



US009760160B2

(12) **United States Patent**
Weissmann et al.

(10) **Patent No.:** **US 9,760,160 B2**
(45) **Date of Patent:** **Sep. 12, 2017**

(54) **CONTROLLING PERFORMANCE STATES OF PROCESSING ENGINES OF A PROCESSOR**

(71) Applicant: **Intel Corporation**, Santa Clara, CA (US)

(72) Inventors: **Eliezer Weissmann**, Haifa (IL); **Efraim Rotem**, Haifa (IL); **Hisham Abu Salah**, Majdal Shams (IL); **Yoni Aizik**, Haifa (IL); **Doron Rajwan**, Rishon Le-Zion (IL); **Nir Rosenzweig**, Givat Ella (IL); **Gal Leibovich**, Kiryat Yam (IL); **Yevgeni Sabin**, Haifa (IL); **Shay Levy**, Zicjron Yaakov (IL)

(73) Assignee: **Intel Corporation**, Santa Clara, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 138 days.

(21) Appl. No.: **14/722,518**

(22) Filed: **May 27, 2015**

(65) **Prior Publication Data**
US 2016/0349828 A1 Dec. 1, 2016

(51) **Int. Cl.**
G06F 1/26 (2006.01)
G06F 1/32 (2006.01)
(Continued)

(52) **U.S. Cl.**
CPC **G06F 1/3287** (2013.01); **G06F 1/324** (2013.01); **G06F 1/3243** (2013.01); **G06F 1/3296** (2013.01); **G06F 9/5094** (2013.01)

(58) **Field of Classification Search**
CPC G06F 1/32; G06F 1/26
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,163,153 A 11/1992 Cole et al.
5,522,087 A 5/1996 Hsiang
(Continued)

FOREIGN PATENT DOCUMENTS

EP 1 282 030 A1 5/2003

OTHER PUBLICATIONS

International Searching Authority, "Notification of Transmittal of the International Search Report and the Written Opinion of the International Searching Authority," mailed Jul. 26, 2016, in International application No. PCT/US2016029921.

(Continued)

Primary Examiner — Phil Nguyen

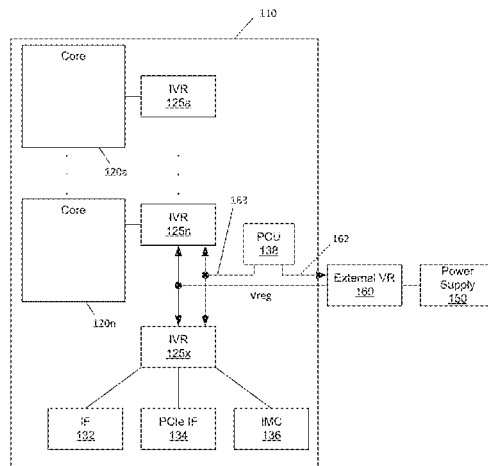
(74) *Attorney, Agent, or Firm* — Trop, Pruner & Hu, P.C.

(57) **ABSTRACT**

In one embodiment, a processor comprises: a plurality of processing engines including a first processing engine and a second processing engine to independently execute instructions; and a power controller including a performance state control logic to control a performance state of at least one of the processing engines, and a first logic to determine an average number of active processing engines over a first window, an estimated activity level of the processor for the first window, and adjust at least one of a window length at which the performance state control logic is to perform a performance state determination and at least one activity level threshold, based at least in part on a comparison of the estimated activity level and the average number of active processing engines. Other embodiments are described and claimed.

22 Claims, 22 Drawing Sheets

100



- (51) **Int. Cl.**
G06F 1/00 (2006.01)
G06F 9/50 (2006.01)

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,590,341 A 12/1996 Matter
 5,621,250 A 4/1997 Kim
 5,931,950 A 8/1999 Hsu
 6,748,546 B1 6/2004 Mirov et al.
 6,792,392 B1 9/2004 Knight
 6,823,516 B1 11/2004 Cooper
 6,829,713 B2 12/2004 Cooper et al.
 6,996,728 B2 2/2006 Singh
 7,010,708 B2 3/2006 Ma
 7,043,649 B2 5/2006 Terrell
 7,093,147 B2 8/2006 Farkas et al.
 7,111,179 B1 9/2006 Girson et al.
 7,194,643 B2 3/2007 Gonzalez et al.
 7,272,730 B1 9/2007 Acquaviva et al.
 7,412,615 B2 8/2008 Yokota et al.
 7,434,073 B2 10/2008 Magklis et al.
 7,437,270 B2 10/2008 Song et al.
 7,454,632 B2 11/2008 Kardach et al.
 7,529,956 B2 5/2009 Stufflebeam
 7,539,885 B2 5/2009 Ma
 7,730,340 B2 6/2010 Hu et al.
 2001/0044909 A1 11/2001 Oh et al.
 2002/0194509 A1 12/2002 Plante et al.
 2003/0061383 A1 3/2003 Zilka
 2004/0064752 A1 4/2004 Kazachinsky et al.
 2004/0098560 A1 5/2004 Storvik et al.
 2004/0139356 A1 7/2004 Ma
 2004/0268166 A1 12/2004 Farkas et al.
 2005/0022038 A1 1/2005 Kaushik et al.
 2005/0033881 A1 2/2005 Yao
 2005/0132238 A1 6/2005 Nanja
 2006/0050670 A1 3/2006 Hillyard et al.
 2006/0053326 A1 3/2006 Naveh
 2006/0059286 A1 3/2006 Bertone et al.
 2006/0069936 A1 3/2006 Lint et al.
 2006/0117202 A1 6/2006 Magklis et al.
 2006/0123253 A1* 6/2006 Morgan G06F 1/3228
 713/300
 2006/0184287 A1 8/2006 Belady et al.
 2007/0005995 A1 1/2007 Kardach et al.
 2007/0016817 A1 1/2007 Albonesi et al.
 2007/0079294 A1 4/2007 Knight
 2007/0106827 A1 5/2007 Boatright et al.
 2007/0156992 A1 7/2007 Jahagirdar
 2007/0214342 A1 9/2007 Newburn
 2007/0239398 A1 10/2007 Song et al.
 2007/0245163 A1 10/2007 Lu et al.
 2008/0028240 A1 1/2008 Arai et al.
 2008/0250260 A1 10/2008 Tomita
 2009/0006871 A1 1/2009 Liu et al.
 2009/0150695 A1 6/2009 Song et al.
 2009/0150696 A1 6/2009 Song et al.
 2009/0158061 A1 6/2009 Schmitz et al.
 2009/0158067 A1 6/2009 Bodas et al.
 2009/0172375 A1 7/2009 Rotem et al.
 2009/0172428 A1 7/2009 Lee
 2009/0235105 A1 9/2009 Branover et al.
 2009/0249094 A1 10/2009 Marshall et al.
 2010/0115309 A1 5/2010 Carvalho et al.
 2010/0146513 A1 6/2010 Song
 2010/0191997 A1 7/2010 Dodeja et al.

2011/0154090 A1 6/2011 Dixon et al.
 2012/0079290 A1 3/2012 Kumar
 2012/0144215 A1 6/2012 Naffziger et al.
 2012/0144217 A1 6/2012 Sistla et al.
 2012/0185709 A1* 7/2012 Weissmann G06F 1/3234
 713/320
 2012/0246506 A1 9/2012 Knight
 2013/0061064 A1 3/2013 Ananthkrishnan et al.
 2013/0080803 A1 3/2013 Ananthkrishnan et al.
 2013/0080804 A1 3/2013 Ananthkrishnan et al.
 2013/0111120 A1 5/2013 Ananthkrishnan et al.
 2013/0111121 A1 5/2013 Ananthkrishnan et al.
 2013/0111226 A1 5/2013 Ananthkrishnan et al.
 2013/0111236 A1 5/2013 Ananthkrishnan et al.
 2013/0346774 A1 12/2013 Bhandaru et al.
 2014/0068290 A1 3/2014 Bhandaru et al.
 2014/0082380 A1 3/2014 Schluessler et al.
 2014/0189399 A1 7/2014 Govindaraju et al.
 2014/0195829 A1 7/2014 Bhandaru et al.
 2014/0208141 A1 7/2014 Bhandaru et al.
 2016/0224100 A1* 8/2016 Chen G06F 9/5094

OTHER PUBLICATIONS

Intel Developer Forum, IDF2010, Opher Kahn, et al., "Intel Next Generation Microarchitecture Codename Sandy Bridge: New Processor Innovations," Sep. 13, 2010, 58 pages.
 SPEC-Power and Performance, Design Overview V1.10, Standard Performance Information Corp., Oct. 21, 2008, 6 pages.
 Intel Technology Journal, "Power and Thermal Management in the Intel Core Duo Processor," May 15, 2006, pp. 109-122.
 Anoop Iyer, et al., "Power and Performance Evaluation of Globally Asynchronous Locally Synchronous Processors," 2002, pp. 1-11.
 Greg Semeraro, et al., "Hiding Synchronization Delays in A GALS Processor Microarchitecture," 2004, pp. 1-13.
 Joan-Manuel Parcerisa, et al., "Efficient Interconnects for Clustered Microarchitectures," 2002, pp. 1-10.
 Grigorios Magklis, et al., "Profile-Based Dynamic Voltage and Frequency Scalling for a Multiple Clock Domain Microprocessor," 2003, pp. 1-12.
 Greg Semeraro, et al., "Dynamic Frequency and Voltage Control for a Multiple Clock Domain Architecture," 2002, pp. 1-12.
 Greg Semeraro, "Energy-Efficient Processor Design Using Multiple Clock Domains with Dynamic Voltage and Frequency Scalling," 2002, pp. 29-40.
 Diana Marculescu, "Application Adaptive Energy Efficient Clustered Architectures," 2004, pp. 344-349.
 L. Benin!, et al., "System-Level Dynamic Power Management," 1999, pp. 23-31.
 Ravindra Jejurikar, et al., "Leakage Aware Dynamic Voltage Scalling for Real-Time Embedded Systems," 2004, pp. 275-280.
 Ravindra Jejurikar, et al., "Dynamic Slack Reclamation With Procrastination Scheduling in Real-Time Embedded Systems," 2005, pp. 13-17.
 R. Todling, et al., "Some Strategies for Kalman Filtering and Smoothing," 1996, pp. 1-21.
 R.E. Kalman, "A New Approach to Linear Filtering and Prediction Problems," 1960, pp. 1-12.
 Intel Corporation, "Intel 64 and IA-32 Architectures Software Developer's Manual," vol. 3 (3A, 3B & 3C): System Programming Guide, Feb. 2014, Chapter 14 Power and Thermal Management (14.1-14.9.5), 44 pages.
 U.S. Appl. No. 14/484,649, filed Sep. 12, 2014, entitled "Optimization of Energy Usage in a Processor," by Yoni Aizik.

