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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
3 July 2014 (03.07.2014)

WIPO | PCT

(10) International Publication Number
WO 2014/105177 A1

(51) International Patent Classification:
G06F 1/32 (2006.01)

(21) International Application Number:
PCT/US2013/047778

(22) International Filing Date:
26 June 2013 (26.06.2013)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
13/730,074 28 December 2012 (28.12.2012) US

(71) Applicant *(for all designated States except US)*: **INTEL CORPORATION** [US/US]; 2200 Mission College Boulevard, Santa Clara, California 95054 (US).

(72) Inventors; and

(71) Applicants *(for US only)*: **GORBATOV, Eugene** [US/US]; 5802 NE Farnham Street, Hillsboro, Oregon 97124 (US). **DIEFENBAUGH, Paul, S.** [US/US]; 9016 NW Bartholomew Dr., Portland, Oregon 97229 (US). **CRAWFORD, John H.** [US/US]; 20128 Chateau Drive, Saratoga, California 95070 (US). **KUMAR, Anil K.** [US/US]; 6312 W Linda Lane, Chandler, Arizona 85226 (US). **GRECO, Richard J.** [US/US]; 1360 SW Turner Road, West Linn, Oregon 97068 (US).

(74) Agent: **JORDAN, B. Delano**; Jordan IP Law, LLC, c/o CPA Global, P.O. Box 52050, Minneapolis, Minnesota 55402 (US).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, 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, 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, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, 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, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, 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))*

(54) Title: PERIODIC ACTIVITY ALIGNMENT

FIG. 1

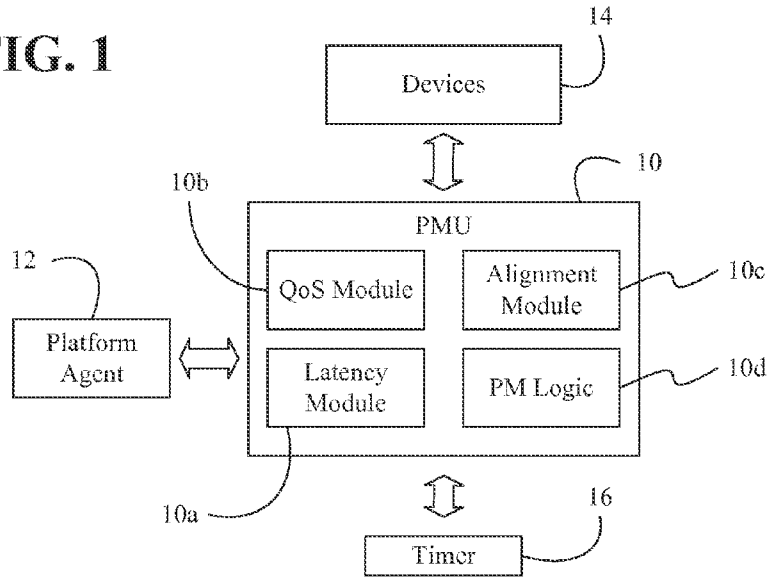

(57) Abstract: Methods and systems may provide for determining a latency constraint associated with a platform and determine an idle window based on the latency constraint. In addition, a plurality of devices on the platform may be instructed to cease one or more activities during the idle window. In one example, the platform is placed in a sleep state during the idle window.

WO 2014/105177 A1

PERIODIC ACTIVITY ALIGNMENT

BACKGROUND

Technical Field

Embodiments generally relate to power management in computing platforms. More particularly, embodiments relate to periodic activity alignment to enhance power efficiency in computing platforms.

Discussion

Computing platforms may enter various sleep states during periods of idleness in order to reduce power consumption. Sleep state entry may be effectively disabled, however, during semi-active workloads due to key elements of the platform such as processor/graphics cores, main memory, system interconnects, etc., becoming even lightly active. Accordingly, conventional platforms may experience a steep drop in energy efficiency when processing typical semi-active workloads.

BRIEF DESCRIPTION OF THE DRAWINGS

The various advantages of the embodiments of the present invention will become apparent to one skilled in the art by reading the following specification and appended claims, and by referencing the following drawings, in which:

FIG. 1 is a block diagram of an example of a power management unit according to an embodiment;

FIG. 2 is a flowchart of an example of a method of managing power in a computing platform according to an embodiment;

FIGs. 3A and 3B are timelines of examples of idle windows according to embodiments; and

FIG. 4 is a block diagram of an example of a platform according to an embodiment.

