

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
17 February 2011 (17.02.2011)

(10) International Publication Number  
**WO 2011/020053 A1**

- (51) International Patent Classification:  
*G06F 9/48* (2006.01)
- (21) International Application Number:  
PCT/US2010/045517
- (22) International Filing Date:  
13 August 2010 (13.08.2010)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
12/540,749 13 August 2009 (13.08.2009) US
- (71) Applicant (for all designated States except US): **QUALCOMM INCORPORATED** [US/US]; ATTN: INTERNATIONAL IP ADMINISTRATION, 5775 Morehouse Drive, San Diego, California 92121 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **KOHLENZ, Mathias** [DE/US]; 5775 Morehouse Drive, San Diego, California 92121 (US). **KHAN, Irfan Anwar** [PK/US]; 5775 Morehouse Drive, San Diego, California 92121 (US). **SATHYANARAYAN, Madhusudan** [IN/US]; 5775 Morehouse Drive, San Diego, California 92121 (US). **MAHESHWARI, Shailesh** [IN/US]; 5775 Morehouse Drive, San Diego, California 92121 (US). **KRISH-**

**NAMOORTHY, Srividhya** [US/US]; 5775 Morehouse Drive, San Diego, California 92121 (US). **URGAONKAR, Sandeep** [IN/US]; 5775 Morehouse Drive, San Diego, California 92121 (US). **KLINGENBRUNN, Thomas** [DK/US]; 5775 Morehouse Drive, San Diego, California 92121 (US). **LIOU, Tim Tynghuei** [US/US]; 5775 Morehouse Drive, San Diego, California 92121 (US). **MIR, Idreas** [US/US]; 5775 Morehouse Drive, San Diego, California 92121 (US).

(74) Agent: **SIMON, Darren M.**; ATTN: INTERNATIONAL IP ADMINISTRATION, 5775 Morehouse Drive, San Diego, California 92121 (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, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH,

[Continued on next page]

(54) Title: APPARATUS AND METHOD FOR EFFICIENT DATA PROCESSING

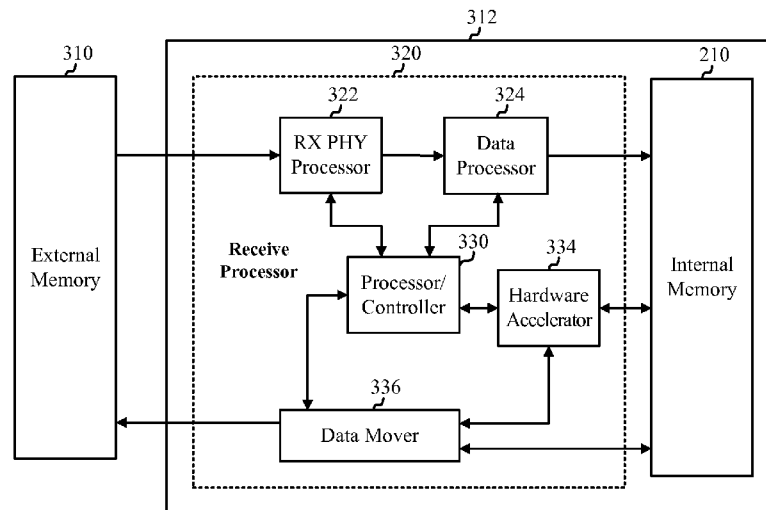


FIG. 3

(57) Abstract: Efficient data processing apparatus and methods include hardware components which are pre-programmed by software. Each hardware component triggers the other to complete its tasks. After the final pre-programmed hardware task is complete, the hardware component issues a software interrupt.

WO 2011/020053 A1



GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, 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, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

— 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))

**Published:**

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

**Declarations under Rule 4.17:**

## **APPARATUS AND METHOD FOR EFFICIENT DATA PROCESSING**

### **BACKGROUND**

#### **Field**

[0001] This disclosure relates generally to data processing, and more specifically to techniques for more efficiently processing data using the hardware in a communication system.

#### **Background**

[0002] Generally, the data processing functions within a communication system employ the use of software to implement numerous tasks. Such software normally provides the intelligence for certain hardware operations. Consequently, there must be close interaction between the software and hardware, thus requiring, in some instances, many steps to implement certain tasks. Dependence on software results in a number of disadvantages, including increased latencies, wasted bandwidth, increased processing load on the system microprocessor, and the like. As such, there is a need for a more efficient way of processing data while improving upon or eliminating the disadvantages brought about by the current paradigm where there is a need for close interaction between hardware and the software providing most of the intelligence for controlling such hardware.

### **SUMMARY**

[0003] The following presents a simplified summary of one or more aspects in order to provide a basic understanding of such aspects. This summary is not an extensive overview of all contemplated aspects, and is intended to neither identify key or critical elements of all aspects nor delineate the scope of any or all aspects. Its sole purpose is to present some concepts of one or more aspects in a simplified form as a prelude to the more detailed description that is presented later.

[0004] According to some aspects, a data processing method comprises receiving, by a first hardware component having associated therewith a first task queue, instructions to execute a first set of tasks; receiving, by a second hardware component having

associated therewith a second task queue, instructions to execute a second set of tasks; receiving, at the first hardware component, a trigger message to execute the first set of tasks; and upon completion of the first set of tasks, transmitting, by the first hardware component, a trigger message to the second hardware component to execute the second set of tasks.

[0005] According to some aspects, a data processing apparatus comprises a processor; a first memory; a second memory; a first hardware component having associated therewith a first task queue; and a second hardware component having associated therewith a second task queue, wherein a plurality of tasks are pre-stored in the first tasks queue, a plurality of second tasks are pre-stored in the second task queue, and the first hardware component is configured to receive a trigger to begin performing the plurality of first tasks and to issue a trigger to the second hardware component to begin performing the plurality of second tasks.

[0006] To the accomplishment of the foregoing and related ends, the one or more aspects comprise the features hereinafter fully described and particularly pointed out in the claims. The following description and the annexed drawings set forth in detail certain illustrative features of the one or more aspects. These features are indicative, however, of but a few of the various ways in which the principles of various aspects may be employed, and this description is intended to include all such aspects and their equivalents.

#### **BRIEF DESCRIPTION OF THE DRAWINGS**

[0007] The disclosed aspects will hereinafter be described in conjunction with the appended drawings, provided to illustrate and not to limit the disclosed aspects, wherein like designations denote like elements, and in which:

[0008] Fig. 1 depicts a wireless communication system, in accordance with various disclosed aspects;

[0009] Fig. 2 depicts a user equipment and Node B, in accordance with various disclosed aspects;

[0010] Fig. 3 depicts a receiver, in accordance with various disclosed aspects;

[0011] Fig. 4 is a data flow diagram depicting a data processing operation, in accordance with various disclosed aspects;

[0012] Fig. 5 is a simplified system diagram, in accordance with various disclosed aspects;

[0013] Fig. 6 is a data flow diagram depicting a data processing operation, in accordance with various disclosed aspects;

[0014] Fig. 7 is a flowchart depicting a data processing operation, in accordance with various disclosed aspects;

[0015] Fig. 8 is a flowchart depicting a threshold processing operation, in accordance with various disclosed aspects;

[0016] Fig. 9 is a system diagram implementing template processing, in accordance with various disclosed aspects;

[0017] Fig. 10 is a system diagram implementing multiple memory pools, in accordance with various disclosed aspects;

[0018] Fig. 11 depicts protocol data unit storage, in accordance with various disclosed aspects;

[0019] Fig. 12 is a simplified block diagram illustrating data processing and storage, in accordance with various disclosed aspects;

[0020] Fig. 13 is a block diagram depicting a typical wireless device;

[0021] Fig. 14 is a block diagram depicting an exemplary wireless device, in accordance with various disclosed aspects;

[0022] Fig. 15 is a simplified block diagram implementing a mini-processor, in accordance with various disclosed aspects;

[0023] Fig. 16 is a flowchart depicting a data processing method using a mini-processor, in accordance with various disclosed aspects;

[0024] Fig. 17 is a block diagram implementing a copy engine, in accordance with various disclosed aspects; and

Fig. 18 depicts an exemplary data frame, in accordance with various disclosed aspects.

#### **DETAILED DESCRIPTION**

[0025] Various aspects are now described with reference to the drawings. In the following description, for purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of one or more aspects. It may be evident, however, that such aspect(s) may be practiced without these specific details.

[0026] As used in this application, the terms “component,” “module,” “system” and the like are intended to include a computer-related entity, such as but not limited to

hardware, firmware, a combination of hardware and software, software, or software in execution. For example, a component may be, but is not limited to being, a process running on a processor, a processor, an object, an executable, a thread of execution, a program, and/or a computer. By way of illustration, both an application running on a computing device and the computing device can be a component. One or more components can reside within a process and/or thread of execution and a component may be localized on one computer and/or distributed between two or more computers. In addition, these components can execute from various computer readable media having various data structures stored thereon. The components may communicate by way of local and/or remote processes such as in accordance with a signal having one or more data packets, such as data from one component interacting with another component in a local system, distributed system, and/or across a network such as the Internet with other systems by way of the signal.

[0027] Furthermore, various aspects are described herein in connection with a terminal, which can be a wired terminal or a wireless terminal. A terminal can also be called a system, device, subscriber unit, subscriber station, mobile station, mobile, mobile device, remote station, remote terminal, access terminal, user terminal, terminal, communication device, user agent, user device, or user equipment (UE). A wireless terminal may be a cellular telephone, a satellite phone, a cordless telephone, a Session Initiation Protocol (SIP) phone, a wireless local loop (WLL) station, a personal digital assistant (PDA), a handheld device having wireless connection capability, a computing device, or other processing devices connected to a wireless modem. Moreover, various aspects are described herein in connection with a base station. A base station may be utilized for communicating with wireless terminal(s) and may also be referred to as an access point, a Node B, or some other terminology.

[0028] Moreover, the term “or” is intended to mean an inclusive “or” rather than an exclusive “or.” That is, unless specified otherwise, or clear from the context, the phrase “X employs A or B” is intended to mean any of the natural inclusive permutations. That is, the phrase “X employs A or B” is satisfied by any of the following instances: X employs A; X employs B; or X employs both A and B. In addition, the articles “a” and “an” as used in this application and the appended claims should generally be construed to mean “one or more” unless specified otherwise or clear from the context to be directed to a singular form.

