




- (51) International Patent Classification:
G06F 1/26 (2006.01)
- (21) International Application Number:
PCT/US2013/047361
- (22) International Filing Date:
24 June 2013 (24.06.2013)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
13/537,260 29 June 2012 (29.06.2012) 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): **MACIOCCO, Christian** [US/US]; 3625 NW Chapin Dr., Portland, Oregon 97229 (US). **FALIK, Ohad** [IL/IL]; 4b, Shizaf, St., 44418 Kfar-Saba (IL). **WANG, Ren** [US/US]; 9137 NW Esson Court, Portland, Oregon 97229 (US). **SHULMAN, Nadav** [IL/IL]; Hasayfan 19 Street, POB 447, 40600 Tel Mond, M (IL). **DIEFENBAUGH, Paul** [US/US]; 9016 NW Bartholomew Drive, Portland, Oregon 97129 (US). **TAI, Tsung-Yuan Charles** [US/US]; 12709 NW Majestic Sequoia Way, Portland, Oregon 97229 (US).
- (74) Agents: **JORDAN, B. Delano** et al.; Jordan IP Law, PC, 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: USING DEVICE IDLE DURATION INFORMATION TO OPTIMIZE ENERGY EFFICIENCY



(57) Abstract: Systems and methods may provide for aggregating a first idle duration from a first device associated with a platform and a second idle duration from a second device associated with the platform. Additionally, an idle state may be selected for the platform based at least in part on the first idle duration and the second idle duration. In one example, the idle durations are classified as deterministic, estimated or statistical.

USING DEVICE IDLE DURATION INFORMATION TO OPTIMIZE ENERGY EFFICIENCY

BACKGROUND

Technical Field

5 Embodiments generally relate to power management in computing platforms. More particularly, embodiments relate to the use of device idle durations to select idle states for computing platforms.

Discussion

10 In conventional mobile computing platforms, idle states may be used to reduce power consumption and extend battery life, wherein deeper idle states may generally conserve more power. Idle state depth may be limited, however, in order to guarantee the quality of service (QoS) that platform devices require. Unduly limiting idle state depth may have a negative impact on energy efficiency and performance in certain circumstances.

BRIEF DESCRIPTION OF THE DRAWINGS

15 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:

20 FIG. 1 is a plot of an example of idle state selections over time according to an embodiment;

 FIG. 2 is a block diagram of an example of an aggregator architecture according to an embodiment;

 FIG. 3 is a block diagram of an example of a platform having an idle duration information reporting scheme according to an embodiment;

25 FIG. 4 is a flowchart of an example of a method of determining idle duration actions according to an embodiment;

 FIG. 5 is a flowchart of an example of a method of managing platform idle states according to an embodiment;

30 FIG. 6 is a flowchart of an example of a method of managing platform idle states near the end of an idle duration according to an embodiment; and

 FIG. 7 is a plot of an example of a clock drift compensation scheme according to an embodiment.

DETAILED DESCRIPTION

Turning now to FIG. 1, an idle state selection plot is shown for a platform such as a computing device. The platform may have a platform latency tolerance requirement (LTR) 12, which can indicate the maximum acceptable delay in exiting an idle state selected for the platform to guarantee correctness. In this regard, transitioning out of a particular idle state is typically not instantaneous, wherein the platform LTR 12 may provide a mechanism of controlling the latency overhead associated with idle state transitions by determining which lower idle state can be reached to meet the exit latency requirements. As will be discussed in greater detail, the platform LTR 12 may be dependent upon the expected activity of the internal and external devices associated with the platform. In the illustrated example, idle duration information is aggregated from the devices associated with the platform, wherein idle duration information from a particular device may generally indicate the amount of time that the device is expected to remain idle.

