


(51) International Patent Classification:
G06F 1/32 (2006.01)(71) Applicant (for JP only): **TEXAS INSTRUMENTS JAPAN LIMITED** [JP/JP]; 24-1, Nishi-Shinjuku 6-chome, Shinjuku-Ku Tokyo 160-8366 (JP).(21) International Application Number:
PCT/US2013/024474(72) Inventors: **LEE, Sejoong**; 952 Redbird Lane, Allen, TX 75013 (US). **CHOI, Soon-hyeok**; 1126 Nick Circle, Allen, TX 75013 (US). **LU, Xiaolin**; 4567 Kentucky Drive, Plano, TX 75024 (US).(22) International Filing Date:
1 February 2013 (01.02.2013)(74) Agents: **FRANZ, Warren, L. et al.**; **TEXAS INSTRUMENTS INCORPORATED**, Deputy General Patent Counsel, P.O. Box 655474, Mail Station 3999, Dallas, TX 75265-5474 (US).(25) Filing Language:
English

(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, 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, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ,

(26) Publication Language:
English(30) Priority Data:
13/363,949 1 February 2012 (01.02.2012) US(63) Related by continuation (CON) or continuation-in-part (CIP) to earlier application:
US **13/363,949 (CON)**
Filed on **1 February 2012 (01.02.2012)**(71) Applicant: **TEXAS INSTRUMENT INCORPORATED** [US/US]; P.o. Box 655474, Mail Station 3999, Dallas, TX 75265-5474 (US).*[Continued on next page]*

(54) Title: DYNAMIC POWER MANAGEMENT IN REAL-TIME SYSTEMS



(57) Abstract: A Dynamic Sleep Controller (220) reduces power consumption by a processor (CPU 202) in a computer system by determining a maximum number of times (token count) that the processor (CPU 202) can incur a start-up delay after being placed into a low-power mode during a token period of time when executing a task for a token period of time. The processor (CPU 202) may be placed into the low-power mode by a sleep request signal (222) sent to a power and clock controller (230) while the processor (CPU 202) is executing the task in response to an idle indicator only if a current value of the token count assigned to the task is greater than zero. The current value of the token count is decremented each time the processor incurs a start-up delay in response to being awakened from the low-power mode. The current token count is reset to match the assigned token count at the end of each token period. Furthermore, wakeup may be anticipated to allow the processor (CPU 202) to be awakened preemptively by a preemptive wakeup signal (224) sent to an interrupt controller (210) which sends a wakeup request (214) to a power and clock controller (230).

FIG. 2



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

Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

- *with international search report (Art. 21(3))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

DYNAMIC POWER MANAGEMENT IN REAL-TIME SYSTEMS

[0001] This invention generally relates to microcontrollers for real-time control systems, and in particular to reducing power used by the microcontroller.

BACKGROUND

[0002] Computer processors used in real-time control systems are typically fabricated using a complimentary metal-oxide semiconductor (CMOS) process. In CMOS circuits, power consumption is proportional to the product of the frequency and the square of the supply voltage. A reduction in the operating frequency of the processor and/or its supply voltage can lead to significant savings in energy consumption and heat dissipation. In a real-time system, the processor may be placed into a low power mode (LPM) during idle periods when a task is not being executed. While in the LPM, the processor consumes less power. If the processor receives a request to resume execution, the processor may be awakened from the LPM and the processor will return to a normal power state and resume processing.

BRIEF DESCRIPTION OF THE DRAWINGS

[0003] FIGS. 1A-1C compare CPU activity with conventional and dynamic power management;

[0004] FIG. 2 is a block diagram illustrating dynamic sleep controller (DSC) based CPU power management;

[0005] FIGS. 3A-3C are timing diagrams illustrating use of tokens for dynamic power management;

[0006] FIG. 4 is an illustration of a sleep decision map;

[0007] FIG. 5 is an illustration of a wake-up time-slot decision map;

[0008] FIG. 6 illustrates formation of a wake-up request histogram;

[0009] FIG. 7 is a block diagram of a performance monitoring unit that may be included with the microcontroller of FIG. 2;

[0010] FIG. 8 is a block diagram of an exemplary SoC that includes dynamic power management; and

[0011] FIG. 9 is a flow diagram illustrating operation of dynamic power management.

DETAILED DESCRIPTION OF EXAMPLE EMBODIMENTS

[0012] A method is described herein that efficiently utilizes a low-power mode (LPM) of a processor in order to reduce the power consumption of the processor. When a processor is idle, the processor may be put into LPM during which the processor provides low performance or even zero performance, while consuming less power. When the processor receives a request to resume execution, the processor is awakened from the LPM and the processor will return to the normal power state and resume the processing.

[0013] A problem of using the LPM is that it takes some time to wake up the processor from the LPM to normal state. Such latency causes processing delay of tasks being executed by the processor. In a real-time system with hard task deadlines, the processing delays can affect adversely the overall system performance.