[0029] The techniques described herein may be used for various wireless communication systems such as CDMA, TDMA, FDMA, OFDMA, SC-FDMA and other systems. The terms “system” and “network” are often used interchangeably. A CDMA system may implement a radio technology such as Universal Terrestrial Radio Access (UTRA), cdma2000, etc. UTRA includes Wideband-CDMA (W-CDMA) and other variants of CDMA. Further, cdma2000 covers IS-2000, IS-95 and IS-856 standards. A TDMA system may implement a radio technology such as Global System for Mobile Communications (GSM). An OFDMA system may implement a radio technology such as Evolved UTRA (E-UTRA), Ultra Mobile Broadband (UMB), IEEE 802.11 (Wi-Fi), IEEE 802.16 (WiMAX), IEEE 802.20, Flash-OFDM, etc. UTRA and E-UTRA are part of Universal Mobile Telecommunication System (UMTS). 3GPP Long Term Evolution (LTE) is a release of UMTS that uses E-UTRA, which employs OFDMA on the downlink and SC-FDMA on the uplink. UTRA, E-UTRA, UMTS, LTE and GSM are described in documents from an organization named “3rd Generation Partnership Project” (3GPP). Additionally, cdma2000 and UMB are described in documents from an organization named “3rd Generation Partnership Project 2” (3GPP2). Further, such wireless communication systems may additionally include peer-to-peer (e.g., mobile-to-mobile) *ad hoc* network systems often using unpaired unlicensed spectrums, 802.xx wireless LAN, BLUETOOTH and any other short- or long- range, wireless communication techniques.

[0030] Various aspects or features will be presented in terms of systems that may include a number of devices, components, modules, and the like. It is to be understood and appreciated that the various systems may include additional devices, components, modules, etc. and/or may not include all of the devices, components, modules etc. discussed in connection with the figures. A combination of these approaches may also be used.

[0031] Figure 1 shows a wireless communication system 100, which includes (i) a radio access network (RAN) 120 that supports radio communication for user equipments (UEs) and (ii) network entities that perform various functions to support communication services. RAN 120 may include any number of Node Bs and any number of Radio Network Controllers (RNCs). For simplicity, only one Node B 130 and one RNC 132 are shown in FIG. 1. A Node B is generally a fixed station that communicates with the UEs and may also be referred to as an evolved Node B, a base station, an access point,

etc. RNC 132 couples to a set of Node Bs and provides coordination and control for the Node Bs under its control.

**[0032]** An Internet Protocol (IP) gateway 140 supports data services for UEs and may also be referred to as a Serving GPRS Support Node or Gateway GPRS Support Node (SGSN or GGSN), an Access Gateway (AGW), a Packet Data Serving Node (PDSN), etc. IP gateway 140 may be responsible for establishment, maintenance, and termination of data sessions for UEs and may further assign dynamic IP addresses to the UEs. IP gateways 140 may couple to data network(s) 150, which may comprise a core network, the Internet, etc. IP gateway 140 may be able to communicate with various entities such as a server 160 via data network(s) 150. Wireless system 100 may include other network entities not shown in FIG. 1.

**[0033]** A UE 110 may communicate with RAN 120 to exchange data with other entities such as server 160. UE 110 may be stationary or mobile and may also be referred to as a mobile station, a terminal, a subscriber unit, a station, etc. UE 110 may be a cellular phone, a personal digital assistant (PDA), a wireless communication device, a wireless modem, a handheld device, a laptop computer, etc. UE 110 may communicate with Node B 120 via air interfaces, such as a downlink and an uplink. The downlink (or forward link) refers to the communication link from the Node B to the UE, and the reverse link (or uplink) refers to the communication link from the UE to the Node B. The techniques described herein may be used for data transmission on the downlink as well as the uplink.

**[0034]** In some aspects, UE 110 may be coupled to a terminal equipment (TE) device 112 via a wired connection (as shown in FIG. 1) or a wireless connection. UE 110 may be used to provide or support wireless data services for TE device 112. TE device 112 may be a laptop computer, a PDA, or some other computing device.

**[0035]** Fig. 2 shows a block diagram of an exemplary design of UE 110 and Node B 130 in FIG. 1. In this exemplary design, UE 110 includes a data processor 210, an external memory 260, a radio transmitter (TMTR) 252, and a radio receiver (RCVR) 254. Data processor 210 includes a controller/processor 220, a transmit processor 230, a receive processor 240, and an internal memory 250. A bus 212 may facilitate data transfers between processors 220, 230 and 240 and memory 250. Data processor 210 may be implemented on an application specific integrated circuit (ASIC), and memory 260 may be external to the ASIC.



[0036] For data transmission on the uplink, transmit processor 230 may process traffic data in accordance with a set of protocols and provide output data. Radio transmitter 252 may condition (e.g., convert to analog, filter, amplify, and frequency upconvert) the output data and generate an uplink signal, which may be transmitted to Node B 230. For data reception on the downlink, radio receiver 254 may condition (e.g., filter, amplify, frequency downconvert, and digitize) a downlink signal received from Node B 230 and provide received data. Receive processor 240 may process the received data in accordance with a set of protocols and provide traffic data. Controller/processor 220 may direct the operation of various units at UE 110. Internal memory 250 may store program codes and data for processors 220, 230 and 240. External memory 260 may provide bulk storage for data and program codes for UE 110.

[0037] Fig. 2 also shows a block diagram of an exemplary design of Node B 130. A radio transmitter/receiver 264 may support radio communication with UE 110 and other UEs. A controller/processor 270 may process data for data transmission on the downlink and data reception on the uplink. Controller/processor 270 may also perform various functions for communication with the UEs. Memory 272 may store data and program codes for Node B 130. A communication (Comm) unit 274 may support communication with other network entities.

[0038] Figure 3 depicts an exemplary design of a receiver 300, which may be part of UE 110. Receiver 300 includes an external memory 310 that provides bulk storage of data, a receive processor 320 that processes received data, and an internal memory 340 that stores processed data. Receive processor 320 and internal memory 340 may be implemented on an ASIC 312.

[0039] Within receive processor 320, an RX PHY processor 322 may receive PHY frames from a transmitter (e.g., Node B 130), process the received PHY frames in accordance with the radio technology (e.g., CDMA2000 1X or 1xEV-DO or LTE) used by the RAN, and provide received frames. A data processor 324 may further process (e.g., decode, decipher, and de-frame) the received frames.

[0040] Data transfer typically involves three tasks. First, data is moved in from some location, such as external memory 310. The data is then operated on by hardware, and moved back out. Data mover 336 may be configured to move data in from one location to another, and to move the data back out after processing by hardware accelerator 334.

Processor/controller 330 may be configured to control the operation of hardware accelerator 334 and data mover 336.

[0041] Figure 4 is a data flow diagram depicting typical data processing operations occurring in software and hardware. As depicted in Figure 4, a software component 402 may interface with one or more hardware components such as a first hardware component 404 and a second hardware component 406. As depicted at step 408, software component 402 may issue a first set of instructions to first hardware component 404. The instructions may include instructions to perform one or more tasks. The first hardware component 404 then responds to the initial set of instructions, as depicted at 410. To begin processing with the second hardware component 406, software component 402 may then send a set of instructions to second hardware component 406, as depicted at 412, and await a response, as depicted at 414. Each time a task is to be executed, software 402 must issue an instruction or set of instructions to the designated hardware component, and await a response, as depicted at 416, 418, 420 and 422.

[0042] In the method depicted in Figure 4, software interrupts must be used after every hardware task or set of tasks. That is, after completing the provided instructions, a hardware component must send an interrupt to software. These interrupts may be costly, as the program context must be stalled, and built anew to perform the interrupt, and then to resume hardware processing. Moreover, interrupts often become locked, resulting in a latent period until the interrupt lock is resolved.

[0043] Figure 5 is a simplified block diagram illustrating various disclosed aspects. Processor 502 may be configured to control the operation of various hardware components, such as first hardware component 504 and second hardware component 506. Like in Figure 4, each of first hardware component 504 and second hardware component 506 may be a dedicated hardware block, such as, for example, a data mover or a hardware accelerator. Processor 502 may have stored thereon software which instructs the hardware components to perform various tasks.

[0044] First hardware component 504 and second hardware component 506 may be communicatively coupled to each other. As such, these components may provide information to each other without interrupting the software stored on processor 502. For example, first hardware component 504 and second hardware component 506 may

be pre-programmed to perform one or more tasks by processor 502. This process is depicted in greater detail in Figure 6.

[0045] As depicted in Figure 6, processor 502 may interface with first hardware component 504 and second hardware component 506. According to some aspects, the first hardware component 504 may be one suited for moving data into a local buffer (e.g., a data mover) from a variety of sources such as external memory. The second hardware component 506 may be a custom hardware block, such as a data accelerator.

[0046] As depicted at 610, processor 502 may issue a first set of commands to first hardware component 504 to perform a plurality of tasks. For example, processor 502 may instruct hardware component 504 to move in a block of data from an external storage location. Processor 502 may then program second hardware component 506 to perform a plurality of tasks, as depicted at 612. For example, processor 502 may instruct the second hardware component 506 to operate on the block of data moved in by the first hardware component 504. While the first hardware component and the second hardware component may be pre-programmed to execute the designated tasks, operation does not begin until processor 502 sends a trigger message, as depicted at 614.

[0047] Upon receipt of the trigger message, the first hardware component 504 may begin executing the pre-programmed tasks. According to some aspects, the final instruction may direct the first hardware component to send a trigger message to the second hardware component 504, as depicted at 616. The trigger may be, for example a hardware interrupt. As depicted at steps 618, 620 and 622, each hardware component triggers the other component to execute an instruction upon completion of its own processing. There is no need for software interrupts during this processing. When the last instruction has been processed, the hardware component processing the last instruction (here, the first hardware component) may then send an interrupt to processor 502, as depicted at 624. As depicted in Figure 6, the number of software interrupts may be significantly decreased. In some aspects, only a single software interrupt is required.

[0048] Figure 7 is a flowchart depicting a data processing operation in further detail, according to various disclosed aspects. As depicted at 702, the software component may program two or more hardware components. Programming the hardware components may include providing a series of instructions to be executed by each hardware component. The series of instructions may be stored in a task queue associated with the hardware components. According to some aspects, more than one

hardware component may share a single task queue, wherein the instructions for each hardware component are interleaved together.

[0049] According to an exemplary aspect, the two or more hardware components may include a data mover and a hardware accelerator. Other hardware components may also be used, alternatively or in addition to the data mover and the hardware accelerator. While the remaining processes depicted in Figure 7 will be described in relation to a data mover and a hardware accelerator, it is noted that this is only exemplary.

