1-520.8 each corresponding to a data output, for example, counter 520.1 for data output 504.1, counter 520.2 for data output 504.2, counter 520.3 for data output 506.1, counter 520.4 for data output 506.2, counter 520.5 for data output 508.1, counter 520.6 for data output 508.2, counter 520.7 for data output 510.1, and counter 520.8 for data output 510.2.

**[00289]** Each data input of the SB 500 may be coupled to some data outputs. For example, the data input 502.1 may be coupled to the data outputs 506.1, 506.2, 508.2, 510.1 and 510.2; the data input 502.2 may be coupled to the data outputs 506.1, 506.2, 508.1, 510.1 and 510.2; the data input 512.1 may be coupled to the data outputs 504.1, 504.2, 506.1, 506.2, and 508.1; the data input 512.2 may be coupled to the data outputs 504.1, 504.2, 506.1, 506.2, and 508.2; the data input 514.1 may be coupled to the data outputs 504.1, 506.1, 506.2, 508.1, and 510.2; and the data input 514.2 may be coupled to the data outputs 504.2, 506.1, 506.2, 508.2, and 510.1.

**[00290]** Externally, depending on the location of the SB 500 in the PE array 214, the data inputs 502.1 and 502.2, and data outputs 504.1 and 504.2 may be coupled to a MP 220, or another SB 222 (e.g., in a multi-row PE array). The data inputs 514.1 and 514.2 may be coupled to a PE 218 or the gasket memory 216. The data inputs 512.1 and 512.2, and data outputs

510.1 and 510.2 may be coupled to another SB 222 (e.g., in a multi-row PE array) or an ICSB 224. The data outputs 506.1, 506.2, 508.1 and 508.2 may be coupled to a PE 218. Data signals output from the data outputs 506.1, 506.2, 508.1 and 508.2 may be denoted as  $A, B, C, D$ , and data signals input from the data inputs 514.1 and 514.2 may be denoted as  $X, Y$ . These data signals  $A, B, C, D$ , and  $X, Y$  may be the input data signals to a PE 218 and output data signals from a PE 218 as described herein.

**[00291]** Each of the counters 520.1-520.8 at the data outputs may be independently responsible for counting data passed. When one or more configurations may be loaded into the C-FIFO 518, each configuration may specify a number of counts. During execution of one configuration, all counters may independently count how many times data has passed through. When all the counters reach the number of counts specified in the configuration, a next configuration may be applied. A similar approach may be applied inside an ICSB 224, a PE 218, the gasket memory 216 and a memory port 220. Because these counters may facilitate configuration and reconfiguration of each component that may have a such counter, these counters may be referred to as reconfiguration counters and a component that has such a counter may be referred to as a reconfigurable unit. An embodiment of a processor 200 may provide massive parallel data processing using the various reconfigurable units and may be referred to as a reconfigurable parallel processor (RPP).

**[00292]** FIG. 6 schematically shows a processing element (PE) 600 in accordance with an embodiment of the present disclosure. The PE 600 may be an embodiment of a PE 218. The PE 600 may comprise an arithmetic logic unit (ALU) 602, a plurality of data buffers (e.g., D-FIFO 604.1, 604.2, 604.3 and 604.4), a counter 606, a plurality of data outputs (e.g., 608.1 and

608.2), a plurality of data inputs (e.g., 610.1, 610.2, 610.3 and 610.4), a configuration input 612 and a configuration buffer (e.g., C-FIFO) 614). In one embodiment, the ALU 602 may be one ALU (e.g., one ALU that is configured to process one piece of data at a time and may be referred as a scalar ALU). In some other embodiments, the ALU 602 may be a vector of ALUs (or referred to as a vector ALU), for example, N ALUs (with N may be referred to as the vector size of ALU), and Same Instruction Multiple Data (SIMD) operation may be applied to all ALUs of the vector. Please note that a scalar ALU may be a special case of a vector ALU with a vector size being one .

**[00293]** Data signals received from the data inputs 610.1, 610.2, 610.3 and 610.4 may be denoted as  $A, B, C, D$ , and data signals output from the data outputs 608.1 and 608.2 may be denoted as  $X, Y$ . In an embodiment in which the ALU 602 may be one ALU, each data input 610.1, 610.2, 610.3 or 610.4 and each data output 608.1 or 608.2 may have a width of M bits that may match the width of the ALU. For example, for an 8-bit ALU, each input and output may be 8-bit; for a 16-bit ALU, each input and output may be 16-bit; for a 32-bit ALU, each input and output may be 32-bit; and so on. And each input data signal  $A, B, C, D$  and each output signal  $X, Y$  may be M bits. In an embodiment in which the ALU 602 may be a vector of ALUs, each data input 610.1, 610.2, 610.3 or 610.4 may be a vector of N M-bit inputs, and each data output 608.1 or 608.2 may be a vector of N M-bit outputs. And each input data signal  $A, B, C, D$  and each output data signal  $X, Y$  may be NxM bits.

**[00294]** The data buffers 604.1-604.4 may be coupled to the inputs 610.1, 610.2, 610.3 and 610.4 to temporarily store data pieces. In some embodiments, however, the data buffers may be located as the output. The D-FIFOs 604.1-604.4 may be used to decouple the timing of

PEs to allow PEs to work in dependently. In one embodiment, the buffers may be implemented as FIFOs (e.g., a D-FIFO for a data buffer, a C-FIFO for a configuration buffer).

**[00295]** The configuration buffer C-FIFO 614 may receive configurations from the configuration input 612, which may be coupled externally to the sequencer 206 via the configuration bus, and store the received configurations before any execution of a data path starts. The configurations for the PE 600 may be referred to as PE configurations. The PE 600 may be statically programmed with instructions for one configuration, e.g., the instructions may be programmed to the PE 600 as a stage of pipeline. No instructions may be changed during one configuration. Once configured, the operation of the ALU 602 (e.g., one ALU or vector of ALUs depending on a particular embodiment) may be triggered if D-FIFOs 610.1, 610.2, 610.3 and 610.4 have data and output ports 608.1 and 608.2 are not busy. One of the configuration parameter may be a number for a specified number of executions for an instruction. The counter 606 may be programmed with the specified number and used to count the number of times data has been processed by executing an instruction. When the number of executions has reached the specified number, a new configuration may be applied. Therefore, reconfiguration capability may be provided in each PE. In one embodiment, the specified number of execution for an instruction may be referred to as NUM\_EXEC and this NUM\_EXEC may be used across a data path for one configuration.

**[00296]** In one embodiment with a multi-row PE array 214, the PEs within each column may be functionally different from each other but the PEs along each row follow a repetitive pattern (e.g., functionally duplicative). For example, ALUs in a first row of PEs may implement a first set of instructions and ALUs in a second row of PEs may implement a second set of

instructions that may be different from the first set. That is, ALU 602 in different embodiments of the PE 600 may comprise different structures or different functional components. In some embodiments, one or more rows of PEs of a processor may comprise ALUs that may be relatively simple and use less space and another row of PEs of the same processor may comprise ALUs that may be relatively more complex and use more space. The relatively simple ALUs may implement a set of instructions that may be different from a set of instructions implemented by the relatively more complex ALUs. For example, one embodiment of PE 600 may have an ALU 602 (e.g., one ALU or a vector of ALUs) that implements a set of instructions that require a relatively simple structure, such as, but not limited to, ADDITION (e.g.,  $A+B$ ), SUBTRACTION (e.g.,  $A-B$ ), etc.; while another embodiment of PE 600 may have an ALU 602 that implements instructions that require a relatively more complex structure, such as, but not limited to, MULTIPLICATION (e.g.,  $A$  times  $B$  ( $A*B$ )), MAD (for multiply–accumulate (MAC) operation) (e.g.,  $A*B + C$ ).

**[00297]** FIG. 6B schematically shows another processing element (PE) 620 in accordance with another embodiment of the present disclosure. The processing element 620 may be alternative embodiment of a PE 218. As shown in FIG. 6B, the PE 620 may have a structure similar to the processing element 600 and differ by having a plurality of indicators “S” 622.1 through 622.4 that each of the indicators may be associated with one respective D-FIFO. For example, the indicator 622.1 may be associated with D-FIFO 604.1, the indicator 622.2 may be associated with D-FIFO 604.2, the indicator 622.3 may be associated with D-FIFO 604.3 and the indicator 622.4 may be associated with D-FIFO 604.4. These indicators may be used for static read configuration (also referred to as static load configuration). For example, during one

configuration, an instruction may be set to be executed a specified number of times (e.g., NUM\_EXEC = 3) by a PE according to the configuration. Executing the instruction may use data from all four D-FIFOs 604.1 through 604.4. The indicator 622.2 may be set while other indicators 622.1, 622.3 and 622.4 may not be set. The data in the D-FIFO 604.2 may be re-used while the instruction is repeated, but new data in the D-FIFOs 604.1, 604.3 and 604.4 may be used for each repetition of the instruction.

**[00298]** FIG. 7 schematically shows an inter-column switch box (ICSB) 700 in accordance with an embodiment of the present disclosure. The ICSB 700 may be an embodiment of an ICSB 224 and may comprise a plurality of data inputs and data outputs, and inter-connections that couple the data outputs to data inputs for data switching. The data inputs of ICSB 700 may comprise data inputs 704.1, 704.2, and 710.1-710.4. The data outputs of ICSB 700 may comprise data outputs 506.1, 506.2, 506.3, and 708.1-708.4. The ICSB 700 may further comprise a configuration buffer 702 and a corresponding configuration input 712. The configuration buffer 702 may be implemented as a First-In-First-Out buffer and referred to as C-FIFO 702. The configuration input 712 may be coupled externally to the configuration bus that's coupled to the sequencer 206. Moreover, the ICSB 700 may further comprise a plurality of counters 714.1-714.6 each corresponding to a data output, for example, counter 714.1 for data output 708.1, counter 714.2 for data output 708.2, counter 714.3 for data output 708.3, counter 714.4 for data output 708.4, counter 714.5 for data output 706.1, and counter 714.6 for data output 706.2.

**[00299]** Each data input of the ICSB 700 may be coupled to some selected data outputs. For example, the data input 704.1 may be coupled to the data outputs 708.1-708.4; the data

input 704.2 may be coupled to the data outputs 708.1-708.4; the data input 710.1 may be coupled to the data outputs 706.1-706.2, and 708.1; the data input 710.2.2 may be coupled to the data outputs 706.1-706.2, and 708.2; the data input 710.3 may be coupled to the data outputs 706.1-706.2, and 708.3; and the data input 710.4 may be coupled to the data outputs 706.1-706.2, and 708.4.

**[00300]** Externally, the data inputs 704.1 and 704.2, and data outputs 706.1 and 706.2 may be coupled to a SB 222. The data inputs 710.1-710.4 may be coupled to a neighboring ICSB 224 or the gasket memory 216. The data outputs 708.1-708.4 may be coupled to another neighboring ICSB 224 or the gasket memory 216.

**[00301]** Each of the counters 714.1-714.6 at the data outputs may be independently responsible for counting data passed. When one or more configurations may be loaded into the C-FIFO 702, each configuration may specify a number of counts. The configurations for the ICSB 700 may be referred to as ICSB configurations. During execution of one configuration of the PE array 214, all counters may independently count how many times data has passed through. When all the counters reach the number of counts specified in the configuration, a next configuration may be applied. This implementation may be similar to what may be applied inside a SB 222, a PE 218, the gasket memory 216 and a memory port 220.

**[00302]** FIG. 8 schematically shows a gasket memory 800 in accordance with an embodiment of the present disclosure. The gasket memory 800 may be an embodiment of the gasket memory 216 shown in FIG. 2. The gasket memory 800 may comprise a plurality of data buffers for temporary storage of data. The data buffers may be implemented as First-In-First-Out (FIFO) buffers and referred to as D-FIFOs (e.g., D-FIFO 802.1-802.F, F being an integer equal

to or larger than 8). In addition, the gasket memory 800 may comprise a plurality of data inputs (e.g., 814.1-814.2, 816.1-816.2, and 818.1-818.4), a plurality of data outputs (e.g., 808.1-808.2, 810.1-810.2, and 812.1-812.4), an input configuration buffer (e.g., Input C-FIFO 804), an output configuration buffer (e.g., Output C-FIFO 806), a configuration input 816, a plurality of input counters 820.1-820.L (L being an integer equal to or larger than 4), and a plurality of output counters 822.1-822.4.

**[00303]** External connections from the MP 220.N, PE 218.N and ICSB 224.N may be taken as inputs at the data inputs 814.1-814.2, 816.1-816.2, and 818.1-818.4, respectively. And external connections to the MP 220.1, SB 222.1 and ICSB 224.1 may generate outputs at the data outputs 808.1-808.2, 810.1-810.2, and 812.1-814.4, respectively. The configuration input 816 may be externally coupled to the sequencer 206 via the configuration bus for the gasket memory 800 to receive configurations from the sequencer 206. The configurations for the gasket memory 800 may be referred to as gasket memory configurations. Two types of configurations may be received from the sequencer 206: input configurations and output configurations. The input C-FIFO 804 may store input configurations for input ICSB ports 818.1-818.4 to be coupled to some Data FIFOs selected from L D-FIFOs 802.5-802.F as inputs to these selected D-FIFOs. The output C-FIFO 806 may store configurations for some data FIFOs selected from L D-FIFOs 802.5-802.F to be coupled to the ICSB ports 812.1-812.4.

**[00304]** The number of gasket D-FIFOs 802.5 through 802.F storing ICSB inputs may be greater or equal to the number of input or output ICSB ports. In some embodiments, as described herein, there may be a data connection that may bypass at least a portion of a physical data path. For example, an execution result generated by one PE 218 may not be

needed by another PE 218 in the same physical data path configuration but may be used in a future configuration. These data signals for the execution result may be routed via a SB 222 and an ICSB 224 to the gasket memory 216 and stored in the D-FIFOs of the gasket memory 216 for the future configuration. Therefore, in some embodiments, the gasket memory 800 may have more D-FIFOs than the number of input or output ports.

**[00305]** Each of the input counters 820.1-820.L at the data inputs and each of the output counters 822.1-822.4 at the data outputs may be independently responsible for counting data passed. When one or more input configurations and output configurations may be loaded into the input C-FIFO 804 and output C-FIFO 806, each configuration may specify a number of counts. During execution of one configuration, all counters may independently count how many times data has passed through. When all the counters reach the number of counts specified in the configuration, a next configuration may be applied.

**[00306]** FIG. 9A schematically shows an example dependency graph of an execution kernel 900 in accordance with an embodiment of the present disclosure. In one embodiment, a sequence of instructions that may be run in one loop (e.g., one routine or one sub-routine) may be referred to as an execution kernel or simply a kernel. The instructions in the group may have some dependency. For example, the kernel 900 may have instructions denoted as A through K in the dependency graph. Instruction A may be the first instruction in the dependency graph of the kernel 900. Instructions B and C may need an output from the instruction A. Instruction B's output may be needed by both instructions D and K. Instruction D also may need an output from instruction C. Instruction D's output may be needed by both instructions E and F. Instruction F's output may be needed by instructions G, H and I. In

addition to the output from instruction F, instruction G may also need an output from instruction E. Instruction J may need outputs from instructions G, H, I, respectively. And finally, instruction K may need outputs from instructions B and I. The dependency graph of kernel 900 may be mapped to data paths to be executed by a processor in accordance with an embodiment of the present disclosure.

**[00307]** FIG. 9B schematically shows the dependency graph of execution kernel 900 of FIG. 9A mapped into a virtual data path (VDP) for a processor in accordance with an embodiment of the present disclosure. In some embodiments, mapping of a dependency graph of a kernel to a processor may be constrained by multiple factors, for example, size of a PE array, connections between PEs, memory access capability etc. It should be noted that the processor 200 in FIG. 2 schematically illustrates one embodiment of a processor with one row of PEs, but in the embodiment shown in FIG. 9B, the virtual data path for the kernel 900 may be mapped to a processor with two (“2”) rows of PEs. In general, a processor in accordance with an embodiment of the present disclosure may utilize a 1-D or 2-D data path (including a processing element (PE) array and interconnections) to process massive parallel data. Each data path may be segmented into sections. In a 1-D data path, a section may include a memory port, a switch box, a PE and an ICSB in one column; and in a 2-D data path, a section may include a memory port, two or more switch boxes, two or more PEs and an ICSB in one column. The data path may be made identical in each section. This allows the dependency graph of the kernel to be mapped to a virtual data path which may include as many repetition as needed (e.g., unlimited or even theoretically infinite) of the data path sections in one dimension. For

example, instructions of a kernel may be mapped to columns of PEs and extended in the row direction repetitively (e.g., loop back from the beginning of the row if needed).