* cited by examiner

100

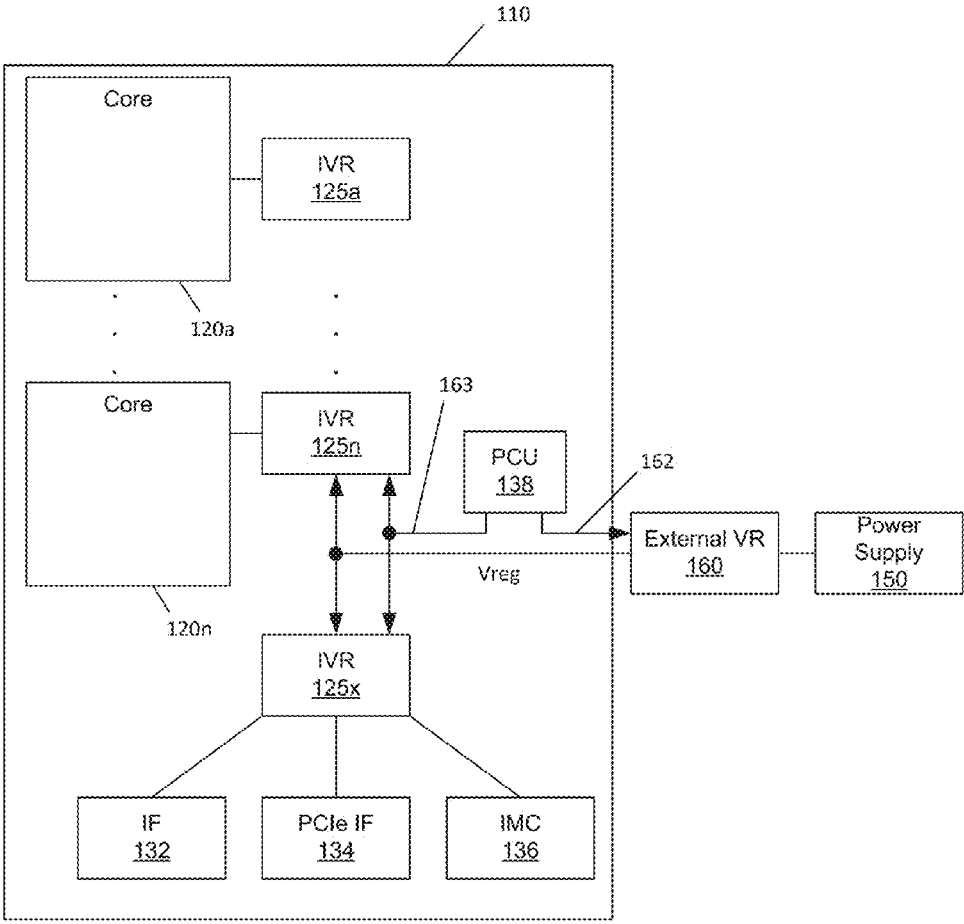


FIG. 1

200

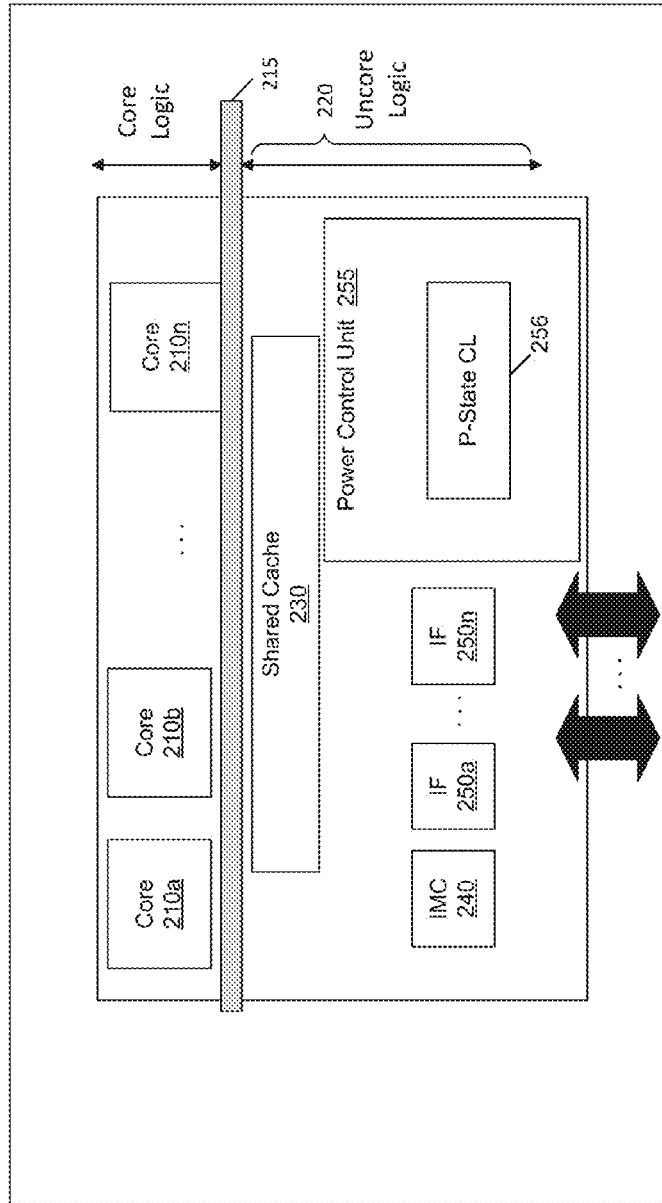


FIG. 2

300

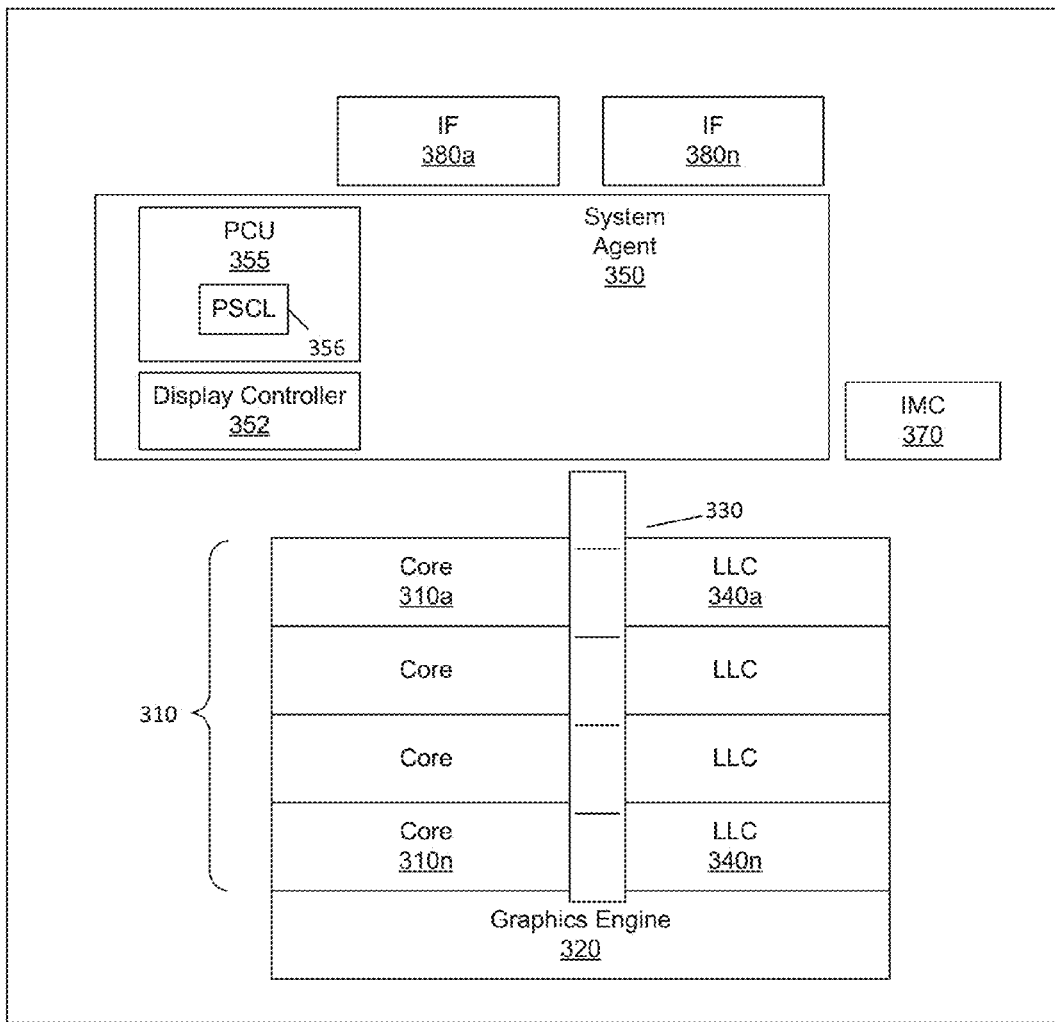


FIG. 3

