



- (51) **International Patent Classification:**
G06F 9/54 (2006.01) G06F 9/32 (2006.01)
- (21) **International Application Number:**
PCT/US2020/012455
- (22) **International Filing Date:**
07 January 2020 (07.01.2020)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
16/275,917 14 February 2019 (14.02.2019) US
- (71) **Applicant:** **AXIS SEMICONDUCTOR, INC.** [US/US];
62 Sherwood Drive, Methuen, Massachusetts 01844 (US).
- (72) **Inventors:** **WANG, Xiaolin**; Axis Semiconductor, Inc.,
62 Sherwood Drive, Methuen, Massachusetts 01844 (US).
WU, Qian; Axis Semiconductor, Inc., 62 Sherwood Drive,
Methuen, Massachusetts 01844 (US).
- (74) **Agent:** **BURUM, Douglas P.** et al.; Maine Cernota &
Rardin, 547 Amherst St., 3rd Floor, Nashua, New Hamp-
shire 03063 (US).
- (81) **Designated States** (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,

AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, 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).

Declarations under Rule 4.17:

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

(54) **Title:** METHOD OF SYNCHRONIZING HOST AND COPROCESSOR OPERATIONS VIA FIFO COMMUNICATION

(57) **Abstract:** A method of synchronizing thread execution of a host and one or more coprocessors includes writing by the host of an event command and at least one coprocessor instruction to a FIFO and comparing of the event command with a current event register of the coprocessor until they match, whereupon the FIFO entries are popped and the instructions are forwarded to the coprocessor for execution. A plurality of entry groups can be written to the FIFO, each beginning with an event command. The instructions can direct the coprocessor to exchange data with shared memory and apply its thread to the received data. The processors and shared memory can be linked by a ring-type bus having a controller that performs the comparison, popping, and instruction forwarding. The coprocessor clears the current event register during thread execution, and then writes an event command to the register when processing is complete.




Fig. 2



Published:

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

METHOD OF SYNCHRONIZING HOST AND COPROCESSOR OPERATIONS VIA FIFO COMMUNICATION

RELATED APPLICATIONS

[0001] This application claims the benefit of U.S. Patent 16/275,917 filed on February 14, 2019, which is herein incorporated by reference in its entirety for all purposes.

FIELD OF THE INVENTION

[0001] The invention relates to multi-thread program execution, and more particularly, to methods for synchronizing the actions of a host and at least one coprocessor during multi-thread program execution.

BACKGROUND OF THE INVENTION

[0002] As the availability of multi-core and other multiple processor computing devices has continued to rise, and the cost thereof has continued to fall, multi-thread processing has become an increasingly effective approach to optimizing the speed and efficiency with which computing devices are able to process data, especially for applications such as image processing and voice encoding, wherein a specified algorithm is applied repeatedly to a large number of data frames or “batches” in a fixed timing order, with little or no interdependence between the processing of the separate data frames.

[0003] In many cases, the functions performed by a given application can be divided into “host” or “control” functions that have an unknown timing order and may have interactive inputs, and one or more “helper” functions that do not have any interactive inputs, and for which the input order is strictly defined. An example would be a computer game, where some functions have an unknown

timing order and may have interactive inputs, such as functions that are directly associated with user interaction, while other functions are directed to performing certain background tasks associated with the game, such as image processing or voice encoding, that accept and process data frames in a defined order without user interaction. Another example would be a voice-enabled application running on a smart appliance or a mobile device such as a cellular telephone.

[0004] In such cases, multi-thread processing can be implemented by assigning the function(s) that perform operations with an unknown timing, including those that support user interaction, to a “host” or “control” thread, while at least some of the helper functions are assigned to one or more “helper” threads. Often, when coprocessors are available, the control functions are executed on a host processor, while some or all of the helper threads are offloaded to one or more coprocessors.

[0005] While this multi-thread, host/coprocessor approach can be very powerful, successful implementation requires a robust and efficient method of synchronizing and coordinating the program executions performed by the host processor and the coprocessors.

[0006] One approach is to use hardware interrupts to synchronize the actions of the host processor and coprocessors. However, this approach is “costly” in terms of hardware utilization, and may be limited if the hardware platform does not provide sufficient interrupts with suitable functionality.

[0007] Another approach is to use a “message passing interface” (“MPI”) protocol implemented in shared memory. However, this approach is costly in terms of execution time.

[0008] What is needed, therefore, is an efficient method of synchronizing and coordinating the program executions of a host processor and one or more coprocessors without dependence on interrupts, and while avoiding the execution speed penalty of an MPI implementation.

SUMMARY OF THE INVENTION

[0009] The present invention is a method of synchronizing and coordinating the program executions of a host processor and one or more coprocessors without dependence on interrupts, and while avoiding the execution speed penalty of an MPI implementation. The invention achieves these goals by implementing at least one FIFO that can store “event” commands as well as coprocessor instructions such as data flow and synchronization instructions. According to the disclosed method, the host processor writes a series of entries to the FIFO that begins with an event command stored in the “head” of the FIFO. In embodiments, the entries can comprise a plurality of entry groups, each of which includes an event command followed by an associated group of coprocessor instructions.

[0010] In some embodiments that include a plurality of coprocessors, a FIFO is dedicated to each of the coprocessors. In addition, each of the coprocessors maintains a “current event” register where the coprocessor stores an event command when it is ready to receive new input.

[0011] According to the disclosed method, during program operation the head entries of the FIFOs and the current event registers of the associated coprocessors are compared with each other on a continuous or pseudo-continuous basis. When the event command stored in the current event register of a coprocessor is found to be equal to, or otherwise “agrees with” the event command stored as the head entry of its associated FIFO, the coprocessor instructions that immediately follow the event command in the FIFO are sequentially “popped” from the FIFO until the next event command is encountered, or until the last entry written to the FIFO has been popped. In embodiments, the FIFO entries are cleared as they are popped from the FIFO, and the remaining entries are advanced so that if a subsequent event command is encountered it is located in the head of the FIFO.

[0012] The popped coprocessor instructions are forwarded to the coprocessor for execution. Typically, execution by the coprocessor of these popped instructions causes the coprocessor to obtain input data of the type and in the

order that is known to be required by the helper thread of the coprocessor. In embodiments, the coprocessor instructions can be data transfer commands, for example between shared DRAM and the coprocessor, and/or synchronization commands between the host and coprocessor. After executing the popped coprocessor instructions that are received from the FIFO, the coprocessor typically proceeds to execute applicable commands that are included in the helper thread assigned to the coprocessor. For example, the popped instructions from the FIFO may direct the coprocessor to exchange data with shared DRAM, after which the newly received data is encoded or otherwise processed by the coprocessor, as directed by its helper thread.