DETAILED DESCRIPTION

FIG. 1 shows a power management unit (PMU) 10 (10a-10d) that may be used to achieve greater energy efficiency in a computing platform such as, for example, a mobile device having computing functionality (e.g., personal digital assistant/PDA, laptop, smart tablet), communications functionality (e.g., wireless smart phone), imaging functionality,

media playing functionality (e.g., smart television/TV), or any combination thereof (e.g., mobile Internet device/MID). The illustrated PMU 10 may also be used in a fixed platform such as, for example, a server, desktop computer, workstation, and so forth.

In the illustrated example, a quality of service (QoS) module 10a (e.g., latency module) determines latency constraints associated with the platform, wherein the latency constraints may be determined based on service requests (e.g., QoS requests, service level agreement/SLA requests), performance level requests (e.g., ACPI/Advanced Configuration and Power Interface performance state requests, e.g., ACPI Specification, Rev. 5.0a, December 6, 2011), and so forth. The latency-related requests may be made by, for example, a platform agent 12 associated with one or more applications and/or an operating system (OS) executing on the platform. As will be discussed in greater detail, the latency constraints may generally establish maximum latencies that are tolerated by a workload of the platform. For example, an SLA for a server workload might stipulate that the workload be completed within x milliseconds, a particular ACPI performance state (P-state) may be associated with a particular maximum latency (e.g., via lookup table), and so forth. The requests may be issued directly to the PMU 10 or indirectly to the PMU 10 by, for example, changing one or more register values, configuration settings, and so forth.

The PMU 10 may also include an idleness module 10b that determines idle windows based on the latency constraints. In one example, idle windows are a function of the latency tolerance and the predicted idleness (e.g., idle window = predicted idleness + latency tolerance). The PMU 10 may also include an alignment module 10c that instructs a plurality of devices 14 to cease one or more activities (e.g., external communication) during the idle window. The devices 14 may include, for example, input output (IO) devices, processors, system controllers, etc., that prevent the platform from entering sleep states when one or more of the devices 14 are issuing interrupts, DMA (direct memory access) requests, fabric access requests, etc., (e.g., active). The PMU 10 may therefore force an alignment of all devices on a platform and create a platform wide idle window. In one example, the PMU 10 uses a timer 16 to determine when the idle window has closed/expired, wherein the alignment module 10c may initiate active windows in response to expirations of the timer 16. The illustrated PMU 10 also includes power management (PM) logic 10d that places the platform in sleep states during the idle windows. Of particular note is that by forcing the devices 14 to cease communication/activity, the alignment module 10c can create platform wide idle windows rather than merely extend idle windows that may (or may not) occur naturally. Moreover, the sleep states selected by

the PM logic 10d may be deeper, more frequent and longer lasting using the illustrated approach, which may in turn yield greater energy efficiency, longer battery life and enhanced performance.

FIG. 2 shows a method 20 of managing power in a computing platform. The method 20 may be implemented as a set of logic instructions stored in a machine- or computer-readable storage medium such as random access memory (RAM), read only memory (ROM), programmable ROM (PROM), firmware, flash memory, etc., in configurable logic such as, for example, programmable logic arrays (PLAs), field programmable gate arrays (FPGAs), complex programmable logic devices (CPLDs), in fixed-functionality logic hardware using circuit technology such as, for example, application specific integrated circuit (ASIC), complementary metal oxide semiconductor (CMOS) or transistor-transistor logic (TTL) technology, or any combination thereof. For example, computer program code to carry out operations shown in method 20 may be written in any combination of one or more programming languages, including an object oriented programming language such as Java, Smalltalk, C++ or the like and conventional procedural programming languages, such as the "C" programming language or similar programming languages.

Illustrated processing block 22 provides for determining a latency constraint associated with a platform. As already noted, the latency constraint may be determined based on a service request, a performance level request, etc., wherein the latency constraint may define a maximum latency for the platform and/or a workload being executed on the platform. Block 24 may determine an idle window based on the latency constraint, and a plurality of devices on the platform may be instructed at block 26 to cease one or more activities during the idle window. Ceasing activity may involve buffering data such as DMA requests and interrupts, so that no IO activity will occur, entering a device sleep state, and so forth. Accordingly, the illustrated approach provides for forced alignment of all devices/controllers/ blocks in the platform to create system wide idle windows. Such an approach can provide significant advantages over conventional solutions that merely wait for idle windows to occur naturally.

