

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

(19) World Intellectual Property Organization

International Bureau



(10) International Publication Number

WO 2015/167777 A1

(43) International Publication Date
5 November 2015 (05.11.2015)

(51) International Patent Classification:
G06F 9/38 (2006.01)

mirovich; 5775 Morehouse Drive, San Diego, California 92121-1714 (US).

(21) International Application Number:
PCT/US2015/025362

(74) Agent: EVANS, Matthew J.; Shumaker & Sieffert, P.A., 1625 Radio Drive, Suite 300, Woodbury, Minnesota 55125 (US).

(22) International Filing Date:
10 April 2015 (10.04.2015)

(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, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, 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.


(25) Filing Language: English
(26) Publication Language: English
(30) Priority Data:
14/268,215 2 May 2014 (02.05.2014) US

(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,

(71) Applicant: QUALCOMM INCORPORATED [US/US];
ATTN: International IP Administration, 5775 Morehouse Drive, San Diego, California 92121-1714 (US).
(72) Inventors: GRUBER, Andrew Evan; 5775 Morehouse Drive, San Diego, California 92121-1714 (US). CHEN, Lin; 5775 Morehouse Drive, San Diego, California 92121-1714 (US). DU, Yun; 5775 Morehouse Drive, San Diego, California 92121-1714 (US). BOURD, Alexei Vlad-

[Continued on next page]

(54) Title: TECHNIQUES FOR SERIALIZED EXECUTION IN A SIMD PROCESSING SYSTEM



(57) Abstract: A SIMD processor may be configured to determine one or more active threads from a plurality of threads, select one active thread from the one or more active threads, and perform a divergent operation on the selected active thread. The divergent operation may be a serial operation.



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).

— *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

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

Declarations under Rule 4.17:

— *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

TECHNIQUES FOR SERIALIZED EXECUTION IN A SIMD PROCESSING SYSTEM

TECHNICAL FIELD

[0001] The disclosure relates to multi-threaded processing and, more particularly, to techniques for serialized execution in a single instruction multiple data (SIMD) processing system.

BACKGROUND

[0002] A single instruction, multiple data (SIMD) processing system is a class of parallel computing systems that includes multiple processing elements which execute the same instruction on multiple pieces of data. A SIMD system may be a standalone computer or a sub-system of a computing system. For example, one or more SIMD execution units may be used in a graphics processing unit (GPU) to implement a programmable shading unit that supports programmable shading.

[0003] A SIMD processing system allows multiple threads of execution for a program to execute synchronously on the multiple processing elements in a parallel manner, thereby increasing the throughput for programs where the same set of operations needs to be performed on multiple pieces of data. In some situations it may be desirable to perform one or more threads in serial. For instance, in a situation where the output of one thread is needed as an input for one or more other threads. Executing serial operations on a SIMD can be resource intensive and inefficient.

SUMMARY

[0004] This disclosure describes techniques for performing a serialized operation in a SIMD processing system. In examples of the disclosure, a hardware-executed instruction is proposed that selects one thread of a plurality of active threads to execute a serial instruction. The instruction may select a thread without checking thread IDs or using instructions to lock a shared variable. As such, serial operations may be performed quickly and efficiently, with minimal impact to system resources.

[0005] In one example of the disclosure, a multi-threaded processor configured to perform divergent operations comprises a control unit configured to determine one or more active threads from a plurality of threads, and to select one active thread from the one or more active threads, and a plurality of processing elements, wherein one

processing element of the plurality of processing elements associated with the selected active thread is configured to perform a divergent operation.

[0006] In another example of the disclosure, an apparatus configured to perform divergent operations on a multi-threaded processor comprises means for determining one or more active threads from a plurality of threads, means for selecting one active thread from the one or more active threads, and means for performing a divergent operation on the selected active thread.

[0007] In another example of the disclosure, an apparatus configured to perform divergent operations comprises a multi-threaded processor comprising a control unit configured to determine one or more active threads from a plurality of threads, and to select one active thread from the one or more active threads, and a plurality of processing elements, wherein one processing element of the plurality of processing elements associated with the selected active thread is configured to perform a divergent operation.

[0008] In another example of the disclosure, a method of performing divergent operations on a multi-threaded processor comprises determining one or more active threads from a plurality of threads, selecting one active thread from the one or more active threads, and performing a divergent operation on the selected active thread.

[0009] The details of one or more examples of the disclosure are set forth in the accompanying drawings and the description below. Other features, objects, and advantages of the disclosure will be apparent from the description and drawings, and from the claims.

BRIEF DESCRIPTION OF DRAWINGS

[0010] FIG. 1 is a block diagram illustrating a device that may be configured to implement the techniques of this disclosure.

[0011] FIG. 2 is a block diagram illustrating an example graphics processing unit (GPU) that may be used to implement the techniques of this disclosure.

[0012] FIG. 3 is a block diagram illustrating the control unit in the example GPU of FIG. 2 in greater detail.

[0013] FIG. 4 is a flow diagram showing the serialized operation techniques of this disclosure.

[0014] FIG. 5 is a flowchart illustrating an example method of the techniques of this disclosure.

DETAILED DESCRIPTION

[0015] This disclosure describes techniques for performing serial operations in a SIMD processing system (e.g., a GPU). Conventional techniques for performing serial operations in SIMD processing systems can be resource intensive and inefficient. This disclosure proposes an instruction (e.g., executed by hardware in a SIMD processor) that selects one active thread of a plurality of active threads to perform the serial operation, without looping over all inactive threads. This may lead to more efficient serial operation in a SIMD processing system.

[0016] FIG. 1 is a block diagram illustrating an example device that may implement the techniques of this disclosure for serialized execution in a SIMD processing system.

FIG. 1 illustrates device 10 that includes GPU 12, system memory 14, and processor 16, which may be a central processing unit (CPU). Examples of device 10 include, but are not limited to, video devices such as media players, set-top boxes, wireless handsets such as mobile telephones, personal digital assistants (PDAs), desktop computers, laptop computers, gaming consoles, video conferencing units, tablet computing devices, and the like. Device 10 may include components in addition to those illustrated in FIG. 1.