**[00308]** FIG. 9C schematically shows the virtual data path of FIG. 9B divided into physical data paths for a processor in accordance with an embodiment of the present disclosure. In one embodiment, in order to simplify the compiler work, a physical data path (PDP) may be made to have repetitive structure. For example, each column may be identical and each PDP may comprise same amount of repetitive columns. As shown in FIG. 9C, the VDP of FIG. 9B may be divided into three PDPs (e.g., PDP1, PDP2 and PDP3) for a 2x2 PE array and thus, the three PDPs may have the same structure. The 2x2 PE array may be the whole PE array of an embodiment of a RPP, or may be part of a NxN (e.g., N being 32 as an example) PE array of another embodiment of a RPP. There may be many connections between PEs in one PDP (e.g., A to B and C, B to D, C to D, etc.), between two consecutive PDPs (e.g., D to E and F, G to J, H to J, F to I, etc.) and between non-consecutive PDPs (e.g., B to K). Within one PDP, all PEs in the PE array may be applied a configuration for the current PDP, and data from one instruction may flow to another instruction directly according to the dependency. For example, a PE may be configured to execute instruction A in PDP1 and data from this PE may directly flow from this PE to the PEs configured to execute instructions B and C. Between PDPs, data from one PDP may flow into a gasket memory for temporal storage. The data may flow out from the gasket memory as input to the next physical data path in a new configuration. For example, a PE may be configured to execute instruction B in PDP1 in one configuration, and data from this PE may be stored to a gasket memory and used as input for a PE configured to execute instruction K in PDP3 in a future configuration.

**[00309]** FIG. 10 schematically shows a pipeline operation for a processor in accordance with an embodiment of the present disclosure. The example of pipeline operation shown in FIG. 10 may be the physical data paths of FIG. 9C mapped to a 2x2 PE array of four PEs: PE0, PE1, PE2 and PE3. The instructions of A, B, C, ... K of the virtual data path may be chopped into PDPs: PDP1, PDP2 and PDP3 as shown in FIG. 9C. PE0 may be configured to execute instruction A in PDP1, instruction E in PDP2 and instruction I in PDP3. PE1 may be configured to execute instruction B in PDP1, instruction F in PDP2 and instruction J in PDP3. PE2 may be configured to execute instruction C in PDP1, instruction G in PDP2 and instruction K in PDP3. PE3 may be configured to execute instruction D in PDP1, instruction H in PDP2 and no instruction in PDP3.

**[00310]** During operation, all concurrent threads in one PE may execute the same instruction and each instruction may be executed multiple times in one PE as a pipeline stage. That is, each PE may be configured to execute an instruction NUM\_EXEC times as a pipeline stage. For example, in an embodiment that each PE may comprise a ALU vector with a vector size of one, each instruction may be configured to execute 4 times by the ALU vector at each PE. The 4 times of execution may be represented by four threads processed with each thread in a different shade. For example, in PDP1, PE0 may be configured to execute instruction A four times, PE1 may be configured to execute instruction B four times, PE2 may be configured to execute instruction C four times and PE3 may be configured to execute instruction D four times. In PDP2, PE0 may be configured to execute instruction E four times, PE1 may be configured to execute instruction F four times, PE2 may be configured to execute instruction G four times and PE3 may be configured to execute instruction H four times. In PDP3, PE0 may be configured to execute instruction I four times, PE1 may be configured to execute instruction J four times, PE2

may be configured to execute instruction K four times and PE3 may have no instruction configured. In this embodiment, because there may be data dependency between different instructions, a thread executing an instruction that depends on another instruction may be executed later in time. For example, instruction B may depend on data from instruction A's execution result and therefore, the first thread executing instruction B may follow the first thread executing instruction A in a later cycle, the second thread executing instruction B may follow the second thread executing instruction A in a later cycle, the third thread executing instruction B may follow the third thread executing instruction A in a later cycle, and the fourth thread executing instruction B may follow the fourth thread executing instruction A in a later cycle. Due to static reconfiguration scheme and dependency of the instructions, there could be some time lost during DPD reconfiguration, e.g., PE2 may have one idle cycle during PDP1 to PDP2 transition. In an embodiment in which each PE has a vector ALU with the vector size N larger than 1, each PE may execute N concurrent threads at a time, and each shaded thread in FIG. 10 may represent N concurrent threads.

**[00311]** In various embodiments, the gasket memory may provide a way to reduce the efficiency loss during reconfiguration. For example, even there may be some idle slots during reconfiguration (e.g., reconfiguration of PE2 between instruction C of PDP1 and instruction G in PDP2), if a larger number of thread are used, the idle slots may be insignificant compared to the total busy cycles.

**[00312]** FIG. 11A-11I schematically show a configuration process for a processor in accordance with an embodiment of the present disclosure. The configuration process may illustrate a data path (DP) configuration and reconfiguration. The Sequencer unit (SEQ) 1102

may be an embodiment of the sequencer 206 and may be responsible to schedule the instructions / configurations for each PE (e.g. PE\_0 1104, PE\_1 1106, and PE\_2 1108), and switch box (SB) (e.g., SB\_1 1110, SB\_2 1112, ICSB\_1 1114 and ICSB\_2 1116). Each of PEs 1104, 1106 and 1108 may be an embodiment of a PE 218. Each of the SBs 1110 and 1112 may be an embodiment of a SB 222. Each of the ICSBs 1114 and 1116 may be an embodiment of an ICSB 224. Each of the PEs and SBs (including the SBs and ICSBs) of FIG. 11A may be coupled to the SEQ 1102 via a configuration bus (shown as the dotted lines from the SEQ 1102). In this example, a Virtual Data Path (VDP) of an execution kernel may be mapped into two Physical Data Paths (PDPs). It should be noted that this example kernel may be different from the data path examples shown in FIGs. 9A-9C and FIG. 10. The PEs may be denoted with an underline in front of a numeral, such as PE\_0, PE\_1 and PE\_2 in FIGs. 11A-11I to distinguish the example PE0, PE1, PE2, and PE3 in FIG. 10.

**[00313]** In the example configuration process in FIGs. 11A-11I, the first PDP (e.g., which may be referred to as PDP\_1 with an underline in front of a numeral to distinguish from the data path example of FIG. 9C) may use the data routing PE\_0 → PE\_1, PE\_1 → PE\_2, and the second PDP (e.g., which may be referred to as PDP\_2 to distinguish from the data path example of FIG. 9C) may use a different routing PE\_0→PE\_1, (PE\_0,PE\_1)→ PE\_2. That is, in the first PDP, PE\_1 may rely on PE\_0's output for its operation, PE\_2 may rely on PE\_1's output for its operation; and in the second PDP, the PE\_1 may rely on PE\_0's output for its operation, PE\_2 may rely on both PE\_1's output and PE\_0's output for its operation.

**[00314]** **Error! Reference source not found.** FIG. 11A illustrates that, initially, the entire DP is not configured yet. All PEs may be in a default status. In one embodiment, this default

status may be implemented by a default operation instruction, called STALL. Although PEs may be programmed STALL, but their respective input data FIFOs may be empty and ready to receive data. The busy signals, however, may be set to be 0. Therefore, all the SBs may be stalled and no signal may be routed. For example, all data connections between the components may be shown in dotted lines to indicate that no data may pass through (e.g., the data connections between PE\_0 1104 and SB\_1 1110, between SB\_1 1110 and ICSB\_1 1114, between SB\_1 1110 and PE\_1 1106, between PE\_1 1106 and SB\_2 1112, between SB\_2 1112 and ICSB\_2 1116, and between SB\_2 1112 and PE\_2 1108).

**[00315]** FIG. 11B illustrates that the SEQ 1102 may start to load the first instruction INS1 into PE\_0 1104, but SB\_1 1110 and ICSB\_1 1114 may still be stalled. For example, the configuration connection from SEQ 1102 to PE\_0 1104 may be open and illustrated as dashed lines, in contrast, the other configuration connections may be closed and shown in dotted lines. Because INST1 may be the first instruction in this execution kernel and not dependent on any other conditions, PE\_0 1104 may be ready to generate outputs but blocked by the output busy signal from SB\_1 1110. At this stage, no data may pass through any data connections between the components and the data connections may be shown in dotted lines.

**[00316]** As shown in FIG. 11C, at this stage, the SEQ 1102 may have programmed both SB\_1 1110 and ICSB\_1 1114, and the data route from PE\_0 → PE\_1 may be opened. For example, the configuration connection from SEQ 1102 to SB\_1 1110 and ICSB\_1 1114 may be open and illustrated as dashed lines, in contrast, the other configuration connections may be closed and shown in dotted lines. PE\_0 1104 may generate outputs and these outputs may be routed via SB\_1 1110 (e.g., a dash dotted line inside the SB\_1 1110 to show internal connection

open) to PE\_1 1106. PE\_1 1106 may receive the outputs from PE\_0 1104 and store these data signals in its data buffer (e.g., D-FIFOs) even when PE\_1 1106 has not been configured yet. The data connections between PE\_0 1104 and SB\_1 1110 and between SB\_1 1110 and PE\_1 1106 may be shown in solid lines to indicate data may pass through. A double-lined arrow 1118 may illustrate this data flow.

**[00317]** Because the output from PE\_0 1104 may only be needed by PE\_1 1106 in the first PDP, at this moment, no data may need to pass through ICSB\_1 1114. Thus, although ICSB\_1 1114's configuration may be programmed already (e.g., its internal connection shown in a dash dotted line), but there is no data coming to ICSB\_1 1114 (e.g., its connection to SB\_1 1110 in a dotted line) and ICSB\_1 1114 may stay still.

**[00318]** At the stage shown in FIG. 11D, the configuration connection from SEQ 1102 to PE\_1 1106 may be open and illustrated as dashed lines. In contrast, the other configuration connections may be closed and shown in dotted lines. SEQ 1102 may now configure a second instruction INS2 to PE\_1 1106, and PE\_1 1106 may be ready to execute and generate outputs but blocked by the output busy signal from SB\_2 1112. At the same time, the data generated from PE\_0 1104 may continuously be sent to D-FIFO of PE\_1 1106. At this stage, all data connections between the components may be the same as in FIG. 11C.

**[00319]** In FIG. 11E, at this stage, the configuration connection from SEQ 1102 to SB\_2 1112 and ICSB 1116 may be open and illustrated as dashed lines. In contrast, the other configuration connections may be closed and shown in dotted lines. When SB\_2 1112 may be configured and open, PE\_1 1106 may start executing INS2 and generate outputs. These outputs may be routed via SB\_2 1112 (e.g., a dash dotted line inside the SB\_2 1112 to show

internal connection open) to PE\_2 1108. PE\_2 1108 may receive the outputs from PE\_1 1106 and store these data signals in its data buffer (e.g., D-FIFOs) even when PE\_2 1108 has not been configured yet. The data connections between PE\_1 1106 and SB\_2 1112 and between SB\_2 1112 and PE\_2 1108 may be shown in solid lines to indicate data may pass through. A double-lined arrow 1120 may illustrate this data flow.

**[00320]** Because in the first PDP, the input to PE\_2 1108 may only come from PE\_1 1106, at this moment, no data may need to pass through ICSB\_2 1116. Thus, although ICSB\_1 1116's configuration may be programmed already (e.g., its internal connection shown in a dash dotted line), but there is no data passing through ICSB\_2 1116 (e.g., its connection to SB\_2 1112 shown in a dotted line) and ICSB\_2 1116 may stay still.

**[00321]** At the stage shown in FIG. 11F, the configuration connection from SEQ 1102 to PE\_2 1108 may be open and illustrated as dashed lines. In contrast, the other configuration connections may be closed and shown in dotted lines. And the last instruction INS3 of the first PDP may be programmed to PE\_2 1108. Although not shown, at this stage, the configuration connection to a gasket memory (e.g., an embodiment of the gasket memory 216) may also be opened, and the data connection of the gasket memory may also be programmed. Now PE\_2 1108 may execute INS3 and result may be stored into the gasket memory. In one embodiment, by now, the entire PDP\_1 configuration may be completed and each PE (e.g., PE\_0 1104, PE\_1 1106 and PE\_2 1108) may independently execute the respective instructions a specified number (e.g., NUM\_EXEC) of times. Each SB (e.g., SB\_1 1110 and SB\_2 1112) and each ICSB (e.g., ICSB\_1 1114 and ICSB\_2 1116) may also execute their respective configurations the specified number (e.g., NUM\_EXEC) times for PDP\_1.

**[00322]** In some embodiments, configurations for PDPs of a VDP (e.g., for a dependency graph of an execution kernel) may be sent to the components independently while each component may be operating according to a current configuration. For example, while the PEs (e.g., PE\_0 1104, PE\_1 1106 and PE\_2 1108), SBs (e.g., SB\_1 1110 and SB\_2 1112) and ICSBs (e.g., ICSB\_1 1114 and ICSB\_2 1116) may be operating under their respective first configuration for PDP\_1, subsequent configurations of other PDPs of the same VDP for each of these components may be received from the SEQ 1102. In one embodiment, a plurality of configurations for one component may be sent via the configuration bus from a sequencer 206 in a batch as long as sending multiple configurations for one component will not slow down or block the operation of any other components.

**[00323]** Therefore, while PDP\_1 may be carried out, all the configurations for PDP\_2 may have been received by the components. As shown in FIG. 11G, at this stage, by the time when all threads (e.g., all ALUs in an ALU vector 602) on PE\_0 1104 have completed PDP\_1, PE\_0 1104 may be reconfigured to a fourth instruction INS4. For example, PE\_0 1104 may reconfigure by itself by applying the PDP\_2 configuration already in its configuration buffer C-FIFO 614. Similarly, SB\_1 1114 may also be reconfigured once the last data from INS1 has passed it. The first output data from INS4 may be delivered to ICSB\_1 1114 due to the new PDP\_2 configuration on SB\_1 1110 (e.g., a new dash dotted line connecting to ICSB\_1 1114). The data signal to ICSB\_1 1114 may be blocked, however, by a busy signal from ICSB\_2 1116 because ICSB\_2 1116 may have not been reconfigured yet.

**[00324]** At the stage shown in FIG. 15H, the PE\_1 1106, SB\_2 1112, ICSB\_2 1116 may have reached the specified number of executions in PDP\_1, and may be reconfigured. For

example, PE\_1 1106 may be reconfigured to a fifth instruction INS5, and SB\_2 1112 and ICSB\_2 1116 may also be reconfigured so that the data from PE\_0 1104 may also reach the D-FIFO of PE\_2 1108 via ICSB\_1 1114. A new dash dotted line inside SB\_2 1112 may be shown to connect ICSB\_2 1116 to PE\_2 1108 via SB\_2 1112. The data flow from PE\_0 1104 to PE\_2 1108 may be illustrated by the double-lined arrows 1122 (e.g., from PE\_0 1104 to ICSB\_1 1114) and 1124 (e.g., from ICSB\_1 1114 to PE\_2 1108).

**[00325]** At the stage shown in FIG. 15I, at last, PE\_2 1108 may reach the specified number of execution for PDP\_1 and may be reconfigured to a last instruction INS6 for PDP\_2. The entire PDP\_2 configuration may be applied to the data path components. The components for PDP\_2 configuration may each run a specified number of times for the PDP\_2 configuration to complete the entire execution kernel.

**[00326]** FIG. 12A schematically shows a pipeline operation for executing a stream of instructions using a processor in accordance with an embodiment of the present disclosure. In a shared memory access mode, one memory address may be accessed by all threads executing same instruction. To load data from the same address for all the threads, the pipeline may keep on fetching the same data for all these threads and creates unnecessary traffic in the memory unit. The instructions in this example may be denoted with an underline in front of a numeral to distinguish the example instructions shown in FIGs. 11A-11I. In the example pipelined instruction execution shown in FIG. 12A, instruction Ins\_0 may be a data loading instruction "Load  $x[j]$ " and a memory port may be configured to execute Ins\_0 three times (e.g., NUM\_EXEC being 3 for a memory port configured to execute the load instruction) as a pipeline stage 1202. The data piece  $x[j]$  may be common to all threads and loaded from the same

address. For example, the data piece  $x[j]$  may be a  $j$ -th data piece in a vector  $x$ , and this  $j$ -th data piece may be used by all threads. In an embodiment with  $N$  ALUs in one PE, there may be  $N$  concurrent threads represented by one block (e.g., one block of threads) and the pipeline stage 1202 may be configured to execute  $Ins\_0$   $3 \times N$  times.