With specific regard to the requests that are delayed during the idle window, an IO request may originate upstream (e.g., from a device) or downstream (e.g., from the OS/software) from the perspective of the power management unit. Both could be delayed during the idle window, but the latency tolerance for each may be treated separately as another optimization. For example, an IO device may have a much tighter latency constraint for accessing main memory than software has for receiving interrupts.

Block 26 may also involve setting a timer such as the timer 16 (FIG. 1), already discussed. Illustrated block 28 processes any pending requests from the platform devices and places the platform and/or processor in a sleep state such as an ACPI low power state, wherein a determination may be made at block 30 as to whether the idle window is closed/over. If so, an active window is initiated at illustrated block 32 and the platform/processor may be brought out of the sleep state. Initiating the active window may include instructing the devices to resume activity, wherein any devices having data and/or instructions, transactions, messages, DMA requests, interrupts, etc., for processing, may issue the appropriate signals to the platform to conduct the activity in question.

In addition, portions of the method 20 may be repeated for relatively long periods of time until conditions change. Moreover, the method 20 may also be disabled if it is determined that the platform utilization is above a certain threshold (e.g., above 20%, low core C6 utilization, etc.). For example, the platform utilization is above a certain percentage, it may be inferred that forced creation of platform wide idle windows will not be productive. Similarly, if is determined that a relatively large number of cores have not been entering deep sleep states, the illustrated approach might be bypassed.

FIG. 3A shows a timeline 34 in which idle windows 36 (Δt_I) are created by instructing platform devices to cease one or more activities. In the illustrated example, a latency constraint window (Δt_L) is determined, wherein a timer arming 40 may be conducted to mark the beginning of the latency constraint window. One or more pending requests 42 from the platform devices may be processed before placing the platform and/or processor in a sleep state during the idle windows 36. Timer expirations 44 may take place to mark the end of the idle windows 36, wherein the platform/processor may be brought out of the sleep state to process information during an active window occurring between the latency constraint windows.

FIG. 3B shows another timeline 35 in which an idle window (Δt_I) is created by instructing platform devices to cease one or more activities. In the illustrated example, at some time t_0 after a final IO request 41 and during the idle window, an IO request is acknowledged but the corresponding IO transfer is delayed because the IO request occurs during the idle window. Also at time t_0, a timer arming 37 may be conducted to mark the beginning of the latency constraint window (Δt_L), wherein a timer expiration 39 can later take place to mark the end of the idle window. The illustrated approach therefore enables allows the idle window to consume the entire latency tolerance. Simply put, the idle window may be longer or shorter than the latency constraint window. If the idle window is

longer than the latency constraint window, the latency tolerance may still be observed – the system just doesn't begin counting until the first access is attempted in such a scenario.

Turning now to FIG. 4, a mobile platform 46 is shown. The platform 46 may be part of a device having computing functionality (e.g., server, workstation, desktop computer, PDA, notebook computer, smart tablet), communications functionality (e.g., wireless smart phone), imaging functionality, media playing functionality (e.g., smart TV), or any combination thereof (e.g., MID). In the illustrated example, the platform 46 includes a processor 48, system memory 50, a network controller 52, an audio IO device 54, a solid state disk (SSD) 56, and an IO module 58. The processor 48 may include an integrated memory controller (IMC) 60, a power management unit (PMU) 62 and a core region with one or several processor cores 64. The processor 48 and the IO module 58 may alternatively be implemented on the semiconductor die as a system on chip (SoC), depending upon the circumstances. Additionally, the platform 46 may also include other components such as dedicated graphics components (not shown), and so forth.

The illustrated IO module 58, sometimes referred to as a Southbridge or South Complex of a chipset, functions as a host controller and communicates with the network controller 52, which could provide off-platform communication functionality for a wide variety of purposes such as, for example, cellular telephone (e.g., Wideband Code Division Multiple Access/W-CDMA (Universal Mobile Telecommunications System/UMTS), CDMA2000 (IS-856/IS-2000), etc.), WiFi (Wireless Fidelity, e.g., Institute of Electrical and Electronics Engineers/IEEE 802.11-2007, Wireless Local Area Network/LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications), 4G LTE (Fourth Generation Long Term Evolution), Bluetooth (e.g., IEEE 802.15.1-2005, Wireless Personal Area Networks), WiMax (e.g., IEEE 802.16-2004, LAN/MAN Broadband Wireless LANS), Global Positioning System (GPS), spread spectrum (e.g., 900 MHz), and other radio frequency (RF) telephony purposes. The IO module 58 may also include one or more wireless hardware circuit blocks to support such functionality.