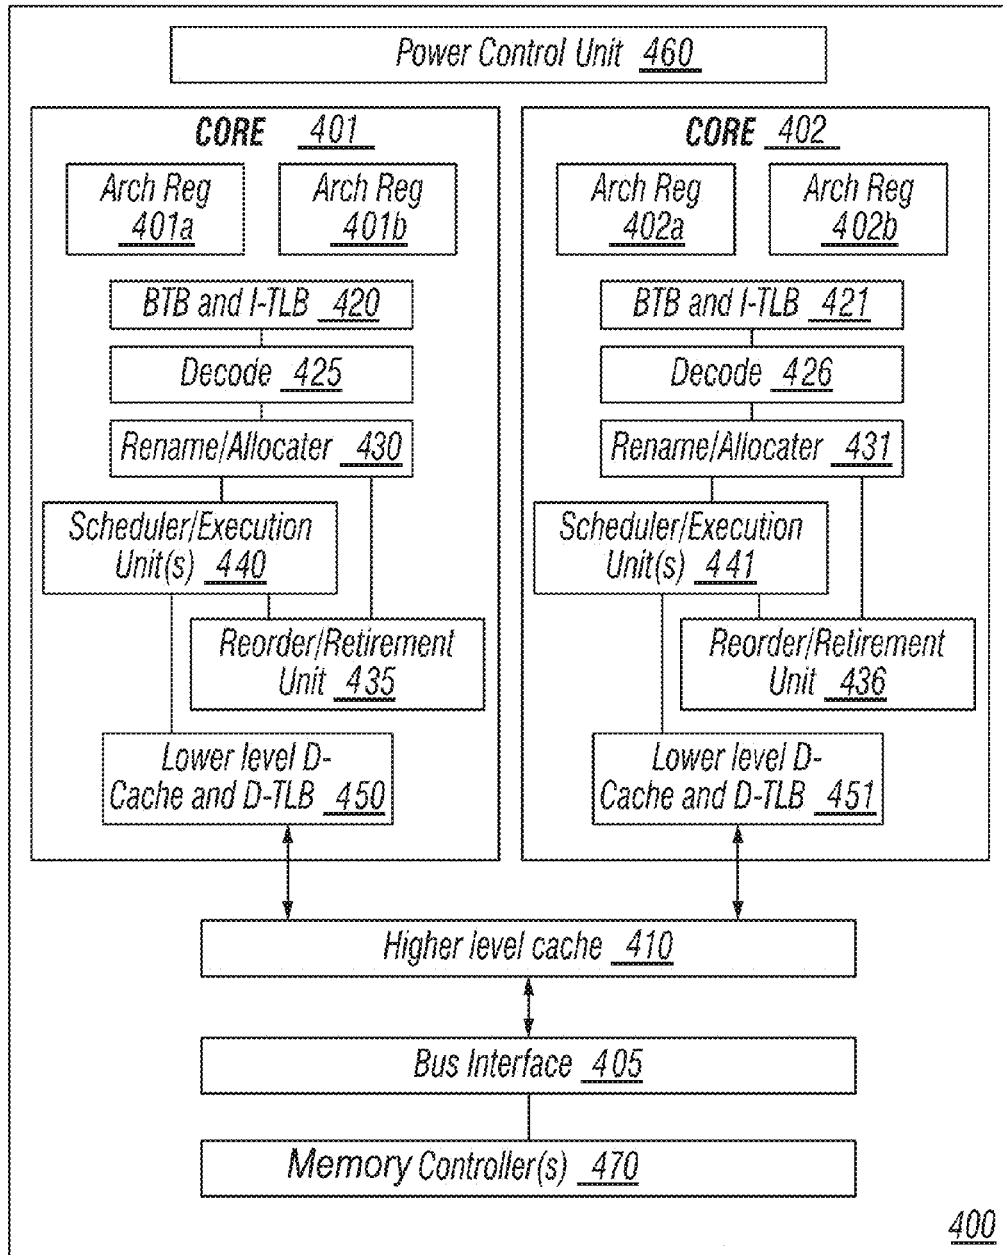


FIG. 4

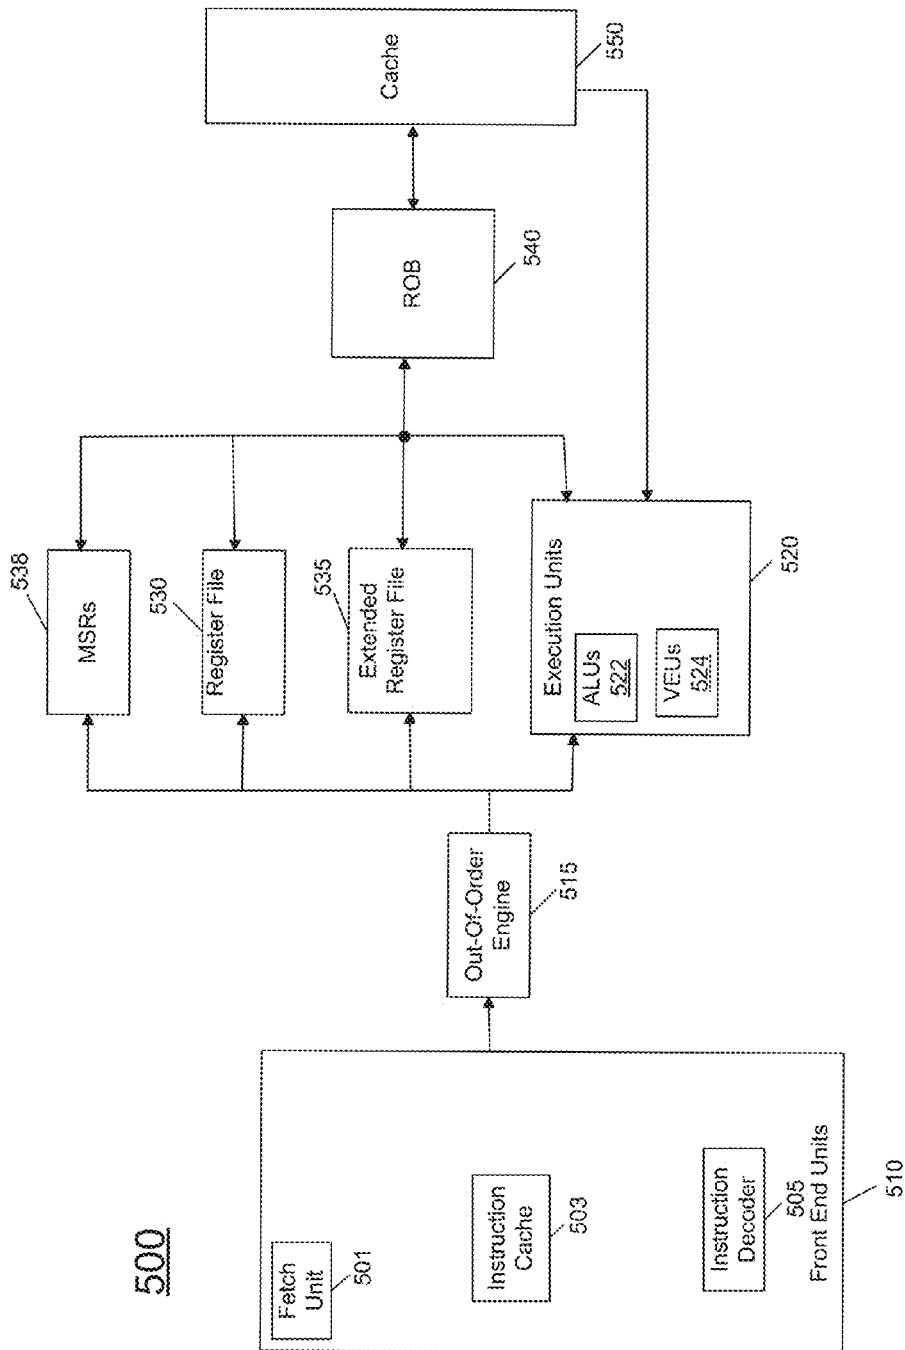


FIG. 5

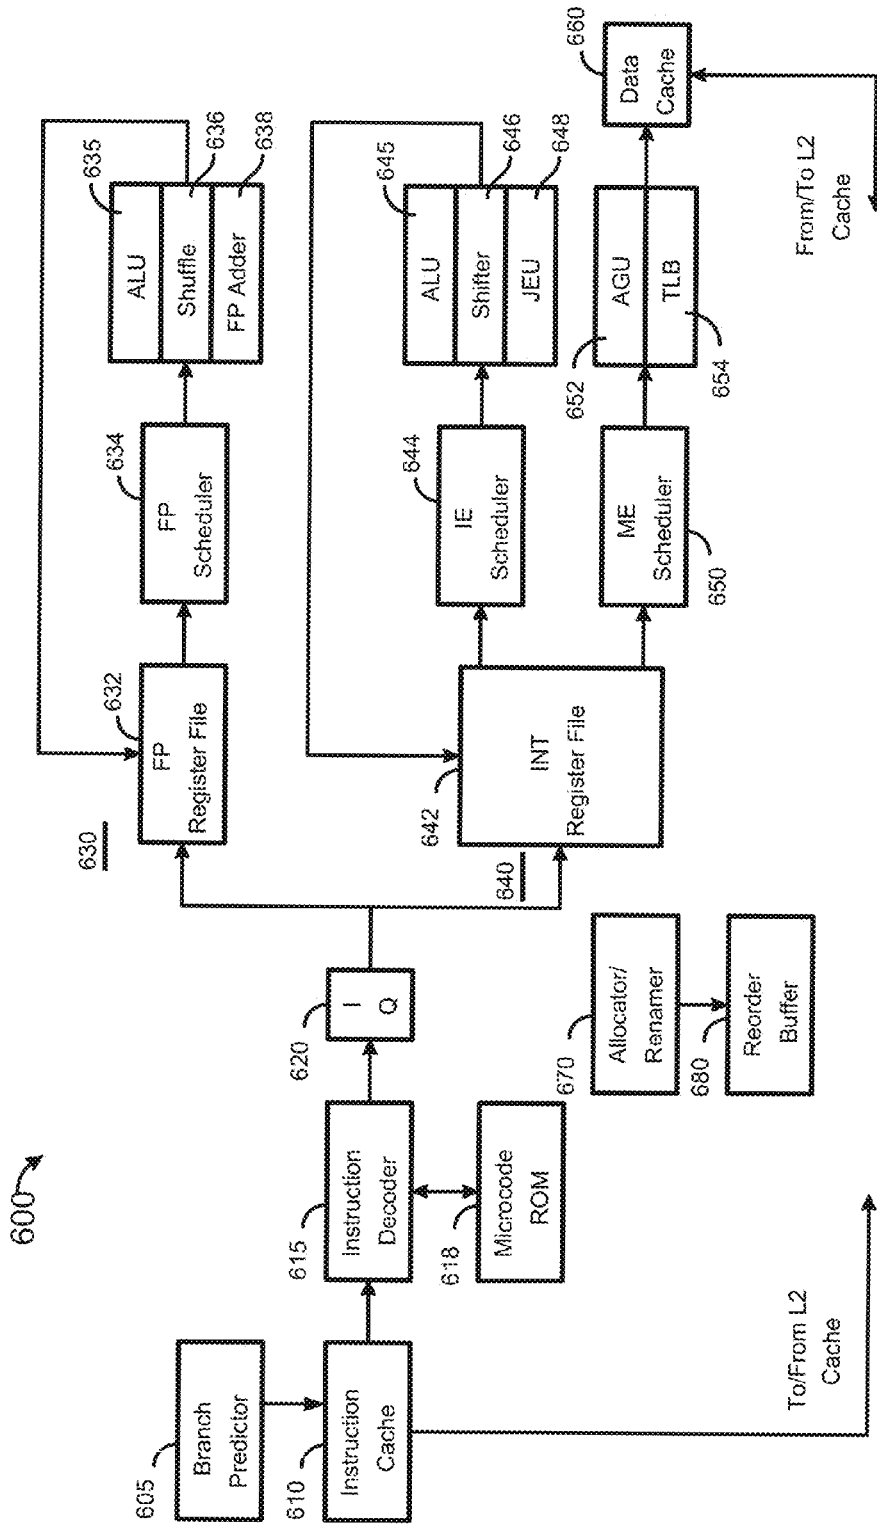


FIG. 6

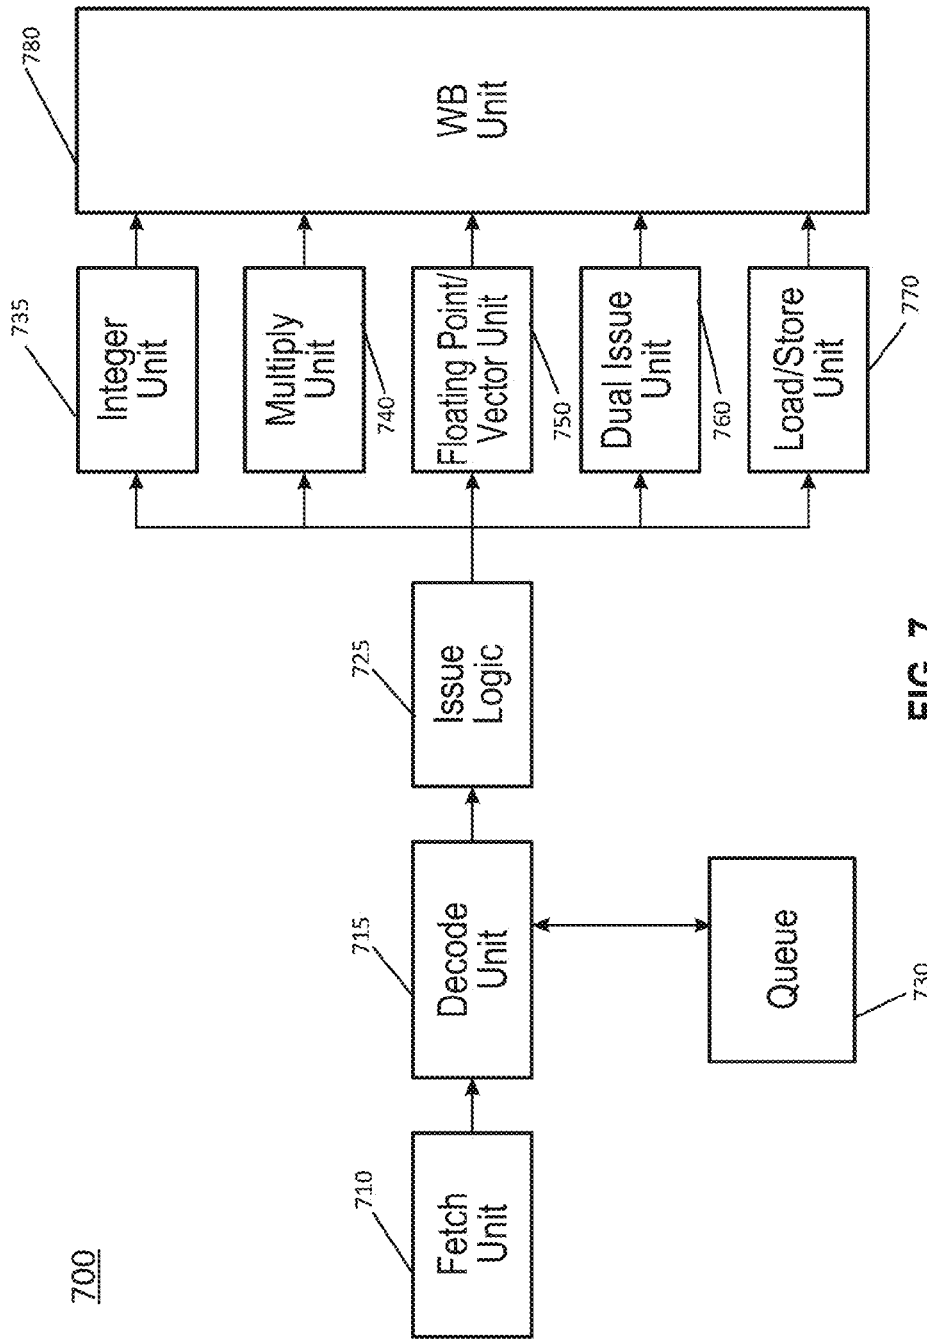