[0017] System memory 14 may be considered as the memory for device 10. System memory 14 may comprise one or more computer-readable storage media. Examples of system memory 14 include, but are not limited to, a random access memory (RAM), an electrically erasable programmable read-only memory (EEPROM), flash memory, or any other medium that can be used to carry or store desired program code in the form of instructions and/or data structures and that can be accessed by a computer or a processor.

[0018] In some aspects, system memory 14 may include instructions that cause processor 16 and/or GPU 12 to perform the functions ascribed to processor 16 and GPU 12 in this disclosure. Accordingly, system memory 14 may be a computer-readable storage medium having instructions stored thereon that, when executed, cause one or more processors (e.g., processor 16 and GPU 12) to perform various functions.

[0019] System memory 14 may, in some examples, be considered as a non-transitory storage medium. The term “non-transitory” may indicate that the storage medium is not embodied in a carrier wave or a propagated signal. However, the term “non-transitory” should not be interpreted to mean that system memory 14 is non-movable or that its contents are static. As one example, system memory 14 may be removed from device 10, and moved to another device. As another example, memory, substantially similar to system memory 14, may be inserted into device 10. In certain examples, a non-transitory storage medium may store data that can, over time, change (e.g., in RAM).

[0020] In some examples, such as examples where device 10 is a wireless handset communication device, processor 16 and GPU 12 may be formed in an integrated circuit (IC). For example, the IC may be considered as a processing chip within a chip package. In some examples, processor 16 and GPU 12 may be housed in different integrated circuits (i.e., different chip packages) such as examples where device 10 is a desktop or laptop computer. However, it may be possible that processor 16 and GPU 12 are housed in different integrated circuits in examples where device 10 is a wireless handset communication device.

[0021] Examples of processor 16 and GPU 12 include, but are not limited to, a digital signal processor (DSP), general purpose microprocessor, application specific integrated circuit (ASIC), field programmable logic array (FPGA), or other equivalent integrated or discrete logic circuitry. In some examples, GPU 12 may be specialized hardware that includes integrated and/or discrete logic circuitry that provides GPU 12 with massive parallel processing capabilities suitable for graphics processing. In some instances, GPU 12 may also include general purpose processing capabilities, and may be referred to as a general purpose GPU (GPGPU) when implementing general purpose processing tasks (i.e., non-graphics related tasks).

[0022] Processor 16 may execute various types of applications. Examples of the applications include web browsers, e-mail applications, spreadsheets, video games, or other applications that generate viewable objects for display. Instructions for execution of the one or more applications may be stored within system memory 14. Processor 16 may transmit graphics data of the viewable objects to GPU 12 for further processing.

[0023] For instance, processor 16 may offload processing tasks to GPU 12, such as tasks that require massive parallel operations. As one example, graphics processing requires massive parallel operations, and processor 16 may offload such graphics processing tasks to GPU 12. Processor 16 may communicate with GPU 12 in

accordance with a particular application processing interface (API). Examples of such APIs include the DirectX® API by Microsoft®, the OpenGL® API by the Khronos group, and the OpenCL™ API; however, aspects of this disclosure are not limited to the DirectX, the OpenGL, or the OpenCL APIs, and may be extended to other types of APIs. Moreover, the techniques described in this disclosure are not required to function in accordance with an API, and processor 16 and GPU 12 may utilize any technique for communication.

[0024] To perform graphics operations, GPU 12 may implement a graphics processing pipeline. The graphics processing pipeline includes performing functions as defined by software or firmware executing on GPU 12 and performing functions by fixed-function units that are hardwired to perform very specific functions. The software or firmware executing on the GPU 12 may be referred to as shader programs (or simply shaders), and the shader programs may execute on one or more shader cores of GPU 12. Shader programs provide users with functional flexibility because a user can design the shader program to perform desired tasks in any conceivable manner. The fixed-function units, however, are hardwired for the manner in which the fixed-function units perform tasks. Accordingly, the fixed-function units may not provide much functional flexibility.

[0025] For example, processor 16 may execute an application, such as a video game, and processor 16 may generate graphics data as part of the execution. Processor 16 may output the graphics data for processing by GPU 12. GPU 12 may then process the graphics data in the graphics pipeline. In some examples, to process the graphic data, GPU 12 may need to execute one or more shader programs. For example, the application executing on processor 16 may cause processor 16 to instruct GPU 12 to retrieve a shader program from system memory 14 and instruct GPU 12 to execute the shader program.

[0026] Device 10 may also optionally include display 60, user interface 62, and transceiver module 64. Device 10 may include additional modules or units not shown in FIG. 1 for purposes of clarity. For example, device 10 may include a speaker and a microphone, neither of which are shown in FIG. 1, to effectuate telephonic communications in examples where device 10 is a mobile wireless telephone. Furthermore, the various modules and units shown in device 10 may not be necessary in every example of device 10. For example, user interface 62 and display 60 may be external to device 10 in examples where device 10 is a desktop computer. As another

example, user interface 62 may be part of display 60 in examples where display 60 is a touch-sensitive or presence-sensitive display of a mobile device.

[0027] Examples of user interface 62 include, but are not limited to, a trackball, a mouse, a keyboard, and other types of input devices. User interface 62 may also be a touch screen and may be incorporated as a part of display 60. Transceiver module 64 may include circuitry to allow wireless or wired communication between device 10 and another device or a network. Transceiver module 64 may include modulators, demodulators, amplifiers and other such circuitry for wired or wireless communication. Display 60 may comprise a liquid crystal display (LCD), a cathode ray tube (CRT) display, a plasma display, a touch-sensitive display, a presence-sensitive display, or another type of display device.

[0028] As will be explained in more detail below, in accordance with the techniques of the disclosure, GPU 12 may be configured to perform a method of performing divergent operations on a multi-threaded processor. GPU 12 may be configured to determine one or more active threads from a plurality of threads, select one active thread from the one or more active threads, and perform a divergent operation on the selected active thread.