[0014] A method will now be described that bounds the overall processing delay in order to limit the performance degradation while maximizing the use of LPM of a processor. The dynamic power management system described herein allows use of LPM even while a real-time task is actively running. A token scheme is used that can bound the performance degradation by wake-up latency such that the tasks running in the CPU still meet their deadlines. A prediction scheme may also be used that can maximize the use of LPM by monitoring the activity pattern of the CPU.

[0015] Use of a low-power mode (LPM) involves state transition overhead. The latency taken by a state change from LPM to normal state as referred to as wake-up latency. The wake-up latency depends on what kinds of mechanism are used in the LPM. For example, if the LPM uses a clock-gating feature only, the wake-up process involves releasing a clock gating signal; therefore, the wake-up latency can be as low as a couple of clock cycles. On the other hand, if the LPM uses power-gating, in which the processor saves internal contexts in another secure or non-volatile memory space during the transition from normal to the LPM, the wake-up latency involves the context-

restoring time which may take a couple of milliseconds. In general, a LPM that consumes less power requires longer wake-up latency.

[0016] If the wake-up latency of a LPM is negligible, the processor may be put into LPM whenever the processor is idle, and wake it up in the event of a wake-up request, such as a wake-up interrupt. In general, however, the wake-up latency is not negligible, and the latency can adversely impact the overall performance of the system as the latency causes delay in task execution time of the processor.

[0017] As a result, previously a typical scenario for power conservation put the processor into LPM only when the application running in the processor is in such state that it is obvious that the wake-up latency can be tolerated at the time of wake-up request. In other words, the LPM is used only when a realtime task is not actively running.

[0018] FIGS. 1A-1C compare CPU activity with conventional and dynamic power management. FIG. 1A illustrates an execution pattern for an illustrative realtime task. During idle periods 102, 106, the CPU is not executing an active task. During active period 104, the CPU is executing a task, but there may still be periods of idle time during the active period 104.

[0019] FIG. 1B illustrates an example of conventional use of LPM, in which LPM is used only when the application is obviously inactive, such as during idle periods 102, 106. While the task is active, the processor remains in the normal power state 114.

[0020] FIG. 1C illustrates a method that makes use of LPM modes even when an application is actively running in the processor, as indicated generally at 124.

[0021] FIG. 2 is a block diagram of an example embodiment of the invention. A Dynamic Sleep Controller (DSC) 220, which can be implemented in software or hardware, monitors the *cpu_idle* signal 204 which is asserted when CPU 202 is idle. This signal is typically asserted when the processor executes an idle instruction that is included in the task being executed or by an operating system that is scheduling task execution. When the *cpu_idle* is asserted, DSC 220 determines if it is good time to go to LPM or it is better to stay in normal state. If DSC 220 determines that it is a good time to go to LPM, it asserts *sleep_req* signal 222. Then, power and clock controller 230 will put

the CPU into LPM by changing clock signal 232 and/or power supply 234 voltage levels accordingly.

[0022] If interrupt controller 210 detects any unmasked functional interrupt 212, it asserts wakeup_req signal 214, and power and clock controller 230 puts CPU 202 back to a normal power state by restoring clock signal 232 and/or power supply levels 234. As mentioned earlier, CPU 202 may also need to perform a context restore or other housekeeping task before returning to execution of the task. CPU 202 may also be preemptively awakened by DSC 220. DSC 220 can wake up CPU 202 by asserting preemptive wakeup_int signal 224. If the CPU is awakened by preemptive wakeup interrupt 224, and if it then receives another functional interrupt 212 after that, the CPU will be able to process the functional interrupt without incurring any wakeup delay since the wakeup latency occurred in response to the preemptive wakeup request before receiving the functional wakeup request. The DSC may use this proactive wake-up mechanism when it expects a functional interrupt to occur soon.

[0023] DSC 220 may perform several dynamic power management functions. When cpu_idle is asserted 204, the DSC determines whether to go to LPM or not. If the DSC determines to choose LPM, it may also determine when it will wake up the CPU. The CPU will experience processing delay due to wake-up latency if a functional interrupt 212 occurs before preemptive wakeup_int 224 is asserted, but the CPU will not experience wake-up latency if a functional interrupt 212 occurs after preemptive wakeup_int 224 is asserted.

[0024] DSC 220 manages accumulation of processing delay caused by the wake-up latency so that wake-up latency does not exceed a specified threshold. In this manner, tasks running in the processor meet their deadlines. A Quality-of-Service (QoS) Policy 226 that is available to DSC 220 contains information specific to the application.

[0025] The DSC may also try to maximize the use of LPM. MinEnergy Policy 228 contains a set of instructions that attempt to use LPM as much as possible as long as the QoS Policy is met. As discussed above, the instructions may preemptively wake up the CPU ahead of a next functional wake-up request in order to eliminate wake-up latency. Also, the instructions may try to wake up the CPU as late as possible in order to

maximize the sleep time. The MinEnergy Policy is generated during run-time based on an activity pattern of the CPU and QoS Policy.

