



(51) International Patent Classification:
G06F 9/38 (2018.01)

(21) International Application Number:
PCT/IB2021/052105

(22) International Filing Date:
14 March 2021 (14.03.2021)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
16/860,070 28 April 2020 (28.04.2020) US
16/889,159 01 June 2020 (01.06.2020) US

(71) Applicant: **SPEEDATA LTD.** [IL/IL]; Matam Park, building no. 6, P.O. Box 15006, 31905 Haifa (IL).

(72) Inventors: **ETSION, Yoav**; 22 Hatzerot Ha'ikarim Street, 3030000 Atlit (IL). **VOITSECHOV, Dani**; 5 Shahaf Street, 3030000 Atlit (IL). **KRIMER, Evgeni**; 17/48 Almog Street, 3542001 Haifa (IL). **FRIEDMANN, Jonathan**; 5 Ha'etrog Street, P.O. Box 5068, 4053260 Even Yehuda (IL).

(74) Agent: **KLIGLER, Daniel**; KLIGLER & ASSOCIATES PATENT ATTORNEYS LTD., P.O. Box 57651, 6157601 Tel Aviv (IL).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, IT, 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,

(54) Title: COARSE-GRAIN RECONFIGURABLE ARRAY PROCESSOR WITH CONCURRENT HANDLING OF MULTIPLE GRAPHS ON A SINGLE GRID

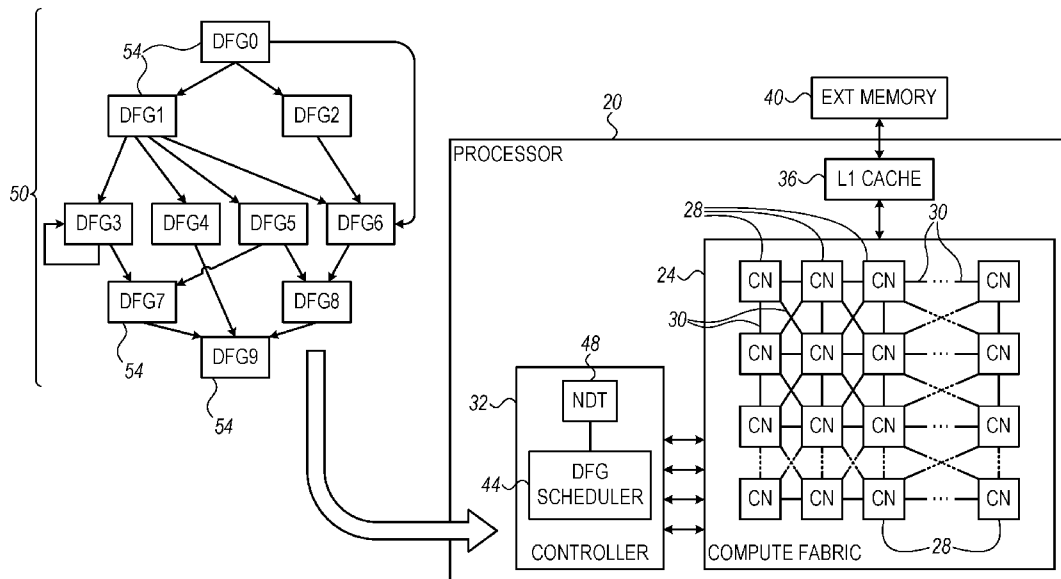


FIG. 1

(57) Abstract: A processor (20) includes a compute fabric (24) and a controller (32). The compute fabric includes an array of compute nodes (28) and interconnects (30) that configurably connect the compute nodes. The controller is configured to receive at least first and second different Data-Flow Graphs (DFGs - 54), each specifying code instructions, and to configure at least some of the compute nodes and interconnects in the compute fabric to concurrently execute the code instructions specified in the first and second DFGs, and send to the compute fabric multiple first threads that execute the code instructions specified in the first DFG and multiple second threads that execute the code instructions specified in the second DFG, thereby causing the compute fabric to execute, at least during a given time interval, both code instructions specified in the first DFG and code instructions specified in the second DFG.



SA, SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM, TN,
TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

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

Published:

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

COARSE-GRAIN RECONFIGURABLE ARRAY PROCESSOR WITH CONCURRENT
HANDLING OF MULTIPLE GRAPHS ON A SINGLE GRID

FIELD OF THE INVENTION

The present invention relates generally to processor architectures, and particularly to
5 Coarse-Grain Reconfigurable Array (CGRA) processors.

BACKGROUND OF THE INVENTION

Coarse-Grain Reconfigurable Array (CGRA) is a computing architecture in which an
array of interconnected computing elements is reconfigured on demand to execute a given
program. Various CGRA-based computing techniques are known in the art. For example, U.S.
10 Patent Application Publication 2018/0101387, whose disclosure is incorporated herein by
reference, describes a GPGPU-compatible architecture that combines a coarse-grain
reconfigurable fabric (CGRF) with a dynamic dataflow execution model to accelerate
execution throughput of massively thread-parallel code. The CGRA distributes computation
across a fabric of functional units. The compute operations are statically mapped to functional
15 units, and an interconnect is configured to transfer values between functional units.

SUMMARY OF THE INVENTION

An embodiment of the present invention that is described herein provides a processor
including a compute fabric and a controller. The compute fabric includes an array of compute
nodes and interconnects that configurably connect the compute nodes. The controller is
20 configured to receive at least first and second different Data-Flow Graphs (DFGs), each
specifying code instructions, and to configure at least some of the compute nodes and
interconnects in the compute fabric to concurrently execute the code instructions specified in
the first and second DFGs, and send to the compute fabric multiple first threads that execute
the code instructions specified in the first DFG and multiple second threads that execute the
25 code instructions specified in the second DFG, thereby causing the compute fabric to execute,
at least during a given time interval, both code instructions specified in the first DFG and code
instructions specified in the second DFG.

In some embodiments, at least in a given clock cycle of the compute fabric, one or
more compute nodes execute code instructions specified in the first DFG and one or more
30 other compute nodes execute code instructions specified in the second DFG. In some
embodiments, at least during the given time interval, one or more of the compute nodes switch
from execution of code instructions specified in the first DFG to execution of code instructions
specified in the second DFG.

In a disclosed embodiment, a first compute node in the compute fabric is configured to switch-over from the first DFG to the second DFG at a first time, and a second compute node in the compute fabric is configured to switch-over from the first DFG to the second DFG at a second time, different from the first time. In an example embodiment, a given compute node is
5 configured to switch-over from the first DFG to the second DFG in response to completing processing all the threads associated with the first DFG. In an embodiment, the controller is configured to send to one or more of the compute nodes a control signal that permits the compute nodes to switch-over from the first DFG to the second DFG, and a given compute node in the compute fabric is configured to switch-over from the first DFG to the second DFG
10 not before receiving the control signal.

In another embodiment, a given compute node, which participates in executing the code instructions specified in the first DFG but does not participate in executing the code instructions specified in the second DFG, is configured to self-deactivate in response to completing processing all the threads associated with the first DFG. Additionally or
15 alternatively, the controller may be configured to activate a given compute node, which does not participate in executing the code instructions specified in the first DFG but participates in executing the code instructions specified in the second DFG.

In some embodiments, the controller is configured to instruct the compute nodes synchronously to alternate between at least the first and second DFGs. In other embodiments,
20 a given compute node is configured to autonomously schedule alternation between at least the first and second DFGs. In an example embodiment, the given compute node is configured to receive first input tokens by communicating with one or more first compute nodes that precede the given compute node in accordance with the first DFG, to receive second input tokens by communicating with one or more second compute nodes that precede the given compute node
25 in accordance with the second DFG, and to execute the first input tokens and the second input tokens in accordance with the scheduled alternation between the first and second DFGs.

