



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

(21) International Application Number:
PCT/US201 1/067377

(22) International Filing Date:
27 December 2011 (27. 12.2011)

(25) Filing Language: English

(26) Publication Language: English

(71) Applicant (for all designated States except US): INTEL CORPORATION [US/US]; 2200 Mission College Blvd., Santa Clara, California 95052 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): MIN, Andrew, W. [US/US]; 5566 NW Primino Avenue, Portland, Oregon 9972 (US). WANG, Ren [US/US]; 9137 NW Esson Ct., Portland, Oregon 97229 (US). TSAI, Jr-Shian [US/US]; 6657 NW 165th Ave., Portland, Oregon 97229 (US). ERGIN, Mesut A. [TR/US]; 15720 NW Nightshade Dr., Portland, Oregon 97229 (US). TAI, Tsung-Yuan C. [US/US]; 12709 NW Majestic Sequoia Way, Portland, Oregon 97229 (US). MAVILA, Rajith K. [US/IN]; K 705

Purva Fountain Square, Near Marathahalli Flyover, Varthur Road, Marathahalli, Bangalore 560-037 (IN). IYER, Prakash N. [US/US]; 12535 NW Lilywood Drive, Portland, Oregon 97229 (US).

(74) Agent: JORDAN, B., Delano; 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, 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, 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, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE,

[Continued on next page]

(54) Title: POWER MANAGEMENT USING REWARD-BASED SLEEP STATE SELECTION

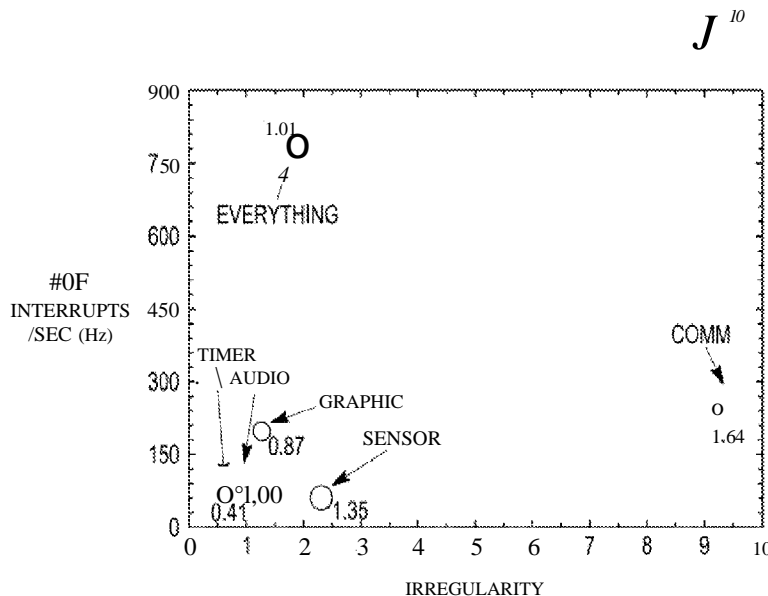


FIG. 1

(57) Abstract: Systems and methods may provide for conducting a reward determination for a plurality of sleep states to obtain a plurality of reward determinations with respect to a device. In addition, a sleep state may be selected for the device from the plurality of sleep states based at least in part on the plurality of reward determinations. In one example, false entry and missed opportunity probabilities may be determined for stochastic interrupts, wherein the reward determination is conducted based at least in part on the false entry and missed opportunity probabilities.

WO 2013/100901 A1

W O 2013/100901

A i 

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, ML, MR, NE, SN, TD, TG).

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

POWER MANAGEMENT USING REWARD-BASED SLEEP STATE SELECTION**BACKGROUND****5 Technical Field**

Embodiments generally relate to power management in computing platforms. More particularly, embodiments relate to the use of reward-based metrics to select sleep states for computing platforms.

10 Discussion

In conventional mobile computing platforms, sleep states may be used to reduce power consumption and extend battery life. The determination of whether to enter a given sleep state may be limited to making a prediction as to whether a known timer interrupt will occur prior to the energy break even time of the sleep state, wherein the prediction can be corrected
15 based on the historical success of past predictions. Such an approach may have a negative impact on energy efficiency and performance in certain circumstances.