**[00327]** Instruction  $Ins\_1$  may be a data loading instruction “Load  $a[k][j]$ ” and a memory port may be configured to be execute  $Ins\_1$  three times as a pipeline stage 1204. The data piece to be load by  $Ins\_1$  may be different for different threads and may be loaded from different addresses for different threads. For example,  $a[k][j]$  may be a  $j$ -th data piece for a  $k$ -th thread, with  $k$  may be an integer between 0 to  $N-1$  (inclusive) for each thread in the first block of threads, between  $N$  to  $2N-1$  (inclusive) for each thread in the second block of threads, and between  $2N$  to  $3N-1$  (inclusive) for each thread in the third block of threads.

**[00328]** In one embodiment, the pipeline stages 1202 and 1204 may be performed at a same memory port if the memory port is configured to carry out two data loading instructions in parallel. For example, two parallel read data lines and two parallel write data lines between each of the MPs 220 and the memory unit 212 are shown in FIG. 2. Moreover, the switch box 500 shows how the parallel input data lines may be switched to a PE (e.g., 502.1 and 502.2 switched to 506.1, 506.2, 508.1 or 508.2). In another embodiment, the pipeline stages 1202 and 1204 may be performed at two different memory ports.

**[00329]** Instruction  $Ins\_2$  may be a multiplication instruction “ $y = a[k][j] * x[j]$ ” with the data piece  $x[j]$  being loaded by  $Ins\_0$  and  $a[k][j]$  being loaded by  $Ins\_1$  and a PE may be configured to execute  $Ins\_2$  three times (e.g.,  $NUM\_EXEC$  being 3 with totally  $3 \times N$  times for all

threads) as a pipeline stage 1206. Therefore, each PE or MP may be configured to execute NUM\_EXEC amount of instructions as a pipeline stage.

**[00330]** Instruction *Ins\_4* may be a data loading instruction “Load  $x[j+1]$ ” and a memory port may be configured to execute *Ins\_4* three times as a pipeline stage 1208. The data piece  $x[j+1]$  may be common to all threads and loaded from the same address. For example, the data piece  $x[j+1]$  may be a  $j+1$ -th data piece in the vector  $x$ , and this  $j+1$ -th data piece may be used by all threads. Instruction *Ins\_5* may be a data loading instruction “Load  $a[k][j+1]$ ” and a memory port may be configured to execute *Ins\_5* three times as a pipeline stage 1210. The data piece to be load by *Ins\_5* may be different for different threads and may be loaded from different addresses for different threads. For example,  $a[k][j+1]$  may be a  $j+1$ -th data piece for a  $k$ -th thread, with  $k$  may be an integer between 0 to  $N-1$  (inclusive) for each thread in the first block of threads, between  $N$  to  $2N-1$  (inclusive) for each thread in the second block of threads, and between  $2N$  to  $3N-1$  (inclusive) for each thread in the third block of threads. In one embodiment, the pipeline stages 1208 and 1210 may be performed at a same memory port if the memory port is configured to carry out two data loading instructions in parallel. In another embodiment, the pipeline stages 1208 and 1210 may be performed at two different memory ports.

**[00331]** Instruction *Ins\_6* may be a multiplication instruction “ $y = a[k][j+1] * x[j+1]$ ” with the data piece  $x[j+1]$  being loaded by *Ins\_4* and  $a[k][j+1]$  being loaded by *Ins\_5* and a PE may be configured to execute *Ins\_6* three times as a pipeline stage 1212.

**[00332]** In the example pipelined instruction execution of FIG. 12A, *Ins\_0* and *Ins\_4* may be repeated although many repeated memory reads may not be necessary. FIG. 12B

schematically shows a reduced pipeline operation for executing a stream of instructions using a processor in accordance with an embodiment of the present disclosure. In one embodiment, a compiler or other method (e.g., the sequencer) may identify the static read in the sequence of instructions in FIG. 12A and the repeated operations may be reduced to one time. For example, the static load instruction (e.g., Ins\_0 and Ins\_4) may be programmed with NUM\_EXEC = 1 (e.g., this may be applied to the MP that will perform data fetching). As shown in FIG. 12B, the pipelined stages 1202A and 1208A may be one blocks respectively. In one embodiment, indicators S may be implemented in the PEs and MPs next to PEs' or MPs' D-FIFO (e.g., indicators 622 in a PE 620), and for any PEs and MPs that that receive data from the static load, the indicator S next to PEs' or MPs' D-FIFO that receive data from the static load may be set. For example, the PE configured to execute Ins\_2 at pipeline stage 1206 may comprise D-FIFOs that receive data from the MP(s) configured for pipeline stages 1202A and MP 1204 respectively, and an indicator S may be set next to the D-FIFO that receives data from pipeline stage 1202A. Similarly, the PE configured to execute Ins\_6 at pipeline stage 1212 may comprise D-FIFOs that receive data from the MP(s) configured for pipeline stages 1208A and 1210 respectively, and an indicator S may be set next to the D-FIFO that receives data from the pipeline stage 1208A. With the NUM\_EXEC =1, the data load operation in pipeline stages 1202A and 1208A may be executed only once for this configuration. Given the static indication S, the PEs configured to perform pipeline stages 1206 and 1212 may execute its operation three times (e.g., their NUM\_EXEC is still equal to 3) but data from the D-FIFO with indication S may be reused for NUM\_EXEC times.

**[00333]** In some embodiments, this operation mode with reduced pipeline stage may be generalized to other instructions. In one embodiment, for an instruction that may generate the same result for different threads, the same approach can be used to reduce power consumptions. For example, a result from one PE may be used as an input for different threads in another PE in the same physical data path, or a result from a PE of one physical data path may be used as an input for different threads in a PE in another physical data path, the result may be loaded only once with the indication S set for a corresponding D-FIFO and reused.

**[00334]** Referring to FIG. 13, a flowchart of a process 1300 for executing an execution kernel is depicted in accordance with an embodiment of the present disclosure. At block 1302, an execution kernel may be mapped into a virtual data path at a processor. For example, as shown in FIG. 9B, an execution kernel may be mapped into a virtual data path by an example processor 200. The execution kernel may include a sequence of instructions to be executed by the processor. In one embodiment, the processor may comprise various reconfigurable units that include a gasket memory. Moreover, in one embodiment, the process may comprise various reconfigurable units that form repetitive columns. At block 1304, the virtual data path may be chopped into one or more physical data paths. For example, as shown in FIG. 9C, a virtual data path may be chopped into three physical data paths. In one embodiment, the various reconfigurable units may form the one or more physical data paths to execute the sequence of instructions. Moreover, in one embodiment, each of the one or more physical data paths may be fitted into the repetitive columns respectively. For example, a first physical data path and a second physical data path may be fitted into the repetitive columns respectively. At block 1306, configurations may be delivered to various reconfigurable units of

the processor. The various reconfigurable units may form the one or more physical data paths according to the configurations to execute the sequence of instructions. At block 1308, the processor may be executed to complete the one or more physical data paths by operating the various reconfigurable units according to the configurations. In one embodiment, data from one physical data path may be routed to the gasket memory to be used in a future physical data path as input.

**[00335]** Referring to FIG. 14, a flowchart of a process 1400 for reconfiguring a processor is depicted in accordance with an embodiment of the present disclosure. At block 1402, a plurality of configurations may be delivered to a plurality of reconfigurable units of a processor. The plurality of configurations may be for the plurality of reconfigurable units to form a plurality of physical data paths to execute a sequence of instructions. In one embodiment, each of the plurality of configurations may include a specified number (e.g., the NUM\_EXEC number as described herein). At block 1404, a respective operation may be repeated at each of the plurality of reconfigurable units for a specified number of times. For example, the plurality of reconfigurable units may include a first reconfigurable processing element (PE), and the first reconfigurable PE may execute a first instruction of the sequence of instructions the specified number of times in a first physical data path of the plurality of physical data paths. At block 1406, each reconfigurable unit may be reconfigured to a new configuration. In one embodiment, after each reconfigurable unit having repeated its respective operation the specified number of times, each reconfigurable unit may be reconfigured. For example, the first reconfigurable PE may be reconfigured to execute a second instruction of the sequence of

instructions the specified number of times according to a second configuration in a second physical data path of the plurality of physical data paths.

**[00336]** Referring to FIG. 15, a flowchart of a process 1500 for reconfiguring a processor is depicted in accordance with an embodiment of the present disclosure. At block 1502 a first configuration and a second configuration may be received at a reconfigurable unit of a processor. For example, PE\_0 1104 in FIGs. 11A-11I may receive a first configuration that may comprise the first instruction INS1 and a second configuration that may comprise the fourth instruction INS4. PE\_0 1104 may store these configurations in its configuration buffer (e.g., C-FIFO 614). At block 1504, a first operation may be executed a first number of times according to the first configuration. For example, PE\_0 1104 may execute the instruction INS1 a number of times according to the NUM\_EXEC in the first configuration. The first configuration may be part of a first physical data path (e.g., PDP\_1) for executing a first part of a sequence of instructions. At block 1506, the reconfigurable unit may be reconfigured to execute a second operation a second number of times according to the second configuration. For example, PE\_0 1104 may be reconfigured by applying the second configuration to execute the instruction INS4 a number of times according to the NUM\_EXEC in the second configuration. The second configuration may be part of a second physical data path (e.g., PDP\_2) to execute a second part of the sequence of instructions.

**[00337]** Referring to FIG. 16, a flowchart of process 1600 for reconfiguring a processor is depicted in accordance with an embodiment of the present disclosure. At block 1602, a first instruction may be executed at a reconfigurable processing element a number of times. The execution may be performed according to a first configuration as part of a first physical data

path. For example, PE\_2 may be configured to execute INS3 NUM\_EXEC times according to a configuration as part of PDP\_1. At block 1604, an execution result from the reconfigurable unit may be delivered to a gasket memory. For example, the execution result from PE\_2 for execution of INS3 may be delivered to a gasket memory. In one embodiment, if the first instruction is to be executed a number of times at the reconfigurable unit, the execution result may be delivered to the gasket memory after each execution of the first instruction. For example, the gasket memory may have store the successive execution results in one of its data buffers (e.g., one of D-FIFOs 802.3 through 802.F). At block 1606, the execution result stored in the gasket memory may be fed from the gasket memory to a second physical data path. For example, the execution result from PE\_2 executing the instruction INS3 may be used as input for an instruction to be executed by a PE in the second physical data path. The execution result stored in the gasket memory may then be delivered from the gasket memory to that PE for executing the second physical data path.

**[00338]** Referring to FIG. 17, a flowchart of a process 1700 for accessing memory is depicted in accordance with an embodiment of the present disclosure. At block 1702, a plurality of memory addresses for a memory unit may be generated by an address calculation unit in a memory port for parallel threads. For example, the memory port may provide memory access to a memory unit for a processing element (PE) that may have a plurality of arithmetic logic units (ALUs) configured to execute a same instruction in parallel threads. In one embodiment, the memory port may be one of a plurality of memory ports that provide access to a memory unit for a plurality of processing elements (PEs). Each of the plurality of PEs may have a plurality of arithmetic logic units (ALUs) configured to execute a same instruction in

parallel threads. At block 1704, a plurality of memory banks in the memory unit may be accessed with each thread accessing a different memory bank. In one embodiment, in a private memory access mode, each thread may access a different memory bank for its data.

**[00339]** Referring to FIG. 18, a flowchart of a process 1800 for accessing memory is depicted in accordance with an embodiment of the present disclosure. At block 1802, a plurality of memory addresses for a memory unit may be generated by an address calculation unit in a memory port for parallel threads. The memory port may provide memory access for a processing element that has a plurality of arithmetic logic units (ALUs) configured to execute a same instruction in parallel threads. In one embodiment, the memory port may be one of a plurality of memory ports that provide access to a memory unit for a plurality of processing elements (PEs) each having a plurality of arithmetic logic units (ALUs) that are configured to execute a same instruction in parallel threads. At block 1804, a plurality of memory banks in the memory unit may be accessed with all threads accessing a common area in the memory unit. In one embodiment, in a shared memory access mode, each of the parallel executing by a PE may have access anywhere in the common area. That is, the common area may be a shared memory space for all threads.

**[00340]** Referring to FIG. 19, a flowchart of a process 1900 for reusing a piece of data is depicted in accordance with an embodiment of the present disclosure. At block 1902, it may be determined that a piece of data is to be shared and reused by all threads at a processing element (PE) of a processor during one configuration applied at the PE. For example, in one embodiment, a shared memory access mode may be applied during one configuration for a PE. A piece of data (e.g., a constant) may be common to all threads, and may be shared and reused

by all threads. At block 1904, the piece of data may be loaded once into a data buffer of the PE. In an embodiment, because the piece of data may be shared and reused. The data load operation may need to be performed only once. At block 1906, an indicator associated with the data buffer may be set to indicate that the piece of data is to be reused. In one embodiment, the data buffer inside the PE may have indicator bit "S" and this indicator "S" may be set if a piece of data in the data buffer should be reused. At block 1908, a same instruction may be executed a number of times by reusing the piece of data as an input. In one embodiment, the PE may repeatedly execute the same instruction as a pipeline stage according to a number (e.g., NUM\_EXEC) specified by the configuration.

**[00341]** The present disclosure provides apparatus, systems and methods for reconfigurable parallel processing. For example, an embodiment of a RPP may utilize a 1-D or 2-D data path which consists of a processing element (PE) array and interconnections to process massive parallel data. The data path may be made identical in each section (e.g., one column of PE(s), MP and data routing units), which may allow the dependency graph of a kernel to be mapped to a virtual data path that may be an infinite repetition of the data path sections in one dimension.

**[00342]** An embodiment of a RPP may also utilize a gasket memory to temporally store data output of data paths as where the virtual data path is segmented into physical data paths. The gasket memory may function like a data buffer (e.g., FIFO) to feed data back into physical data path of the next configuration.

**[00343]** An embodiment of a RPP may also have a one-dimensional memory unit with memory ports (MPs) connected to each column of data path. All data accessed throughout the

virtual data path may be stored in the memory unit. Each time for a new configuration, a MP may be reconfigured to access the memory unit differently while the data could stay the same. An embodiment of a RPP may separate types of memory access to private memory access and shared memory access. Private memory access may be dedicated to a particular thread with no overlapping access allowed between different threads. Shared memory access may allow all threads to access common area. Instead of defining different memories for shared and private types. An embodiment of a RPP may store data into the same memory space but provides different access method. This eliminates unnecessary data movement from private memory to shared memory and vice visa.

**[00344]** Embodiments of a RPP may be optimized to allow massive parallelism for multithread processing. In one example, with one row of 32 PEs and each PE having 32 arithmetic and logic units (ALUs), 1024 ALUs may be included in one RPP core. In some embodiments, a multi-core processor may comprise multiple RPPs.

**[00345]** Embodiments of a RPP may be reconfigured according to a reconfiguration mechanism. The various components of a RPP that include one or more reconfiguration counters may be referred to as reconfigurable units. For example, each of the PEs (e.g., PE 218), the switching units (e.g., SB 222 and ICSB 224) and memory units (e.g., MP 220, gasket memory 216), may comprise one or more reconfiguration counters, such as the counter 606 in a PE, the counters 520 in a SB, the counters 714 in an ICSB, the counters 820 and 822 in a gasket memory, and similar counters in a MP (not shown in FIG. 4A or 4B). Data processing may be pipelined when there may be no dependency between threads. Identical instruction may be executed multiple times until all threads (e.g., four blocks of threads for one physical

data path) are processed. When counters in a reconfigurable unit reach a programmed number and the reconfigurable unit may replace its configuration to a new context. This reconfiguration may be done in the same way in each PE, switch boxes and memory access units. Self-reconfiguration may be achieved with minimum idle time for switching.

**[00346]** The exemplary reconfiguration mechanism may reduce the power spent on configuration because the configuration is only switched once after all threads have been processed. This may also reduce idle time between configurations by switching each PE independently at its earliest time. By doing that, the memory required to store intermediated data may also be reduced.

**[00347]** In some embodiments, all threads may load data use the same address in a shared memory access mode. Due to the pipelined nature of operation, only the first data load instruction of all threads may need to be performed. The data loaded may be shared with all threads to reduce the memory access traffic and power consumption.

**[00348]** The techniques described herein may be implemented in one or more application specific integrated circuits (ASICs) in digital logic gates, or by a processor that execute instructions stored in a tangible processor readable memory storage media.

**[00349]** In one embodiment, any of the disclosed methods and operations may be implemented in software comprising computer-executable instructions stored on one or more computer-readable storage media. The one or more computer-readable storage media may include non-transitory computer-readable media (such as removable or non-removable magnetic disks, magnetic tapes or cassettes, solid state drives (SSDs), hybrid hard drives, CD-ROMs, CD-RWs, DVDs, or any other tangible storage medium), volatile memory components

(such as DRAM or SRAM), or nonvolatile memory components (such as hard drives)). The computer-executable instructions may be executed on a processor (e.g., a microcontroller, a microprocessor, a digital signal processor, etc.). Moreover, an embodiment of the present disclosure may be used as a general-purpose processor, a graphics processor, a microcontroller, a microprocessor, or a digital signal processor.