There is additionally provided, in accordance with an embodiment of the present invention, a computing method including, in a processor that includes a compute fabric including an array of compute nodes and interconnects that configurably connect the compute
30 nodes, receiving at least first and second different Data-Flow Graphs (DFGs), each specifying code instructions. At least some of the compute nodes and interconnects in the compute fabric are configured to concurrently execute the code instructions specified in the first and second DFGs. Multiple first threads that execute the code instructions specified in the first DFG, and multiple second threads that execute the code instructions specified in the second DFG, are

sent to the compute fabric, thereby causing the compute fabric to execute, at least during a given time interval, both code instructions specified in the first DFG and code instructions specified in the second DFG.

5 The present invention will be more fully understood from the following detailed description of the embodiments thereof, taken together with the drawings in which:

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram that schematically illustrates a Coarse-Grain Reconfigurable Array (CGRA) processor, in accordance with an embodiment of the present invention;

10 Fig. 2 is a block diagram that schematically illustrates an FPU/ALU node in the CGRA processor of Fig. 1, in accordance with an embodiment of the present invention;

Fig. 3 is a block diagram that schematically illustrates an LSU node in the CGRA processor of Fig. 1, in accordance with an embodiment of the present invention;

Fig. 4 is a flow chart that schematically illustrates a method for concurrent execution of multiple DFGs, in accordance with an embodiment of the present invention; and

15 Fig. 5 is a flow chart that schematically illustrates a method for concurrent execution of multiple DFGs, in accordance with an alternative embodiment of the present invention.

DETAILED DESCRIPTION OF EMBODIMENTS

OVERVIEW

20 Embodiments of the present invention that are described herein provide improved methods and devices for Coarse-Grain Reconfigurable Array (CGRA) computing. An example CGRA-based processor comprises a coarse-grain compute fabric, and a controller that configures and controls the compute fabric. The compute fabric comprises (i) an array of compute nodes and (ii) interconnects that configurably connect the compute nodes.

25 The controller receives compiled software code, also referred to as a software program, which is represented as a Data-Flow Graph (DFG). The controller configures the compute nodes and the interconnects in accordance with the DFG, thereby configuring the fabric to execute the program. The controller then sends to the compute fabric multiple threads, which the fabric executes in accordance with the DFG in a pipelined manner.

30 Multiple threads typically execute the same code, as represented by the DFG, on different data. (Strictly speaking, different threads differ slightly from one another in that they depend on the thread number, e.g., in specifying different memory addresses from which to fetch the different data. Nevertheless, in the present context such threads are regarded herein

as executing the same code. While different threads execute the same static code, dynamic execution may differ from one thread to another due to differences in data.)

In many practical cases, however, the software program is far too large to be mapped onto the compute fabric in its entirety. Therefore, in some embodiments of the present invention, the software program provided to the controller is represented as a set of interconnected DFGs, rather than as a single DFG. Each DFG specifies code instructions that perform a respective portion of the program. Each DFG has a single input and one or more exits. The DFGs are interconnected in accordance with a predefined interconnection topology that follows the high-level flow control of the program. This structure is typically provided by a compiler. Note that, even when the compute fabric is large enough to map the entire code, there may be other reasons for dividing the code into multiple DFGs, e.g., due to branches.

In the disclosed embodiments, the controller comprises a scheduler that schedules execution of the various DFGs on the fabric in time alternation. When a certain DFG is scheduled, the scheduler “maps the DFG onto the fabric,” i.e., configures at least some of the compute nodes and interconnects in the fabric to execute the code instructions specified in the DFG. The scheduler sends to the fabric multiple threads for execution, and the fabric executes the threads in accordance with the DFG, typically beginning with reading initial values for the various threads from memory. During the execution of the threads on the currently-scheduled DFG, the fabric saves the results of the various threads to memory, and the scheduler reconfigures the fabric to execute the next DFG. Scheduling, reconfiguration and execution proceeds in this manner until all threads complete executing the software program, or until the controller decides to suspend processing for any reason.

Further aspects of running multiple threads on a CGRA fabric are addressed in U.S. Patent Application Publication 2018/0101387 and U.S. Patent Application 16/860,070, both cited above and incorporated herein by reference.

In some embodiments of the present invention, the CGRA processor, including the compute fabric and the controller, is configured to concurrently execute instructions belonging to multiple different DFGs. The disclosed techniques refer to concurrent handling of multiple *different* DFGs, as opposed to multiple replicas of the same DFG that may run concurrently on different sets of CNs. Hybrid cases, in which the fabric concurrently handles multiple different DFGs, one or more of which having multiple replicas, are also supported. The different DFGs may belong to the same software program or to different programs or tasks, e.g., different database queries. Different programs or tasks may be associated with the same user or with different users.

In the present context, the phrase “concurrent execution of code instructions specified in multiple different DFGs” refers to the compute fabric as a whole, and not necessarily to each individual compute node. In some embodiments that are described herein, the fabric switches-over from one DFG to another progressively. In these embodiments, during a certain transition interval, one or more compute nodes execute instructions specified in the current DFG, while at the same time (e.g., same clock cycle) one or more other compute nodes execute instructions specified in another DFG that has threads ready for execution. Gradual switch-over is highly effective in reducing the high latency overhead, which would otherwise be incurred by switching from one DFG to another.

In other embodiments, one or more of the compute nodes in the fabric alternate (e.g., from one clock cycle to the next) between execution of instructions specified in two or more different DFGs. This technique provides improved processing efficiency and throughput, since the compute-node resources are utilized more efficiently.

SYSTEM DESCRIPTION

Fig. 1 is a block diagram that schematically illustrates a Coarse-Grain Reconfigurable Array (CGRA) processor 20, in accordance with an embodiment of the present invention. Processor 20 is based on a reconfigurable compute fabric 24, which comprises multiple Compute Nodes (CNs) 28 arranged in an array. CNs 28 are also referred to simply as “nodes” for brevity. In an example embodiment, fabric 24 has a total of 256 CNs 28, arranged in a 16-by-16 array, although any other suitable number and arrangement of CNs can be used.

CNs 28 in fabric 24 are connected to one another using interconnects 30, for transferring data and control information (referred to as “tokens”) between CNs. In the present embodiment, each CN 28 comprises a switch (not seen in Fig. 1, examples shown in Figs. 2 and 3 below) that is configurable to connect the CN with one or more selected neighboring CNs. In some embodiments, CNs 28 may send and/or receive additional control information, e.g., configurations and/or exceptions. Such information may be transferred over separate interconnections, e.g., a control network that is separate from interconnects 30. See, for example, the interconnects labeled “TO/FROM CONTROLLER” in Figs. 2 and 3 below.

Processor 20 further comprises a Level-1 (L1) cache 36 (or, alternatively, a multi-level cache), which is connected to an external memory 40 and to fabric 24. Cache 36 and external memory 40 are referred to collectively as the memory system of the processor. At least some of CNs 28 in fabric 24 are configured to store data in the memory system and to load data from the memory system.

In some embodiments, all CNs 28 in fabric 24 are similar in structure. In other embodiments, fabric 24 comprises CNs of two or more different types. Figs. 2 and 3 below show examples of a Floating Point Unit / Arithmetic Logic Unit (FPU/ALU) CN that is configured to perform mathematical computations, and a Load-Store Unit (LSU) CN that is configured to perform memory load and store operations. In some embodiments, the LSU CNs are located along the perimeter of fabric 24, in order to simplify their interconnection to the memory system. FPU/ALU CNs may be located at any suitable locations in fabric 24. Other types of CNs, not discussed in detail herein, may comprise control CNs, and CNs that perform high-complexity computations such as division or square-root computations.

