



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

(21) International Application Number:
PCT/US2018/035251

(22) International Filing Date:
31 May 2018 (31.05.2018)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
15/640,537 01 July 2017 (01.07.2017) US

(71) Applicant: MICROSOFT TECHNOLOGY LICENSING, LLC [US/US]; One Microsoft Way, Redmond, WA 98052-6399 (US).

(72) Inventors: HALVERSON, Alan, Dale; Microsoft Technology Licensing, LLC, One Microsoft Way, Redmond, WA 98052-6399 (US). SEN, Rathijit; Microsoft Technology Licensing, LLC, One Microsoft Way, Redmond, WA 98052-6399 (US).

(74) Agent: MINHAS, Sandip, S. et al.; Microsoft Technology Licensing, LLC, One Microsoft Way, Redmond, WA 98052-6399 (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, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, 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, SA, 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, ST, 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).

(54) Title: CORE FREQUENCY MANAGEMENT USING EFFECTIVE UTILIZATION FOR POWER-EFFICIENT PERFORMANCE

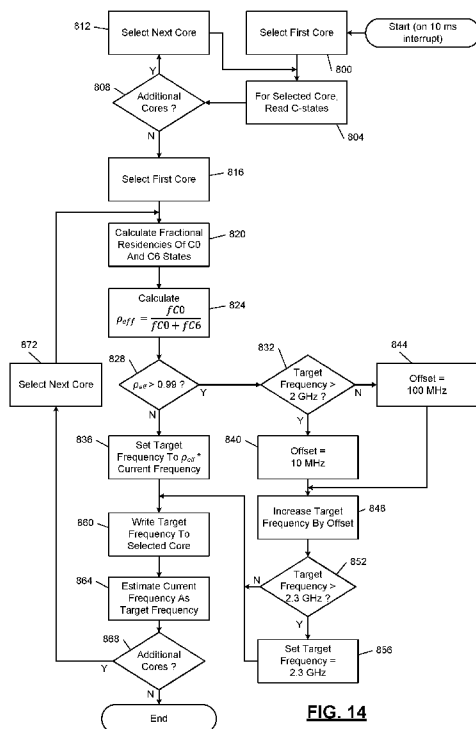


FIG. 14

(57) Abstract: A frequency governing method for a processor includes reading power management information from the processor. The processor operates in an active state or one of multiple power saving states. The power management information includes first information indicating occupancy of a first core of the processor in the active state and second information indicating occupancy of the first core in a first power saving state. The method includes generating an effective utilization based on the first information and the second information. The method includes setting a target frequency by performing a selected action of a first action including increasing the target frequency and a second action including decreasing the target frequency with respect to a current frequency of the processor based on the effective utilization. The selected action is selected according to the effective utilization. The method includes instructing the processor to operate the first core at the target frequency.

WO 2019/009973 A1

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))*

CORE FREQUENCY MANAGEMENT USING EFFECTIVE UTILIZATION FOR POWER-EFFICIENT PERFORMANCE

FIELD

5 [0001] The present disclosure relates to processor frequency control, and more particularly to frequency control of processors having multiple power saving states.

BACKGROUND

[0002] Modern processors may be able to operate in different power saving states. In some implementations, these power saving states are defined by an Advanced
10 Configuration and Power Interface (ACPI) specification, such as Section 8 of the ACPI specification version 6.1. Taking a simplistic example of a one core processor, the processor may execute instructions while in an active state, which may be labeled C0.

[0003] If the processor is waiting on another system component, which will be generally slower than the processor (for example, main memory or secondary storage), the processor
15 may enter a power saving state. For example, when a halt instruction, rather than a substantive instruction, is being executed by the processor, the processor may enter a first power saving state, labeled C1. The clock of the execution unit may be gated so that in the power saving state, the clock is not passed through to the execution units and the execution units are no longer consuming power. The processor can resume operation by
20 re-allowing the clock to be passed through.

[0004] If the processor will be idle for a longer period of time, the processor may enter another, deeper power saving state. For example, this power saving state may be labeled C3. In this second power saving state, the level 1 cache, and in some instances, the level 2
25 cache, of the processor is flushed so that power to the caches can be removed. The level 1 and 2 caches are generally volatile, meaning that removing power will cause the contents to be erased. To refill the cache takes more time, so the processor requires additional time to return to the active state from the second power saving state. However, when the processor will be idle for a longer period of time, such as when requesting data from main memory, entering the second power saving state may be worthwhile.

30 [0005] Still further power saving states are available in various processor architectures. For example, a third power saving state may save off state data from the processor's volatile storage and then remove power to most of the functional units of the processor. For example, caches, pipeline registers, architectural registers, clock distribution circuitry, branch predictors, arithmetic units, etc. may be powered down. The clock generation

circuitry may remain powered to avoid the additional latency of generation of the clock resuming and stabilizing. In order to resume execution, the saved state is reloaded into the processor.

5 [0006] The third power saving state may be worthwhile when the processor is waiting for a much higher-latency task, such as a disk access or when the processor is idle/inactive such as while waiting for the next transaction to arrive. The tradeoffs between the latency and power savings of power saving states may be determined by the designers of the processors and may be systematized in a power management unit of the processor.

10 [0007] In the active state (C0), there may be multiple performance levels. For example, a highest performance state (labeled P0) may be the highest frequency at which the processor can operate. If the processor is lightly loaded, the processor may operate at a lower performance state, such as a state labeled P1. Different processor architectures, and even different processor models sharing an architecture, may have different sets of performance states. For example, defined performance states from P0 to Pn, where n is an integer greater than one, may be defined by the processor. In some processors, a range of
15 frequencies are available and may not necessarily map to specific P numbers. Instead, the frequency of the processor may be adjustable anywhere within the range subject to a set increment, or granularity. For example, the frequency of the processor may be adjusted in increments of 100 MHz.

20 [0008] In more modern processors, the processor includes more than one core. Further, each core may, depending on the processor model, be able to execute multiple threads. Except in the very rare instances in which all of the threads executing across all the cores of the processor consistently present a uniform load, certain cores of the processor may be more lightly loaded at certain times than others. Therefore, the processor power
25 management may adjust the power saving state of each core separately.

[0009] Further, some recent processors allow the performance state of the processor to be adjusted per core. In other words, in addition to being able to enter individual cores into separate power saving states, the frequency of each core may be adjusted independently in the active state. In other implementations, groups of cores may be controlled together. For
30 example, in an eight-core processor, each pair of cores may be set to a frequency that is independent of the other pairs of cores.

[0010] Simply as a graphical illustration and not representative of any specific processor, FIG. 1 depicts a processor with four power states: three power saving states and an active (C0) power state 10. The C0 power state includes five performance states: performance

state 0 (20-0), performance state 1 (20-1), performance state 2 (20-2), and performance state 3 (20-3). Simply for illustration, the performance states 20-0, 20-1, 20-2, and 20-3 correspond to frequencies of 2.3 GHz, 2.1 GHz, 1.8 GHz, and 1.2 GHz, respectively.

[0011] Further, in some multi-core processors, when fewer than all cores are in the C0 state 10, one or more higher performance states (sometimes referred to as turbo states) are available. The ability to run a subset of cores at a higher frequency is generally due to thermal management. The processor and any associated heat sink may only be able to dissipate the heat generated from a subset of cores operating at a higher frequency. In FIG. 1, performance state 24 may allow a single core to operate at 3.0 GHz when the other cores are in a power saving state.

[0012] In FIG. 1, inactive power saving states include an idle state 32-1 (C1, also called a first power saving state), a second power saving state 32-2 (C3), and a third power saving state 32-3 (C6). Although performance cannot be negative, the depth of the graphical bars for the power saving states are a graphical reminder of the latency involved in returning from a power saving state.

[0013] In other words, the amount of time required to return to the active state 10 from the first power saving state 32-1 is much less than the latency to return from the second power saving state 32-2, which, in turn, is less than the time required to resume execution from the inactive state 32-3. For example only, the second power saving state 32-2 may involve flushing the L1 and L2 caches, while the power saving state 32-3 may include saving the state of the core and removing power from most functional units. To be clear, the latency depicted for the power saving states 32 is a different, not comparable, scale to the frequency depicted for the performance states 20.

[0014] The labels C1, C3, and C6 were shown in FIG. 1 simply to illustrate the fact that certain processors implement power saving states that are a subset of the power saving states defined by the manufacturer or by the ACPI specification. In addition, a processor manufacturer may define additional power saving states that are not industry standard (though they may be defined by the processor manufacturer according to the format of the ACPI specification).

[0015] In FIG. 2, a more granular array 40 of performance states may be available. For example, the frequency of the processor may be adjusted arbitrarily by predefined increments. In some implementations, higher performance states 44 of the array 40 require that at least some of the cores be in one of the power saving states 32.

[0016] Operating systems include frequency governors designed with various tradeoffs

between power and performance. These frequency governors may instruct the processor at which frequency to operate each processor core. For example, Linux distribution Ubuntu, from Canonical Ltd., includes several frequency governors grouped into two classes. See the following table:

CLASS	GOVERNOR	TYPE	SHORT DESCRIPTION
Cpufreq	Performance	static	All cores at maximum (turbo) frequency.
	Powersave	static	All cores at minimum frequency.
	Ondemand	dynamic	Governor changes core frequencies depending on core utilization.
	Conservative	dynamic	Similar to Ondemand governor, but changes frequencies more gradually.
	Userspace	-	Frequencies can be reconfigured by the administrator.
Intel P-State	Performance	static	Similar to Cpufreq Performance governor.
	Powersave	dynamic	Similar in strategy to Cpufreq Ondemand governor, but with a different implementation.

- 5 [0017] An operating system user with root privileges may be able to switch governors within a class while the operating system is running. However, changing the class may require modifying a kernel perimeter and rebooting the system. These prior art frequency governors often fail to save power at high loads or suffer significant performance loss. Failing to save power increases the energy footprint of the device, the negative effects of which are magnified at cloud scale. On the other hand, failing to maximize performance may require additional systems or more stringent hardware specifications in order to handle a given load. For a fixed hardware environment, failing to maximize performance may lead to worse application-level performance for customers, such as slower response times or lower transaction throughput.
- 10
- 15 [0018] The background description provided here is for the purpose of generally presenting the context of the disclosure. Work of the presently named inventors, to the

extent it is described in this background section, as well as aspects of the description that may not otherwise qualify as prior art at the time of filing, are neither expressly nor impliedly admitted as prior art against the present disclosure.

SUMMARY

5 [0019] A frequency governing method for a processor includes reading power management information from the processor. The processor is configured to selectively operate in a plurality of power saving states. The power management information includes (i) first information indicating occupancy of a first core of the processor in an active state and (ii) second information indicating occupancy of the first core of the processor in a first
10 power saving state of the plurality of power saving states. The method includes generating an effective utilization based on the first information and the second information. The method includes setting a target frequency by performing a selected action of a first action including increasing the target frequency and a second action including decreasing the target frequency with respect to a current frequency of the processor based on the effective
15 utilization. The selected action is selected according to the effective utilization. The method includes instructing the processor to operate the first core at the target frequency.

[0020] In other features, the first action is selected if the effective utilization is greater than a threshold, and the second action is selected if the effective utilization is less than the threshold. In other features, decreasing the target frequency includes setting the target
20 frequency to a product of the effective utilization and the current frequency of the processor.

[0021] In other features, the effective utilization is generated based on a first number and a second number. The first number is based on the first information. The second number is based on the second information. The first number indicates a fraction of time the first
25 core occupied the active state since a prior analysis. The second number indicates a fraction of time the first core occupied a subset of the plurality of power saving states since the prior analysis. The subset of the plurality of power saving states is a proper subset and includes at least the first power saving state. In other features, the effective utilization is generated based on a division of the first number by a sum of the first number
30 and the second number.

[0022] In other features, the processor includes multiple cores including the first core. The method includes instructing the processor to operate all of the multiple cores at the target frequency. The method includes one of (i) determining an individual effective utilization for each core of the multiple cores and averaging the individual effective

utilizations to generate the effective utilization and (ii) averaging the first information across the multiple cores, averaging the second information across the multiple cores, and generating the effective utilization based on the averaged first information and the averaged second information.

5 [0023] In other features, the processor includes multiple cores including the first core. The method includes performing the reading, the generating, the setting, and the instructing separately for each core of the multiple cores. In other features, from among the plurality of power saving states, a latency of the first core to return to the active state is greatest for the first power saving state.

10 [0024] A frequency governor system for a processor having a first core includes a power state data collection subsystem configured to read power management information from the processor. The processor is configured to selectively operate in a plurality of power saving states. The power management information includes (i) first information indicating occupancy of the first core of the processor in an active state and (ii) second information
15 indicating occupancy of the first core of the processor in a first power saving state of the plurality of power saving states. The system includes a utilization analyzer configured to generate an effective utilization based on the first information and the second information. The system includes a frequency control circuit configured to set a target frequency by performing a selected action of a first action including increasing the target frequency and
20 a second action including decreasing the target frequency with respect to a current frequency of the processor based on the effective utilization. The selected action is selected according to the effective utilization. The frequency control circuit is configured to instruct the processor to operate the first core at the target frequency.

[0025] In other features, the frequency control circuit is configured to select the first
25 action if the effective utilization is greater than a threshold and select the second action if the effective utilization is less than the threshold. In other features, the frequency control circuit is configured to decrease the target frequency by setting the target frequency to a product of the effective utilization and the current frequency of the processor.