[0050] As described above, each hardware component may have associated with it a task queue. As such, the data mover and the hardware accelerator each have a task queue associated therewith. The task queues may be configured in advance by software. The task queue associated with the hardware accelerator may be gated by a write pointer which is controlled by software. A read pointer, also associated with the hardware accelerator, may be controlled by hardware.

[0051] When execution of the programmed tasks is desired, the software component may send a trigger signal to the data mover, as depicted at 704. Upon receipt of the trigger signal, the data mover may execute its first task, as depicted at 706. For example, the data mover may be instructed to read a first block of data from an external memory location, and move the data into an internal memory location. The data mover may be further instructed to write to the write pointer register associated with the hardware accelerator, thereby triggering the hardware accelerator to begin executing its assigned instructions. The data mover may also be instructed to halt processing until receipt of an event.

[0052] Upon receipt of the trigger from the data mover, hardware accelerator may begin executing its instructions, as depicted at 708. This may include, for example, performing various types of operations on the data loaded by the data mover. Upon completion of the queued tasks, the hardware accelerator may issue a hardware interrupt, which triggers the data mover to continue operating, as depicted at 710. Upon receipt of a trigger, the data mover may determine whether it has additional tasks to complete, as depicted at 712. If so, processing returns to step 706. If the data mover does not have additional tasks to perform, it may send an interrupt to software, as depicted at 714. Enabling the data mover and hardware accelerator to trigger each other by way of hardware interrupts reduces interrupt latency and reduces the number of

context switches. Moreover, the data mover and hardware accelerator are used more efficiently, thereby reducing overall latency.

[0053] As described above, hardware components, such as the hardware accelerator and data mover, may be pre-programmed by software to perform a plurality of tasks. In the case of certain hardware accelerators, the length of the output from the accelerator may be random. However, because the data mover is also pre-programmed, the data move-out instruction is typically programmed for the worst case output from the hardware accelerator. Such a configuration may lead to an unnecessary waste of bandwidth because most of the time the output length would be less than the pre-programmed length.

[0054] According to some aspects of the present apparatus and methods, a threshold may be defined which accommodates the data length of the most common outputs of the hardware accelerator. In some aspects, it may be preferable that the threshold be defined such that a majority of all outputs from the hardware accelerator would not exceed the threshold. However, the threshold may be configured according to any other desired parameters. For example, the threshold may be configured such that at least a predefined percentage of all outputs of the accelerator would not exceed the threshold, or such that at least some of the outputs of the accelerator would not exceed the threshold. Figure 8 is a flowchart illustrating threshold processing in further detail.

[0055] As depicted at 810, a hardware component, such as a hardware accelerator receives and processes data according to pre-programmed instructions. Upon completion of the processing, the hardware accelerator determines whether the result of the instruction, which is to be output by a second hardware component, such as a data mover, exceeds a predefined threshold, as depicted at 812. More particularly, the hardware accelerator may include a threshold determination module which is pre-configured by software. The threshold may set a maximum value for the size of the data processed by the hardware accelerator.

[0056] As depicted at 814, if the threshold is exceeded, the hardware accelerator may generate an interrupt indicating that the move-out task was under provisioned. That is, the interrupt indicates that the size of the data move-out, which was preprogrammed in the data mover, is not large enough to accommodate the data just processed by the hardware accelerator. Upon receipt of an interrupt, software may read the excess data beyond the limits of the programmed data move-out size, as depicted at 816. The

hardware accelerator may be stalled during this interrupt, and may be given a “go” by the software once the interrupt is resolved. Processing may then continue. For example, the hardware accelerator or the software component may send a trigger to the data mover, and the data mover may move out data according to the pre-programmed instructions, as depicted at 818. If it is determined at step 812 that the threshold is not exceeded, then processing may also continue at step 818, where the data mover moves the data out to main memory as programmed.

[0057] According to some aspects, the threshold described in reference to Figure 8 may be adaptively adjusted. For example, if the frequency of interrupts resulting from an under-provisioned data move-out instruction is large over a given period of time, software may increase the size of the threshold. Likewise, the size of the threshold may be decreased if a large number of interrupts are not received over a given period of time. Thus, software may be configured with a timer and a counter which keep track of all interrupts received as a result of un-provisioned data-move out instructions. Controlling the amount of data moved out while providing a mechanism to account for under-provisioning may reduce power consumption and system bus bandwidth. Additionally, software complexity may be reduced by getting a fewer number of interrupts. Moreover, hardware is kept simple by passing exceptional cases to software for handling. Adaptively adjusting the threshold may help to ensure a good trade-off between maintaining a low interrupt rate and avoiding moving data unnecessarily.

[0058] According to some aspects, operations typically performed in software may be moved over to hardware to reduce software processing and latency, as well as to reduce use of bus bandwidth. For example, a hardware accelerator may be configured with logic to process data frame headers. Typically, hardware simply deciphers data and then forwards the data to software for further processing. Software then forwards the data back to hardware with instructions for processing the data. This may reduce software processing, latency, and bus bandwidth.

[0059] In some exemplary aspects, for example, a hardware accelerator may be configured with logic to parse a data frame header, compare the header to a plurality of predefined templates, and determine a next processing step based on the headers. Software interaction is needed only if a header match is not found, or if template reconfiguration is necessary.

[0060] Figure 9 is an example of an environment wherein the template processing may be implemented. As depicted in Figure 9, hardware accelerator 922, controlled by a control processor 920, may receive input data, and process the input data via a filtering module 922. The filtering module 922 may provide its output to one or more of processor 914, memory 916, or interfaces 918-1 through 918-N. Memory 916 may be an internal or an external memory. Interfaces 918-1 through 918-N may provide connections to external devices, such as, for example, laptops, PDAs, or any other external electronic device. For example, interfaces 918-1 through 918-N may include a USB port, Bluetooth, SDIO, SDCC, or other wired or wireless interfaces.

[0061] Filtering module 922 may include a plurality of predefined templates which may be used to make routing decisions as to where output data should be routed. For example, ciphered data may be received by hardware accelerator 912 and deciphered. Based on the template programmed in the filtering module 922, the deciphered data may be forwarded to processor 914, memory 916, or to an external device via interfaces 918-1 through 918-N. According to some aspects, processor 914 may be used as a back-up filter for more complex operations which may be difficult for hardware accelerator 912 to process.

[0062] According to some aspects, each traffic template may include one or more parameters and a specific value for each parameter. Each parameter may correspond to a specific field of a header for IP, TCP, UDP, or some other protocol. For example, IP parameters may include source address, destination address, address range, and protocol. TCP or UDP parameters may include source port, destination port, and port range. The traffic template specifies the location of each parameter in the header. Thus, the hardware accelerator need not have any actual knowledge of the protocols in use. Rather, the hardware performs a matching of the header to the template parameters.

[0063] For example, traffic templates may be defined to detect for TCP frames for destination port x and sent in IPv4 packets. The template may include three parameters that may be set as follows: version = IPv4, protocol = TCP, and destination port = x. In general, any field in any protocol header may be used as a traffic template parameter. Any number of templates may be defined, and each template may be associated with any set of parameters. Different templates may be defined for different application, sockets, etc., and may be defined with different sets of parameters.

[0064] Each template may also be associated with an action to perform if there is a match and an action to perform if there is no match. Upon receipt of a data frame, the received values of the frame may be compared against the specified values of a template. A match may be declared if the received values match the specified values, and no match may be declared otherwise. If no match is defined, hardware may issue a software interrupt, and software may then process the frame.

[0065] As described herein, data processing typically involves moving data in from a first location, operating on the data, and moving the data back out. Typically, a single large memory pool is defined in a low cost, high latency memory. Data is moved in from this high latency memory pool, operated on, and then moved back to the memory pool. However, data that has been recently accessed is often re-used. Thus, moving data back out to the high latency memory pool after every access unnecessarily increases system bus bandwidth. According to some aspects of the present apparatus and methods, multiple memory pools may be defined in physical memories. Memory allocation may be based on the best memory pool available, or may depend on how often the data is likely to be accessed.

[0066] Figure 10 is an example of a system 1000 implementing multiple memory pools, in accordance with some aspects. System 1000 may include an ASIC 910 comprising low latency memory 1012, processor 1014, data mover 1016, and hardware accelerator 1018. A high latency memory 1020 may also be provided. While low latency memory 912 is depicted as an internal memory and high latency memory 1020 is depicted as an external memory, this configuration is merely exemplary. Either or both memories may be internal or external. Processor 1004 may include a memory controller 1022 which controls access to high latency memory 1020 and low latency memory 1012. If most operations use low latency memory 1012, data can be processed more efficiently. As such, memory controller 1022 may be configured to limit the number of accesses to high latency memory 1020.

[0067] As in normal operation, data mover 1020 may be configured to move data in and out, for example, between high latency memory 1020 and low latency memory 912. In accordance with some exemplary aspects, hardware accelerator 1018 may be configured to operate on data directly from low latency memory 1012. For example, in some aspects, data may be maintained in low latency memory 1012 as long as space is available. In other aspects, data may be stored and maintained in low latency memory



1012 based on specific data transmission characteristics, such as quality of service requirements, communication channel properties, and/or other factors.

[0068] Providing multiple memory pools may reduce hardware costs, as a small, fast pool and a large, slower pool can be defined. Moreover, system bus bandwidth and power may be reduced.

[0069] According to some aspects, data stored in low latency memory 1012 may be accessed by both hardware and software. The payload may be processed by hardware accelerator 1018, ensuring low-latency access. The low latency memory 1012 may be placed in close proximity to hardware accelerator 1018. As such, data transfers do not need to cross a system bus, saving power and system bus bandwidth.

[0070] When processing data, the data is often copied multiple times while passing through the data stack in order to simplify implementation at each layer (e.g., when removing headers, multiplexing data from multiple flows, segmentation/reassembly, etc.). According to some aspects of the present apparatus and methods, repeated copying may be prevented by leaving data at the same location and having the different layers operate on the same data. Each data operation instruction, whether performed by hardware or software, may point to the same location, such as a local hardware buffer. For example, after UMTS deciphering has been performed, the deciphered data may be copied back into the local memory, such as a local hardware buffer. After software concatenates payloads by evaluating protocol headers, software may instruct a hardware accelerator to perform either TCP checksum calculation or PPP framing. Data need not be moved back and forth from an external memory location during such processing.

[0071] For data transmission, data frames are typically partitioned into smaller units, depending on the protocol in use. As such, related payloads are often divided into multiple transmissions. A memory block for each unit is typically allocated, and a linked list of these units is formed to concatenate the payload into larger data units as it is passed to higher layers.