The system memory 50 may include, for example, double data rate (DDR) synchronous dynamic random access memory (SDRAM, e.g., DDR3 SDRAM JEDEC Standard JESD79-3C, April 2008) modules. The modules of the system memory 50 may be incorporated into a single inline memory module (SIMM), dual inline memory module (DIMM), small outline DIMM (SODIMM), and so forth. The SSD 56 may include one or more NAND (negated AND) chips and might be used to provide high capacity data storage and/or a significant amount of parallelism. There may also be solutions that include NAND

controllers implemented as separate ASIC controllers being connected to the IO module 58 on standard buses such as a Serial ATA (SATA, e.g., SATA Rev. 3.0 Specification, May 27, 2009, SATA International Organization/SATA-IO) bus, or a PCI Express Graphics (PEG, e.g., Peripheral Components Interconnect/PCI Express x16 Graphics 150W-ATX Specification 1.0, PCI Special Interest Group) bus. The SSD 56 could also be used as a USB (Universal Serial Bus, e.g., USB Specification 3.0, USB Implementers Forum) flash storage device.

The illustrated cores 64 of the processor 48, system memory 50, network controller 52, IO module 58, audio IO device 54 and SSD 56 may therefore be considered devices of the platform 46, wherein the PMU 62 may generally have functionality similar to that of the PMU 10 (FIG. 1), already discussed. Accordingly, the PMU 62 may be configured to determine latency constraints associated with the platform 46, determine idle windows based on the latency constraints, and instruct the cores 64, system memory 50, network controller 52, audio IO device 54, SSD 56 and/or IO module 58 to cease one or more activities during the idle windows. Moreover, the PMU 62 may place the platform 46 and/or processor 48 in sleep states during the idle windows, wherein the sleep states may be deeper, more frequent and longer lasting than under conventional solutions. The platform devices themselves may also enter deeper sleep states, depending upon the circumstances.

Thus, techniques described herein may provide significant benefits to both client and server systems. For example, semi-active workloads such as web browsing and video conferencing may be executed on mobile platforms with much greater energy efficiency and battery life. In servers, the opportunity to enter platform idle states may be preserved even as the number of cores increases. For example, in dual-processor server systems with twelve cores (i.e., twenty-four threads per socket), package C-states and platform idle states (e.g., ACPI) may still be used by transforming long “dribbles” of activity into short bursts of activity followed by quasi-deterministic idle windows/periods. Moreover, QoS constraints and PLA commitments may be met in a manner that is transparent to the OS and applications. Indeed, energy efficiency savings may be two to three times greater for semi-active workloads under the techniques described herein.

Embodiments may therefore include a method that provides for determining a latency constraint associated with a platform and determining an idle window based on the latency constraint. Additionally, a plurality of devices on the platform may be instructed to cease one or more activities during the idle window.

Embodiments may also include a non-transitory computer readable storage medium having a set of instructions which, if executed by a processor, cause a platform to determine a latency constraint associated with a platform. The instructions, if executed, may also cause the platform to determine an idle window based on the latency constraint and instruct a plurality of devices on the platform to cease one or more activities during the idle window.

Embodiments may also include an apparatus having a latency module to determine a latency constraint associated with a platform and an idleness module to determine an idle window based on the latency constraint. The apparatus may also have an alignment module to instruct a plurality of devices on the platform to cease one or more activities during the idle window.

Embodiments may also include a platform having a plurality of devices, wherein the plurality of devices includes one or more of input output (IO) devices and system controllers. The platform may also include a power management unit having a latency module to determine a latency constraint associated with the platform, an idleness module to determine an idle window based on the latency constraint, and an alignment module to instruct the plurality of devices to cease one or more activities during the idle window.

Embodiments of the present invention are applicable for use with all types of semiconductor integrated circuit ("IC") chips. Examples of these IC chips include but are not limited to processors, controllers, chipset components, PLAs, memory chips, network chips, SoCs, SSD/NAND controller ASICs, and the like. In addition, in some of the drawings, signal conductor lines are represented with lines. Some may be different, to indicate more constituent signal paths, have a number label, to indicate a number of constituent signal paths, and/or have arrows at one or more ends, to indicate primary information flow direction. This, however, should not be construed in a limiting manner. Rather, such added detail may be used in connection with one or more exemplary embodiments to facilitate easier understanding of a circuit. Any represented signal lines, whether or not having additional information, may actually comprise one or more signals that may travel in multiple directions and may be implemented with any suitable type of signal scheme, e.g., digital or analog lines implemented with differential pairs, optical fiber lines, and/or single-ended lines.