More particularly, the idle duration information can be used to transition the platform more quickly into the deepest idle state that satisfies the platform LTR 12 along a state selection path 10. Indeed, the illustrated idle duration scheme would also be advantageous to platforms that do not have an LTR. Of particular note is that a conventional state selection path 14 that does not aggregate idle duration information from platform devices could result in the additional energy usage shown in a region 16 of the plot. In one example, the idle states are ACPI (Advanced Configuration and Power Interface, e.g., ACPI Specification, Rev. 4.0a, April 5, 2010) low power states, although other types of idle states may also be used.

FIG. 2 shows an aggregator architecture in which a subordinate aggregator 18 collects idle duration information from a first set of devices 20, and a root aggregator 22 collects idle duration information from the subordinate aggregator 18 and a second set of devices 24. The first and second sets of devices 20, 24 may include a wide variety of internal and/or external platform devices such as, for example, timers, audio devices, displays, cameras, mice, keyboards, sensors, network devices, communication devices, and so forth. The idle duration information may be carried from the devices 20, 24 to the aggregators 18, 22 using, for example, specific bus primitives such as a PCIe (Peripheral Components Interconnect Express, e.g., PCI Express x16 Graphics 150W-ATX Specification 1.0, PCI Special Interest Group) VDM (vendor defined message) or writes to specific memory addresses. Additionally, the aggregators 18, 22 can track the source of

idle duration information via, for example, identifying the downstream port from which idle duration messages originate or using global enumeration schemes such as the source BDF (bus, device, function) information that is reported in PCIe messages.

Turning now to FIG. 3 a platform is shown, wherein the platform may be part of a 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 platform includes a more detailed example of an idle duration reporting scheme in which the first set of devices 20 (20a-20k) report idle duration information to the subordinate aggregator 18, and the second set of devices 24 (24a-24e) report idle duration information to the root aggregator 22. In one example, the root aggregator 22 executes on a processor such as central processing unit (CPU) of the platform, and the subordinate aggregator 18 executes on an input/output (IO) subsystem of the platform, wherein the idle durations (e.g., idle duration information) from the devices 20, 24 may be classified by the aggregators 18, 22 as deterministic, estimated, or statistical.

In general, deterministic devices may be devices that are not expected to generate traffic for a specified period and are known to wake and begin conducting activity when the specified period ends. Such a device may generate traffic before the specified period ends, in which case the device in question may expect such traffic to be serviced according to a maximum latency (e.g., LTR) requested by the device. The LTR may therefore establish a QoS “floor” for each deterministic device. Delayed servicing, however, may be relatively infrequent (e.g., an exception to normal operation) in many circumstances. As will be discussed in greater detail, the deterministic devices can also report idle durations, wherein in one example, the deterministic idle durations are reported as a period of time from the time a message was last sent by the device.

Thus, the illustrated subordinate aggregator 18 receives deterministic idle durations and LTR information from internal devices such as a timer 20e (e.g., high performance timer/HPET) and an audio device 20i (e.g., audio direct memory access/DMA controller), and the illustrated root aggregator 22 receives deterministic idle durations from internal devices such as one or more timers 24e (e.g., local advanced programmable interrupt controller/LAPIC timers). Deterministic idle durations may also be received from external devices such as displays, cameras, and so forth.

In one example, the subordinate aggregator 18 maintains a deterministic counter for each of the timer 20e and the audio device 20i, wherein the deterministic counters decrement (e.g., count down) from the reported deterministic idle durations and the subordinate aggregator 18 selects the minimum deterministic counter value as a
5 deterministic idle duration to be reported to the root aggregator 22 when it is later queried by the root aggregator 22 or decides to push the information to root aggregator 22. Similarly, the root aggregator 22 may maintain a deterministic counter for the each of the timers 24e and the deterministic idle duration reported by the subordinate aggregator 18, and select the minimum deterministic counter value as a deterministic idle duration for the
10 platform.

FIG. 4 shows a method 28 of determining idle duration actions. The method 28 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
15 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.