[0013] While it is executing the popped instructions and helper thread commands, the coprocessor sequencer clears its current event register or writes some other registry entry to the current event register that does not correspond to or “agree with” any event command entries in the FIFO, so that additional coprocessor instructions are not popped from the FIFO until the coprocessor is ready for them. Once the coprocessor has executed all of the popped instructions received from the FIFO, and has performed any and all processing steps as directed by its helper thread, the coprocessor writes an event command to its current event register indicating that it is ready to receive further popped commands from the FIFO. In embodiments, the coprocessor also transmits a “done” message to the host processor, which allows the host processor to monitor the progress of the coprocessor in executing the instructions that it previously wrote to the FIFO. When an expected number of “done” commands is received by the host processor, it can, for example, exchange additional data between the DRAM and associated input and output channels, and then write a new string of entries to the FIFO.

[0014] In embodiments, the host processor and the one or more coprocessors are interconnected via a ring-type bus, such as the bus architecture disclosed in US patent 8,181,003 (incorporated herein by reference in its entirety for all purposes), for which the present inventors are co-inventors. In some of these embodiments,

the bus includes a bus controller, which can be implemented as a finite state machine (FSM) in hardware and/or software. In embodiments, the bus controller is responsible for comparing the contents of the FIFO header(s) and coprocessor current event register(s), for popping coprocessor instructions from the FIFO(s), and for transmitting the popped coprocessor instructions to the coprocessor(s) for execution.

[0015] Embodiments include additional FIFOs, wherein the FIFO described above can be referred to as the “scheduled” FIFO, and the additional FIFOs can include, for example, one or more host read FIFOs, host write FIFOs, and/or coprocessor write FIFOs. In various embodiments, the host read FIFO(s) can be used by the host to read registers of the coprocessor(s), the host write FIFO(s) can be used by the host to write to coprocessor registers, and/or the coprocessor write FIFO(s) can be used by the coprocessor(s) to write to the host without a “host read” command. In particular, in some embodiments the “done” messages can be sent by any of the coprocessors to the host via a coprocessor write FIFO.

[0016] The present invention is a method of synchronizing program executions of a host processor and of a coprocessor. The method includes the following steps:

- A) dividing an application into a host portion and a helper portion;
- B) implementing of the host portion on the host processor and implementing the helper portion on the coprocessor;
- C) writing to a FIFO by the host processor of FIFO entries that include an event command followed by at least one coprocessor instruction, whereby the event command occupies a head of the FIFO;
- D) comparing the event command in the FIFO head with a registry entry in a current event register of the coprocessor;
- E) repeating step D) until the registry entry in the current event register agrees with the event command in the FIFO head;
- F) popping the event command and the at least one coprocessor instruction from the FIFO and forwarding the popped at least one coprocessor

instruction to the coprocessor, said popping being continued until all of the FIFO entries have been popped, or until a FIFO entry is encountered that is a subsequent event command;

- G) writing by the coprocessor to the current event register of an entry that does not agree with any remaining FIFO entry event commands;
- H) executing by the coprocessor of the helper portion of the application;
- I) transmitting by the coprocessor to the host processor of a message indicating that execution of the helper portion has been completed; and
- J) writing by the coprocessor to the current event register of an event command registry entry.

[0017] In embodiments, the FIFO and the coprocessor are linked by a bus having a bus controller, and wherein steps D), E), and F) are executed by the bus controller. In some of these embodiments, the bus has a ring architecture.

[0018] In any of the above embodiments, the registry entry in the current event register can be deemed to agree with the event command in the FIFO head if it is identical with the event command in the FIFO head.

[0019] Any of the above embodiments can further include a coprocessor memory as part of the coprocessor and a shared memory that is accessible to the coprocessor, and wherein the at least one popped coprocessor instruction includes an instruction directing the coprocessor to read data from the shared memory into the coprocessor memory. In some of these embodiments, the at least one popped coprocessor instruction includes a command directing the coprocessor to exchange data with the shared memory. In any of these embodiments, the shared memory can be accessible to the host processor, and the host portion of the application can include commands directing the host processor to exchange data between the shared memory and at least one input and/or output channel.

[0020] In any of the above embodiments, the at least one popped coprocessor instruction can include at least one of a data flow command and a synchronization command.

[0021] In any of the above embodiments, in step C) the FIFO entries written by the host to the FIFO can include a plurality of entry groups, each entry group comprising an event command followed by at least one coprocessor instruction. In some of these embodiments, step F) further includes advancing the FIFO entries in the FIFO so that when a FIFO entry is encountered that is a subsequent event command, the subsequent event command occupies the head of the FIFO.

[0022] Any of the above embodiments can further include reading by the host processor of the event command in the FIFO head and writing by the host processor of the event command to the current event register of the coprocessor.

[0023] In any of the above embodiments, the helper portion can be free of any interactive inputs and can have an input order that is strictly defined.

[0024] In any of the above embodiments, the host portion can be subject to timing that is not predetermined. In some of these embodiments, the host portion is subject to interactive timing that is not predetermined,

[0025] In any of the above embodiments, the application can be a pipelined multi-batch execution program.

[0026] In any of the above embodiments, the message indicating that execution of the helper portion has been completed can be transmitted by the coprocessor to the host processor via a coprocessor write FIFO.

[0027] And any of the above embodiments can further include at least one of reading by the host of a register of the coprocessor via a host read FIFO, writing by the host to a register of the coprocessor via a host write FIFO, and writing by the coprocessor to the host via a coprocessor write FIFO.

[0028] The features and advantages described herein are not all-inclusive and, in particular, many additional features and advantages will be apparent to one of ordinary skill in the art in view of the drawings, specification, and claims.

Moreover, it should be noted that the language used in the specification has been

principally selected for readability and instructional purposes, and not to limit the scope of the inventive subject matter.