[0029] FIG. 2 is a block diagram illustrating an example configuration of GPU 12 that may be used to implement the techniques for serial processing of this disclosure. GPU 12 is configured to execute instructions for a program in a parallel manner. GPU 12 includes a control unit 72, processing elements 74A–74D (collectively “processing elements 74”), an instruction store 76, a data store 78, and communication paths 80, 82, 84, 86A–86D. Communication paths 86A–86D may be referred to collectively as “communication paths 86.” In some examples, GPU 12 may be configured as a single-instruction, multiple-data (SIMD) processing system that is configured to execute a plurality of threads of execution for a program using processing elements 74. In such a SIMD system, processing elements 74 may together process a single instruction at a time with respect to different data items. The program may retire after all of the threads associated with the program complete execution.

[0030] Control unit 72 is configured to control GPU 12 to execute instructions for a program stored in instruction store 76. For each instruction of the program, control unit 72 may retrieve the instruction from instruction store 76 via communication path 80, and process the instruction. In some examples, control unit 72 may process the instruction by causing an operation associated with the instruction to execute on one or more of processing elements 74. For example, the instruction retrieved by control unit

72 may be an arithmetic instruction that instructs GPU 12 to perform an arithmetic operation with respect to data items specified by the instruction, and control unit 72 may cause one or more of processing elements 74 to perform the arithmetic operation on the specified data items. In further examples, control unit 72 may process the instruction without causing an operation to be performed on processing elements 74.

[0031] Control unit 72 may cause an operation to be performed on one or more of processing elements 74 by providing an instruction to processing elements 74 via communication path 82. The instruction may specify the operation to be performed by processing elements 74. The instruction provided to the one or more of processing elements 74 may be the same as or different than the instruction retrieved from instruction store 76. In some examples, control unit 72 may cause the operation to be performed on a particular subset of processing element 74 (including by a single processing element) by one or both of activating a particular subset of processing elements 74 upon which the operation should be performed and deactivating another subset of processing elements 74 upon which the operation should not be performed. Control unit 72 may activate and/or deactivate processing elements 74 by providing respective activation and/or deactivation signals to each of processing elements 74 via communication path 82. In some examples, control unit 72 may activate and/or deactivate processing elements 74 by providing activation and/or deactivation signals to processing elements 74 in conjunction with providing an instruction to processing elements 74. In further examples, control unit 72 may activate and/or deactivate processing elements 74 prior to providing an instruction to processing elements 74.

[0032] Control unit 72 may execute a plurality of threads of execution for a program using processing elements 74. A plurality of threads to be executed in parallel is sometimes called a wave. Each of processing elements 74 may be configured to process instructions of the program for a respective thread of the plurality of threads. For example, control unit 72 may assign each thread of execution to an individual one of processing elements 74 for processing. The threads of execution for the program may execute the same set of instructions with respect to different data items in a set of data items. For example, processing element 74A may execute a first thread of execution for a program stored in instruction store 76 with respect to a first subset of data items in a plurality of data items, and processing element 74B may execute a second thread of execution for the program stored in instruction store 76 with respect to a second subset of data items in the plurality of data items. The first thread of execution may include

the same instructions as the second thread of execution, but the first subset of data items may be different than the second subset of data items.

[0033] In some examples, control unit 72 may activate and deactivate individual threads in the plurality of threads of execution. When control unit 72 deactivates a thread, control unit 72 may also deactivate and/or disable the processing element 74A–74D that is assigned to execute the thread. Such deactivated threads may be referred to as inactive threads. Similarly, when control unit 72 activates a thread, control unit 72 may also activate the processing element 74A–74D that is assigned to execute the thread. Such activated threads may be referred to as active threads. As will be explained in more detail below, control unit 72 may be configured to select an active thread to execute a divergent operation (e.g., a serial operation), without needing to consider other active or inactive threads.

[0034] As used herein, an active thread may refer to a thread that is activated, and an inactive thread may refer to a thread that is deactivated. For a plurality of threads executing on GPU 12 during a given processing cycle, each of the active threads may be configured to process an instruction of the program identified by a global program counter register for the plurality threads during the processing cycle. For example, control unit 72 may activate processing elements 74 that are assigned to active threads in order to configure such processing elements 74 to process the instruction of the program during the processing cycle. On the other hand, for a plurality of threads executing on GPU 12 during a given processing cycle, each of the inactive threads may be configured to not process the instruction of the program during the processing cycle. For example, control unit 72 may deactivate processing elements 74 that are assigned to inactive threads to configure such processing elements 74 to not process the instruction of the program during the processing cycle.

[0035] In some examples, a processing cycle may refer to the time interval between successive loads of the program counter. For example, a processing cycle may refer to the time between when the program counter is loaded with a first value and when the program counter is loaded with a second value. The first and second values may be the same or different values. In examples where the program counter is loaded in an asynchronous manner due to resume check techniques, such asynchronous loads may not, in some examples, serve to differentiate processing cycles. In other words, in such examples, a processing cycle may refer to the time interval between successive

synchronous loads of the program counter. A synchronous load of the program counter may, in some examples, refer to a load that is triggered by a clock signal.

[0036] Sometime prior to the retrieval of the next instruction, control unit 72 determines a next instruction to be processed by GPU 12. The manner in which control unit 72 determines the next instruction to be processed is different depending on whether the instruction previously retrieved by GPU 12 is a control flow instruction. If the instruction previously retrieved by GPU 12 is not a control flow instruction, then control unit 72 may determine that the next instruction to be processed by GPU 12 corresponds to a next sequential instruction stored in instruction store 76. For example, instruction store 76 may store the instructions for a program in an ordered sequence, and the next sequential instruction may be an instruction that occurs immediately after the previously retrieved instruction.