Processor 20 further comprises a controller 32 that, among other functions, configures fabric 24 to carry out desired computational tasks. Controller 32 typically configures the functionality of each CN, and the interconnection topology between CNs. Controller 32 typically configures fabric 24 by sending control registers to CNs 28 over the control network.

Controller 32 typically receives a compiled version of the software code, produced by a suitable compiler. The compiled code comprises a representation of the software code as a structure 50 of interconnected Data Flow Graphs (DFGs) 54. Each DFG 54 represents a respective section of the code, and has a single entry point and one or more exit points. Each DFG is a directed graph, in which vertices represent code instructions and arcs represent the data flow between instructions. An exit point of a DFG (e.g., a conditional or unconditional branch, a function call or a return from a function, etc.) comprises an instruction that leads to an input of a (different or the same) DFG, or terminates the program.

In order to execute a given DFG, controller 32 configures selected CNs 28, and the interconnections between them, to execute the code instructions and the flow-control specified in the DFG. This operation is also referred to as “mapping the DFG onto the fabric.”

A typical DFG begins with readout of initial values (live values) and possibly data from the memory system, then proceeds to perform a sequence of instructions that process the data and possibly store it, and ends with storage of the live values. Generally, however, loading of values from memory (both live values and data) and/or storage of values to memory may be performed at any stage of the DFG, not necessarily at the beginning/end. In an example implementation, controller 32 configures the fabric to execute such a DFG by assigning one or more LSU CNs to read the initial values, assigning appropriate CNs (e.g., FPU/ALU CNs) to execute the instructions specified in the DFG, and assigning one or more LSU CNs to store the results. Controller 32 also configures the switches of one or more CNs, in order to interconnect the CNs according to the flow specified in the DFG. Examples of such

configurations can be found in U.S. Patent Application Publication 2018/0101387 and U.S. Patent Application 16/860,070, both cited above and incorporated herein by reference.

Once configured with a given DFG, fabric 24 is able to execute multiple threads, each thread performing the instructions of the given DFG on different data, in a pipelined manner. In such a pipelined operation, a given CN executes the same instruction in the DFG (with different arguments), for different threads. As soon as the CN completes execution of the instruction for one thread, it is ready to execute it for the next thread. Typically, different threads process different data since load instructions that load data for the thread specify addresses that depend on the thread ID.

Typically, the size of fabric 24 is not sufficient for executing the entire structure 50, i.e., all DFGs, simultaneously. Thus, in some embodiments, controller 32 comprises a DFG scheduler 44 that schedules execution of the various DFGs 54 in time alternation. Scheduler 44 schedules the DFGs using a data structure referred to as a Next DFG Table (NDT) 48. The operation of scheduler 44, and the various data structures used for scheduling, are described in detail in U.S. Patent Application 16/860,070.

Fig. 2 is a block diagram that schematically illustrates an FPU/ALU node 60, in accordance with an embodiment of the present invention. The configuration of node 60 can be used for implementing at least some of CNs 28 in fabric 24 of CGRA processor 20 of Fig. 1.

FPU/ALU node 60 comprises a compute element 64, a switch 68 and a token buffer 72. Switch 68 is configured to connect node 60 to one or more selected neighbor CNs in fabric 24, for receiving input tokens from neighbor CNs and for sending output tokens to neighbor CNs. Input tokens may comprise, for example, data and meta-data (control information) from neighbor CNs to be used as operands in the computation of node 60. Output tokens may comprise, for example, results of the computation of node 60 that are provided to neighbor CNs for subsequent processing. Input and output tokens may also comprise control information.

In the present example, node 60 can be connected to any of eight nearest neighbor CNs 28 in fabric 24, denoted CN0...CN7. As such, node 60 has eight interconnects 30 connected to switch 68. Alternatively, any other suitable interconnection topology can be used.

Token buffer 72 is configured to store input tokens that were received from neighbor CNs and are pending for processing, and/or output tokens that were processed and are pending for outputting to neighbor CNs. Compute element 64 comprises the circuitry that performs the computations of node 60, e.g., an FPU, an ALU, or any other suitable computation.

Consider a scenario in which controller 32 configures fabric 24 to execute a certain DFG 54. As part of this configuration, FPU/ALU node 60 is assigned to execute a certain code instruction in the DFG. To perform this role, controller 32 configures compute element 64 to perform the computation specified in the instruction, e.g., an addition or multiplication.

5 Controller 32 also configures switch 68 to connect the appropriate interconnects 30 to the CNs that supply inputs to the code instruction and consume the result of the code instruction. This switch may also be configured to transfer values to other CNs without going through compute element 64.

In the present example, node 60 further comprises auxiliary logic 76 and a decoder 80.

10 Auxiliary logic 76 comprises configuration registers that hold the configuration parameters of switch 68 and compute element 64, as configured by controller 32. In addition, auxiliary logic 76 may perform functions such as communication with controller 32, interrupt handling, counters management and the like.

In a typical configuration process, controller 32 sends to FPU/ALU node 60 one or more instructions that configure the configuration registers in auxiliary logic 76. Decoder 80 decodes the instructions received from controller 32, and configures the configuration registers accordingly. As will be explained below, in some embodiments node 60 comprises dual configuration registers (also referred to as “shadow registers”) for enabling fast switch-over between DFGs.

20 Fig. 3 is a block diagram that schematically illustrates an LSU node 84, in accordance with an embodiment of the present invention. The configuration of node 84 can be used for implementing at some of CNs 28 in fabric 24 of CGRA processor 20 of Fig. 1.

LSU node 84 is configured to execute load and store instructions, i.e., load data from memory 40 via cache 36, and store data to memory 40 via cache 36. As noted above, LSU nodes are typically located along the perimeter of fabric 24, to simplify their interconnection to the memory system.

LSU node 84 comprises a switch 88, a token buffer 92 and a decoder 98, which are similar in structure and functionality to switch 68, buffer 72 and decoder 80 of FPU/ALU node 60, respectively. Node 84 additionally comprises a compute element 94 that is used, for example, for memory address computations. Node 84 further comprises auxiliary logic 96, which performs functions such as communication with controller 32, communication with L1 cache 36, storage of configuration registers, interrupt handling, counters management, NDT update, as well as initiator and terminator functions (to be explained below).

In alternative embodiments, initiator and terminator functions may be implemented in separate CNs, not in LSU nodes 84. Thus, in the present context, the terms “initiator” and “terminator” refer to any CN that performs initiator and terminator functions, respectively. Initiator and terminator functions may be implemented in dedicated CNs, or combined with other functions in the same CN, e.g., in an LSU node.

In some embodiments, when mapping a given DFG onto fabric 24, scheduler 44 may deactivate (e.g., power-off) one or more of CNs 28 that do not participate in executing the code instructions specified in the given DFG, in order to reduce power consumption. The scheduler may deactivate only parts of a CN, e.g., only a switch (e.g., switch 68 or 88) or only a compute element (e.g., element 64 or 94).

In some embodiments, when configuring fabric 24 to execute a certain DFG, scheduler 44 configures the fabric with multiple replicas of the DFG that will run concurrently. This feature is useful, for example, for DFGs that are considerably smaller than the fabric. In these embodiments, the DFG replicas execute the same section of the program, but are mapped to different subsets of CNs 28. In particular, each replica of the DFG has its own initiator and terminator nodes.