BRIEF DESCRIPTION OF THE DRAWINGS

[0029] Fig. 1 is a block diagram illustrating an embodiment of the present invention that includes a single coprocessor and FIFO;

[0030] Fig. 2 is a block diagram of an embodiment of the present invention that implements communication between the host, DRAM, and a coprocessor using a ring-bus and bus controller;

[0031] Fig. 3 is a flow diagram illustrating the operation and interactions between a host, bus controller, and coprocessor in an embodiment of the present invention;

[0032] Fig. 4 is a block diagram illustrating an embodiment of the present invention that includes a plurality of FIFOs;

[0033] Fig. 5 is a block diagram that illustrates the embodiment of Fig. 4; and

[0034] Fig. 6 is a block diagram similar to Fig. 5 but directed to an embodiment of the invention that includes a plurality of coprocessors and FIFOs;

DETAILED DESCRIPTION

[0035] With reference to Fig. 1, the present invention is a method of synchronizing and coordinating the program executions of a host processor 100 and one or more coprocessors 102 without dependence on interrupts, and while avoiding the execution speed penalty of an MPI implementation. The invention achieves these goals by implementing at least one FIFO 104 that can store “event” commands as well as coprocessor instructions such as data flow and synchronization instructions. According to the disclosed method, the host processor 100 writes a series of entries to the FIFO 104 that begins with an “event” command stored in the “head” 106 of the FIFO 104. In addition, each of

the coprocessors 102 maintains a “current event” register 108 where the coprocessor 102 stores an event command when it is ready to receive new input.

[0036] With reference to Fig. 2, in embodiments the host processor and all of the coprocessors are interconnected via a ring-type bus 200, such as the bus architecture disclosed in US patent 8,181,003 (incorporated herein by reference in its entirety for all purposes), for which the present inventors are co-inventors. In some of these embodiments, the bus 200 includes a bus controller 202, which can be implemented as a finite state machine (FSM) in hardware and/or software. In embodiments, the bus controller 202 is responsible for comparing the contents of the FIFO header(s) 106 and coprocessor current event register(s) 108, for popping coprocessor instructions from the FIFO(s) 104, and for transmitting the coprocessor instructions from the FIFO(s) 104 to the coprocessor(s) 102 for execution. In the illustrated embodiment, the FIFO 104 is implemented in a DRAM interface unit 204 that is accessible to both the host 100 and the coprocessor 102.

[0037] Fig. 3 is a flow diagram that illustrates an embodiment of the present method wherein the program executions of a host 100 and coprocessor 102 are synchronized by a FIFO 104 and bus controller 202 so as to process a plurality of data frames, as might be the case for a voice encoding application. According to the illustrated example, the host 100 receives N frames of digitized voice data 300 from an input channel and stores them in DRAM 204 that is accessible to both the host 100 and coprocessor 102. In addition, the host 100 outputs N frames of processed voice data 302 from the DRAM 204 to an output channel. Since there are now N frames of input data in DRAM that are ready to be encoded or otherwise processed, the host 100 then writes N sequential groups of entries 304 to the FIFO 104, where each entry group begins with an event command “Event = 1.” The entry groups are essentially identical, in that each entry group includes coprocessor instructions that direct the coprocessor to input a frame of input data from the DRAM 204 into internal memory of the coprocessor, encode or otherwise process the data, and then write the processed data back to the DRAM 204 as

output data. At this point the host 100 proceeds with other tasks, or simply waits 306 to receive $N \times$ “done” messages from the coprocessor 102, which will indicate that all of the N frames of input data have been processed by the coprocessor 102.

[0038] Note that in various embodiments the host 100 is able to write event commands to the current event register 108 of the coprocessor 102, for example to force the popping of entries when handling error conditions. Meanwhile the bus controller 202 continuously or pseudo continuously compares 308 the head of the FIFO 106 with the contents of the coprocessor current event register 108 to determine if they are equal to each other 310, or otherwise “agree” with each other, where agreement can be any defined relationship between the two entries. When the event command stored in the current event register 108 of the coprocessor 102 is found to be equal to, or otherwise agree with, the event command stored as the head entry 106 of its associated FIFO, the event command and coprocessor instructions that immediately follow the event command in the FIFO are “popped” 312 from the FIFO 104 until the next event command is encountered 314, or until the last entry written to the FIFO 104 has been popped. The popped coprocessor instructions are transmitted to the coprocessor for execution. In embodiments, the FIFO entries are cleared as they are popped from the FIFO 104, and the remaining FIFO entries are advanced, so that if a subsequent event command is encountered and the popping process is halted, the subsequent event command occupies the head 106 of the FIFO 104.

[0039] The presence of an event command such as “Event = 1” in the current event register 108 at the time it is matched 310 with the head entry 108 of the FIFO 104 indicates that the coprocessor 102 is ready and waiting for instructions 316 from the bus controller 202. Upon receiving the popped coprocessor instructions, the coprocessor 102 writes 318 “Event = 0” or some other entry to the current event register 108 that will not match any of the event command entries in the FIFO 100, so as to indicate that the coprocessor 102 is busy and is not prepared to receive any further coprocessor instructions from the FIFO 104.

[0040] In the illustrated example, the popped coprocessor instructions direct the coprocessor 102 to read an input data frame 320 from DRAM 204 to location $n+1$ of the coprocessor internal memory, and to output a processed data frame 322 from location $n-1$ to the DRAM 204. The coprocessor 102 also proceeds to encode or otherwise process the input data frame 322 at location $n+1$ according to code that is included in the helper thread that is assigned to the coprocessor 102. Once processing of the data frame is complete, the coprocessor 102 sends a “done” message 326 to the host 100, increments 328 the value of n , and sets the current event register 330 back to $\text{Event} = 1$, thereby indicating that it is ready to receive more popped instructions from the FIFO 104.

[0041] When an expected number of “done” messages is received 306 by the host processor 100, it can, for example, read 300 and write 302 additional data to and from the DRAM 204, and then write a new string of entries 304 to the FIFO 102.

[0042] With reference to Fig. 4, embodiments include additional FIFOs 402 - 404, wherein the FIFO 100 described above can be referred to as the “scheduled” FIFO 100, and the additional FIFOs can include, for example, one or more host read FIFOs 400, host write FIFOs 402, and/or coprocessor write FIFOs 404. In various embodiments, the host read FIFO(s) 400 can be used by the host 100 to read registers of the coprocessor(s) 102, the host write FIFO(s) 402 can be used by the host 100 to write to coprocessor registers, and/or the coprocessor write FIFO(s) 404 can be used by the coprocessor 102 to write to the host 100 without a “host read” command. In some embodiments, the “done” messages can be sent 326 by any of the coprocessors 102 to the host 100 via a coprocessor write FIFO 404.