FIG. 7

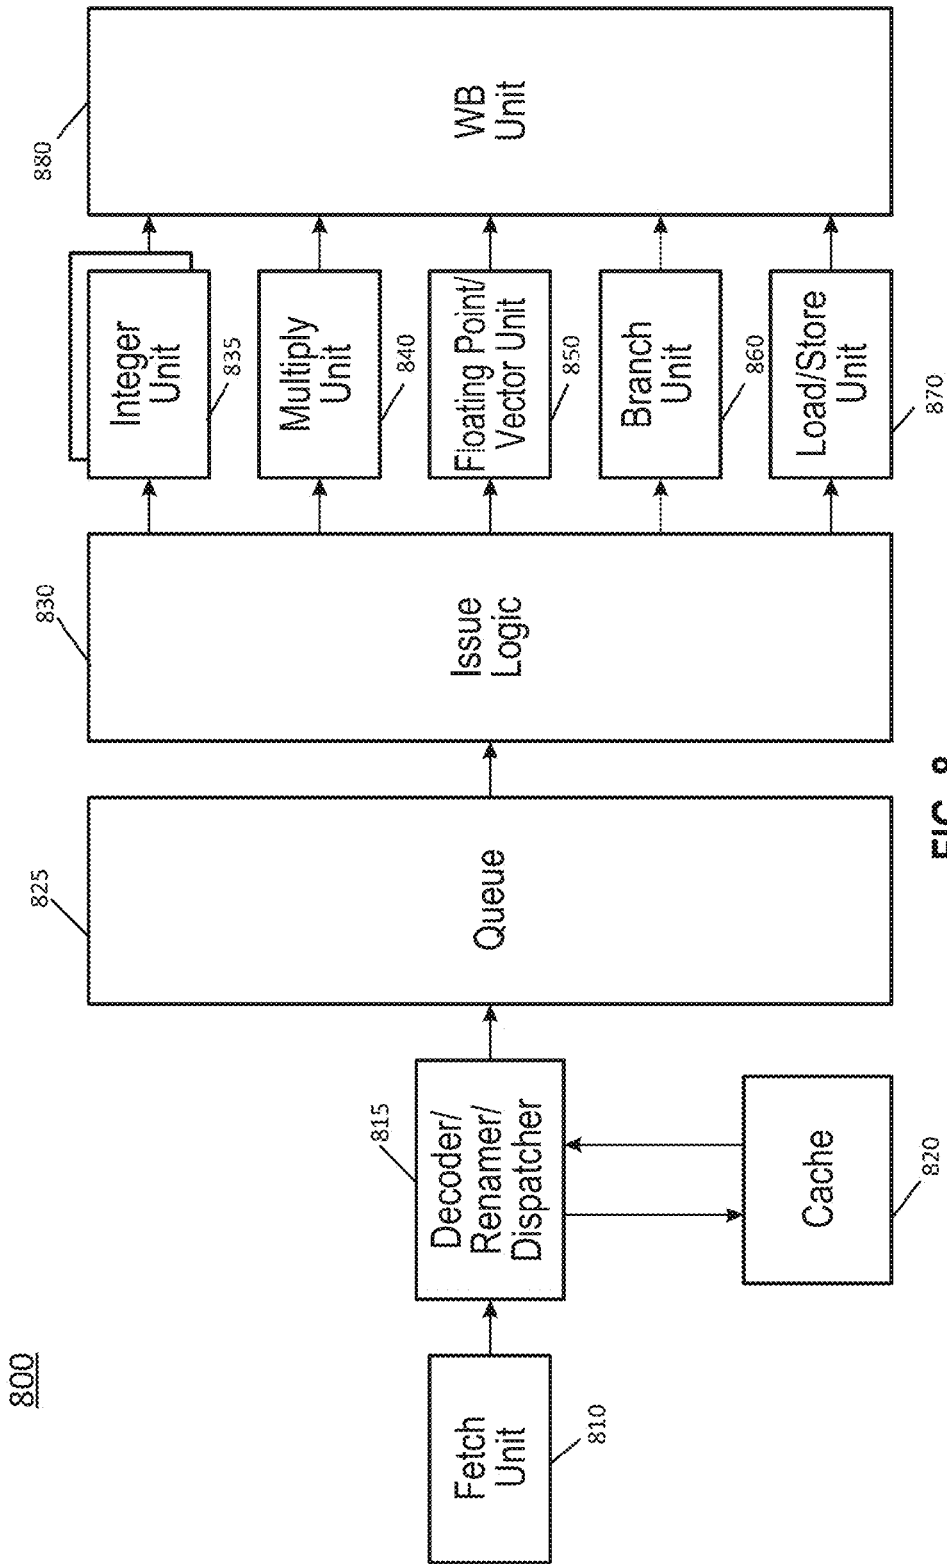


FIG. 8

800

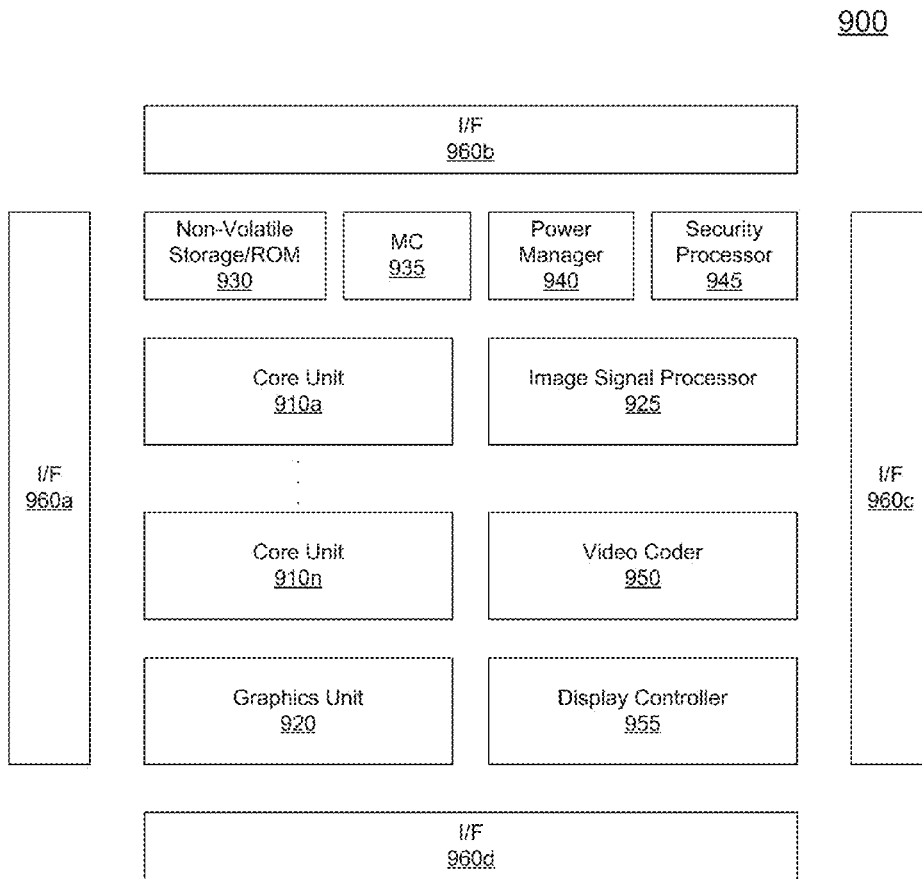


FIG. 9

1000

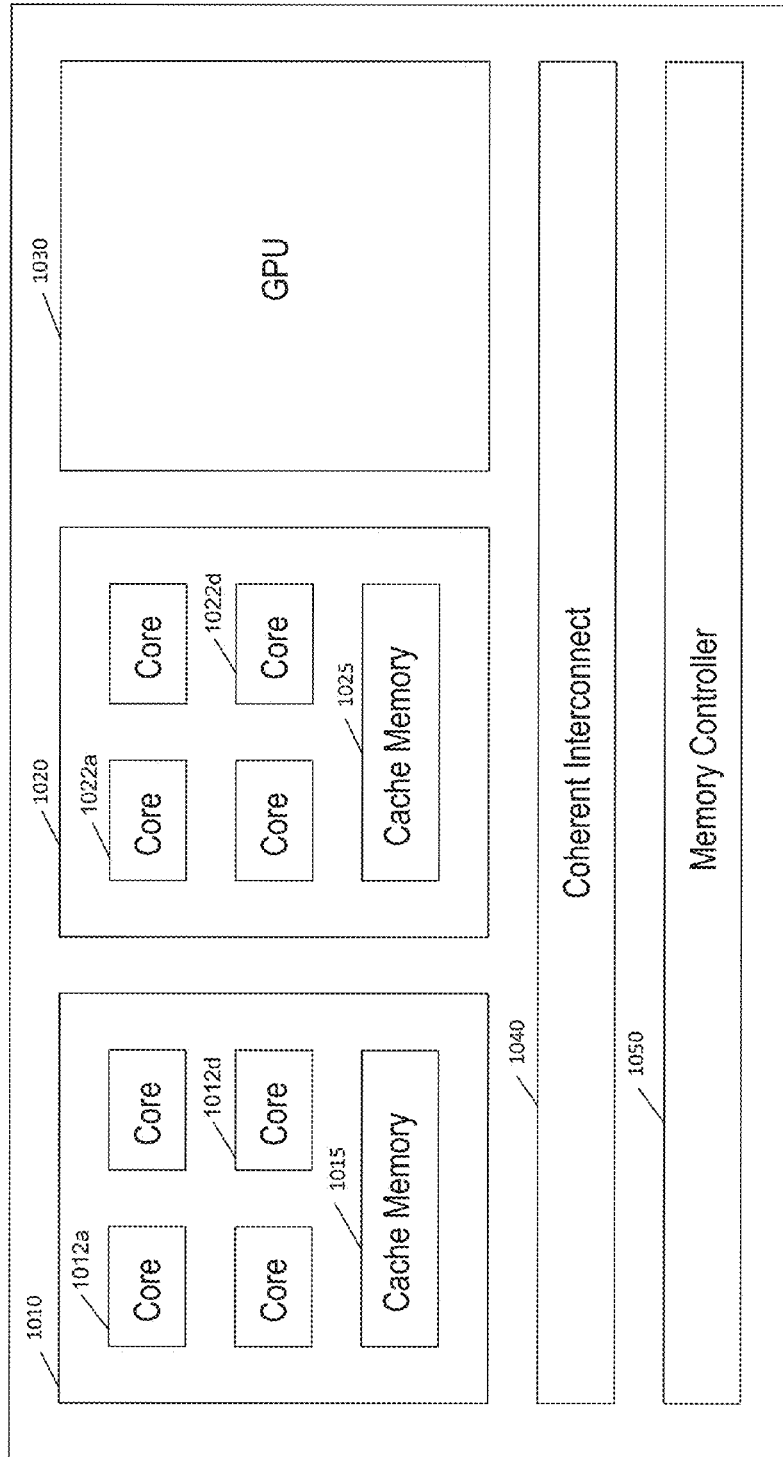


FIG. 10

1100