The configurations of CGRA processor 20, FPU/ALU node 60 and LSU node 84, and their various components, as shown in Figs. 1-3, are example configurations that are depicted purely for the sake of conceptual clarity. In alternative embodiments, any other suitable configurations can be used. Additional aspects of multi-threaded CGRA processors are addressed in U.S. Patent Application Publication 2018/0101387 and in U.S. Patent Application 16/860,070, cited above.

In various embodiments, processor 20 and its elements may be implemented using any suitable hardware, such as in an Application-Specific Integrated Circuit (ASIC) or Field-Programmable Gate Array (FPGA). In some embodiments, some of the elements of processor 20, e.g., controller 32, may be implemented using a programmable processor that is programmed in software to carry out the functions described herein. The software may be downloaded to the processor in electronic form, over a network, for example, or it may, alternatively or additionally, be provided and/or stored on non-transitory tangible media, such as magnetic, optical, or electronic memory.

INITIATOR AND TERMINATOR OPERATION

In some embodiments, each of the LSU nodes in fabric 24 (e.g., node 84 of Fig. 3) can serve as an initiator or as a terminator. In other embodiments, the initiators and terminators are

implemented as dedicated CNs. An initiator is a CN that begins execution of threads in a DFG. A terminator is a CN that completes execution of threads in a DFG.

An initiator typically receives thread batches from scheduler 44 for execution. The initiator then computes the Thread IDs (tid's) for the initiated threads and sends the tid's to its
5 successor CNs 28 in fabric 24 (one after another). Upon receiving a tid, each CN 28 executes its configured computation, and forwards the result to the next CN according to the DFG. Thus, each tid travels through fabric 24 as a data token.

Typically, a thread batch is communicated between scheduler 44 and the initiator as a (Thread_batch_id, Thread_batch_bitmap, Thread_set_ID) tuple. When a thread batch arrives
10 in an initiator, the initiator analyzes the bitmap to identify the Thread IDs it should initiate (e.g., by adding the indices of the set bits in the bitmap to Thread_batch_ID). To avoid stalls, the initiator may use prefetching and double buffering ("shadowing") of thread batches. Whenever a thread batch is received, the initiator immediately requests the next thread batch from scheduler 44.

When an entire thread set completes execution on the present DFG (or when the scheduler decides to end the thread set for some reason, e.g., to perform context switch), i.e.,
15 when the scheduler has no remaining thread batches to send, the initiator sends a "final token" down the fabric. This token is not permitted to bypass any other tokens. The final token is sent to all replicas of the DFG, and is used by terminator nodes for verifying that the thread set is
20 completed.

A terminator executes the final branch instruction that exits the present DFG. In the disclosed embodiments, each exit point from a DFG is implemented as a branch instruction, referred to herein as an "exit branch" for brevity. The target address of an exit branch is the start address of the next DFG. An exit branch may be a conditional or unconditional branch
25 instruction. An exit branch may be, for example, a direct branch instruction, an indirect branch instruction, a predicated branch instruction, or a "thread halt" instruction. A DFG may comprise any suitable number of exit points, possibly a single exit point. An exit branch can be located anywhere in the DFG, not necessarily the last instruction.

For exit branches implemented as direct branches, the identities of the successor DFGs
30 are indicated to the terminator node when it is initially configured (as part of configuring the fabric to execute the present DFG). At runtime, the input token (or tokens) to the terminator node, or the terminator node itself, determines the successor DFG according to the branch condition. The successor DFG is determined per thread, depending on the outcome (taken/not taken) of the branch condition for each thread. In an embodiment, the terminator node holds

four 64-bit bitmaps. As the threads of the batch complete execution, the terminator node gradually constructs a pair of 64-bit bitmaps, one bitmap for each possible successor DFG (one bit map for each possible outcome of the branch decision).

5 When all the threads in the batch completed execution, the terminator node holds two bitmaps. In one bitmap, the set bits indicate the threads for which the successor DFG corresponds to the “taken” outcome. In the other bitmap, the set bits indicate the threads for which the successor DFG corresponds to the “not taken” outcome.

10 When the next batch of thread begins execution, the terminator node uses the other pair of 64-bit bitmaps. This mechanism enables the terminator to return to the first pair of bitmaps in case of out-of-order thread batches. If another batch of threads arrives (i.e., a new address) the terminator updates the least-recently used 64-bit bitmap in the NDT.

15 For exit branches implemented as indirect branches, the terminator node typically calculates the target addresses of the successor DFGs at runtime using the above-described mechanism. The terminator node computes the address of each successor DFG and updates the relevant 64-bit bitmap. Once a new 64-bit bitmap is needed (e.g., when the present batch of threads is a new batch, or when the successor DFG is different) the terminator node updates the least-recently used 64-bit bitmap in the NDT.

20 In some embodiments, the terminator node may additionally receive a respective predicate per thread, i.e., a Boolean value that specifies whether the branch is the actual exit point of the DFG or not. The terminator may receive the predicate in an input token. Since a DFG may have several exit points, in the multiple-exit case each exit point will typically be predicated. Per thread, only the actual exit point will receive a true predicate, and the other exit points will receive false predicates.

25 Upon receiving a final token, the terminator node updates all its active 64-bit bitmaps in the NDT (i.e., updates the NDT with the identities of the successor DFG of each completed thread). Once the NDT update is known to be complete, the terminator node notifies scheduler 44 of the completion by sending a final token. This reporting mechanism enables the scheduler to ensure that all terminators, of all the replicas of the present DFG, have completed executing their respective thread batches.

30 In an embodiment, instead of using final tokens, scheduler 44 counts the number of thread batches sent to the initiators (e.g., the number of bitmaps sent to the initiators) and compares this number to the number of completed thread batches reported by the terminators (e.g., to the number of bitmaps reported by the terminators). When the number of sent batches (e.g., bitmaps) equals the number of received batches (e.g., bitmaps), the scheduler concludes

that execution of the present DFG is completed. At this stage, controller 32 may regard the program state as complete (fully defined), save the state, and switch to a new DFG (of the same program/user or of a different program/user).

CONCURRENT EXECUTION OF MULTIPLE DFGs

5 In some embodiments of the present invention, fabric 24 of processor 20 is configured to concurrently execute instructions belonging to multiple different DFGs. The embodiments described below refer mainly to concurrent handling of two DFGs, for the sake of clarity. The disclosed techniques, however, can be used in a similar manner to handle a larger number of DFGs concurrently.

10 The different DFGs may belong to the same software program or to different programs or tasks, e.g., different database queries. Different programs or tasks may be associated with the same user or with different users. When different DFGs belong to the same software program or task, information such as NDT, live values, as well as buffers and cache content, may be shared. When different DFGs belong to different software programs or tasks, such
15 information is typically maintained separately per program or task.

In the present context, the phrase “concurrent execution of code instructions specified in multiple different DFGs” refers to fabric 24 as a whole, and not necessarily to each individual CN 28 in the fabric. In some embodiments CNs 28 switch-over from one DFG to another gradually. In other embodiments, one or more CNs 28 alternate between execution of
20 instructions specified in two or more different DFGs. All of these possibilities are regarded herein as concurrent execution of code instructions specified in multiple different DFGs.

Gradual DFG switch-over

In some embodiments, CNs 28 in fabric 24 switch-over from the current DFG to the
25 next scheduled DFG gradually. For example, a given CN 28 may switch to the next DFG in response to detecting that it has completed its part in executing all the threads in the current thread set. Such CN-specific switch-over reduces the latency overhead incurred by initiating a new DFG.