[0037] If the instruction previously retrieved by GPU 12 is a control flow instruction, then control unit 72 may determine the next instruction to be processed by GPU 12 based on information specified in the control flow instruction. For example, the control flow instruction may be an unconditional control flow instruction, e.g., an unconditional branch instruction or a jump instruction, in which case control unit 72 may determine the next instruction to be processed by GPU 12 is a target instruction identified by the control flow instruction. As another example, the control flow instruction may be a conditional control flow instruction, e.g., a conditional branch instruction, in which case control unit 72 may select one of a target instruction identified by the control flow instruction or a next sequential instruction stored in instruction store 76 as the next instruction to process from instruction store 76.

[0038] As used herein, a control flow instruction may refer to an instruction that includes information that identifies a target instruction in instruction store 76. For example, the control flow instruction may include a target value indicative of a target program counter value for the control flow instruction. The target program counter value may be indicative of a target address for the target instruction. The target instruction may, in some examples, be different than the next sequential instruction stored in instruction store 76. High-level program code may include control flow statements such as, e.g., if, switch, do, for, while, continue, break, and goto statements. A compiler may translate the high-level control flow statements into low-level, e.g., machine-level, control flow instructions. An instruction that is not a control flow

instruction may be referred to herein as a sequential instruction. For example, a sequential instruction may not include information that identifies a target instruction.

[0039] For control flow instructions, the information that identifies the target instruction may be a value indicative of a target instruction stored in instruction store 76. In some examples, the value indicative of the target instruction in instruction store 76 may be a value indicative of the instruction address of the target instruction in instruction store 76. The value indicative of the instruction address of the target instruction may, in some cases, be the address of the target instruction in instruction store 76. The value indicative of the instruction address of the target instruction may, in additional cases, be a value used to calculate the address of the target instruction. In further examples, the value indicative of the instruction address of the target instruction may be a value indicative of a target program counter value that corresponds to the target instruction. The value indicative of the target program counter value may, in some cases, be the target program counter value that corresponds to the target instruction. The value indicative of the target program counter value may, in additional cases, be a value used to calculate the target program counter value. The target program counter value that corresponds to the target instruction may, in some examples, be equal to the address of the target instruction.

[0040] A control flow instruction may be a forward control flow instruction or a backward control flow instruction. A forward control flow instruction may be a control flow instruction where the target instruction occurs after the control flow instruction in the ordered sequence of instructions stored in instruction store 76. A backward control flow instruction may be a control flow instruction where the target instruction occurs prior to the next sequential instruction in the ordered sequence of instructions stored in instruction store 76. The next sequential instruction may occur immediately after the control flow instruction in the ordered sequence of instructions.

[0041] A control flow instruction may be a conditional control flow instruction or an unconditional control flow instruction. A conditional control flow instruction includes information that specifies a condition for jumping to the target instruction associated with the control flow instruction. When processing a conditional control flow instruction, if control unit 72 determines that the condition is satisfied, then control unit 72 may determine that the next instruction to be processed is the target instruction. On the other hand, if control unit 72 determines that the condition is not satisfied, then control unit 72 may determine that the next instruction to be processed is the next

sequential instruction stored in instruction store 76. An unconditional control flow instruction does not include information that specifies a condition for jumping to the target instruction associated with the control flow instruction. When processing an unconditional control flow instruction, control unit 72 may unconditionally determine that the next instruction to process is the target instruction identified by the control flow instruction. In other words, the determination in such a case is not conditioned upon any condition specified in the unconditional control flow instruction. As used herein, a condition control flow instruction may be referred to herein as a branch instruction unless the branch instruction is otherwise designated as an unconditional branch instruction. Also, an unconditional control flow instruction may be referred to herein as a jump instruction.

[0042] A conditional branch instruction may include conditions that are specified with respect to one or more data item values. For example, one type of condition may be a comparison condition that compares a first data item value to a second data item value for each active thread executing in GPU 12. Comparing the data item values may include, e.g., determining whether the first data item value is greater than, less than, not greater than, not less than, equal to, or not equal to the second data item value. Another type of condition may be a zero check condition that determines whether a data item value for each active thread executing on GPU 12 is equal to or not equal to zero.

Because each of processing elements 74 operates on different data items, the result of evaluating the condition may be different for each active thread executing on GPU 12.

[0043] If either all of the active threads executing on GPU 12 satisfy the branch condition or all of the active threads executing on GPU 12 do not satisfy the branch condition, then a uniform branching condition occurs and the branching divergence for the threads is said to be uniform. On the other hand, if at least one of the active threads executing on GPU 12 satisfies the branch condition and at least one of the active threads executing on GPU 12 does not satisfy the branch condition, then a divergent branching condition occurs and the branching divergence for the threads is said to be divergent.

[0044] One example of a divergent operation is a serial operation. A serial operation (or series of serial operations) is an operation where the output of one thread (e.g., from the serial operation) is needed as the input for one or more other threads. That is other threads may not continue to process in parallel into the output of one or more other threads is completed. In this sense, threads in a SIMD device may perform in serial for a period of time.

[0045] Typically, the threads executing on GPU 12 may execute the same instruction in a lockstep fashion. In other words, each of processing elements 74 may together execute the same instruction for all active threads during a processing cycle. However, when a divergent branch condition occurs, the threads that satisfy that branch condition may be scheduled to execute next instructions that are different than the next instructions scheduled to be executed by the threads that do not satisfy the branch condition. As discussed above, one such divergent branch condition may be a serial operation.

[0046] Control unit 72 is communicatively coupled to instruction store 76 via communication path 800, to processing elements 74 via communication path 82, and to data store 78 via communication path 84. Control unit 72 may use communication path 800 to send read instructions to instruction store 76. A read instruction may specify an instruction address in instruction store 76 from which an instruction should be retrieved. Control unit 72 may receive one or more program instructions from instruction store 76 in response to sending the read instruction. Control unit 72 may use communication path 82 to provide instructions to processing elements 74, and in some examples, to receive data from processing elements 74, e.g., the result of a comparison instruction for evaluating a branch condition. In some examples, control unit 72 may use communication path 84 to retrieve data items values from data store 78, e.g., to determine a branch condition. Although FIG. 2 illustrates GPU 12 as including a communication path 84, in other examples, GPU 12 may not include a communication path 84.