[0026] FIG. 2 illustrates an embodiment in which the DSC manages only one CPU. In other embodiments, a DSC may manage multiple CPUs as long as the DSC has enough bandwidth and resources to handle that many CPUs. In some multi-CPU architectures, it may be beneficial to implement signals such as sleep_req 222, preemptive wakeup_int 224, and cpu_idle 204 using bus transaction commands delivered through a bus interconnect rather than dedicated hardwire signals in order to minimize signal routing and to provide better scalability.

Quality of Service Policy

[0027] An embodiment of QoS Policy and MinEnergy Policy will now be described in detail. The case where the CPU is in LPM and is awakened up by a functional interrupt rather than a preemptive wakeup_int will be referred to herein as mis-prediction. In each mis-prediction occurrence, the DSC has failed to predict the correct wake-up time and the DSC did not awaken the CPU promptly. When mis-prediction occurs, the task processing will be delayed by the amount of wake-up latency. The QoS Policy describes how many times and how often such mis-prediction can be tolerated by the system.

[0028] Table 1 shows an example of QoS Policy 226. It includes multiple token constraints. A token constraint includes two parameters: the number of tokens and a period. One token means that the application running in the CPU can tolerate the processing delay caused by one occurrence of wake-up latency during each specified period. For example, token constraint 1 in Table 2 indicates that the system can tolerate two occurrences of wake-up latency for every 1ms token period. Therefore, if the system has experienced two occurrences of wake-up latency within 1ms, the DSC will not use LPM again during that token period in order to avoid any additional processing delay. During the next 1ms token period, the DSC may use the LPM again. However, it can now tolerate only one occurrence of wake-up latency during the next 1ms because of the token constraint 2. Thus, a hierarchical token constraint policy allows a number

of wake-up latencies to be specified for a longer period such as 10ms, while preventing all of the allowable wake-up latencies from being incurred in a short period.

Token constraint ID	Number of tokens	Token Period
1	2	1ms
2	3	2ms
...
N	5	10ms

Table 1 - Token Constraints

[0029] The token-based QoS Policy described above is effective to bound the performance degradation caused by the use of LPM. Soft real-time applications, which have throughput requirements, can use the token-based QoS Policy while maintaining the performance above a certain threshold.

[0030] The token-based QoS Policy is also very effective to guarantee deadline of real-time tasks. In real-time systems, a task τ_i is typically modeled by (C, T, D) parameters, where C is worst-case execution time, T is the period of the task, and D is the relative-deadline of τ_i respectively. This modeling may be done using known simulation and analysis techniques.

[0031] Once the C, T, and D parameters are determined for each task, a token constraint may be computed for each task by analyzing slack time of each task. Slack time is the amount of time left over after the worst-case execution time. The slack time of each task can be directly translated into the number of tokens by dividing the slack time by the LPM wake-up latency. As discussed earlier, the LPM wake-up latency depends on how deeply the processor has been put to sleep. The token period may be selected to be the same as the T parameter of the task. Each task will produce one token constraint. After the analysis of a given set of tasks, a set of tuple (TK_i, T_i) is obtained, where TK_i = the number of Tokens for task i, and T_i = the parameter T of task i. If $TK_i \leq TK_j$ and $T_i \geq T_j$, (TK_i, T_i) is more strict constraint than (TK_j, T_j) , then (TK_j, T_j) may be ignored.

Minimum Energy Policy

[0032] FIGS. 3A-3C are timing diagrams illustrating use of tokens for dynamic power management. MinEnergy Policy 228 describes an optimal timing for use of LPM in order to maximize sleep time of CPU 202. For this example, assume a token constraint (3, 10ms). FIG. 3A illustrates one 10ms token period of task execution. During this token period, there are four bursts of task execution activity 301-304 interspersed with idle periods 311-313. The Task activity must be completed prior to deadline time 320. Slack time 314 is the idle period between the end of task processing 304 and deadline time 320. In this example, the wake-up latency can occur at least three times without deadline time 320 being exceeded.

[0033] In FIG. 3B, three occurrences of wake-up latency 331-333 have occurred in the middle of a 10ms token period, therefore, the CPU cannot be placed in LPM again during this token period and therefore remains in the normal power state during the remainder 335 of the 10ms duration.

[0034] On the other hand, FIG. 3C shows another way of using LPM, in which the DSC is able to use LPM for a longer amount of time and thereby save additional power. In this example, when first activity burst 301 finishes, the DSC determines to stay at the normal state. As a result, second activity burst 302 does not cause wake-up latency. As a result, at the end of second activity burst, the number of remaining tokens is still two, since one token was used during the wakeup for activity pattern 301. A second token is used 343 during the wake-up for third activity pattern 303.

[0035] DSC 220 successfully wakes up 344 the CPU ahead of fourth activity burst 304. Therefore, wake-up latency did not cause processing delay. This means that the number of remaining tokens is still one.

[0036] When the fourth activity burst 305 finishes, DSC 220 still has one remaining token. As a result, it can use LPM during the rest of 10ms period 345.