BRIEF DESCRIPTION OF THE DRAWINGS

The various advantages of the embodiments of the present invention will become
20 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 plot of an example of irregularity versus frequency for various types of interrupts according to an embodiment;

FIG. 2 is a plot of an example of deterministic and stochastic interrupts over time
25 according to an embodiment;

FIG. 3 is a block diagram of an example of a scheme of decomposing deterministic sources according to an embodiment;

FIG. 4 is a chart of an example of prediction error for grouped and ungrouped
deterministic sources according to an embodiment;

FIG. 5 is a flowchart of an example of a method of selecting sleep states according to
30 an embodiment; and

FIG. 6 is a block diagram of an example of a computing platform according to an
embodiment.

35

DETAILED DESCRIPTION

Reducing platform power consumption to extend battery life can be particularly advantageous for small form factor platforms with active workloads, such as smartphones and tablets. A number of processors and/or platforms may provide multiple low power sleep states (e.g., Cx, SOix), wherein deeper sleep states can typically consume less power but may require more time to enter and exit the sleep states. Ensuring efficient use of sleep states may therefore be a key to power management, particularly in platforms having a wide variety of interrupts that can cause sleep state exits. As will be discussed in greater detail, an enhanced approach to sleep state selection may involve taking into account the probability of making sub-optimal decisions and their associated overhead (e.g., reward based sleep state selection), as well as distinguishing between the deterministic and stochastic nature of system interrupts that can impact the effectiveness of sleep state selections. In particular, the following discussion first addresses the deterministic/stochastic distinction, and then discusses reward based sleep state selections.

Turning now to FIG. 1, a plot 10 shows the irregularity and prediction error (indicated by the size the circles) of various types of interrupts in a typical mobile platform streaming video versus the frequency (number of interrupts per second) of the interrupts. In general, the plot 10 demonstrates that certain types of interrupts may be considered deterministic, whereas other types of interrupts may be considered stochastic. In the example shown, timer (e.g., operating system/OS programmed future events) interrupts occur with the most regularity, are relatively infrequent, and have the lowest prediction error (i.e., 0.41). Thus, such interrupts may be considered deterministic. Communications (e.g., network packet related) interrupts, on the other hand, may occur with the least regularity, be more frequent, and have the highest prediction error (i.e., 1.64) when a simple moving average estimator is used - accordingly, these types of interrupts may be considered stochastic.

Simply put, more predictable types of interrupts, such as timer and audio interrupts, may be designated as deterministic interrupts, whereas less predictable types of interrupts, such as graphic, sensor and communications interrupts, may be designated as stochastic interrupts. As will be discussed in greater detail, separating and isolating interrupt patterns into the deterministic and stochastic classes described herein can improve the probabilistic accuracy of idle duration predictions, and may ultimately reduce false sleep state decisions caused by stochastic interrupts. FIG. 2 shows a deterministic (e.g., timer subsystem) interrupt pattern 12 in comparison to a stochastic (e.g., communications subsystem) interrupt pattern 14.

Deterministic Interrupts

FIG. 3 shows a scheme 16 of decomposing a composite source of interrupts in order to improve interrupt prediction accuracy. Generally, deterministic interrupts can be identified at design time (e.g., by consulting device specifications), at run-time via estimation, or any

combination thereof. The three stages of the illustrated scheme 16 may be conducted as follows.

Decomposition: A platform may have multiple interrupt sources during active workloads, wherein the overall composition of interrupts can exhibit more random behavior than their individual contributing sources. Because such a phenomenon may make idle duration predictions difficult, stage one of the illustrated scheme 16 decomposes the multiple interrupt sources and considers their associated idle periods individually. Information pertaining to interrupt sources may be generally available in platform/system firmware as well as in OS device drivers.

Grouping: Certain interrupt sources may generate a burst of interrupts with very small inter-arrival times occurring within the burst. Stage two of the illustrated scheme 16 groups this known burst with a single system event, and introduces a "holding period". In particular, during the holding period, idle duration prediction may pause until the rest of the interrupt burst completes. Such an approach can save computation cycles related to estimation, and may in turn increase accuracy.

Dependency: In addition, some interrupts from different sources/devices may have timing dependencies on one another. Accordingly, stage three of the illustrated scheme 16 exploits such dependencies among interrupt sources to improve the prediction accuracy while minimizing overhead. For example, most slave device controller interrupts may be likely to be followed by an interrupt from the controlled device. Moreover, holding periods may be implemented in the dependency stage as well to further enhance prediction performance. Simply put, the scheme 16 can provide the basis of more accurate prediction of deterministic interrupts.

