

(12) STANDARD PATENT
(19) AUSTRALIAN PATENT OFFICE

(11) Application No. **AU 2020213829 B2**

(54) Title
Handling an input/output store instruction

(51) International Patent Classification(s)
G06F 9/30 (2018.01) **G06F 13/20** (2006.01)

(21) Application No: **2020213829** (22) Date of Filing: **2020.01.14**

(87) WIPO No: **WO20/156798**

(30) Priority Data

(31) Number	(32) Date	(33) Country
19154737.1	2019.01.31	EP

(43) Publication Date: **2020.08.06**

(44) Accepted Journal Date: **2022.09.15**

(71) Applicant(s)
International Business Machines Corporation

(72) Inventor(s)
RAISCH, Christoph;KRAEMER, Marco;LEHNERT, Frank;KLEIN, Matthias;BRADBURY, Jonathan;JACOBI, Christian;BELMAR, Brenton;DRIEVER, Peter

(74) Agent / Attorney
Spruson & Ferguson, GPO Box 3898, Sydney, NSW, 2001, AU

(56) Related Art
US 2011/0320764 A1
US 2015/0378737 A1
US 6286095 B1
US 4947316 A



- (51) International Patent Classification:
 G06F 9/30 (2018.01) G06F 13/20 (2006.01)
- (21) International Application Number:
 PCT/EP2020/050760
- (22) International Filing Date:
 14 January 2020 (14.01.2020)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
 19154737.1 31 January 2019 (31.01.2019) EP
- (71) Applicant: **INTERNATIONAL BUSINESS MACHINES CORPORATION** [US/US]; New Orchard Road, Armonk, New York 10504 (US).
- (71) Applicant (for MG only): **IBM DEUTSCHLAND GMBH** [DE/DE]; c/o IBM Deutschland Management & Business Support GmbH, Patentwesen und Urheberrecht, IBM Allee 1, 71139 Ehningen (DE).
- (72) Inventors: **RAISCH, Christoph**; c/o IBM Deutschland Research & Development GmbH, Schoenaicher Strasse 220, 71032 Boeblingen (DE). **KRAEMER, Marco**; c/o IBM Deutschland Research & Development GmbH, Schoenaicher Strasse 220, 71032 Boeblingen (DE). **LEHN-**

ERT, Frank; c/o IBM Deutschland Research & Development GmbH, Schoenaicher Strasse 220, 71032 Boeblingen (DE). **KLEIN, Matthias**; c/o IBM Corp., 2455 South Road, Poughkeepsie, New York 12601 (US). **BRADBURY, Jonathan**; c/o IBM Corp., 2455 South Road, Poughkeepsie, New York 12601 (US). **JACOBI, Christian**; c/o IBM Corp., 2455 South Road, Poughkeepsie, New York 12601 (US). **BELMAR, Brenton**; c/o IBM Corp., 2455 South Road, Poughkeepsie, New York 12601 (US). **DRIEVER, Peter**; c/o IBM Corp., 2455 South Road, Poughkeepsie, New York 12601 (US).

(74) Agent: **DOEHLER, Denis**; c/o IBM Deutschland Management & Business Support GmbH, Patentwesen und Urheberrecht, 71139 Ehningen (DE).

(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,

(54) Title: HANDLING AN INPUT/OUTPUT STORE INSTRUCTION




Fig. 1

(57) Abstract: A data processing system (210) and a method for handling an input/output store instruction (30), comprising a system nest (18) communicatively coupled to at least one input/output bus (22) by an input/output bus controller (20). The data processing system (210) further comprises at least a data processing unit (216) comprising a core (12), a system firmware (10) and an asynchronous core-nest interface (14). The data processing unit (216) is communicatively coupled to the system nest (18) via an aggregation buffer (16). The system nest (18) is configured to asynchronously load from and/or store data to at least one external device (214) which is communicatively coupled to the input/output bus (22). The data processing unit (216) is configured to complete the input/output store instruction (30) before an execution of the input/output store instruction (30) in the system nest (18) is completed. The asynchronous core-nest interface (14) comprises an input/output status array (44) with multiple input/output status buffers (24).

WO 2020/156798 A1

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

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

Published:

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

HANDLING AN INPUT/OUTPUT STORE INSTRUCTION

[0001] The present invention relates in general to data processing systems, in particular to a method for handling an input/output store instruction to multiple external devices as well as a computer program product and a data processing system.

BACKGROUND

10 [0002] A computing environment may include one or more types of input/output devices, including various types of adapters. One type of adapter is a Peripheral Component Interconnect (PCI) or Peripheral Component Interconnect Express (PCIe) adapter. This adapter includes one or more address spaces used in communicating data between the adapter and the system to which the adapter is attached.

15

[0003] In some systems, a portion of an address space of the central processing unit (CPU) coupled to the adapter is mapped to an address space of the adapter enabling CPU instructions that access storage to directly manipulate the data in the adapter's address space.

20 [0004] Communication with adapters, such as PCI or PCIe adapters can be facilitated by control instructions specifically designed for communicating data to and from adapters and used for communication.

[0005] In the state of the art, a store instruction for storing data in an adapter includes, for instance, obtaining a machine instruction for execution, the machine instruction being defined for computer execution according to a computer architecture, the machine instruction including, for instance, an opcode field identifying a store to adapter instruction. A first field identifies a first location that includes data to be stored in an adapter. A second field identifies a second location, the contents of which include a function handle identifying the adapter, a designation of an address space within the adapter in which data is to be stored, and an offset within the address space. The machine instruction is executed, the executing including using the function handle to obtain a function table entry associated with the adapter. A data address of the adapter is obtained using at least one of information

30

in the function table entry and the offset. Data are stored from the first location in a specific location in the address space identified by the designation of the address space, the specific location identified by the data address of the adapter.

5 [0006] An existing feature in a large multi-processor system is the ability to quiesce all processors within a target zone. Quiesce functions operate to temporarily pause or alter the state of a processor or group of processors to perform, e.g., system updates or backups. In some instances, a quiesce interruption is applicable to only a subset of the system resources. In such instances, the system can be divided into different zones. For a quiesce operation
 10 applicable to one zone (a target zone), processors outside of the target zone are permitted to continue running, although new translations may be blocked. Typically, at least one system controller or other mechanism broadcasts the quiesce to all physical processors in the system, handles collecting quiesce state information and indicates to a requesting processor when all processors have started, or are ignoring (filtering) the quiesce request.

15 [0007] A quiesce controller can be communicatively coupled to a processor in a multi-processor system, and a quiesce state machine configured to receive a quiesce request. The computer system is configured to perform a method that includes receiving a quiesce request at the quiesce controller from a requesting processor, the requesting processor being
 20 one of a plurality of processors in a multi-processor system, and determining that the quiesce request is not accepted based on a state of the quiesce state machine. The method also includes, based on the request being not accepted, generating a reject message configured to indicate that the quiesce request has been rejected, holding the reject message
 25 until a quiesce command is broadcast to the multi-processor system, the quiesce command based on a different quiesce request, and sending the reject message to the requesting processor based on the broadcast of the quiesce command being detected by the quiesce controller.

SUMMARY

30 [0007A] According to a first aspect of the present invention, there is provided a data processing system for handling an input/output store instruction, the data processing system comprising:

a data processing unit including a core and system firmware, the data processing unit configured to perform a method, the method comprising:

identifying an input/output function by an address specified using the input/output store instruction, the input/output store instruction specifying at least the input/output function with an offset through the address, at least one of data to be transferred and a pointer to data to be transferred, and a length of the data;

5 verifying whether access to the input/output function is allowed on an address space and on a guest instance level;

 completing the input/output store instruction before an execution of the input/output store instruction in a selected component of the data processing system different from the data processing unit is completed, the selected component configured to asynchronously load from and store data to at least one external device, the selected component being a system nest communicatively coupled to the data processing unit via an aggregation buffer, and communicatively coupled to asynchronously load data from and store data to the at least one external device via an input/output bus controller, the completing including:

15 issuing a system message to the aggregation buffer to forward the data to be transferred via the system nest as a single nest message to the input/output bus controller; and

 receiving a free for reuse message after the nest message has been injected into the system nest to enable an early continuation of a transfer of data before a completion message is issued by the system nest;

20 providing notification through an interrupt, based on detecting an error during an asynchronous execution of the input/output store instruction in the data processing unit;

 obtaining a barrier instruction which when executed assists in ordering of asynchronous store instructions; and

25 delaying, using the barrier instruction, processing of at least the input/output store instruction, until previous input/output store instructions have been completed, the barrier instruction completing execution after the previous input/output store instructions have been completed.