[0072] Figure 11 depicts the typical storage allocation of received packet data units. As depicted in Figure 11, an incoming data frame 1100 may include a plurality of protocol data units (PDUs), each comprising a header (H1, H2, H3) and a payload (P1, P2, P3). Typically, the data frame is partitioned into a plurality of segments, wherein each segment is stored in a separate data service memory (DSM) unit. As shown in Figure

11, PDU1 (H1 + P1) is stored in a first DSM unit 1104, PDU2 (H2 + P2) is stored in a second DSM unit 1106, and PDU3 (H3 + P3) is stored in a third DSM unit 1108.

[0073] Storing data in this manner is inefficient for various reasons. For example, each DSM unit includes its own header H, which adds to overhead. Additionally, each DSM unit within a DSM pool is identically sized. As such, space is wasted when a PDU is smaller than the pre-configured DSM unit size. As depicted at 1104 and 1108 padding data P is added after the PDU to fill the DSM unit. Moreover, to later concatenate the segmented PDUs, linked lists must be maintained that indicated where each PDU is stored and how it relates to the others.

[0074] According to some exemplary aspects of the present apparatus and methods, received data may be aggregated per channel from different transmission into larger blocks of contiguous memory by removing headers and thereby directly concatenating continuous payloads for higher layers. This may result in application data coherency as well as a reduced need for additional memory allocations due to concatenation. Moreover, memory is reduced due to a reduction of padding overhead, memory allocation operations, and memory overhead.

[0075] As described above, when hardware receives a ciphered data frame, the frame is deciphered, and then forwarded to hardware for further processing. According to some exemplary aspects, hardware components, such as hardware accelerators, may be further configured to concatenate related payloads without forwarding the frames to software.

[0076] Figure 12 is a simplified block diagram illustrating data processing and storage, according to some aspects. As depicted in figure 12, an incoming data frame 1201 may include a plurality of protocol data units, each comprising a header and a payload. Partitioning unit 1202 may be configured to process the incoming data frame 1201. Processing may include, for example, separating the headers and payloads. The headers (H1, H2, H3) may be stored in memory 1203, along with a pointer to the associated payload.

[0077] Hardware concatenation logic 1204 may then combine the payloads and store them in a single DSM. Both partitioning unit 1202 and hardware concatenation logic 1204 may be programmed and/or controlled by software logic 1206. For example, software logic 1206 may program the partitioning unit 1202 to remove the headers and to generate pointer information. Additionally, software logic 1206 may direct the concatenation logic 1204 to fetch certain payloads or headers.

[0078] According to some aspects, data from multiple transmissions may be combined in a single DSM. Data may be passed along to other layers once stored in the DSM, even if the DSM is not full. Additionally data may be added to the end of the DSM unit.

[0079] Typically headers and payloads are not separated from each other when data is received. Rather, the header and payload are typically stored together. However, most decisions on how to process a packet are based only on the packet header. According to some aspects of the present apparatus and methods, only the header may be moved from layer to layer. Therefore, the payload may be moved only when necessary for processing, thus improving payload data coherency and cache efficiency, and reducing bus utilization.

[0080] In a typical wireless device, a main processor handles all modem related functionality. Figure 13 is a block diagram depicting a typical configuration. Main processor 1304 is communicatively coupled to hardware 1308 via a plurality of buses 1306. Memory 1302 is provided for storing data operated on by hardware 1308. In operation, whenever main processor 1304 issues an instruction to hardware 1308, the instruction must traverse the plurality of buses 1306. Moreover, the hardware 1308 must issue a software interrupt back to main processor 1304 after completing each instruction. Traversing the buses and issuing software interrupts adds substantial processing latency.

[0081] In accordance with various exemplary aspects, a mini-processor may be provided which is closely located to the hardware, thereby reducing latency. Figure 14 is a block diagram depicting such a configuration. Main processor 1404 is communicatively coupled to hardware 1408 via a plurality of buses 1406. Memory 1402 is provided for storing data processed by hardware 1408.

[0082] Mini-processor 1410 is provided in close proximity to hardware 1408. Mini-processor 1410 may be a flexibly programmable processor having access to the same memory 1402 as the main processor 1404. The mini-processor 1410 may be programmed to interact directly with the hardware 1408. The mini-processor 1410 may be programmed to direct hardware 1408 to perform tasks such as header extraction, ciphering, deciphering, data movement, contiguous storage, IP filtering, header insertion, PPP framing, and/or other tasks.

[0083] According to some aspects, when new information is received at a wireless device, the main processor 1404 may direct the mini-processor 1410 to begin processing the data and to store the results, after processing by hardware 1408, into memory 1402. Thus, hardware no longer has to interrupt back to the main processor 1414 after each task. Additionally, because the mini-processor 1410 is programmable, it can easily be re-programmed if there is a change in processing requirements. For example, if there is a change in the air interface, or if a new protocol release is available, the mini-processor 1410 can be reprogrammed to implement the change.

[0084] In accordance with the configuration depicted in Figure 14, latency is reduced as compared to a pure software implementation. Additionally, there may be a reduction in interrupts, memory access, and bus latency. Moreover, power consumption may be lowered, as the mini-processors may be implemented close to the hardware memory than in a pure software application.

[0085] According to some aspects, both a main processor and a mini-processor may have access to one or more dedicated hardware blocks. This is depicted in Figure 15. A main processor 1502 and mini-processor 1504 are each communicatively coupled to a plurality of dedicated hardware blocks. The plurality of dedicated blocks include a first hardware block 1506, a second hardware block 1508, and a third hardware block 1510. The hardware blocks may include, for example, a hardware accelerator, a data mover, a ciphering engine, a deciphering engine, and/or any other dedicated hardware blocks.

[0086] Main processor 1502 and mini-processor 1504 may each simultaneously access a different one of hardware blocks 1504, 1506, and 1508. This allows for parallel processing. For example, main processor 1502 may be configured to access first hardware block 1506, which may be dedicated to deciphering, while mini-processor 1504 may be configured to simultaneously access second hardware block 1508, which may be dedicated to ciphering. Thus, uplink and downlink processing can occur simultaneously.

[0087] The mini-processor (such as mini-processor 1504 depicted in Fig. 15), may be unable to handle complex tasks. For example, in dealing with retransmission of packets, there is complex logic involved in determining when sequence numbers rollover. According to some aspects, the main processor may be configured to serve as a back-up to the mini-processor and handle those complex tasks which cannot be

properly handled by the mini-processor. Figure 16 is a flowchart depicting a process for handling such complex logic.

[0088] As depicted at 1602, the process begins when the mini-processor receives data to be processed. According to some aspects, the mini-processor may be programmed to always process a certain type of task using a pre-defined processing method. Thus, as depicted at 1604, the mini-processor may be configured to instruct hardware to process data using a first processing method. For example, in the case of determining whether sequence numbers should roll over, the mini processor may be configured to always assume that the sequence numbers do not roll over.

[0089] As depicted at 1606, the main processor receives the data processed by hardware. The main processor may also receive an indication of how the data was processed. As depicted at 1608, the main processor determines whether the data was processed correctly. If so, the main processor simply forwards the data to its destination, such as memory, as depicted at 1610. If, however, the main processor detects that the data was not processed correctly, the main processor can reverse the actions performed by the hardware and re-program the hardware to process the data correctly, as depicted at 1612.

[0090] In typical data moving operations, a processor programs a data mover to perform a copy task for each chunk of data to be moved. According to various exemplary aspects, a copy engine may be included which is closely located to hardware. Figure 17 is a block diagram of a system 1700 incorporating a copy engine. System 1700 comprises data mover 1702, main memory 1704, copy engine 1710, copy engine memory 1706, and hardware 1708. Copy engine 1710 may be programmed to operate directly on hardware 1708 to facilitate copy and/or data move operations. Copy engine 1710 may have associated therewith its own dedicated memory 1706 which reduces latency associated with storing and retrieving data from main memory 1704. Data mover 1702 may be programmed to fetch data from the copy engine memory 1706 and store it in main memory 1704.

[0091] Including a copy engine allows for bit-level granularity to support protocols with bit-level widths. Additionally, programming overhead may be reduced by allowing the copying of evenly scattered source data to evenly scattered destination locations using a single programming task. Moreover, the copy engine may be used for any type of operation used by software, for example, header extraction and insertions, data

concatenation or segmentation, byte/word alignment, as well as regular data-mover tasks for data manipulation.

[0092] According to some aspects, copy engine 1710 may be programmed to copy data from multiple PDUs in a single tasks. As depicted in Figure 18, a data frame may include a plurality of PDUs, each comprising a header (H1, H2, and H3), and a payload (P1, P2, P3). The size of the headers and payloads may be known *a priori*, and a copy engine may be programmed to know these sizes. For example, as depicted in Figure 18, all headers, such as header 1802 may be of size X, while all payloads, such as payload 1804, may be of size Y. A copy engine may be programmed to copy N headers or N payloads in a single task based on knowledge of the header and payload sizes.

[0093] The various illustrative logics, logical blocks, modules, and circuits described in connection with the embodiments disclosed herein may be implemented or performed with a general purpose processor, a digital signal processor (DSP), an application specific integrated circuit (ASIC), a field programmable gate array (FPGA) or other programmable logic device, discrete gate or transistor logic, discrete hardware components, or any combination thereof designed to perform the functions described herein. A general-purpose processor may be a microprocessor, but, in the alternative, the processor may be any conventional processor, controller, microcontroller, or state machine. A processor may also be implemented as a combination of computing devices, e.g., a combination of a DSP and a microprocessor, a plurality of microprocessors, one or more microprocessors in conjunction with a DSP core, or any other such configuration. Additionally, at least one processor may comprise one or more modules operable to perform one or more of the steps and/or actions described above.

[0094] Further, the steps and/or actions of a method or algorithm described in connection with the aspects disclosed herein may be embodied directly in hardware, in a software module executed by a processor, or in a combination of the two. A software module may reside in RAM memory, flash memory, ROM memory, EPROM memory, EEPROM memory, registers, a hard disk, a removable disk, a CD-ROM, or any other form of storage medium known in the art. An exemplary storage medium may be coupled to the processor, such that the processor can read information from, and write information to, the storage medium. In the alternative, the storage medium may be integral to the processor. Further, in some aspects, the processor and the storage

medium may reside in an ASIC. Additionally, the ASIC may reside in a user terminal. In the alternative, the processor and the storage medium may reside as discrete components in a user terminal. Additionally, in some aspects, the steps and/or actions of a method or algorithm may reside as one or any combination or set of codes and/or instructions on a machine readable medium and/or computer readable medium, which may be incorporated into a computer program product.