**[00350]** While various aspects and embodiments have been disclosed herein, other aspects and embodiments will be apparent to those skilled in the art. The various aspects and embodiments disclosed herein are for purposes of illustration and are not intended to be limiting, with the true scope and spirit being indicated by the following claims.

What is claimed is:

1. A processor, comprising:
  - a plurality of processing elements (PEs) each comprising a configuration buffer;
  - a sequencer coupled to the configuration buffer of each of the plurality of PEs and configured to distribute one or more PE configurations to the plurality of PEs; and
  - a gasket memory coupled to the plurality of PEs and being configured to store at least one PE execution result to be used by at least one of the plurality of PEs during a next PE configuration.
2. The processor of claim 1, further comprising a plurality of switch boxes coupled to the sequencer to receive switch box configurations from the sequencer, each of the plurality of switch boxes being associated with a respective PE of the plurality of PEs and configured to provide input data switching for the respective PE according to the switch box configurations.
3. The processor of claim 2, wherein the plurality of switch boxes and their associated PEs are arranged in a plurality of columns, a first switch box in a first column of the plurality of columns is coupled between the gasket memory and a first PE in the first column of the plurality of columns, and a second PE in a last column of the plurality of columns is coupled to the gasket memory.
4. The processor of claim 3, further comprising:
  - a memory unit for providing data storage for the plurality of PEs; and
  - a plurality of memory ports each arranged in a separate column of the plurality of columns for the plurality of PEs to access the memory unit.

5. The processor of claim 4, further comprising a plurality of inter-column switch boxes (ICSBs) coupled to the sequencer to receive ICSB configurations from the sequencer, the plurality of ICSBs being configured to provide data switching between neighboring columns of the plurality of columns according to the ICSB configurations.
6. The processor of claim 5, wherein the plurality of memory ports (MPs) are coupled to the sequencer to receive MP configurations from the sequencer and configured to operate in a private access mode or a shared access mode during one MP configuration.
7. The processor of claim 6, wherein a piece of data stored in the memory unit is accessed through the private access mode and the shared access mode in different part of a program without the piece of data being moved in the memory unit.
8. The processor of claim 6, wherein each of the plurality of columns comprises one PE, the plurality of PEs are identical and form one row of repetitive identical PEs.
9. The processor of claim 6, wherein each of the plurality of columns comprises two or more PEs and the plurality of PEs form two or more rows.
10. The processor of claim 9, wherein a first row of PEs are configured to implement a first set of instructions and a second row of PEs are configured to implement a second set of instructions, at least one instruction of the second set of instructions is not in the first set of instructions, wherein the of plurality of columns are identical and form repetitive columns.
11. The processor of claim 6, wherein each of the plurality of PEs comprises a plurality of arithmetic logic units (ALUs) that are configured to execute same instruction in parallel threads.
12. The processor of claim 11, wherein each of the plurality of PEs comprises a plurality of data buffers for the plurality of ALUs and is configured to operate independently.

**13.** The processor of claim 7, wherein each of the plurality of memory ports is configured to access the memory unit using a vector address, wherein in the private access mode, one address in the vector address is routed to one memory bank of the memory unit according to a thread index and all private data for one thread are located in a same memory bank.

**14.** The processor of claim 7, wherein each of the plurality of memory ports is configured to access the memory unit using a vector address, wherein in the shared access mode, one address in the vector address is routed in a defined region across memory banks regardless of the thread index and data shared to all threads are spread in all memory banks.

**15.** A method, comprising:

mapping an execution kernel into a virtual data path at a processor, wherein the execution kernel includes a sequence of instructions to be executed by the processor, and the processor comprises various reconfigurable units that include a gasket memory;

chopping the virtual data path into one or more physical data paths;

delivering configurations to various reconfigurable units of the processor for the various reconfigurable units to form the one or more physical data paths to execute the sequence of instructions; and

executing the processor to complete the one or more physical data paths by operating the various reconfigurable units according to the configurations, including routing data from one physical data path to the gasket memory to be used in a future physical data path as input.

**16.** The method of claim 15, wherein the various reconfigurable units further comprise a plurality of processing elements, a plurality of switch boxes each associated with a separate processing element, a plurality of memory ports that provide access to a memory unit for the

plurality of processing elements, and a plurality of inter-column switch boxes, where each of the various reconfigurable units are reconfigured by applying a next configuration independently from other reconfigurable units.

**17.** The method of claim 16, wherein each of the plurality of PEs comprises a plurality of arithmetic logic units (ALUs) that are configured to execute same instruction in parallel threads.

**18.** The method of claim 17, wherein each of the plurality of memory ports is configured to operate in a private access mode or a shared access mode during one configuration.

**19.** The method of claim 18, further comprising accessing a piece of data stored in the memory unit through the private access mode and the shared access mode in different physical data paths without the piece of data being moved in the memory unit.

**20.** The method of claim 18, wherein each of the memory ports is configured to access the memory unit using a vector address, wherein in the private access mode, one address in the vector address is routed to one memory bank of the memory unit according to a thread index and all private data for one thread are located in a same memory bank, and wherein in the shared access mode, one address in the vector address is routed in a defined region across memory banks regardless of the thread index and data shared to all threads are spread in all memory banks.

**21.** The method of claim 17, wherein each of the plurality of PEs comprises a plurality of data buffers for the plurality of ALUs and is configured to operate independently during one physical data path.

**22.** The method of claim 16, wherein the plurality of PEs form a PE array and the execution kernel is mapped into one or more physical data paths on the processor based on a size of the PE array, connections between the plurality of PEs, and memory access capability.

**23.** The method of claim 15, wherein the various reconfigurable units form multiple repetitive columns, and each of the one or more physical data paths is fitted into the multiple repetitive columns, and data flows between the repetitive columns are in one direction.

**24.** A system, comprising:

a processor, comprising:

a sequencer configured to map an execution kernel to be executed by the processor into a virtual data path and chop the virtual data path into one or more physical data paths;

a plurality of processing elements (PEs) coupled to the sequencer, each of the plurality of PEs comprising a configuration buffer configured to receive PE configurations for the one or more physical data paths from the sequencer; and

a gasket memory coupled to the plurality of PEs and being configured to store data from one of the one or more physical data paths to be used by another physical data path of the one or more physical data paths as input.

**25.** The system of claim 24, wherein the processor further comprises a plurality of switch boxes (SBs) coupled to the sequencer to receive SB configurations for the one or more physical data paths from the sequencer, each of the plurality of switch boxes being associated with a respective PE of the plurality of PEs and configured to provide input data switching for the respective PE according to the SB configurations.

**26.** The system of claim 25, wherein the plurality of switch boxes and their associated PEs are arranged in a plurality of columns, a first switch box in a first column of the plurality of columns is coupled between the gasket memory and a first PE in the first column of the plurality of columns, and a second PE in a last column of the plurality of columns is coupled to the gasket memory.

**27.** The system of claim 26, wherein the processor further comprises:  
a memory unit for providing data storage for the plurality of PEs; and  
a plurality of memory ports each arranged in a separate column of the plurality of columns for the plurality of PEs to access the memory unit.

**28.** The system of claim 27, wherein the processor further comprises a plurality of inter-column switch boxes (ICSBs) coupled to the sequencer to receive ICSB configurations for the one or more physical data paths from the sequencer, the plurality of ICSBs being configured to provide data switching between neighboring columns of the plurality of columns according to the ICSB configurations.

**29.** The system of claim 28, wherein the plurality of memory ports (MPs) are coupled to the sequencer to receive MP configurations for the one or more physical data paths from the sequencer and configured to operate in a private access mode or a shared access mode during one MP configuration.

**30.** The system of claim 29, wherein a piece of data stored in the memory unit is accessed through the private access mode and the shared access mode in different physical data paths of the one or more physical data paths without the piece of data being moved in the memory unit.

- 31.** The system of claim 29, wherein each of the plurality of columns comprises one PE and the plurality of PEs are identical and form one row of repetitive identical PEs.
- 32.** The system of claim 29, wherein each of the plurality of columns comprises two or more PEs and the plurality of PEs form two or more rows.
- 33.** The system of claim 30, wherein a first row of PEs are configured to implement a first set of instructions and a second row of PEs are configured to implement a second set of instructions, at least one instruction of the second set of instructions is not in the first set of instructions, wherein the of plurality of columns are identical and form repetitive columns.
- 34.** The system of claim 29, wherein each of the plurality of PEs comprises a plurality of arithmetic logic units (ALUs) that are configured to execute same instruction in parallel threads.
- 35.** The system of claim 34, wherein each of the plurality of PEs comprises a plurality of data buffers for the plurality of ALUs and is configured to operate independently.
- 36.** The system of claim 29, wherein each of the plurality of memory ports is configured to access the memory unit using a vector address, wherein in the private access mode, one address in the vector address is routed to one memory bank of the memory unit according to a thread index and all private data for one thread are located in a same memory bank.
- 37.** The system of claim 29, wherein each of the plurality of memory ports is configured to access the memory unit using a vector address, wherein in the shared access mode, one address in the vector address is routed in a defined region across memory banks regardless of the thread index and data shared to all threads are spread in all memory banks.

**38.** The system of claim 24, wherein the plurality of PEs form a PE array and the execution kernel is mapped into one or more physical data paths on the processor based on a size of the PE array, connections between the plurality of PEs, and memory access capability.

**39.** A processor, comprising

a plurality of processing elements (PEs);

a plurality of switch boxes arranged in a plurality of columns, each of the plurality of switch boxes being associated with a respective PE and configured to provide input data switching for the respective PE;

a plurality of memory ports arranged in the plurality of columns and being coupled to a memory unit and a top switch box in each column of the plurality of columns, each of the plurality of memory port being configured to provide data access to the memory unit for one or more switch boxes in a respective column;

a plurality of inter-column switch boxes (ICSBs) each coupled to a bottom switch box in each column of the plurality of columns; and

a gasket memory with its input coupled to a memory port, a PE, one or more switch boxes and an ICSB in a last column of the plurality of columns, and its output coupled to a memory port, one or more switch boxes and an ICSB in a first column of the plurality of columns.

**40.** The processor of claim 39, further comprising a sequencer coupled to the plurality of PEs, the plurality of switch boxes, the plurality of ICSBs, the plurality of memory ports and the gasket memory to deliver configurations to these components.

- 41.** The processor of claim 40, further comprising a configuration memory coupled to the sequencer to store compiled configurations for the sequencer to decode and deliver.
- 42.** The processor of claim 39, further comprising a memory unit for providing data storage for the processor.
- 43.** A processor, comprising:  
a plurality of processing elements (PEs) each comprising a configuration buffer and a plurality of arithmetic logic units (ALUs), and each configured to operate independently according to respective PE configurations stored in the configuration buffer; and  
a gasket memory coupled to the plurality of PEs and being configured to store at least one PE execution result to be used by at least one of the plurality of PEs during a next PE configuration.
- 44.** The processor of claim 43, further comprising a plurality of switch boxes each comprising a configuration buffer configured to store switch box configurations, each of the plurality of switch boxes being associated with a respective PE of the plurality of PEs and configured to provide input data switching for the respective PE according to the switch box configurations.
- 45.** The processor of claim 44, wherein the plurality of switch boxes and their associated PEs are arranged in a plurality of columns, a first switch box in a first column of the plurality of columns is coupled between the gasket memory and a first PE in the first column of the plurality of columns, and a second PE in a last column of the plurality of columns is coupled to the gasket memory.
- 46.** The processor of claim 45, further comprising:

a memory unit for providing data storage for the plurality of PEs; and  
a plurality of memory ports each arranged in a separate column of the plurality of columns for the plurality of PEs to access the memory unit.

**47.** The processor of claim 46, further comprising a plurality of inter-column switch boxes (ICSBs) each comprising a configuration buffer configured to store ICSB configurations, the plurality of ICSBs being configured to provide data switching between neighboring columns of the plurality of columns according to the ICSB configurations.

**48.** The processor of claim 46, wherein each of the plurality of memory ports (MPs) comprises a configuration buffer to store MP configurations and configured to operate in a private access mode or a shared access mode during one MP configuration.

**49.** The processor of claim 48, wherein a piece of data stored in the memory unit is accessed through the private access mode and the shared access mode in different part of a program without the piece of data being moved in the memory unit.

**50.** The processor of claim 46, wherein each of the plurality of columns comprises one PE, the plurality of PEs are identical and form one row of repetitive identical PEs.

**51.** The processor of claim 46, wherein each of the plurality of columns comprises two or more PEs and the plurality of PEs form two or more rows.

**52.** The processor of claim 51, wherein a first row of PEs are configured to implement a first set of instructions and a second row of PEs are configured to implement a second set of instructions, at least one instruction of the second set of instructions is not in the first set of instructions, wherein the of plurality of columns are identical and form repetitive columns.

**53.** The processor of claim 48, wherein each of the plurality of memory ports is configured to access the memory unit using a vector address, wherein in the private access mode, one address in the vector address is routed to one memory bank of the memory unit according to a thread index and all private data for one thread are located in a same memory bank.

**54.** The processor of claim 48, wherein each of the plurality of memory ports is configured to access the memory unit using a vector address, wherein in the shared access mode, one address in the vector address is routed in a defined region across memory banks regardless of the thread index and data shared to all threads are spread in all memory banks.

**55.** The processor of claim 43, wherein each of the plurality of PEs comprises a plurality of data buffers for the plurality of ALUs and is configured to operate independently.

**56.** A processor, comprising:

a plurality of processing elements (PEs) arranged in a plurality of columns;

a plurality of switch boxes (SBs) each associated with a separate PE of the plurality of PEs to providing data switching; and

a gasket memory coupled to the plurality of PEs and being configured to store at least one PE execution result to be delivered to at least one of the plurality of PEs via a switch box for the PE execute result to be used as input data during a next PE configuration.

**57.** A method, comprising:

mapping an execution kernel into a virtual data path, wherein the execution kernel includes a sequence of instructions to be executed by a processor, and the processor comprises various reconfigurable units that form repetitive columns;

chopping the virtual data path into one or more physical data paths to fit each physical data path into the repetitive columns respectively; and

delivering configurations to various reconfigurable units of the processor for the various reconfigurable units to form the one or more physical data paths to execute the sequence of instructions.

**58.** A method, comprising:

mapping an execution kernel into a virtual data path for a processor to execute, the processor comprising various reconfigurable units that form repetitive columns;

chopping the virtual data path into a plurality of physical data paths including a first physical data path that fits into the repetitive columns and a second physical data path that fits into the repetitive columns; and

delivering configurations to various reconfigurable units for the repetitive columns to form the first physical data path to execute a first part of the execution kernel and form the second physical data path to execute a second part of the execution kernel.

**59.** A processor, comprising:

a plurality of reconfigurable units including a plurality of processing elements (PEs) and a plurality of memory ports (MPs) for the plurality of PEs to access a memory unit, each of the plurality of reconfigurable units comprising a configuration buffer and a reconfiguration counter; and

a sequencer coupled to the configuration buffer of each of the plurality of reconfigurable units and configured to distribute a plurality of configurations to the plurality of

reconfigurable units for the plurality of PEs and the plurality of memory ports to execute a sequence of instructions.

**60.** The processor of claim 59, wherein each of the plurality of configurations includes a specified number, and the reconfigurable counter of each of the plurality of PEs and the plurality of memory ports is configured to count for a respective PE or MP to repeat an instruction of the sequence of instructions the specified number of times.

**61.** The processor of claim 60, wherein the plurality of reconfigurable units further include a plurality of data switching units, each of the plurality of data switching units is configured to a data switching setting according to a current data switching configuration the specified number of times.

**62.** The processor of claim 60, wherein the plurality of reconfigurable units further include a gasket memory, the gasket memory comprises a plurality of data buffers, an input configuration buffer, an output configuration buffer, a plurality of input reconfigurable counters and a plurality of output reconfigurable counters, and wherein the gasket memory is configured to perform reconfiguration for input and output independently.

**63.** The processor of claim 62, wherein the plurality of configurations include a first set of configurations for the plurality of reconfigurable units to form a first physical data path and a second set of configurations for the plurality of reconfigurable units to form a second physical data path, the gasket memory is configured to store data from the first physical data path to be used as input to the second physical data path.

**64.** The processor of claim 59, wherein each of the plurality of reconfigurable units is configured to switch to a next configuration independently after its reconfigurable counter reaches the specified number.