30 [0007B] According to a second aspect of the present invention, there is provided a computer program product for handling an input/output store instruction, the computer program product comprising:

 at least one computer readable storage medium readable by at least one

processing circuit and storing instructions for performing a method comprising:

- identifying an input/output function by an address specified using the input/output store instruction, the input/output store instruction specifying at least the input/output function with an offset through the address, at least one of data to be transferred and a pointer to data to be transferred, and a length of the data;
- verifying whether access to the input/output function is allowed on an address space and on a guest instance level;
- completing, by a data processing unit of a data processing system, the input/output store instruction before an execution of the input/output store instruction in a selected component of the data processing system different from the data processing unit is completed, the selected component configured to asynchronously load from and store data to at least one external device, the data processing unit including a core and system firmware, the selected component being a system nest communicatively coupled to the data processing unit via an aggregation buffer, and communicatively coupled to asynchronously load data from and store data to the at least one external device via an input/output bus controller, the completing including:
 - issuing a system message to the aggregation buffer to forward the data to be transferred via the system nest as a single nest message to the input/output bus controller; and
 - receiving a free for reuse message after the nest message has been injected into the system nest to enable an early continuation of a transfer of data before a completion message is issued by the system nest;
 - providing notification through an interrupt, based on detecting an error during an asynchronous execution of the input/output store instruction in the data processing unit;
 - obtaining a barrier instruction which when executed assists in ordering of asynchronous store instructions; and
 - delaying, using the barrier instruction, processing of at least the input/output store instruction, until previous input/output store instructions have been completed, the barrier instruction completing execution after the previous input/output store instructions have been completed.

[0007C] According to a third aspect of the present invention, there is provided a computer-implemented method of handling an input/output store instruction, the computer-

implemented method comprising:

identifying an input/output function by an address specified using the input/output store instruction, the input/output store instruction specifying at least the input/output function with an offset through the address, at least one of data to be transferred and a pointer to data to be transferred, and a length of the data;

verifying whether access to the input/output function is allowed on an address space and on a guest instance level;

completing, by a data processing unit of a data processing system, the input/output store instruction before an execution of the input/output store instruction in a selected component of the data processing system different from the data processing unit is completed, the selected component configured to asynchronously load from and store data to at least one external device, and the data processing unit including a core and system firmware, the selected component being a system nest communicatively coupled to the data processing unit via an aggregation buffer, and communicatively coupled to asynchronously load data from and store data to the at least one external device via an input/output bus controller, the completing including:

issuing a system message to the aggregation buffer to forward the data to be transferred via the system nest as a single nest message to the input/output bus controller; and

receiving a free for reuse message after the nest message has been injected into the system nest to enable an early continuation of a transfer of data before a completion message is issued by the system nest;

providing notification through an interrupt, based on detecting an error during an asynchronous execution of the input/output store instruction in the data processing unit;

obtaining a barrier instruction which when executed assists in ordering of asynchronous store instructions; and

delaying, using the barrier instruction, processing of at least the input/output store instruction, until previous input/output store instructions have been completed, the barrier instruction completing execution after the previous input/output store instructions have been completed.

[0008] A data processing system is proposed for handling an input/output store instruction, comprising a system nest communicatively coupled to at least one input/output bus by an input/output bus controller. The data processing system further comprises at least

a data processing unit comprising a core,
a system firmware and an asynchronous core-nest interface. The data processing unit is
communicatively coupled to the system nest via an aggregation buffer. The system nest is
configured to asynchronously load from and/or store data to at least one external device
5 which is communicatively coupled to the input/output bus. The asynchronous core-nest
interface comprises an input/output status array with multiple input/output status buffers, as
well as an array management and access logic.

[0009] The data processing system is configured to perform: (i) an operating system
10 running on the data processing system issues the input/output store instruction specifying at
least an input/output function with an offset through an address, data to be transferred and/or
a pointer to data to be transferred, and a length of the data; (ii) the data processing unit is
configured to identify the input/output function by the address specified in the input/output
store instruction; (iii) the data processing unit is configured to verify if access to the
15 input/output function is allowed on an address space and on a guest instance level, the guest
running on the data processing system; (iv) the data processing unit is configured to
complete the input/output store instruction before an execution of the input/output store
instruction in the system nest is completed; (v) the system firmware is configured to notify
the operating system through an interrupt, if during the asynchronous execution of the
20 input/output store instruction an error is detected by the data processing unit, transmitting
the data of the failed asynchronous execution; (vi) a barrier instruction delaying a
processing of a store instruction, until all previous store instructions have been completed.

[0010] Favorably a barrier instruction may be introduced for ordering of asynchronous
25 store instructions. Advantageously, the barrier instruction may be a PCI SMT thread barrier
instruction, which completes when all previous PCI store instructions have been sent to a
PCIe bus for that SMT thread.

[0011] The data processing system according to first embodiment of the invention
30 comprises instructions loading from and storing to at least one external device of the data
processing system via an input/output bus. Asynchronous instructions complete before data
has been stored to the external device while synchronous instructions complete after data
has been stored to the external device. Within the embodiments described here, PCI will be
used interchangeably for any other input/output technology, thus not restricting the

embodiment of the invention to PCI.

[0012] Embodiments of the invention describe an input/output store instruction execution in a strictly ordered way as observable from above the architecture boundary while the actual execution may be out of order within the hardware of the data processing unit (CPU).

[0013] According to embodiments of the invention a PCI store instruction may be executed with an asynchronous execution of the PCIe store effect and an asynchronous status handling. Asynchronous reliable execution is based on reliable forwarding mechanisms in microarchitecture of the inventive data processing system.

[0014] An existing PCI store and store block instruction is usually synchronous up to the point where the PCI store data has been delivered to the PCIe interface and completion returned to a processing unit.

[0015] PCI standard only requires an asynchronous send command of PCI information, which is typically implemented through a store queue in the processor aggregating data with asynchronous send-out.

[0016] Advantageously, according to embodiments of the invention, an improvement concerning cycles per instruction may be achieved by replacing a synchronous PCI instruction by a reliable asynchronous send process of an input/output store instruction.

[0017] Alternatively or additionally of data to be transferred, the store instruction according to an embodiment of the invention may also specify a pointer to a main memory which should be used to fetch data from, instead of containing the data directly.

[0018] Guest instance level may also mean that a single guest or host may be running on the data processing system.

[0019] The address of the offset of the input/output function itself can be a virtual, physical, logical address. Virtual and logical addresses typically get translated through a memory management unit (MMU) into a physical address, and the physical address then

allows to identify which function and offset is meant.

[0020] Physical address in this context means "lowest address in the address translation hierarchy accessible from within a guest/operating system".

5

[0021] Advantageously, the input/output status buffers may collect returned states from the system nest and/or from the input/output bus controller, in particular a completion message from the system nest. These input/output status buffers may collect the returned states acting as an asynchronous system message buffer supporting the asynchronous transmit process. Advantageously the input/output status buffers may be integrated directly in the asynchronous core-nest interface for quick response.

10

[0022] According to a favourable embodiment of the inventive data processing system, the barrier instruction may restrict a serial processing of store instructions to a single input/output bus controller. Thus an efficient ordering of the store instructions aiming at this input/output bus controller is possible in order to guarantee that a message will reach an specific external device.

15

[0023] According to a favourable embodiment of the inventive data processing system, the barrier instruction may restrict a serial processing of store instructions to a single input/output function, in particular a single external input/output device. Thus an efficient ordering of the store instructions aiming at this input/output bus controller is possible in order to guarantee that a message will reach an specific external device.

20

[0024] According to a favourable embodiment of the inventive data processing system, a summary query capability for all input/output status buffers of the input/output status array may be provided by an array management and access logic, returning aggregated information of a status of the input/output status buffers. By this way information may gathered in an efficient manner from the input/output status buffers.

25

30

[0025] According to a favourable embodiment of the inventive data processing system, the system firmware may comprise an array management logic, which allocates/deallocates input/output status buffers in the input/output status array and/or initiates a start of a new store instruction. Thus idle status buffers may be attributed to further store instructions. An

ordered processing of store instructions may be handled in an efficient and time saving way.

[0026] According to a favourable embodiment of the inventive data processing system, the aggregation buffer may be communicatively coupled to the asynchronous core-nest interface via an asynchronous bus. Thus the aggregation buffer can handle data directly sent by the asynchronous core-nest interface consecutively until all data to be transferred to the external device are stored in the aggregation buffer. By this way the asynchronous transmit mechanism for data transfer from the asynchronous core-nest interface may be favourably supported.

10

[0027] According to a favourable embodiment of the inventive data processing system, the data may be transferred by the input/output store instruction through an asynchronous transmit mechanism with an early completion message in multiple data packets to the aggregation buffer, if the length of the source data exceeds eight bytes, else the data may be transferred in one data packet. The asynchronous transmit mechanism is favourable because the sending device is free for reuse at an earlier state.

15

[0028] According to a favourable embodiment of the inventive data processing system, the system firmware may comprise an asynchronous input/output driver code for handling the input/output store instruction. Thus an asynchronous transmit mechanism may be used for transferring data from the data processing unit to the external device. Further, advantageously, the core may comprise an asynchronous setup code for handling memory requirements for status information of the asynchronous input/output driver code. This asynchronous setup code may further facilitate the asynchronous transmit mechanism through the aggregation buffer to the system nest and the input/output bus controller.

20

25

[0029] According to a favourable embodiment of the inventive data processing system, the asynchronous core-nest interface may comprise an asynchronous core-nest interface forwarding component for forwarding the data with local completion. This component may be implemented in hardware in the asynchronous core-nest interface. Thus a favourable asynchronous transmit mode for sending the data in data packets to the aggregation buffer may be supported.