[0095] In one or more aspects, the functions described may be implemented in hardware, software, firmware, or any combination thereof. If implemented in software, the functions may be stored or transmitted as one or more instructions or code on a computer-readable medium. Computer-readable media includes both computer storage media and communication media including any medium that facilitates transfer of a computer program from one place to another. A storage medium may be any available media that can be accessed by a computer. By way of example, and not limitation, such computer-readable media can comprise RAM, ROM, EEPROM, CD-ROM or other optical disk storage, magnetic disk storage or other magnetic storage devices, or any other medium that can be used to carry or store desired program code in the form of instructions or data structures and that can be accessed by a computer. Also, any connection may be termed a computer-readable medium. For example, if software is transmitted from a website, server, or other remote source using a coaxial cable, fiber optic cable, twisted pair, digital subscriber line (DSL), or wireless technologies such as infrared, radio, and microwave, then the coaxial cable, fiber optic cable, twisted pair, DSL, or wireless technologies such as infrared, radio, and microwave are included in the definition of medium. Disk and disc, as used herein, includes compact disc (CD), laser disc, optical disc, digital versatile disc (DVD), floppy disk and blu-ray disc where disks usually reproduce data magnetically, while discs usually reproduce data optically with lasers. Combinations of the above should also be included within the scope of computer-readable media.

[0096] While the foregoing disclosure discusses illustrative aspects and/or embodiments, it should be noted that various changes and modifications could be made herein without departing from the scope of the described aspects and/or embodiments as defined by the appended claims. Furthermore, although elements of the described aspects and/or embodiments may be described or claimed in the singular, the plural is contemplated unless limitation to the singular is explicitly stated. Additionally, all or a

portion of any aspect and/or embodiment may be utilized with all or a portion of any other aspect and/or embodiment, unless stated otherwise.



**CLAIMS****WHAT IS CLAIMED IS:**

1. A data processing method, comprising:
  - receiving, by a first hardware component having associated therewith a first task queue, instructions to execute a first set of tasks;
  - receiving, by a second hardware component having associated therewith a second task queue, instructions to execute a second set of tasks;
  - receiving, at the first hardware component, a trigger message to execute the first set of tasks; and
  - upon completion of the first set of tasks, transmitting, by the first hardware component, a trigger message to the second hardware component to execute the second set of tasks.
2. The method of claim 1, wherein the first set of tasks includes a task of reading a block of data from a first memory and storing the block of data in a second memory.
3. The method of claim 2, wherein the second set of tasks includes a task of processing the block of data stored in the second memory.
4. The method of claim 1, wherein the first hardware component triggers the second hardware component by moving a write pointer associated with the second task queue.
5. The method of claim 1, further comprising:
  - transmitting, by the second hardware component, a trigger message to the first hardware component after completing the second set of tasks;
  - determining, by the first hardware components, whether any further tasks are in the first task queue; and
  - if no further tasks are in the first task queue, transmitting a software interrupt to a software component.

6. The method of claim 3, further comprising:  
determining, by the second hardware component, when a length of the processed block of data exceeds a threshold length, wherein the threshold length is less than a maximum possible length of the processed data; and  
generating, by the second hardware component, a software interrupt indicating that the length of the processed block of data exceeds the threshold length.

7. The method of claim 6, further comprising:  
moving, by a software component, a portion of the block of data exceeding the threshold length;  
triggering, by the software component, the first hardware component to move the remaining portions of the block of data.

8. The method of claim 6, wherein the threshold length is dynamically adjusted depending on the frequency of software interrupts.

9. The method of claim 3, further comprising:  
providing a traffic template to the second hardware component, wherein the traffic template comprise a set of conditions for the second hardware component to take action.

10. The method of claim 9, wherein the second hardware component compares the block of data with the traffic template.

11. The method of claim 10, wherein the second hardware component acts on the block of data when the block of data matches one of the sets of conditions in the traffic template.

12. The method of claim 10, wherein the second hardware component sends the block of data to a software module when the block of data does not match any of the set of conditions in the traffic template.

13. An data processing apparatus, comprising:

a processor;  
a first memory;  
a second memory;  
a first hardware component having associated therewith a first task queue; and  
a second hardware component having associated therewith a second task queue,  
wherein a plurality of tasks are pre-stored in the first task queue, a plurality of  
second task are pre-stored in the second task queue, and the first hardware component is  
configured to receive a trigger to begin performing the plurality of first tasks, and to  
issue a trigger to the second hardware component to being performing the plurality of  
second tasks.

14. The data processing apparatus of claim 13, wherein at least one of the  
plurality of first tasks is a task of reading a block of data from a first memory and  
storing the block of data in the second memory.

15. The data processing apparatus of claim 14, wherein at least one of the  
plurality of second tasks is a task of processing the block of data stored in the second  
memory.

16. The data processing apparatus of claim 13, wherein the first hardware  
component is a data mover.

17. The data processing apparatus of claim 13, wherein the second hardware  
component is a hardware accelerator.

18. The data processing apparatus of claim 15, wherein the second hardware  
component comprises a comparator for comparing a length of the processed block of  
data to a threshold length, and

wherein the second hardware component generates a software interrupt when the  
length of the processed block of data is greater than the threshold length.

19. The data processing apparatus of claim 18, wherein the threshold length  
is dynamically adjusted depending on the frequency of software interrupts.

20. The data processing apparatus of claim 15, further comprising a traffic template comprising a set of conditions, wherein

the second hardware component compares the block of data with the traffic template.

21. At least on processor, comprising:

a first module for receiving instructions to execute a first set of tasks by a first hardware component;

a second module for receiving instructions to execute a second set of tasks by a second hardware component;

a third module for receiving a trigger message to execute the first set of tasks; and

a fourth module for issuing, from the first hardware component, a trigger message to the second hardware component to execute the second set of tasks.

22. A computer program recorded on a computer-readable medium, comprising:

receiving instructions to execute a first set of tasks by a first hardware component;

receiving instructions to execute a second set of tasks by a second hardware component;

receiving a trigger message to execute the first set of tasks; and

issuing, from the first hardware component, a trigger message to the second hardware component to execute the second set of tasks.

23. An apparatus, comprising:

means for receiving, by a first hardware component having associated therewith a first task queue, instructions to execute a first set of tasks;

means for receiving, by a second hardware component having associated therewith a second task queue, instructions to execute a second set of tasks;

means for receiving, at the first hardware component, a trigger message to execute the first set of tasks; and

upon completion of the first set of tasks, means for transmitting, by the first hardware component, a trigger message to the second hardware component to execute the second set of tasks.