**65.** The processor of claim 59, wherein each of the plurality of memory ports are configured to operate in a private memory access mode or a shared memory access mode during one configuration.

**66.** The processor of claim 65, wherein a piece of data stored in the memory unit is accessed through the private memory access mode and the shared memory access mode in configurations for different physical data paths without the piece of data being moved in the memory unit.

**67.** The processor of claim 65, wherein each of the plurality of memory ports is configured to access the memory unit using a vector address, wherein in the private memory access mode, one address in the vector address is routed to one memory bank of the memory unit according to a thread index and all private data for one thread are located in a same memory bank.

**68.** The processor of claim 65, wherein each of the plurality of memory ports is configured to access the memory unit using a vector address, wherein in the shared memory access mode, one address in the vector address is routed in a defined region across memory banks regardless of thread index and data shared to all threads are spread in all memory banks.

**69.** The processor of claim 59, wherein each of the plurality of PEs comprises a plurality of arithmetic logic units (ALUs) that are configured to execute a same instruction in parallel threads.

- 70.** The processor of claim 69, wherein each of the plurality of PEs comprises a plurality of data buffers and is configured to operate independently.
- 71.** The processor of claim 59, wherein the plurality of PEs form two or more rows.
- 72.** The processor of claim 71, wherein a first row of PEs are configured to implement a first set of instructions and a second row of PEs are configured to implement a second set of instructions, at least one instruction of the second set of instructions is not in the first set of instructions
- 73.** The processor of claim 59, wherein the plurality of PEs and the plurality of memory ports (MPs) are arranged in repetitive columns.
- 74.** The processor of claim 59, wherein each of the sequence of instructions is executed by one of the plurality of PEs or one of the plurality of memory ports as a stage of pipeline according to a respective configuration.
- 75.** A method, comprising:
- delivering a plurality of configurations to a plurality of reconfigurable units of a processor for the plurality of reconfigurable units to form a plurality of physical data paths to execute a sequence of instructions, each of the plurality of configurations including a specified number;
  - repeating a respective operation at each of the plurality of reconfigurable units for the specified number of times, including executing a first instruction of the sequence of instructions at a first reconfigurable processing element (PE) the specified number of times according to a first configuration in a first physical data path; and

reconfiguring each of the plurality of reconfiguration units to a new configuration after repeating the respective operation the specified number of times, including executing a second instruction of the sequence of instructions at the first reconfigurable PE the specified number of times according to a second configuration in a second physical data path.

**76.** The method of claim 75, wherein the plurality of reconfiguration units includes a plurality of PEs and a plurality of memory ports, and at least one instruction of the sequence of instructions is a memory access instruction and executed by a memory port the specified number of times before the memory port is reconfigured by applying a next memory port configuration.

**77.** The method of claim 76, wherein the plurality of reconfigurable units further include a plurality of data switching units, and each of the plurality of data switching units is configured to repeat a respective operation by applying a data switching setting according to a current data switching configuration the specified number of times.

**78.** The method of claim 75, further comprising a gasket memory, wherein the gasket memory comprises a plurality of data buffers, an input configuration buffer, an output configuration buffer, a plurality of input reconfigurable counters and a plurality of output reconfigurable counters, and wherein the gasket memory is configured to perform reconfiguration for input and output independently.

**79.** The method of claim 78, further comprising storing data from the first physical data path in the gasket memory to be used as input to the second physical data path.

**80.** The method of claim 76, wherein each of the plurality of PEs comprises a plurality of arithmetic logic units (ALUs) that are configured to execute a same instruction in parallel threads.

**81.** The method of claim 76, wherein each of the plurality of memory ports is configured to operate in a private memory access mode or a shared memory access mode during one configuration.

**82.** The method of claim 81, further comprising accessing a piece of data stored in a memory unit through the private memory access mode and the shared memory access mode in different physical data paths without the piece of data being moved in the memory unit.

**83.** The method of claim 81, wherein each of the plurality of memory ports is configured to access the memory unit using a vector address, wherein in the private memory access mode, one address in the vector address is routed to one memory bank of a memory unit according to a thread index and all private data for one thread are located in a same memory bank.

**84.** The method of claim 81, wherein each of the plurality of memory ports is configured to access the memory unit using a vector address, and wherein in the shared memory access mode, one address in the vector address is routed in a defined region across memory banks regardless of thread index and data shared to all threads are spread in all memory banks.

**85.** The method of claim 76, wherein each of the plurality of PEs comprises a plurality of data buffers and is configured to operate independently during one physical data path configuration.

**86.** The method of claim 75, wherein the plurality of PEs form a PE array and the sequence of instructions is mapped into one or more physical data paths on the processor based on a size of the PE array, connections between the plurality of PEs, and memory access capability.

**87.** A method, comprising:

delivering a first set of configurations to a plurality of reconfigurable units of a processor for the plurality of reconfigurable units to form a first physical data path to execute a first part of a sequence of instructions, each of the first set of configurations including a specified number;

delivering a second set of configurations to the plurality of reconfigurable units for the plurality of reconfigurable units to form a second physical data path to execute a second part of a sequence of instructions, each of the second set of configurations including the specified number;

applying the first set of configurations at the plurality of reconfigurable units for each of the plurality of reconfiguration units to repeat a respective operation the specified number of times to execute the first physical data path;

storing data from the first physical data path to a gasket memory; and

applying the second set of configurations at the plurality of reconfigurable units for each of the plurality of reconfiguration units to repeat a respective operation the specified number of times to execute the second physical data path, with the data stored in the gasket memory as input to the second physical data path.

**88.** The method of claim 87, wherein the gasket memory comprises a plurality of data buffers, an input configuration buffer, an output configuration buffer, a plurality of input

reconfigurable counters and a plurality of output reconfigurable counters, and wherein the gasket memory is configured to perform reconfiguration for input and output independently.

**89.** The method of claim 87, wherein the plurality of reconfiguration units includes a plurality of PEs and a plurality of memory ports, and at least one instruction of the sequence of instructions is a memory access instruction and executed by a memory port the specified number of times before the memory port is reconfigured by applying a next memory port configuration.

**90.** The method of claim 89, wherein the plurality of reconfigurable units further include a plurality of data switching units, and each of the plurality of data switching units is configured to repeat a respective operation by applying a data switching setting according to a current data switching configuration the specified number of times.

**91.** The method of claim 89, wherein each of the plurality of PEs comprises a plurality of arithmetic logic units (ALUs) that are configured to execute a same instruction in parallel threads.

**92.** The method of claim 89, wherein each of the plurality of memory ports is configured to operate in a private memory access mode or a shared memory access mode during one configuration.

**93.** The method of claim 92, further comprising accessing a piece of data stored in a memory unit through the private memory access mode and the shared memory access mode in different physical data paths without the piece of data being moved in the memory unit.

**94.** The method of claim 92, wherein each of the plurality of memory ports is configured to access the memory unit using a vector address, wherein in the private memory access mode,

one address in the vector address is routed to one memory bank of a memory unit according to a thread index and all private data for one thread are located in a same memory bank.

**95.** The method of claim 94, wherein each of the plurality of memory ports is configured to access the memory unit using a vector address, and wherein in the shared memory access mode, one address in the vector address is routed in a defined region across memory banks regardless of thread index and data shared to all threads are spread in all memory banks.

**96.** The method of claim 89, wherein each of the plurality of PEs comprises a plurality of data buffers and is configured to operate independently during one physical data path configuration.

**97.** The method of claim 97, wherein the plurality of PEs form a PE array and the sequence of instructions is mapped into one or more physical data paths on the processor based on a size of the PE array, connections between the plurality of PEs, and memory access capability.

**98.** A processor, comprising:

a plurality of processing elements (PEs) each having a plurality of arithmetic logic units (ALUs) that are configured to execute a same instruction in parallel threads; and

a plurality of memory ports (MPs) for the plurality of PEs to access a memory unit, each of the plurality of MPs comprising an address calculation unit configured to generate respective memory addresses for each thread to access a different memory bank in the memory unit.

**99.** The processor of claim 98, wherein the address calculation unit has a first input coupled to a base address input that provides a base address common to all threads, a second input coupled to a vector address that provides address offsets for each thread individually, and a third input coupled to a counter that is configured to provide thread indexes.

- 100.** The processor of claim 99, wherein one address in the vector address is routed to one memory bank according to a thread index.
- 101.** The processor of claim 99, wherein the memory unit comprises a plurality of memory caches each associated with a different memory bank.
- 102.** The processor of claim 101, wherein each of the plurality of memory ports are coupled to the plurality of memory caches.
- 103.** The processor of claim 101, wherein each memory bank comprises a plurality of memory words and a cache miss in a memory cache causes a word to be fetched from a memory bank associated with the memory cache.
- 104.** The processor of claim 99, wherein each of the plurality of PEs comprises a plurality of data buffers to store data for each thread separately.
- 105.** The processor of claim 99, further comprising a sequencer coupled to the plurality of memory ports, wherein each of the plurality of memory ports comprises a configuration buffer to receive one or more configurations from the sequencer, and each memory port is configured to provide a same memory access pattern during one configuration.
- 106.** The processor of claim 99, wherein consecutive data pieces for one thread are located in one word of a memory bank and continue in a next word of the memory bank.
- 107.** The processor of claim 99, wherein consecutive data pieces for one thread are located in a same position of consecutive words of a memory bank.
- 108.** The processor of claim 99, wherein the plurality of MPs are configured in a one column mode, in which one MP is configured to access the memory unit for all concurrent threads in one PE and the address offsets are independent for each thread.

**109.** The processor of claim 99, wherein the plurality of MPs are configured in a linear mode, in which multiple MPs are configured to access the memory unit, wherein a first MP is configured to access the memory unit for all concurrent threads in a first PE, and a second MP is configured to access the memory unit for all concurrent threads in a second PE, the address offsets in the second MP are linear to the address offsets in the first MP.

**110.** The processor of claim 99, wherein the plurality of MPs are configured in a reverse linear mode, in which multiple MPs are configured to access the memory unit, wherein a first MP is configured to access the memory unit for all concurrent threads in a first PE, and a second MP is configured to access the memory unit for all concurrent threads in a second PE, the address offsets in the second MP are reverse linear to the address offsets in the first MP.

**111.** The processor of claim 99, wherein the plurality of MPs are configured in an overlap mode, in which multiple MPs are configured to access the memory unit, wherein a first MP is configured to access the memory unit for all concurrent threads in a first PE, and a second MP is configured to access the memory unit for all concurrent threads in a second PE, the address offsets in the second MP have overlap with the address offsets in the first MP.

**112.** The processor of claim 99, wherein the plurality of MPs are configured in a non-unity stride mode, in which multiple MPs are configured to access the memory unit, wherein a first MP is configured to access the memory unit for all concurrent threads in a first PE, and a second MP is configured to access the memory unit for all concurrent threads in a second PE, the address offsets in the second MP and the address offsets in the first MP are spaced by a stride.

**113.** The processor of claim 99, wherein the plurality of MPs are configured in a random mode, in which multiple MPs are configured to access the memory unit, and address offsets in different MPs are random numbers.

**114.** The processor of claim 113, wherein the memory unit comprises a plurality of memory caches each associated with a different memory bank and the random numbers are within a range depending on a size of a memory cache.

**115.** The processor of claim 98, wherein the memory unit is configured to be used as registers to store spilled variables for register spilling.

**116.** A method, comprising:

generating a plurality of memory addresses by an address calculation unit in a memory port of a plurality of memory ports, wherein the plurality of memory ports provide access to a memory unit for a plurality of processing elements (PEs) each having a plurality of arithmetic logic units (ALUs) that are configured to execute a same instruction in parallel threads; and

accessing a plurality of memory banks in the memory unit using the plurality of memory addresses with each thread accessing a different memory bank in the memory unit.

**117.** The method of claim 116, wherein the address calculation unit has a first input coupled to a base address input that provides a base address common to all threads, a second input coupled to a vector address that provides address offsets for each thread individually, and a third input coupled to a counter that is configured to provide thread indexes, and the address calculation unit is configured to generate the plurality of memory addresses using the first input, the second input and the third input.

- 118.** The method of claim 117, wherein one address in the vector address is routed to one memory bank according to a thread index.
- 119.** The method of claim 117, wherein the memory unit comprises a plurality of memory caches each associated with a different memory bank, and accessing the plurality of memory banks in the memory unit comprises accessing the plurality of memory caches.
- 120.** The method of claim 119, wherein each of the plurality of memory ports are coupled to the plurality of memory caches.
- 121.** The method of claim 120, further comprising fetching a word from a plurality of words of a memory bank when there is a cache miss in a memory cache associated with the memory bank.
- 122.** The method of claim 117, further comprising storing data for each thread in a separate data buffer in each of the plurality of PEs.
- 123.** The method of claim 117, further comprising receiving one or more configurations by the memory port from a sequencer, wherein the memory port is configured to provide a same memory access pattern during one configuration.
- 124.** The method of claim 117, wherein consecutive data pieces for one thread are located in one word of a memory bank and continue in a next word of the memory bank.
- 125.** The method of claim 117, wherein consecutive data pieces for one thread are located in a same position of consecutive words of a memory bank.
- 126.** The method of claim 117, wherein accessing the plurality of memory banks in the memory unit is in a one column mode, in which one MP is configured to access the memory

unit for all concurrent threads in one PE and the address offsets are independent for each thread.

**127.** The method of claim 117, wherein accessing the plurality of memory banks in the memory unit is in a linear mode, in which multiple MPs are configured to access the memory unit, wherein a first MP is configured to access the memory unit for all concurrent threads in a first PE, and a second MP is configured to access the memory unit for all concurrent threads in a second PE, address offsets in the second MP are linear to address offsets in the first MP.

**128.** The method of claim 117, wherein accessing the plurality of memory banks in the memory unit is in a reverse linear mode, in which multiple MPs are configured to access the memory unit, wherein a first MP is configured to access the memory unit for all concurrent threads in a first PE, and a second MP is configured to access the memory unit for all concurrent threads in a second PE, address offsets in the second MP are reverse linear to address offsets in the first MP.

**129.** The method of claim 117, wherein accessing the plurality of memory banks in the memory unit is in an overlap mode, in which multiple MPs are configured to access the memory unit, wherein a first MP is configured to access the memory unit for all concurrent threads in a first PE, and a second MP is configured to access the memory unit for all concurrent threads in a second PE, address offsets in the second MP have overlap with address offsets in the first MP.

**130.** The method of claim 117, wherein accessing the plurality of memory banks in the memory unit is in a non-unity stride mode, in which multiple MPs are configured to access the memory unit, wherein a first MP is configured to access the memory unit for all concurrent threads in a first PE, and a second MP is configured to access the memory unit for all concurrent

threads in a second PE, address offsets in the second MP and address offsets in the first MP are spaced by a stride.

**131.** The method of claim 117, wherein accessing the plurality of memory banks in the memory unit is in a random mode, in which multiple MPs are configured to access the memory unit, and address offsets in different MPs are random numbers.

**132.** The method of claim 131, wherein the memory unit comprises a plurality of memory caches each associated with a different memory bank and the random numbers are within a range depending on a size of a memory cache.

**133.** The method of claim 116, further comprising storing variables for register spilling.

**134.** A processor, comprising:

a memory unit comprising a plurality of memory banks;

a plurality of processing elements (PEs) each having a plurality of arithmetic logic units (ALUs) that are configured to execute a same instruction in parallel threads; and

a plurality of memory ports (MPs) for the plurality of PEs to access the memory unit, each of the plurality of MPs comprising an address calculation unit configured to generate respective memory addresses for each thread to access a different memory bank in the memory unit.

**135.** The processor of claim 134, wherein the address calculation unit has a first input coupled to a base address input that provides a base address common to all threads, a second input coupled to a vector address that provides address offsets for each thread individually, and a third input coupled to a counter that is configured to provide thread indexes.

**136.** A processor, comprising:

a processing element (PE) having a plurality of arithmetic logic units (ALUs) that are configured to execute a same instruction in parallel threads; and

a memory port (MP) for the PE to access a memory unit, the MP comprising an address calculation unit configured to generate respective memory addresses for each thread to access a different memory bank in the memory unit.

**137.** The processor of claim 136, wherein the PE is one of a plurality of PEs that each has a plurality of ALUs configured to execute a same instruction in parallel threads.

**138.** The processor of claim 137, wherein the MP is one of a plurality of MPs that each has an address calculation unit configured to generate respective memory addresses for each thread in one of plurality of PEs to access a different memory bank in the memory unit.