[0026] In other features, the utilization analyzer is configured to generate the effective
30 utilization based on a first number and a second number. The first number is based on the first information. The second number is based on the second information. The first number indicates a fraction of time the first core occupied the active state since a prior analysis by the utilization analyzer. The second number indicates a fraction of time the first core occupied a subset of the plurality of power saving states since the prior analysis. The

subset of the plurality of power saving states is a proper subset and includes at least the first power saving state. In other features, the utilization analyzer is configured to generate the effective utilization based on a division of the first number by a sum of the first number and the second number.

5 [0027] In other features, the processor includes multiple cores including the first core. The frequency control circuit is configured to instruct the processor to operate all of the multiple cores at the target frequency. The utilization analyzer is configured to generate the effective utilization by one of (i) determining an individual effective utilization for each core of the multiple cores and averaging the individual effective utilizations to
10 generate the effective utilization and (ii) averaging the first information across the multiple cores, averaging the second information across the multiple cores, and generating the effective utilization based on the averaged first information and the averaged second information.

[0028] A frequency governor system for a processor having a first core includes power
15 state data collection means for reading power management information from the processor. The processor is configured to selectively operate in a plurality of power saving states. The power management information includes (i) first information indicating occupancy of the first core of the processor in an active state and (ii) second information indicating occupancy of the first core of the processor in a first power saving state of the
20 plurality of power saving states. The system includes utilization analysis means for generating an effective utilization based on the first information and the second information. The system includes frequency control means for setting a target frequency by performing a selected action of a first action including increasing the target frequency and a second action including decreasing the target frequency with respect to a current
25 frequency of the processor based on the effective utilization. The selected action is selected according to the effective utilization. The frequency control means instructs the processor to operate the first core at the target frequency.

[0029] In other features, the frequency control means selects the first action if the effective utilization is greater than a threshold and selects the second action if the effective
30 utilization is less than the threshold. In other features, the frequency control means decreases the target frequency by setting the target frequency to a product of the effective utilization and the current frequency of the processor.

[0030] In other features, the utilization analysis means generates the effective utilization based on a first number and a second number. The first number is based on the first

information. The second number is based on the second information. The first number indicates a fraction of time the first core occupied the active state since a prior analysis by the utilization analysis means. The second number indicates a fraction of time the first core occupied a subset of the plurality of power saving states since the prior analysis. The subset of the plurality of power saving states is a proper subset and includes at least the first power saving state. In other features, the utilization analysis means generates the effective utilization based on a division of the first number by a sum of the first number and the second number.

[0031] In other features, the processor includes multiple cores including the first core. The frequency control means instructs the processor to operate all of the multiple cores at the target frequency. The utilization analysis means generates the effective utilization by one of (i) determining an individual effective utilization for each core of the multiple cores and averaging the individual effective utilizations to generate the effective utilization and (ii) averaging the first information across the multiple cores, averaging the second information across the multiple cores, and generating the effective utilization based on the averaged first information and the averaged second information.

[0032] Further areas of applicability of the present disclosure will become apparent from the detailed description, the claims, and the drawings. The detailed description and specific examples are intended for purposes of illustration only and are not intended to limit the scope of the disclosure.

BRIEF DESCRIPTION OF THE DRAWINGS

[0033] The present disclosure will become more fully understood from the detailed description and the accompanying drawings.

[0034] FIG. 1 is a graphical illustration of processor performance in various power states.

[0035] FIG. 2 is a graphical representation of processor performance in various power states for processor models having finer granularity of frequency control.

[0036] FIG. 3 is a functional block diagram of a distributed computing system that includes client devices and servers.

[0037] FIG. 4 is a functional block diagram of an example client device that may implement the frequency control of the present disclosure.

[0038] FIG. 5 is a functional block diagram of an example server implementing frequency control according to the present disclosure.

[0039] FIG. 6 is a high-level functional block diagram of a computing system in which

the operating system controls multiple cores of processor hardware.

[0040] FIG. 7 is a functional block diagram of operating system control of an example implementation of processor hardware.

5 [0041] FIG. 8 is a functional block diagram depicting user software control of processor power management.

[0042] FIG. 9 is a functional block diagram of operating system control of processor power management.

[0043] FIG. 10 is a functional block diagram of an example implementation of processor hardware incorporating frequency control of the present disclosure.

10 [0044] FIG. 11 is a functional block diagram of an example implementation of a frequency governor system according to the principles of the present disclosure.

[0045] FIG. 12 is a flowchart of an example frequency control according to the principles of the present disclosure.

15 [0046] FIG. 13A and FIG. 13B are flowcharts of frequency control according to the principles of the present disclosure in which all processor cores in a set are operated at the same frequency.

[0047] FIG. 14 is a flowchart of a more specific example of the operation of FIG. 12 in which specific values for certain variables are presented as examples.

20 [0048] FIG. 15A is plot of system power versus load for three different frequency governing systems.

[0049] FIG. 15B is a plot of system power versus load for five frequency governing systems.

[0050] In the drawings, reference numbers may be reused to identify similar and/or identical elements.

25

DETAILED DESCRIPTION

INTRODUCTION

30 [0051] Frequency governors attempt to minimize power consumption while still handling the load presented to the processor. Some frequency governors attempt to analyze load by looking at application layer metrics, such as query response time or transactions per second. However, these metrics may not be fully correlated with the activity level of the processor at any given point in time. Further, in environments where virtual machines or containers are being hosted by the operating system, the operating system and the underlying hardware may not have visibility into these application layer metrics (see FIG. 6).

[0052] One performance-aware power saving technique that does not rely on application layer metrics is to set processor frequencies so that each core is fully utilized. However, identifying full utilization is difficult. If the active (C0) state is used as the only measurement of active time, a frequency governor may underestimate the load. This is
5 because the processor may transition out of the C0 state into a power saving state (such as the C1 inactive state) due to delays in other systems, not inactivity.

[0053] For example, the C1 state may be entered due to slower storage or network delays or other scheduling delays. It may also be caused by C-state demotion, where a deeper sleep state is not entered because it is speculated that wakeup may happen shortly. If
10 power saving state transitions are performed quickly enough, and/or if main memory latency is slow enough relative to the processor clock, fetching data from main memory may even cause a transition to C1.

[0054] The time that the processor spends in the inactive state is not necessarily reflective of a decreased load; instead, a lower percentage of time being spent in the C0
15 state may be an indication that the processor is waiting on other subsystems, such as I/O (input/output). In other words, observing that the C0 state is occupied only 82% of the time may not necessarily mean that the frequency should be decreased so that the percentage of C0 time approaches 100%. By considering only C0 occupancy, an aggressive power saving governor may lose half of the possible throughput by erroneously
20 assuming that lower C0 occupancy means lower load.

[0055] The present disclosure describes that certain power saving states (such as C1 and C3) are indicative of delays imposed on the processor, while other inactive states (such as C6) represent actual underutilization of the processor. Therefore, instead of controlling to maximize C0 occupancy, a frequency governor may instead control to minimize the
25 occupancy in the power saving states that indicate underutilization. In other words, for systems where state C6 is assessed to indicate underutilization, the frequency of a core may be decreased to drive the percentage of time spent in state C6 close to 0%. Alternatively, this can be phrased as decreasing the frequency to drive the percentage of time spent in other states (C0, C1 and C3) close to 100%.

[0056] The present disclosure describes a frequency governor that, in certain operating
30 regimes, decreases the amount of power used compared to traditional frequency governors and, in other operating regimes, increases performance with respect to other frequency governors. This new frequency governor improves on governors that only consider core utilization (the percentage of time spent in the C0 state). Such governors either do not

aggressively save power at high loads or suffer large throughput losses if C1 or C3 states are significantly occupied by the processor cores.

EXAMPLE IMPLEMENTATION

[0057] In the following, the notation fC_n denotes the fraction of time (expressed as a real number between 0 and 1) spent in power state C_n , where n is an integer. In some implementations, n can be any string, which may include an integer or a combination of integer and letter (such as “1e”).

[0058] The fraction of time that a core is active, fC_0 , scales inversely with the core’s frequency. This scaling is generally not linear because some contributors to execution time, such as memory hierarchy stall times, are not affected by core frequency. The value fC_6 may indicate underutilization because power state C_6 is only entered when the processor expects to be idle for a long period of time. Decreasing fC_6 by decreasing processor frequency increases fC_0 . The sum of fC_n for n between 1 and 5, inclusive, may be generally unaffected by core frequency. This is because these states are generally caused directly or indirectly by I/O (including storage and networking) or other scheduling events whose latency does not directly depend on processor core frequency.

[0059] With C_{freq} denoting the current frequency for a core, the governor selects a target frequency, T_{freq} , such that:

$$\left(\frac{C_{freq}}{T_{freq}}\right) \times fC_0 + \sum_{n=1}^5 fC_n \leq 1$$

[0060] Recognizing that the sum of fC_n for all n is equal to 1, the sum of fC_n for n from 1 to 5 can be rewritten as 1 minus the sum of fC_n for n outside of the range 1–5. In the current example, states C_0 and C_6 are the only power saving states outside of the range 1–5. The above equation can therefore be rewritten as follows:

$$\left(\frac{C_{freq}}{T_{freq}}\right) \times fC_0 + (1 - fC_0 - fC_6) \leq 1$$

Solving for T_{freq} in this inequality leads to:

$$T_{freq} \geq \left(\frac{fC_0}{fC_0 + fC_6}\right) \times C_{freq}$$

[0061] The ratio of fC_0 to the sum of fC_0 and fC_6 is referred to as the effective utilization, ρ_{eff} . Mathematically, ρ_{eff} must be less than or equal to 1 and if the target frequency is based on multiplying the current frequency by a number less than or equal to 1, this inequality indicates only how to lower operating frequency.

[0062] To correct for overutilization, the frequency governor increases T_{freq} when ρ_{eff}

exceeds a threshold. For example, the governor may increase T_{freq} by a predetermined increment whenever ρ_{eff} exceeds 0.99. The predetermined increment may vary based on the current frequency of the processor core. For example, below a threshold (such as 2 GHz), a 100 MHz increment may be used, while above that threshold, a 10 MHz
5 increment may be used. The smaller value of the frequency increment in the higher frequencies is motivated by the expensive power consumption of higher frequencies. For example, the threshold may correspond to the area of turbo operation (that is, frequencies at which not all processor cores can operate simultaneously). Instead of two increments, an expression may be used that decreases the increment as the current frequency increases.
10 The expression may include a floor specifying the smallest increment as well as a ceiling specifying the largest increment.

[0063] These particular parameter choices may slightly favor power savings over performance losses, and therefore, particular implementations may adjust these values based on the priorities of the system design. The frequency governor may schedule an
15 iteration of T_{freq} calculation after every 10 milliseconds. However, with processing time and scheduling delays, the frequency governor may be invoked fewer than the expected 100 times a second ($1 \text{ s} / 10 \text{ milliseconds}$). For example, the governor may actually be invoked approximately 70 times per second.

[0064] The frequency governor may be implemented in user software (as shown in FIG. 8), as part of the operating system (FIG. 9) either in the kernel or a driver, as part of the BIOS (basic input/output system), or by the processor itself (FIG. 10). The effective
20 utilization based on the power states may be measured per core and the frequency controlled per core for processors that allow different frequencies for different cores. In processors or systems where frequencies of multiple cores are controlled together, the
25 effective utilization may be averaged across the cores. For example, the effective utilization may be calculated for a single core, for a subset of cores of one processor, for all of the cores of one processor, or for cores of multiple processors in a system.

[0065] The frequency governor of the present disclosure may be implemented in servers (such as is shown in FIG. 5) whether those servers are standalone or in a distributed
30 computing or cloud environment. The frequency governor may also be used in client devices (see FIG. 4), such as desktop and portable computers. Further, personal devices, such as smartphones and wearable devices, may implement the frequency governor of the present disclosure. For cloud systems, the decrease in power may save on energy costs and greenhouse gas emissions as well as operating costs for a data center. For portable devices,

the frequency governor may extend battery life.

EXPERIMENT

- 5 [0066] To test the frequency governor on a server task, an online transaction processing (OLTP) workload was used for testing purposes. In one test, a ThinkStation P710 desktop computer from Lenovo Group Ltd. includes dual-socket Xeon Broadwell processors (E5-2620 v4) from Intel Corp. These processors have a total of 32 logical cores with simultaneous multi-threading (also known as hyper-threading) enabled. Further, the computer includes 64GB of double data rate (DDR4) memory, two 512GB solid state disks (SSDs) configured as a single logical volume and two 1TB magnetic disks.
- 10 [0067] A database server, such as SQL server from Microsoft Corp., is configured on a distribution of Linux such as Ubuntu 16.04.1 LTS from Canonical Ltd. The database server is configured to store data and logs on the logical volume formed by the SSDs. These processors support dynamic voltage and frequency scaling in which core frequency scaling is accompanied by a change in voltage, which results in a cubic or nearly cubic
- 15 change in dynamic power or at least a linear change in static power.
- [0068] These processors support per-core frequency control, which may be actuated by writing frequency values to model specific registers (MSRs). The maximum frequency when all cores are active is 2.3 GHz. The entire power used by the system can be measured by a power meter placed in line with the mains AC power supply.
- 20 [0069] The test workload includes a database schema of six tables and a mix of nine types of transactions that involve reads, inserts, updates, and deletes of records. The transactions are initiated by many connection threads running on a separate client networked with the test machine. A “pacing delay” is defined as the average time in seconds between completion of one transaction and the start of the next transaction for
- 25 each thread. The pacing delay can be varied over time to observe the response of the frequency governor to a change in load. Specifically, in this test the time between the completion of one transaction and the start of the next transaction for a thread is an exponentially-distributed random variable with a mean set to the pacing delay and a maximum capped to 10 times the pacing delay.
- 30 [0070] In the plots of FIGs. 15A and 15B, the pacing delay is decreased from 1 through 0.5, 0.2, 0.1, 0.05, 0.04, all the way through 0 (in which there is no delay between one transaction completing and the request for the next to begin). For each value of pacing delay, the benchmark is run for 15 minutes of steady state following a startup period, such as two minutes. Because the way in which database workloads interact with the hardware

may evolve over time, the database was restored to a known state before each frequency governor was tested.