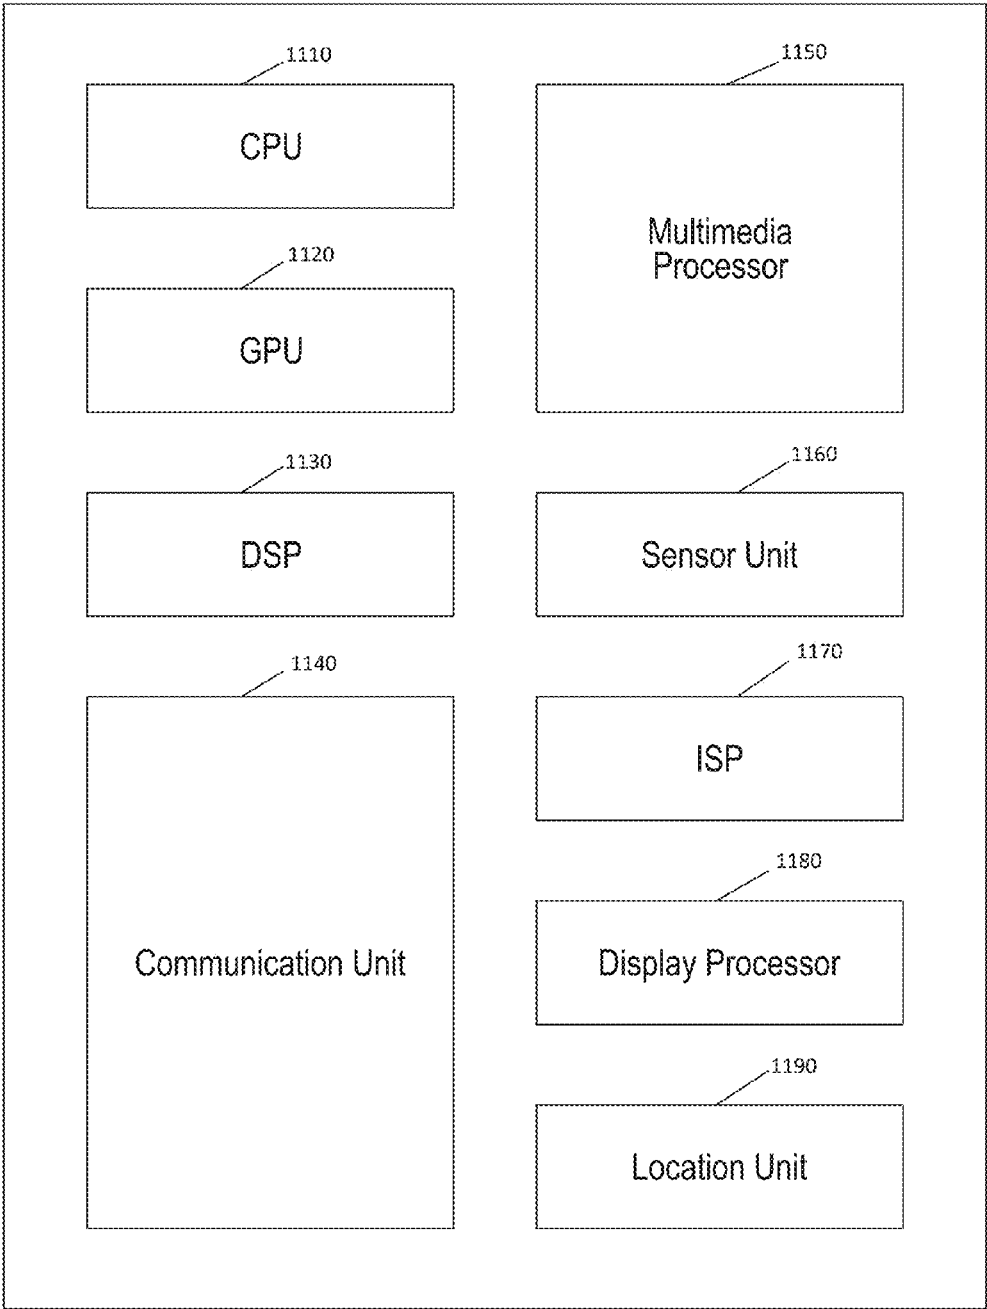


FIG. 11

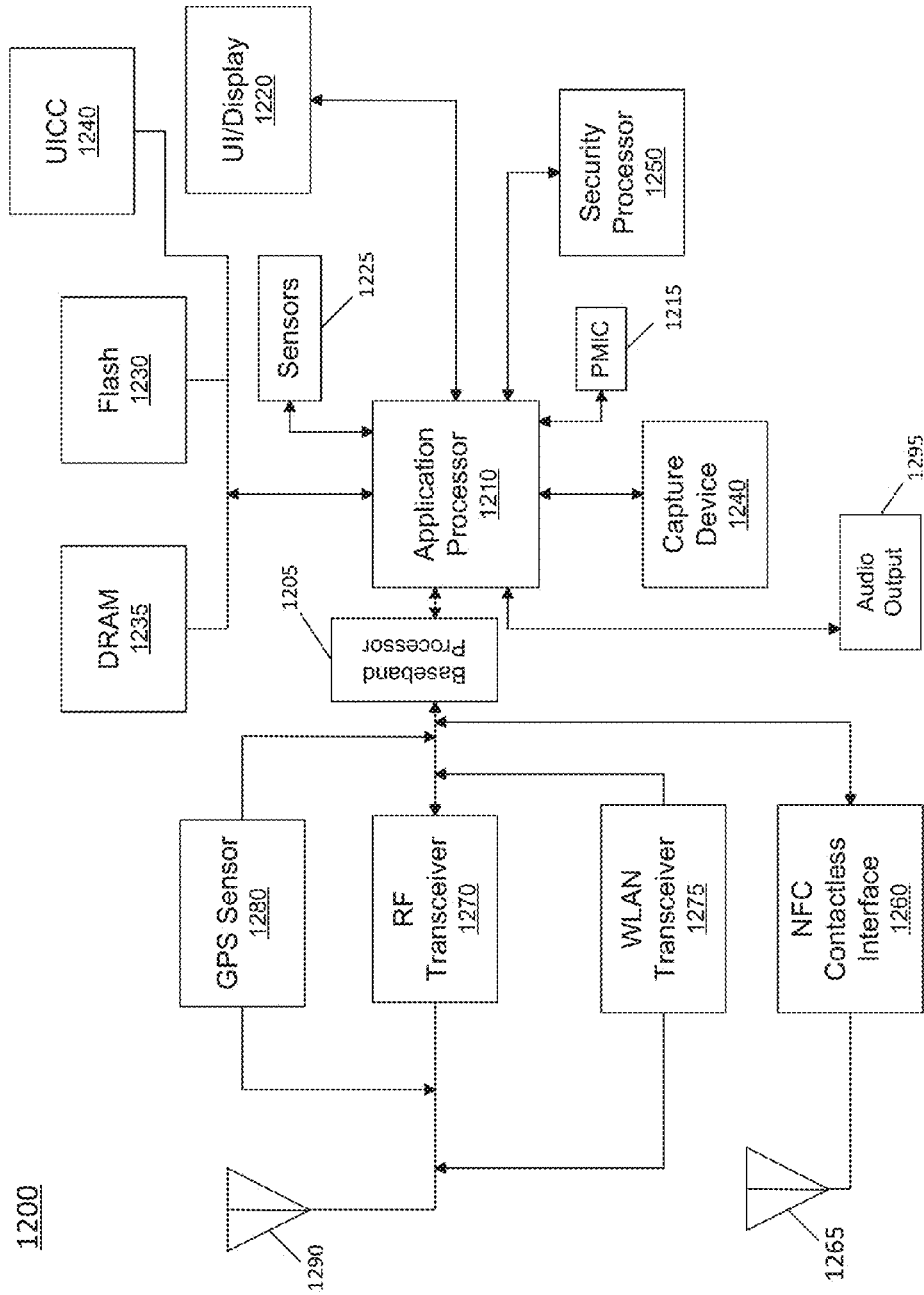


FIG. 12

1300

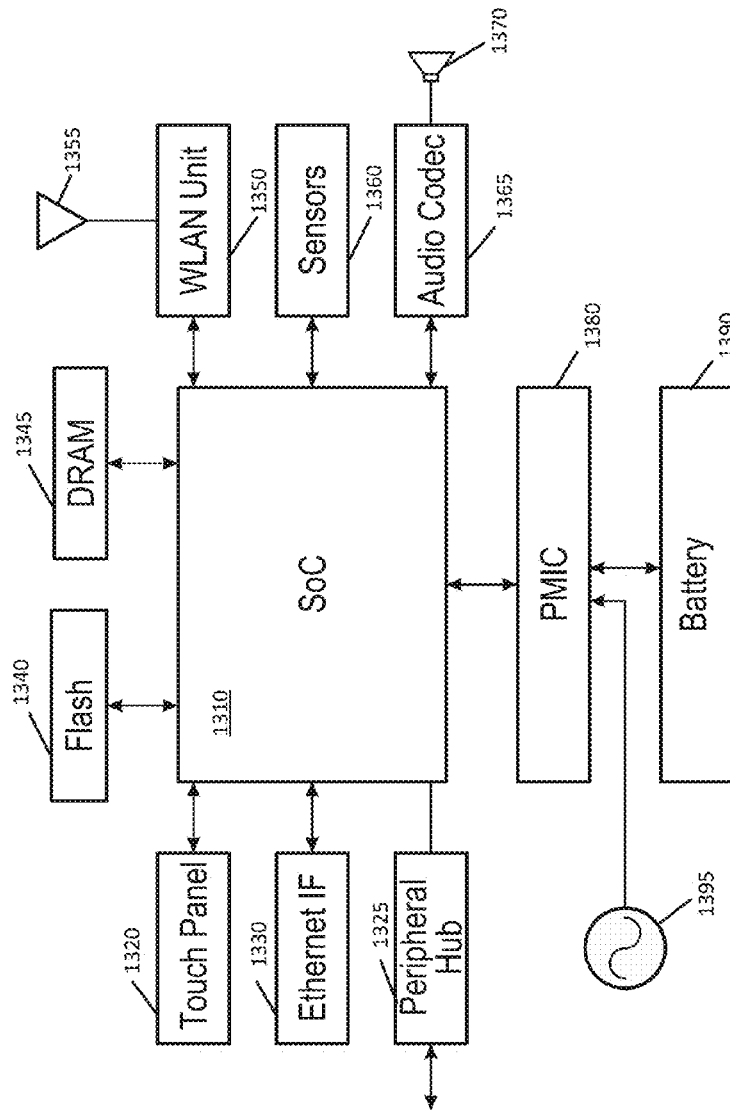


FIG. 13

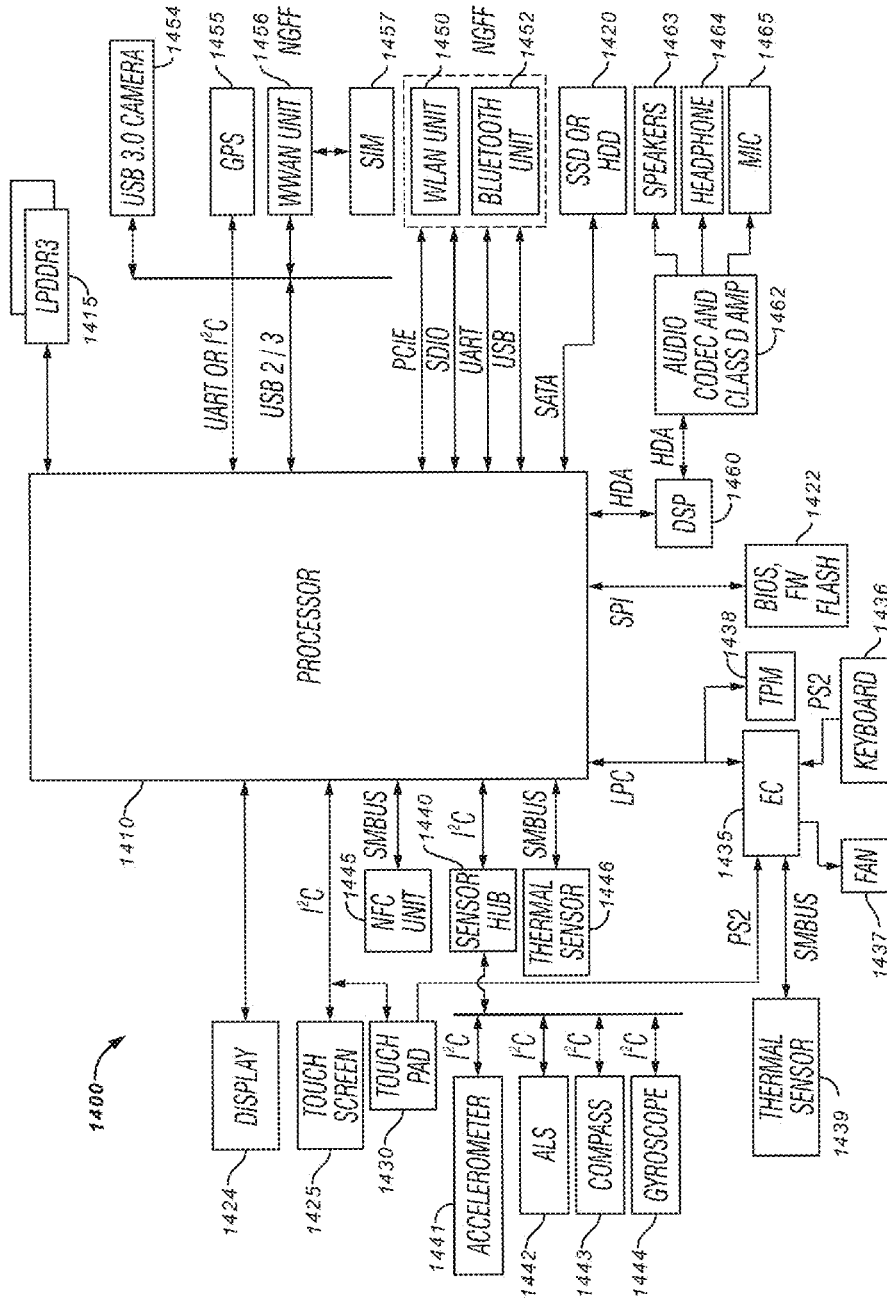