**139.** A method, comprising:

generating a plurality of memory addresses by an address calculation unit in a memory port, wherein the memory port provides access to a memory unit for a processing element (PE) having a plurality of arithmetic logic units (ALUs) configured to execute a same instruction in parallel threads; and

accessing a plurality of memory banks in the memory unit using the plurality of memory addresses with each thread accessing a different memory bank in the memory unit.

**140.** The method of claim 139, wherein the PE is one of a plurality of PEs that each has a plurality of ALUs configured to execute a same instruction in parallel threads.

**141.** The method of claim 140, wherein the MP is one of a plurality of MPs that each has an address calculation unit configured to generate respective memory addresses for each thread in one of plurality of PEs to access a different memory bank in the memory unit.

**142.** A processor, comprising:

a plurality of processing elements (PEs) each having a plurality of arithmetic logic units (ALUs) that are configured to execute a same instruction in parallel threads; and

a plurality of memory ports (MPs) for the plurality of PEs to access a memory unit, each of the plurality of MPs comprising an address calculation unit configured to generate respective memory addresses for each thread to access a common area in the memory unit.

**143.** The processor of claim 142, wherein the address calculation unit has a first input coupled to a base address input that provides a base address common to all threads, and a second input coupled to a vector address that provides address offsets for each thread individually.

**144.** The processor of claim 143, wherein the address calculation unit is configured to generate a number of memory addresses that match a number of threads in a PE.

**145.** The processor of claim 144, wherein each of the plurality of MPs further comprises a plurality of selection units coupled to the number of memory addresses, each of the plurality of selection units is configured to select zero or more memory addresses to be routed to one memory bank of the memory unit.

**146.** The processor of claim 145, wherein each selection unit is configured with a mask for a different memory bank of the memory unit.

**147.** The processor of claim 143, wherein one MP is configured to access the memory unit for all threads in one PE and the address offsets are same for all threads.

- 148.** The processor of claim 143, wherein multiple MPs are configured to access the memory units for threads in different PEs, the address offsets are the same within one MP but different for different MPs.
- 149.** The processor of claim 143, wherein one MP is configured to access the memory unit for all threads in one PE and the address offsets are sequential in the MP.
- 150.** The processor of claim 143, wherein multiple MPs are configured to access the memory unit for threads in different PEs, the address offsets are the sequential within each MP respectively.
- 151.** The processor of claim 143, wherein one MP is configured to access the memory unit for all threads in one PE, the address offsets are sequential with discontinuity.
- 152.** The processor of claim 143, wherein the plurality of MPs are configured to access the memory unit for different threads in different PEs, the address offsets are sequential with discontinuity in each of the MPs respectively.
- 153.** The processor of claim 143, wherein one MP is configured to access the memory unit for all threads in one PE, the address offsets are linear with non-unity stride.
- 154.** The processor of claim 143, wherein multiple MPs are configured to access the memory unit for all threads in one PE, the address offsets are random but within a small range  $C$  to  $C+R$  dependent on a size of the memory cache.
- 155.** The processor of claim 143, wherein multiple MPs are configured to access the memory unit for threads in different PEs, the address offsets are random but with a small range  $C$  to  $C+R$  dependent on a size of the memory cache.

- 156.** The processor of claim 142, wherein the common area includes all memory banks of the memory unit.
- 157.** The processor of claim 142, wherein the memory unit comprises a plurality of memory caches each associated with a different memory bank.
- 158.** The processor of claim 157, wherein each of the plurality of memory ports are coupled to the plurality of memory caches.
- 159.** The processor of claim 157, wherein each memory bank comprises a plurality of memory words and a cache miss in a memory cache causes a memory word to be fetched from a memory bank associated with the memory cache.
- 160.** The processor of claim 142, wherein each of the plurality of PEs comprises a plurality of data buffers to store data for each thread separately.
- 161.** The processor of claim 142, further comprising a sequencer coupled to the plurality of memory ports, wherein each of the plurality of memory ports comprises a configuration buffer to receive one or more configurations from the sequencer, and each memory port is configured to provide a same memory access pattern during one configuration.
- 162.** A method, comprising:
- generating a plurality of memory addresses by an address calculation unit in a memory port of a plurality of memory ports, wherein the plurality of memory ports provide access to a memory unit for a plurality of processing elements (PEs) each having a plurality of arithmetic logic units (ALUs) that are configured to execute a same instruction in parallel threads; and
  - accessing a plurality of memory banks in the memory unit using the plurality of memory addresses with all threads accessing a common area in the memory unit.

**163.** The method of claim 162, wherein the address calculation unit takes a base address common to all threads as a first input and a vector address that provides address offsets for each thread individually as a second input to generate the plurality of memory address.

**164.** The method of claim 163, wherein the address calculation unit is configured to generate a number of memory addresses that match a number of threads in a PE.

**165.** The method of claim 164, wherein accessing the plurality of memory banks comprises selecting zero or more memory addresses to be routed to one memory bank of the memory unit using a plurality of selection units respectively.

**166.** The method of claim 165, wherein each selection unit is configured with a mask for a different memory bank of the memory unit.

**167.** The method of claim 163, wherein one MP is configured to access the memory unit for all threads in one PE and the address offsets are same for all threads.

**168.** The method of claim 163, wherein multiple MPs are configured to access the memory units for threads in different PEs, the address offsets are the same within one MP but different for different MPs.

**169.** The method of claim 163, wherein one MP is configured to access the memory unit for all threads in one PE and the address offsets are sequential in the MP.

**170.** The method of claim 163, wherein multiple MPs are configured to access the memory unit for threads in different PEs, the address offsets are the sequential within each MP respectively.

**171.** The method of claim 163, wherein one MP is configured to access the memory unit for all threads in one PE, the address offsets are sequential with discontinuity.

**172.** The method of claim 163, wherein the plurality of MPs are configured to access the memory unit for different threads in different PEs, the address offsets are sequential with discontinuity in each of the MPs respectively.

**173.** The method of claim 163, wherein one MP is configured to access the memory unit for all threads in one PE, the address offsets are linear with non-unity stride.

**174.** The method of claim 163, wherein multiple MPs are configured to access the memory unit for all threads in one PE, the address offsets are random but within a small range  $C$  to  $C+R$  dependent on a size of the memory cache.

**175.** The method of claim 163, wherein multiple MPs are configured to access the memory unit for threads in different PEs, the address offsets are random but with a small range  $C$  to  $C+R$  dependent on a size of the memory cache.

**176.** The method of claim 162, wherein the common area includes all memory banks of the memory unit.

**177.** The method of claim 162, wherein the memory unit comprises a plurality of memory caches each associated with a different memory bank.

**178.** The method of claim 177, wherein each of the plurality of memory ports are coupled to the plurality of memory caches.

**179.** The method of claim 177, wherein each memory bank comprises a plurality of memory words and a cache miss in a memory cache causes a memory word to be fetched from a memory bank associated with the memory cache.

**180.** The method of claim 162, wherein each of the plurality of PEs comprises a plurality of data buffers to store data for each thread separately.

**181.** The method of claim 162, further comprising receiving one or more configurations from a sequencer for each of the plurality of memory ports, wherein each memory port is configured to provide a same memory access pattern during one configuration.

**182.** A processor, comprising:

a memory unit comprising a plurality of memory banks;

a plurality of processing elements (PEs) each having a plurality of arithmetic logic units (ALUs) that are configured to execute a same instruction in parallel threads; and

a plurality of memory ports (MPs) for the plurality of PEs to access the memory unit, each of the plurality of MPs comprising an address calculation unit configured to generate respective memory addresses for each thread to access a common area across the plurality of memory banks in the memory unit.

**183.** A processor, comprising:

a processing element (PE) having a plurality of arithmetic logic units (ALUs) that are configured to execute a same instruction in parallel threads; and

a memory port (MP) for the PE to access a memory unit, the MP comprising an address calculation unit configured to generate respective memory addresses for each thread to access a common area in the memory unit.

**184.** The processor of claim 183, wherein the PE is one of a plurality of PEs that each has a plurality of ALUs configured to execute a same instruction in parallel threads.

**185.** The processor of claim 184, wherein the MP is one of a plurality of MPs that each has an address calculation unit configured to generate respective memory addresses for each thread in one of plurality of PEs to access the common area in the memory unit.

**186.** A method, comprising:

generating a plurality of memory addresses by an address calculation unit in a memory port, wherein the memory port provides access to a memory unit for a processing element (PE) having a plurality of arithmetic logic units (ALUs) configured to execute a same instruction in parallel threads; and

accessing a plurality of memory banks in the memory unit using the plurality of memory addresses with each thread accessing a common area in the memory unit.

**187.** The method of claim 186, wherein the PE is one of a plurality of PEs that each has a plurality of ALUs configured to execute a same instruction in parallel threads.

**188.** The method of claim 187, wherein the MP is one of a plurality of MPs that each has an address calculation unit configured to generate respective memory addresses for each thread in one of plurality of PEs to access the common area in the memory unit.

**189.** A processor, comprising:

a plurality of processing elements (PEs) each comprising:

an arithmetic logic unit (ALU);

a data buffer associated with the ALU; and

an indicator associated with the data buffer to indicate whether a piece of data inside the data buffer is to be reused for repeated execution of a same instruction as a pipeline stage.

**190.** The processor of claim 189, further comprising a plurality of memory ports (MPs) for the plurality of PEs to access a memory unit, each of the plurality of MPs comprising an address calculation unit configured to generate respective memory addresses for each thread to access

a common area in the memory unit, wherein one MP of the plurality of MPs responsible for loading a piece of data from the memory unit for the piece of data to be reused at a PE is configured to load the piece of data only once.

**191.** The processor of claim 190, wherein the MP responsible for loading the piece of data to be reused is configured to determine that the piece of data is to be reused by determining that multiple threads to be executed at the PE are loading the piece of data using a same memory address.

**192.** The processor of claim 189, wherein at least one piece of data to be reused is an execution result generated by one PE of the plurality of PEs.

**193.** The processor of claim 189, wherein each PE of the plurality of PEs further comprises a configuration buffer to store configurations for each PE and a reconfiguration counter to count a number of repeated execution, each configuration specifies an instruction to be executed by a respective PE and the number for the instruction to be repeated during a respective configuration.

**194.** The processor of claim 189, wherein the ALU is a vector ALU and the data buffer is a vector data buffer, each data buffer of the vector data buffer is associated with one ALU of the vector ALU.

**195.** The processor of claim 189, further comprising a plurality of memory ports (MPs) for the plurality of PEs to access a memory unit, each of the plurality of MPs comprising an address calculation unit configured to generate respective memory addresses for each thread to access a common area in the memory unit, wherein each MP of the plurality of MPs comprises at least one data buffer for temporarily storing data loaded from the memory unit and each of the at

least one data buffer has an indicator associated therewith to indicate whether a piece of data stored therein is to be reused for other load operation.

**196.** The processor of claim 189, wherein each PE comprises a plurality of data buffers associated with the ALU, each of the plurality data buffers is configured to store a separate input for the ALU and has an associated indicator to indicate whether a respective input is to be reused for repeated execution.

**197.** A method, comprising:

determining that a piece of data is to be shared and reused by all threads at a processing element (PE) of a processor during one configuration applied at the PE;

loading the piece of data once into a data buffer of the PE;

setting an indicator associated with the data buffer to indicate that the piece of data is to be reused; and

executing a same instruction with the piece of data as an input at the PE repeatedly a number of times as a pipeline stage, the same instruction and the number being specified by the configuration.

**198.** The method of claim 197, further comprising loading the piece of data from a memory unit for the piece of data to be loaded into the data buffer of the PE, wherein the processor comprises a plurality of PEs and a plurality of memory ports (MPs) for the plurality of PEs to access the memory unit, wherein one MP of the plurality of MPs responsible for loading the piece of data from the memory unit is configured to load the piece of data only once.

**199.** The method of claim 197, further comprising generating the piece of data as an execution result by one of a plurality of PEs of the processor.

**200.** The method of claim 197, further comprising receiving the configuration and storing the configuration in a configuration buffer of the PE, wherein the configuration specifies an instruction to be executed by the PE and the number for the instruction to be repeated.

**201.** The method of claim 197, wherein determining that the piece of data is to be shared and reused by all threads at the PE comprises determining that the all threads are using a same memory address to access the piece of data.

**202.** The method of claim 197, further comprising:

loading the piece of data once into a data buffer of a memory port, the memory port providing access to a memory unit for the PE; and

setting an indicator associated with the data buffer of the memory port to indicate that the piece of data is to be reused for other load operations accessing a same memory address.

**203.** A processor, comprising:

a plurality of processing elements (PEs) each comprising:

a vector arithmetic logic unit (ALU) including a plurality of ALUs;

a plurality of data buffer associated with each of the plurality of ALU; and

a plurality of indicators each associated with a separate data buffer to indicate

whether a piece of data inside a respective data buffer is to be reused for repeated execution of a same instruction as a pipeline stage at a respective PE.

**204.** The processor of claim 203, further comprising a plurality of memory ports (MPs) for the plurality of PEs to access a memory unit, each of the plurality of MPs comprising an address calculation unit configured to generate respective memory addresses for each thread to access a common area in the memory unit, wherein one MP of the plurality of MPs responsible for

loading a piece of data from the memory unit for the piece of data to be reused at a PE is configured to load the piece of data only once.

**205.** The processor of claim 204, wherein the MP responsible for loading the piece of data to be reused is configured to determine that the piece of data is to be reused by determining that multiple threads to be executed at the PE are loading the piece of data using a same memory address.

**206.** The processor of claim 203, wherein at least one piece of data to be reused is an execution result generated by one PE of the plurality of PEs.

**207.** The processor of claim 203, further comprising a plurality of memory ports (MPs) for the plurality of PEs to access a memory unit, each of the plurality of MPs comprising an address calculation unit configured to generate respective memory addresses for each thread to access a common area in the memory unit, wherein each MP of the plurality of MPs comprises at least one data buffer for temporarily storing data loaded from the memory unit and each of the at least one data buffer has an indicator associated therewith to indicate whether a piece of data stored therein is to be reused for other load operation.

**208.** The processor of claim 203, wherein each of the plurality of data buffers is a vector data buffer having a plurality of data buffer units and a piece of data to be reused for repeated execution is duplicated in all data buffer units of one vector data buffer.

**209.** The processor of claim 203, wherein each of the plurality of data buffers is a vector data buffer having a plurality of data buffer units and a piece of data to be reused for repeated execution is stored only in one data buffer unit of one vector data buffer.

**210.** A method, comprising:

receiving a first configuration and a second configuration at a reconfigurable unit of a processor, the reconfigurable unit having a configuration buffer to store the first configuration and the second configuration;

executing a first operation a first number of times according to the first configuration, the first configuration being part of a first physical data path for executing a first part of a sequence of instructions; and

reconfiguring the reconfigurable unit to execute a second operation a second number of times according to the second configuration, the second configuration being part of a second physical data path for executing a second part of a sequence of instructions.

**211.** A method, comprising:

executing a first instruction at a reconfigurable processing element a number of times according to a first configuration, the reconfigurable processing element being configured to be part of a first physical data path during the first configuration;

delivering an execution result from the reconfigurable processing element to a gasket memory to temporarily store the execution result after each execution of the first instruction; and

feeding the execution result stored in the gasket memory to a second physical data path.