[0043] Fig. 5 is a simplified block diagram that illustrates the method of Fig. 4 as implemented in a voice encoding application. In the illustrated example, the voice data frames are received 500 from an input channel at a rate of one frame every 20 ms, and are stored 502 in DRAM 204. Each frame is considered a

“batch,” whereby a pipeline multi-batch execution process is applied to a plurality of input frames received from a single input channel. As discussed above, the host 100 is responsible for receiving 300, 500 the input frames from an input channel and storing them 502 in DRAM 204, and for outputting data frames 302, 504 from DRAM 204 to an output channel after they have been encoded by the coprocessor 102.

[0044] Tables 1-3 illustrate the code (Table 1) that is executed by the host 100 and the code (Table 2) that is executed by the coprocessor 102, as well as the entries (Table 3) that are written to the FIFO 104 by the host 100 before they are popped by the bus controller 202 in the embodiment of Fig. 5.

[0045] Table 1: Host CPU Program

```
while (! Done) {
    receive input of F frames from
input channels
    send output of F frames to
output channels
    write F frames of input/output
pair each starting with event #1 to
scheduled FIFO as shown
    receive F frames of “end” from
coprocessor }
```

Note: $f = \text{loop time} / 20 \text{ ms}$

[0046] Table 2: Coprocessor Program

1. Set current Event = 1
2. Start encoding algorithm
3. Write to CPU to indicate frame processing “done”
4. Goto 1

[0047] Table 3: Scheduled FIFO Content

Entries from Head of Scheduled
FIFO -

001: Event = 1
002: Command for input of frame n+1 data transfer from dram Addr Xn+1 to coprocessor
003: Command for output of frame n-1 data transfer from coprocessor to DRAM Addr Yn-1
004: Event = 1
005: Command for input of frame n+2 data transfer from DRAM Addr Xn+2 to coprocessor
006: Comand for output of frame n data transfer from coprocessor to DRAM Addr Yn
.....

[0048] Note that the batch rate for the illustrated example is the loop time divided by the frame time. For example, if the loop time is 100 ms and the frames are received at a rate of 20 ms per frame, then for each loop, the host 100 will be able to receive and process $100/20 = 5$ frames of voice data, and each string of entries that is written by the host 100 to the FIFO 104 will include five batches of entries, each beginning with an event command and followed by data read/write instructions to be executed by the coprocessor 102.

[0049] Accordingly, in the illustrated example, the host 100 and coprocessor 102 are able to operate separately and independently from each other, including writing to and reading from the DRAM 204, because the events and associated coprocessor commands that are stored in the FIFO 104 serve to ensure that there are no read-after-write or other synchronization errors, or other problems. As such, no interrupts are required, and no added protocol is needed for synchronization and data transfer.

[0050] In the illustrated embodiment, the event commands are used essentially as “ready” flags. In other embodiments, a plurality of different types of event commands can be implemented, for example “Event = 1,” “Event = 2,” etc. This can be helpful if, for example, the coprocessor requires input from the FIFO at different stages of data processing, and/or for identifying and handling error conditions.

[0051] In a multi-core processor environment, the memory requirement, and thus the power consumption, of a pipelined multi-batch execution can be further reduced by distributing the “helper” data processing among a plurality of sub-functions or helper “threads,” each having a similar time of execution, where each of the sub-functions is assigned to a separate coprocessor core. According to this approach, each of the coprocessor cores is only required to contain the code for a single one of the sub-functions.

[0052] If, for example, there are n cores, then the total “helper” processing time is reduced by $1/n$ as compared to an implementation with only one core, such as in the previous example. Accordingly, the clock rate required for multi-core embodiments is much slower than for single core embodiments, and also the power consumption can be greatly reduced. Or, if the clock rate is maintained, then multi-core embodiments can perform encoding or other data processing for a larger number of channels at a given rate as compared to single-core embodiments.

[0053] It should be noted, however, that in multi-core embodiments data flow is sometimes required between the sub-functions, i.e. between the cores, as well as input and output data flow between shared DRAM and the cores, which causes the networking requirements for some multi-core embodiments to be much more complex than for single-core embodiments. The networking structure describe in US Patent 8,811,387 (incorporated herein by reference in its entirety for all purposes) can be a good solution in some of these embodiments.

[0054] Fig. 6 illustrates an embodiment of the present invention that includes a plurality of coprocessors 324a – k with a FIFO 104a – k dedicated to each of the coprocessors 324a – k, as can be applied for example to a partitioned voice encoder that implements pipelined multi-batch execution on a multi-core coprocessor. Tables 4, 5, and 6 illustrate the code (Table 4) that is executed in the embodiment of Fig. 6 by the host 100 and the code (Table 5) that is executed by

the coprocessor 102, as well as the entries (Table 6) that are written to the FIFOs 104a-k by the host 100 before they are popped by the bus controller 202.

[0055] Table 4: Host CPU Program

```

while (! Done) {
  receive input of F frames from input channels
  send output of F frames to output channels
  write F frames of input/output pair each starting with event #1 to
  scheduled FIFO as shown
  receive F frames of "end" from coprocessor}

```

Note: $f = \text{loop time} / 20 \text{ ms}$

[0056] Table 5: Coprocessor Programs

Program in core 1

1. Set current Event = 1
2. Start subvn k algorithm
3. If core == 1. wrote to CPU to indicate batch processing "done"
4. Goto 1

....

Program in core k

1. Set current Event = 1
2. Start subvn k algorithm
3. If core == 1. wrote to CPU to indicate batch processing "done"
4. Goto 1

[0057] Table 6: Scheduled FIFO Content

Entries from head of Scheduled FIFO of Coprocessor Core 1 -

- 001: Event = 1
- 002: Command for input of frame n-1 data transfer from Coprocessor Core k to DRAM Addr Y_{n-1}
- 003: Event = 1
- 004: Command for input of frame n data transfer from Coprocessor Core k to DRAM Addr Y_n
-

....

Entries from head of Scheduled FIFO of Coprocessor Core k -

- 001: Event = 1

002: Command for input of frame n-1 data transfer from Coprocessor Core k to DRAM Addr Y_{n-1}

003: Event = 1

004: Command for input of frame n data transfer from Coprocessor Core k to DRAM Addr Y_n

.....