Example sizes/models/values/ranges may have been given, although embodiments of the present invention are not limited to the same. As manufacturing techniques (e.g., photolithography) mature over time, it is expected that devices of smaller size could be

manufactured. In addition, well known power/ground connections to IC chips and other components may or may not be shown within the figures, for simplicity of illustration and discussion, and so as not to obscure certain aspects of the embodiments of the invention. Further, arrangements may be shown in block diagram form in order to avoid obscuring embodiments of the invention, and also in view of the fact that specifics with respect to implementation of such block diagram arrangements are highly dependent upon the platform within which the embodiment is to be implemented, i.e., such specifics should be well within purview of one skilled in the art. Where specific details (e.g., circuits) are set forth in order to describe example embodiments of the invention, it should be apparent to one skilled in the art that embodiments of the invention can be practiced without, or with variation of, these specific details. The description is thus to be regarded as illustrative instead of limiting.

The term "coupled" may be used herein to refer to any type of relationship, direct or indirect, between the components in question, and may apply to electrical, mechanical, fluid, optical, electromagnetic, electromechanical or other connections. In addition, the terms "first", "second", etc. are used herein only to facilitate discussion, and carry no particular temporal or chronological significance unless otherwise indicated.

Those skilled in the art will appreciate from the foregoing description that the broad techniques of the embodiments of the present invention can be implemented in a variety of forms. Therefore, while the embodiments of this invention have been described in connection with particular examples thereof, the true scope of the embodiments of the invention should not be so limited since other modifications will become apparent to the skilled practitioner upon a study of the drawings, specification, and following claims.

CLAIMS

1. A platform comprising:

a plurality of devices including one or more of input output (IO) devices and system controllers;

a power management unit including,

a latency module to determine a latency constraint associated with the platform,

an idleness module to determine an idle window based on the latency constraint, and

an alignment module to instruct the plurality of devices to cease an activity during the idle window.

2. The platform of claim 1, wherein the power management unit further includes power management logic to place one or more components of the platform in a sleep state during the idle window.

3. The platform of claim 1, wherein the activity is to include external communication, and wherein the idle window is to be a platform wide idle window.

4. The platform of claim 1, further including a timer, wherein the alignment module is to set the timer based on the latency constraint.

5. The platform of claim 4, wherein the alignment module is to initiate an active window in response to an expiration of the timer.

6. The platform of claim 5, wherein the alignment module is to instruct the plurality of devices to resume activity to initiate the active window.

7. The platform of any one of claims 1 to 6, wherein the latency constraint is to be determined based on one or more of a service request and a performance level request.

8. An apparatus comprising:

a latency module to determine a latency constraint associated with a platform;

an idleness module to determine an idle window based on the latency constraint; and

an alignment module to instruct a plurality of devices on the platform to cease an activity during the idle window.

9. The apparatus of claim 8, further including power management logic to place the platform in a sleep state during the idle window.

10. The apparatus of claim 8, wherein the activity is to include external communication, and wherein the idle window is to be a platform wide idle window.

11. The apparatus of claim 8, further including a timer, wherein the alignment module is to set the timer based on the latency constraint.

12. The apparatus of claim 11, wherein the alignment module is to initiate an active window in response to an expiration of the timer.

13. The apparatus of claim 12, wherein the alignment module is to instruct the plurality of devices to resume activity to initiate the active window.

14. The apparatus of any one of claims 8 to 13, wherein the latency constraint is to be determined based on one or more of a service request and a performance level request.

15. A method comprising:
determining a latency constraint associated with a platform;
determining an idle window based on the latency constraint; and
instructing a plurality of devices on the platform to cease an activity during the idle window.

16. The method of claim 15, further including placing one or more components of the platform in a sleep state during the idle window.

17. The method of claim 16, wherein the activity includes external communication, and wherein the idle window is a platform wide idle window.

18. The method of claim 15, further including setting a timer based on the latency constraint.

19. The method of claim 18, further including initiating an active window in response to an expiration of the timer.

20. The method of claim 19, wherein initiating the active window includes instructing the plurality of devices to resume activity.

21. The method of any one of claims 15 to 20, wherein the latency constraint is determined based on one or more of a service request and a performance level request.

22. A non-transitory computer readable storage medium comprising a set of instructions which, if executed by a processor, cause a platform to:

determine a latency constraint associated with the platform;

determine an idle window based on the latency constraint; and

instruct a plurality of devices on the platform to cease an activity during the idle window.