[0071] Observations include that $fC6$ decreases as load increases due to decreasing opportunity for cores to reside in the deep sleep state. $fC0$ (active state) increases with load but only reaches approximately 82% at full load, with nearly all of the remainder of time being spent in C1. For a fixed load, $fC6$ increases and $fC0$ decreases as core frequency goes up since the core finishes processing work faster.

[0072] FIGs. 15A and 15B compare the frequency governor of the present disclosure (labeled as “New”) compared to existing frequency governors. The load, expressed as transactions per second (TPS), is increased by decreasing the pacing delay. In FIG. 15A, the new frequency governor is compared to the P-state Performance and P-state Powersave frequency governors. In FIG. 15B, the new frequency governor is compared to the Cpufreq Performance, Powersave, On-demand, and Conservative frequency governors. Transactions per second and power consumed both increase with load (decreasing values of pacing delay).

[0073] Generally, a resource governor should scale power with load served as well as aim to select power-performance Pareto-optimal configurations. In other words, the frequency governor should be configured to consume the least amount of power for the loads served. Compared to the P-state Performance and Cpufreq Performance governors, the Cpufreq Powersave governor saves the most power (approximately 34 watts at $D=0.1$), but can serve only low to intermediate loads. The Cpufreq on-demand governor saves less power but can serve a wider range of loads. However, it significantly loses throughput compared to the Cpufreq performance governor at the full $D=0$ load.

[0074] The remaining governors are able to achieve the full throughput, though their power efficiency varies greatly over the range of loads. As seen, the new frequency governor consumes no more power, and in some cases, consumes much less power, at each level of load compared to the other frequency governors that can serve the full range of loads. The new frequency governor can be adjusted to achieve different tradeoffs between performance and power. For example, by increasing the frequency increment for lower frequencies, the throughput of the new frequency governor can be increased at the expense of slightly higher power consumption.

DEVICE IMPLEMENTATIONS

[0075] Below are simplistic examples of a distributed computing environment in which the systems and methods of the present disclosure can be implemented. Throughout the

description, references to terms such as servers, client devices, applications and so on are for illustrative purposes only. The terms server and client device are to be understood broadly as representing computing devices with one or more processors and memory configured to execute machine readable instructions. The terms application and computer
5 program are to be understood broadly as representing machine readable instructions executable by the computing devices.

[0076] FIG. 3 shows a simplified example of a distributed computing system 100. The distributed computing system 100 includes a distributed communications system 110, one or more client devices 120-1, 120-2, ..., and 120-M (collectively, client devices 120), and
10 one or more servers 130-1, 130-2, ..., and 130-M (collectively, servers 130). N and M are integers greater than or equal to one. The distributed communications system 110 may include a local area network (LAN), a wide area network (WAN) such as the Internet, or other type of network. For example, the servers 130 may be located at different geographical locations. The client devices 120 and the servers 130 communicate with each
15 other via the distributed communications system 110. The client devices 120 and the servers 130 connect to the distributed communications system 110 using wireless and/or wired connections.

[0077] The client devices 120 may include smartphones, personal digital assistants (PDAs), laptop computers, personal computers (PCs), etc. The servers 130 may provide
20 multiple services to the client devices 120. For example, the servers 130 may execute software applications developed by one or more vendors. The server 130 may host multiple databases that are relied on by the software applications in providing services to users of the client devices 120.

[0078] FIG. 4 shows a simplified example of the client device 120-1. The client device
25 120-1 may typically include a central processing unit (CPU) or processor 150, one or more input devices 152 (e.g., a keypad, touchpad, mouse, touchscreen, etc.), a display subsystem 154 including a display 156, a network interface 158, memory 160, and bulk storage 162.

[0079] The network interface 158 connects the client device 120-1 to the distributed
30 computing system 100 via the distributed communications system 110. For example, the network interface 158 may include a wired interface (for example, an Ethernet interface) and/or a wireless interface (for example, a Wi-Fi, Bluetooth, near field communication (NFC), or other wireless interface). The memory 160 may include volatile or nonvolatile memory, cache, or other type of memory. The bulk storage 162 may include flash

memory, a magnetic hard disk drive (HDD), and other bulk storage devices.

[0080] The processor 150 of the client device 120-1 executes an operating system (OS) 164 and one or more client applications 166. The client applications 166 include an application that accesses the servers 130 via the distributed communications system 110.

5 [0081] FIG. 5 shows a simplified example of the server 130-1. The server 130-1 typically includes one or more CPUs or processors 170, a network interface 178, memory 180, and bulk storage 182. In some implementations, the server 130-1 may be a general-purpose server and include one or more input devices 172 (e.g., a keypad, touchpad, mouse, and so on) and a display subsystem 174 including a display 176.

10 [0082] The network interface 178 connects the server 130-1 to the distributed communications system 110. For example, the network interface 178 may include a wired interface (e.g., an Ethernet interface) and/or a wireless interface (e.g., a Wi-Fi, Bluetooth, near field communication (NFC), or other wireless interface). The memory 180 may include volatile or nonvolatile memory, cache, or other type of memory. The bulk storage
15 182 may include flash memory, one or more magnetic hard disk drives (HDDs), or other bulk storage devices.

[0083] The processor 170 of the server 130-1 executes an operating system (OS) 184 and one or more server applications 186, which may be housed in a virtual machine hypervisor or containerized architecture. The bulk storage 182 may store one or more databases 188
20 that store data structures used by the server applications 186 to perform respective functions.

BLOCK DIAGRAMS

[0084] In FIG. 6, simplified processor hardware 200 includes four cores 204-1, 204-2, 204-3, and 204-4 (collectively, cores 204). A hypervisor layer 208 for an operating system
25 kernel and hypervisor includes a power management subsystem 212 that manages the power states of the cores 204. For example, the power management subsystem 212 may control the operating frequencies of the cores 204.

[0085] Virtual machines 220-1, 220-2, ... 220-n (virtual machines 220) execute on the operating system kernel and the hypervisor layer 208. User software 224-1, 224-2, ... 224-
30 n (user software 224) is shown schematically. The user software 224 may include end user software, such as office productivity applications and web browsers, and may include server applications, such as web servers and database servers. The operating system kernel and the hypervisor layer 208 may not have insight into the application level data within the virtual machines 220. As a result, the power management subsystem 212 may have to rely

solely on the data provided by the cores 204 regarding their relative utilizations.

[0086] In FIG. 7, simplified processor hardware 300 includes cores 304-1 and 304-2. A voltage and frequency regulator 308-1 controls the operating voltage and frequency of the core 304-1 while a voltage and frequency regulator 308-2 controls the operating voltage and frequency of the core 304-2. For example, each of the voltage and frequency
5 regulators 308 may include a voltage regulator as well as a frequency divider that vary in tandem according to instructions from a power management unit 312.

[0087] The power management unit 312 may communicate with the cores 304 to determine when the cores 304 should enter a power saving state and controls the
10 frequency of the cores by sending appropriate commands to the voltage and frequency regulators 308. The power management unit 312 may operate according to settings made in control registers 316. In various implementations, the power management unit 312 may write certain data to the control registers 316, such as current operating frequencies of the cores 304.

[0088] Performance management counters 320 measure performance metrics of the cores 304. For example, the performance management counters 320 may measure the residency of each of the cores 304 within each of the power states. Based on data from the performance management counters, a power management subsystem 324 of an operating system kernel 328 may update values in the control registers 316 to instruct the power
15 management unit 312 to adjust the frequency of the cores 304. As an example, the performance management counters 320 may measure the residency of the cores 304 in an active state (C0) as well as in a deep sleep (C6) state. The power management subsystem 324 periodically reads these values and adjusts the frequencies of the cores 304 to decrease the proportion of time spent in the deep sleep states.

[0089] In FIG. 8, an example of how user software may be used to manage processor frequencies is described. The processor hardware 300 of FIG. 7 is shown, while the voltage and frequency regulators 308 are omitted for clarity. An operating system kernel 350 includes a frequency governor system 354, which communicates with the performance management counters 320 and the control registers 316 through a processor
20 hardware access layer 358. For example, the processor hardware access layer 358 may be a driver and may include the basic input/output system (BIOS).

[0090] User software 362 may include privileged user software 366. For example, the privileged user software 366 may operate with administrator (or, root) credentials. The privileged user software 366 may include a frequency governor system 370 according to
30

the principles of the present disclosure. The frequency governor system 370 may access the control registers 316 and the performance management counters 320 by the processor hardware access layer 358. The frequency governor system 370 may need to disable the frequency governor system 354 and/or instruct the frequency governor system 354 that manual frequency control should be enabled.

5 [0091] For example, the Cpubfreq Userspace frequency governor mentioned above may allow the privileged user software 366 to access the control registers 316 and the performance management counters 320. In some circumstances, the Cpubfreq Userspace frequency governor imposes certain restrictions, such as larger granularity on values. As a result, the frequency governor system 370 may bypass the Userspace interface and instead directly write values to the control registers 316. This may allow for finer granularity of control of core frequencies. In the absence of the privileged user software 366, the frequency governor system 354 may control core frequencies, such as using one of the governors in the table above.

15 [0092] The frequency governor system 370 may be able to interact with the control registers 316 and the performance management counters 320 through the processor hardware access layer 358 even if the frequency governor system 354 were omitted. In other words, if the operating system kernel exposes hardware access to the privileged user software 366, the operating system kernel 350 may not need to include a frequency governor or understand the information exchanged between the frequency governor system 370 and the control registers 316 and the performance management counters 320.

20 [0093] In FIG. 9, an operating system kernel 400 may include a frequency governor system 404 according to the principles of the present disclosure, which directly accesses the control registers 316 and the performance management counters 320. The frequency governor system 404 may be implemented as part of the operating system kernel 400, as a user mode or protected mode driver, as part of the basic input/output system (BIOS) of a computer, or as a package of the operating system integrated with or executing on the operating system kernel 400.

25 [0094] In FIG. 10, processor hardware 420 may include cores 424-1 and 424-2 and voltage and frequency regulators 428-1 and 428-2, which are similar to those of FIG. 7. Further, a power management unit 432 may access performance management counters 436, which are similar to the performance management counters 320 of FIG. 7. In various implementations, the power management unit 432 may be activated or enabled by a power management subsystem 440 of an operating system kernel 444. For example, the power

management subsystem 440 may handle power management including suspend, sleep and hibernate.

[0095] The power management subsystem 440 may write to control registers 448 which are used to control the power management unit 432. The power management unit 432 may
5 include a power management coprocessor 452 and nonvolatile storage 456 that stores microcode. The microcode is executed by the power management coprocessor 452 to control the frequencies of the cores 424 according to the principles of the present disclosure. The microcode may also make decisions about when to enter each of the cores 424 into different power saving modes.

10 [0096] In various implementations, hardware implementations of the frequency governor of the present disclosure may be implemented in the processor hardware 420. However, for flexibility, the frequency governor may be represented as microcode.

[0097] In FIG. 11, a functional block diagram of an example frequency governor system 500 includes a power state data collection subsystem 504. For example, the power state
15 data collection subsystem 504 may acquire power state residency information, such as from performance management counters described above. The power state data collection subsystem 504 may clear the relevant counters upon reading them so that each read is incremental to the last. In other implementations, the power state data collection subsystem 504 may subtract the current reading from the prior reading to determine the
20 incremental change.

[0098] A statistical analysis subsystem 508 converts whatever form the power state data is collected in to a fractional representation, such as a real number between 0 and 1, for each power state for each core. Equivalently, the real number may be a percentage between 0 and 100.

25 [0099] A utilization analyzer 512 determines an effective utilization of each core. For example, this may be calculated as the ratio of fC0 to the sum of fC0 and fC6. In some implementations, the power states that indicate underutilization of the processor, such as C6, may be specified by power state desirability control 516. For example, for certain processors, a C5 state and/or a C7 state may also indicate underutilization of the processor.
30 In the case where C5 and C7 also indicate underutilization, the power state desirability control 516 would indicate that the sum of fC5, fC6, and fC7 should be minimized. The utilization analyzer 512 may then include fC5 and fC7 in the sum of fC0 and fC6.

[0100] A closed-loop frequency control circuit 520 receives a current operating frequency for the cores from an operating frequency determination circuit 524. For

example, the operating frequency determination circuit 524 may read the current operating frequencies of the cores from processor registers or performance counters. In other implementations, the operating frequency determination circuit 524 may store and recall the prior commanded target frequency with the assumption that the target frequency will
5 now be the operating frequency.