20 Illustrated processing block 30 determines an idle duration for the device. As already noted, the idle duration may be considered relative in time to a message sent by the device, wherein the device could plan to enter a sleep/idle state until a response to the message needs to be processed. A determination may be made at block 32 as to whether the platform is configured to conduct one or more pre-wake activities, wherein pre-wake
25 activities might include, for example, power and/or clock related tasks that enable the platform to exit the idle state in gradual steps more efficiently and quicker. The determination at block 32 could be made, for example, by accessing an appropriate register or programmable memory location. If the platform is configured to conduct pre-wake activities, illustrated block 34 increases (e.g., lengthens) the LTR of the
30 deterministic device. Increasing the LTR can permit more latency during exits from idle states and may effectively enable the platform to enter deeper idle states.

If the platform is not configured to conduct pre-wake activities, on the other hand, the LTR is left unchanged in the illustrated example. Block 36 can provide for outputting the idle duration and LTR to an aggregator such as the root aggregator 22 (FIGs. 2 and 3)

or the subordinate aggregator 18 (FIGs. 2 and 3), as already discussed. As will be discussed in greater detail, the deterministic device may also report another, shorter, LTR to be used for “pop-ups” near the end of idle periods.

Estimated devices, on the other hand, may be devices in which there is a relatively high likelihood that the device will not generate traffic for a specified period of time, wherein the device might not have traffic when the specified period of time expires. Examples of such devices include, for example, network devices (e.g., extensible host controller interface/XHCI, wired/wireless devices) and other communication devices having buffers that fill up based on incoming traffic or scheduled windows of operation during which activity may occur. These devices may also report idle duration as a period of time from when a message was sent.

Thus, returning to FIG. 3, the illustrated subordinate aggregator 18 receives estimated idle durations from external devices such as expansion devices 20a, 20b that are coupled to an external bus via an expansion port (“Expansion Port 1”) and an external hub (“Ex. Hub”), and from a card device 20c that is coupled to the external bus via an expansion port (“Expansion Port 2”). In one example, the expansion ports and external hub use a PCIe protocol to communicate with one another. The subordinate aggregator 18 may also receive estimated idle durations from internal devices such as a local area network (LAN) device 20d, a remote interface (I/F, e.g., management engine/ME interface) 20f, an input/output (IO) port 20h (e.g., Universal Serial Bus/USB port), a storage I/F 20j (e.g., Serial Advanced Technology Attachment/SATA interface), and so forth.

In the illustrated example, the platform also includes certain devices 21 (21a-21g) that do not report idle duration information (e.g., “no hint” devices). In such a case, one or more interfaces may be used to estimate idle durations based on the traffic pattern of communications with the devices 21. For example, the illustrated IO port 20h estimates the idle duration based on the traffic pattern of connected devices such as an HID (human interface device) device 21a (e.g., mouse, keyboard, sensor, etc.), an isochronous device 21b (e.g., audio, camera), a bulk device 21c (e.g., storage, printer), and an SSD (solid state disk) 21d. More particularly, the traffic pattern of the isochronous device 21b could be closer to deterministic in nature, as when the next slot will happen may be known exactly. Idle durations for the bulk device 21c, on the other hand, might be estimated based on previous activity. The illustrated storage I/F 20j is also coupled to the SSD 21d, which

may in turn be coupled to a hub 21e and one or more devices 21f, 21g, wherein the storage I/F 21j can estimate idle durations.

For example, the storage I/F 20j may know if there are any IO requests pending, in which case it might estimate the idle duration to be something like “soon < 1ms”. If, on the other hand, there are no IO requests pending, the storage I/F 21j may estimate the idle duration to be “long ~ at least 5 ms”. Other values and implementations may also be used, depending upon the circumstances. Moreover, the illustrated root aggregator 22 receives estimated idle durations from an external communications device 24a, an external graphics controller 24b, an integrated display 24c, and an internal graphics controller 24d.