30

[0030] According to a favourable embodiment of the inventive data processing system,

the aggregation buffer may comprise an early completion logic for delivering a free for reuse message after sending a request. This enables an early continuation of the transmit process of the data via the aggregation buffer to the system nest and the input/output bus controller.

5

[0031] Advantageously, the system firmware may comprise an array management logic, which allocates/deallocates input/output status buffers in the input/output status array and/or initiates a start of a new store instruction. Thus idle status buffers may be attributed to further store instructions. An ordered processing of store instructions may be handled in an efficient and time saving way.

10

[0032] According to a favourable embodiment of the inventive data processing system, a system message may comprise one of - a hierarchical physical target address, - sourcing an SMT (simultaneous multithreading) thread or an aggregate buffer identifier, - a length of data, an input/output bus address, or - an input/output status buffer index. Thus an advantageous passing of relevant information through the data processing system can be guaranteed.

15

[0033] Further a method is proposed for handling an input/output store instruction to at least one external device of a data processing system, the data processing system comprising a system nest communicatively coupled to at least one input/output bus by an input/output bus controller. The data processing system further comprises at least a data processing unit comprising a core, a system firmware and an asynchronous core-nest interface. The data processing unit is communicatively coupled to the system nest via an aggregation buffer. The external device is communicatively coupled to the input/output bus. The asynchronous core-nest interface comprises an input/output status array with multiple input/output status buffers, as well as an array management and access logic.

20

25

[0034] The method comprises: (i) an operating system running on the data processing system issuing the input/output store instruction specifying at least an input/output function with an offset through an address, data to be transferred and/or a pointer to data to be transferred, and a length of the data; (ii) the data processing unit being configured to identify the input/output function by the address specified in the input/output store instruction; (iii) the data processing unit being configured to verify if access to the

30

input/output function is allowed on an address space and on a guest instance level, the guest running on the data processing system; (iv) the data processing unit being configured to complete the input/output store instruction before an execution of the input/output store instruction in the system nest is completed; (v) the system firmware being configured to
5 notify the operating system through an interrupt, if during the asynchronous execution of the input/output store instruction an error is detected by the data processing unit, transmitting the data of the failed asynchronous execution; (vi) a barrier instruction delaying a processing of a store instruction, until all previous store instructions have been completed.

10 [0035] Favorably a barrier instruction may be introduced for ordering of asynchronous store instructions. Advantageously, the barrier instruction may be a PCI SMT thread barrier instruction, which completes when all previous PCI store instructions have been sent to a PCIe bus for that SMT thread.

15 [0036] The method according to a further embodiment of the invention comprises instructions loading from and storing to at least one external device of the data processing system via an input/output bus. Asynchronous instructions complete before data has been stored to the external device while synchronous instructions complete after data has been stored to the external device. Within the embodiments described here, PCI will be used
20 interchangeably for any other input/output technology, thus not restricting the embodiment of the invention to PCI.

[0037] Embodiments of the inventive method describe an input/output store instruction execution in a strictly ordered way as observable from above the architecture boundary
25 while the actual execution may be out of order within the hardware of the data processing unit (CPU).

[0038] According to embodiments of the inventive method a PCI store instruction may be executed with an asynchronous execution of the PCIe store effect and an asynchronous
30 status handling. Asynchronous reliable execution is based on reliable forwarding mechanisms in microarchitecture of the inventive data processing system.

[0039] An existing PCI store and store block instruction is usually synchronous up to the point where the PCI store data has been delivered to the PCIe interface and completion

returned to a processing unit.

[0040] PCI standard only requires an asynchronous send command of PCI information, which is typically implemented through a store queue in the processor aggregating data with asynchronous send-out.

[0041] Advantageously, according to embodiments of the inventive method, an improvement concerning cycles per instruction may be achieved by replacing a synchronous PCI instruction by a reliable asynchronous send process of an input/output store instruction.

[0042] Alternatively or additionally of data to be transferred, the store instruction according to an embodiment of the invention may also specify a pointer to a main memory which should be used to fetch data from, instead of containing the data directly.

[0043] Guest instance level may also mean that a single guest or host may be running on the data processing system.

[0044] The address of the offset of the input/output function itself can be virtual, physical, logical address. Virtual and logical addresses typically get translated through a memory management unit (MMU) into a physical address, and the physical address then allows to identify which function and offset is meant.

[0045] Physical address in this context means "lowest address in the address translation hierarchy accessible from within a guest/operating system".

[0046] According to a favourable embodiment of the inventive method, the barrier instruction may restrict a serial processing of store instructions to a single input/output bus controller. Thus an efficient ordering of the store instructions aiming at this input/output bus controller is possible in order to guarantee that a message will reach an specific external device.

[0047] According to a favourable embodiment of the inventive method, the barrier instruction may restrict a serial processing of store instructions to a single input/output function, in particular a single external input/output device. Thus an efficient ordering of the

store instructions aiming at this input/output bus controller is possible in order to guarantee that a message will reach an specific external device.

[0048] According to a favourable embodiment of the inventive method, a summary query capability for all input/output status buffers of the input/output status array may be provided by an array management and access logic, returning aggregated information of a status of the input/output status buffers. By this way information may gathered in an efficient manner from the input/output status buffers.

[0049] According to a favourable embodiment of the inventive method, the system firmware may comprise an array management logic, which allocates/deallocates input/output status buffers in the input/output status array and/or initiates a start of a new store instruction. Thus idle status buffers may be attributed to further store instructions. An ordered processing of store instructions may be handled in an efficient and time saving way.

[0050] Advantageously, the input/output status buffers may collect message states from the system nest and/or from the input/output bus controller, in particular a completion status from the system nest, wherein the message states and/or the completion status are numbered by an input/output status buffer index. By this way information about the completion status of different store instructions may be handled in an ordered and efficient manner. The numbering enables the possibility of handling messages, and particularly completion states in an ordered and efficient way for further processing other store instructions.

[0051] According to a favourable embodiment, the method may further comprise: (i) the operating system issuing the input/output store instruction; (ii) the system firmware allocating a free input/output status buffer index; if there is no free input/output status buffer index available, then waiting for a free input/output status buffer index ; (iii) the system firmware injecting the store instruction into the asynchronous send engine; if blocked by another store instruction waiting until the store instruction has been completed; (iv) depending on the length of the data: if a length of the data exceeds eight bytes, the system firmware issuing repeatedly a system message to send a data packet to the aggregation buffer until all data of a store block have been forwarded to the aggregation buffer, while the system firmware waiting until the data have been sent by the system message; else the system firmware issuing a system message to send the data to the aggregation buffer; further

independent of the length of the data, (v) the system firmware issuing a system message to the aggregation buffer to forward the data asynchronously as single nest message to the input/output bus controller, while waiting for the aggregation buffer to send a completion message; (vi) the aggregation buffer injecting the nest message into the system nest, wherein
5 the aggregation buffer is free for reuse right after the send operation, signaling back to the system firmware; then the aggregation buffer sending a free for reuse message; (vii) the system nest forwarding the message to the target location; (viii) the input/output bus controller receiving the message and forwarding data in a data frame to the input/output bus; (ix) the input/output bus controller sending a completion message to the system nest; (x) the
10 system nest forwarding the completion message to the originating aggregation buffer; (xi) the aggregation buffer forwarding completion to the asynchronous core-nest interface; (xii) the asynchronous core-nest interface storing the completion status in the input/output status buffer for the input/output status buffer index and signalling completion of operation to the system firmware; (xiii) the system firmware updating an input/output status buffer tracking
15 by the input/output status buffer index; and (xiv) the system firmware signalling asynchronously defects to the operating system in case of an error.

[0052] Only step (ii) is dependent on the length of the data and is different for the length of the data exceeding eight bytes from for the length of the data not exceeding eight
20 bytes.

[0053] According to the embodiment of the inventive method, the data are transmitted in slices to the aggregation buffer until all data of a store block are forwarded to the aggregation buffer, wherein the system firmware is waiting until the data has been sent by
25 the asynchronous core-nest interface.

[0054] Thus if data is less than eight bytes the filling process of the aggregation buffer in slices with data packets may be skipped and the transmit process of the data to the external device can be completed in a single step.

30

[0055] According to a favourable embodiment of the inventive method, the data may be transferred by the input/output store instruction through an asynchronous transmit mechanism with an early completion message in multiple data packets to the aggregation buffer, if the length of the data exceeds eight bytes. The asynchronous transmit mechanism

is favourable because the sending device is free for reuse at an earlier state.

[0056] According to a favourable embodiment of the inventive method, the system firmware may use an asynchronous input/output driver code for handling the input/output store instruction. Thus an asynchronous transmit mechanism may be used for transferring data from the data processing unit to the external device.

[0057] According to a favourable embodiment of the inventive method, the core may use an asynchronous setup code for handling memory requirements for status information of the asynchronous input/output driver code. This asynchronous setup code may further facilitate the asynchronous transmit mechanism through the aggregation buffer to the system nest and the input/output bus controller.

[0058] According to a favourable embodiment of the inventive method, the asynchronous core-nest interface may use an asynchronous core-nest interface forwarding component for forwarding the data with local completion. Thus a favourable asynchronous transmit mode for sending the data in data packets to the aggregation buffer may be supported.