FIG. 4 demonstrates the advantages that may be associated with a composite interrupt source decomposition scheme as described herein. In particular, a chart 18 shows average interrupt interval prediction error for an audio subsystem, when a weighted moving average predictor is employed. For example, three audio subsystems, denoted by "A1", "A2", and "A3", share one interrupt, and by decomposing this composite interrupt to identify interrupt instances corresponding to these subsystems, we can improve the prediction accuracy. The first bar ("Audio") indicates the prediction error when the interrupts from three subsystems are considered as one composite source (i.e., similar to conventional Linux kernel idle governor solutions). The following three bars indicate the prediction error when the interrupts from different subsystems are considered separately. The illustrated results reveal that the prediction error significantly decreases by decomposing the interrupts, except for the interrupt A3. Accordingly, the underlying audio function tied to interrupt A3 may be bursty in nature. Applying the aforementioned grouping solution to interrupt A3, however, enables each burst to be considered as a single event. Accordingly, the prediction accuracy may improve to a level

similar to the other interrupts, as demonstrated in the fifth bar.

Stochastic Interrupts

Unlike deterministic and pseudo-deterministic interrupts, stochastic interrupts may be more challenging to predict due to their inherent randomness. For interrupts that fall into the stochastic class, a probabilistic idle duration estimator may be deployed. In addition, an efficient approach to using the stochastic interrupt prediction may be deployed for processor and/or sleep state selection.

The interrupt arrivals from a stochastic source can be modeled by using an appropriate stochastic process, such as a Poisson process, where the probability of occurrence for at least one stochastic interrupt during the next time period τ can be estimated as

$$P_{stc}(r) \approx 1 - e^{-\hat{\lambda}_{stc}\tau} \quad (1)$$

where $\hat{\lambda}_{stc}$ is the interrupt arrival rate, which can be estimated using an adaptive filter (e.g., moving average filter) at run-time. This approach enables the prediction mechanism to evaluate the possibility of a stochastic interrupt arrival before the energy break-even time of the targeted sleep state, and use this information on state selection accordingly.

In most practical cases where idle period/duration prediction is deployed, it may be sufficient to accurately predict if the upcoming idle duration is longer than a certain threshold (i.e., instead of accurately predicting an absolute idle duration). This is due to the fact that the decision for a processor or platform to enter a sleep state may be based on the comparison of the predicted idle duration with the energy break-even time (EBT) of each sleep state.

In particular, a device might not be able to enter a sleep state if a projected idle period is shorter than the energy break-even time for the low power state in question. In this regard, there may be transition-related power consumption overhead as well as latency overhead that could offset the advantages associated with the reduced power consumption of a sleep state if the idle period is not long enough. Thus, the energy break-even time can generally indicate the amount of time the component would have to remain idle in order to justify transitioning to and from the low power state under consideration. Techniques described herein use a probability-based idle-duration prediction approach that considers the threshold-based nature of the sleep state selection process.

Idle Duration Predictions

First, the expected arrival time for the next deterministic interrupt can be calculated as the minimum of the arrival time for all deterministic interrupts considered

$$\hat{t}_r = \min\{\hat{t}_1, \hat{t}_2, \dots, \hat{t}_n\} \quad (2)$$

where \hat{t}_k , $k=1, \dots, n$ is the expected time to arrival for the deterministic interrupt from

source k.

Then, the probability of having an interrupt within the next τ time period, denoted as $\text{Prob}(x)$, can be calculated by considering the following two cases:

1. If τ is larger than the time to arrival for the nearest deterministic interrupt, then the probability is one:

$$\text{If } \hat{t}_k < \tau, \text{ then } \text{Prob}(x) = 1$$

2. If τ is smaller than the time to arrival for the nearest deterministic interrupt, the probability is defined by the probability of the stochastic interrupt assuming a Poisson process is deemed appropriate for modeling

$$\text{If } \hat{t}_k > \tau, \text{ then } \text{Prob}(x) = 1 - e^{-\lambda_{ste}\tau}$$

The outlined probabilistic estimation technique can help transform the more challenging problem of precise idle duration estimation into a more relaxed sleep state selection problem. Finding a near-optimal solution to the latter may be more feasible, because of the differences in the potential of making an erroneous estimation. In the following subsection, a reward-based sleep state selection scheme is described that relies on the described probabilistic estimation approach.