FIG. 14

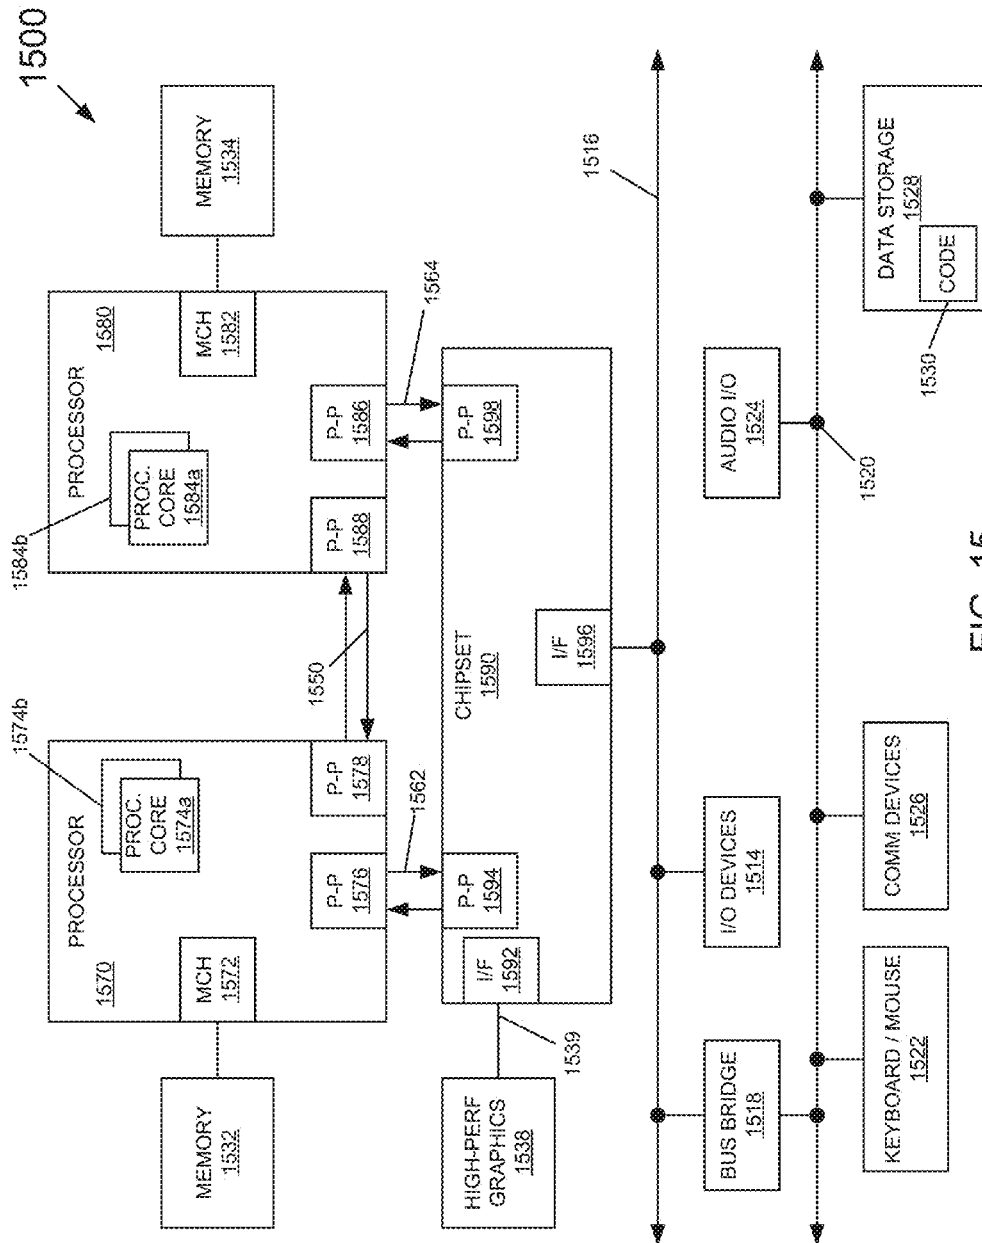


FIG. 15

1600

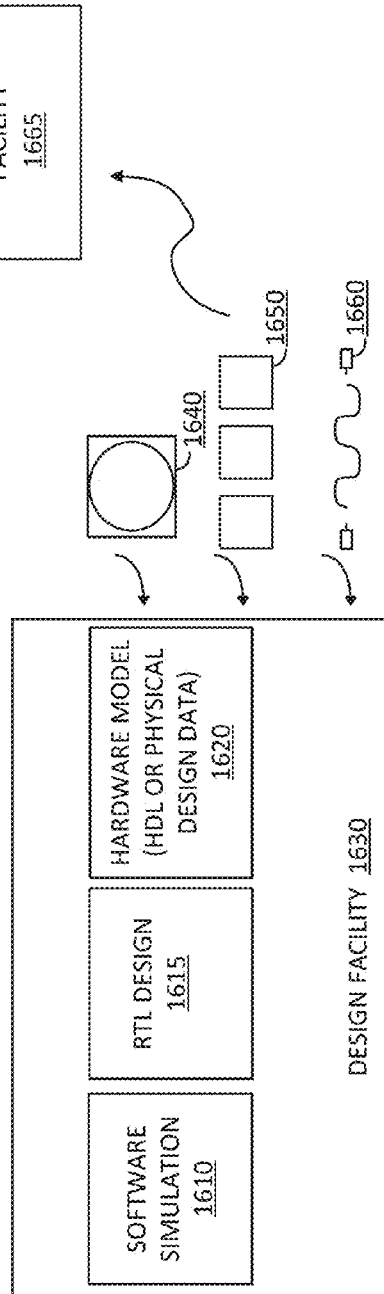


FIG. 16

1700

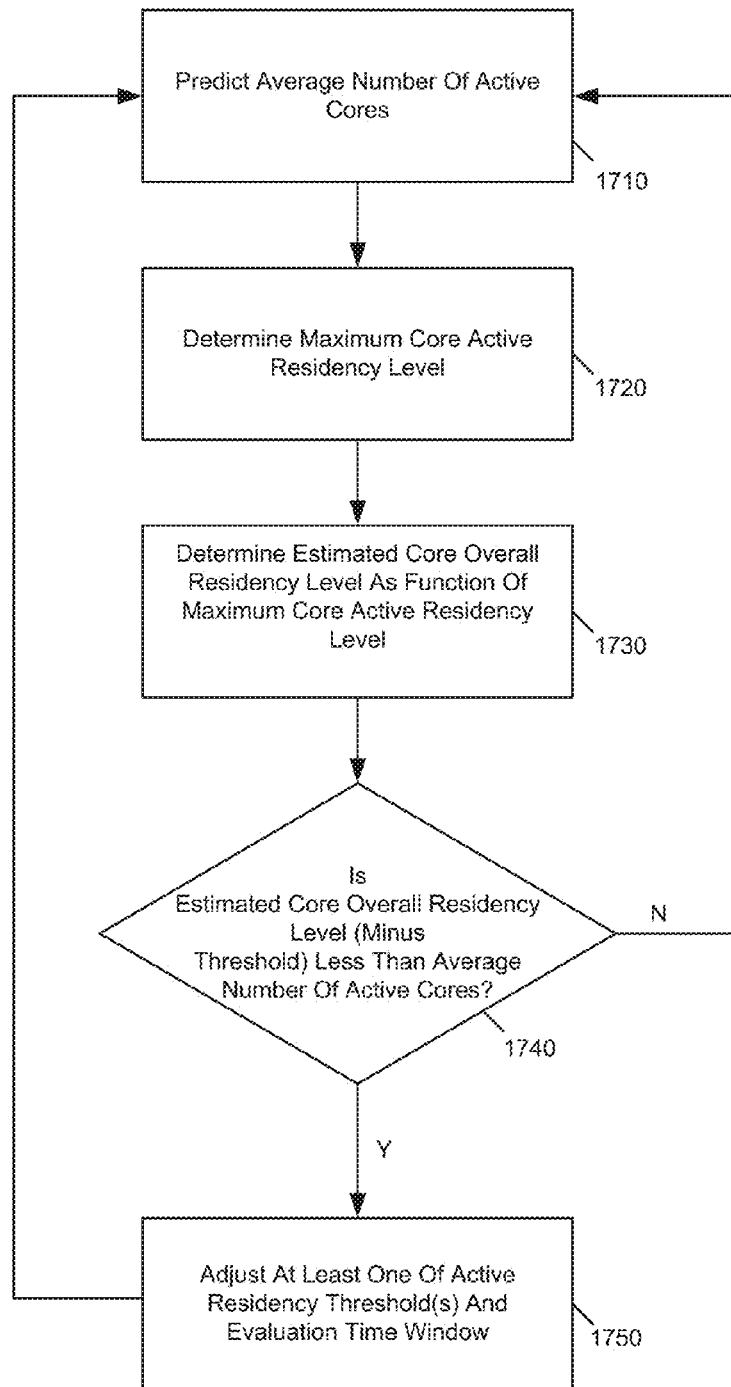


FIG. 17