In one example, the subordinate aggregator 18 maintains an estimated counter for each of the expansion devices 20a, 20b, card device 20c, LAN device 20d, remote I/F 20f, IO port 20h, and storage I/F 20j, wherein the estimated counters decrement from the reported estimated idle durations. The subordinate aggregator 18 may therefore identify the minimum estimated counter value as an estimated idle duration to be reported to the root aggregator 22. Similarly, the illustrated root aggregator 22 maintains an estimated counter for each of the external communications device 24a, external graphics controller 24b, integrated display 24c, and internal graphics controller 24d, and the estimated idle duration reported by the subordinate aggregator 18, and selects the minimum estimated counter value as an estimated idle duration for the platform.

Statistical devices may be devices having traffic patterns that are only quantifiable in terms of a specific average rate. Examples of such devices include, for example, mice, keyboards, sensors, and so forth. These devices may report idle duration as an expected period of time between events. Thus, the illustrated subordinate aggregator 18 also receives statistical idle durations from internal devices such as an HID device 20g and an HID device 20k via an IO port (“IO Port 1”, e.g., USB port). In one example, the subordinate aggregator 18 maintains a statistical counter for each of the HID devices 20g, 20k, wherein the statistical counters decrement from the reported statistical idle durations. The subordinate aggregator 18 may therefore identify the minimum statistical counter value as a statistical idle duration to be reported to the root aggregator 22.

If the minimum statistical counter value has multiple values close to it, a compensation approach may be used in which a shorter duration is assumed because the probability that one of the events will occur increases. Thus, the illustrated subordinate aggregator 18 could reduce the statistical idle duration and/or the platform idle duration if the idle duration from the HID device 20g and the idle duration from the HID device 20k

are within a predetermined range of one another. The platform may also include platform management logic 26 that selects idle states for the platform based on the idle duration information from the aggregators 18, 22.

In particular, FIG. 5 shows a method 38 of managing platform idle states. The method 38 may be implemented as a set of logic instructions stored in a machine- or computer-readable storage medium such as RAM, ROM, PROM, firmware, flash memory, etc., in configurable logic such as, for example, PLAs, FPGAs, CPLDs, in fixed-functionality logic hardware using circuit technology such as, for example, ASIC, CMOS or TTL technology, or any combination thereof. Illustrated processing block 40 aggregates idle durations from platform devices, wherein one or more deterministic idle durations and latency tolerances may be identified at block 42. Block 44 may provide for determining a minimum deterministic idle duration among the deterministic idle durations. Additionally, one or more statistical idle durations and latency requirements can be identified at block 46, wherein illustrated block 48 determines a minimum statistical idle duration among the statistical idle durations. Block 50 may identify one or more estimated idle durations and latency tolerances, and illustrated block 52 determines a minimum estimated idle duration among the estimated idle durations. In addition, a platform idle duration may be determined at block 54 based on the minimum deterministic idle duration, the minimum statistical idle duration, and the minimum estimated idle duration. In one example, the minimum value of these minimums is used as the platform idle duration. A platform idle state may be selected at block 56 based on the platform idle duration and the latency tolerances.

For example, the device corresponding to the minimum value used for the platform idle duration might be identified using downstream port information, global enumeration schemes, etc., as already noted. Accordingly, the LTR of that device may also be identified and used as a platform LTR, such as the platform LTR 12 (FIG. 1), to select an idle state for the platform.

Turning now to FIG. 6, a method 58 of managing platform idle states near the end of an idle duration is shown. The method 58 may be implemented as a set of logic instructions stored in a machine- or computer-readable storage medium such as RAM, ROM, PROM, firmware, flash memory, etc., in configurable logic such as, for example, PLAs, FPGAs, CPLDs, in fixed-functionality logic hardware using circuit technology such as, for example, ASIC, CMOS or TTL technology, or any combination thereof. Illustrated processing block 60 determines whether an idle platform is near the end of the platform

idle duration. If so, a “pop-up duration” may be determined at block 62 based on a relatively short latency tolerance. In this regard, a process may be used to briefly exit the platform from the idle state and determine whether the same or a deeper idle state is achievable. Thus, the process could be considered a “pop-up process” that takes place
5 over a period of time that may be referred to as a “pop-up duration” or window (e.g., temporary “wake duration”).