[0047] Each of processing elements 74 may be configured to process instructions for the program stored in instruction store 76. In some examples, each of processing elements 74 may be configured to perform the same set of operations. For example, each of processing elements 74 may implement the same instruction set architecture (ISA). In additional examples, each of processing elements 74 may be an arithmetic logic unit (ALU). In further examples, GPU 12 may be configured as a vector processor, and each of processing elements 74 may be a processing element within the vector processor. In additional examples, GPU 12 may be a SIMD execution unit, and each of processing elements 74 may be a SIMD processing element within the SIMD execution unit.

[0048] The operations performed by processing elements 74 may include arithmetic operations, logic operations, comparison operations, etc. Arithmetic operations may

include operations such as, e.g., an addition operation, a subtraction operation, a multiplication operation, a division operation, etc. The arithmetic operations may also include, e.g., integer arithmetic operations and/or floating-point arithmetic operations. The logic operations may include operations, such as, e.g., a bit-wise AND operation, a bit-wise OR operation, a bit-wise XOR operation, etc. The comparison operations may include operations, such as, e.g., a greater than operation, a less than operation, an equal to zero operation, a not equal to zero operation, etc. The greater than and less than operations may determine whether a first data item is greater than or less than a second data item. The equal to zero and not equal to zero operations may determine whether a data item is equal to zero or not equal to zero. The operands used for the operations may be stored in registers contained in data store 78.

[0049] Each of processing elements 74 may be configured to perform an operation in response to receiving an instruction from control unit 72 via communication path 82. In some examples, each of processing elements 74 may be configured to be activated and/or deactivated independently of the other processing elements 74. In such examples, each of processing elements 74 may be configured to perform an operation in response to receiving an instruction from control unit 72 when the respective processing element 74A–74D is activated, and to not perform the operation in response to receiving the instruction from control unit 72 when the respective processing element 74A–74D is deactivated, i.e., not activated.

[0050] Each of processing elements 14A–14D may be communicatively coupled to data store 78 via a respective communication path 86A–86D. Processing elements 74 may be configured to retrieve data from data store 78 and store data to data store 78 via communication paths 86. The data retrieved from data store 18 may, in some examples, be operands for the operations performed by processing elements 74. The data stored to data store 78 may, in some examples, be the result of an operation performed by processing elements 74.

[0051] Instruction store 76 is configured to store a program for execution by GPU 12. The program may be stored as a sequence of instructions. In some examples, each instruction may be addressed by a unique instruction address value. In such examples, instruction address values for later instructions in the sequence of instructions are greater than instruction address values for earlier instructions in the sequence of instructions. The program instructions, in some examples, may be machine-level instructions. That is, in such examples, the instructions may be in a format that

corresponds to the ISA of GPU 12. Instruction store 76 is configured to receive a read instruction from control unit 72 via communication path 80. The read instruction may specify an instruction address from which an instruction should be retrieved. In response to receiving the read instruction, instruction store 76 may provide an instruction corresponding to the instruction address specified in the read instruction to control unit 72 via communication path 80.

[0052] Instruction store 76 may be any type of memory, cache or combination thereof. When instruction store 76 is a cache, instruction store 76 may cache a program that is stored in a program memory external to GPU 12. Although instruction store 76 is illustrated as being within GPU 12, in other examples, instruction store 76 may be external to GPU 12.

[0053] Data store 78 is configured to store data items used by processing elements 74. In some examples, data store 78 may comprise a plurality of registers, each register being configured to store a respective data item within a plurality of data items operated on GPU 12. Data store 78 may be coupled to one or more communication paths (not shown) that are configured to transfer data between the registers in data store 78 and a memory or cache (not shown).

[0054] Although FIG. 2 illustrates a single data store 78 for storing data used by processing elements 74, in other examples, GPU 12 may include separate, dedicated data stores for each of processing elements 74. GPU 12 illustrates four processing elements 74 for exemplary purposes. In other examples, GPU 12 may have many more processing elements in the same or a different configuration.

[0055] FIG. 3 is a block diagram illustrating control unit 72 in the example GPU 12 of FIG. 2 in greater detail. Control unit 72 includes a program counter 88, a fetch module 90, a decode module 92 and a control flow module 94. Control flow module 94 may be alternatively referred to herein as a control flow unit.

[0056] Program counter 288 is configured to store a program counter value. In some examples, program counter 88 may be a hardware register, such as, e.g., a program counter register. The program counter value may be indicative of an instruction stored in instruction store 76. The program counter value may, in some cases, be equal to the instruction address of the instruction stored in instruction store 76. In additional cases, the program counter value may be used to compute the instruction address of the instruction stored in instruction store 76. For example, the program counter value may be added to an offset value to generate the instruction address. Program counter 88 may

be referred to herein as a “global program counter” or a “global program counter register” because program counter 88 may be used as a single program counter for all of processing elements 74.

[0057] Fetch module 90 is configured to fetch, e.g., retrieve, an instruction from control unit 72 based on the program counter value stored in program counter 88. For example, fetch module 90 may fetch an instruction from an instruction address identified by the program counter value stored in program counter 88. Fetch module 90 may provide the fetched instruction to decode module 92 for further processing.

[0058] Decode module 92 is configured to decode the instruction received from fetch module 90. Decoding the instruction may involve determining whether the instruction is a type of instruction that can be processed by processing elements 74. If the instruction is a type of instruction that can be processed by processing elements 74, decode module 92 may cause the instruction to execute on one or more of processing elements 74. In some examples, decode module 92 may cause the instruction to execute on all of processing elements 74. In other examples, decode module 92 may cause the instruction to execute on less than all of processing elements 74. Causing the instruction to execute on one or more of processing elements 74 may, in some cases, include issuing the instruction to one or more of processing elements 74 for execution. For example, fetch module 90 may fetch a sequential instruction identified by program counter 88, and issue the sequential instruction to all processing elements 74 that correspond to active threads for processing. If the instruction is not the type of instruction that can be processed by processing elements 74, then control unit 72 may process the instruction without issuing the instruction to any of processing elements 74 for processing. For example, the instruction may be a control flow instruction of the type that does not require processing by processing elements 74, in which case control unit 72 may process the instruction without issuing the instruction any of processing elements 74.