23. The medium of claim 22, wherein the instructions, if executed, cause the platform to place one or more components of the platform in a sleep state during the idle window.

24. The medium of claim 23, wherein the activity is to include external communication, and wherein the idle window is to be a platform wide idle window.

25. The medium of claim 22, wherein the instructions, if executed, cause the platform to set a timer based on the latency constraint.

26. The medium of claim 25, wherein the instructions, if executed, cause the platform to initiate an active window in response to an expiration of the timer.

27. The medium of claim 26, wherein the instructions, if executed, cause the platform to instruct the plurality of devices to resume activity to initiate the active window.

28. The medium of any one of claims 22 to 27, wherein the latency constraint is to be determined based on one or more of a service request and a performance level request.

FIG. 1

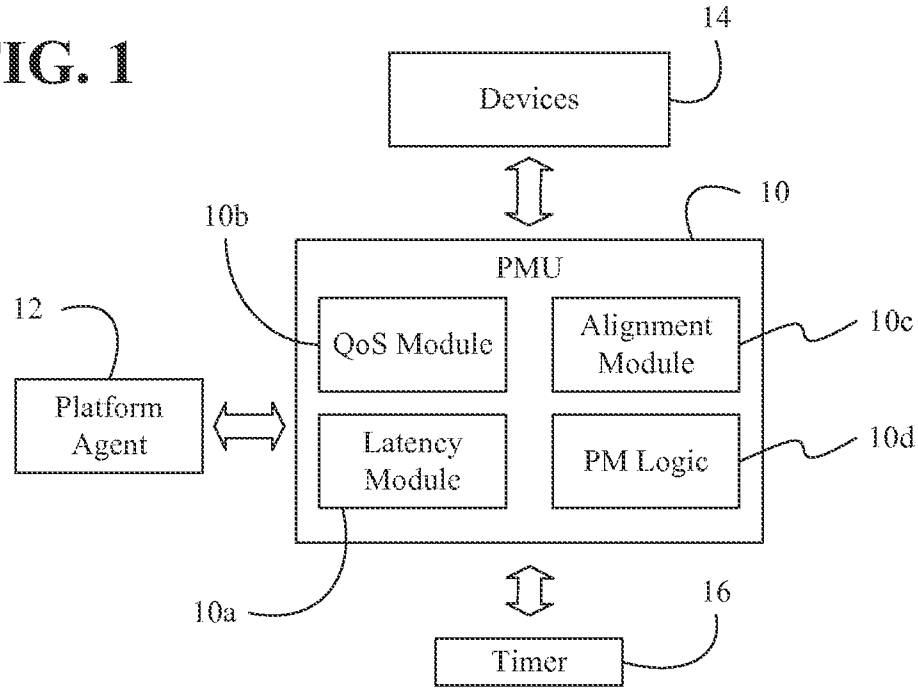
14
Devices
10
10b
PMU
12
QoS Module
Alignment Module
10c
Platform Agent
Latency Module
PM Logic
10d
10a
16
Timer

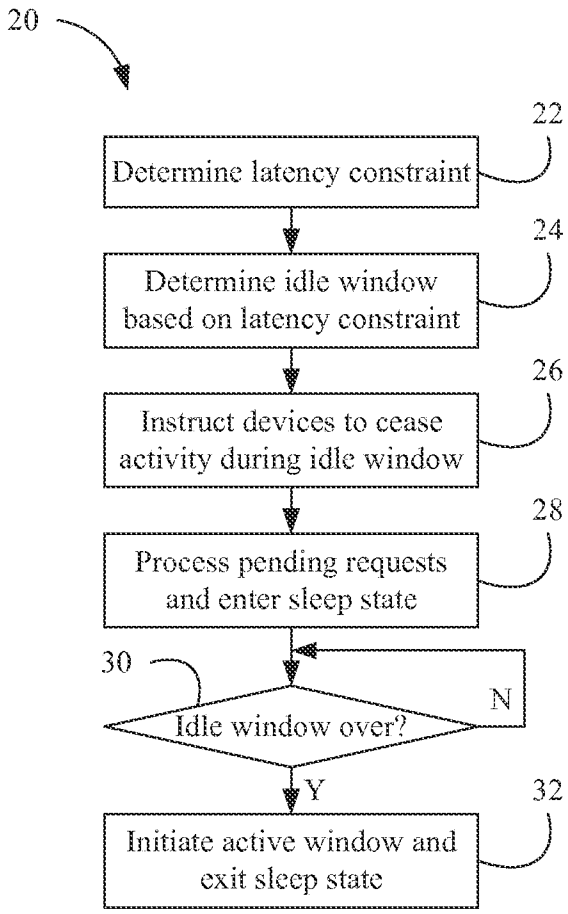
20
22
Determine latency constraint
24
Determine idle window based on latency constraint
26
Instruct devices to cease activity during idle window
28
Process pending requests and enter sleep state
30
Idle window over?
N
Y
32
Initiate active window and exit sleep state