Moreover, when the platform idle duration is associated with a deterministic device, the illustrated approach determines the length of the pop-up duration based on a relatively short LTR in the example shown, in order to ensure/guarantee proper servicing
10 of deterministic device traffic (because deterministic devices can report longer LTRs when the platform is configured for pre-wake operation). The shorter LTR may therefore be stored to a deterministic guaranteed LTR register that defines the responsiveness of the platform at the end of a deterministic idle duration, wherein the shorter LTR can be guaranteed for the pop-up duration.

Thus, illustrated block 64 conducts one or more pre-wake activities, wherein a
15 determination can be made at block 66 as to whether the deterministic device (e.g., associated with the platform idle duration) has generated any traffic. Thus, block 66 may involve monitoring the device associated with the platform idle duration for activity for a period of time equal to the pop-up duration. If no such activity has occurred, block 68
20 determines whether the pop-up duration has expired. If the pop-up duration has not expired the illustrated loop checks again for device activity. When the pop-up duration expires without any device activity being detected, illustrated block 70 places the platform back into the same or a deeper idle state in order to further conserve energy and extend battery life. If device activity is detected at block 66, illustrated block 72 provides for
25 processing the data/traffic associated with the device in question.

In cases where the device and the aggregators do not use a common time base (e.g., clock), clock drift compensation may be used. In this regard, FIG. 7 shows a clock drift compensation scheme in which device/aggregator action is plotted against time. In the illustrated example, an idle message is transmitted from a first device to an aggregator,
30 wherein the first device uses a clock that is slightly faster than the clock being used by the aggregator (e.g., negative drift). A device line segment 76 begins at the opening of an idle message propagation window 78 and an aggregator line segment 74 begins at the closing of the idle message propagation window 78. Thus, the device line segment 76 represents the time period during which the first device is in the idle state, and the aggregator line

segment 74 represents the time period during which the aggregator is in the idle state, in the example shown. At time T_1 , the idle duration expires. Because the first device has a clock that is ahead of the clock being used by the aggregator, however, the first device wakes before the aggregator. In order to ensure that messages from the first device to the aggregator are not lost, the illustrated approach creates a device wait window 80 during which the device waits for the aggregator to exit the idle state.

Similarly, the aggregator could use a clock that is slightly faster than the clock being used by a second device (e.g., positive drift). For example, an aggregator line segment 84 and a device line segment 82 for the second device might exit an idle state at time T_2 , wherein the clock drift (e.g., skew) causes the aggregator to actually wake before the second device. Accordingly, the illustrated approach implements an aggregator wait window 86 in order to prevent the aggregator from incorrectly assuming that the second device has no data/traffic to be processed. In one example, the device wait window 80 and the aggregator wait window 86 may be implemented by compensating the platform idle duration and/or the pre-wake activity for clock drift. For example, a limit may be placed on the platform idle duration so that the clock drift during the idle duration will be shorter than the pop-up duration. Additionally, pre-wake activities can be scheduled to occur before the device wakes, taking into account maximum negative drift. For example, the pre-wake activities may be scheduled to occur before the end of the device line segment 76, in the example shown.

Embodiments may therefore include a deterministic device having logic to determine an idle duration for the device, and determine a latency tolerance based on a pre-wake configuration of a platform associated with the device. The logic may also output the idle duration and the latency tolerance. In one example, the logic increases the latency tolerance if the platform is configured to conduct one or more pre-wake activities.

Embodiments may also include an apparatus having first aggregator logic to aggregate a first idle duration from a first device and a second idle duration from a second device. Additionally, the apparatus may have power management logic to select an idle state for the platform based at least in part on the first idle duration and the second idle duration.

Embodiments can also include a platform having a first device to output a first idle duration, a second device to output a second idle duration, and root aggregator logic to aggregate the first idle duration and the second idle duration. The platform may also have

power management logic to select an idle state for the platform based at least in part on the first idle duration and the second idle duration.