[0101] The closed-loop frequency control circuit 520 attempts to maximize the effective utilization from the utilization analyzer 512. As described in more detail below, however, if the effective utilization becomes too high, overutilization may be occurring and therefore the frequency is increased to increase performance. A frequency limiter 528 may
10 apply an upper limit to the commanded frequency from a closed-loop frequency control circuit 520. The frequency limiter 528 then sends frequency commands to the processor for each of the cores. Together, the closed-loop frequency control circuit 520 and the frequency limiter 528 may be referred to as a frequency control circuit 532.

FLOWCHARTS

[0102] In FIG. 12, control begins at 600, where a first core of the processor is selected. Control may begin periodically, such as at a scheduled 10 millisecond interval. In 604, for the selected core, control determines residency of the core in the various power states since the prior execution of the control of FIG. 12. Control also determines the current
15 frequency of the selected core.

[0103] At 608, control calculates a value labeled “A” indicating a relative residency of the core in desired power states—that is, power states that indicate utilization of the processor, such as C0. Control also calculates a value labeled “B” indicating the relative residency of the core in undesired power states that indicate underutilization of the core. For example, C6 may be an undesired power state.
20

[0104] As 612, control calculates an effective utilization based on the ratio of A to the sum of A and B. At 616, if the effective utilization exceeds the threshold, control transfers to 620; otherwise, control transfers to 624. If the effective utilization is above a threshold, that may indicate overutilization of the core—in other words, that the frequency of the core should be increased to accommodate the load. As 620, control increases the target
25 frequency by a specified offset. The offset may vary based on the target frequency.
30

[0105] Control continues at 628, where an upper limit is applied to the target frequency. Control then continues at 632. At 624, control sets the target frequency to the product of the effective utilization and the current frequency of the core. Control then continues at 632. At 632, control instructs the selected core to operate at the target frequency. Control

continues at 636 where, if there are additional cores remaining to evaluate, control transfers to 640; otherwise, control ends. At 640, control selects the next core and returns to 604.

5 [0106] In FIG. 13A and FIG. 13B, the frequencies of a set of cores are controlled together and therefore are evaluated together in each control iteration. For example, a processor may not have individualized frequency control of cores and therefore all the cores of the processor are part of a single set. In another example, each pair of cores of the processor may form a set. When there are multiple sets, the control of FIG. 13A or FIG. 13B is performed for each set.

10 [0107] In various implementations, multiple cores of a processor may be assigned to a certain virtual machine or container. These multiple cores may form a set. Then, different parameters may be used for the frequency governor based on different requirements for different virtual machines. For example, the cores dedicated to a virtual machine with more demanding service level objectives may be controlled by the new frequency
15 governor with parameters favoring performance over power savings.

[0108] Control begins at 700 and selects the first core of the set of cores. At 704, control determines the current frequency for the selected core, as well as determining residency in the available power states since the prior execution of this control. At 708, if there are
20 additional cores in the set to evaluate, control transfers to 712; otherwise, control transfers to 716. At 712, control selects the next core in the set and returns to 704.

[0109] At 716, control calculates (A) the relative residency of the cores in the desired power states as well as (B) the relative residency of the cores in undesired power states. Specifically, the relative residencies in 716 are averaged across all of the cores in the set. At 720, control calculates an average effective utilization of the cores based on the ratio of
25 A to the sum of A and B.

[0110] At 724, if the average effective utilization is greater than the threshold, control transfers to 728; otherwise, control transfers to 732. At 728, control increases the target frequency by an offset and continues at 736. At 736, control applies an upper limit to the target frequency and continues at 740. At 732, control sets the target frequency of the set
30 of cores to the product of the average effective utilization and the current frequency. Control continues at 740 and instructs the processor to operate the set of cores at the target frequency. Control then ends.

[0111] FIG. 13B is similar to FIG. 13A, with the exception being that the effective utilization of each core is determined separately and the individual effective utilizations

are averaged to determine average effective utilization. Elements of FIG. 13B that are similar to those of FIG. 13A are labeled with the same reference numerals as FIG. 13A.

5 [0112] After 704, control continues at 752, where control calculates the relative residency of the core in the desired power states as well as the relative residency of the core in undesired power states. At 756, control calculates the effective utilization for the selected core. Control then continues at 708, where if there are additional cores in the set, control transfers to 712; otherwise, control transfers to 760. At 760, control averages the effective utilizations of the cores to create an average effective utilization. Control then continues at 724.

10 [0113] In FIG. 14, control begins on a 10-millisecond schedule at 800. Control selects the first core and continues at 804. For the selected core, control reads the C-states in which the selected core resided since the last execution of the FIG. 14 control. At 808, if there are additional cores, control transfers to 812; otherwise, control transfers to 816. At 812, control selects the next core and returns to 804.

15 [0114] At 816, control selects the first core. At 820, control calculates the fractional residencies of the selected core in the C0 and C6 power states. At 824, control calculates an effective utilization based on the ratio of fC0 to the sum of fC0 and fC6. At 828, if the effective utilization is greater than a threshold, such as 0.99, control transfers to 832. Otherwise, control transfers to 836.

20 [0115] At 832, if the target frequency is over a threshold indicating turbo (high power) operation, control transfers to 840; otherwise, control transfers to 844. At 840, control sets the offset to 10 MHz and continues at 848. At 844, control sets the offset to 100 MHz and continues at 848. At 848, control increases the target frequency by the offset and continues at 852. If the target frequency is greater than a maximum upper limit, control transfers to 25 856; otherwise, control transfer to 860.

[0116] At 856, control sets the target frequency to the upper limit and continues at 860. At 860, control writes the target frequency to the model specific register of the selected core. At 864, control estimates the current frequency of the selected core as the target frequency. This current frequency is then used in the next iteration of the loop in 836.

30 Control continues at 868 where, if there are additional cores, control transfers to 872; otherwise, control ends. At 872, control selects the next core and returns to 820.

CONCLUSION

[0117] The foregoing description is merely illustrative in nature and is in no way intended to limit the disclosure, its application, or uses. The broad teachings of the

disclosure can be implemented in a variety of forms. Therefore, while this disclosure includes particular examples, the true scope of the disclosure should not be so limited since other modifications will become apparent upon a study of the drawings, the specification, and the following claims. It should be understood that one or more steps
5 within a method may be executed in different order (or concurrently) without altering the principles of the present disclosure. Further, although each of the embodiments is described above as having certain features, any one or more of those features described with respect to any embodiment of the disclosure can be implemented in and/or combined
10 with features of any of the other embodiments, even if that combination is not explicitly described. In other words, the described embodiments are not mutually exclusive, and permutations of one or more embodiments with one another remain within the scope of this disclosure.

[0118] Spatial and functional relationships between elements (for example, between modules) are described using various terms, including “connected,” “engaged,”
15 “interfaced,” and “coupled.” Unless explicitly described as being “direct,” when a relationship between first and second elements is described in the above disclosure, that relationship encompasses a direct relationship where no other intervening elements are present between the first and second elements, and also an indirect relationship where one or more intervening elements are present (either spatially or functionally) between the first
20 and second elements. As used herein, the phrase at least one of A, B, and C should be construed to mean a logical (A OR B OR C), using a non-exclusive logical OR, and should not be construed to mean “at least one of A, at least one of B, and at least one of C.”

[0119] In the figures, the direction of an arrow, as indicated by the arrowhead, generally
25 demonstrates the flow of information (such as data or instructions) that is of interest to the illustration. For example, when element A and element B exchange a variety of information but information transmitted from element A to element B is relevant to the illustration, the arrow may point from element A to element B. This unidirectional arrow does not imply that no other information is transmitted from element B to element A.
30 Further, for information sent from element A to element B, element B may send requests for, or receipt acknowledgements of, the information to element A.

[0120] In this application, including the definitions below, the term ‘module’ or the term ‘controller’ may be replaced with the term ‘circuit.’ The term ‘module’ may refer to, be part of, or include processor hardware (shared, dedicated, or group) that executes code and

memory hardware (shared, dedicated, or group) that stores code executed by the processor hardware.

5 [0121] The module may include one or more interface circuits. In some examples, the interface circuits may include wired or wireless interfaces that are connected to a local area network (LAN), the Internet, a wide area network (WAN), or combinations thereof. The functionality of any given module of the present disclosure may be distributed among multiple modules that are connected via interface circuits. For example, multiple modules may allow load balancing. In a further example, a server (also known as remote, or cloud) module may accomplish some functionality on behalf of a client module.

10 [0122] The term code, as used above, may include software, firmware, and/or microcode, and may refer to programs, routines, functions, classes, data structures, and/or objects. Shared processor hardware encompasses a single microprocessor that executes some or all code from multiple modules. Group processor hardware encompasses a microprocessor that, in combination with additional microprocessors, executes some or all code from one or more modules. References to multiple microprocessors encompass
15 multiple microprocessors on discrete dies, multiple microprocessors on a single die, multiple cores of a single microprocessor, multiple threads of a single microprocessor, or a combination of the above.

[0123] Shared memory hardware encompasses a single memory device that stores some
20 or all code from multiple modules. Group memory hardware encompasses a memory device that, in combination with other memory devices, stores some or all code from one or more modules.

[0124] The term memory hardware is a subset of the term computer-readable medium. The term computer-readable medium, as used herein, does not encompass transitory
25 electrical or electromagnetic signals propagating through a medium (such as on a carrier wave); the term computer-readable medium is therefore considered tangible and non-transitory. Non-limiting examples of a non-transitory computer-readable medium are nonvolatile memory devices (such as a flash memory device, an erasable programmable read-only memory device, or a mask read-only memory device), volatile memory devices
30 (such as a static random access memory device or a dynamic random access memory device), magnetic storage media (such as an analog or digital magnetic tape or a hard disk drive), and optical storage media (such as a CD, a DVD, or a Blu-ray Disc).

[0125] The apparatuses and methods described in this application may be partially or fully implemented by a special purpose computer created by configuring a general purpose

computer to execute one or more particular functions embodied in computer programs. The functional blocks and flowchart elements described above serve as software specifications, which can be translated into the computer programs by the routine work of a skilled technician or programmer.

5 [0126] The computer programs include processor-executable instructions that are stored on at least one non-transitory computer-readable medium. The computer programs may also include or rely on stored data. The computer programs may encompass a basic input/output system (BIOS) that interacts with hardware of the special purpose computer, device drivers that interact with particular devices of the special purpose computer, one or
10 more operating systems, user applications, background services, background applications, etc.

[0127] The computer programs may include: (i) descriptive text to be parsed, such as HTML (hypertext markup language), XML (extensible markup language), or JSON (JavaScript Object Notation) (ii) assembly code, (iii) object code generated from source
15 code by a compiler, (iv) source code for execution by an interpreter, (v) source code for compilation and execution by a just-in-time compiler, etc. As examples only, source code may be written using syntax from languages including C, C++, C#, Objective-C, Swift, Haskell, Go, SQL, R, Lisp, Java®, Fortran, Perl, Pascal, Curl, OCaml, Javascript®, HTML5 (Hypertext Markup Language 5th revision), Ada, ASP (Active Server Pages),
20 PHP (PHP: Hypertext Preprocessor), Scala, Eiffel, Smalltalk, Erlang, Ruby, Flash®, Visual Basic®, Lua, MATLAB, SIMULINK, and Python®.

[0128] None of the elements recited in the claims are intended to be a means-plus-function element within the meaning of 35 U.S.C. §112(f) unless an element is expressly recited using the phrase “means for” or, in the case of a method claim, using the phrases
25 “operation for” or “step for.”

CLAIMS

1. A frequency governor system for a processor having a first core, the system comprising:
 - a power state data collection subsystem configured to read power management information from the processor, wherein the processor is configured to selectively operate in a plurality of power saving states, and wherein the power management information includes (i) first information indicating occupancy of the first core of the processor in an active state and (ii) second information indicating occupancy of the first core of the processor in a first power saving state of the plurality of power saving states;
 - a utilization analyzer configured to generate an effective utilization based on the first information and the second information; and
 - a frequency control circuit configured to:
 - set a target frequency by performing a selected action of:
 - a first action including increasing the target frequency; and
 - a second action including decreasing the target frequency with respect to a current frequency of the processor based on the effective utilization,
 - wherein the selected action is selected according to the effective utilization; and
 - instruct the processor to operate the first core at the target frequency.
2. The system of claim 1, wherein the frequency control circuit is configured to:
 - select the first action if the effective utilization is greater than a threshold; and
 - select the second action if the effective utilization is less than the threshold.
3. The system of claim 1, wherein the frequency control circuit is configured to decrease the target frequency by setting the target frequency to a product of the effective utilization and the current frequency of the processor.
4. The system of claim 1, wherein:
 - the utilization analyzer is configured to generate the effective utilization based on a first number and a second number;
 - the first number is based on the first information;
 - the first number indicates a fraction of time the first core occupied the active state since a prior analysis by the utilization analyzer;
 - the second number is based on the second information;
 - the second number indicates a fraction of time the first core occupied a subset of the plurality of power saving states since the prior analysis; and

the subset of the plurality of power saving states is a proper subset and includes at least the first power saving state.

5. The system of claim 4, wherein the utilization analyzer is configured to generate the effective utilization based on a division of the first number by a sum of the first number and the second number.

6. The system of claim 1, wherein:

the processor includes multiple cores including the first core; and

the frequency control circuit is configured to instruct the processor to operate all of the multiple cores at the target frequency.