[0059] According to a favourable embodiment of the inventive method, the aggregation buffer may use an early completion logic for delivering a free for reuse message after sending a request. This enables an early continuation of the transmit process of the data via the aggregation buffer to the system nest and the input/output bus controller.

[0060] Advantageously, the input/output status buffers may collect returned states from the system nest and/or from the input/output bus controller, in particular a completion message from the system nest. These input/output status buffers may collect the returned states acting as an asynchronous system message buffer supporting the asynchronous transmit process.

[0061] According to a favourable embodiment of the inventive method, a system message may comprise one of - a hierarchical physical target address, - sourcing an SMT thread or an aggregate buffer identifier, - a length of data, - an input/output bus address, or - an input/output status buffer index. Thus an advantageous passing of relevant information

through the data processing system can be guaranteed.

[0062] Further, a favorable computer program product is proposed for handling an input/output store instruction to at least one external device of a data processing system, the data processing system comprising a system nest communicatively coupled to at least one input/output bus by an input/output bus controller. The data processing system further comprises at least a data processing unit comprising a core, a system firmware and an asynchronous core-nest interface. The data processing unit is communicatively coupled to the system nest via an aggregation buffer. The external device is communicatively coupled to the input/output bus. The asynchronous core-nest interface comprises an input/output status array with multiple input/output status buffers, as well as an array management and access logic.

[0063] The computer program product comprises a computer readable storage medium having program instructions embodied therewith, the program instructions executable by the computer system to cause the computer system to perform a method comprising: (i) an operating system running on the data processing system issuing the input/output store instruction specifying at least an input/output function with an offset through an address, data to be transferred and/or a pointer to data to be transferred, and a length of the data; (ii) the data processing unit being configured to identify the input/output function by the address specified in the input/output store instruction; (iii) the data processing unit being configured to verify if access to the input/output function is allowed on an address space and on a guest instance level, the guest running on the data processing system; (iv) the data processing unit being configured to complete the input/output store instruction before an execution of the input/output store instruction in the system nest is completed; (v) the system firmware being configured to notify the operating system through an interrupt, if during the asynchronous execution of the input/output store instruction an error is detected by the data processing unit, transmitting the data of the failed asynchronous execution; (vi) a barrier instruction delaying a processing of a store instruction, until all previous store instructions have been completed.

[0064] Further, a data processing system for execution of a data processing program is proposed, comprising computer readable program instructions for performing the method described above.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

5 [0065] The present invention together with the above-mentioned and other objects and advantages may best be understood from the following detailed description of the embodiments, but not restricted to the embodiments.

[0066] Figure 1 depicts a block diagram of a data processing system for handling an
10 input/output store instruction to an external device according to an embodiment of the invention.

[0067] Figure 2 depicts a message sequence chart of a method for handling an
15 input/output store instruction to an external device according to an embodiment of the invention.

[0068] Figure 3 depicts a first part of a flow chart for handling an input/output store instruction to an external device according to an embodiment of the invention.

20 [0069] Figure 4 depicts a second part of a flow chart for handling an input/output store instruction to an external device according to an embodiment of the invention.

[0070] Figure 5 depicts an example embodiment of a data processing system for
25 executing a method according to the invention.

DETAILED DESCRIPTION

[0071] In the drawings, like elements are referred to with equal reference numerals. The
30 drawings are merely schematic representations, not intended to portray specific parameters of the invention. Moreover, the drawings are intended to depict only typical embodiments of the invention and therefore should not be considered as limiting the scope of the invention.

[0072] The illustrative embodiments described herein provide a data processing system

and a method for handling an input/output store instruction, comprising a system nest communicatively coupled to at least one input/output bus by an input/output bus controller. The data processing system further comprises at least a data processing unit comprising a core, a system firmware and an asynchronous core-nest interface. The data processing unit is communicatively coupled to the system nest via an aggregation buffer. The system nest is configured to asynchronously load from and/or store data to an external device which is communicatively coupled to the input/output bus. The asynchronous core-nest interface comprises an input/output status array with multiple input/output status buffers, as well as an array management and access logic.

10

[0073] The illustrative embodiments may be used for the method comprising: (i) an operating system running on the data processing system issues the input/output store instruction specifying at least an input/output function with an offset through an address, data to be transferred and/or a pointer to data to be transferred, and a length of the data; (ii) the data processing unit is configured to identify the input/output function by the address specified in the input/output store instruction; (iii) the data processing unit is configured to verify if access to the input/output function is allowed on an address space and on a guest instance level, the guest running on the data processing system; (iv) the data processing unit is configured to complete the input/output store instruction before an execution of the input/output store instruction in the system nest is completed; (v) the system firmware is configured to notify the operating system through an interrupt, if during the asynchronous execution of the input/output store instruction an error is detected by the data processing unit, transmitting the data of the failed asynchronous execution; (vi) a barrier instruction delaying a processing of a store instruction, until all previous store instructions have been completed.

25

[0074] Alternatively or additionally of data to be transferred, the store instruction according to an embodiment of the invention may also specify a pointer to main memory which should be used to fetch data from, instead of containing the data directly.

30

[0075] Guest instance level may also mean that a single guest or host may be running on the data processing system.

[0076] The address of the offset of the input/output function itself can be virtual,

physical, logical address. Virtual and logical addresses typically get translated through a memory management unit (MMU) into a physical address, and the physical address then allows to identify which function and offset is meant.

5 [0077] Physical address in this context means "lowest address in the address translation hierarchy accessible from within a guest/operating system".

[0078] Figure 1 depicts a block diagram of a data processing system 210 for handling an input/output store instruction 30 to at least one external device 214 according to an
10 embodiment of the invention. The data processing system 210 comprises a system nest 18 communicatively coupled to an input/output bus 22 by an input/output bus controller 20, a data processing unit 216 comprising a core 12, a system firmware 10 and an asynchronous core-nest interface 14. The input/output bus controller 20 may also be coupled via multiple input/output busses 22 to multiple external devices 214.

15

[0079] The data processing unit 216 is communicatively coupled to the system nest 18 via an aggregation buffer 16. The system nest 18 is configured to asynchronously load from and/or store data to the external device 214 which is communicatively coupled to the input/output bus 22 via a buffer-input/output bus controller interface 28 as part of the system
20 nest 18 and the input/output bus controller 20.

[0080] The aggregation buffer 16 is communicatively coupled to the asynchronous core-nest interface 14. The system firmware 10 comprises an asynchronous input/output driver code 32 for handling the input/output store instruction 30. The core 12 comprises an
25 asynchronous setup code 34 for handling memory requirements for status information of the asynchronous input/output driver code 32. The asynchronous core-nest interface 14 comprises an asynchronous core-nest interface forwarding component 36 for forwarding the data with local completion. The aggregation buffer 16 comprises an early completion logic 26 for delivering a free for reuse message after sending a request. The aggregation buffer 16
30 is coupled to the asynchronous core-nest interface 14 via an asynchronous bus 38. The asynchronous core-nest interface 14 comprises an input/output status array 44 with multiple input/output status buffers 24, as well as an array management and access logic 46. The input/output status buffers 24 collect returned states from the system nest 18 and/or from the input/output bus controller 20, in particular a completion message from the system nest 18.

The input/output status buffers 24 are integrated directly in the asynchronous core-nest interface 14. A message 48 with an identification of an array entry, e.g. a completion message to one of the input/output status buffers 24 may be received by the system nest 18.

5 [0081] The system firmware 10 comprises an array management logic 42, which allocates/deallocates input/output status buffers 24 in the input/output status array 44 and/or initiates a start of a new store instruction 30.

[0082] According to an embodiment of the inventive method, an operating system
10 running on the data processing system 210 issues the input/output store instruction 30 specifying at least an input/output function with an offset through an address, data to be transferred and/or a pointer to data to be transferred, and a length of the data. The data processing unit 216 is hereby configured to identify the input/output function by the address specified in the input/output store instruction 30. The data processing unit 216 is configured
15 to verify if access to the input/output function is allowed on an address space and on a guest instance level, the guest running on the data processing system 210. The data processing unit 216 is configured to complete the input/output store instruction 30 before an execution of the input/output store instruction 30 in the system nest 18 is completed. The system firmware 10 is configured to notify the operating system through an interrupt, if during the
20 asynchronous execution of the input/output store instruction 30 an error is detected by the data processing unit 216, transmitting the data of the failed asynchronous execution. A barrier instruction delays a processing of a store instruction 30, until all previous store instructions 30 have been completed.

25 [0083] The barrier instruction restricts a serial processing of store instructions 30 to a single input/output bus controller 20. The barrier instruction may further restrict a serial processing of store instructions 30 to a single input/output function, in particular a single external input/output device 214. A summary query capability for all input/output status buffers 24 of the input/output status array 44 is provided by an array management and
30 access logic 46, returning aggregated information of a status of the input/output status buffers 24.

[0084] The input/output status buffers 24 collect message states from the system nest 18 and/or from the input/output bus controller 20, in particular a completion status from the

system nest 18. The message states and/or the completion status may be favorably numbered by an input/output status buffer index.

[0085] The input/output store instruction 30 is located in the data processing system
5 210 on the side of the user interface 40 across the architecture boundary which separates the system hardware/firmware 50 from the user side 40.