In these embodiments, the switch-over from one DFG to the next is performed at
30 different times for different CNs 28. Therefore, during a transition interval, some CNs will still execute instructions specified in the current DFG, while other CNs already execute instructions specified in the next DFG.

In the present context, the term “switch-over time” refers to the time at which the CN hardware (e.g., compute element, switch) is reconfigured in accordance with the next DFG. The configuration setting for the next DFG is typically provided to the CN in advance, and stored in the CN’s shadow configuration registers.

5 In an example embodiment, as noted above, the initiator generates a final token after initiating the last batch of threads in the thread set. CNs 28 pass the final token from one to another through fabric 24, according to the data flow specified by the DFG, until reaching the terminator. In some cases, when the data flow splits into multiple paths and later merges, a given CN 28 will receive multiple final tokens. Each CN is aware of the number of final
10 tokens that should pass through it, *en route* to the terminator. Thus, each CN 28 is able to verify whether it has completed its part in executing all the threads in the current thread set, by counting the final tokens passing through the CN.

In some embodiments, a given CN 28 detects that it has forwarded the last final token, and in response switches-over to the configuration of the next DFG. This detection and switch-
15 over is typically independent of any other CN. When CNs 28 operate in this manner, the switch-over to the next DFG will typically occur gradually through fabric 24, with CNs closer to the initiator switching earlier, and CNs further away from the initiator switching later.

In such embodiments, the initiator may begin sending the threads to run on the next DFG via the fabric sooner, before all the CNs have switched over (i.e., while some of the CNs,
20 closer to the terminator, are still busy executing the last threads of the previous DFG). As a result, the latency incurred by initiating the new DFG is reduced.

In one embodiment, once verifying that the last final token has passed, the auxiliary logic of the CN reconfigures all of the CN components (e.g., both the compute element and the switch) to the configuration of the next DFG. In other embodiments, the CN reconfigures only
25 partially when some of the final tokens have passed, e.g., reconfigures only the compute element and part of the switch but not all of it. The remaining CN components are still reconfigured at the same time, e.g., once the last final token reaches the terminator.

Fig. 4 is a flow chart that schematically illustrates a method for concurrent execution of multiple DFGs, in accordance with an embodiment of the present invention. The method is
30 described from the perspective of an individual CN 28. Typically, different CNs 28 carry out this method independently of one another.

The method begins with the CN reading the next token from its token buffer, at a token readout step 100. At a checking step 104, the CN checks whether the token is the final token

of the current DFG. (If the CN expects to receive N final tokens for the current DFG, the CN checks whether the token is the Nth final token.)

5 If the token is not the final token (or not the Nth final token, as applicable), the CN executes the token, at an execution step 108. The method then loops back to step 100 for reading the next token from the token buffer. (In the present context, the phrase “a CN executes a token” means that the CN applies the code instruction it is configured to execute in accordance with the applicable DFG, to an argument provided in the token.)

10 If the token is the final token (or the Nth final token, as applicable), the CN passes the token to the next CN or CNs along the data flow of the current DFG, toward the terminator, at a token passing step 112. The CN then reads the configuration of the next DFG from the shadow registers, and applies the new configuration, at a CN reconfiguration step 116. The method then loops back to step 100 for reading the next token from the token buffer. At this stage, the next token will belong to the newly-configured DFG.

15 The flow of Fig. 4 is a simplified example flow that is depicted for the sake of clarity. In alternative embodiments, gradual switch-over from one DFG to another can be performed in any other suitable way.

20 In some embodiments, the switch-over time is determined entirely by each CN 28. In other embodiments, the switch-over is regulated, at least in part, by controller 32. For example, controller 32 may decide it is time to begin switching-over from the current DFG to the next DFG, and in response send a suitable control signal to CNs 28. In such an embodiment, each CN 28 switches-over to the next DFG after it has received both (i) the last final token for the current DFG, and (ii) the control signal from controller 32.

25 Typically, switching-over a CN 28 to a different DFG involves reconfiguration of various CN elements, e.g., compute element (64, 94) and switch (68, 88). Consider a CN that, in accordance with the current DFG, is configured to receive multiple final tokens via the switch. In one embodiment, the CN waits until the last final token has passed, and then reconfigures the switch in accordance with the next DFG. In an alternative embodiment, the CN attempts to partially-reconfigure the switch, to the extent possible, after passing every final token (not only the last final token).

30 In some embodiments, a CN that is not active in the next DFG may self-deactivate in response to detecting it has forwarded the last final token (or otherwise detecting it has completed its part in handling the current DFG). When using this technique, power

consumption is reduced since CNs power-off independently of one another as soon as they determine they will not be used in the next DFG.

In some embodiments, controller 32 identifies CNs 28 that were inactive in the current DFG but need to be active in the next DFG, and activates them. Such CNs should be treated differently because, for example, they do not receive any final token of the current DFG. In one embodiment, controller 32 powers-on these DFGs once fabric 24 begins switching-over to the next DFG, e.g., upon sending the control signal that permits CNs 28 to switch-over when they are ready. In another embodiment, controller 32 powers-on these DFGs upon sending them the configuration of the next DFG.

In any of the above embodiments, activation and de-activation of a CN may be implemented in any suitable way, e.g., using clock gating or connection/disconnection of supply voltage.

In practice, it is possible that two DFGs that run concurrently on fabric 24 follow one another in the data-flow of the program. In such a case, controller 32 typically ensures that no thread will begin executing on the next DFG before it has completed execution (e.g., reached the terminator) in the current DFG. In one embodiment, this condition is guaranteed by resetting the NDT bits of the threads that execute on the current DFG, as explained above. Alternatively, any other suitable means can be used for this purpose.

As noted above, processor 32 may configure fabric 24 with multiple replicas of a certain DFG that run concurrently. In some embodiments, processor 32 configures the fabric with both multiple replicas of a certain DFG, and multiple different DFGs, all running concurrently. Generally, fabric 24 may concurrently execute any number of different DFGs, each of these DFGs having a single replica or any number of replicas.

Alternation between DFG configurations in CN

In some embodiments, at least some of CNs 28 in fabric 24 alternate between execution of instructions specified in two or more different DFGs. The alternation may be performed, for example, from one clock cycle to the next, or with any other suitable time granularity.

In order to alternate between N different DFGs, a CN 28 typically holds the N configurations of the N DFGs in its configuration registers. Controller 32 sends to the initiator (or initiators) thread batches associated with all the N DFGs being executed concurrently. The initiator typically sends each thread to the relevant CN or CNs, along with information that

associates the thread with the DFG on which it should execute. Alternatively, information that associates each thread with the appropriate DFG may be provided to the CNs in any other suitable way.

Typically, a CN 28 may hold in its token buffer tokens belonging to various different DFGs. (As explained above, a token typically comprises an argument that is provided by a preceding CN in the DFG, to be used as an operand in the code instruction executed by the present CN.) The terms “threads” and “tokens” are sometimes used interchangeably. The CN may schedule the threads in accordance with any suitable scheduling criterion that serves the various DFGs. In one example embodiment, the CN performs Round-Robin scheduling among all the DFGs that have threads ready to execute. Alternatively, the CN may apply suitable priorities among the DFGs that have threads ready to execute.

In some embodiments, the alternation between DFGs is centralized, and is controlled by controller 32. In these embodiments, controller 32 typically instructs the various CNs synchronously, to switch to a specified DFG in a certain clock cycle, then to another specified DFG in the next clock cycle, and so on. In these embodiments, in a given cycle all the CNs execute the same DFG. Centralized control of this sort is relatively simple to implement, control and debug, but is sub-optimal in terms of performance.