[0059] In either case, decode module 92 may forward control information to control flow module 94 for further processing. In some examples, the control information may be the instruction itself. In further examples, the control information may include information, such as, e.g., information indicative of whether the instruction is a control flow instruction or a sequential instruction; if the instruction is a control flow instruction, information indicative of whether the instruction is a branch instruction or a jump instruction; if the instruction is a branch or jump instruction, information

indicative of whether the branch or jump instruction is a forward or backward branch or jump instruction, and if the instruction is a branch instruction, information specifying the branch condition.

[0060] Instructions that are of a type that can be processed by processing elements 74 may include arithmetic instructions and logic instructions. An arithmetic instruction may refer to an instruction that instructs processing elements 74 to perform an arithmetic operation, and a logic instruction may refer to an instruction that instructs processing elements 74 to perform a logic operation. In some examples, a control flow instruction may be an instruction that can be processed by processing elements 74, e.g., the control flow instruction may include a branch condition that is evaluated by processing elements 74. Instructions that are not of a type that can be processed by processing elements 74 may include control flow instructions where the branch condition is evaluated by control unit 72 and/or control flow instructions that do not have a branch condition.

[0061] Control flow module 94 may determine a program counter value associated with a next instruction to be processed by control unit 72, and load the program counter value into program counter 88. If the previously fetched instruction is a sequential instruction, then control flow module 94 may select a program counter value that is indicative of a next sequential instruction stored in instruction store 76 to load into program counter 88. If the previously fetched instruction is a control flow instruction, then control flow module 94 may utilize resume counter techniques.

[0062] Control flow module 94 may store a resume counter value for each thread executing on GPU 12. For example, the number of resume counter values stored in control flow module 94 may be equal to the number of processing elements 74 contained GPU 12. For each resume counter value, if the thread corresponding to the respective resume counter value is inactive, then the resume counter value may be indicative of a program counter value at which the inactive thread should be activated or reactivated. Otherwise, if the thread corresponding to the respective resume counter value is active, then the resume counter value may be, in some examples, set to a maximum value, i.e., a value that is the largest value that can be represented in the storage slot or register for the resume counter.

[0063] Control flow module 94 may store an active flag for each thread executing on GPU 12. For example, the number of active flags stored in control flow module 94 may be equal to the number of processing elements 74 contained in GPU 12. Each active

flag may indicate whether or not the thread associated with the active flag is active or inactive. In some examples, the active flag may be a single bit that is set to indicate that the thread associated with the active flag is active, and reset to indicate that the thread associated with the active flag is inactive.

[0064] In some examples, control flow module 94 may also store a program state. For example, a first program state may indicate that all threads are active, a second program state may indicate that at least one thread is active and at least one thread is inactive and a third program state may indicate that all threads are inactive. The program state may be used in such examples, to select a program counter value to load into program counter 88.

[0065] Control flow module 94 may be configured, in some examples, to activate and deactivate one or more of processing elements 74 via communication path 82. In additional examples, control flow module 94 may instruct decode module 92 to activate and deactivate particular processing elements 74. In further examples, control flow module 94 may receive the results of a comparison instruction from one or more of processing elements 74 via communication path 82. The results of the comparison instruction in some examples may be used to evaluate a branch condition. In yet further examples, control flow module 94 may retrieve one or more data items from data store 78, via communication path 84, for purposes of evaluating a branch condition.

[0066] As discussed above, within a parallel machine (e.g., a SIMD processing system), there is often the need for performing divergent operations, such as a serialized operation(s). For example, GPU 12 may be instructed to perform a 'prefix sum,' where the input to a given thread(s) executing on processing elements 74 is dependent on the output of another thread executing on processing elements 74.

[0067] One conventional approach to performing serialized operations on a SIMD processing system involves using atomic native instructions whereby all threads would compete to perform a lock on a shared variable (i.e., so only one of processing elements 74 has access to the variable at a certain time). However, the use atomic instructions and locks can be inefficient, as they are a limited resource and typically require a long latency for execution.

[0068] Another approach for performing a serial operation involves generating unique value for each thread (e.g., a thread_id). GPU 12, through control unit 72, would then be configured to loop through all threads, selecting only a single thread (by Thread ID)

for each iteration of the loop. Example pseudocode for such an operation is shown below:

```
For (i = 0; i)           // i may represent the total number of threads in a wave
{ if (thread_id == i) {perform serialized operation} }
```

[0069] The above approach requires additional testing and test variables, as well forcing a loop iteration even for those threads that might be inactive. This is both inefficient in resources and time.

[0070] In view of these drawbacks, this disclosure proposes techniques for performing serialized operations in a parallel processing system (e.g., a SIMD processing system). In examples below, this disclosure describes an instruction (e.g., a hardware-executed instruction) whereby an active thread from a plurality of threads may be selected to perform a serialized operation, without the need to access any inactive threads.

[0071] The solution proposed is in this disclosure is a “GetOne” instruction (GetOne_and_branch()) that may be executed by control unit 72 of GPU 12. Control unit 72 may configured to select a single active thread out of a plurality of threads (e.g., the threads in wave) for execution. Control unit 72 may determine which threads are currently active using the active flag stored by control flow module 94. Once the single active thread is selected, all other active threads in the wave may execute a branch condition, while the selected thread performs a divergent operation (e.g., the serial operation). The “GetOne” instruction can be then used to execute a serial operation as shown by the following pseudo code:

```
LOOP: GetOne_and_branch TRY AGAIN
    {perform serialized operation}
    goto CONTINUE
TRY AGAIN: goto LOOP
```

CONTINUE

[0072] In the example above, control unit 72 is initially instructed to go to the LOOP instruction of the code. The LOOP instruction includes the “GetOne” instruction (GetOne_and_branch) that selects a single active thread from a plurality of active threads (e.g., a wave) and proceeds to the serial instruction ({perform serialized

operation}) and then exits the control block (goto CONTINUE) and goes inactive. Unselected threads perform the branch instruction of the GetOne_and_branch instruction. In the above example, the branch instruction is a “TRY AGAIN” instruction, which essentially sends all unselected active threads back to the Loop instruction. Then, another of the remaining active threads may be selected to perform the serialized operation.