[0058] Another example of an application that can be addressed by embodiments of the present invention is a deep learning inference application. These applications are similar to voice encoding applications, in that for voice and video inference the frame structure is similar to the frame structure for a voice encoding application that has a fixed rate of input. Often, the rate of input can be fixed for other inference programs as well. As such, the present invention can be implemented in at least three different ways using scheduled FIFOs 104 to provide the synchronization and data flow mechanism in a pipeline multi-batch execution:

- a single channel inference program using a single coprocessor;
- a multi-channel inference program using a plurality of coprocessors; or
- a single channel or multi-channel inference program for which the program is divided into sub-functions, and each of the sub-functions is assigned to a coprocessor core.

[0059] In embodiments, the current event register is accessible to the host. In some of these embodiments, the host is able, if necessary, to initiate execution by a coprocessor of its helper thread by reading the event command that is currently in the FIFO head and writing it to the current event register of the coprocessor.

[0060] The foregoing description of the embodiments of the invention has been presented for the purposes of illustration and description. Each and every page of this submission, and all contents thereon, however characterized, identified, or numbered, is considered a substantive part of this application for all purposes, irrespective of form or placement within the application. This specification is not intended to be exhaustive or to limit the invention to the precise form disclosed. Many modifications and variations are possible in light of this disclosure.

[0061] Although the present application is shown in a limited number of forms, the scope of the invention is not limited to just these forms, but is amenable to various changes and modifications without departing from the spirit thereof. The disclosure presented herein does not explicitly disclose all possible combinations of features that fall within the scope of the invention. The features disclosed herein for the various embodiments can generally be interchanged and combined into any combinations that are not self-contradictory without departing from the scope of the invention. In particular, the limitations presented in dependent claims below can be combined with their corresponding independent claims in any number and in any order without departing from the scope of this disclosure, unless the dependent claims are logically incompatible with each other.

CLAIMS

I claim:

- 1 1. A method of synchronizing program executions of a host processor and of a
2 coprocessor, the method comprising:
 - 3 A) dividing an application into a host portion and a helper portion;
 - 4 B) implementing of the host portion on the host processor and
5 implementing the helper portion on the coprocessor;
 - 6 C) writing to a FIFO by the host processor of FIFO entries that
7 include an event command followed by at least one coprocessor instruction,
8 whereby the event command occupies a head of the FIFO;
 - 9 D) comparing the event command in the FIFO head with a registry
10 entry in a current event register of the coprocessor;
 - 11 E) repeating step D) until the registry entry in the current event
12 register agrees with the event command in the FIFO head;
 - 13 F) popping the event command and the at least one coprocessor
14 instruction from the FIFO and forwarding the popped at least one
15 coprocessor instruction to the coprocessor, said popping being continued
16 until all of the FIFO entries have been popped, or until a FIFO entry is
17 encountered that is a subsequent event command;
 - 18 G) writing by the coprocessor to the current event register of an
19 entry that does not agree with any remaining FIFO entry event commands;
 - 20 H) executing by the coprocessor of the helper portion of the
21 application;
 - 22 I) transmitting by the coprocessor to the host processor of a message
23 indicating that execution of the helper portion has been completed; and
 - 24 J) writing by the coprocessor to the current event register of an event
25 command registry entry.
- 1 2. The method of claim 1, wherein the FIFO and the coprocessor are
2 linked by a bus having a bus controller, and wherein steps D), E), and F)
3 are executed by the bus controller.

- 1 3. The method of claim 2, wherein the bus has a ring architecture.
- 1 4. The method of any preceding claim, wherein the registry entry in the
2 current event register is deemed to agree with the event command in the
3 FIFO head if it is identical with the event command in the FIFO head.
- 1 5. The method of any preceding claim, further comprising a coprocessor
2 memory included in the coprocessor and a shared memory that is accessible
3 to the coprocessor, and wherein the at least one popped coprocessor
4 instruction includes an instruction directing the coprocessor to read data
5 from the shared memory into the coprocessor memory.
- 1 6. The method of claim 5, wherein the at least one popped coprocessor
2 instruction includes a command directing the coprocessor to exchange data
3 with the shared memory.
- 1 7. The method of claim 5, wherein the shared memory is accessible to
2 the host processor, and wherein the host portion of the application includes
3 commands directing the host processor to exchange data between the shared
4 memory and at least one input and/or output channel.
- 1 8. The method of any preceding claim, wherein the at least one popped
2 coprocessor instruction includes at least one of a data flow command and a
3 synchronization command.
- 1 9. The method of any preceding claim, wherein in step C) the FIFO
2 entries written by the host to the FIFO include a plurality of entry groups,
3 each entry group comprising an event command followed by at least one
4 coprocessor instruction.
- 1 10. The method of claim 9, wherein step F) further includes advancing
2 the FIFO entries in the FIFO so that when a FIFO entry is encountered that
3 is a subsequent event command, the subsequent event command occupies
4 the head of the FIFO.

1 11. The method of any preceding claim, further comprising reading by
2 the host processor of the event command in the FIFO head and writing by
3 the host processor of the event command to the current event register of the
4 coprocessor.

1 12. The method of any preceding claim, wherein the helper portion does
2 not have any interactive inputs and has an input order that is strictly defined.

1 13. The method of any preceding claim, wherein the host portion is subject to
2 timing that is not predetermined.

1 14. The method of claim 13, wherein the host portion is subject to interactive
2 timing that is not predetermined.

1 15. The method of any preceding claim, wherein the application is a pipelined
2 multi-batch execution program.

1 16. The method of any preceding claim, wherein the message indicating
2 that execution of the helper portion has been completed is transmitted by
3 the coprocessor to the host processor via a coprocessor write FIFO.

1 17. The method of any preceding claim, further comprising at least one
2 of:

3 reading by the host of a register of the coprocessor via a host read
4 FIFO;

5 writing by the host to a register of the coprocessor via a host write FIFO;

6 and

7 writing by the coprocessor to the host via a coprocessor write FIFO.