In alternative embodiments, the alternation between DFGs is distributed, allowing each CN 28 autonomy in choosing the DFG it will execute in each cycle. For example, each CN 28 may comprise a separate token buffer per DFG, and a separate mechanism for coordinating the reception of tokens from preceding CNs per DFG. In the present context, separate regions in the same buffer are also considered to be separate token buffers for different DFGs. Allocation of the separate buffer spaces may be predefined or performed dynamically at runtime.

In one embodiment, per DFG, the CN sends a “ready” signal to the preceding CNs when it is ready to receive a new input token into the respective token buffer. (Note that the same CN 28 may receive input tokens from different preceding CNs for different DFGs, and send different output tokens to different CNs for different DFGs. Therefore, delivering a token typically involves reconfiguring both the switch of the CN sending the token and the switch of the CN receiving the token. The switch reconfiguration should therefore be coordinated between the two neighbor CNs, e.g., using the “ready” signal).

In other words, the CN carries multiple buffer population processes, e.g., using multiple separate hardware units. Each buffer population process corresponds to a respective DFG. Each buffer-population process interacts with the preceding CN or CNs, e.g., using a

respective “ready” signal, and coordinates the transferal of input tokens from the preceding CN or CNs to the respective token buffer allocated to the respective DFG.

In parallel to of the buffer population processes, the CN autonomously schedules the alternation between DFGs using a suitable scheduling criterion. For example, the CN may perform Round-Robin scheduling among the DFGs having non-empty token buffers. Alternatively, the CN may apply suitable priorities among the DFGs having non-empty token buffers.

Fig. 5 is a flow chart that schematically illustrates a method for concurrent execution of multiple DFGs, in accordance with an alternative embodiment of the present invention. The method is again described from the perspective of an individual CN 28. In this simplified example, the CN alternates between two DFGs denoted “DFG X” and “DFG Y”.

The method begins with the CN receiving new tokens from the preceding CN (or CNs) of DFG X, and from the preceding CN (or CNs) of DFG Y, at token reception steps 120 and 124, respectively. The new tokens are buffered in the token buffers of the CN, separately per DFG.

At a scheduling step 128, the CN selects the DFG whose token will be executed in the next cycle. Any suitable scheduling criterion can be used for this purpose, e.g., Round-Robin scheduling or some prioritization among the non-empty token buffers.

At a token readout step 132, the CN reads a token from the token buffer of the selected DFG. At an execution step 136, the CN executes the read token. The method then loops back to steps 120 and 124 above. As noted above, the phrase “a CN executes a token” means that the CN applies the code instruction it is configured to execute in accordance with the applicable DFG, to an argument provided in the token.

The flow of Fig. 5 is a simplified example flow that is depicted for the sake of clarity. In alternative embodiments, alternation between DFGs can be performed in any other suitable way.

In these embodiments, too, it is possible that two DFGs that run concurrently on fabric 24 follow one another in the data-flow of the program. In such a case, controller 32 typically ensures that no thread will begin executing on the next DFG before it has completed execution (e.g., reached the terminator) in the current DFG. In one embodiment, this condition is guaranteed by resetting the NDT bits of the threads that execute on the current DFG, as explained above. Alternatively, any other suitable means can be used for this purpose.

In some embodiments, processor 32 configures the fabric with both multiple replicas of a certain DFG, and multiple different DFGs, all running concurrently. Generally, fabric 24

may concurrently execute any number of different DFGs, each of these DFGs having a single replica or any number of replicas.

It will be appreciated that the embodiments described above are cited by way of example, and that the present invention is not limited to what has been particularly shown and described hereinabove. Rather, the scope of the present invention includes both combinations and sub-combinations of the various features described hereinabove, as well as variations and modifications thereof which would occur to persons skilled in the art upon reading the foregoing description and which are not disclosed in the prior art. Documents incorporated by reference in the present patent application are to be considered an integral part of the application except that to the extent any terms are defined in these incorporated documents in a manner that conflicts with the definitions made explicitly or implicitly in the present specification, only the definitions in the present specification should be considered.

CLAIMS

1. A processor, comprising:
 - a compute fabric, comprising an array of compute nodes and interconnects that configurably connect the compute nodes; and
 - 5 a controller, configured to:
 - receive at least first and second different Data-Flow Graphs (DFGs), each specifying code instructions; and
 - configure at least some of the compute nodes and interconnects in the compute fabric to concurrently execute the code instructions specified in the first and second
 - 10 DFGs, and send to the compute fabric multiple first threads that execute the code instructions specified in the first DFG and multiple second threads that execute the code instructions specified in the second DFG, thereby causing the compute fabric to execute, at least during a given time interval, both code instructions specified in the first DFG and code instructions specified in the second DFG.
- 15 2. The processor according to claim 1, wherein, at least in a given clock cycle of the compute fabric, one or more compute nodes execute code instructions specified in the first DFG and one or more other compute nodes execute code instructions specified in the second DFG.
3. The processor according to claim 1, wherein, at least during the given time interval,
- 20 one or more of the compute nodes switch from execution of code instructions specified in the first DFG to execution of code instructions specified in the second DFG.
4. The processor according to any of claims 1-3, wherein a first compute node in the compute fabric is configured to switch-over from the first DFG to the second DFG at a first time, and a second compute node in the compute fabric is configured to switch-over from the
- 25 first DFG to the second DFG at a second time, different from the first time.
5. The processor according to any of claims 1-3, wherein a given compute node is configured to switch-over from the first DFG to the second DFG in response to completing processing all the threads associated with the first DFG.
6. The processor according to any of claims 1-3, wherein the controller is configured to
- 30 send to one or more of the compute nodes a control signal that permits the compute nodes to switch-over from the first DFG to the second DFG, and wherein a given compute node in the

compute fabric is configured to switch-over from the first DFG to the second DFG not before receiving the control signal.

7. The processor according to any of claims 1-3, wherein a given compute node, which participates in executing the code instructions specified in the first DFG but does not participate in executing the code instructions specified in the second DFG, is configured to self-deactivate in response to completing processing all the threads associated with the first DFG.

8. The processor according to any of claims 1-3, wherein the controller is configured to activate a given compute node, which does not participate in executing the code instructions specified in the first DFG but participates in executing the code instructions specified in the second DFG.

9. The processor according to any of claims 1-3, wherein the controller is configured to instruct the compute nodes synchronously to alternate between at least the first and second DFGs.

10. The processor according to any of claims 1-3, wherein a given compute node is configured to autonomously schedule alternation between at least the first and second DFGs.

11. The processor according to claim 10, wherein the given compute node is configured to: receive first input tokens by communicating with one or more first compute nodes that precede the given compute node in accordance with the first DFG;

receive second input tokens by communicating with one or more second compute nodes that precede the given compute node in accordance with the second DFG; and

execute the first input tokens and the second input tokens in accordance with the scheduled alternation between the first and second DFGs.

12. A computing method, comprising:

in a processor that includes a compute fabric comprising an array of compute nodes and interconnects that configurably connect the compute nodes, receiving at least first and second different Data-Flow Graphs (DFGs), each specifying code instructions; and

configuring at least some of the compute nodes and interconnects in the compute fabric to concurrently execute the code instructions specified in the first and second DFGs, and sending to the compute fabric multiple first threads that execute the code instructions specified in the first DFG and multiple second threads that execute the code instructions specified in the second DFG, thereby causing the compute fabric to execute, at least during a given time

interval, both code instructions specified in the first DFG and code instructions specified in the second DFG.