7. The system of claim 6, wherein the utilization analyzer is configured to generate the effective utilization by one of:

determining an individual effective utilization for each core of the multiple cores and averaging the individual effective utilizations to generate the effective utilization; and

averaging the first information across the multiple cores, averaging the second information across the multiple cores, and generating the effective utilization based on the averaged first information and the averaged second information.

8. A frequency governing method for a processor, the method comprising:

reading power management information from the processor, wherein the processor is configured to selectively operate in a plurality of power saving states, and wherein the power management information includes (i) first information indicating occupancy of a first core of the processor in an active state and (ii) second information indicating occupancy of the first core of the processor in a first power saving state of the plurality of power saving states;

generating an effective utilization based on the first information and the second information;

setting a target frequency by performing a selected action of:

a first action including increasing the target frequency; and

a second action including decreasing the target frequency with respect to a current frequency of the processor based on the effective utilization,

wherein the selected action is selected according to the effective utilization;

and

instructing the processor to operate the first core at the target frequency.

9. The method of claim 8, wherein:

the first action is selected if the effective utilization is greater than a threshold; and

the second action is selected if the effective utilization is less than the threshold.

10. The method of claim 8, wherein:

the effective utilization is generated based on a first number and a second number;

the first number is based on the first information;

the first number indicates a fraction of time the first core occupied the active state since a prior analysis;

the second number is based on the second information;

the second number indicates a fraction of time the first core occupied a subset of the plurality of power saving states since the prior analysis; and

the subset of the plurality of power saving states is a proper subset and includes at least the first power saving state.

11. The method of claim 8, wherein decreasing the target frequency includes setting the target frequency to a product of the effective utilization and the current frequency of the processor.

12. The method of claim 11, wherein the effective utilization is generated based on a division of the first number by a sum of the first number and the second number.

13. The method of claim 8, wherein:

the processor includes multiple cores including the first core;

the method includes instructing the processor to operate all of the multiple cores at the target frequency; and

the method includes one of:

determining an individual effective utilization for each core of the multiple cores and averaging the individual effective utilizations to generate the effective utilization; and

averaging the first information across the multiple cores, averaging the second information across the multiple cores, and generating the effective utilization based on the averaged first information and the averaged second information.

14. The method of claim 8, wherein:

the processor includes multiple cores including the first core; and

the method includes performing the reading, the generating, the setting, and the instructing separately for each core of the multiple cores.

15. The method of claim 8, wherein, from among the plurality of power saving states, a latency of the first core to return to the active state is greatest for the first power saving state.