Reward-based Sleep State Selection

Due to the relatively high likelihood of making an erroneous idle time estimation using standard estimation techniques, a sleep state selection algorithm based on such estimation might perform sub-optimally. Therefore, a new decision metric (called "reward" herein) may be defined, wherein the reward decision metric incorporates the potential cost of false entry and missed opportunity in the sleep state decision, and uses the probabilistic estimation scheme described earlier to optimize the reward. In this context, reward may represent any measure of performance of particular interest to the system designer. For example, false entry or missed opportunity metrics can be proxies for potential loss in energy or throughput performance. In particular, when false entry of a sleep state occurs, the platform may consume more power compared to staying at the active state. Similarly, when a missed opportunity of entering a sleep state occurs, the platform may consume more power compared to entering the sleep state in question. Thus, the proposed reward-based sleep state selection scheme may be highly flexible compared to the conventional threshold-based sleep state selection, which relies solely on idle duration prediction. Once the reward metric is defined, the selection process may chose an optimal sleep state (iOPT) that maximizes the expected reward.

In particular, FIG. 5 shows a method 20 of selecting sleep states for a platform, processor, etc. 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.

Illustrated processing block 22 provides for selecting the next deepest sleep state for consideration. A determination may be made at block 24 as to whether a deterministic interrupt is expected to occur before the energy break even time for the sleep state under consideration. Thus, block 24 may involve determining the idle period for one or more deterministic interrupts, and determining the EBT for the sleep state currently under consideration. If the EBT exceeds the idle period (e.g., a deterministic interrupt is expected before the EBT), the illustrated process returns to block 22 and effectively bypasses the reward determination for the sleep state in question.

If no deterministic interrupts are expected before the EBT, block 26 may make one or more stochastic probability determinations. In particular, block 26 can involve determining a false entry probability (Prob_{FE}) and a missed opportunity probability (Prob_{Mo}). For example, the false entry probability and missed opportunity probability may be calculated for a state "i" as:

$$\text{Prob}_{\text{FE}}(i) = 1 - \exp(-\lambda \cdot \text{EBT}(i)) \quad (3)$$

$$\text{Prob}_{\text{MO}}(i) = \exp(-\lambda \cdot \text{EBT}(i+1)) \quad (4)$$

where λ is the estimated stochastic interrupt arrival rate, $\text{Prob}_{\text{FE}}(i)$ is the percent probability that a stochastic interrupt occurs before the EBT, and $\text{Prob}_{\text{MO}}(i)$ is the percent probability that a stochastic interrupt occurs after the EBT of the next deeper sleep state.

Moreover, block 28 may conduct a reward determination for the sleep state under consideration based on the stochastic probabilities. An example of a reward determination is:

$$\text{Reward}(i) = \text{Prob}_{\text{FE}}(i) \cdot \text{Pen}_{\text{FE}}(i) + \text{Prob}_{\text{MO}}(i) \cdot \text{Pen}_{\text{MO}}(i) + (1 - \text{Prob}_{\text{FE}}(i) - \text{Prob}_{\text{MO}}(i)) \cdot \text{Gain}(i) \quad (5)$$

where $\text{Pen}_{\text{FE}}(i)$ and $\text{Pen}_{\text{MO}}(i)$ are the penalties associated with a false entry and a missed opportunity, respectively, and $\text{Gain}(i)$ is the expected gain from entering state i without the occurrence of a false entry or a missed opportunity (e.g., correct state).

Illustrated block 30 determines whether the deepest sleep state (e.g., state "«") has been reached. If so, the most rewarding sleep state may be selected at block 32. In particular, block 32 may involve selecting the sleep state that has the highest reward metric.

Turning now to FIG. 6, a platform 34 is shown. The platform 34 may be part of 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). In the illustrated example, the platform 34 includes a processor 36, an integrated memory controller (IMC) 38, an IO module 40, system memory 42, a network controller (e.g., network interface card) 44, an audio IO device 46 and a solid state disk (SSD) 48. The processor 36, which may include a core region with one or several processor cores 50, may use a power management unit (PMU) 52 to place its cores 50 and other platform components into one or more active and/or sleep states based on performance and/or power management concerns, as already noted.

The illustrated IO module 40, sometimes referred to as a Southbridge or South Complex of a chipset, functions as a host controller and communicates with the network controller 44, which could provide off-platform communication functionality for a wide variety of purposes such as, for example, cellular telephone (e.g., W-CDMA (UMTS), CDMA2000 (IS-856/IS-2000), etc.), WiFi (e.g., IEEE 802.11, 1999 Edition, LAN/MAN Wireless LANS), 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 40 may also include one or more wireless hardware circuit blocks to support such functionality.