Fig. 1

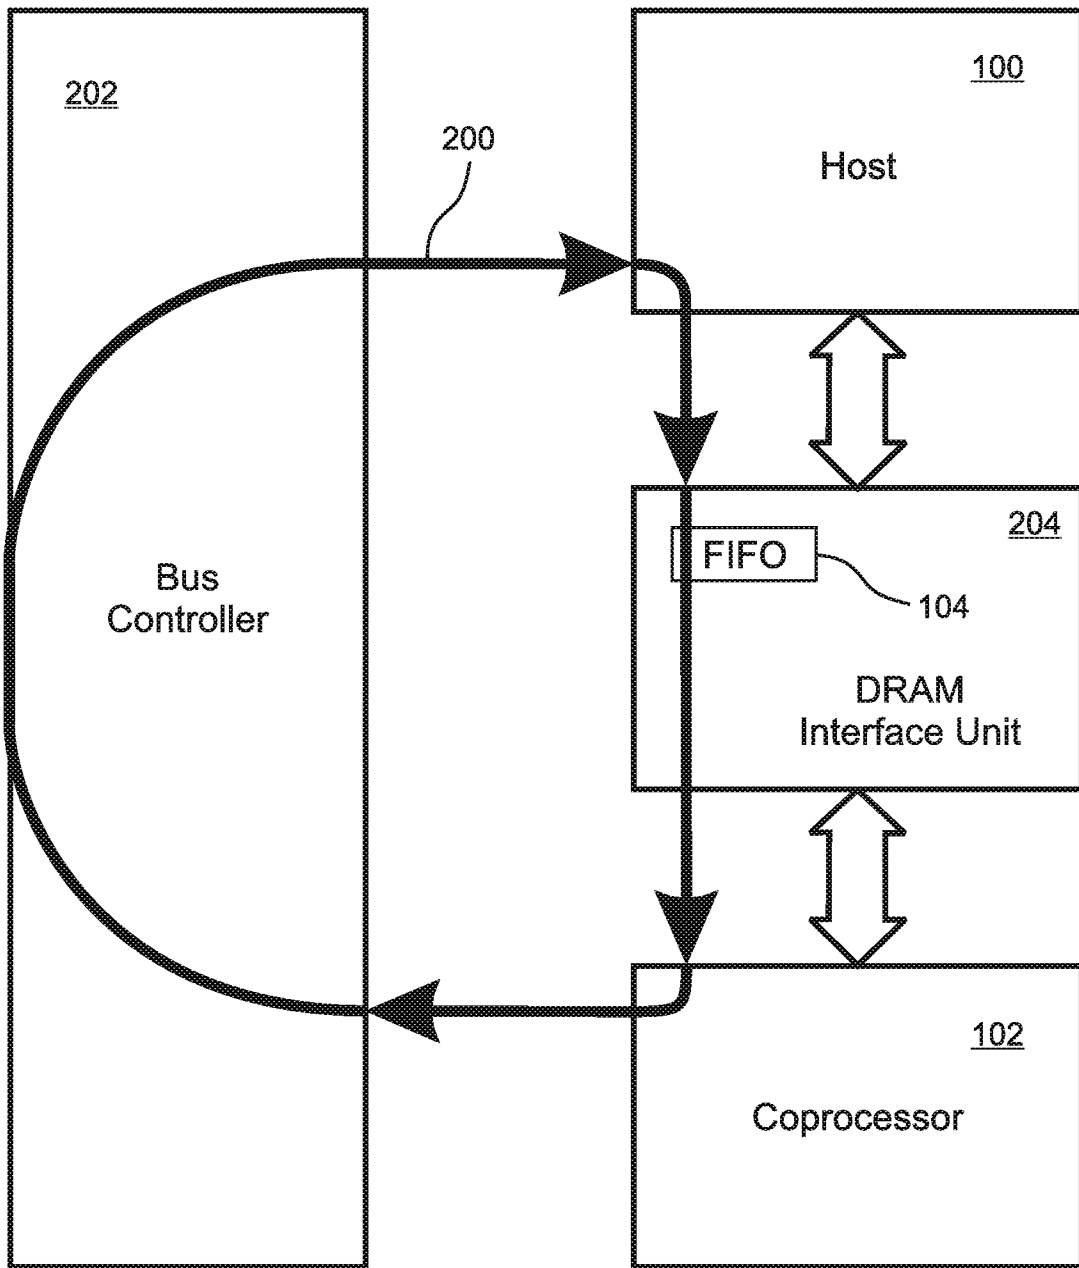


Fig. 2

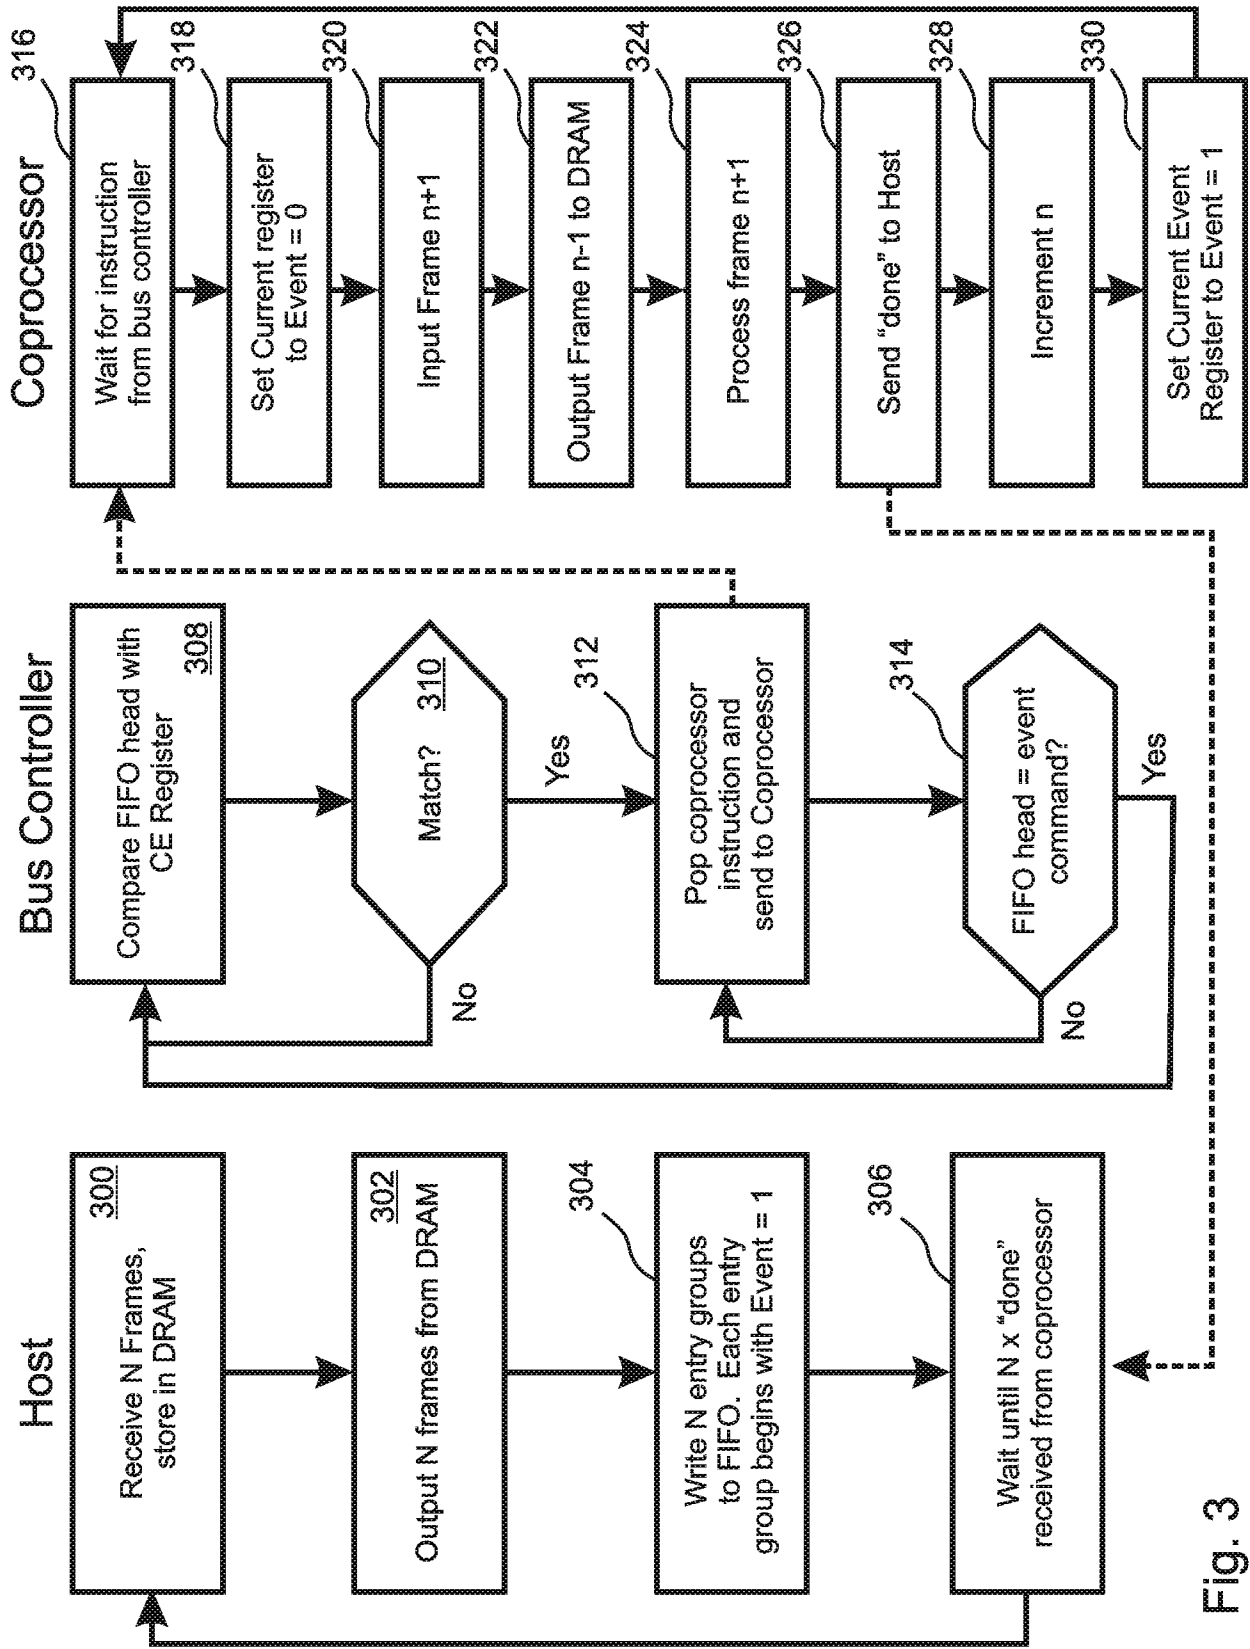


Fig. 3




Fig. 4

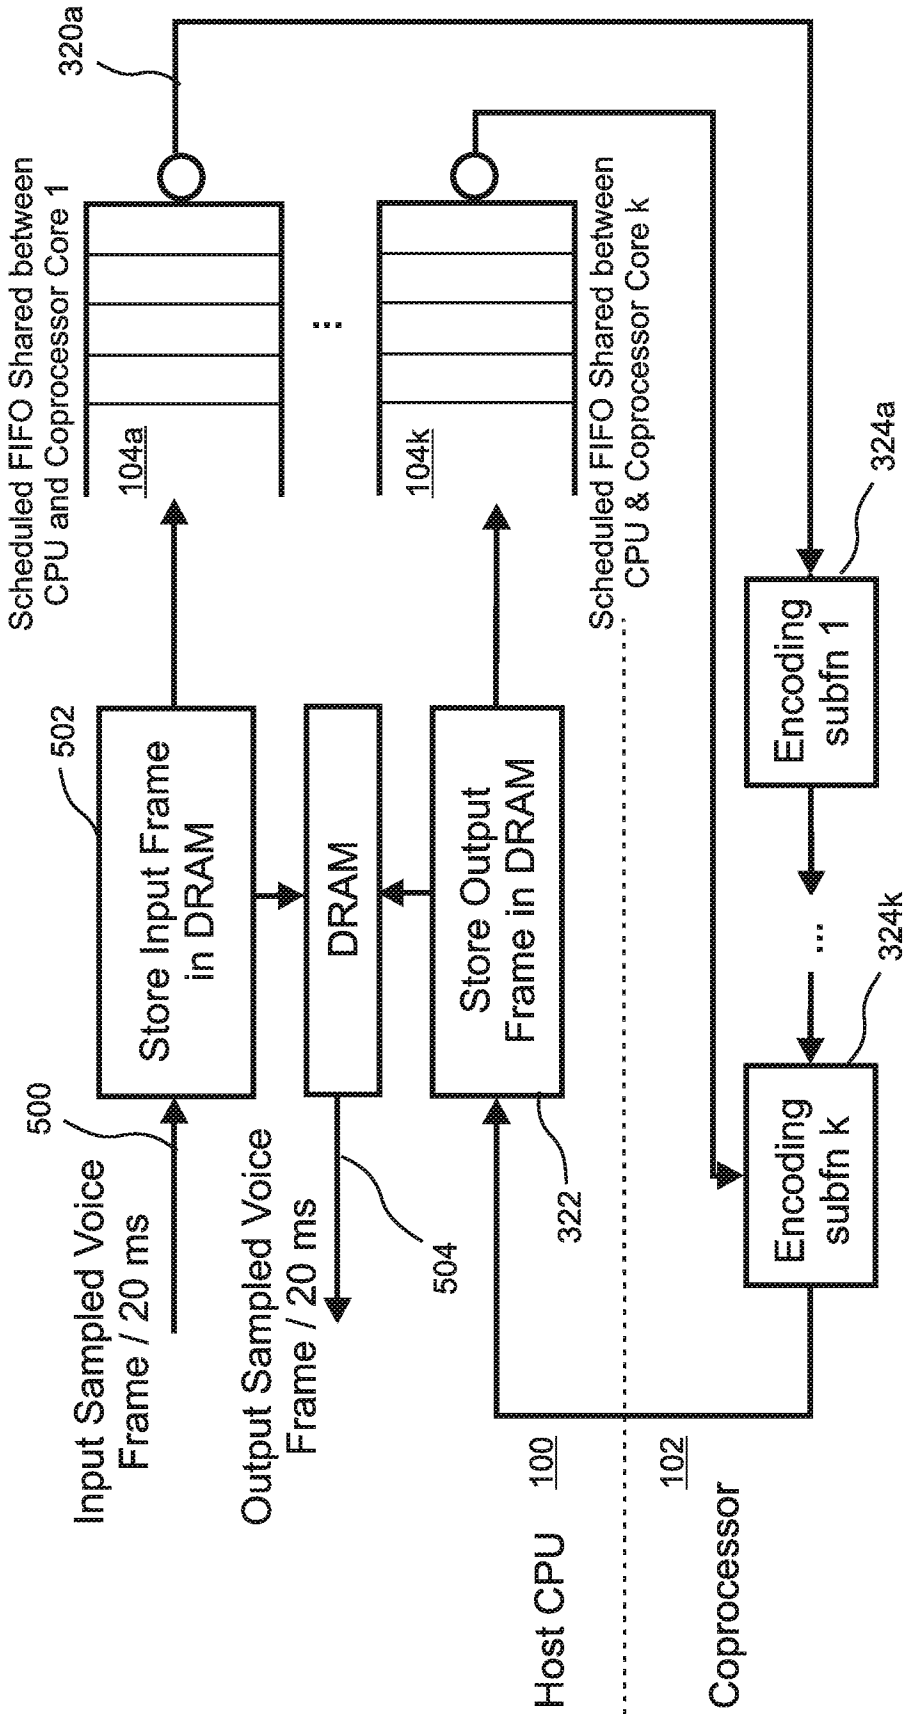


Fig. 6

A. CLASSIFICATION OF SUBJECT MATTER**G06F 9/54(2006.01)i, G06F 9/32(2006.01)i**

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F 9/54; G06F 009/46; G06F 012/00; G06F 013/38; G06F 12/08; G06F 15/76; G06F 9/38; G06F 9/46; G06F 9/48; G06F 9/32

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: multi-core, divide, FIFO(First-Input-First-Output), thread, synchronize

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2005-0223382 A1 (MARK DAVID LIPPETT) 06 October 2005 Paragraphs [0017], [0029], [0068], [0137]; and figure 8	1-17
A	US 2006-0004942 A1 (RICKY C. HETHERINGTON et al.) 05 January 2006 Paragraphs [0031], [0042], [0063]; and figures 4, 8	1-17
A	US 2012-0089812 A1 (GRAEME ROY SMITH) 12 April 2012 Paragraphs [0025], [0055], [0066]; and figure 1	1-17