[0037] Thus, as illustrated in FIG. 3B, it can be seen that using tokens alone to keep track of how often the CPU can be placed in LPM during any token period leads to power savings. However, as illustrated in FIG. 3C, when a minimal energy policy also specifies optimized timing to choose to go to LPM and wake up depending on the

number of remaining tokens, an increased power savings may result. A MinEnergy Policy may be defined as follows:

- for a given remaining token count, the MinEnergy Policy tells whether to stay at normal power state or go to LPM at each time when CPU becomes idle;
- if MinEnergy Policy tells the CPU to go to LPM, it may also tell when to wake the CPU up.

[0038] In one embodiment, a MinEnergy Policy may be implemented with two tables referred to as a 'Sleep Decision Map (SDM)' 400 illustrated in FIG. 4 and a 'Wakeup_time_slot Decision Map (WDM)' 500 illustrated in FIG. 5. The SDM and WDM divide a token period into smaller pieces of time durations called time-slot. The SDM describes the condition to go to LPM for each time-slot. The condition is the number of tokens remaining. For example, if the CPU becomes idle during time-slot 2 (402) and Min_token [2] = 3, the DSC may decide to go to LPM only if the number of tokens remaining is equal to or more than three.

[0039] DSC 220 uses WDM 500 to decide when to wake the CPU up for each time-slot. Once the DSC decides to go to LPM based on SDM 400, it refers to WDM 500 in order to decide when to wake the CPU up. WDM 500 is a two-dimensional array. For each entry in the array, wakeup_time_slot[w][x] is a time-slot value at which the DSC should wake the CPU if the DSC decides to use LPM at time-slot x and the number of remaining tokens is (w+1).

Time-slot based CPU activity monitoring

[0040] The MinEnergy Policy including SDM 400 and WDM 500 depends on knowing a CPU activity pattern. FIG. 6 illustrates a procedure of monitoring the activity pattern and gathering statistics. In the example of FIG. 6, the Idle-to-Active event of the CPU is monitored. This is the event that will wake up the CPU if it was in LPM mode, which is wake-up request 214, referring back to FIG. 2.

[0041] A token period 530 is pre-defined, and the period is divided into time-slots. If a wake-up request occurs at a certain time-slot, a counter corresponding to the time-slot is incremented. For example, wake-up request 545 is recorded for time slot 17 as

indicated at 517a. It should be noted that a finer time-slot means finer resolution of DSC operation.

[0042] The counting process is repeated for every period to form a histogram over a specified number of periods. For example, during next token period 531, wake-up event 550 is recorded for slot 17 by incrementing the count to two, as indicated at 517b. After a certain amount of time, the counting process stops, and the final counter value of each time slot is translated into a probability of a wake-up request occurring in that time slot. The resulting set of wake-up probabilities for the time-slots is referred to herein as a wake-up probability map.

[0043] Note that the wake-up probability map represents the history of the task activity pattern and the hypothesis is that such statistics of a past activity pattern will be somewhat valid for prediction of near future behavior of the system.

[0044] The period of the wake-up probability map should be selected such that the period is somewhat synchronous to wake-up requests, otherwise, the probability distribution of wake-up requests may be flat.

[0045] In case of multiple token constraints having different periods, a wake-up probability map may be built for each token constraint. In this case, the different token constraints may be correlated to different tasks that are being executed on the processor. A task ID may be provided to the DSC by the processor to identify what task is currently being executed, as indicated by task ID signals 206 in FIG. 2. Alternatively, processor 202 may write the task ID to DSC 220 over a bus interconnect, for example.

[0046] In another embodiment, a single wake-up probability map may be built as long as the map captures wake-up requests synchronously. For instance, a least common multiple of token periods may be used.

SDM and WDM determination

[0047] If DSC puts CPU into LPM at time-slot x and sets wake-up for time-slot x+a, the energy-saving that can be expected is given by equation (1).

Energy-saving expectation =

$$\begin{aligned}
 & E_{LPM} * (1-P_x) + \\
 & E_{LPM} * (1-P_x) * (1-P_{x+1}) + \\
 & E_{LPM} * (1-P_x) * (1-P_{x+1}) * (1-P_{x+2}) + \dots + \\
 & E_{LPM} * (1-P_x) * (1-P_{x+1}) * (1-P_{x+2}) * \dots * (1-P_{x+a-1}) \quad (1)
 \end{aligned}$$

where:

E_{LPM} = Energy-saving during one time-slot by putting CPU into LPM

P_k = Probability of wake-up request in time-slot k

[0048] The probability P that the token is not consumed so that it can be used later until the end of the period is given by equation (2).

$$P = (1-P_x) * (1-P_{x+1}) * (1-P_{x+2}) * \dots * (1-P_{x+a-1}) \quad (2)$$

[0049] When multiple tokens are available, various different scenarios are possible which eventually use up all the tokens while maximizing the sleep time.

[0050] Considering all possible scenarios and probabilities associated with the scenarios, an expectation of energy-saving can be computed when the CPU is placed into LPM at time-slot x and a wake-up time-slot is set for time-slot x+a. $MinToken[x]$ and $Wakeup_time_slot[w][x]$ may be determined such that the energy-saving expectation is maximized.

Updating statistics

[0051] The SDM and WDM are computed based on the wake-up probability map. As mentioned above, the wake-up probability map represents the history of the task activity pattern, and it was assumed that such a historic pattern may be somewhat valid for prediction of near future behavior of the system. Depending on the dynamics of the system behavior, such history may valid during long or short period of time.