The SSD 48 may include one or more NAND 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 40 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 48 could also be used as a USB (Universal Serial Bus, e.g., USB Specification 3.0, USB Implementers Forum) flash storage device.

The illustrated IO module 40 includes a reward module 54 having logic configured to conduct a reward determination for a plurality of sleep states to obtain a plurality of reward determinations with respect to the platform 34. The logic of the reward module 54 may also select a sleep state for the system from the plurality of sleep states based at least in part on the plurality of reward determinations, wherein the PMU 52 of the processor 36 may place the platform 34 in the selected sleep state. In addition, the reward module 54 may distinguish one or more deterministic interrupts from one or more stochastic interrupts in making the reward

determinations, wherein devices such as the network controller 44, audio 10 device 46, and/or SSD 48 may be interrupt sources. The reward module 54 may be implemented elsewhere in the platform 34 such as, for example, on the processor 36.

Thus, techniques described herein may capture the characteristics of deterministic
5 interrupts (e.g., timer, audio), and exploit related information available in the system to improve interrupt prediction accuracy. Moreover, the sleep state selection process described herein can make platforms more energy efficient by enabling the platform to make better decisions. In addition, techniques described herein can enable platforms to scale their power consumption based on workload, even when the workloads are very active and have relatively short idle
10 durations. Additionally, idle duration predictions may provide greater flexibility in designing solutions for system performance optimization.

Embodiments may therefore provide for a system having a controller with logic to conduct a reward determination for a plurality of sleep states to obtain a plurality of reward
15 determinations with respect to the system. The logic may also select a sleep state from the system from the plurality of sleep states based at least in part on the reward determinations. In addition, the system may have a processor with a power management unit to place the system in the selected sleep state.

Embodiments may also include a non-transitory computer readable storage medium having a set of instructions which, if executed by a processor, cause a computer to conduct a
20 reward determination for a plurality of sleep states to obtain a plurality of reward determinations with respect to a device. The instructions can also cause a computer to select a sleep state for the device from the plurality of sleep states based at least in part on the plurality of reward determinations.

Additionally, embodiments can include an apparatus having logic to conduct a reward
25 determination for a plurality of sleep states to obtain a plurality of reward determinations with respect to a device. The logic may also select a sleep state for the device from the plurality of sleep states based at least in part on the plurality of reward determinations.

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
30 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
35 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

We claim:

- 5 1. A system comprising:
a controller having logic to,
 conduct a reward determination for a plurality of sleep states to obtain a plurality
of reward determinations with respect to the system, and
 select a sleep state for the system from the plurality of sleep states based at least
10 in part on the plurality of reward determinations; and
a power management unit to place the system in the selected sleep state.
2. The system of claim 1, wherein the logic is to,
 determine a false entry probability, and
15 determine a missed opportunity probability, wherein the reward determination is
based at least in part on the false entry probability and the missed opportunity probability.
3. The system of claim 2, wherein the false entry probability is a probability that one
or more stochastic interrupts will occur before an energy break even time associated with a sleep
20 state.
4. The system of claim 2, wherein the missed opportunity probability is a probability
that all probabilistic interrupts will occur after an energy break even time associated with a sleep
state.
- 25 5. The system of claim 2, wherein the logic is to use a stochastic process to
determine the false entry probability and the missed opportunity probability.
6. The system of claim 1, wherein the logic is to,
30 determine a false entry penalty, and
 determine a missed opportunity penalty, wherein the reward determination is
based at least in part on the false entry penalty and the missed opportunity penalty.
7. The system of claim 1, wherein the logic is to determine a correct state gain,
35 wherein the reward determination is based at least in part on the correct state gain.

8. The system of claim 1, wherein the logic is to,
distinguish one or more deterministic interrupts from one or more stochastic
interrupts,
5 determine an idle period with respect to at least one of the one or more
deterministic interrupts,
determine an energy break even time for a sleep state, and
bypass the reward determination if the energy break even time exceeds the idle
period.

10

9. The system of claim 8, wherein the logic is to identify a plurality of deterministic
interrupt sources, wherein the idle period is to be determined for two or more interrupt sources
individually.