FIG. 1
Prior Art




FIG. 2
Prior Art




FIG. 4

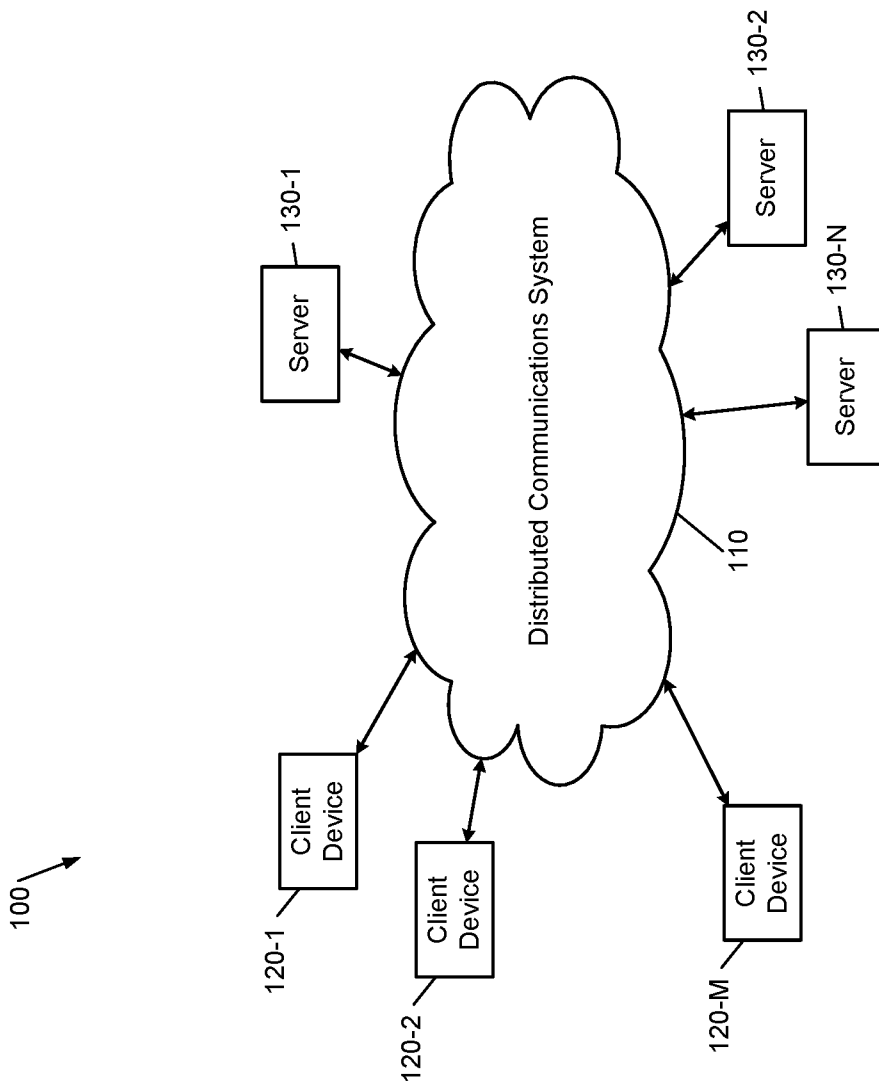


FIG. 3




FIG. 5

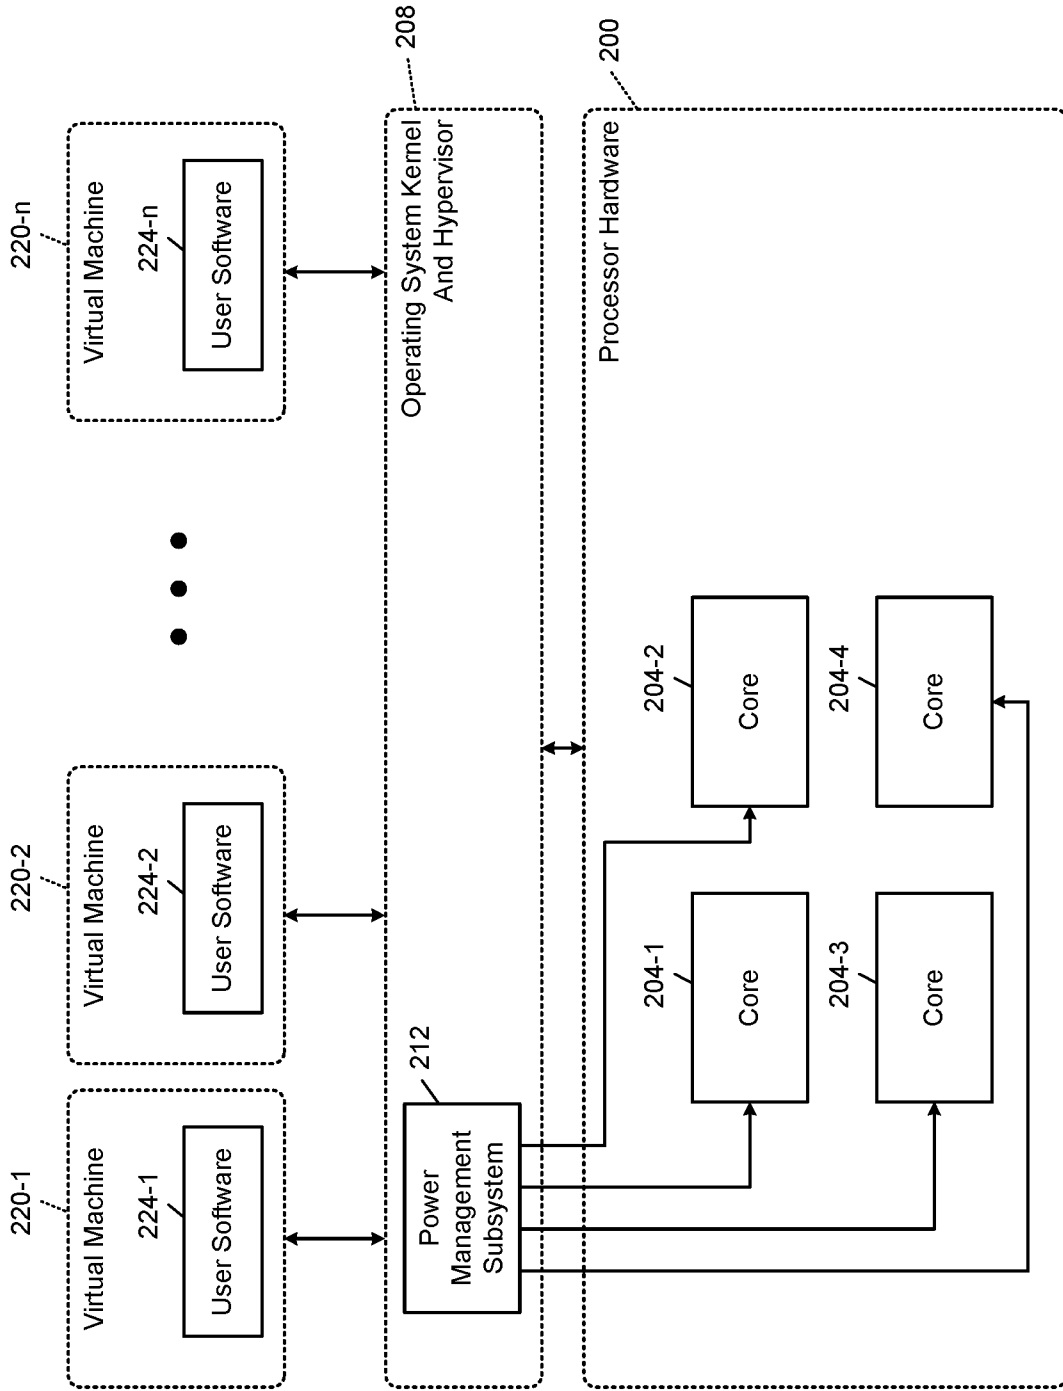


FIG. 6




FIG. 7

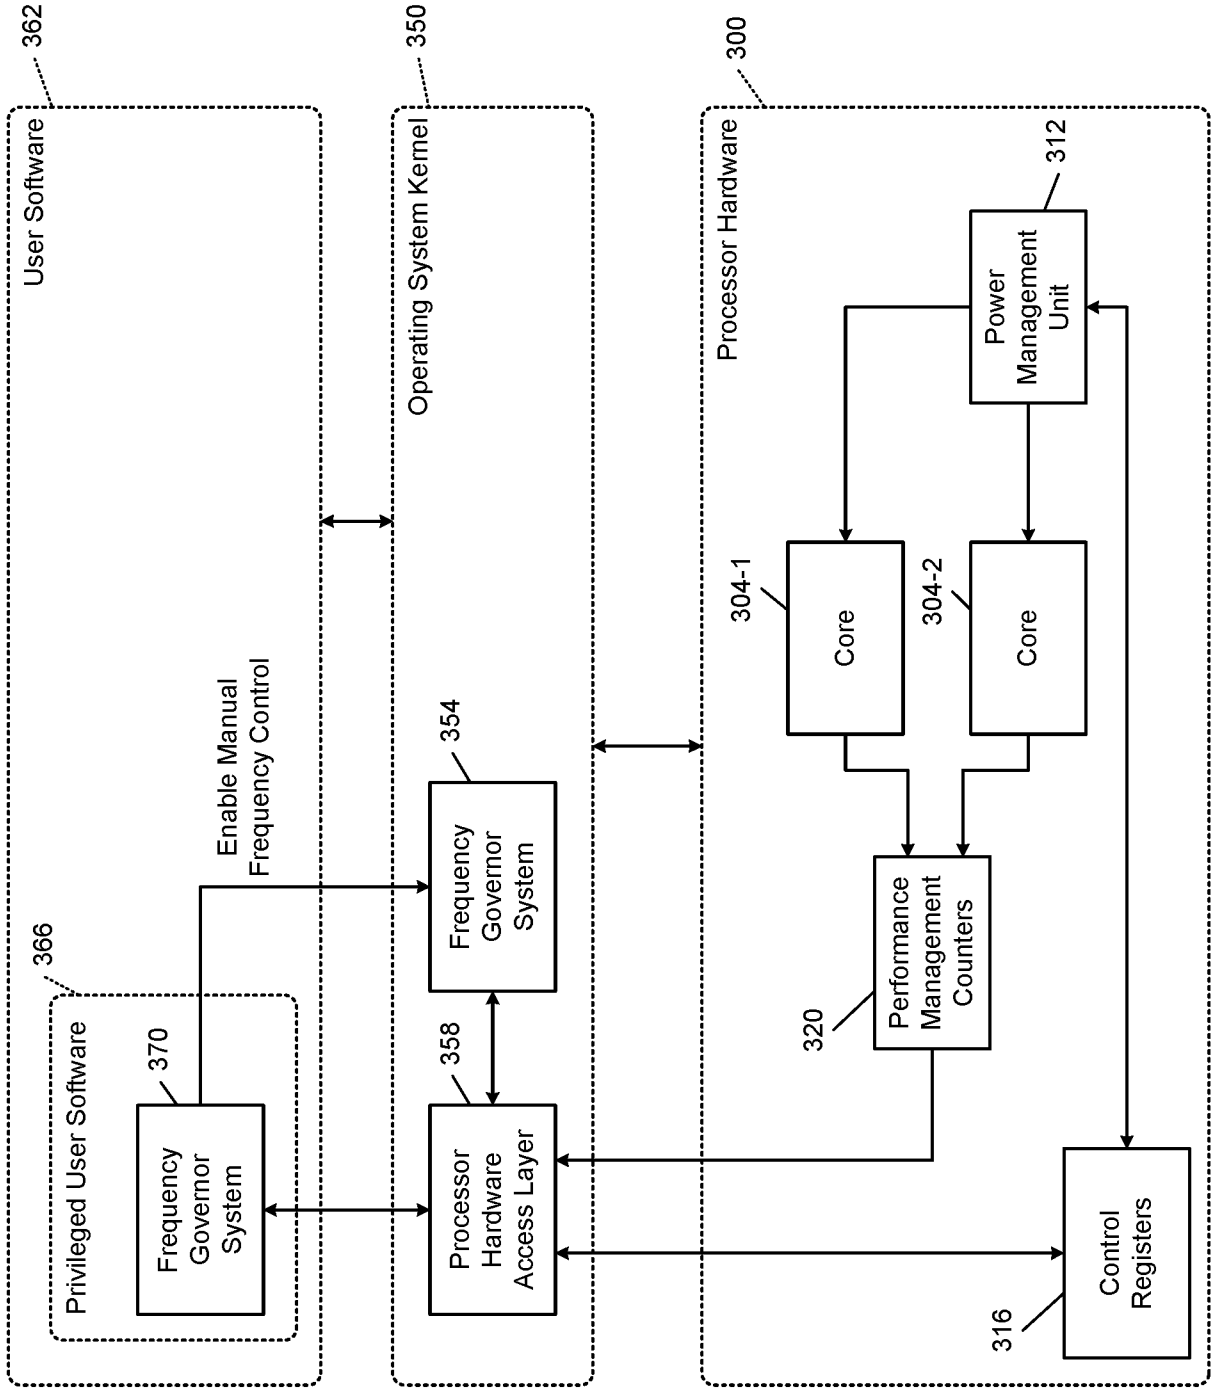


FIG. 8




FIG. 9

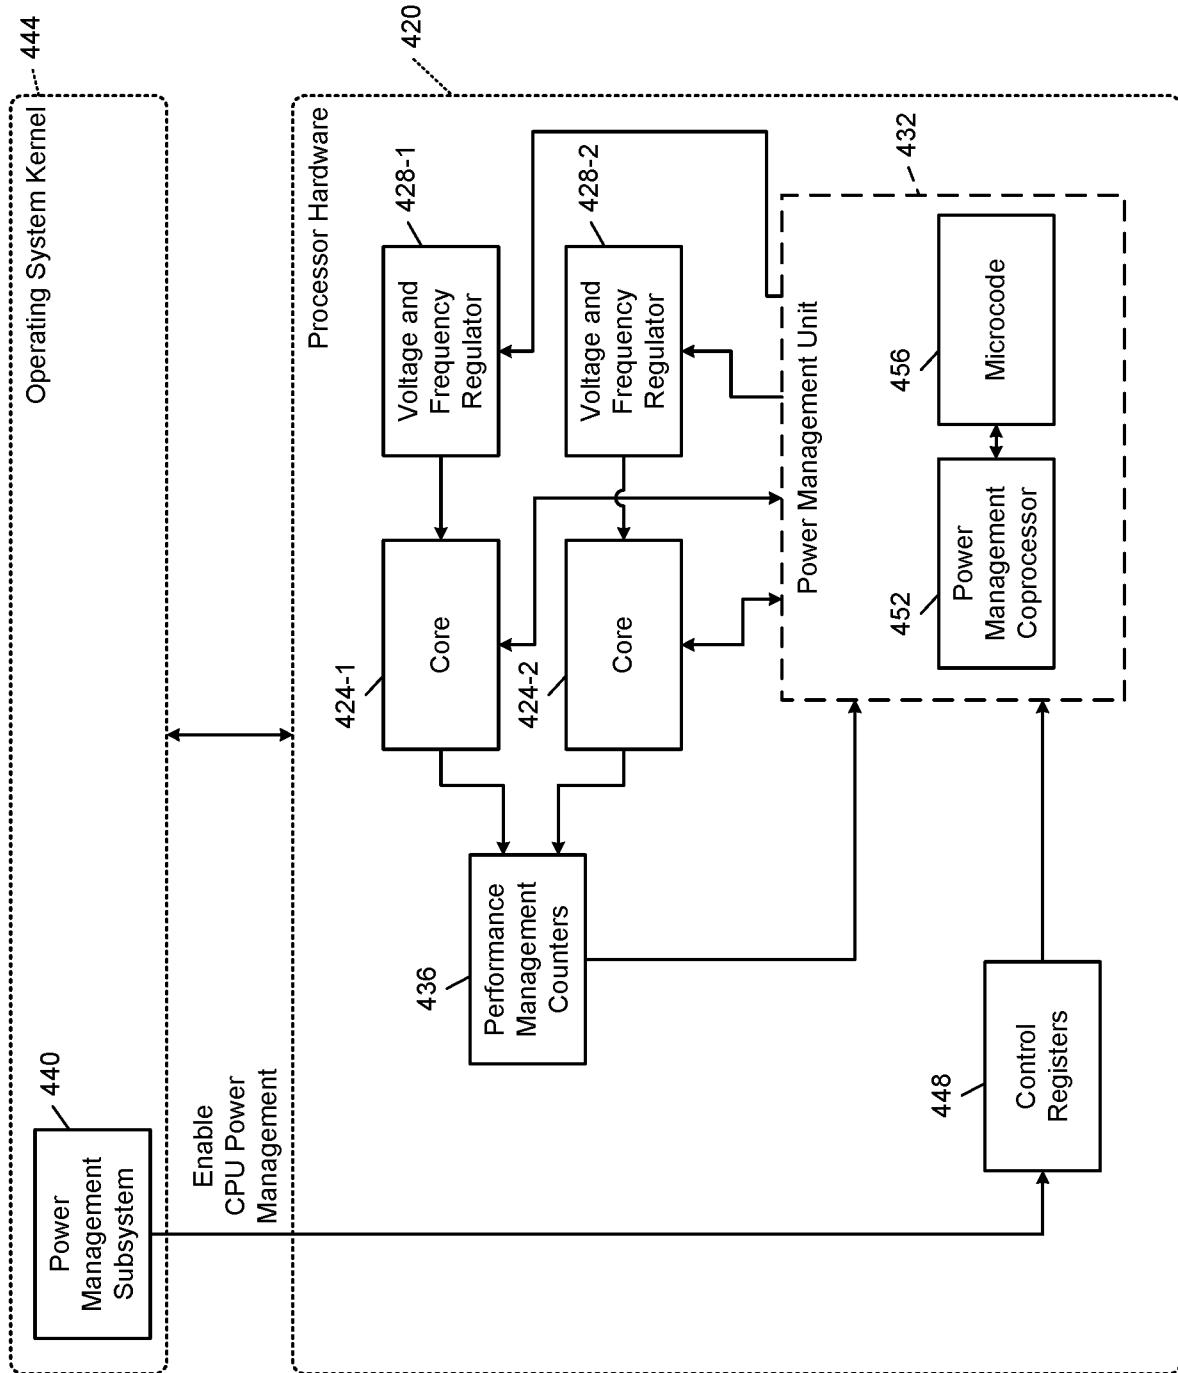


FIG. 10




FIG. 11

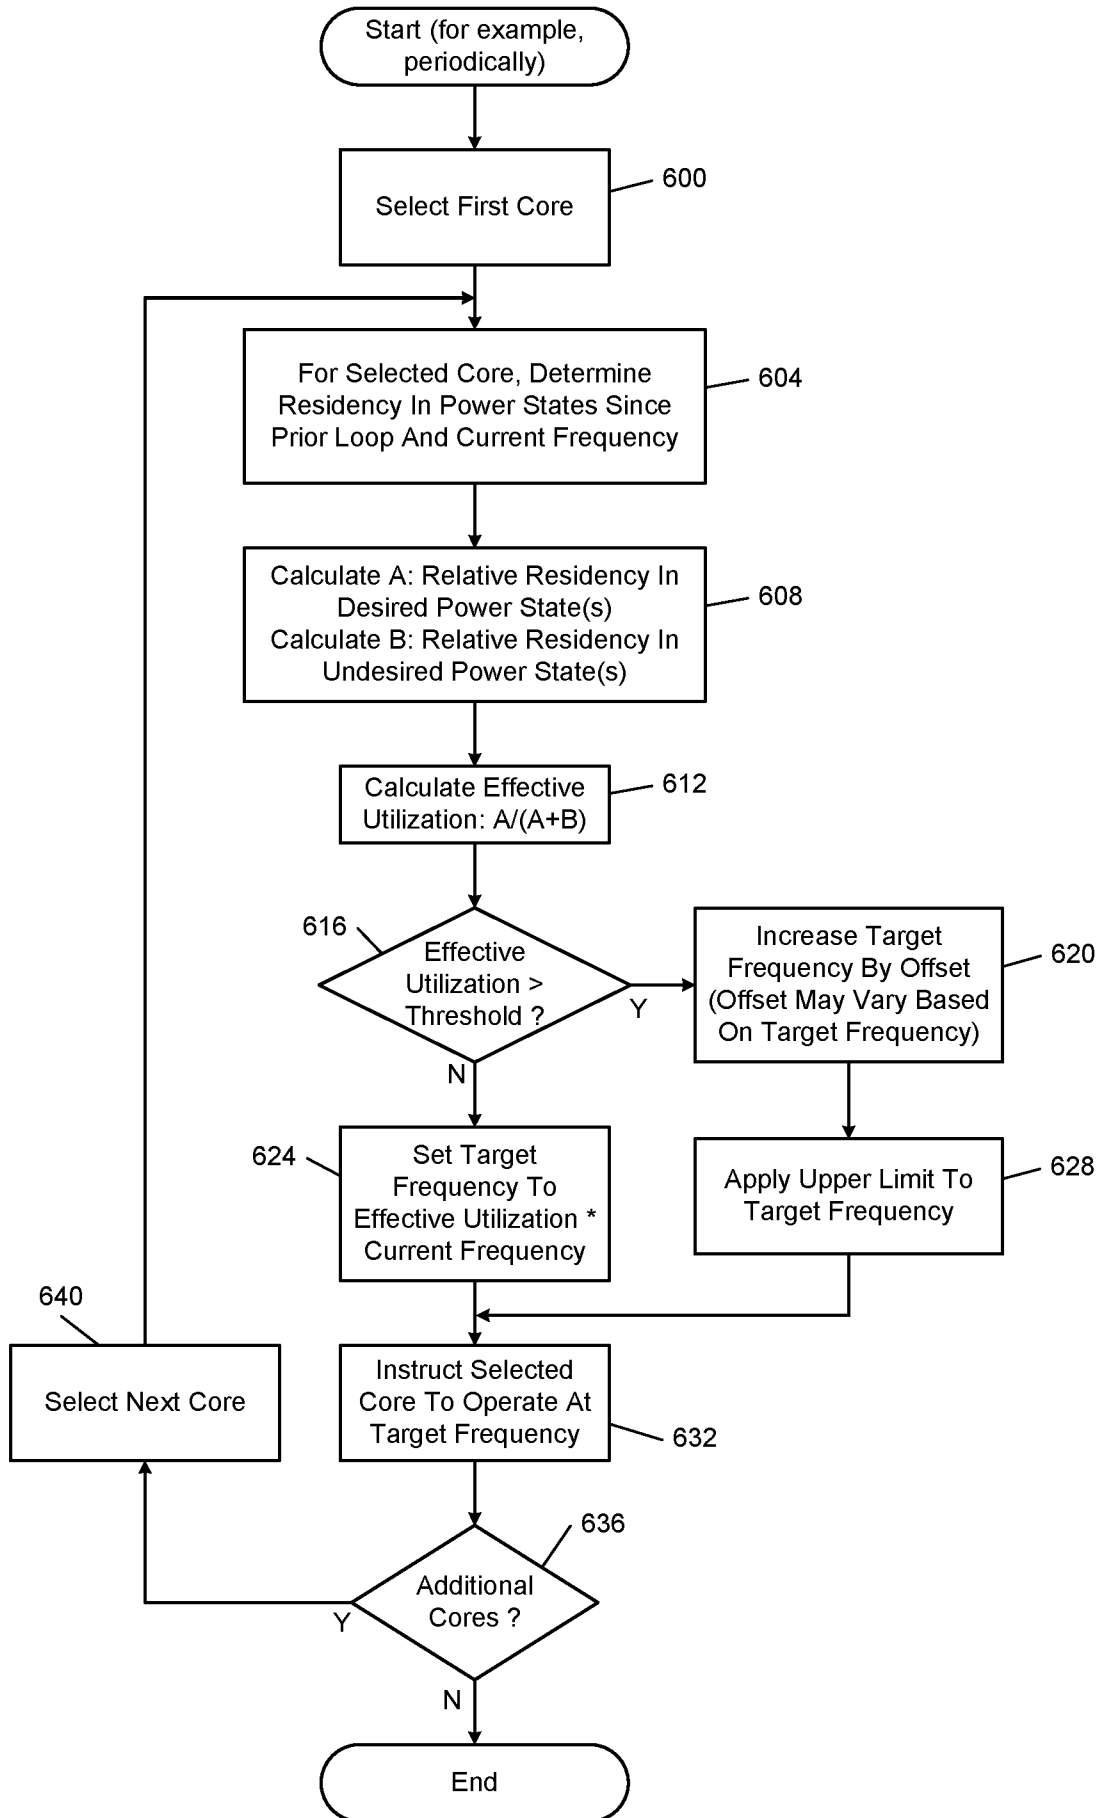


FIG. 12

12/15

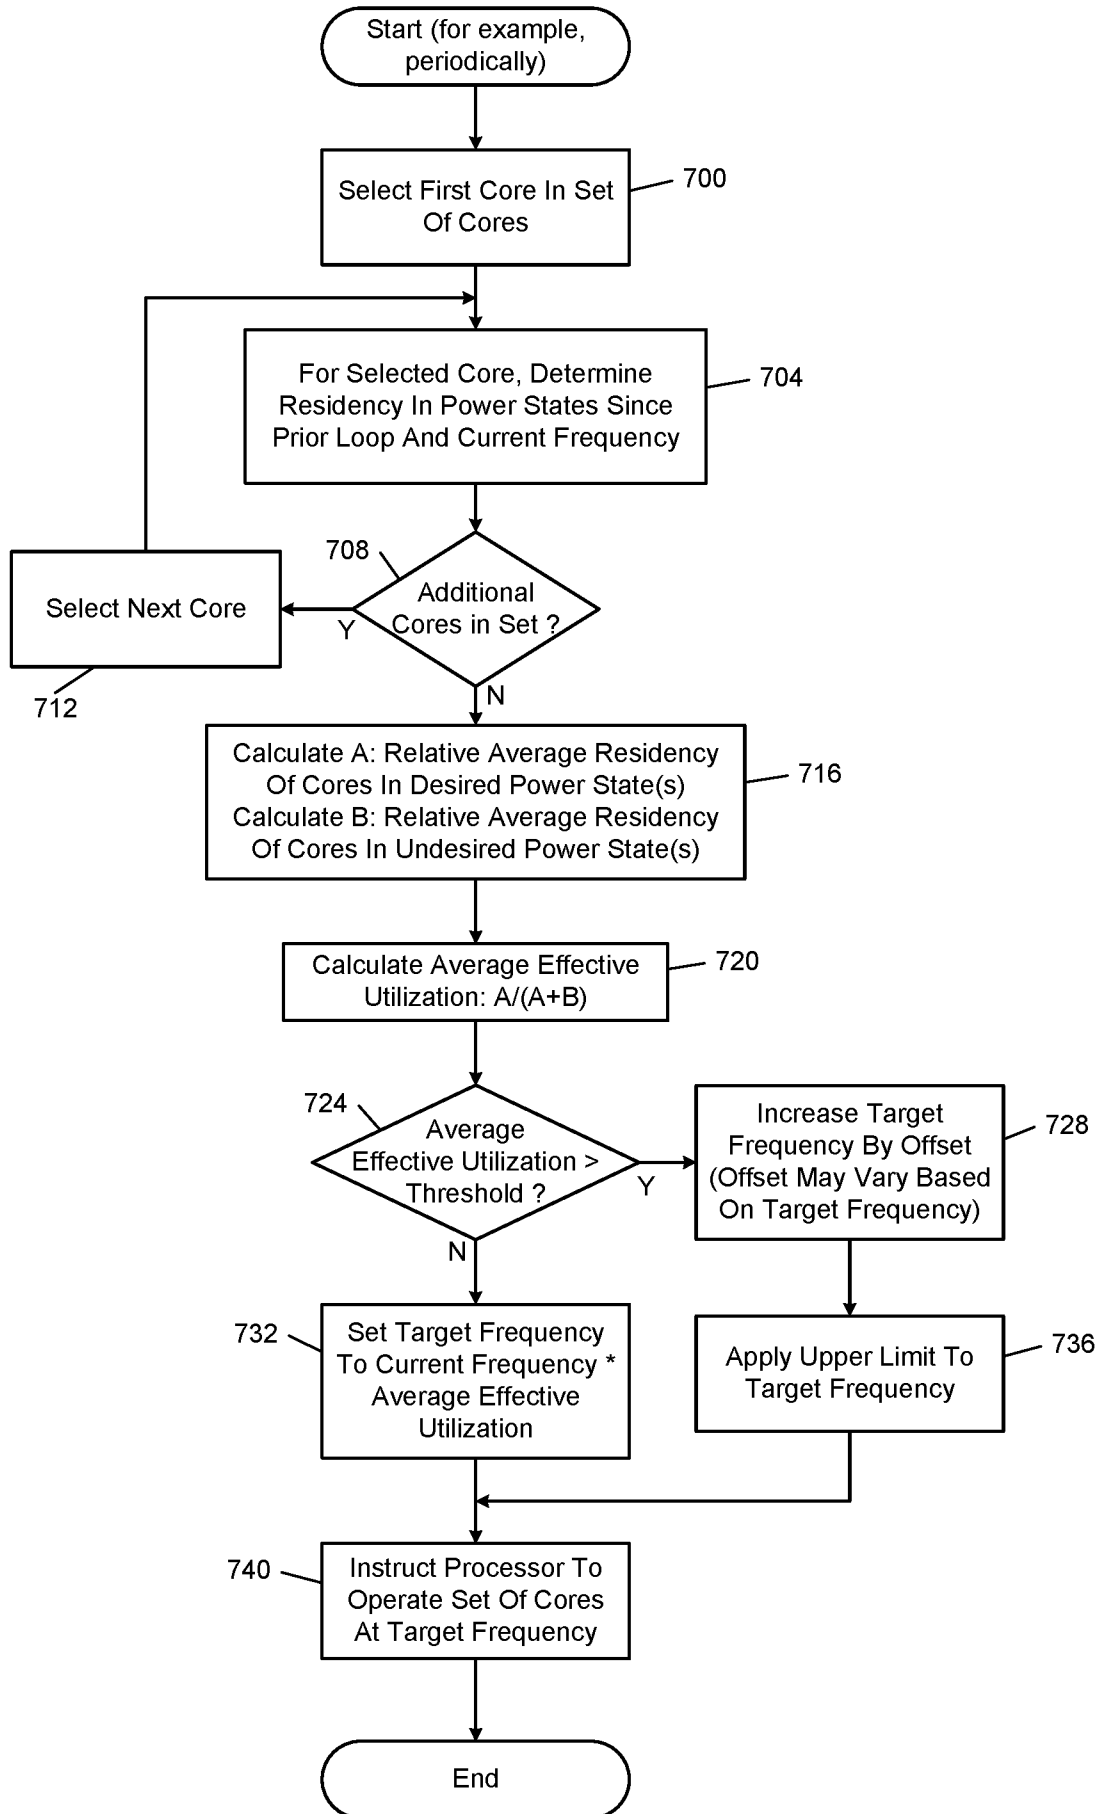


FIG. 13A

13/15

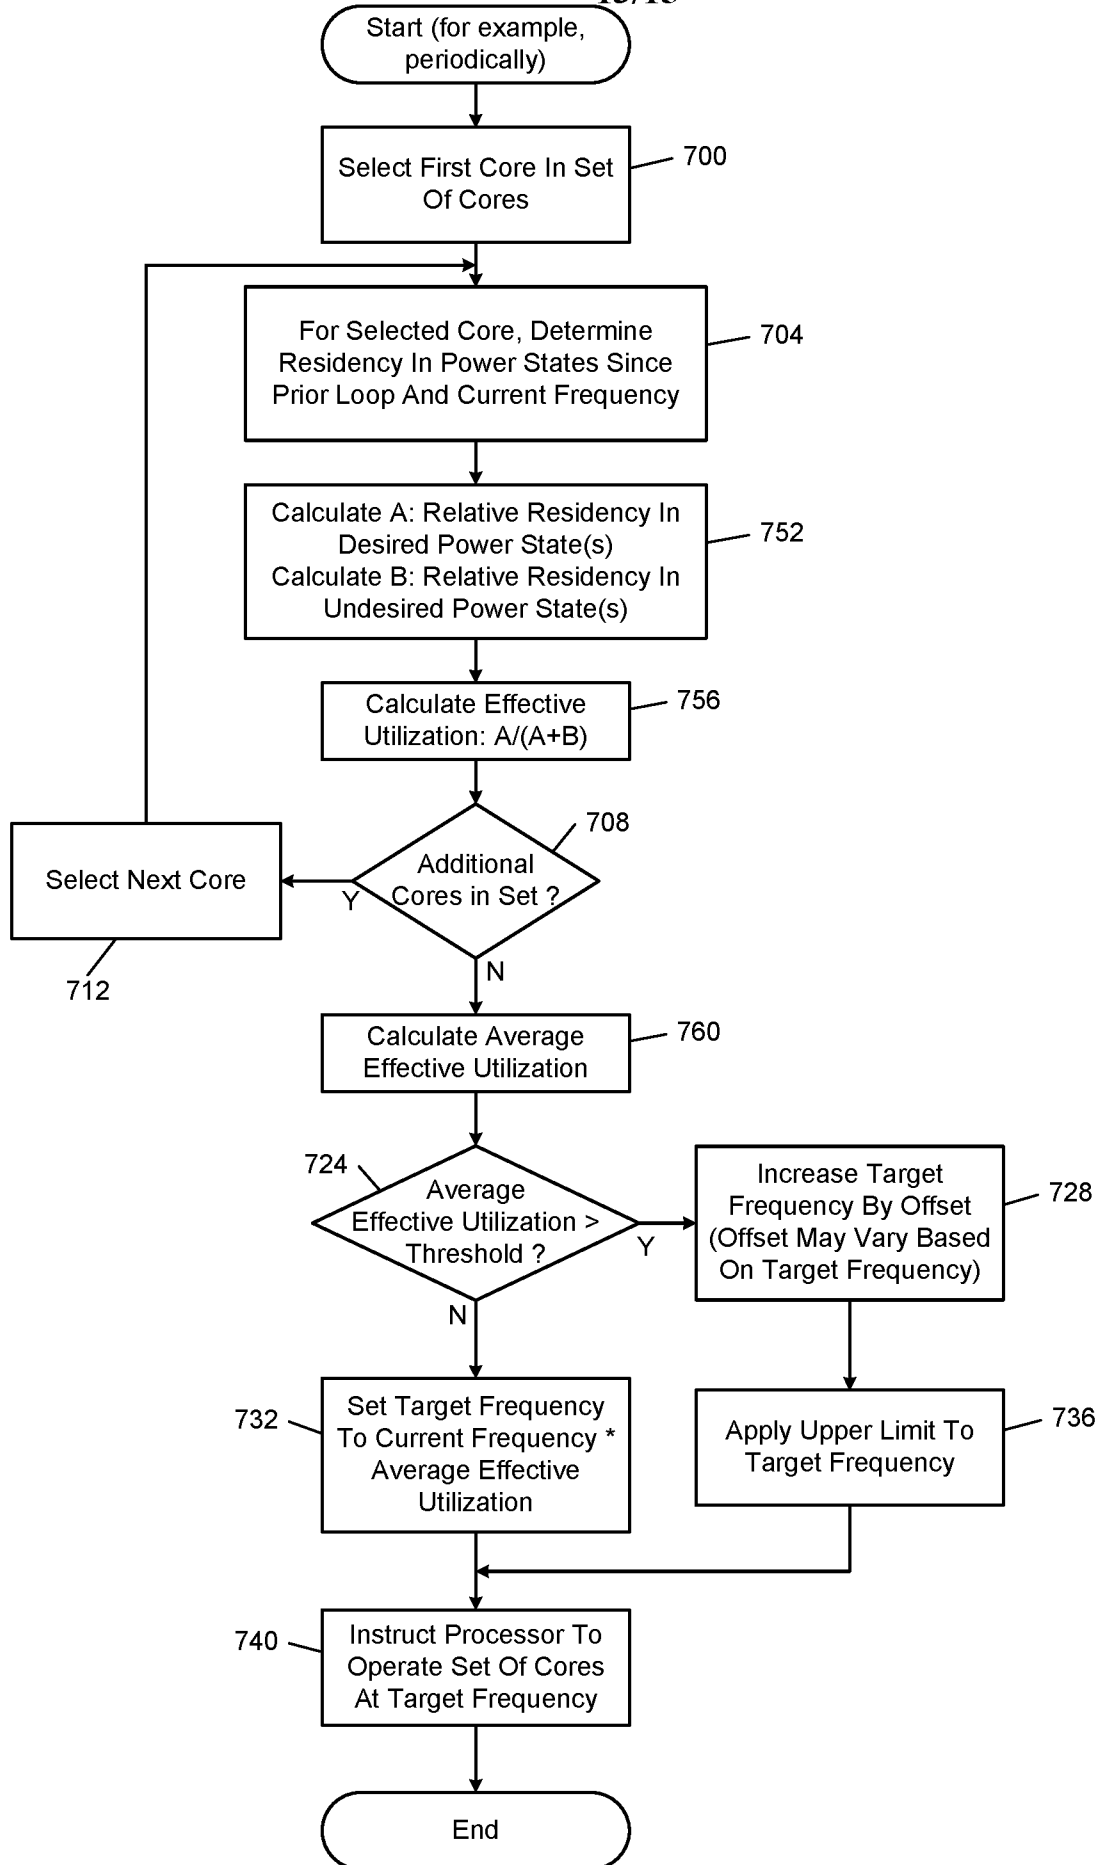


FIG. 13B

14/15

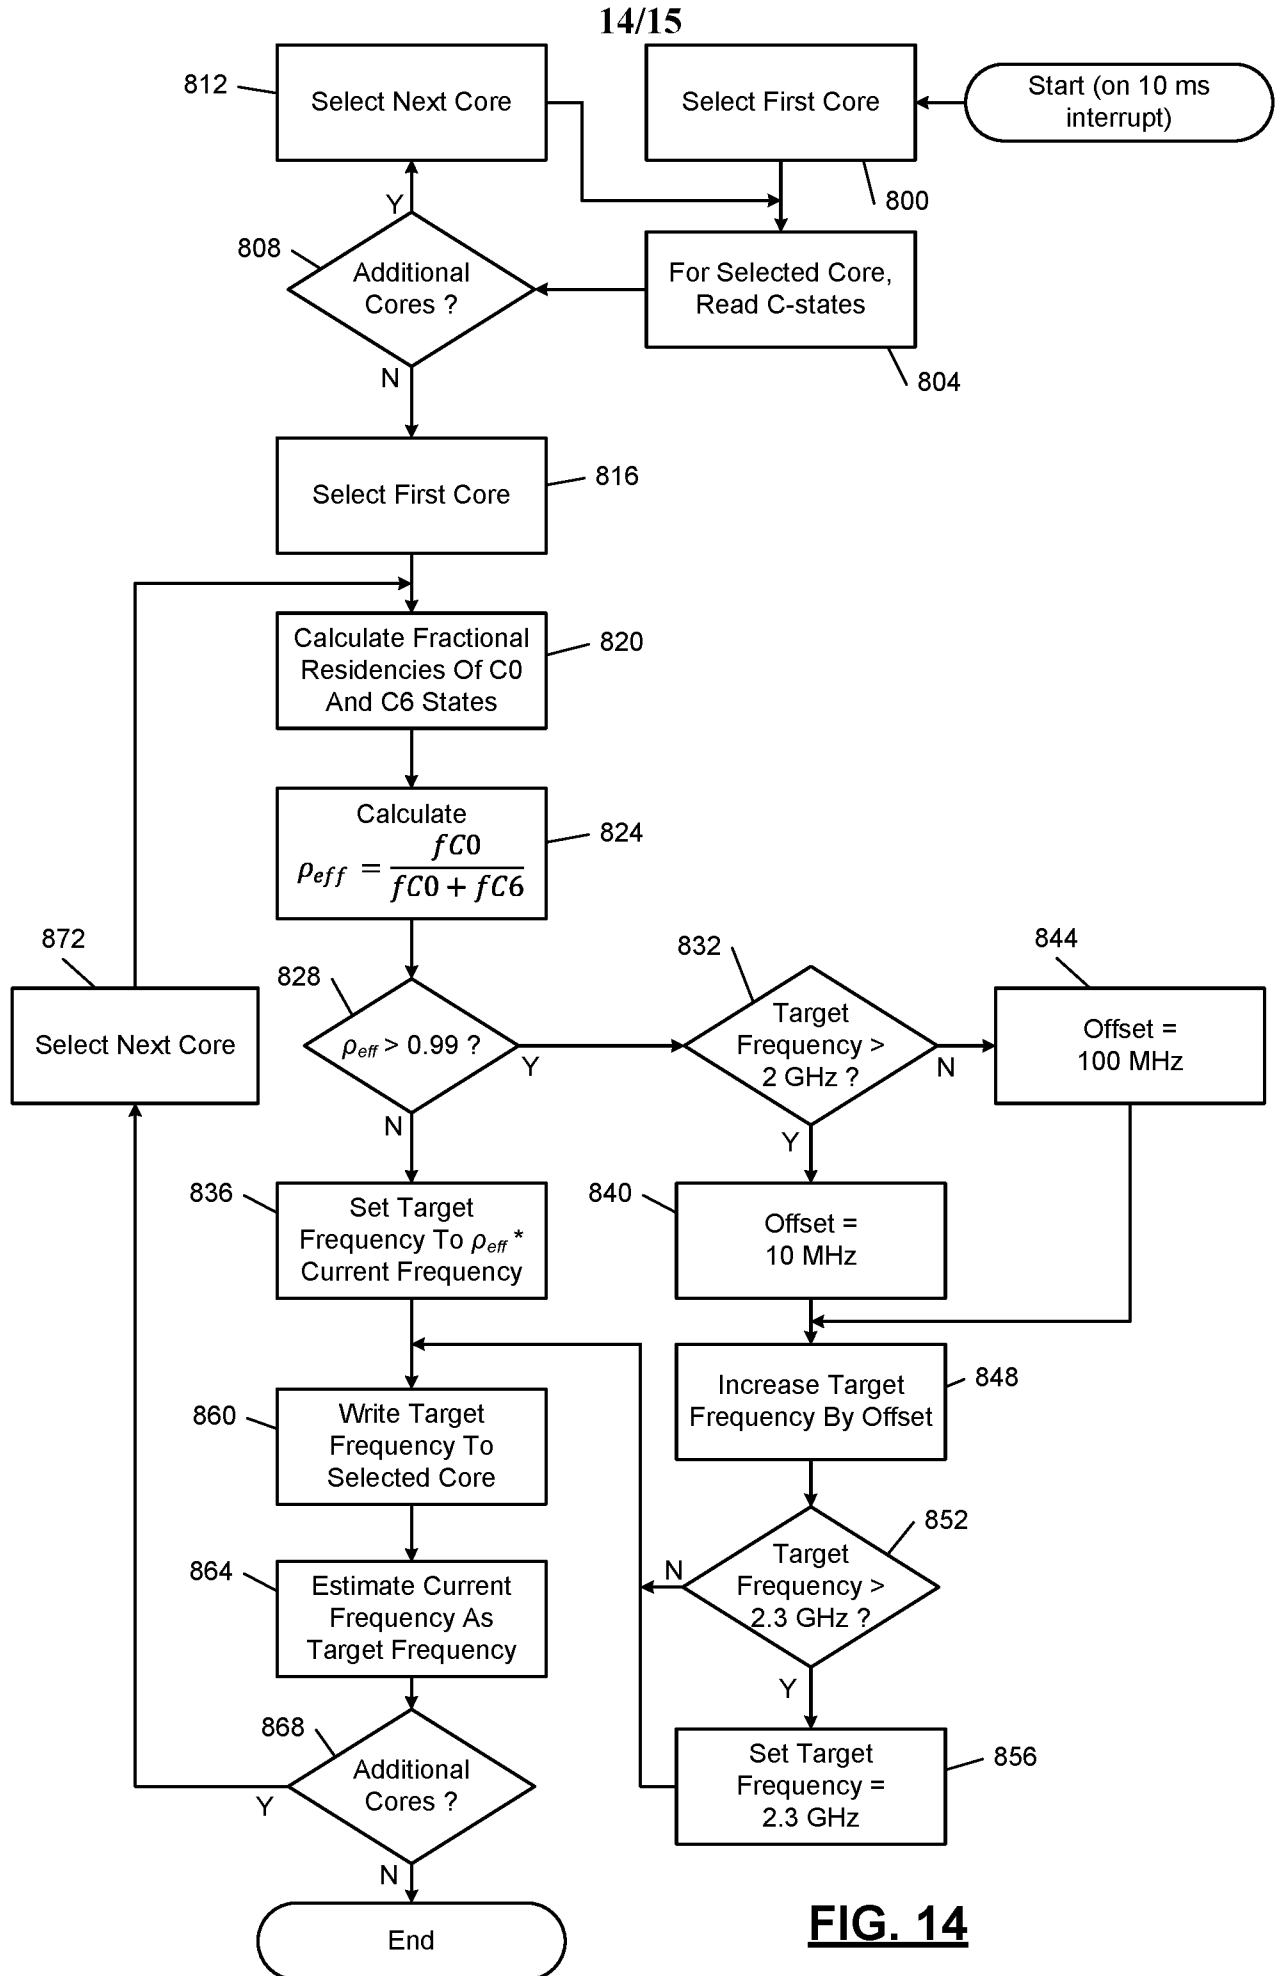


FIG. 14




FIG. 15A

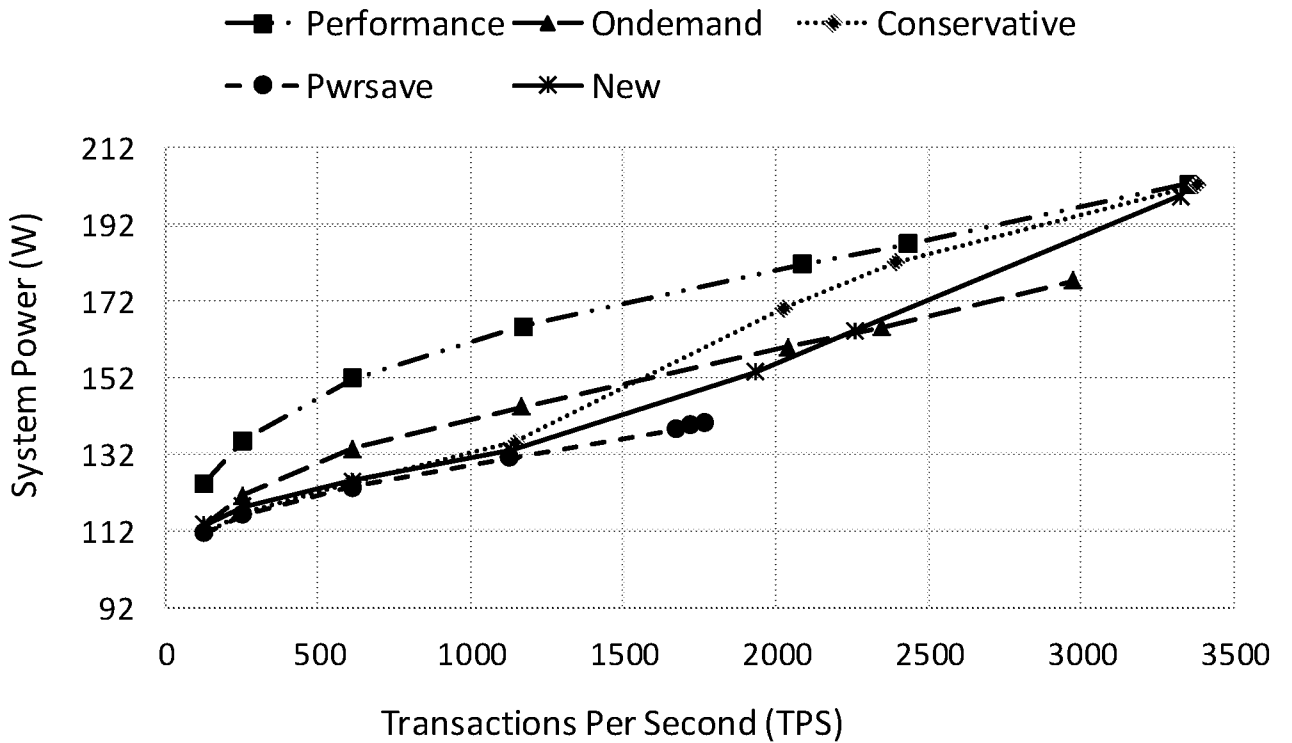


FIG. 15B

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2018/035251

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F1/32
ADD.

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