[0052] Subsequent probability maps and consequent SDM and WDM can be built continuously in order to represent the most up-to-date characteristics of the system behavior. However, such continuous update computation effort will use additional energy, so there may be a tradeoff between actual power savings and attempts to maximize power savings.

[0053] FIG. 7 is a block diagram of a performance monitoring unit 440 that may be included with the system of FIG. 2. In order to update the probability map, it may be desirable to have a performance monitoring unit 440 that monitors how well DSC 220 is performing. If DSC 220 is not performing well, monitoring unit 440 may trigger a reiteration of building a probability map and consequent SDM and WDM.

[0054] The performance monitoring unit may monitor several parameters in order to evaluate the performance of the DSC. The average number of times that CPU 202 receives an interrupt before DSC wakes it up may be monitored. This represents how well the DSC's wake-up mechanism is working. The total sleep time that the CPU spent in LPM vs. the total idle time of CPU may be monitored. This represents how well the DSC's sleep decision mechanism and wake-up mechanism are working.

[0055] The performance monitoring unit may alert the DSC to reiterate the probability map if the parameters described above are lower than a selected threshold by asserting alert signal 442.

[0056] FIG. 8 is a block diagram illustrating an exemplary application specific integrated circuit system on a chip (SOC) 800 with CPU 802. For purposes of this disclosure, the somewhat generic term "microcontroller" (MCU) may be used to apply to any complex digital system on a chip (SOC) that may include one or more processing modules 802, memory 850, and peripherals and/or DMA (direct memory access) controllers 860. At least a portion of memory module 850 may be non-volatile and hold instruction programs that are executed by processing module(s) 802 to perform the system applications. CPU 802 may also be coupled to a data cache and instruction cache, not shown. CPU 802 is coupled to system bus 852 for access to bulk memory 850. Peripherals 860 are also coupled to system bus 852 to allow access and control by CPU 802.

[0057] The topology and configuration of SOC 800 is strictly intended as an example. Other embodiments of the invention may involve various configurations of buses for interconnecting various combinations of memory modules, various combinations of peripheral modules, multiple processors, etc. CPU 802 may be any one of the various types of microprocessors or microcontrollers that are now known or later developed. For example, CPU 802 may be a digital signal processor, a

conventional processor, or a reduced instruction set processor. As used herein, the term "microprocessor" or CPU is intended to refer to any processor that is included within a system on a chip.

[0058] SOC 800 is coupled to realtime subsystem (RTS) 870. RTS 870 may be a motor, for example, in which case SOC 800 controls motor speed and direction by controlling the application of voltage to multiple sets of stator windings based on rotor position. In another example, RTS 870 may be a speaker for playing audio sound or music that is converted from a digital stream by SOC 800. For the purpose of the description herein, RTS 870 is any type of device or component now known or later developed that requires some form of hard real-time control.

[0059] One or more of the peripheral devices 860 may provide control signals or data signals to RTS 870 and may receive status or other information from RTS 870. For example, if RTS 870 is a motor, peripheral device 860 may receive rotor position data from RTS 870 that generates an interrupt for a new stator control setting. SOC 800 may be part of an automotive engine control system, for example and receive engine position and speed information and provide fuel and spark firing control signals.

[0060] As another example, if RTS 870 is a speaker, peripheral device 860 may provide an analog sound signal to RTS 870. Another peripheral module may be accessing a digital stream of audio data and generate an interrupt when a new frame of audio data is available. SOC 800 may be part of a mobile handset and be receiving voice and music digital signals via a cellular telephone network, for example.

[0061] DSC 820 may operate as described above in more detail to control clock and voltage levels 830 to each CPU 802. In one embodiment, there may be sixteen CPUs 802 that are provided with dynamic power management by DSC 820. In this example, DSC 820 receives idle indications from the various CPU 802 via transactions across bus 852. Similarly, DSC 820 sends preemptive wakeup interrupts to the various CPU 802 via transactions across bus 852. DSC 820 may be implemented as a hardwired logic, a software controlled microcontroller, a state machine, or other type of logic that performs the dynamic power management function as described above in more detail.

[0062] A performance monitoring unit 840 may also be coupled to DSC 820 to monitor the performance of DSC 820 for each of the various CPU 802, as described in more detail above.

[0063] FIG. 9 is a flow diagram illustrating operation of dynamic power management for dynamically reducing power consumption by a processor in a computer system. A maximum number of times that the processor can incur a start-up delay after being placed into a low-power mode during a token period of time when executing a task for a token period of time is determined 902. This is generally done while the system is being designed and configured by determining slack time of the task prior to a deadline time and by determining the wakeup latency that will be incurred based on the level of a selected low power mode. In some embodiments, there may be more than one type of low power mode that may be selected.

[0064] A token count is assigned 904 for the task in accordance with the determined maximum number of times that the processor can incur a start-up delay for the task. This is done by dividing the slack time for the task by the wakeup latency time. As described earlier, the start-up delay time may vary depending on what level of low power mode the processor is placed into.