[0086] Thus the data are transferred by the input/output store instruction 30 through an asynchronous transmit mechanism with an early completion message in multiple data
10 packets to the aggregation buffer 16, if the length of the source data exceeds eight bytes, else the data are transferred in one data packet.

[0087] A system message according to an embodiment of the inventive data processing system comprises one of a hierarchical physical target address, sourcing an SMT thread or
15 an aggregate buffer identifier, a length of data, an input/output bus address, or an input/output status buffer index.

[0088] The queueing and ordering semantics for handling store instructions 30 to multiple external devices 214 may advantageously performed as described in the following.
20 For an individual SMT thread versus input/output function relation, all legacy input/output load/store operations may be ordered in respect to a single thread of the processor unit 216. The new input/output store instructions are completely unordered amongst each other. New input/output store instructions are ordered against legacy input/output instructions. All input/output instructions for different input/output functions are not ordered against each
25 other.

[0089] Two types of additional ordering instructions are added to the ruleset of the system firmware 10. First a function barrier separating all previous input/output store instructions from all following input/output store instructions from an individual SMT
30 thread perspective, when the barrier instruction completes all previous input/output store instructions have been sent to the input/output bus 22 where this function resides. Secondly, an input/output SMT thread barrier instruction, which completes when all previous input/output store instructions have been sent to the input/output bus 22 for that SMT thread. For example, a lightweight exit/entry of a z Systems Start Interpretive Execution

(SIE) instruction has no effect on ordering, a full SIE exit/entry, where a virtual data processing unit is changed, is handled by an input/output SMT thread barrier.

[0090] Figure 2 depicts a message sequence chart of the method for handling an
5 input/output store instruction 30 to an external device 214 according to an embodiment of the invention.

[0091] As shown in Figure 2 the method starts with the operating system issuing the
input/output store instruction 30. In step S101, the system firmware 10 allocates a free
10 input/output status buffer index. If there is no free input/output status buffer index available, the system firmware 10 waits. In step S103, the system firmware 10 checks, if the store instruction can be injected into an asynchronous send engine. If this is possible, the process continues. If this is not possible, the store instruction is delayed until the store instructions causing the delay have been completed.

15

[0092] Next, as is indicated by the steps S100 and S104, the system firmware 10 issues repeatedly, if a length of the data exceeds eight bytes, a system message to send a data packet to the aggregation buffer 16 until all data of a store block have been forwarded to the aggregation buffer 16, while the system firmware 10 is waiting until the data have been sent
20 by the system message. In steps S102 and S106 a local completion message is sent back to the system firmware 10.

[0093] Then in step S108, the system firmware 10 issues a system message to the aggregation buffer 16 to forward the data asynchronously as single nest message to the
25 input/output bus controller 20, while waiting for the aggregation buffer 16 to send a completion message.

[0094] Next in step S110, the aggregation buffer 16 injects the nest message into the system nest 18, wherein the aggregation buffer 16 is free for reuse right after the send
30 operation, signaling back to the system firmware 10. Then the aggregation buffer 16 sends a free for reuse message.

[0095] In step S112, the system nest 18 forwards the message to the target location, followed by step S114, the input/output bus controller 20 receiving the message and

forwarding data in a data frame to the input/output bus, followed by the input/output bus controller 20 sending a completion message to the system nest 18 in step S116.

[0096] Next in step S118, the system nest 18 forwards the completion message to the
5 originating aggregation buffer 16, followed by the aggregation buffer 16 forwarding completion to the asynchronous core-nest interface 14 in step S120. Then in step S122 the asynchronous core-nest interface 14 stores the status in the input/output buffer 24 for the respective input/output status buffer index and signals completion of operation to the system
10 firmware 10. Finally in step S123, the system firmware 10 updates the input/output status buffer 24 tracking by the input/output status buffer index. The input/output status buffer 24 is now free again.

[0097] In case of an error occurring during transfer of data, the system firmware 10
15 signals asynchronously defects to the operating system.

[0098] In case, the data to be transferred are less than eight bytes, the repeatedly filling
of the aggregation buffer 16 is skipped.

[0099] Figure 3 depicts a first part of a flow chart for handling an input/output store
20 instruction 30 to an external device 214 according to an embodiment of the invention, whereas Figure 4 depict a second part of the flow chart.

[00100] The system firmware of the data processing unit starts in step S200. In step S202
the system firmware receives via a message a legacy input/output store block instruction. In
25 step S208 it is checked, if there is a new store instruction 30 outstanding with or without an asynchronous response. If this is the case, it is waited until the core-nest interface 14 has no more outstanding asynchronous responses. If this is not the case, the process continues immediately with step S212 by the core-nest interface 14 sending the legacy input/output store instruction 30 as a message. Then the system is waiting, step S214, until it receives a
30 legacy response message in step S216. Then the legacy store instruction 30 is finished in step S218 with a message and the system firmware 10 is ending the process in step S220.

[00101] In parallel the system firmware 10 receives the new input/output store
instruction 30 by a message in step S204. In step S222 it is checked if a free status slot,

namely an input/output status buffer 24 is available. If this is the case the slot is marked in step S224 as used and the store process continues with connection point A, depicted in the second part of the flow chart depicted in Figure 4. If this is not the case then it is waited until the core-nest interface 14 has a free slot available in step S226.

5

[00102] An asynchronous execution complete message may be received in step S206, followed by a update the slot tracking in step S228 with the respective input/output status buffer index. Then the system firmware 10 is ending the process in step S230.

- 10 [00103] In a parallel path the system firmware 10 receives a message with a barrier instruction per store function in step S232. Next it is checked in step S234 if the core-nest interface status buffers 24 are idle. If this is the case, the barrier instruction is finished in step S236 and the system firmware 10 is ending the process in step S238. If this is not the case, the system firmware 10 waits for a tracking update of a status buffer 24 in step S240.
- 15 If any of the status buffers 24 in step S242 is still non-idle for the function the loop continues by waiting in step S240. If any of the status buffers 24 is idle, the process continues with step S236 by finishing the barrier instruction.

- [00104] A barrier instruction is received in step S244, followed by waiting until the core-nest interface 14 has no more outstanding asynchronous responses in step S246, using the summary query capability provided in the core-nest interface 14 by the array management and access logic 46. Then the barrier instruction is finished in step S248 and the system firmware 10 is ending the process in step S250.

- 25 [00105] The second part of the flow chart, beginning with connection point A, is depicted in Figure 4. First in step S304 it is checked if the more than 8 bytes are to be transferred. If this is the case the core-nest interface fills the aggregation buffer with an up to 16 bytes message in step S306. The system firmware is waiting, step S308, until a message of local completion is sent in step S310, returning to step S304. If there are less
- 30 than 8 bytes left in the check of step S304, the flow continues in step S312 with the core-nest interface sending an asynchronous input/output message, followed by waiting in step S314 for a buffer response in step S316. Then in step S318 a finish store block instruction is executed and the flow ends in step S320 with an ending in the system firmware.

[00106] In step S328 the asynchronous core-nest interface logic starts an outbound process loop, followed by receiving an aggregation buffer completion message in step S322 and a forward data message to the aggregation buffer in step S324, followed by a send completion message back to the system firmware in step S326. In step S330 an
5 asynchronous input/output send message is received followed by a forward of the input/output send message to the aggregation buffer.

[00107] In step S338 the aggregation buffer logic starts an outbound process loop followed by a receive data in step S334 and aggregating data in the aggregation buffer in
10 step S336. The aggregation buffer is also receiving an input/output send message in step S340, followed by forwarding data from the aggregation buffer with an input/output send message in step S242. Next in step S344 a response message from the aggregation buffer is sent via the core-nest interface to the system firmware.

[00108] Referring now to Figure 5, a schematic of an example of a data processing system 210 is shown. Data processing system 210 is only one example of a suitable data processing system and is not intended to suggest any limitation as to the scope of use or functionality of embodiments of the invention described herein. Regardless, data processing system 210 is capable of being implemented and/or performing any of the functionality set
15 20 forth herein above.

[00109] In data processing system 210 there is a computer system/server 212, which is operational with numerous other general-purpose or special-purpose computing system environments or configurations. Examples of well-known computing systems,
25 environments, and/or configurations that may be suitable for use with computer system/server 212 include, but are not limited to, personal computer systems, server computer systems, thin clients, thick clients, handheld or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputer systems, mainframe computer systems, and distributed cloud
30 computing environments that include any of the above systems or devices, and the like.

[00110] Computer system/server 212 may be described in the general context of computer system executable instructions, such as program modules, being executed by a computer system. Generally, program modules may include routines, programs, objects,

components, logic, data structures, and so on that perform particular tasks or implement particular abstract data types. Computer system/server 212 may be practiced in distributed cloud computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed cloud computing
5 environment, program modules may be located in both local and remote computer system storage media including memory storage devices.

[00111] As shown in Fig. 5, computer system/server 212 in data processing system 210 is shown in the form of a general-purpose computing device. The components of computer
10 system/server 212 may include, but are not limited to, one or more processors or processing units 216, a system memory 228, and a bus 218 that couples various system components including system memory 228 to processor 216.

[00112] Bus 218 represents one or more of any of several types of bus structures,
15 including a memory bus or memory controller, a peripheral bus, an accelerated graphics port, and a processor or local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnect (PCI) bus.