[0073] The mechanism of the “GetOne” instruction allows for efficient serialization within a wave without additional variables and without requiring a loop iteration for inactive threads. The techniques of this disclosure take advantage of the fact that the parallel threads within a wave already have some level of synchronization due to their shared single instruction. In this way, the “GetOne” instruction of this disclosure allows for one active thread to be selected for serial operation, while simultaneously making all other active threads inactive.

[0074] FIG. 4 is a flow diagram showing the serialized operation techniques of this disclosure in more detail. Initially, control unit 72 of GPU 12 may receive a “GetOne” instruction (100). In response to the “GetOne” instruction, control unit 72 is configured to determine which threads in a plurality of threads (e.g., a wave) are active (102). In one example, control unit 72 may be configured to read an active flag associated with each of the threads to determine which threads are active. If control unit 72 determines that no active threads are remaining (103), the control block ends (return 112).

[0075] If control unit 72 determines that there are active threads remaining (103), control unit 72 may then select one active thread of the plurality of threads (104). The selected thread (YES of 106) then performs the serial operation (108) on one of processing elements 74. After the serial operation is performed, the selected active thread is deactivated (110). The remaining active threads that are not selected may be looped back to the “Get One” instruction (NO of 106).

[0076] GPU 12 configured to execute the “GetOne” instruction for serial execution described above, may also be configured to use the “GetOne” instruction to split a single wave that exhibits divergence into multiple sub-waves, each of which may be uniform. Suppose an example where it is desirable to perform a (possibly divergent) subroutine call. That is, a situation where only a portion of the threads in the wave are to perform the subroutine. One conventional technique for executing a divergent subroutine involves executing each thread uniquely. This may be inefficient.

[0077] Instead, GPU 12 may be configured to execute the “GetOne” instruction to store a subroutine in a shared register that may then be executed by all threads having access to the shared register (i.e., processing elements 74 that have access to the shared register). The shared register may be part of data store 78, instruction store 76, or other memory available to processing elements 74.

[0078] Using the “GetOne” instruction and the shared register, GPU 12 may be configured to select a subroutine target (i.e., a location where subroutine instructions are stored) using the “GetOne” instruction, broadcast that target to all threads (i.e., provide an indication that a subroutine location is stored in a shared register), and execute the subroutine on all matching threads (i.e., execute the subroutine on processing elements 74 having access to the shared register). This process may then be repeated for other threads having access to different shared registers. Example pseudocode for this process is shown below.

LOOP: GetOne_and_branch CONTINUE1

```
shared_R0 = subroutine_target //store target in a 'shared' register
```

CONTINUE1: if (subroutine_target == shared_R0)

```
{      call(subroutine_target)
```

```
      goto CONTINUE2
```

```
}
```

```
      goto LOOP
```

CONTINUE2:

[0079] Initially, control unit 72 would proceed to the LOOP section of the code and execute the “GetOne” instruction as previously described. The “GetOne” instruction selects one active thread to perform an operation. In the example above, a selected active thread would perform an operation that stores the address of a location (subroutine_target) in a shared register (shared_R0). Any active threads that were not selected would perform the CONTINUE1 operation. In the CONTINUE1 operation, each thread (i.e., processing element 74) having access to shared register shared_R0 executes the code located at subroutine_target and exits the control block at CONTINUE 2. If any non-selected active threads do not have access to the shared register, the non-selected active threads returns to the loop operation, where another

sub-routine call may be stored in a different shared register. Again, as with the example of serial operation, any inactive threads are skipped and the minimal number of unique subroutine calls is made.

[0080] FIG. 5 is a flowchart illustrating an example method of the disclosure. The techniques of FIG. 5 may be implemented by GPU 12, including control unit 72 of GPU12. In one example of the disclosure, GPU 12 is configured to perform a method of performing divergent operations on a multi-threaded processor. GPU 12 may be configured to determine one or more active threads from a plurality of threads (152), select one active thread from the one or more active threads (154), and perform a divergent operation on the selected active thread (156).

[0081] In one example of the disclosure, the multi-threaded processor is a SIMD processor. In another example of the disclosure, the plurality of threads form a wave, wherein the wave is executed by the SIMD processor in parallel.

[0082] In another example of the disclosure, GPU 12 is configured to perform the divergent operation on the selected active thread by performing a serialized operation on the selected active thread. In another example of the disclosure, the divergent operation stores a sub-routine target in a shared register, and GPU 12 is further configured to execute code stored at the sub-routine target on the one or more active threads having access to the shared register.

[0083] The techniques described in this disclosure may be implemented, at least in part, in hardware, software, firmware or any combination thereof. For example, various aspects of the described techniques may be implemented within one or more processors, including one or more microprocessors, digital signal processors (DSPs), application specific integrated circuits (ASICs), field programmable gate arrays (FPGAs), or any other equivalent integrated or discrete logic circuitry, as well as any combinations of such components. The term “processor” or “processing circuitry” may generally refer to any of the foregoing logic circuitry, alone or in combination with other logic circuitry, or any other equivalent circuitry such as discrete hardware that performs processing.

[0084] Such hardware, software, and firmware may be implemented within the same device or within separate devices to support the various operations and functions described in this disclosure. In addition, any of the described units, modules or components may be implemented together or separately as discrete but interoperable logic devices. Depiction of different features as modules or units is intended to highlight different functional aspects and does not necessarily imply that such modules

or units must be realized by separate hardware or software components. Rather, functionality associated with one or more modules or units may be performed by separate hardware, firmware, and/or software components, or integrated within common or separate hardware or software components.