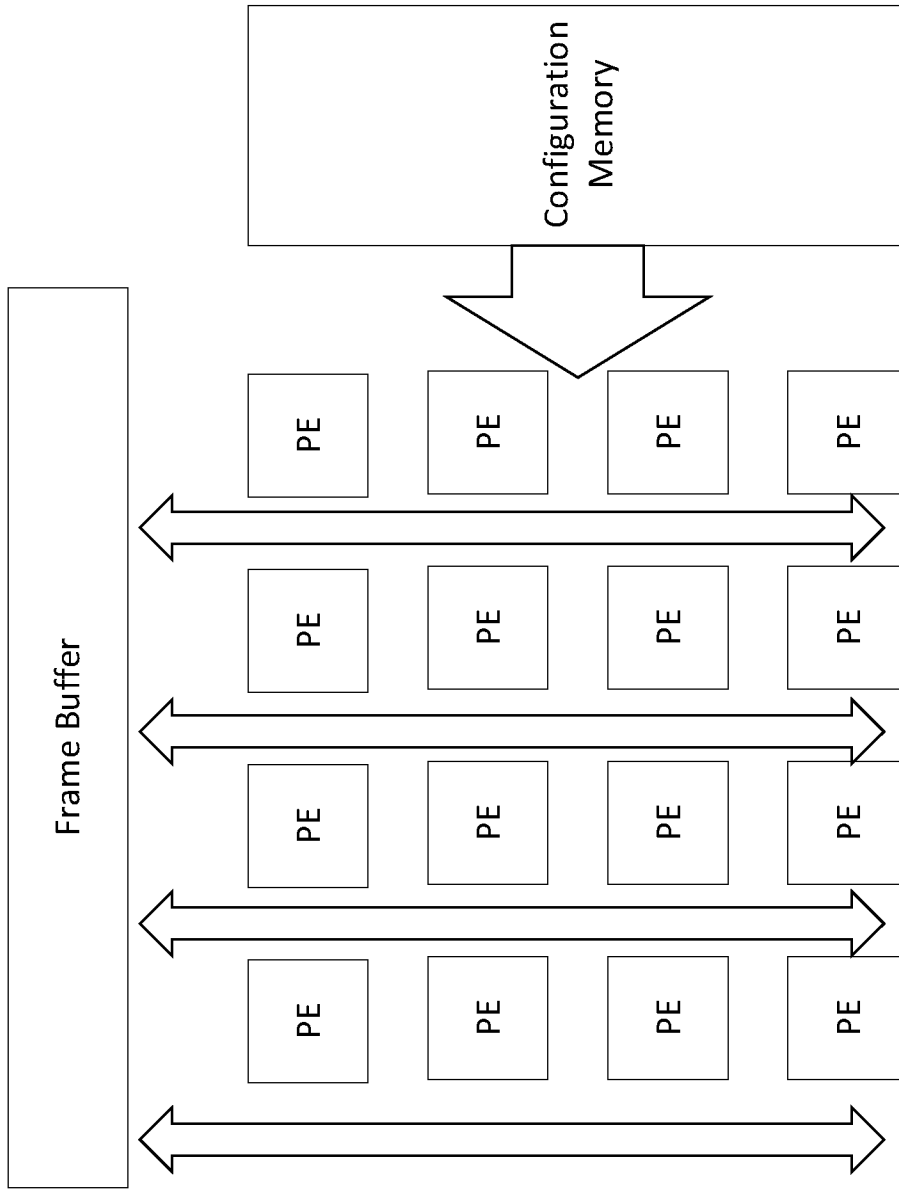


FIG. 1 (Prior Art)