15

10. The system of claim 9, wherein the logic is to,
group two or more of the plurality of deterministic sources together for
consideration, and
identify one or more dependencies associated with the grouped two or more of the
plurality of deterministic sources.

20

11. A non-transitory computer readable storage medium comprising a set of
instructions which, if executed by a controller, cause a computer to:
conduct a reward determination for a plurality of sleep states to obtain a plurality of
reward determinations with respect to a device; and
25 select a sleep state for the device from the plurality of sleep states based at least in part on
the plurality of reward determinations.

30

12. The medium of claim 11, wherein the instructions, if executed, cause a computer
to:
determine a false entry probability; and
determine a missed opportunity probability, wherein the reward determination is
conducted based at least in part on the false entry probability and the missed opportunity
probability.

35

13. The medium of claim 12, wherein the false entry probability is a probability that

one or more stochastic interrupts will occur before an energy break even time associated with a sleep state.

14. The medium of claim 12, wherein the missed opportunity probability is a
5 probability that all probabilistic interrupts will occur after an energy break even time associated with a sleep state.

15. The medium of claim 12, wherein the instructions, if executed, cause a computer
to use a stochastic process to determine the false entry probability and the missed opportunity
10 probability.

16. The medium of claim 11, wherein the instructions, if executed, cause a computer
to:

determine a false entry penalty; and

15 determine a missed opportunity penalty, wherein the reward determination is based at least in part on the false entry penalty and the missed opportunity penalty.

17. The medium of claim 11, wherein the instructions, if executed, cause a computer
to determine a correct state gain, wherein the reward determination is based at least in part on the
20 correct state gain.

18. The medium of claim 11, wherein the instructions, if executed, cause a computer
to:

distinguish one or more deterministic interrupts from one or more stochastic interrupts;

25 determine an idle period with respect to at least one of the one or more deterministic interrupts;

determine an energy break even time for a sleep state; and

bypass the reward determination if the energy break even time exceeds the idle period.

19. The medium of claim 18, wherein the instructions, if executed, cause a computer
to identify a plurality of deterministic interrupt sources, wherein the idle period is to be
determined for two or more interrupt sources individually.

20. The medium of claim 19, wherein the instructions, if executed, cause a computer
35 to:

group two or more of the plurality of deterministic sources together for consideration;
and
identify one or more dependencies associated with the grouped two or more of the
plurality of deterministic sources.

5

21. An apparatus comprising:
logic to,

conduct a reward determination for a plurality of sleep states to obtain a plurality
of reward determinations with respect to a device,

10 select a sleep state for the device from the plurality of sleep states based at least in
part on the plurality of reward determinations.

22. The apparatus of claim 21, wherein the logic is to,
determine a false entry probability, and

15 determine a missed opportunity probability, wherein the reward determination is
conducted based at least in part on the false entry probability and the missed opportunity
probability.

23. The apparatus of claim 22, wherein the false entry probability is a probability that
20 one or more stochastic interrupts will occur before an energy break even time associated with a
sleep state.

24. The apparatus of claim 22, wherein the missed opportunity probability is a
probability that that all probabilistic interrupts will occur after an energy break even time
25 associated with a sleep state.

25. The apparatus of claim 22, wherein the logic is to use a stochastic process to
determine the false entry probability and the missed opportunity probability.

30 26. The apparatus of claim 21, wherein the logic is to,
determine a false entry penalty, and
determine a missed opportunity penalty, wherein the reward determination is
based at least in part on the false entry penalty and the missed opportunity penalty.

35 27. The apparatus of claim 21, wherein the logic is to determine a correct state gain,

wherein the reward determination is based at least in part on the correct state gain.

28. The apparatus of claim 21, wherein the logic is to,
distinguish one or more deterministic interrupts from one or more stochastic
5 interrupts,
determine an idle period with respect to at least one of the one or more
deterministic interrupts,
determine an energy break even time for a sleep state, and
bypass the reward determination if the energy break even time exceeds the idle
10 period.

29. The apparatus of claim 28, wherein the logic is to identify a plurality of
deterministic interrupt sources, wherein the idle period is to be determined for two or more
interrupt sources individually.