[00113] Computer system/server 212 typically includes a variety of computer system readable media. Such media may be any available media that is accessible by computer system/server 212, and it includes both volatile and non-volatile media, removable and non-removable media.
20

[00114] System memory 228 can include computer system readable media in the form of volatile memory, such as random access memory (RAM) 230 and/or cache memory 232. Computer system/server 212 may further include other removable/non-removable, volatile/non-volatile computer system storage media. By way of example only, storage
30 system 234 can be provided for reading from and writing to a non-removable, non-volatile magnetic media (not shown and typically called a "hard drive"). Although not shown, a magnetic disk drive for reading from and writing to a removable, non-volatile magnetic disk (e.g., a "floppy disk"), and an optical disk drive for reading from or writing to a removable, non-volatile optical disk such as a CD-ROM, DVD-ROM or other optical media can be

provided. In such instances, each can be connected to bus 218 by one or more data media interfaces. As will be further depicted and described below, memory 228 may include at least one program product having a set (e.g., at least one) of program modules that are configured to carry out the functions of embodiments of the invention.

5

[00115] Program/utility 240, having a set (at least one) of program modules 242, may be stored in memory 228 by way of example, and not limitation, as well as an operating system, one or more application programs, other program modules, and program data. Each of the operating system, one or more application programs, other program modules, and program data or some combination thereof, may include an implementation of a networking environment. Program modules 242 generally carry out the functions and/or methodologies of embodiments of the invention as described herein.

10

[00116] Computer system/server 212 may also communicate with one or more external devices 214 such as a keyboard, a pointing device, a display 224, etc.; one or more devices that enable a user to interact with computer system/server 212; and/or any devices (e.g., network card, modem, etc.) that enable computer system/server 212 to communicate with one or more other computing devices. Such communication can occur via Input/Output (I/O) interfaces 222. Still yet, computer system/server 212 can communicate with one or more networks such as a local area network (LAN), a general wide area network (WAN), and/or a public network (e.g., the Internet) via network adapter 220. As depicted, network adapter 220 communicates with the other components of computer system/server 212 via bus 218. It should be understood that although not shown, other hardware and/or software components could be used in conjunction with computer system/server 212. Examples, include, but are not limited to: microcode, device drivers, redundant processing units, external disk drive arrays, RAID systems, tape drives, and data archival storage systems, etc.

15

20

[00117] The present invention may be a system, a method, and/or a computer program product. The computer program product may include a computer readable storage medium (or media) having computer readable program instructions thereon for causing a processor to carry out aspects of the present invention.

25

[00118] The computer readable storage medium can be a tangible device that can retain

and store instructions for use by an instruction execution device. The computer readable storage medium may be, for example, but is not limited to, an electronic storage device, a magnetic storage device, an optical storage device, an electromagnetic storage device, a semiconductor storage device, or any suitable combination of the foregoing. A non-
5 exhaustive list of more specific examples of the computer readable storage medium includes the following: a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), a static random access memory (SRAM), a portable compact disc read-only memory (CD-ROM), a digital versatile disk (DVD), a memory stick, a floppy disk, a
10 mechanically encoded device such as punch-cards or raised structures in a groove having instructions recorded thereon, and any suitable combination of the foregoing. A computer readable storage medium, as used herein, is not to be construed as being transitory signals per se, such as radio waves or other freely propagating electromagnetic waves, electromagnetic waves propagating through a waveguide or other transmission media (e.g.,
15 light pulses passing through a fiber-optic cable), or electrical signals transmitted through a wire.

[00119] Computer readable program instructions described herein can be downloaded to respective computing/processing devices from a computer readable storage medium or to an
20 external computer or external storage device via a network, for example, the Internet, a local area network, a wide area network and/or a wireless network. The network may comprise copper transmission cables, optical transmission fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers. A network adapter card or network interface in each computing/processing device receives computer readable program
25 instructions from the network and forwards the computer readable program instructions for storage in a computer readable storage medium within the respective computing/processing device.

[00120] Computer readable program instructions for carrying out operations of the
30 present invention may be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, or either source code or object code written in any combination of one or more programming languages, including an object oriented programming language such as Smalltalk, C++ or the like, and conventional procedural

programming languages, such as the "C" programming language or similar programming languages. The computer readable program instructions may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or
5 server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider). In some embodiments, electronic circuitry including, for example, programmable logic circuitry, field-programmable gate arrays
10 (FPGA), or programmable logic arrays (PLA) may execute the computer readable program instructions by utilizing state information of the computer readable program instructions to personalize the electronic circuitry, in order to perform aspects of the present invention.

[00121] Aspects of the present invention are described herein with reference to flowchart
15 illustrations and/or block diagrams of methods, apparatus (systems), and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions.

20

[00122] These computer readable program instructions may be provided to a processor of a general-purpose computer, special-purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means
25 for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks. These computer readable program instructions may also be stored in a computer readable storage medium that can direct a computer, a programmable data processing apparatus, and/or other devices to function in a particular manner, such that the computer readable storage medium having instructions stored therein comprises an article of
30 manufacture including instructions which implement aspects of the function/act specified in the flowchart and/or block diagram block or blocks.

[00123] The computer readable program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other device to cause a series

of operational steps to be performed on the computer, other programmable apparatus or other device to produce a computer implemented process, such that the instructions which execute on the computer, other programmable apparatus, or other device implement the functions/acts specified in the flowchart and/or block diagram block or blocks.

5

[00124] The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion
10 of instructions, which comprises one or more executable instructions for implementing the specified logical function(s). In some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It
15 will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special-purpose hardware-based systems that perform the specified functions or acts or carry out combinations of special-purpose hardware and computer instructions.

20

[00125] The descriptions of the various embodiments of the present invention have been presented for purposes of illustration, but are not intended to be exhaustive or limited to the embodiments disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the described
25 embodiments. The terminology used herein was chosen to best explain the principles of the embodiments, the practical application or technical improvement over technologies found in the marketplace, or to enable others of ordinary skill in the art to understand the embodiments disclosed herein.

30

REFERENCE NUMERALS

	10	system FW
	12	core
5	14	async core-nest IF
	16	aggregation buffer
	18	system nest
	20	I/O bus controller
	22	I/O bus
10	24	I/O status buffer
	26	early completion logic
	28	buffer-IO bus controller IF
	30	I/O store instruction
	32	async IO driver code
15	34	I/O setup code
	36	async forwarding
	38	async bus
	40	user IF
	42	array management logic
20	44	I/O status array
	46	array management & access logic
	48	message with array entry identification
	50	system HW/FW
25	210	data processing system
	212	computer system/server
	214	external devices
	216	CPU / data processing unit
	218	IO Bus
30	220	network adapter
	222	IO interfaces
	224	display
	228	memory
	230	RAM

- 232 cache
- 234 storage system
- 240 program/utility
- 242 program modules

5

CLAIMS

1. A data processing system for handling an input/output store instruction, the data processing system comprising:

a data processing unit including a core and system firmware, the data processing unit configured to perform a method, the method comprising:

identifying an input/output function by an address specified using the input/output store instruction, the input/output store instruction specifying at least the input/output function with an offset through the address, at least one of data to be transferred and a pointer to data to be transferred, and a length of the data;

verifying whether access to the input/output function is allowed on an address space and on a guest instance level;

completing the input/output store instruction before an execution of the input/output store instruction in a selected component of the data processing system different from the data processing unit is completed, the selected component configured to asynchronously load from and store data to at least one external device, the selected component being a system nest communicatively coupled to the data processing unit via an aggregation buffer, and communicatively coupled to asynchronously load data from and store data to the at least one external device via an input/output bus controller, the completing including:

issuing a system message to the aggregation buffer to forward the data to be transferred via the system nest as a single nest message to the input/output bus controller; and

receiving a free for reuse message after the nest message has been injected into the system nest to enable an early continuation of a transfer of data before a completion message is issued by the system nest;

providing notification through an interrupt, based on detecting an error during an asynchronous execution of the input/output store instruction in the data processing unit;

obtaining a barrier instruction which when executed assists in ordering of asynchronous store instructions; and

delaying, using the barrier instruction, processing of at least the input/output store instruction, until previous input/output store instructions have

been completed, the barrier instruction completing execution after the previous input/output store instructions have been completed.

- 5 2. The data processing system of claim 1, wherein execution of the barrier instruction restricts a serial processing of input/output store instructions to a single input/output bus controller.
- 10 3. The data processing system of claim 1, wherein execution of the barrier instruction restricts a serial processing of input/output store instructions to a single input/output function, wherein the single input/output function comprises a single external input/output device.
- 15 4. The data processing system of claim 1, wherein the method further comprises providing a summary query for one or more input/output status buffers of the input/output status array of the data processing unit, returning aggregated information of a status of the one or more input/output status buffers, the one or more input/output status buffers to collect returned states from the selected component.
- 20 5. The data processing system of claim 1, wherein the data is transferred by the input/output store instruction through an asynchronous transmit mechanism with an early completion message in multiple data packets to the aggregation buffer, based on the length of the data exceeding a defined size.
- 25 6. The data processing system of claim 5, wherein the data is transferred in one data packet, based on the length of the data not exceeding the defined size.
- 30 7. The data processing system according of claim 1, wherein the data processing unit comprises an asynchronous core-nest interface, the asynchronous core-nest interface comprising an input/output status array including a plurality of input/output status buffers to collect returned states from the selected components, the aggregation buffer being further communicatively coupled to the asynchronous core-nest interface.