FIG. 2

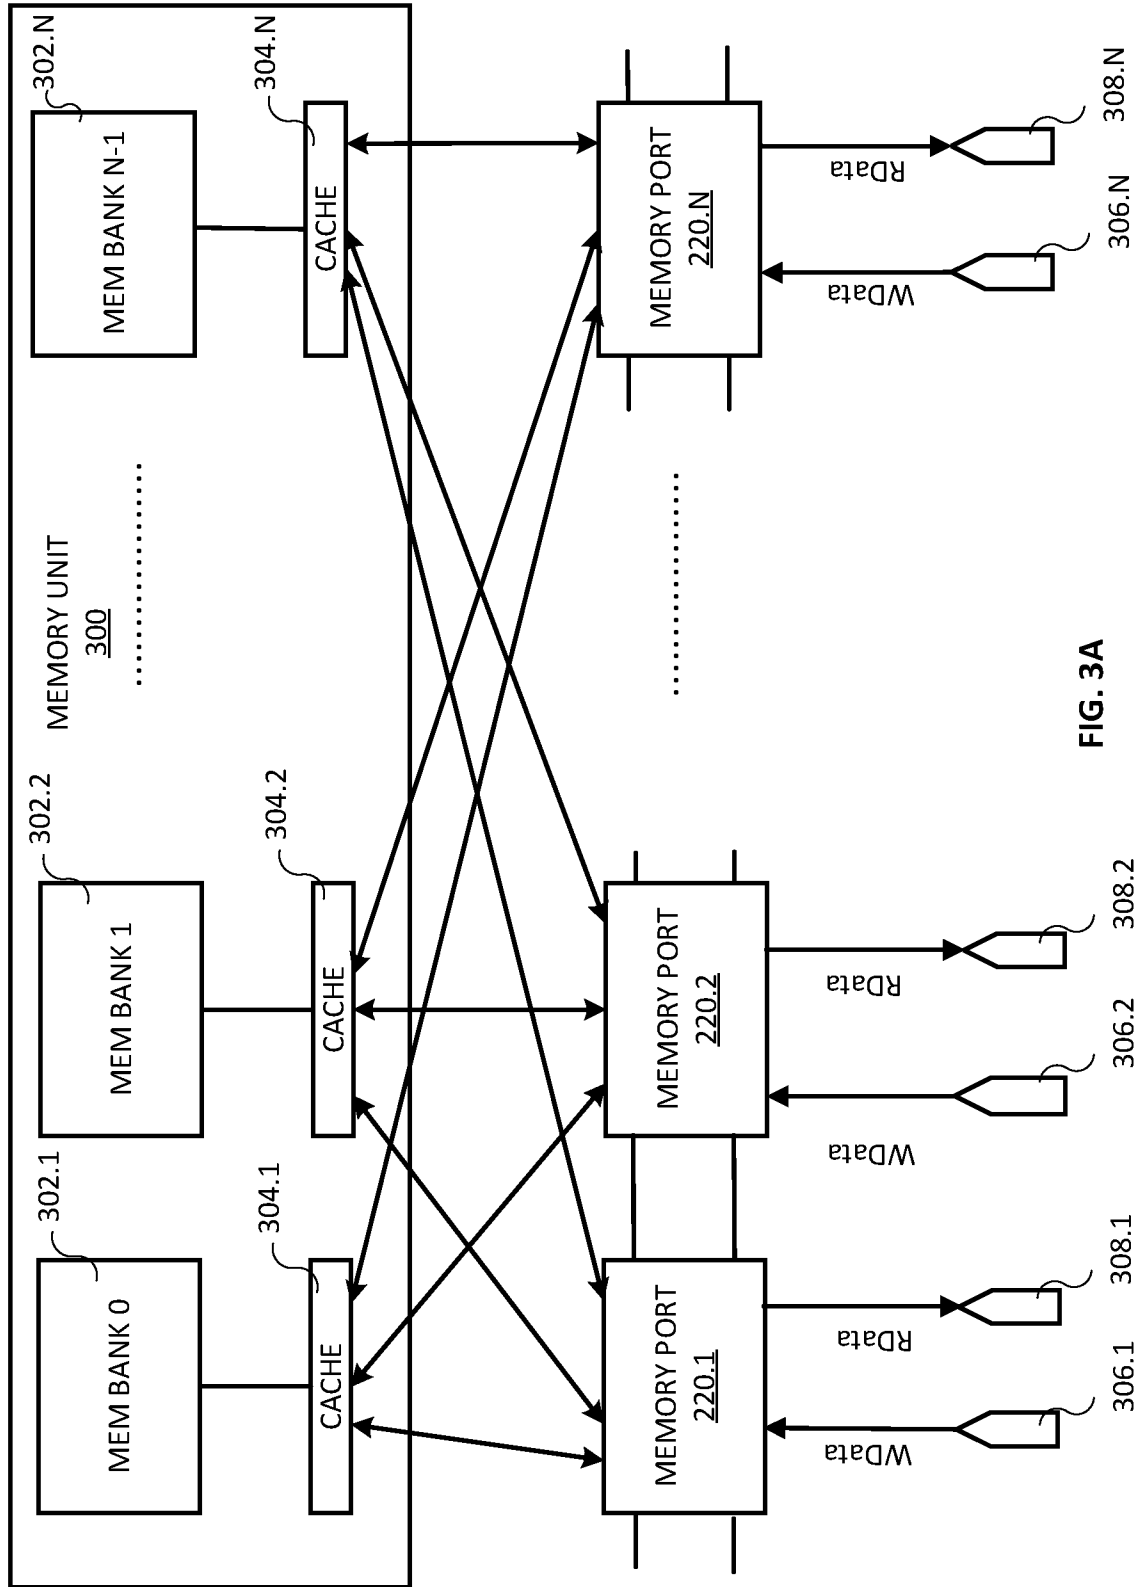


FIG. 3A




FIG. 3B




FIG. 3C




FIG. 3D




FIG. 4A




FIG. 4B






FIG. 6A




FIG. 6B




FIG. 7


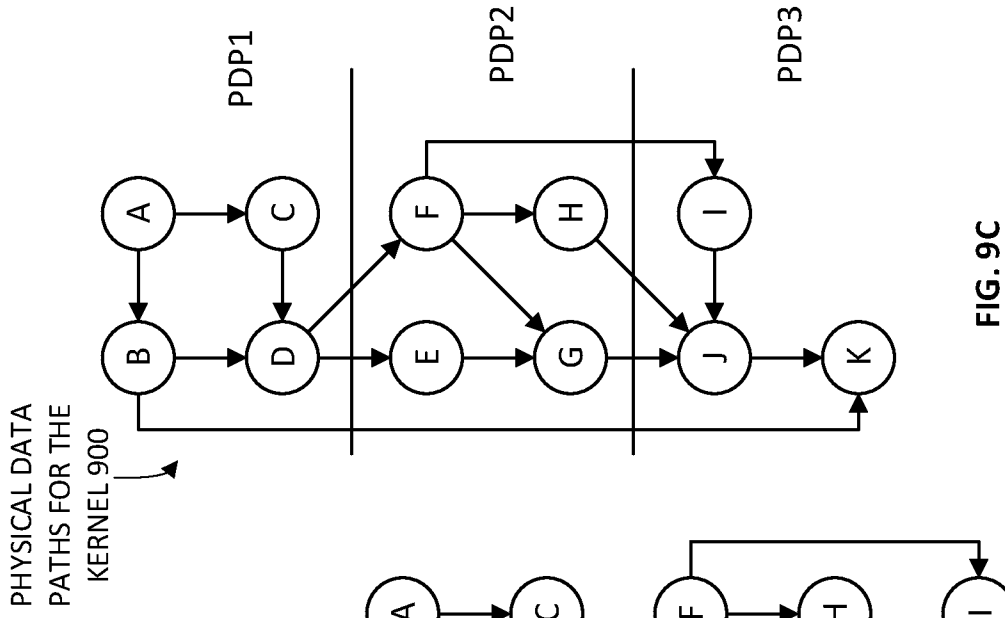




FIG. 8








FIG. 10






FIG. 11C




FIG. 11D




FIG. 11E




FIG. 11F




FIG. 11G




FIG. 11H




FIG. 111




FIG. 12A




FIG. 12B

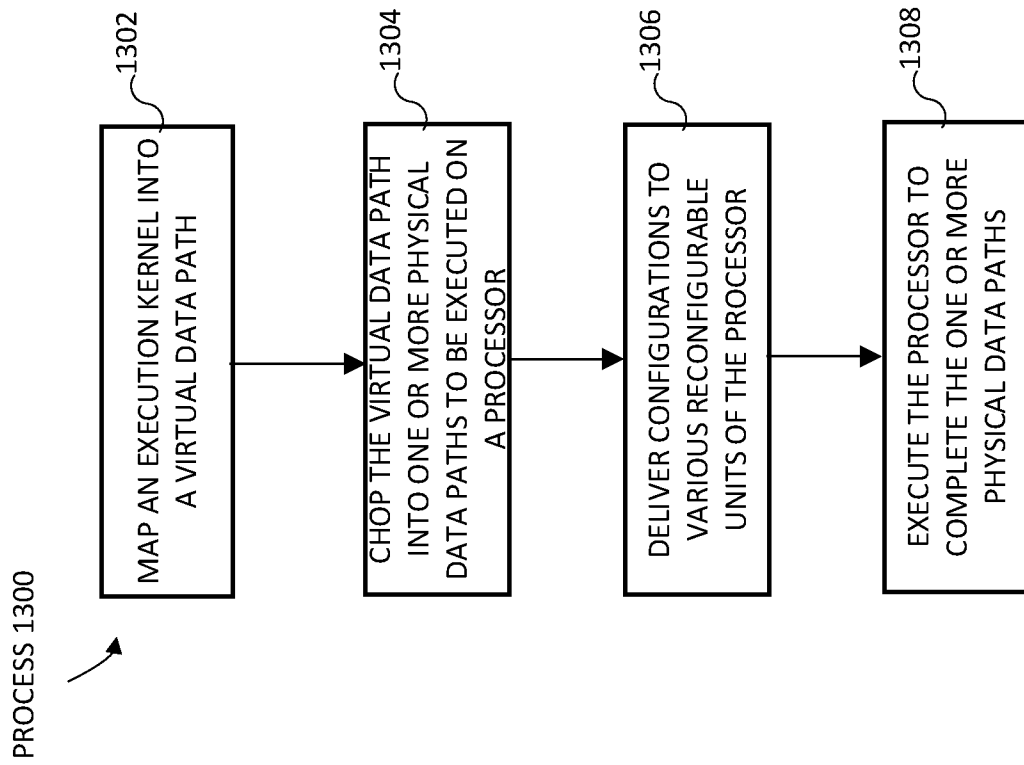


FIG. 13

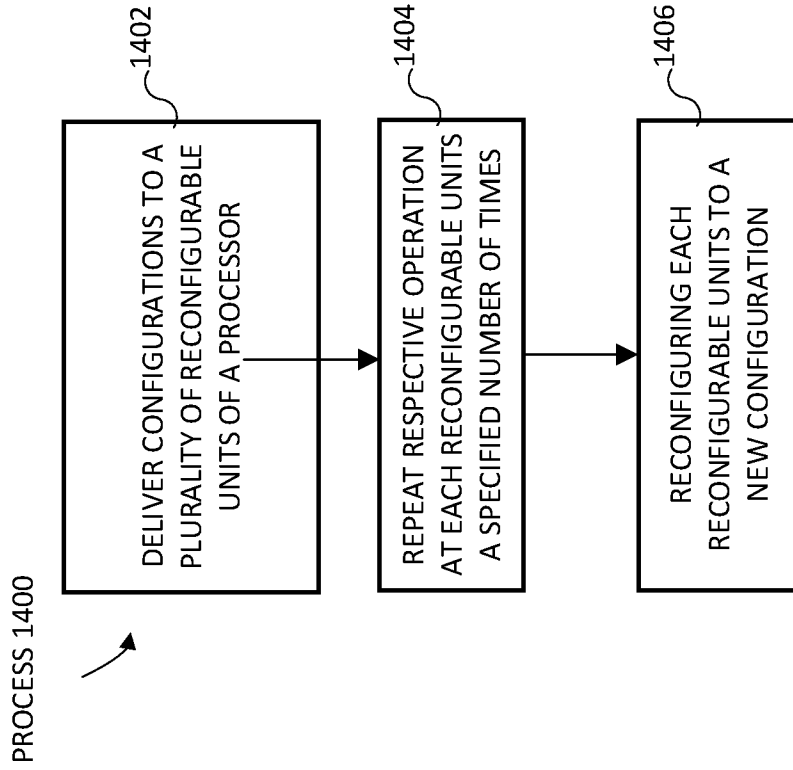


FIG. 14

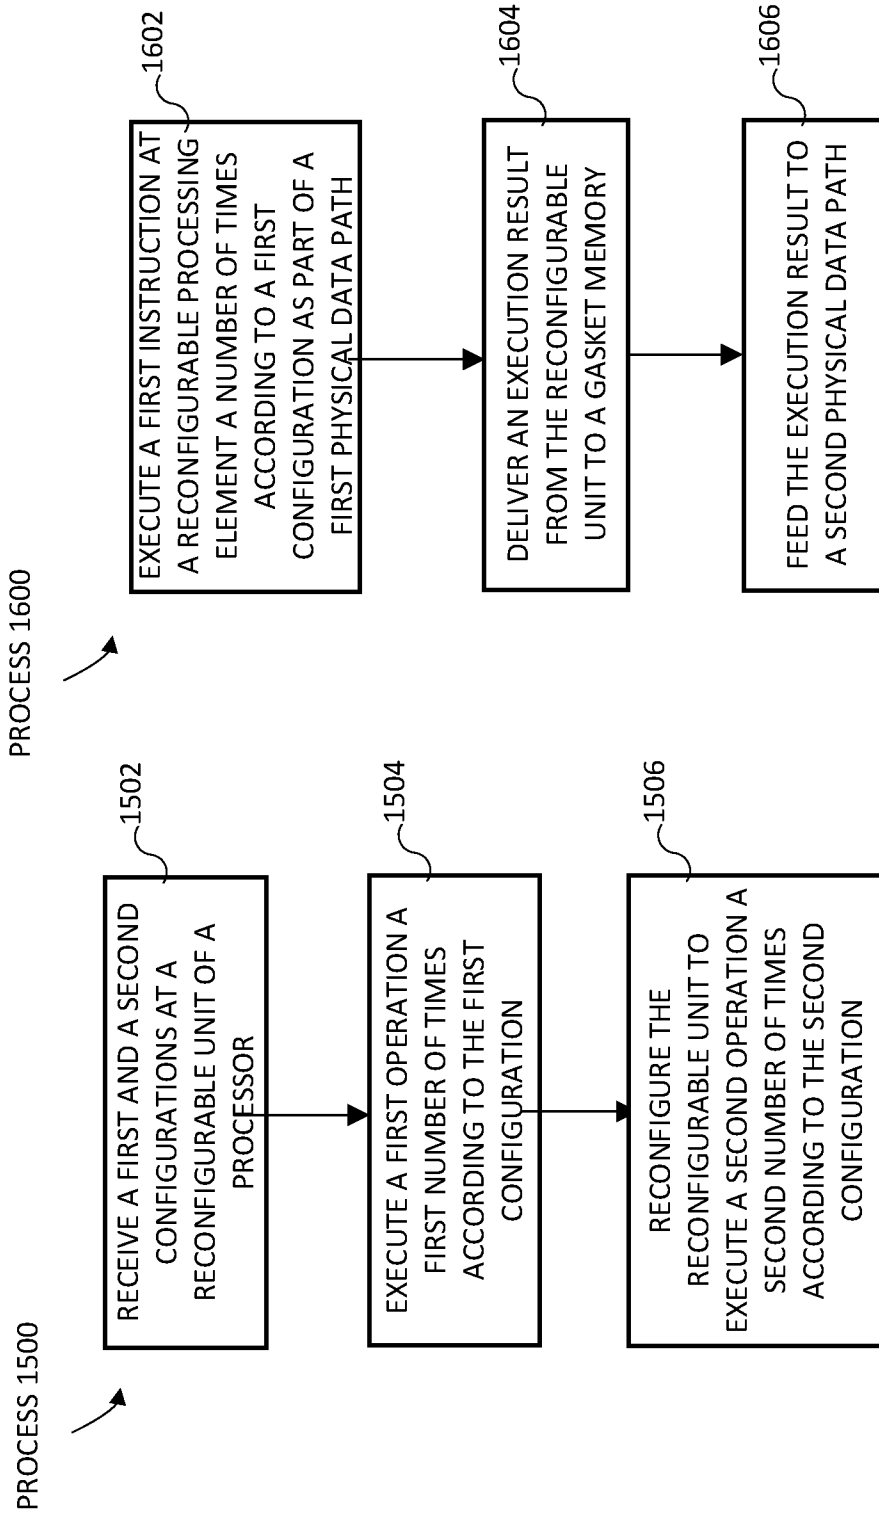
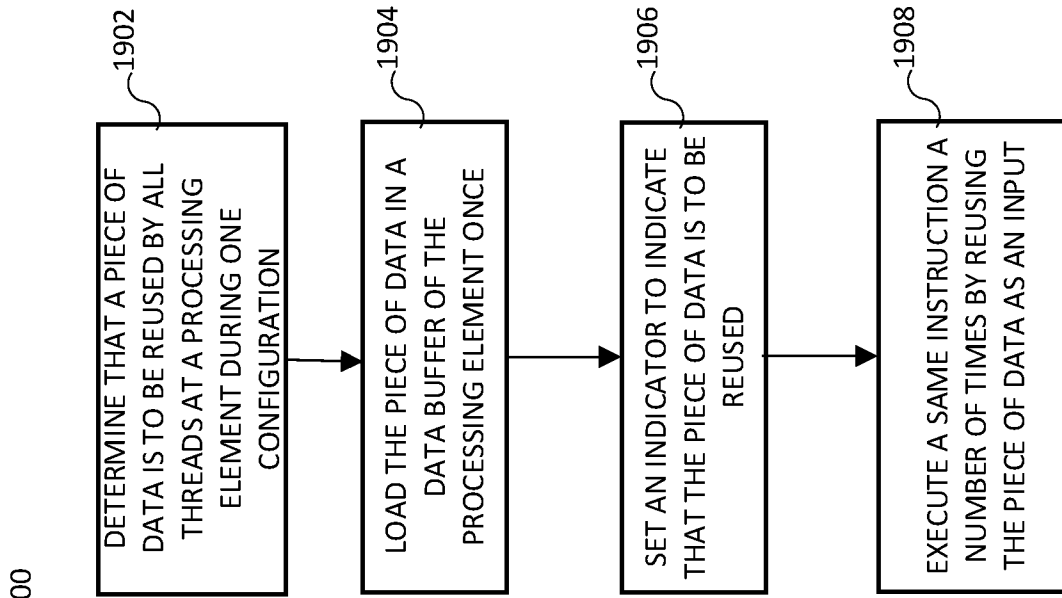


FIG. 15

FIG. 16



PROCESS 1700




FIG. 17

PROCESS 1800

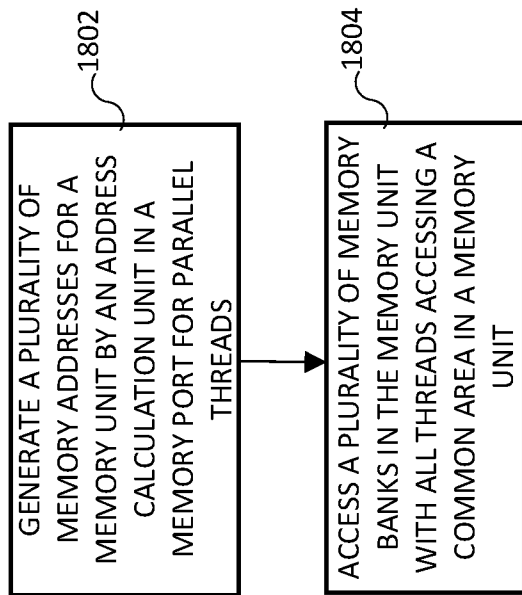


FIG. 18

PROCESS 1900

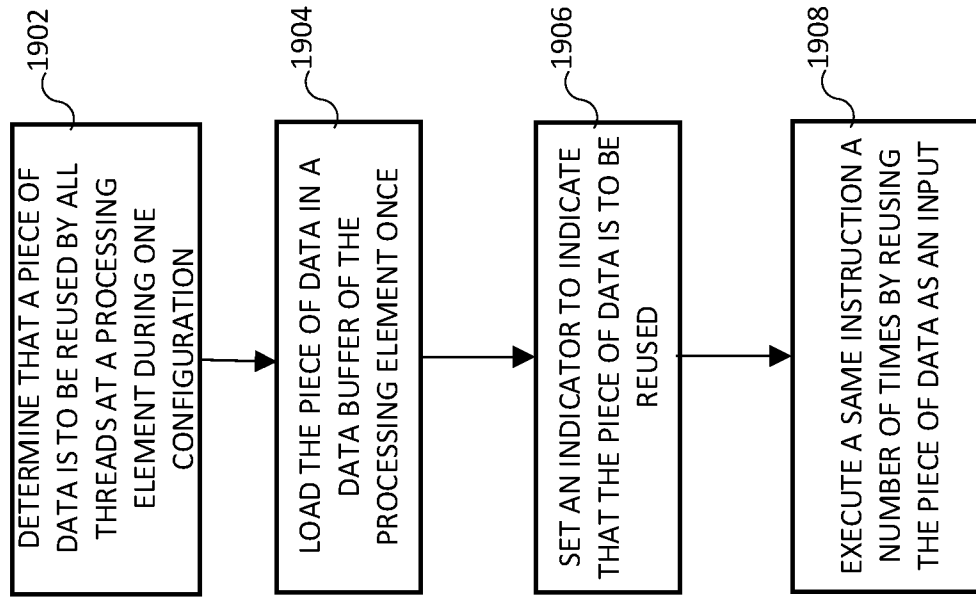


FIG. 19

## INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 18/22106

<b>A. CLASSIFICATION OF SUBJECT MATTER</b> IPC(8) - G06F 15/00, G06F 15/76 (2018.01) CPC - G06F 9/30181, G06F 9/3851, G06F 9/3885, G06F 9/3897, G06F 15/7867, G06F 15/7867, H03K 19/177, Y02D 10/13, H03K 19/17756, Y02D 10/12, H03K 19/17752		
According to International Patent Classification (IPC) or to both national classification and IPC		
<b>B. FIELDS SEARCHED</b>		
Minimum documentation searched (classification system followed by classification symbols) See Search History Document		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched See Search History Document		
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) See Search History Document		
<b>C. DOCUMENTS CONSIDERED TO BE RELEVANT</b>		
<b>Category*</b>	<b>Citation of document, with indication, where appropriate, of the relevant passages</b>	<b>Relevant to claim No.</b>
Y	US 2008/0189514 A1 (McConnell), 07 August 2008 (07.08.2008), entire document, especially Abstract; Para [0027]-[0029], [0035]-[0038]	1-14, 24-38
Y	US 2007/0143577 A1 (Smith), 21 June 2007 (21.06.2007), entire document, especially Abstract; Para [0052], [0056], [0069], [0084]	1-14, 18-20, 29-37
Y	US 2005/0216700 A1 (Honary et al.), 29 September 2005 (29.09.2005), entire document, especially Abstract; para [0019], [0023]-[0024], [0032]	15-23, 57-58
Y	US 8,516,222 B1 (Agarwal et al.), 20 August 2013 (20.08.2013), entire document, especially Abstract; col 1, ln 65 to col 2, ln 10; col 4, ln 24-31; col 11, ln 65 to col 12, ln 15; col 23, ln 20-40	2-38, 57-58
A	US 2016/0202991 A1 (INTERNATIONAL BUSINESS MACHINES CORPORATION), 14 July 2016 (14.07.2016), entire document	1-14, 15-38, 57-58
<input type="checkbox"/> Further documents are listed in the continuation of Box C.		<input type="checkbox"/> See patent family annex.
* Special categories of cited documents: "A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier application or patent but published on or after the international filing date "I" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed		"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art "&" document member of the same patent family
Date of the actual completion of the international search 24 July 2018		Date of mailing of the international search report <b>02 AUG 2018</b>
Name and mailing address of the ISA/US Mail Stop PCT, Attn: ISA/US, Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-8300		Authorized officer: Lee W. Young PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774

**INTERNATIONAL SEARCH REPORT**

International application No.

PCT/US 18/22106

**Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)**

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1.  Claims Nos.:  
because they relate to subject matter not required to be searched by this Authority, namely:
  
2.  Claims Nos.:  
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
  
3.  Claims Nos.:  
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

**Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)**

This International Searching Authority found multiple inventions in this international application, as follows:

This application contains the following inventions or groups of inventions which are not so linked as to form a single general inventive concept under PCT Rule 13.1. In order for all inventions to be examined, the appropriate additional examination fees must be paid.

- Group I - Claims 1-14, 15-38, 57-58 are directed to a method and processor for configuring data path for parallel processing.
- Group II - Claims 1-14, 43-56, 59-74 are directed to a processor architecture for reconfigurable parallel processing.
- Group III - Claims 1-14, 75-97, 189-211 are directed to a method and processor for configuring the processors to perform repeated execution of an instruction.
- Group IV - Claims 39-42 are directed to a processor architecture.
- Group V - Claim 98-188 are directed to a method and processor memory management for configurable parallel processing.

\* Claims 1-14 are generic to Groups I-III

--- ( See Continuation in Supplemental Box ) ---

1.  As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
  
2.  As all searchable claims could be searched without effort justifying additional fees, this Authority did not invite payment of additional fees.
  
3.  As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
  
4.  No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:  
Claims 1-14, 15-38, 57-58

**Remark on Protest**

- The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- No protest accompanied the payment of additional search fees.

## INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 18/22106

Continuation of:

Box III. Observations where unity of invention is lacking

The inventions listed as Groups I-V do not relate to a single general inventive concept under PCT Rule 13.1 because, under PCT Rule 13.2, they lack the same or corresponding special technical features for the following reasons:

Special Technical Features:

The invention of Group I included the features of mapping an execution kernel into a virtual data path at a processor, wherein the execution kernel includes a sequence of instructions to be executed by the processor; chopping the virtual data path into one or more physical data paths; executing the processor to complete the one or more physical data paths by operating the various reconfigurable units according to the configurations, including routing data from one physical data path to the gasket memory to be used in a future physical data path as input, not required by any other groups.

The invention of Group II included the features of each of the plurality of processing elements (PEs) are configured to operate independently according to respective PE configurations stored in the configuration buffer; and each of the plurality of reconfigurable units comprising a reconfiguration counter, not required by any other groups.

The invention of Group III included the features of an indicator associated with the data buffer to indicate whether a piece of data inside the data buffer is to be reused for repeated execution of a same instruction as a pipeline stage; each of the first set of configurations including a specified number; each of the second set of configurations including the specified number; applying the first set of configurations at the plurality of reconfigurable units for each of the plurality of reconfiguration units to repeat a respective operation the specified number of times to execute the first physical data path; applying the second set of configurations at the plurality of reconfigurable units for each of the plurality of reconfiguration units to repeat a respective operation the specified number of times to execute the second physical data path, not required by any other groups.

The invention of Group IV included the features of a plurality of switch boxes (SBs) each associated with a separate PE of the plurality of PEs to providing data switching; and a plurality of inter-column switch boxes (ICSBs) each coupled to a bottom switch box in each column of the plurality of columns; and a gasket memory with its input coupled to a memory port, a PE, one or more switch boxes and an ICSB in a last column of the plurality of columns, and its output coupled to a memory port, one or more switch boxes and an ICSB in a first column of the plurality of columns; store at least one PE execution result to be used by at least one of the plurality of PEs during a next PE configuration, not required by any other groups.

The invention of Group V included the features of a memory unit comprising a plurality of memory banks; a plurality of processing elements (PEs) each configured to execute a same instruction in parallel threads; and each of the plurality of MPs comprising an address calculation unit configured to generate respective memory addresses for each thread to access a different memory bank in the memory unit; access a common area across the plurality of memory banks in the memory unit, not required by any other groups.

Common Technical Features

Groups I-V share the technical features of Processor; a plurality of processing elements (PEs); a gasket memory.

Groups I-III share the technical features of a plurality of processing elements (PEs) each comprising a configuration buffer and a plurality of arithmetic logic units (ALUs), PE configurations stored in the configuration buffer; a sequencer coupled to the configuration buffer of each of the plurality of PEs and configured to distribute one or more PE configurations to the plurality of PEs; and a gasket memory coupled to the plurality of PEs and being configured to store at least one PE execution result.

Groups II and IV-V share the technical feature of a plurality of memory ports (MPs) for the plurality of PEs to access a memory unit.

Groups I and III share the technical feature of delivering configurations to various reconfigurable units of the processor for the various reconfigurable units to form the one or more physical data paths to execute the sequence of instructions.

However, the shared technical features do not represent a contribution over prior art as being obvious over by US 2008/0189514 A1 (McConnell), 07 August 2008 (07.08.2008) in view of US 2007/0143577 A1 (Smith), 21 June 2007 (21.06.2007).

McConnell teaches a processor (Para [0024]- a generic processor 1); a plurality of processing elements (PEs) each comprising a configuration buffer and a plurality of arithmetic logic units (ALUs) (Abstract; Para [0027], [0033] - data processor comprises an array of processing elements (PEn 4); PE 4 includes the usual association of I/O unit 5, local memory 6, register file 7 and arithmetic logic unit (ALU); a small RAM in which is stored a set of configuration data); PE configurations stored in the configuration buffer (Para [0033]- a small RAM in which is stored a set of configuration data); a plurality of memory ports (MPs) for the plurality of PEs to access a memory unit (Fig. 2); a gasket memory coupled to the plurality of PEs and being configured to store at least one PE execution result (Para [0027] - Operands from the register file 7 are connected to the ALU to perform a function as instructed by the control unit 9 and the result fed back into the register file); delivering configurations to various reconfigurable units of the processor for the various reconfigurable units to form the one or more physical data paths to execute the sequence of instructions (Para [0029], [0035]-[0038] data and instruction paths are represented by the various arrows in the drawing. CLUs are connected to the register file in the standard way, i.e. inputs and outputs are of fixed width and fixed location. A number of general purpose microcode bits can be fed into all the CLUs. These can be used to both configure the CLU and to control a configured CLU).

Smith teaches a sequencer coupled to the configuration buffer of each of the plurality of PEs and configured to distribute one or more PE configurations to the plurality of PEs (Para [0052]- master controller 16 is used to configure and reconfigure the processing blocks 2, the micro-code memory 520 of each controller element 22 and reconfigurable interconnect 21. Reconfiguring of these blocks by the master controller 16 means the contents of the memories associated with these blocks are updated with new configuration and program data).

As the common features were known in the art at the time of the invention, this cannot be considered a common technical feature that would otherwise unify the groups. Therefore, Groups I-V lack unity under PCT Rule 13.