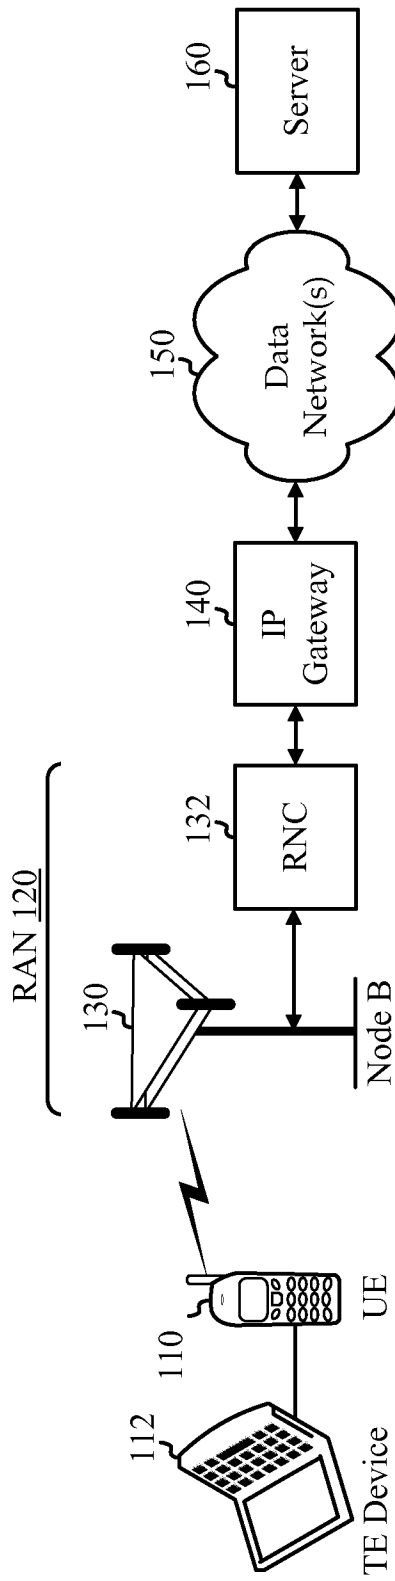


FIG. 1

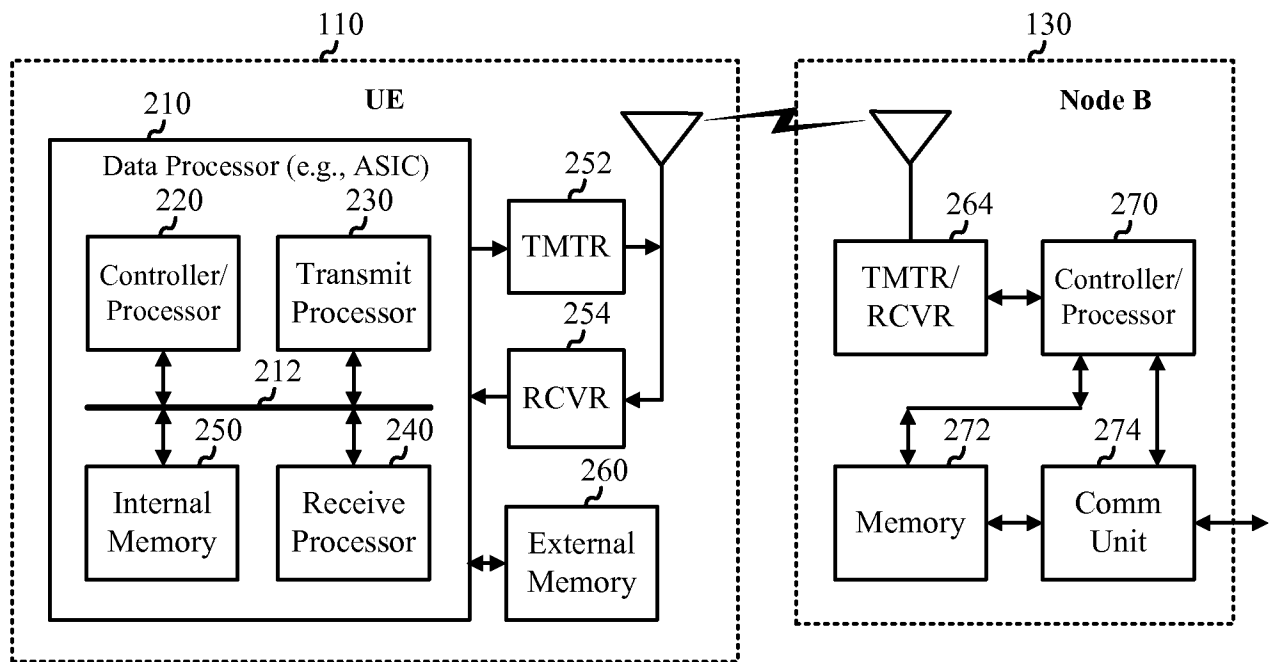


FIG. 2

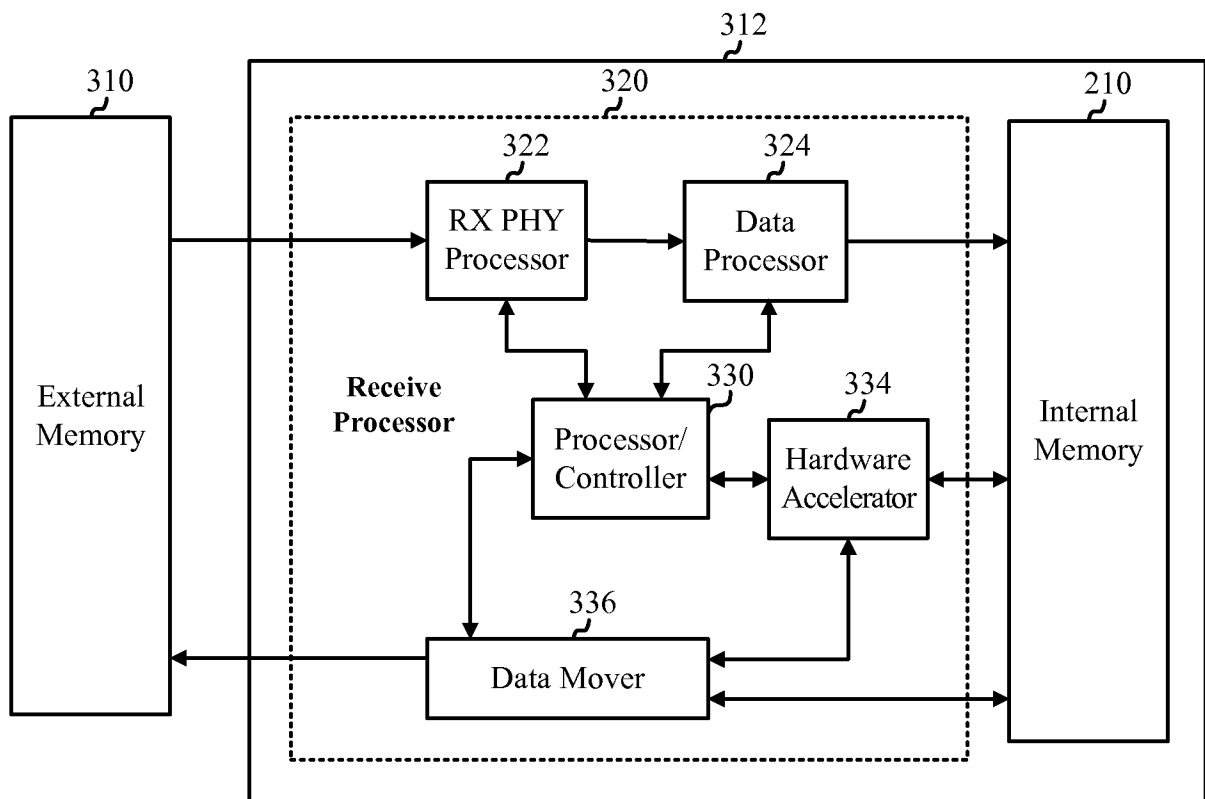
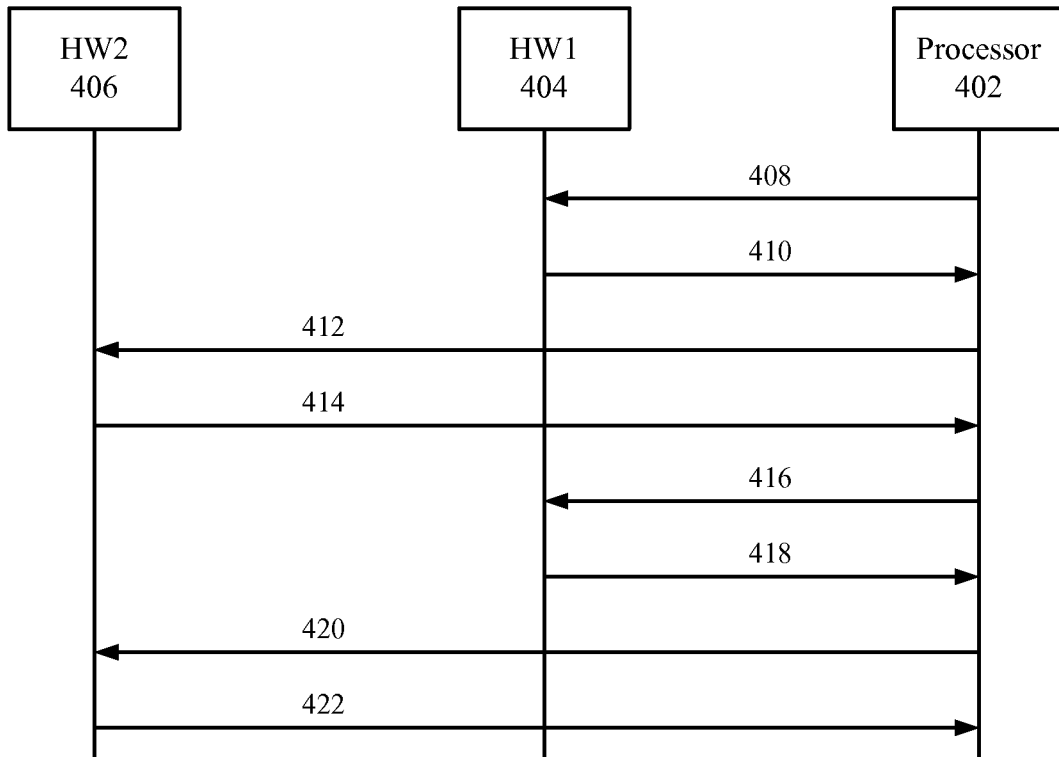