Other embodiments may include at least one machine readable storage medium having a set of instructions which, if executed by at least one processor, cause a platform
5 to aggregate a first idle duration from a first device associated with the platform and a second idle duration from a second device associated with the platform. The instructions, if executed, can further cause the platform to select an idle state for the platform based at least in part on the first idle duration and the second idle duration.

Additionally, embodiments may include a method in which a first idle duration
10 associated with a platform is aggregated with a second idle duration from a second device associated with the platform. The method can also provide for selecting an idle state for the platform based at least in part on the first idle duration and the second idle duration.

Techniques described herein may therefore use knowledge of upcoming deterministic periods of idleness in platform devices such as display devices, audio
15 devices, timers, camera pipelines, etc., to enable processing components to pre-wake, reduce latency and optimize wake energy. Moreover, knowledge of upcoming estimated periods of idleness in platform devices such as communication devices, network devices, etc., can enable processing components to select the most appropriate idle state given the idle duration period and the energy break-even time of the specific state. Simply put, idle
20 duration reporting and aggregation can add determinism to the idle state decision process, which may in turn enable the optimal energy efficiency state to be reached as soon as possible as well as enable residency in this state to be optimized.

In addition, techniques described herein may be scalable, covering both internal and external devices while remaining topology independent. Accordingly, architectures
25 such as laptops, notebooks, ultrabooks, smart phones, etc., may readily benefit from these solutions. Moreover, the data collection schemes described herein can particularly enhance operation in the presence of deep sleep states with relatively long exit times and energy break even times.

Embodiments of the present invention are applicable for use with all types of
30 semiconductor integrated circuit ("IC") chips. Examples of these IC chips include but are not limited to processors, controllers, chipset components, programmable logic arrays (PLAs), memory chips, network chips, systems on chip (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. Moreover, any use of the terms “first”, “second”, etc., does not limit the embodiments discussed to the number of components listed.

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

We claim:

1. An apparatus comprising:
first aggregator logic to aggregate a first idle duration from a first device and a
5 second idle duration from a second device; and
power management logic to select an idle state for the platform based at least in
part on the first idle duration and the second idle duration.
2. The apparatus of claim 1, wherein the first aggregator logic is to,
identify the first idle duration and the second idle duration as deterministic idle
10 durations, and
determine a platform idle duration based at least in part on the deterministic idle
durations, wherein the idle state is to be selected based at least in part on the platform idle
duration.
3. The apparatus of claim 2, wherein the first aggregator logic is to determine
15 a first latency tolerance associated with the platform idle duration, and the power
management logic is to conduct a pre-wake activity prior to an expiration of the platform
idle duration.
4. The apparatus of claim 3, further including a register to store a second
latency tolerance, wherein the first aggregator logic is to,
20 determine a wake duration for the platform based at least in part on the second
latency tolerance, wherein the second latency tolerance is to be less than the first latency
tolerance; and
use the wake duration to monitor a device associated with the platform idle
duration for activity.
- 25 5. The apparatus of claim 3, wherein the power management logic is to
compensate one or more of the platform idle duration and the pre-wake activity for clock
drift.
6. The apparatus of claim 1, wherein the first aggregator logic is to,
identify the first idle duration and the second idle duration as statistical idle
30 durations, and
determine a platform idle duration based at least in part on the statistical idle
durations, wherein the idle state is to be selected based at least in part on the platform idle
duration.

7. The apparatus of claim 6, wherein the first aggregator logic is to reduce the platform idle duration if the first idle duration and the second idle duration are within a predetermined range of one another.

8. The apparatus of claim 1, wherein the first aggregator logic is to,
5 identify the first idle duration and the second idle duration as estimated idle durations, and

determine a platform idle duration based at least in part on the estimated idle durations.

9. The apparatus of claim 1, further including second aggregator logic to
10 output a plurality of aggregated idle durations, wherein the first aggregator logic is to receive the plurality of aggregated idle durations, and wherein the idle state is to be selected further based on the plurality of aggregated idle durations.