[0065] The task is executed 906 until an idle point is reached. At this point, the processor will assert an idle signal to indicate that useful processing has stopped. As mentioned above, this may be in response to executing an idle instruction that is part of the task or part of an operating system that is controlling scheduling of the task.

[0066] The processor is placed 910 into the low-power mode while executing the task in response to the idle indicator only if a current value of the token count assigned to the task is greater than zero 908. The actual threshold number may be different than zero, depending on the activity pattern of the processor which is captured by the histogram 930. The threshold value may also depend on the current time-slot value at which the idle point has been detected. FIG 4 is an example of such threshold policy, where each $\text{Min_token}[x]$ of time-slot x is determined based on the histogram.

[0067] Execution of the task is resumed by waking 916 the processor out of the low-power mode. This generally occurs as a result of a functional interrupt produced by a device or peripheral within the system, or by a counter or timer, for example.

[0068] The current value of the token count of the task is decremented 918 each time the processor incurs a start-up delay in response to being awakened from the low-power mode.

[0069] The current token count is reset 924 to match the assigned token count at the end of each token period 922.

[0070] A time when the processor is likely to be awakened may be anticipated 912 and the processor may then be preemptively awakened 914 by a dynamic sleep controller prior to the anticipated time such that a start-up delay is not incurred. In this case, the current value of the token counter is maintained 915 when a start-up delay is not incurred when the processor starts processing. Even though the processor is likely to be awakened, the dynamic sleep controller may decide to not preemptively awaken the processor if the expected energy-saving from the small probability that the processor may not be awakened is worth possibly using a token.

[0071] A wake-up interrupt signal is asserted 916 to wake the processor out of the low power mode when the processor is in a low-power mode and to restart task execution 920 when the processor did not enter low power mode.

[0072] A histogram of wake-up interrupt occurrences during the token period may be occasionally created 930 for a plurality of time slots. Each reiteration of the histogram may be done in response to a performance alert from monitoring unit 440, for example. A sleep ratio of total sleep time in the low-power mode versus total idle time of the processor during a period of time may be determined by monitoring unit 440 and the histogram may be reset when the sleep ratio is lower than a sleep ratio threshold. The monitoring unit may determine a wakeup ratio of how many times a wakeup interrupt is received prior to preemptively waking the processor during a period of time, referred to as prediction failure, and then reset the histogram when the wakeup ratio is greater than a wakeup ratio threshold. When the monitoring unit counts the prediction failure, it may disregard a situation in which the dynamic sleep controller intentionally did not awaken the processor 912 because it was not worth saving a token.

[0073] Alternatively, the histogram may be reset at a periodic rate that is determined during system design, for example. Typically, the histogram will exclude preemptive wake-up requests generated by the dynamic sleep controller.

[0074] A time slot with a histogram value above a threshold value may be selected as the anticipated time when the processor is likely to be awakened 912.

[0075] After creating a histogram 930, the DSC may determine for each time slot of the histogram a potential energy-saving that may be achieved by saving a token for later use. While the processor is in a low-power mode, a time slot may be selected 912 as a preemptive wake-up time slot and a preemptive wakeup request is asserted 914 by the dynamic sleep controller if the histogram indicates it is likely that the processor will be awakened by a wake-up interrupt during the time-slot and if it is probable that a potential energy-saving that can be achieved by saving a token for later use is higher than an energy-saving that might be achieved by continuing the sleep at the time-slot, as explained in more detail with regard to equations (1) and (2) by using SDM 400 and WDM 500, for example.

Other Embodiments

[0076] While the invention has been described with reference to illustrative embodiments, this description is not intended to be construed in a limiting sense. Various other embodiments of the invention will be apparent to persons skilled in the art upon reference to this description. For example, a single DSC may provide dynamic power control for a single processor or for multiple processor cores within a system on a chip.

[0077] Embodiments of the dynamic sleep manager and methods described herein may be provided on any of several types of digital systems: digital signal processors (DSPs), general purpose programmable processors, application specific circuits, or systems on a chip (SoC) such as combinations of a DSP and a reduced instruction set (RISC) processor together with various specialized accelerators. A stored program in an onboard or external (flash EEPROM) ROM or FRAM may be used to implement aspects of the dynamic power management.

[0078] The techniques described in this disclosure may be implemented in hardware, software, firmware, or any combination thereof. If implemented in software, the software may be executed in one or more processors, such as a microprocessor, application specific integrated circuit (ASIC), field programmable gate array (FPGA), or digital signal processor (DSP). The software that executes the techniques may be

initially stored in a computer-readable medium such as compact disc (CD), a diskette, a tape, a file, memory, or any other computer readable storage device and loaded and executed in the processor. In some cases, the software may also be sold in a computer program product, which includes the computer-readable medium and packaging materials for the computer-readable medium. In some cases, the software instructions may be distributed via removable computer readable media (e.g., floppy disk, optical disk, flash memory, USB key), via a transmission path from computer readable media on another digital system, etc.