13. The method according to claim 12, wherein configuring the compute fabric comprises, at least in a given clock cycle of the compute fabric, causing one or more compute nodes to execute code instructions specified in the first DFG, and causing one or more other compute nodes to execute code instructions specified in the second DFG.

14. The method according to claim 12, wherein configuring the compute fabric comprises, at least during the given time interval, causing one or more of the compute nodes to switch from execution of code instructions specified in the first DFG to execution of code instructions specified in the second DFG.

15. The method according to any of claims 12-14, wherein configuring the compute fabric comprises causing a first compute node in the compute fabric to switch-over from the first DFG to the second DFG at a first time, and a second compute node in the compute fabric to switch-over from the first DFG to the second DFG at a second time, different from the first time.

16. The method according to any of claims 12-14, wherein configuring the compute fabric comprises, in a given compute node, switching-over from the first DFG to the second DFG in response to completing processing all the threads associated with the first DFG.

17. The method according to any of claims 12-14, wherein configuring the compute fabric comprises sending to one or more of the compute nodes a control signal that permits the compute nodes to switch-over from the first DFG to the second DFG, and, in a given compute node in the compute fabric, switching-over from the first DFG to the second DFG not before receiving the control signal.

18. The method according to any of claims 12-14, and comprising self-deactivating a given compute node, which participates in executing the code instructions specified in the first DFG but does not participate in executing the code instructions specified in the second DFG, in response to completing processing all the threads associated with the first DFG.

19. The method according to any of claims 12-14, wherein configuring the compute fabric comprises activating a given compute node, which does not participate in executing the code instructions specified in the first DFG but participates in executing the code instructions specified in the second DFG.

20. The method according to any of claims 12-14, wherein configuring the compute fabric comprises instructing the compute nodes synchronously to alternate between at least the first and second DFGs.

21. The method according to any of claims 12-14, and comprising, in a given compute node, autonomously scheduling alternation between at least the first and second DFGs.

22. The method according to claim 21, wherein scheduling the alternation comprises performing, in the given compute node:

receiving first input tokens by communicating with one or more first compute nodes that precede the given compute node in accordance with the first DFG;

10 receiving second input tokens by communicating with one or more second compute nodes that precede the given compute node in accordance with the second DFG; and

executing the first input tokens and the second input tokens in accordance with the scheduled alternation between the first and second DFGs.