10. A platform comprising:
a first device to output a first idle duration;
15 a second device to output a second idle duration;
first aggregator logic to aggregate the first idle duration and the second idle duration; and

power management logic to select an idle state for the platform based at least in part on the first idle duration and the second idle duration.

20 11. The platform of claim 10, wherein the first aggregator logic is to, identify the first idle duration and the second idle duration as deterministic idle durations, and

determine a platform idle duration based at least in part on the deterministic idle durations, wherein the idle state is to be selected based at least in part on the platform idle
25 duration.

12. The platform of claim 11, wherein the first aggregator logic is to determine a first latency tolerance associated with the platform idle duration, and the power management logic is to conduct a pre-wake activity prior to an expiration of the platform idle duration.

30 13. The platform of claim 12, further including a register to store a second latency tolerance, wherein the first aggregator logic is to,

determine a wake duration for the platform based at least in part on the second latency tolerance, wherein the second latency tolerance is to be less than the first latency tolerance; and

use the wake duration to monitor a device associated with the platform idle duration for activity.

14. The platform of claim 11, wherein the first and second devices include one or more of a timer, an audio device, a display and a camera.

5 15. The platform of claim 10, wherein the first aggregator logic is to, identify the first idle duration and the second idle duration as statistical idle durations, and

determine a platform idle duration based at least in part on the statistical idle durations, wherein the idle state is to be selected based at least in part on the platform idle
10 duration.

16. The platform of claim 15, wherein the first and second devices include one or more of a mouse, a keyboard and a sensor.

17. The platform of claim 10, wherein the first aggregator logic is to, identify the first idle duration and the second idle duration as estimated idle
15 durations, and

determine a platform idle duration based at least in part on the estimated idle durations.

18. The platform of claim 17, wherein the first and second devices include one or more of a network device and a communication device.

20 19. At least one machine readable storage medium comprising a set of instructions which, if executed by at least one processor, cause a platform to:

aggregate a first idle duration from a first device associated with the platform and a second idle duration from a second device associated with the platform; and

select an idle state for the platform based at least in part on the first idle duration
25 and the second idle duration.

20. The at least one machine readable storage medium of claim 19, wherein the instructions, if executed, cause a platform to:

identify the first idle duration and the second idle duration as deterministic idle durations; and

30 determine a platform idle duration based at least in part on the deterministic idle durations, wherein the idle state is to be selected based at least in part on the platform idle duration.

21. The at least one machine readable medium of claim 20, wherein the instructions, if executed, cause a platform to:

determine a first latency tolerance associated with the platform idle duration; and
conduct a pre-wake activity prior to an expiration of the platform idle duration.

22. The at least one machine readable medium of claim 21, wherein the
instructions, if executed, cause a platform to:

5 determine a wake duration for the platform based at least in part on a second
latency tolerance associated with the platform idle duration, wherein the second latency
tolerance is to be less than the first latency tolerance; and

use the wake duration to monitor a device associated with the platform idle
duration for activity.

10 23. The at least one machine readable medium of claim 21, wherein the
instructions, if executed, cause a platform to compensate one or more of the platform idle
duration and the pre-wake activity for clock drift.

24. The at least one machine readable medium of claim 19, wherein the
instructions, if executed, cause a platform to:

15 identify the first idle duration and the second idle duration as statistical idle
durations; and

determine a platform idle duration based at least in part on the statistical idle
durations, wherein the idle state is to be selected based at least in part on the platform idle
duration.

20 25. The at least one machine readable medium of claim 24, wherein the
instructions, if executed, cause a platform to reduce the platform idle duration if the first
idle duration and the second idle duration are within a predetermined range of one another.

26. The at least one machine readable medium of claim 19, wherein the
instructions, if executed, cause a platform to:

25 identify the first idle duration and the second idle duration as estimated idle
durations; and

determine a platform idle duration based at least in part on the estimated idle
durations.