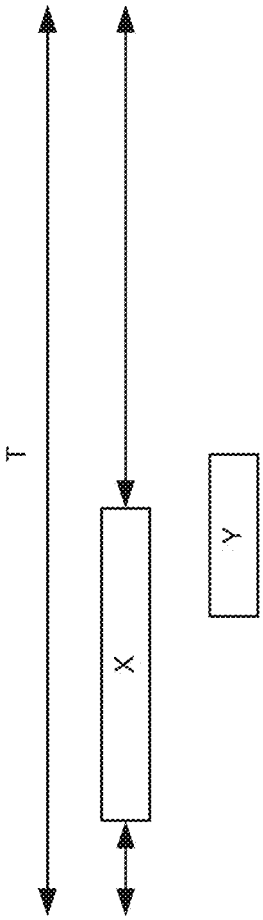


FIG. 18

1800

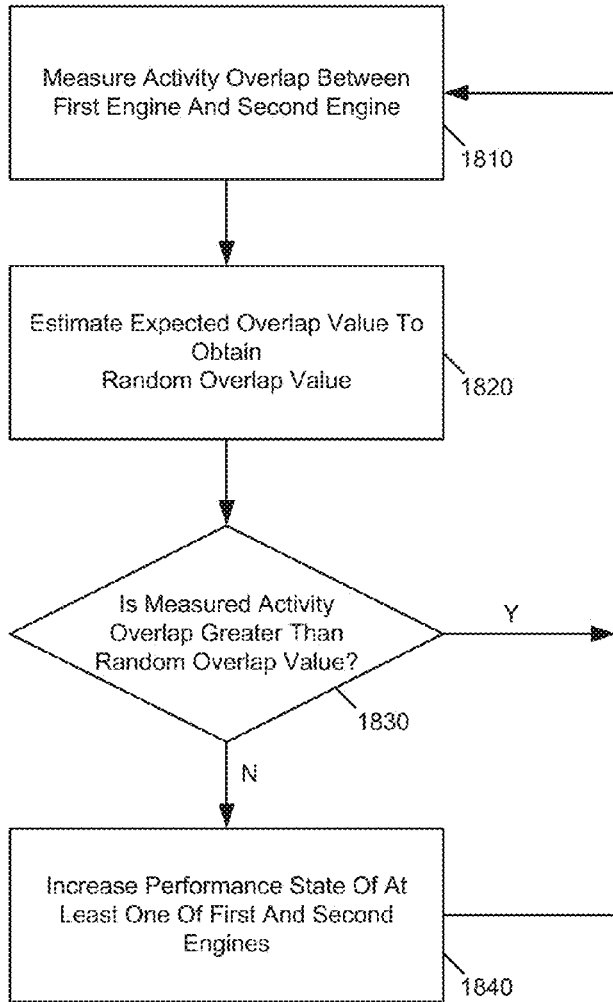


FIG. 19

1900

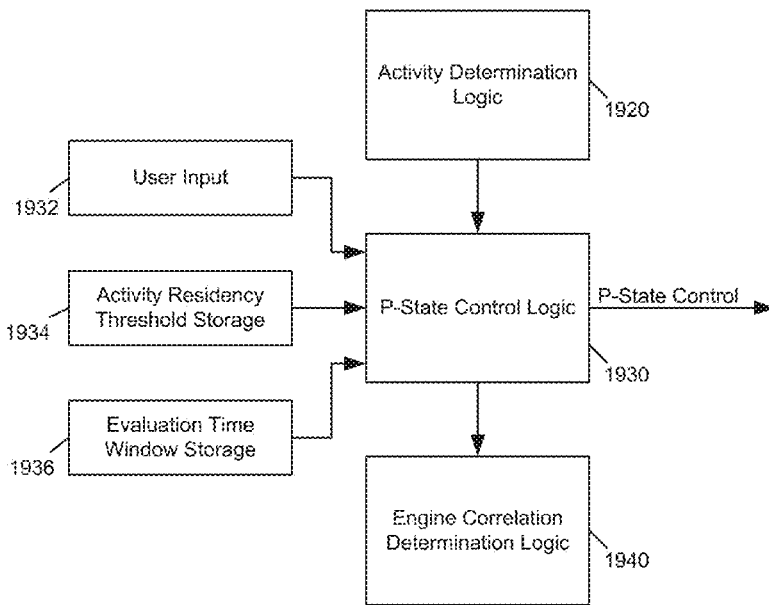
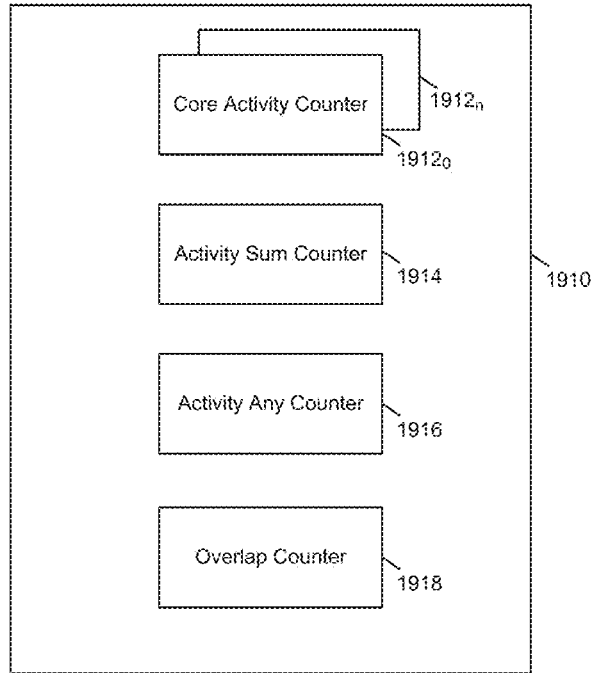