FIG. 3





**FIG. 4**  
**(PRIOR ART)**




FIG. 5

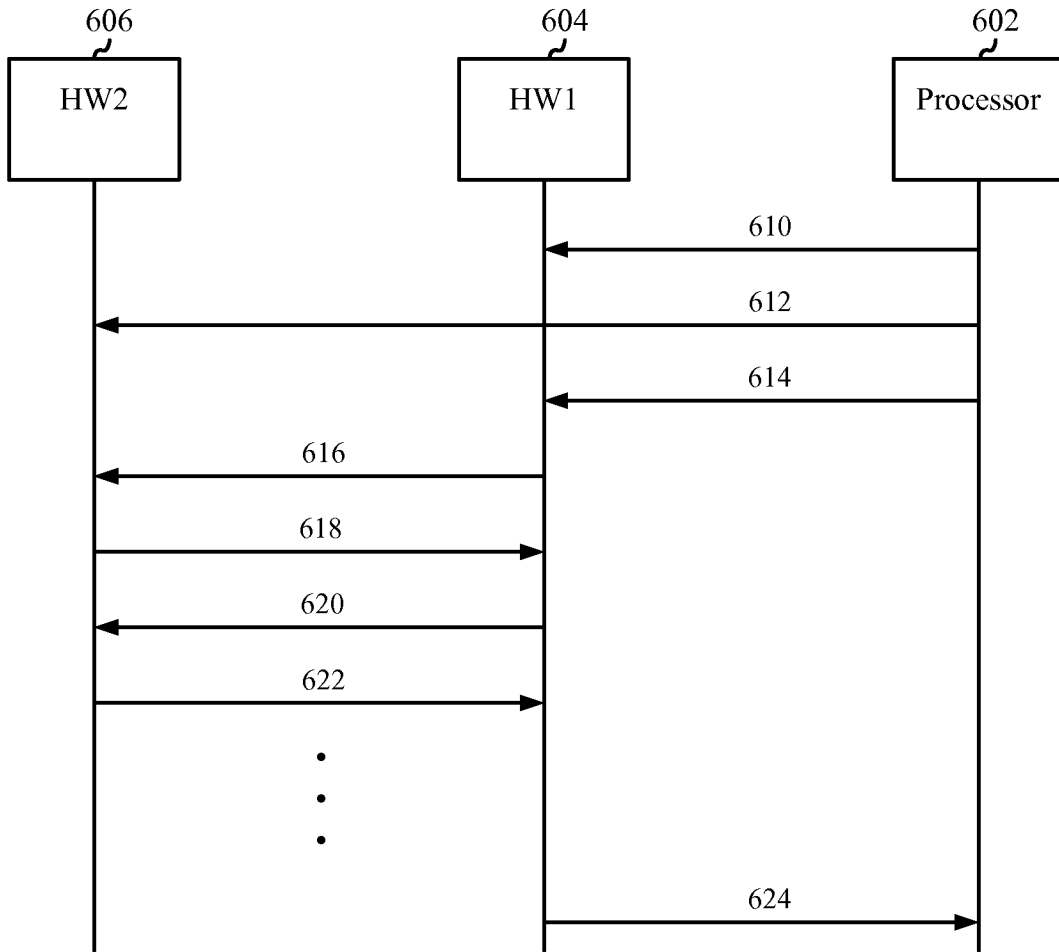


FIG. 6

7/18

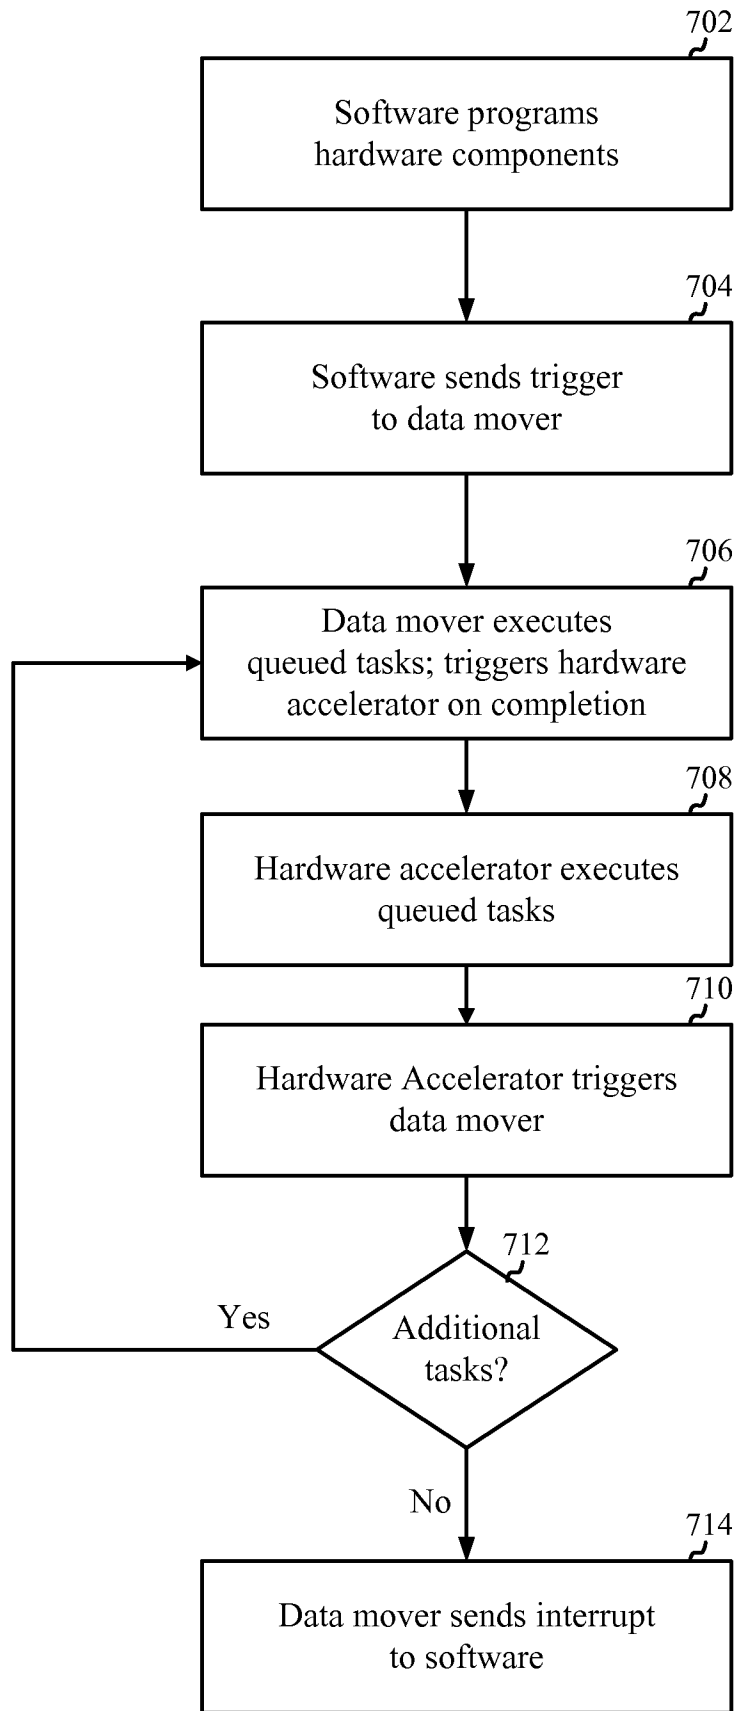


FIG. 7

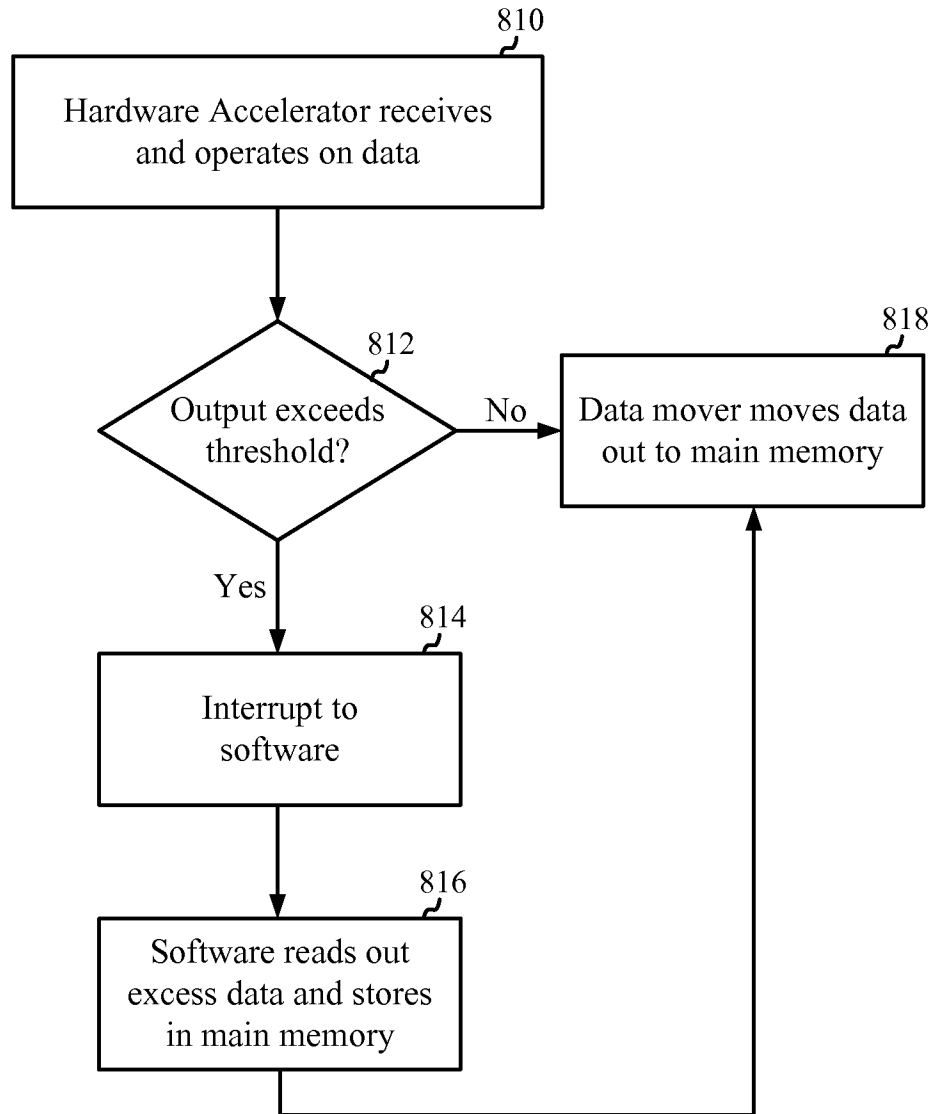


FIG. 8




FIG. 9

10/18

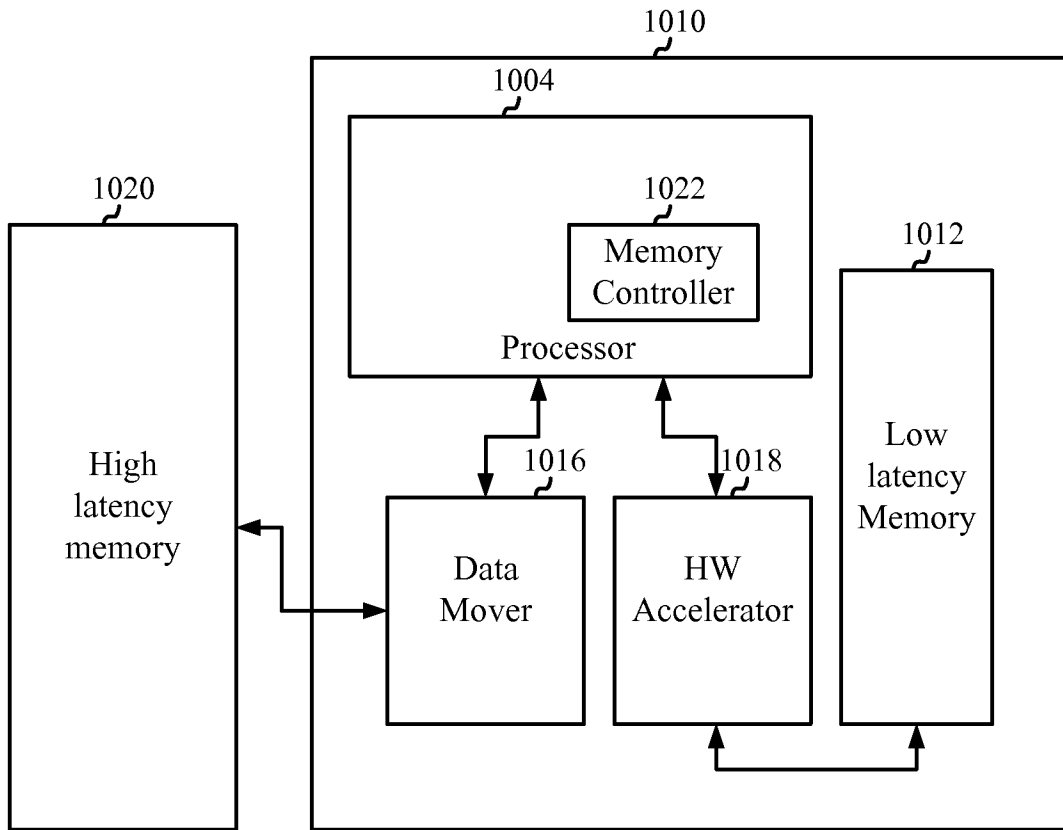


FIG. 10




FIG. 11






FIG. 12

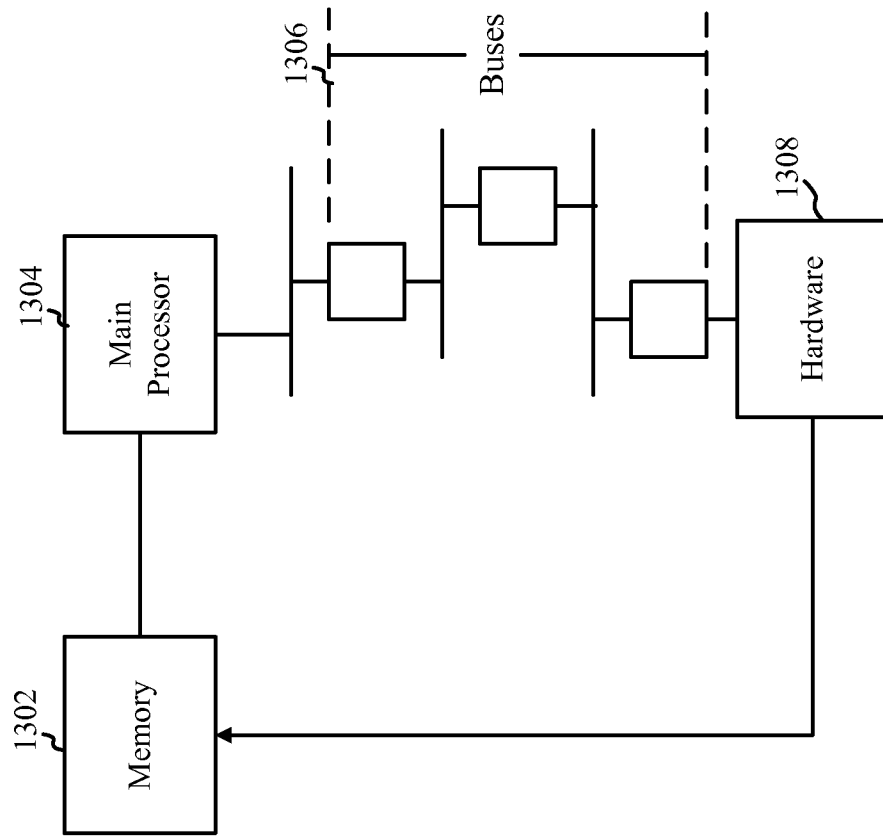


FIG. 13

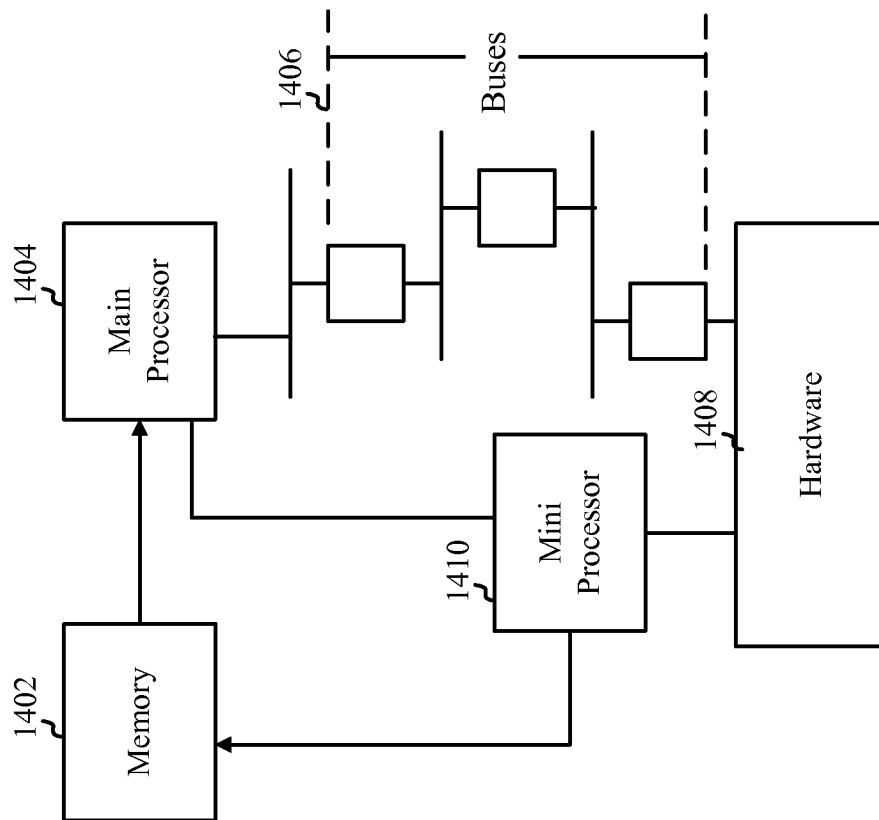


FIG. 14




FIG. 15

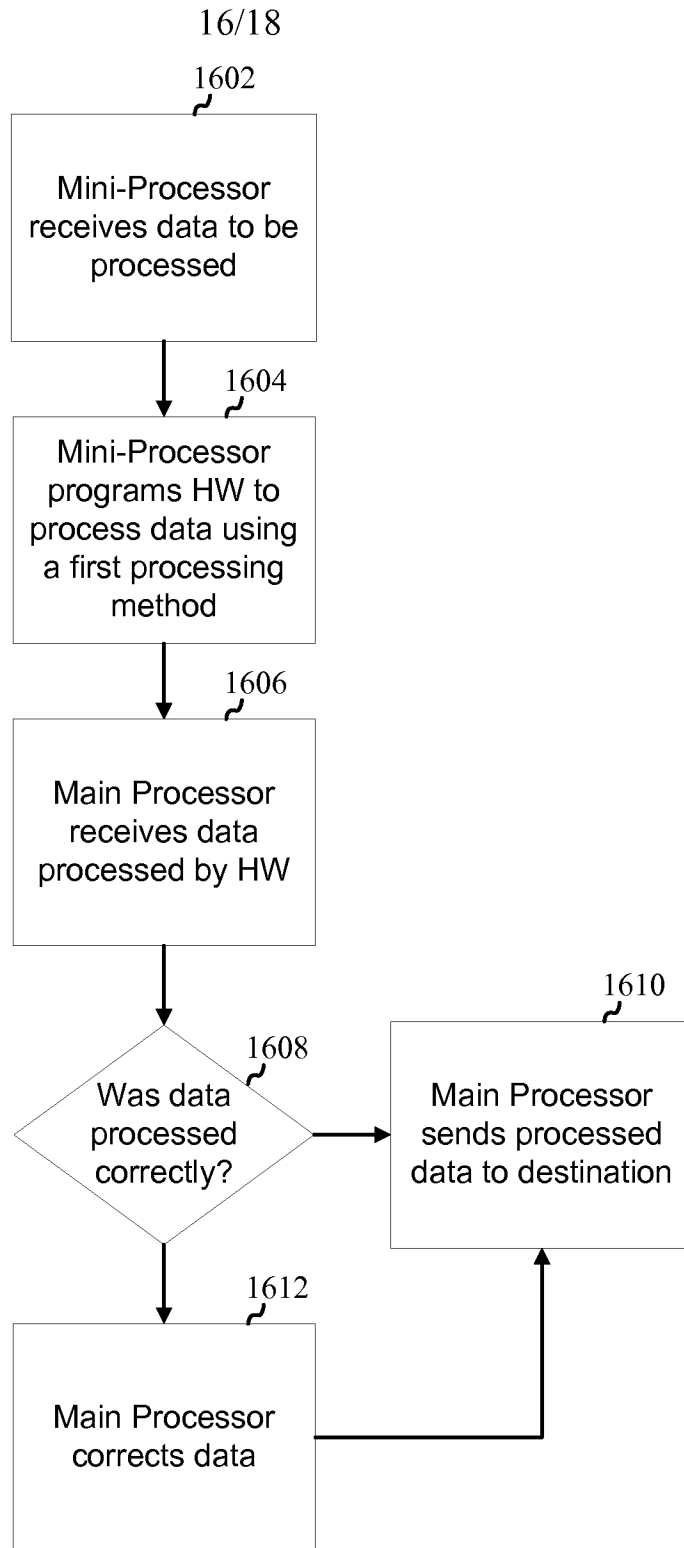


FIG. 16




FIG. 17




FIG. 18

# INTERNATIONAL SEARCH REPORT

International application No  
PCT/US2010/045517

<b>A. CLASSIFICATION OF SUBJECT MATTER</b> INV. G06F9/48 ADD.		
According to International Patent Classification (IPC) or to both national classification and IPC		
<b>B. FIELDS SEARCHED</b>		
Minimum documentation searched (classification system followed by classification symbols) G06F		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched		
Electronic data base consulted during the international search (name of data base and, where practical, search terms used) EPO-Internal, WPI Data, PAJ, INSPEC, COMPENDEX		
<b>C. DOCUMENTS CONSIDERED TO BE RELEVANT</b>		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X A A A A	US 4 980 824 A (TULPULÉ BHALCHANDRA R [US] ET AL) 25 December 1990 (1990-12-25) the whole document ----- US 5 832 262 A (JOHNSON CHRISTOPHER T [US] ET AL) 3 November 1998 (1998-11-03) the whole document ----- US 2007/220517 A1 (LIPPETT MARK D [GB]) 20 September 2007 (2007-09-20) the whole document ----- WO 2008/046716 A1 (IBM [US]; IBM UK [GB]; JU HAI [CN]; LIN GUO HUI [CN]; LIU QIANG [CN];) 24 April 2008 (2008-04-24) the whole document ----- -/--	1-4,6-23  5  1-23  1-23  1-23