27. The at least one machine readable medium of claim 19, wherein the
30 instructions, if executed, cause a platform to receive a plurality of aggregated idle
durations from an aggregator associated with the platform, wherein the idle state is to be
selected further based on the plurality of aggregated idle durations.

28. A method of managing platform idle states comprising:

aggregating a first idle duration from a first device associated with a platform and a second idle duration from a second device associated with the platform; and

selecting an idle state for the platform based at least in part on the first idle duration and the second idle duration.

5 29. The method of claim 28, further including:

identifying the first idle duration and the second idle duration as deterministic idle durations; and

 determining a platform idle duration based at least in part on the deterministic idle durations, wherein the idle state is selected based at least in part on the platform idle
10 duration.

 30. The method of claim 29, further including:

determining a first latency tolerance associated with the platform idle duration; and
conducting a pre-wake activity prior to an expiration of the platform idle duration.

 31. A device comprising:

15 logic to,

 determine an idle duration for the device,
 determine a latency tolerance based on a pre-wake configuration of a
platform associated with the device, and
 output the idle duration and the latency tolerance.

20 32. The device of claim 31, wherein the logic is to increase the latency tolerance if the platform is configured to conduct one or more pre-wake activities.

 33. The device of claim 31, wherein the device includes one or more of a timer, an audio device, a display and a camera.




**FIG. 1****FIG. 2**

FIG. 3







**FIG. 4****FIG. 5**

FIG. 6**FIG. 7**

A. CLASSIFICATION OF SUBJECT MATTER**G06F 1/26(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 1/26; G06F 12/14; G06F 1/00; G06F 1/32; G06F 12/08

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:computer, power, management, idle, duration, state

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2011-0154080 A1 (REN WANG et al.) 23 June 2011 See paragraphs [0006], [0017], [0030], [0037] and [0047]-[0048]; and figures 1-2 and 4	1, 10, 19, 28
Y		2, 6, 8, 11, 14-18, 20, 24, 26, 29
A		3-5, 7, 9, 12-13, 21-23, 25, 27, 30, 31-33
Y	US 2011-0239302 A1 (WON-HA CHOI) 29 September 2011 See paragraph [0016]	2, 6, 8, 11, 14-18, 20, 24, 26, 29
A	US 2008-0098245 A1 (MICHAEL T. HOGAN et al.) 24 April 2008 See paragraphs [0038]-[0044]; and figures 1-3	1-33
A	US 2007-0294554 A1 (SEO KWANG KIM) 20 December 2007 See paragraphs [0039]-[0042]; and figures 1-2 and 5	1-33
A	US 2006-0288240 A1 (JAMES P. KARDACH) 21 December 2006 See paragraphs [0017]-[0019]; and figure 1	1-33



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

08 October 2013 (08.10.2013)

Date of mailing of the international search report

10 October 2013 (10.10.2013)

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



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2013/047361

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2011-0154080 A1	23/06/2011	None	
US 2011-0239302 A1	29/09/2011	KR 10-2011-0108504 A	06/10/2011
US 2008-0098245 A1	24/04/2008	WO 2008-070912 A1	19/06/2008
US 2007-0294554 A1	20/12/2007	CN 100541393 C	16/09/2009
		CN 1499339 A	26/05/2004
		EP 1416358 A2	06/05/2004
		EP 1416358 A3	08/04/2009
		JP 04636302 B2	23/02/2011
		JP 2004-152309 A	27/05/2004
		KR 10-0499611 B1	05/07/2005
		US 07376851 B2	20/05/2008
		US 07707447 B2	27/04/2010
		US 2004-0117679 A1	17/06/2004
US 2006-0288240 A1	21/12/2006	CN 101198923 A	11/06/2008
		CN 101198923 B	20/06/2012
		EP 1891500 A2	27/02/2008
		JP 04660591 B2	30/03/2011
		JP 2008-544736 T	04/12/2008
		TW I330779 A	21/09/2010
		US 07454632 B2	18/11/2008
		WO 2006-138687 A2	28/12/2006
		WO 2006-138687 A3	01/03/2007