[0079] Those skilled in the art to which this relates will appreciate that modifications may be made to the described embodiments, and also that many other embodiments are possible, within the scope of the claimed invention.

CLAIMS

What is claimed is:

1. A method for dynamically reducing power consumption by a processor in a computer system, the method comprising:

 determining a maximum number of times that the processor can incur a start-up delay after being placed into a low-power mode during a token period of time when executing a task for the token period of time;

 assigning a token count for the task in accordance with the determined maximum number of times that the processor can incur a start-up delay for the task;

 executing the task until reaching an idle point;

 placing the processor into the low-power mode while executing the task in response to an idle indicator only if a current value of the token count assigned to the task is greater than zero;

 resuming execution of the task by waking the processor out of the low-power mode;

 decrementing the current value of the token count of the task each time the processor incurs a start-up delay in response to being awakened from the low-power mode; and

 resetting the current token count to match the assigned token count at the end of each token period.

2. The method of Claim 1, further comprising:

 anticipating a time when the processor is likely to be awakened and preemptively waking the processor by a dynamic sleep controller prior to the anticipated time such that a start-up delay is not incurred; and

 maintaining the current value of the token counter when a start-up delay is not incurred when the processor is awakened.

3. The method of Claim 2, further comprising:

asserting a wake-up interrupt signal to wake the processor out of the low power mode when the processor is in a low-power mode;

creating a histogram of wake-up interrupt occurrences during the token period for a plurality of time slots; and

wherein a time slot with a histogram value above a threshold value is selected as the anticipated time when the processor is likely to be awakened.

4 The method of Claim 3 wherein the histogram excludes preemptive wake-up requests generated by the dynamic sleep controller.

5. The method of Claim 3, further comprising:

determining a sleep ratio of total sleep time in the low-power mode versus total idle time of the processor during a period of time; and

resetting the histogram when the sleep ratio is lower than a sleep ratio threshold.

6. The method of claim 3, further comprising:

determining a wakeup ratio of how many times a wakeup interrupt is received prior to preemptively waking the processor during a period of time; and

resetting the histogram when the wakeup ratio is greater than a wakeup ratio.

7. The method of Claim 1, further comprising:

asserting a wake-up interrupt signal to wake the processor out of the low power mode when the processor is in a low-power mode;

creating a histogram of wake-up interrupt occurrences during the token period for a plurality of time slots;

determining for each time slot of the histogram a potential energy-saving that may be achieved by saving a token for later use; and

wherein while the processor is in a low-power mode, a time slot is selected as a preemptive wake-up time slot and a preemptive wakeup request is asserted by a dynamic sleep controller if the histogram indicates it is likely that the processor will be

awakened by a wake-up interrupt during the time-slot and if it is probable that a potential energy-saving that can be achieved by saving a token for later use is higher than an energy-saving that might be achieved by continuing the sleep at the time-slot.

8. The method of Claim 7 wherein the histogram excludes preemptive wake-up requests generated by the dynamic sleep controller.

9. The method of Claim 1, wherein slack time analysis is used to determine the maximum number of times the processor can incur a start-up delay for the task.

10. The method of Claim 1, wherein the computer system comprises a plurality of processor cores coupled to a single dynamic sleep controller, wherein the dynamic sleep controller is responsible for placing each of the plurality of processor cores into a low power mode while executing respective tasks.

11. A computer system, comprising:

a processor coupled to a memory, wherein the processor is configured to execute a software task stored in the memory;

an interrupt controller coupled to the processor, wherein the interrupt controller is configured to provide a wakeup request signal to a power controller, wherein the power controller is configured to place the processor into a low power mode; and

a dynamic sleep controller coupled to receive an idle signal from the processor, wherein the dynamic sleep controller is configured to:

send a signal to the power controller to place the processor into the low-power mode in response to an idle indicator only if a current value of a token count assigned to the task is greater than zero, wherein the token count indicates a maximum number of times that the processor can incur a start-up delay after being placed into a low-power mode during a token period of time when executing a task for the token period of time;

decrement the current value of the token count of the task each time the processor incurs a start-up delay in response to being awakened from the low-power mode; and

reset the current token count to match the assigned token count at the end of each token period.

12. The system of Claim 11, wherein the dynamic sleep controller is further configured to:

anticipate a time when the processor is likely to be awakened and preemptively wake the processor prior to the anticipated time such that a start-up delay is not incurred; and

maintain the current value of the token counter when a start-up delay is not incurred when the processor is awakened.

13. The system of Claim 12, wherein the dynamic sleep controller is further configured to:

create a histogram of wake-up interrupt occurrences during the token period for a plurality of time slots; and

wherein a time slot with a histogram value above a threshold value is selected as the anticipated time when the processor is likely to be awakened.

14. The system of Claim 13, wherein the dynamic sleep controller is further configured to exclude preemptive wake-up requests generated by the dynamic sleep controller while creating the histogram.