<input checked="" type="checkbox"/> Further documents are listed in the continuation of Box C. <span style="margin-left: 200px;"><input checked="" type="checkbox"/> See patent family annex.</span>		
* Special categories of cited documents :		
*A* document defining the general state of the art which is not considered to be of particular relevance *E* earlier document 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  <p style="text-align: center;">21 October 2010</p>	Date of mailing of the international search report  <p style="text-align: center;">28/10/2010</p>	
Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer  <p style="text-align: center;">Beltrán-Escavy, José</p>	



# INTERNATIONAL SEARCH REPORT

International application No  
PCT/US2010/045517

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 7 558 929 B2 (ZOHAR OFIR [IL] ET AL) 7 July 2009 (2009-07-07) the whole document -----	1-23

**INTERNATIONAL SEARCH REPORT**

Information on patent family members

International application No PCT/US2010/045517
---

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 4980824	A	25-12-1990	NONE	
US 5832262	A	03-11-1998	NONE	
US 2007220517	A1	20-09-2007	CN 1955931 A EP 1770509 A2 JP 2007133858 A KR 20070037427 A US 2007220294 A1	02-05-2007 04-04-2007 31-05-2007 04-04-2007 20-09-2007
WO 2008046716	A1	24-04-2008	CN 101165655 A	23-04-2008
US 7558929	B2	07-07-2009	US 2007068061 A1	29-03-2007