8. The data processing system of claim 7, wherein the system firmware comprises an asynchronous input/output driver code to handle the input/output store instruction.
- 5 9. The data processing system of claim 8, wherein the core comprises an asynchronous setup code to handle memory requirements for status information of the asynchronous input/output driver code.
- 10 10. The data processing system of claim 7, wherein the asynchronous core-nest interface comprises an asynchronous core-nest interface forwarding component to forward data with local completion.
- 15 11. The data processing system of claim 7, wherein the aggregation buffer comprises an early completion logic to deliver the free for reuse message after sending a request.
12. A computer-implemented method of handling an input/output store instruction, the computer-implemented method comprising:
- 20 identifying an input/output function by an address specified using the input/output store instruction, the input/output store instruction specifying at least the input/output function with an offset through the address, at least one of data to be transferred and a pointer to data to be transferred, and a length of the data;
- 25 verifying whether access to the input/output function is allowed on an address space and on a guest instance level;
- 30 completing, by a data processing unit of a data processing system, the input/output store instruction before an execution of the input/output store instruction in a selected component of the data processing system different from the data processing unit is completed, the selected component configured to asynchronously load from and store data to at least one external device, the data processing unit including a core and system firmware, the selected component being a system nest communicatively coupled to the data processing unit via an aggregation buffer, and communicatively coupled to asynchronously load data from and store data to the at least one external device via an input/output bus controller, the completing including:
- 35 issuing a system message to the aggregation buffer to forward the data

to be transferred via the system nest as a single nest message to the input/output bus controller; and

receiving a free for reuse message after the nest message has been injected into the system nest to enable an early continuation of a transfer of data before a completion message is issued by the system nest;

providing notification through an interrupt, based on detecting an error during an asynchronous execution of the input/output store instruction in the data processing unit;

obtaining a barrier instruction which when executed assists in ordering of asynchronous store instructions; and

delaying, using the barrier instruction, processing of at least the input/output store instruction, until previous input/output store instructions have been completed, the barrier instruction completing execution after the previous input/output store instructions have been completed.

13. The computer-implemented method of claim 12, wherein execution of the barrier instruction restricts a serial processing of input/output store instructions to a single input/output bus controller.

14. The computer-implemented method of claim 12, the execution of the barrier instruction restricts a serial processing of input/output store instructions to a single input/output function, wherein the single input/output function comprises a single external input/output device.

15. The computer-implemented method of claim 12, further comprising providing a summary query for one or more input/output status buffers of an input/output status array of the data processing unit, returning aggregated information of a status of the one or more input/output status buffers, the one or more input/output status buffers to collect returned states from the selected component.

16. A computer program product for handling an input/output store instruction, the computer program product comprising:
at least one computer readable storage medium readable by at least one processing circuit and storing instructions for performing a method comprising:

identifying an input/output function by an address specified using the input/output store instruction, the input/output store instruction specifying at least the input/output function with an offset through the address, at least one of data to be transferred and a pointer to data to be transferred, and a length of the data;

5 verifying whether access to the input/output function is allowed on an address space and on a guest instance level;

10 completing, by a data processing unit of a data processing system, the input/output store instruction before an execution of the input/output store instruction in a selected component of the data processing system different from the data processing unit is completed, the selected component configured to asynchronously load from and store data to at least one external device, and the data processing unit including a core and system firmware, the selected component being a system nest communicatively coupled to the data processing unit via an aggregation buffer, and communicatively coupled to asynchronously load data from and store data to the at least one external device via an input/output bus controller, the completing including:

15 issuing a system message to the aggregation buffer to forward the data to be transferred via the system nest as a single nest message to the input/output bus controller; and

20 receiving a free for reuse message after the nest message has been injected into the system nest to enable an early continuation of a transfer of data before a completion message is issued by the system nest;

25 providing notification through an interrupt, based on detecting an error during an asynchronous execution of the input/output store instruction in the data processing unit;

obtaining a barrier instruction which when executed assists in ordering of asynchronous store instructions; and

30 delaying, using the barrier instruction, processing of at least the input/output store instruction, until previous input/output store instructions have been completed, the barrier instruction completing execution after the previous input/output store instructions have been completed.

17. The computer program product of claim 16, wherein execution of the barrier instruction restricts a serial processing of input/output store instructions to a single input/output bus controller.
- 5 18. The computer program product of claim 16, wherein execution of the barrier instruction restricts a serial processing of input/output store instructions to a single input/output function, wherein the single input/output function comprises a single external input/output device.
- 10 19. The computer program product of claim 16, wherein the method further comprises providing a summary query for one or more input/output status buffers of an input/output status array of the data processing unit, returning aggregated information of a status of the one or more input/output status buffers, the one or more input/output status buffers to collect returned states from the selected
15 component.

International Business Machines Corporation
Patent Attorneys for the Applicant/Nominated Person
SPRUSON & FERGUSON