FIG. 2

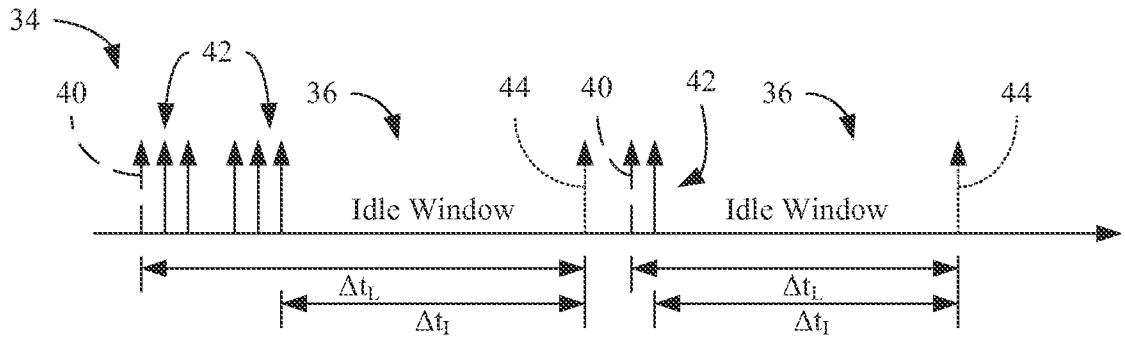

FIG. 3A

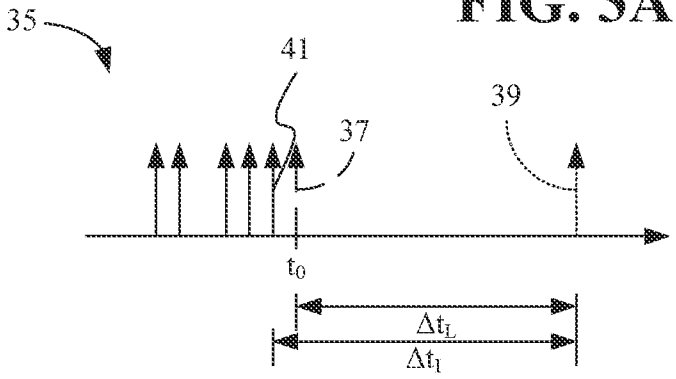

FIG. 3B

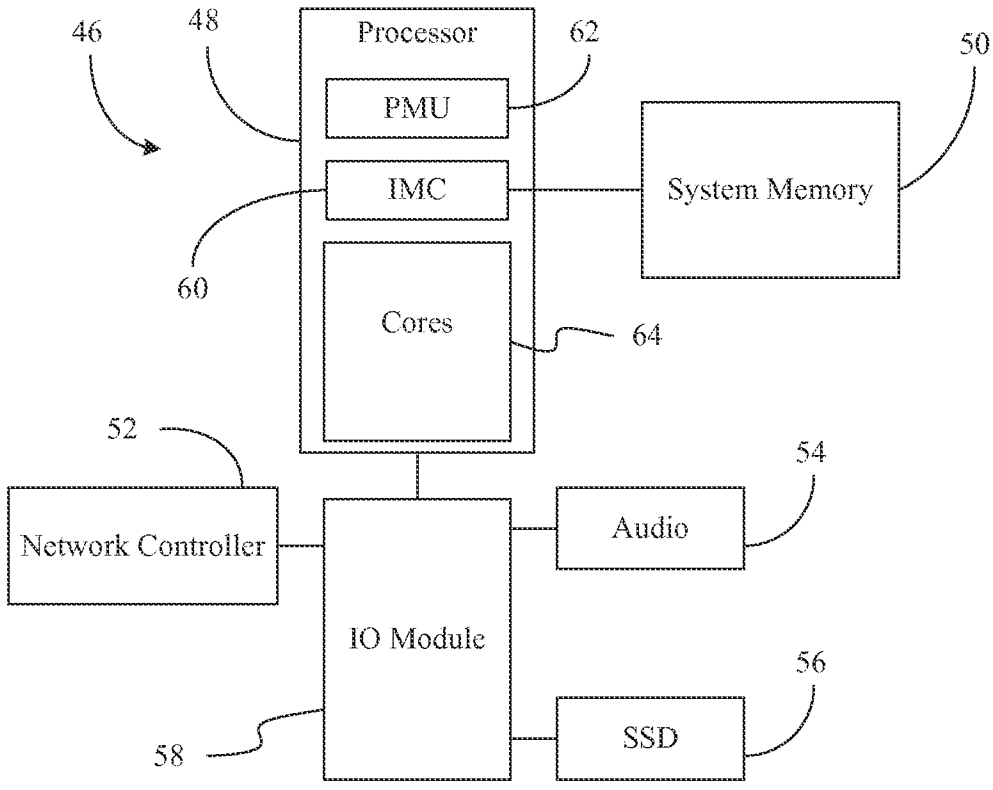

FIG. 4

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2013/047778

A. CLASSIFICATION OF SUBJECT MATTER

G06F 1/32(2006.01)i

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06F 1/32; H04M 1/00; H04B 1/16; G06F 1/26; G06F 1/00

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Korean utility models and applications for utility models
Japanese utility models and applications for utility models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & Keywords: power management unit, max latency, timer, idle window, active window.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2010-0077243 A1 (REN WANG et al.) 25 March 2010 See paragraphs [0013], [0018], [0028]-[0030], [0032]; and figs. 1, 4.	1-2,4-9,11-16,18-23,25-28
A		3,10,17,24
Y	US 2007-0079154 A1 (PAUL S. DIEFENBAUGH et al.) 05 April 2007 See paragraphs [0016], [0021], [0029]; and figs. 1, 4.	1-2,4-9,11-16,18-23,25-28
A	US 2007-0197186 A1 (ALAA MUQATTASH et al.) 23 August 2007 See paragraphs [0002]-[0012], [0095]-[0103]; and fig. 5.	1-28
A	US 2011-0131427 A1 (JOEL A. JORGENSON et al.) 02 June 2011 See paragraphs [0003], [0023]-[0026], [0056]-[0063]; and figs. 1, 8.	1-28
A	US 2010-0083017 A1 (CHRISTOPHER R. FULKERSON et al.) 01 April 2010 See paragraphs [0001]-[0003], [0022]-[0028]; and fig. 4.	1-28

☐ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