A	US 2002-0166017 A1 (JASON SEUNG-MIN KIM et al.) 07 November 2002 Paragraphs [0024], [0035]; and figure 1	1-17
A	KR 10-2013-0093571 A (SYNOPSIS, INC. et al.) 22 August 2013 Paragraphs [0038], [0163]; and figure 2	1-17

 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

28 April 2020 (28.04.2020)

Date of mailing of the international search report

29 April 2020 (29.04.2020)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-481-8578

Authorized officer

KIM, Sung Hee

Telephone No. +82-42-481-3516



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2020/012455

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2005-0223382 A1	06/10/2005	US 10268609 B2	23/04/2019
		US 2014-0025857 A1	23/01/2014
		US 2014-0026141 A1	23/01/2014
		US 8533716 B2	10/09/2013
		US 9779042 B2	03/10/2017
US 2006-0004942 A1	05/01/2006	EP 1766529 A1	28/03/2007
		US 7240160 B1	03/07/2007
		US 7685354 B1	23/03/2010
		US 7873776 B2	18/01/2011
		WO 2006-004875 A1	12/01/2006
US 2012-0089812 A1	12/04/2012	EP 2441013 A1	18/04/2012
		US 9158575 B2	13/10/2015
		WO 2010-142987 A1	16/12/2010
US 2002-0166017 A1	07/11/2002	US 7996592 B2	09/08/2011
		WO 02-088970 A1	07/11/2002
KR 10-2013-0093571 A	22/08/2013	CN 1955931 A	02/05/2007
		EP 1770509 A2	04/04/2007
		EP 2328076 A1	01/06/2011
		EP 2328077 A1	01/06/2011
		JP 2007-133858 A	31/05/2007
		JP 2012-089154 A	10/05/2012
		JP 2013-239199 A	28/11/2013
		JP 5311732 B2	09/10/2013
		JP 5386572 B2	15/01/2014
		JP 5651214 B2	07/01/2015
		KR 10-1369352 B1	04/03/2014
		KR 10-1392934 B1	12/05/2014
		KR 10-2007-0037427 A	04/04/2007
		US 2007-0220294 A1	20/09/2007
		US 2007-0220517 A1	20/09/2007
		US 2014-0068619 A1	06/03/2014
		US 2014-0282593 A1	18/09/2014
		US 2014-0317378 A1	23/10/2014
		US 2015-0378776 A1	31/12/2015
		US 8533503 B2	10/09/2013
US 8732439 B2	20/05/2014		
US 8751773 B2	10/06/2014		
US 9164953 B2	20/10/2015		
US 9286262 B2	15/03/2016		
US 9442886 B2	13/09/2016		