30. The apparatus of claim 29, wherein the logic is to,
group two or more of the plurality of deterministic sources together for
consideration, and
15 identify one or more dependencies associated with the grouped two or more of the
plurality of deterministic sources.
20

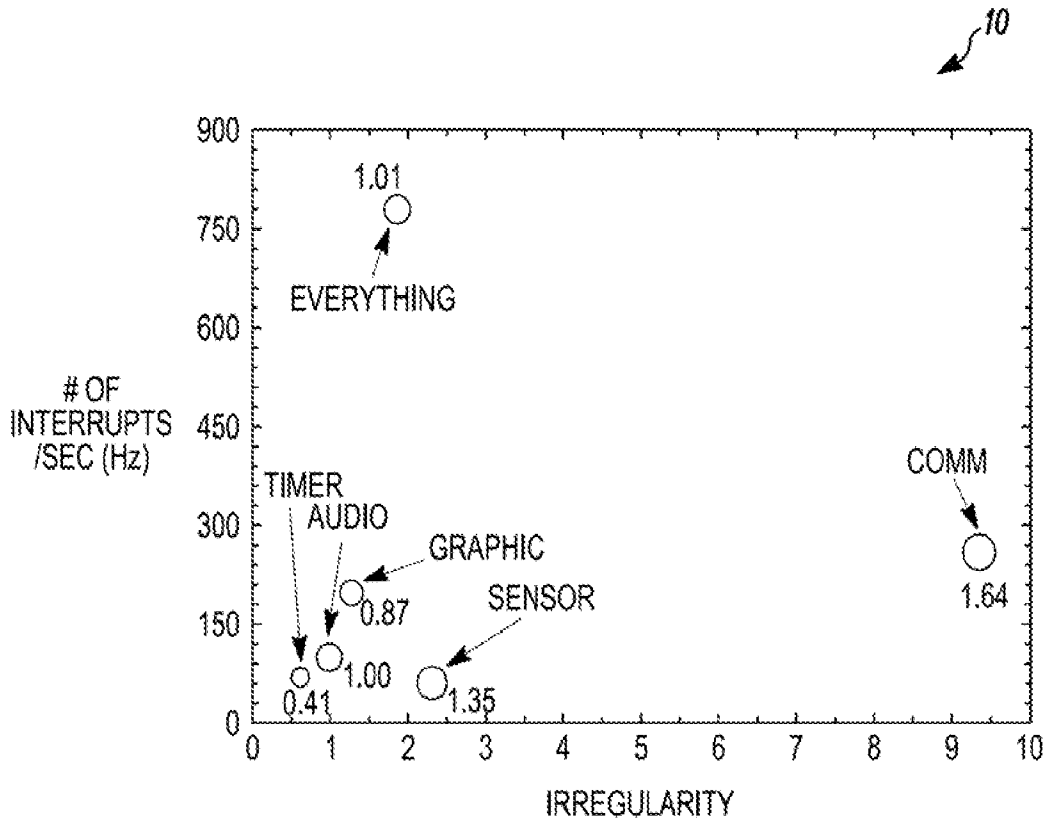


FIG. 1




FIG. 2




FIG. 3




FIG. 4

3/4

20

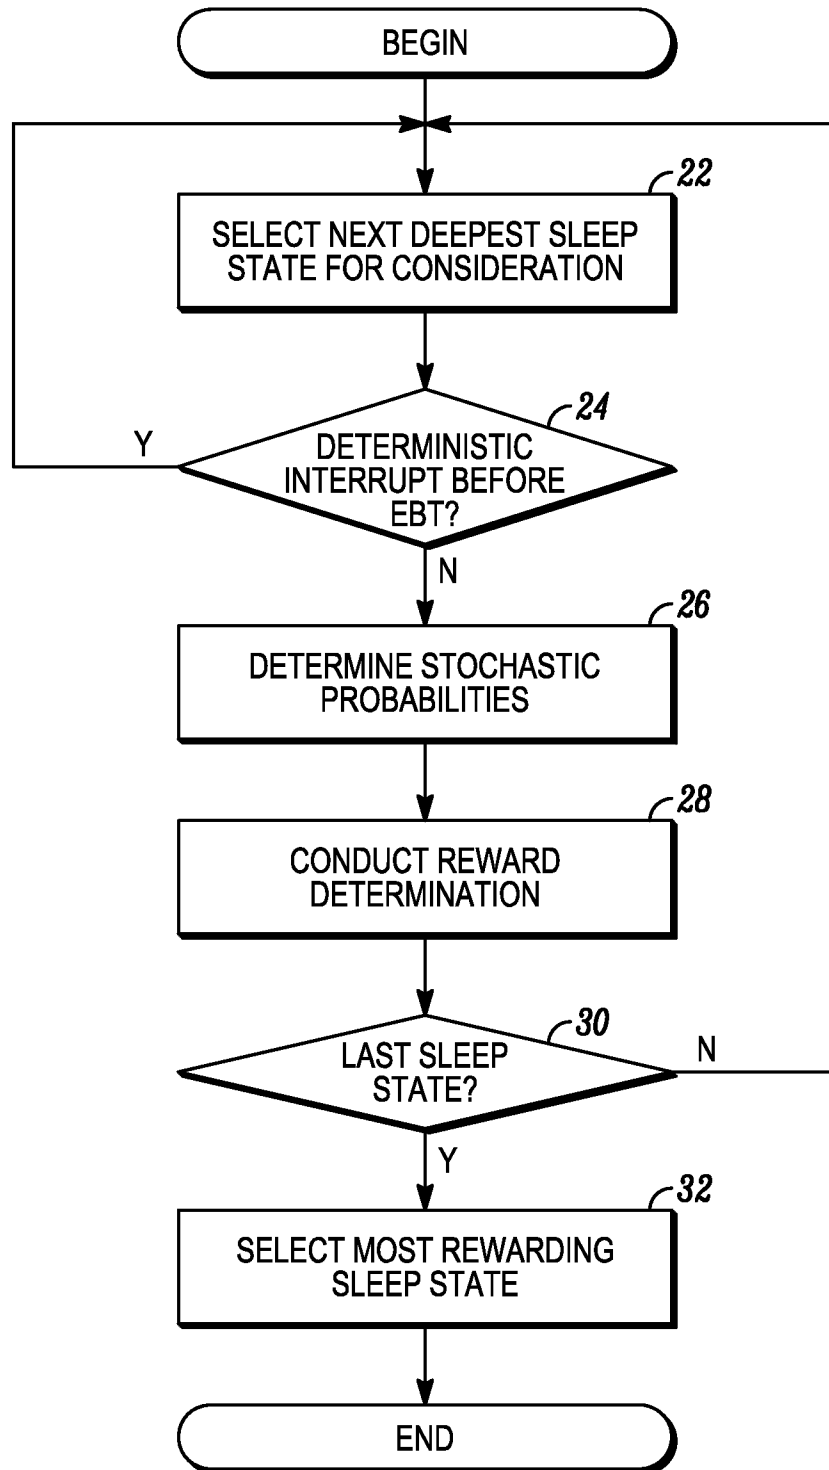


FIG. 5






FIG. 6

A. CLASSIFICATION OF SUBJECT MATTER		
<i>G06F 1/32(2006.01)i</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		
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: sleep state, interrupt, power management, interrupt, reward determination		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EUI-YOUNG CHUNG et al. "Dynamic power management using adaptive learning tree" In: Computer-Aided Design, November 1999, pp. 274-279. See pp. 275 and 279.	1-30
A	QIAN DIAO et al. "Prediction of CPU idle-busy activity pattern." In: High Performance Computer Architecture, 16-20 February 2008, pp. 27-36. See fig. 1 and p. 27.	1-30
A	REN WANG et al. "Reducing power consumption for mobile platforms via adaptive traffic coalescing." In: IEEE Journal on Selected Areas in Communications, September 2011, 29(8), pp. 1618-1629. See p. 1620.	1-30
A	US 2010-0169683 A1 (WANG REN et al.) 01 July 2010 See paras. [0012], [0016], [0023]-[0027], and figs. 1, 3.	1-30
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input checked="" type="checkbox"/> See patent family annex.		
<p>* Special categories of cited documents:</p> <p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"E" earlier application or patent but published on or after the international filing date</p> <p>"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)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p> <p>"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</p> <p>"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</p> <p>"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</p> <p>"&" document member of the same patent family</p>		
Date of the actual completion of the international search 21 SEPTEMBER 2012 (21.09.2012)		Date of mailing of the international search report 24 SEPTEMBER 2012 (24.09.2012)
Name and mailing address of the ISA/KR  Korean Intellectual Property Office 189 Cheongsu-ro, Seo-gu, Daejeon Metropolitan City, 302-701, Republic of Korea Facsimile No. 82-42-472-7140		Authorized officer Park Ji Eun Telephone No. 82-42-481-5696 

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2011/067377

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2010-0169683 A1	01.07.2010	CN 101916137 A JP 2010-191945 A TW 201037506 A	15.12.2010 02.09.2010 16.10.2010