[0085] The techniques described in this disclosure may also be stored, embodied or encoded in a computer-readable medium, such as a computer-readable storage medium that stores instructions. Instructions embedded or encoded in a computer-readable medium may cause one or more processors to perform the techniques described herein, e.g., when the instructions are executed by the one or more processors. Computer readable storage media may include random access memory (RAM), read only memory (ROM), programmable read only memory (PROM), erasable programmable read only memory (EPROM), electronically erasable programmable read only memory (EEPROM), flash memory, a hard disk, a CD-ROM, a floppy disk, a cassette, magnetic media, optical media, or other computer readable storage media that is tangible.

[0086] Computer-readable media may include computer-readable storage media, which corresponds to a tangible storage medium, such as those listed above. Computer-readable media may also comprise communication media including any medium that facilitates transfer of a computer program from one place to another, e.g., according to a communication protocol. In this manner, the phrase “computer-readable media” generally may correspond to (1) tangible computer-readable storage media which is non-transitory, and (2) a non-tangible computer-readable communication medium such as a transitory signal or carrier wave.

[0087] Various aspects and examples have been described. However, modifications can be made to the structure or techniques of this disclosure without departing from the scope of the following claims.

WHAT IS CLAIMED IS:

1. A method of performing divergent operations on a multi-threaded processor, the method comprising:
 - determining one or more active threads from a plurality of threads;
 - selecting one active thread from the one or more active threads; and
 - performing a divergent operation on the selected active thread.
2. The method of claim 1, wherein the multi-threaded processor is a single instruction multiple data (SIMD) processor.
3. The method of claim 2, wherein the plurality of threads form a wave, wherein the wave is executed by the SIMD processor in parallel.
4. The method of claim 1, wherein performing the divergent operation on the selected active thread comprises performing a serialized operation on the selected active thread.
5. The method of claim 1, wherein performing the divergent operation comprises storing a sub-routine target in a shared register, the method further comprising:
 - executing code stored at the sub-routine target on the one or more active threads having access to the shared register.
6. An apparatus configured to perform divergent operations, the apparatus comprising:
 - a multi-threaded processor comprising:
 - a control unit configured to determine one or more active threads from a plurality of threads, and to select one active thread from the one or more active threads; and
 - a plurality of processing elements, wherein one processing element of the plurality of processing elements associated with the selected active thread is configured to perform a divergent operation.

7. The apparatus of claim 6, wherein the multi-threaded processor is a single instruction multiple data (SIMD) processor.
8. The apparatus of claim 7, wherein the plurality of threads form a wave, wherein the wave is executed by the SIMD processor in parallel.
9. The apparatus of claim 6, wherein the processing element associated with the selected active thread is configured to perform a serialized operation.
10. The apparatus of claim 6, wherein the divergent operation stores a sub-routine target in a shared register, and wherein the plurality of processing elements are further configured to execute code stored at the sub-routine target on any processing elements of the plurality of processing elements that are associated with the one or more active threads and that have access to the shared register.
11. An apparatus configured to perform divergent operations on a multi-threaded processor, the apparatus comprising:
 - means for determining one or more active threads from a plurality of threads;
 - means for selecting one active thread from the one or more active threads; and
 - means for performing a divergent operation on the selected active thread.
12. The apparatus of claim 11, wherein the multi-threaded processor is a single instruction multiple data (SIMD) processor.
13. The apparatus of claim 12, wherein the plurality of threads form a wave, wherein the wave is executed by the SIMD processor in parallel.
14. The apparatus of claim 11, wherein the means for performing the divergent operation on the selected active thread comprises means for performing a serialized operation on the selected active thread.

15. The apparatus of claim 11, wherein the divergent operation stores a sub-routine target in a shared register, the apparatus further comprising:

means for executing code stored at the sub-routine target on the one or more active threads having access to the shared register.

16. A multi-threaded processor configured to perform divergent operations, the multi-threaded processor comprising:

a control unit configured to determine one or more active threads from a plurality of threads, and to select one active thread from the one or more active threads; and


a plurality of processing elements, wherein one processing element of the plurality of processing elements associated with the selected active thread is configured to perform a divergent operation.


17. The multi-threaded processor of claim 16, wherein the multi-threaded processor is a single instruction multiple data (SIMD) processor.

18. The multi-threaded processor of claim 17, wherein the plurality of threads form a wave, wherein the wave is executed by the SIMD processor in parallel.

19. The multi-threaded processor of claim 16, wherein the processing element associated with the selected active thread is configured to perform a serialized operation.

20. The multi-threaded processor of claim 16, wherein the divergent operation stores a sub-routine target in a shared register, and wherein the plurality of processing elements are further configured to execute code stored at the sub-routine target on any processing elements of the plurality of processing elements that are associated with the one or more active threads and that have access to the shared register.

**FIG. 1**

**FIG. 2**




FIG. 3





FIG. 4

**FIG. 5**

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2015/025362

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F9/38
ADD.

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)
G06F

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

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

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2011/078690 A1 (FAHS BRIAN [US] ET AL) 31 March 2011 (2011-03-31) paragraph [0056] - paragraph [0057] paragraph [0062] - paragraph [0081]; figures 5A,5B ----- US 7 634 637 B1 (LINDHOLM JOHN ERIK [US] ET AL) 15 December 2009 (2009-12-15) column 20, line 52 - column 22, line 24 ----- US 6 947 047 B1 (MOY SIMON [US] ET AL) 20 September 2005 (2005-09-20) column 5, line 26 - column 6, line 43 column 15, line 19 - line 54 -----	1-20 1-20 1-20
A		



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier application or patent but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
22 June 2015	30/06/2015
Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer Moraiti, Marina

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2015/025362

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2011078690	A1	31-03-2011	NONE
US 7634637	B1	15-12-2009	NONE
US 6947047	B1	20-09-2005	NONE