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	Dirk Brandewie ET AL: "Extract from linux/drivers/cpufreq/intel_pstate.c", 29 June 2017 (2017-06-29), pages 1-8, XP055506271, Retrieved from the Internet: URL:https://github.com/torvalds/linux/blob/fab24dcc395637557a7988a867e7b3a5823917a9/drivers/cpufreq/intel_pstate.c [retrieved on 2018-09-11]	1-3,6-9, 11,13-15
A	page 1, line 2 page 2, lines 17-68 page 3, lines 47-78 page 4, line 50 - page 5, line 38 page 6, lines 27-54 page 7, lines 2-68 page 8, lines 1-69 ----- -/--	4,5,10, 12

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 another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search 13 September 2018	Date of mailing of the international search report 27/09/2018
--	--

Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer Millet, Guillaume
--	---

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2018/035251

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2008/148027 A1 (FENGER RUSSELL J [US] ET AL) 19 June 2008 (2008-06-19) paragraphs [0001] - [0003] paragraphs [0030] - [0033] -----	1-15
A	US 2011/283286 A1 (WU WUXIAN [US] ET AL) 17 November 2011 (2011-11-17) paragraphs [0024] - [0034] -----	1-15
A	US 2010/332856 A1 (SONG JUSTIN J [US]) 30 December 2010 (2010-12-30) paragraphs [0020] - [0029] -----	1-15

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2018/035251

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2008148027	A1	19-06-2008	NONE

US 2011283286	A1	17-11-2011	NONE

US 2010332856	A1	30-12-2010	EP 2270625 A1 05-01-2011
		US 2010332856 A1	30-12-2010