15. The system of Claim 11, wherein the dynamic sleep controller is further configured to:

create a histogram of wake-up interrupt occurrences during the token period for a plurality of time slots;

determine for each time slot of the histogram a potential energy-saving that may be achieved by saving a token for later use; and

wherein while the processor is in a low-power mode, a time slot is selected as a preemptive wake-up time slot and a preemptive wakeup request is asserted by a dynamic sleep controller if the histogram indicates it is likely that the processor will be

awakened by a wake-up interrupt during the time-slot and if it is probable that a potential energy-saving that can be achieved by saving a token for later use is higher than an energy-saving that might be achieved by continuing the sleep at the time-slot.

16. The system of Claim 15, wherein the dynamic sleep controller is further configured to exclude preemptive wake-up requests generated by the dynamic sleep controller while creating the histogram.

17. The system of Claim 15, wherein the dynamic sleep controller is further configured to determine a sleep ratio of total sleep time in the low-power mode versus total idle time of the processor during a period of time; and to reset the histogram when the sleep ratio is lower than a sleep ratio threshold.

18. The system of Claim 15, wherein the dynamic sleep controller is further configured to determine a wakeup ratio of how many times a wakeup interrupt is received prior to preemptively waking the processor during a period of time; and to reset the histogram when the wakeup ratio is greater than a wakeup ratio.

19. The system of Claim 11, wherein a plurality of processors are coupled to the dynamic sleep controller for dynamic power control.

20. The system of Claim 11, further comprising a real-time subsystem coupled to the processor, wherein the task being executed by processor initiates a real-time control signal for the real-time subsystem.





FIG. 1A






FIG. 2

3/7




FIG. 3A




FIG. 3B




FIG. 3C

4/7

TOKEN PERIOD			
Time-slot0	Time-slot1	Time-slot2	Time-slotS-1
Min_token[0]	Min_token[1]	Min_token[2]	...

FIG. 4

NUMBER OF TOKENS REMAINING	TOKEN PERIOD			
	Time-slot0	Time-slot1	Time-slot2	Time-slotS-1
1	Wakeup_time_slot[0][0]	Wakeup_time_slot[0][1]	...	Wakeup_time_slot[0][S-1]
2	Wakeup_time_slot[1][0]	Wakeup_time_slot[1][1]	...	Wakeup_time_slot[1][S-1]
3	Wakeup_time_slot[2][0]	Wakeup_time_slot[2][1]	...	Wakeup_time_slot[2][S-1]
...
N	Wakeup_time_slot[N-1][0]	Wakeup_time_slot[N-1][1]	...	Wakeup_time_slot[N-1][S-1]

FIG. 5




FIG. 6

6/7




FIG. 7




FIG. 8




FIG. 9

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; G06F 1/08; H04M 1/00; H04W 52/02; G06F 1/26; G08C 17/00

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 modelsElectronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & Keywords: power management, real time, interrupt, sensor, dynamic, idle indicator, token.**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2010-0332877 A1 (MARK A. YARCH et al.) 30 December 2010 See paragraphs 23-41; and figures 2-3.	1-20
A	US 2011-0194471 A1 (JI-EUN KIM et al.) 11 August 2011 See paragraphs 56-68, 81-95; and figures 4-5, 8.	1-20
A	US 7792066 B2 (KAZUO FUJII et al.) 07 September 2010 See column 8, line 4 - column 10, line 37; and figures 2-5.	1-20
A	EP 1008030 B1 (INTEL CORPORATION) 25 April 2007 See paragraphs 11-22; and figures 2-4, 7.	1-20
A	KR 10-1997-0066802 A (TRIGEM COMPUTER INC.) 13 October 1997 See claims 1-3; and figures 1-3.	1-20

 Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:	"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
"A" document defining the general state of the art which is not considered to be of particular relevance	"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
"E" earlier application or patent but published on or after the international filing date	"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
"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)	"&" document member of the same patent family
"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	

Date of the actual completion of the international search 30 May 2013 (30.05.2013)

Date of mailing of the international search report 02 June 2013 (02.06.2013)
--

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 KIM, Byoung Sung Telephone No. 82-42-481-8403
--

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2013/024474

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2010-0332877 A1	30.12.2010	CN 101937265 A DE 102010025307 A1 JP 2011-014135 A KR 10-2011-0001984 A TW 201111975 A	05.01.2011 03.03.2011 20.01.2011 06.01.2011 01.04.2011
US 2011-0194471 A1	11.08.2011	CN 102017527 A KR 10-0932920 B1 KR 10-2009-0065152 A WO 2009-078575 A1	13.04.2011 21.12.2009 22.06.2009 25.06.2009
US 7792066 B2	07.09.2010	CN 1601987 A JP 04095501 B2 JP 2005-018377 A US 2005-0047356 A1	30.03.2005 04.06.2008 20.01.2005 03.03.2005
EP 1008030 B1	25.04.2007	AU 5450098 A EP 1008030 A1 EP 1008030 A4 JP 04191254 B2 JP 2001-506788 A KR 10-0370641 B1 TW 388011 A US 6085325 A WO 98-27482 A1	15.07.1998 14.06.2000 02.01.2002 03.12.2008 22.05.2001 05.02.2003 21.04.2000 04.07.2000 25.06.1998
KR 10-1997-0066802 A	13.10.1997	None	