15




FIG. 1

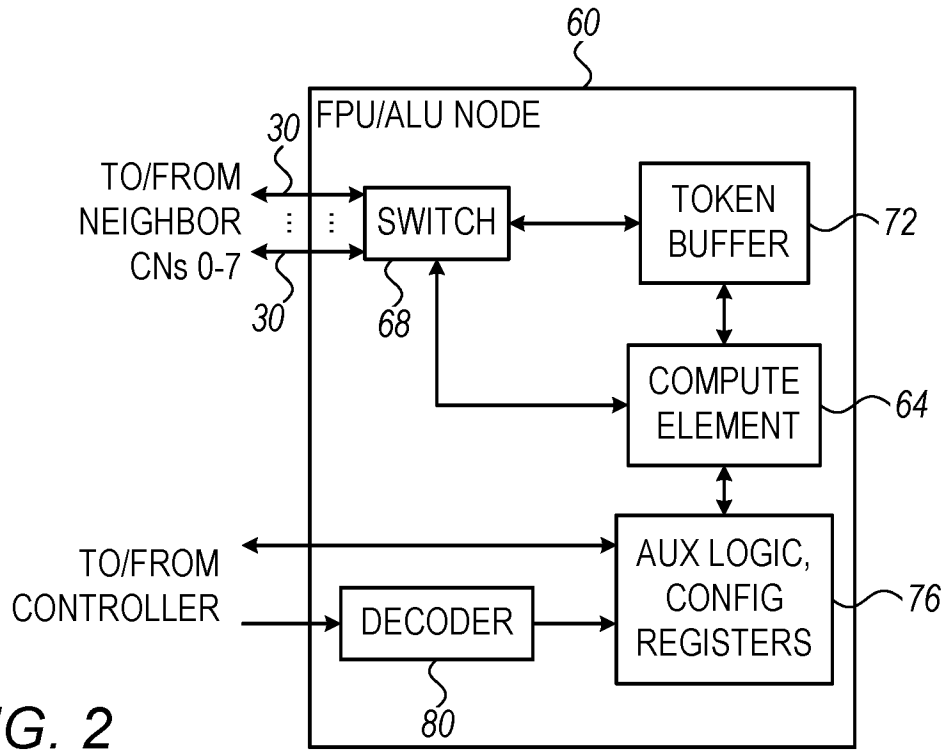


FIG. 2

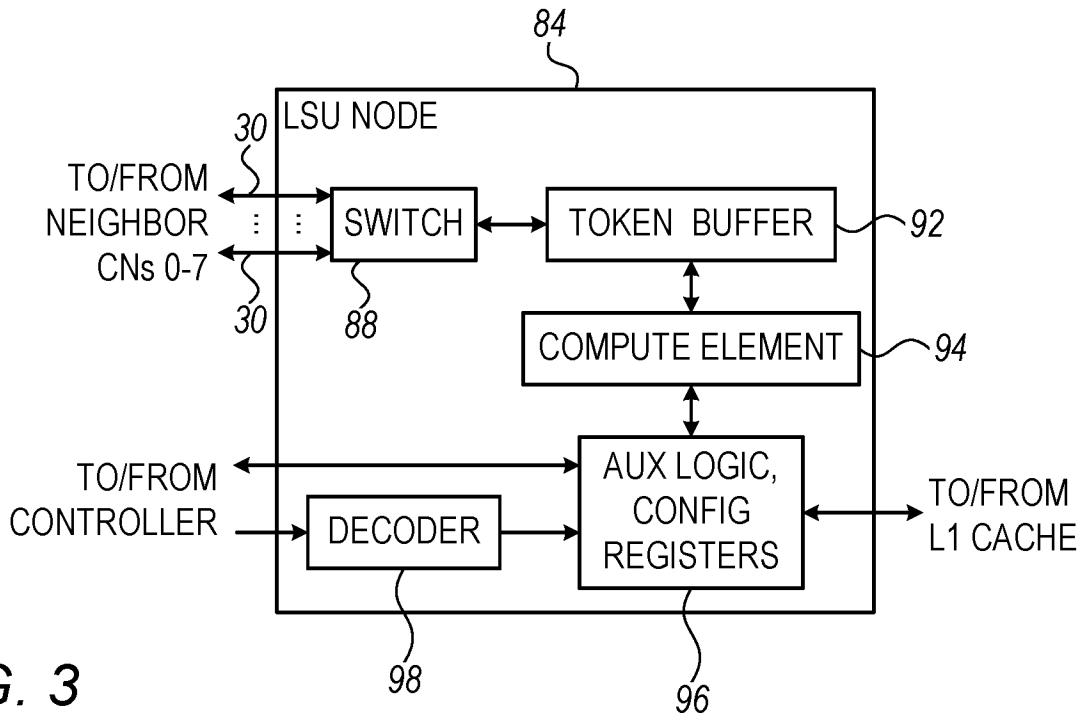


FIG. 3




FIG. 4

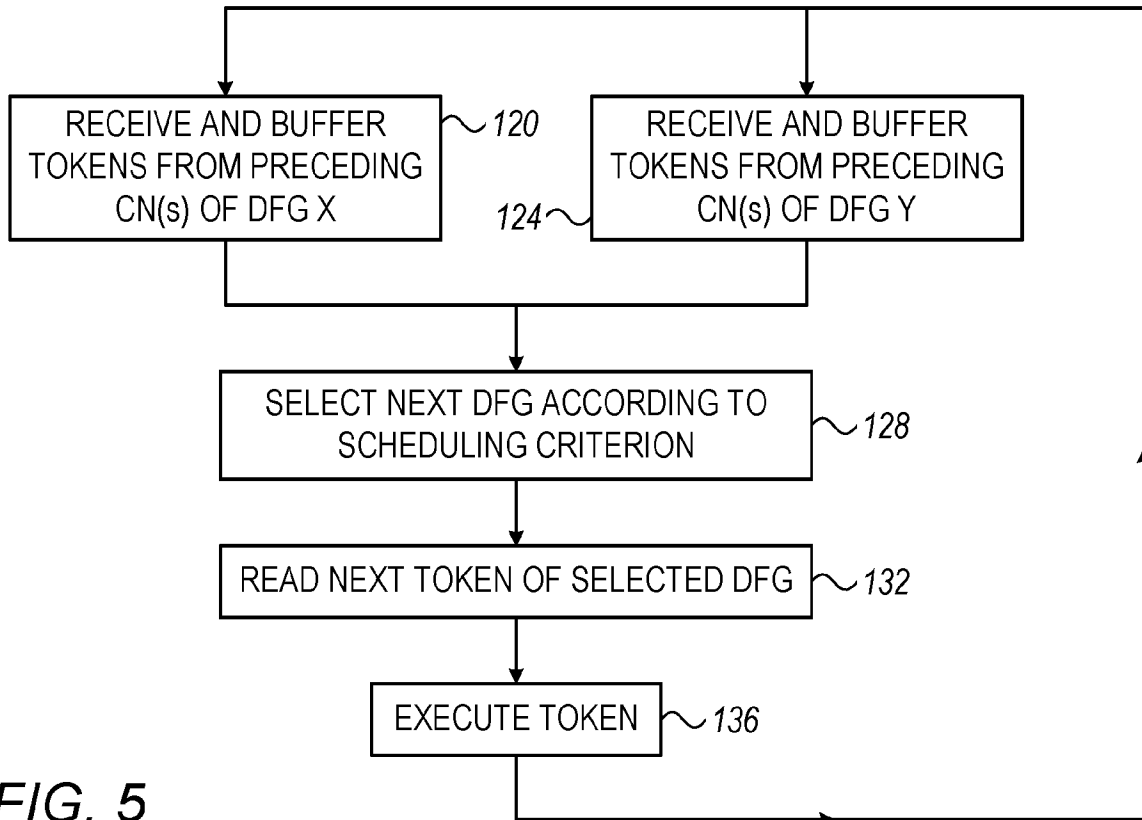


FIG. 5

INTERNATIONAL SEARCH REPORT

International application No.

PCT/IB2021/052105

A. CLASSIFICATION OF SUBJECT MATTER
 IPC (20210101) G06F 9/38
 CPC (20130101) G06F 9/3889, G06F 9/3869, G06F 9/38, G06F 9/3851, G06F 9/3836, G06F 9/3885
 According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
 IPC (20210101) G06F 9/38
 CPC (20130101) G06F 9/3889, G06F 9/3869, G06F 9/38, G06F 9/3851, G06F 9/3836, G06F 9/3885

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

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
 Databases consulted: Esp@cenet, Google Patents, Google Scholar, Orbit, Similari (AI-based)
 Search terms used: different dataflow graphs array compute nodes single grid multiple graphs concurrently threads coarse-grain fabric

C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
|-----------|--|-----------------------|
| X | US 2019004878 A1 ADLER et al. 03 Jan 2019 (2019/01/03) The whole document | 1,2,12,13 |
| A | The whole document | 3-11,14-22 |
| A | US 2017046138 A1 STANFILL et al. 16 Feb 2017 (2017/02/16) The whole document | 1-22 |
| A | US 2019102179 A1 FLEMING et al. 04 Apr 2019 (2019/04/04) The whole document | 1-22 |
| A | US 2018101387 A1 ETSION et al. 12 Apr 2018 (2018/04/12) The whole document | 1-22 |

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

* Special categories of cited documents:
 "A" document defining the general state of the art which is not considered to be of particular relevance
 "D" document cited by the applicant in the international application
 "E" earlier application or patent but published on or after the international filing date
 "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
 "O" document referring to an oral disclosure, use, exhibition or other means
 "P" document published prior to the international filing date but later than the priority date claimed
 "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
 "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
 "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
 "&" document member of the same patent family

| | |
|--|---|
| Date of the actual completion of the international search 31 May 2021 | Date of mailing of the international search report 31 May 2021 |
|--|---|

| | |
|--|---|
| Name and mailing address of the ISA: Israel Patent Office Technology Park, Bldg.5, Malcha, Jerusalem, 9695101, Israel Email address: pctoffice@justice.gov.il | Authorized officer DRUCKER Ekaterina Telephone No. 972-73-9327221 |
|--|---|

INTERNATIONAL SEARCH REPORT
Information on patent family members

| |
|--|
| International application No. PCT/IB2021/052105 |
|--|

| Patent document cited search report | Publication date | Patent family member(s) | Publication Date |
|-------------------------------------|------------------|-------------------------|------------------|
| US 2019004878 A1 | 03 Jan 2019 | US 2019004878 A1 | 03 Jan 2019 |
| | | CN 109213723 A | 15 Jan 2019 |
| US 2017046138 A1 | 16 Feb 2017 | US 2017046138 A1 | 16 Feb 2017 |
| | | US 10037198 B2 | 31 Jul 2018 |
| | | AU 2016306489 A1 | 22 Feb 2018 |
| | | AU 2016306489 B2 | 18 Apr 2019 |
| | | AU 2019204087 A1 | 04 Jul 2019 |
| | | AU 2019204087 B2 | 15 Aug 2019 |
| | | CA 2994420 A1 | 16 Feb 2017 |
| | | CA 2994420 C | 15 Dec 2020 |
| | | CN 108139898 A | 08 Jun 2018 |
| | | CN 108139898 B | 23 Mar 2021 |
| | | EP 3335116 A1 | 20 Jun 2018 |
| | | HK 1256053 A1 | 13 Sep 2019 |
| | | JP 2018530037 A | 11 Oct 2018 |
| | | JP 6598981 B2 | 30 Oct 2019 |
| | | JP 2020013608 A | 23 Jan 2020 |
| | | JP 6763072 B2 | 30 Sep 2020 |
| | | KR 20180034626 A | 04 Apr 2018 |
| | | KR 102182198 B1 | 24 Nov 2020 |
| | | US 2018329696 A1 | 15 Nov 2018 |
| | | US 10423395 B2 | 24 Sep 2019 |
| | | WO 2017027652 A1 | 16 Feb 2017 |
| US 2019102179 A1 | 04 Apr 2019 | US 2019102179 A1 | 04 Apr 2019 |
| | | US 10445098 B2 | 15 Oct 2019 |
| | | CN 109597459 A | 09 Apr 2019 |
| | | DE 102018006889 A1 | 04 Apr 2019 |
| US 2018101387 A1 | 12 Apr 2018 | US 2018101387 A1 | 12 Apr 2018 |

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/IB2021/052105

| Patent document cited search report | Publication date | Patent family member(s) | Publication Date |
|-------------------------------------|------------------|-------------------------|------------------|
| | | US 10579390 B2 | 03 Mar 2020 |
| | | US 2020159539 A1 | 21 May 2020 |
| | | US 11003458 B2 | 11 May 2021 |
| | | US 2015268963 A1 | 24 Sep 2015 |