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

Date of the actual completion of the international search	Date of mailing of the international search report
08 October 2013 (08.10.2013)	**09 October 2013 (09.10.2013)**
Name and mailing address of the ISA/KR Korean Intellectual Property Office 189 Cheongsa-ro, Seo-gu, Daejeon Metropolitan City, 302-701, Republic of Korea Facsimile No. +82-42-472-7140	Authorized officer LEE Dong Yun Telephone No. +82-42-481-8734

Form PCT/ISA/210 (second sheet) (July 2009)

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.
PCT/US2013/047778

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2010-0077243 A1	25/03/2010	US 2012-131370 A1	24/05/2012
		US 8112651 B2	07/02/2012
		WO 2010-036592 A2	01/04/2010
		WO 2010-036592 A3	01/07/2010
US 2007-0079154 A1	05/04/2007	US 7519839 B2	14/04/2009
US 2007-0197186 A1	23/08/2007	EP 1982474 A1	22/10/2008
		JP 2009-521895 A	04/06/2009
		JP 2009-526437 A	16/07/2009
		JP 2009-526438 A	16/07/2009
		JP 2009-526439 A	16/07/2009
		US 2007-0184809 A1	09/08/2007
		US 2007-184870 A1	09/08/2007
		US 7751357 B2	06/07/2010
		US 7920504 B2	05/04/2011
		US 8135427 B2	13/03/2012
		WO 2007-076139 A2	05/07/2007
		WO 2007-076139 A3	13/09/2007
		WO 2007-092423 A1	16/08/2007
		WO 2007-092423 A8	27/09/2007
		WO 2007-092424 A1	16/08/2007
		WO 2007-136435 A2	29/11/2007
		WO 2007-136435 A3	10/04/2008
US 2011-0131427 A1	02/06/2011	TW 201135446 A	16/10/2011
		WO 2011-068833 A1	09/06/2011
US 2010-0083017 A1	01/04/2010	US 8495400 B2	23/07/2013