Fig. 1




Fig. 2

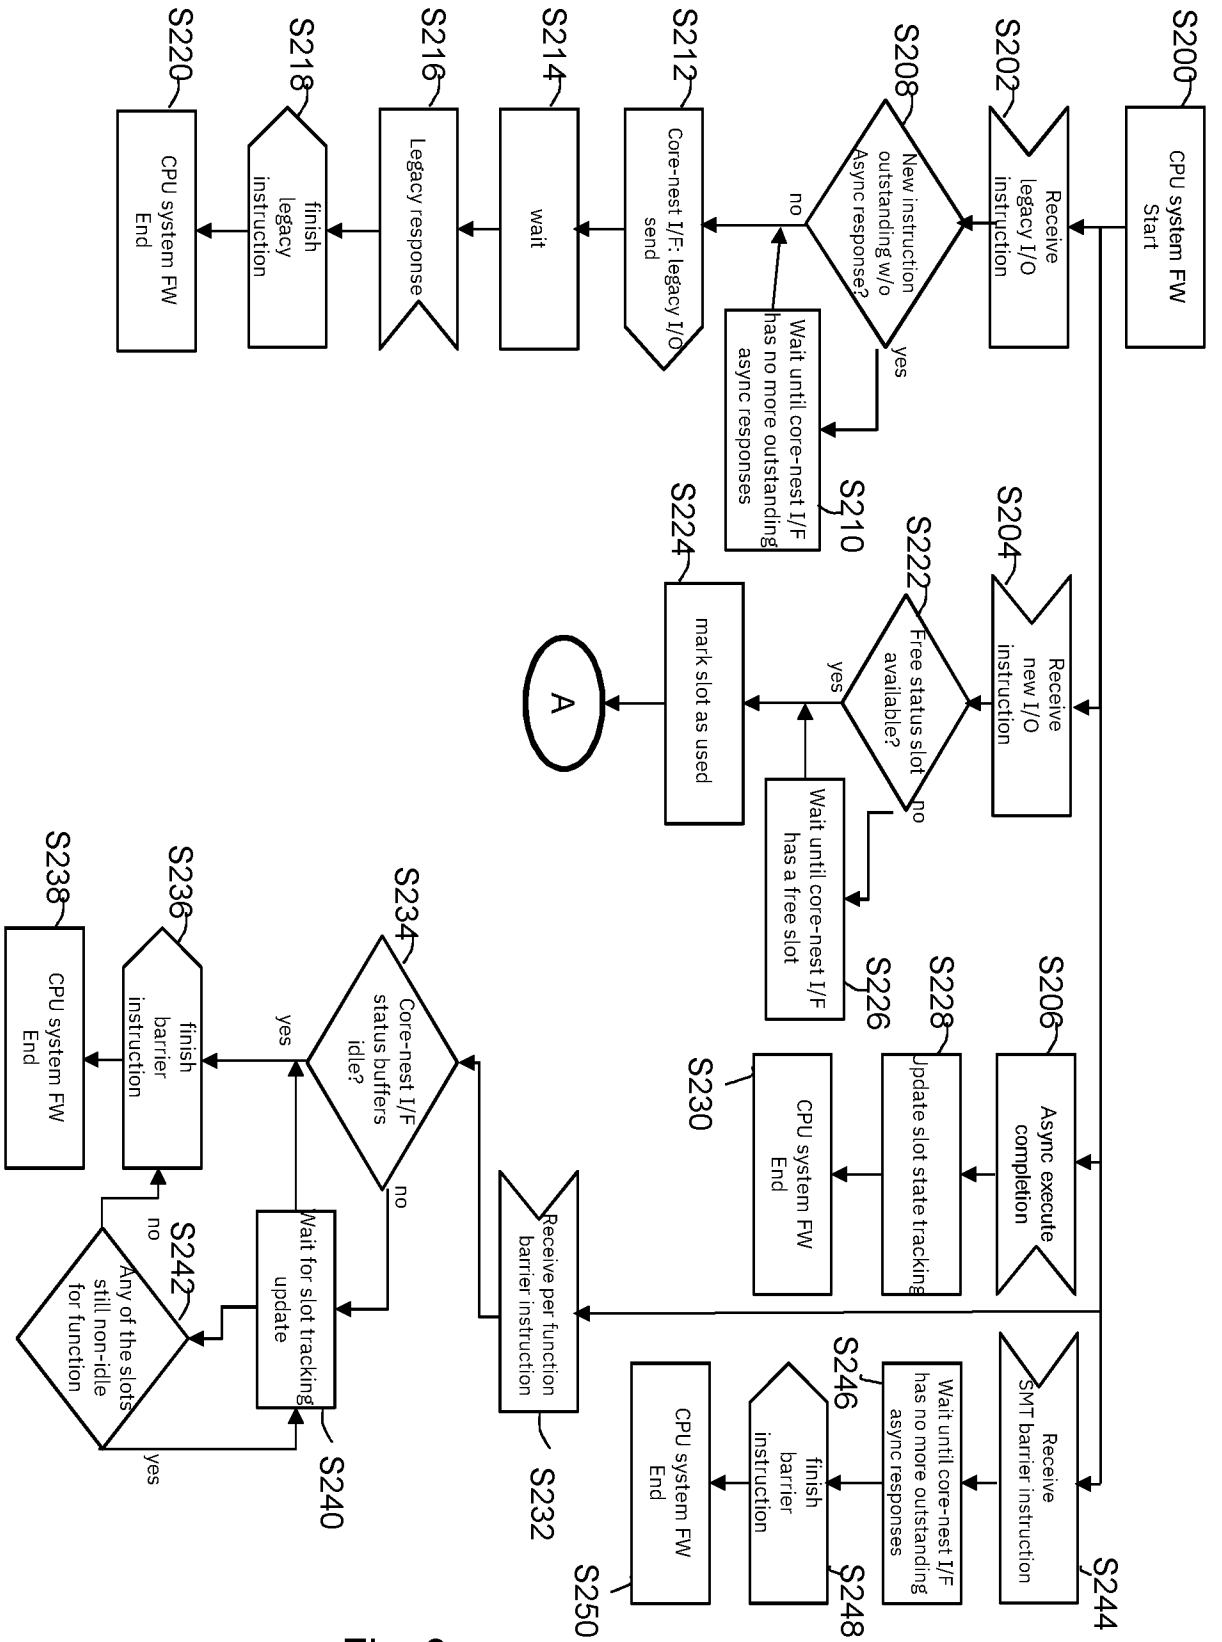


Fig. 3

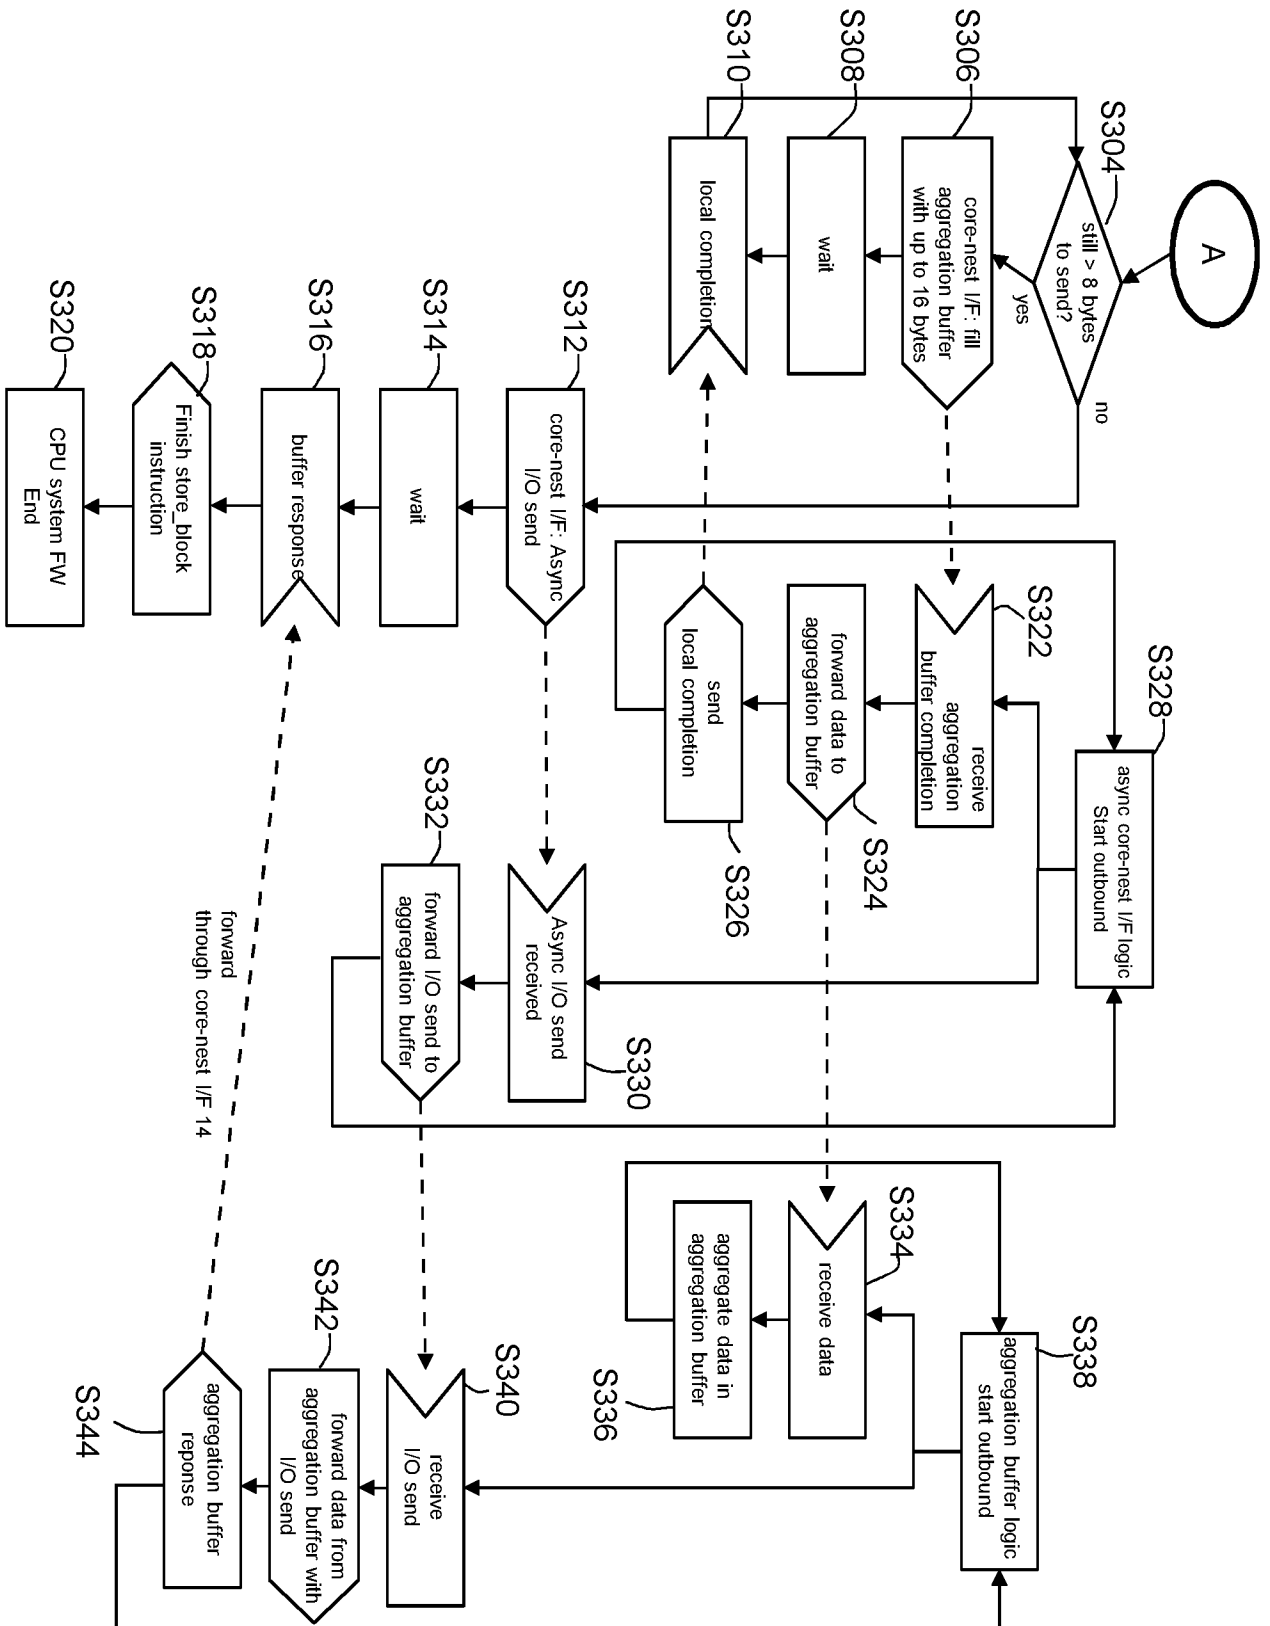


Fig. 4




Fig. 5