FIG. 20

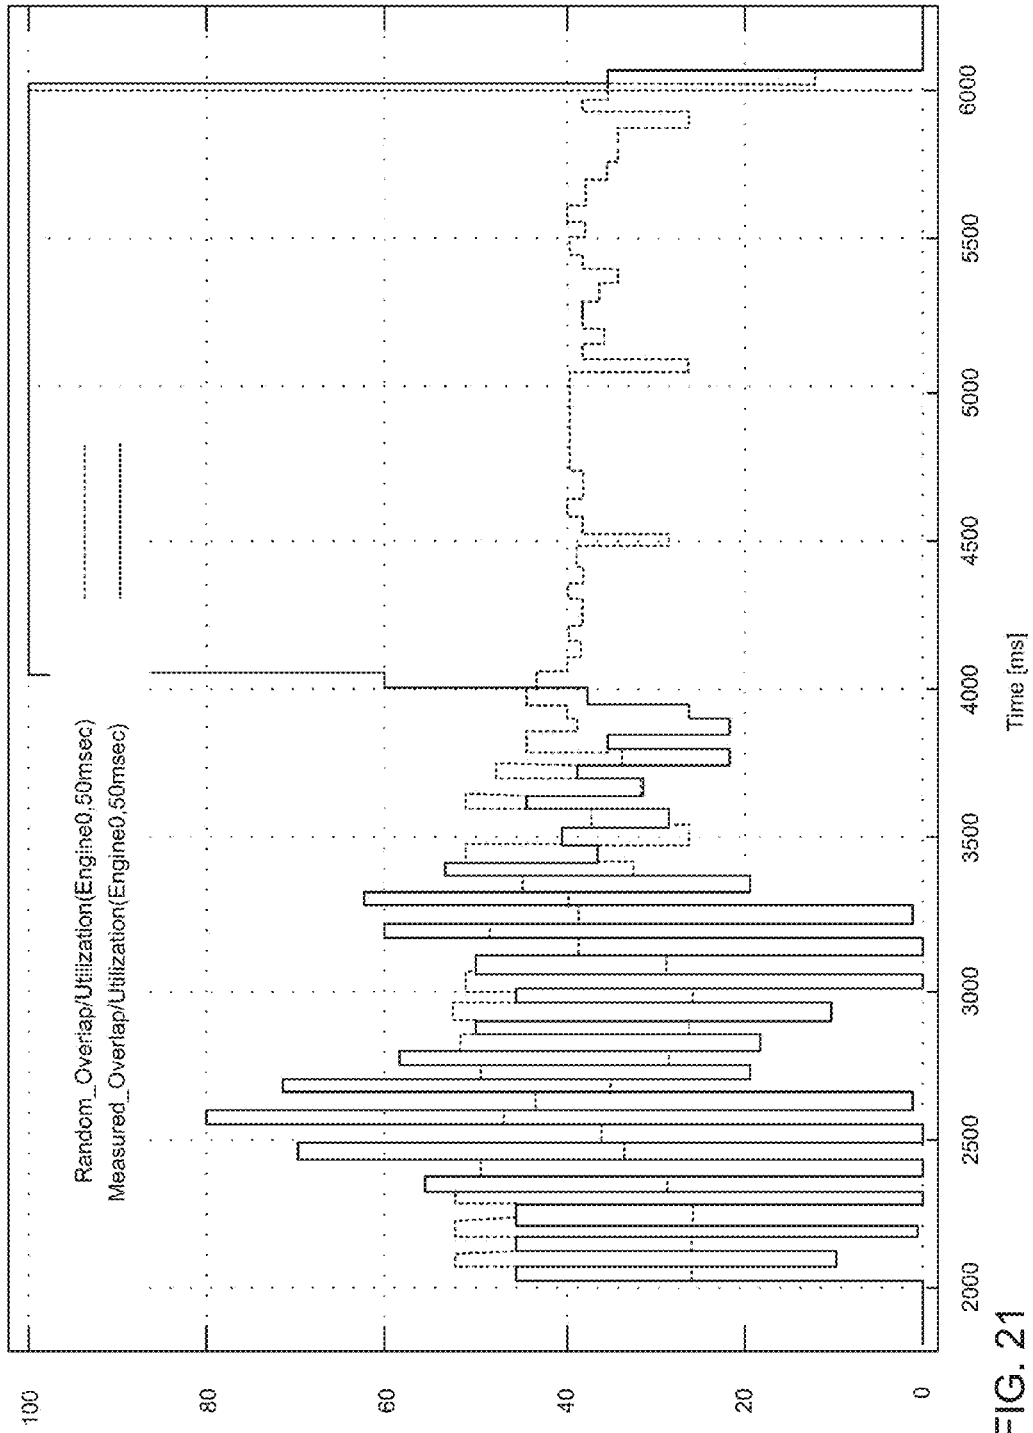


FIG. 21

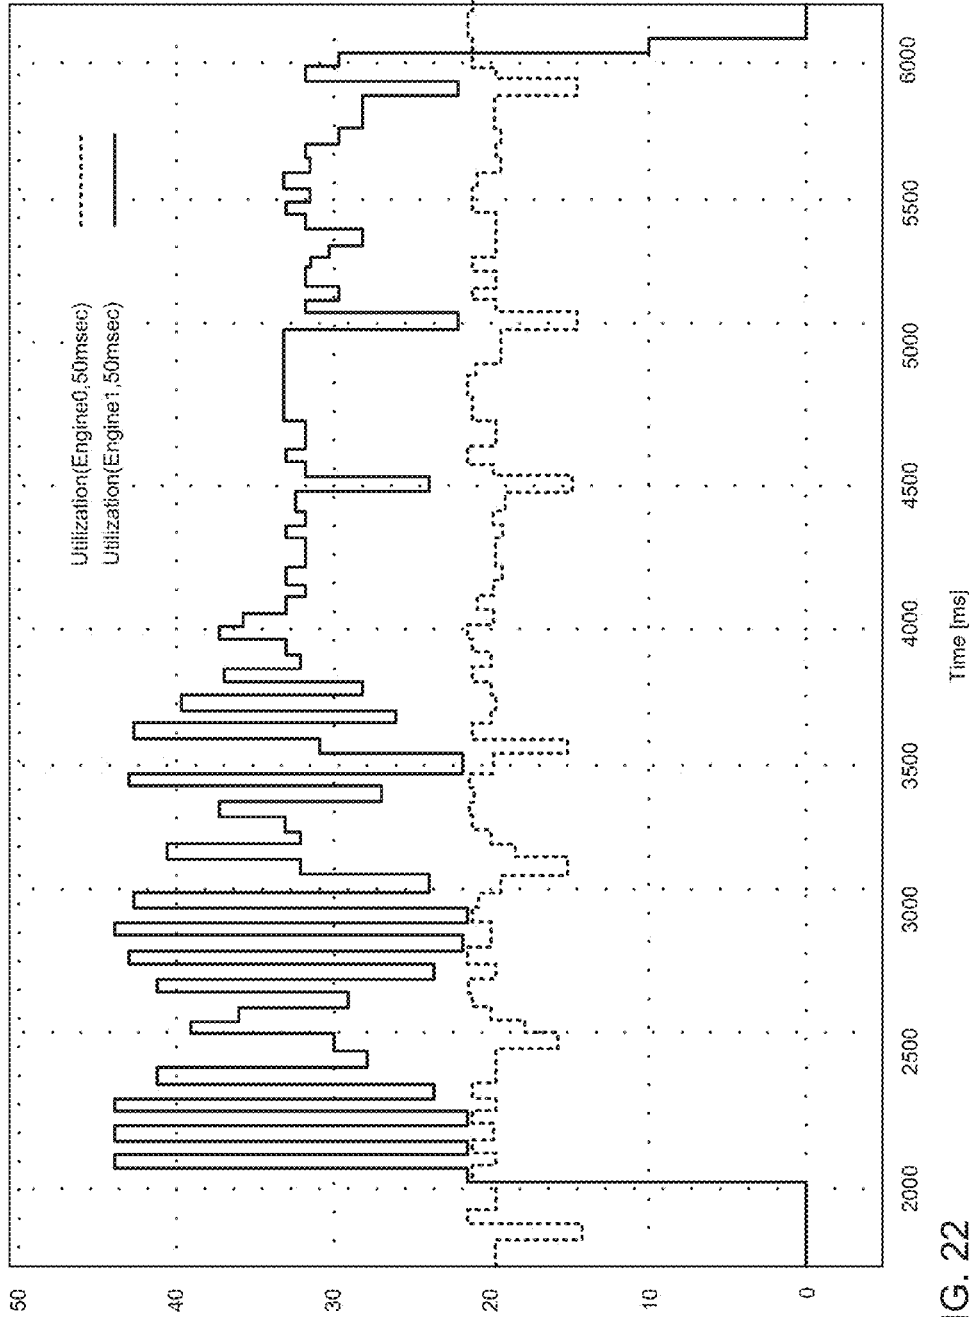


FIG. 22

1

CONTROLLING PERFORMANCE STATES OF PROCESSING ENGINES OF A PROCESSOR

TECHNICAL FIELD

Embodiments relate to power management of a system, and more particularly to power management of a multicore processor.

BACKGROUND

Advances in semiconductor processing and logic design have permitted an increase in the amount of logic that may be present on integrated circuit devices. As a result, computer system configurations have evolved from a single or multiple integrated circuits in a system to multiple hardware threads, multiple cores, multiple devices, and/or complete systems on individual integrated circuits. Additionally, as the density of integrated circuits has grown, the power requirements for computing systems (from embedded systems to servers) have also escalated. Furthermore, software inefficiencies, and its requirements of hardware, have also caused an increase in computing device energy consumption. In fact, some studies indicate that computing devices consume a sizeable percentage of the entire electricity supply for a country, such as the United States of America. As a result, there is a vital need for energy efficiency and conservation associated with integrated circuits. These needs will increase as servers, desktop computers, notebooks, Ultrabooks™, tablets, mobile phones, processors, embedded systems, etc. become even more prevalent (from inclusion in the typical computer, automobiles, and televisions to biotechnology).

Some modern processors provide for autonomous power management; however, power conservation opportunities are often lost as a result of a desire to increase performance as soon as a workload begins.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of a portion of a system in accordance with an embodiment of the present invention.

FIG. 2 is a block diagram of a processor in accordance with an embodiment of the present invention.

FIG. 3 is a block diagram of a multi-domain processor in accordance with another embodiment of the present invention.

FIG. 4 is an embodiment of a processor including multiple cores.

FIG. 5 is a block diagram of a micro-architecture of a processor core in accordance with one embodiment of the present invention.

FIG. 6 is a block diagram of a micro-architecture of a processor core in accordance with another embodiment.

FIG. 7 is a block diagram of a micro-architecture of a processor core in accordance with yet another embodiment.

FIG. 8 is a block diagram of a micro-architecture of a processor core in accordance with a still further embodiment.

FIG. 9 is a block diagram of a processor in accordance with another embodiment of the present invention.

FIG. 10 is a block diagram of a representative SoC in accordance with an embodiment of the present invention.

FIG. 11 is a block diagram of another example SoC in accordance with an embodiment of the present invention.

2

FIG. 12 is a block diagram of an example system with which embodiments can be used.

FIG. 13 is a block diagram of another example system with which embodiments may be used.

5 FIG. 14 is a block diagram of a representative computer system.

FIG. 15 is a block diagram of a system in accordance with an embodiment of the present invention.

10 FIG. 16 is a block diagram illustrating an IP core development system used to manufacture an integrated circuit to perform operations according to an embodiment.

FIG. 17 is a flow diagram of a method in accordance with an embodiment of the present invention.

15 FIG. 18 is a block diagram illustrating activity of two processing engines X and Y during a time interval T in accordance with an embodiment of the present invention.

FIG. 19 is a flow diagram of a method for controlling performance state within a processor in accordance with another embodiment of the present invention.

20 FIG. 20 is a block diagram of a portion of a processor in accordance with an embodiment of the present invention.

FIG. 21 is a timing illustration showing utilization of two processing engines.

25 FIG. 22 is a timing illustration showing a measured overlap and a calculated random overlap of two engines.

DETAILED DESCRIPTION

In various embodiments, power control logic of a processor may be adapted to analyze activity levels of one or more cores and/or other processing engines of the processor to determine the concurrency of processing on such cores or other engines. Based at least in part on such information, the power control logic can determine an appropriate performance state at which one or more of the engines may operate to provide for energy efficient operation levels. For example, when multiple processing engines are executing a parallel workload, a performance state of one or more engines may be less aggressively increased.

40 Embodiments may further analyze overlap of activity states of cores and/or other processing engines. Based on such analysis, a performance state of one or more such cores and/or engines can be increased, if such increase leads to a situation in which a pending workload being executed by such engines is faster completed to allow the processor to enter into a particular, e.g., deep low power state, such as in a race to halt scenario. In contrast, if execution of workloads on different cores and/or other processing engines is correlated, such performance state increases may be prevented, as it is likely that even when an execution of a pending workload is completed on one such engine, execution of a correlated or parallel workload on another engine would prevent entry of the processor into such deep low power state.

Understand also that while as described herein control methods and analysis are based on two cores or other processing engines (or two independent domains, each including one or more processing engines), the scope of the present invention is not limited in this regard and embodiments apply equally to many-core and/or many-domain processors including more than two cores, other processing engines, and/or domains to be analyzed and controlled as described herein.

65 Although the following embodiments are described with reference to energy conservation and energy efficiency in specific integrated circuits, such as in computing platforms or processors, other embodiments are applicable to other

types of integrated circuits and logic devices. Similar techniques and teachings of embodiments described herein may be applied to other types of circuits or semiconductor devices that may also benefit from better energy efficiency and energy conservation. For example, the disclosed embodiments are not limited to any particular type of computer systems. That is, disclosed embodiments can be used in many different system types, ranging from server computers (e.g., tower, rack, blade, micro-server and so forth), communications systems, storage systems, desktop computers of any configuration, laptop, notebook, and tablet computers (including 2:1 tablets, phablets and so forth), and may be also used in other devices, such as handheld devices, systems on chip (SoCs), and embedded applications. Some examples of handheld devices include cellular phones such as smartphones, Internet protocol devices, digital cameras, personal digital assistants (PDAs), and handheld PCs. Embedded applications may typically include a microcontroller, a digital signal processor (DSP), network computers (NetPC), set-top boxes, network hubs, wide area network (WAN) switches, wearable devices, or any other system that can perform the functions and operations taught below. More so, embodiments may be implemented in mobile terminals having standard voice functionality such as mobile phones, smartphones and phablets, and/or in non-mobile terminals without a standard wireless voice function communication capability, such as many wearables, tablets, notebooks, desktops, micro-servers, servers and so forth. Moreover, the apparatuses, methods, and systems described herein are not limited to physical computing devices, but may also relate to software optimizations for energy conservation and efficiency. As will become readily apparent in the description below, the embodiments of methods, apparatuses, and systems described herein (whether in reference to hardware, firmware, software, or a combination thereof) are vital to a 'green technology' future, such as for power conservation and energy efficiency in products that encompass a large portion of the US economy.

Referring now to FIG. 1, shown is a block diagram of a portion of a system in accordance with an embodiment of the present invention. As shown in FIG. 1, system 100 may include various components, including a processor 110 which as shown is a multicore processor. Processor 110 may be coupled to a power supply 150 via an external voltage regulator 160, which may perform a first voltage conversion to provide a primary regulated voltage Vreg to processor 110.

As seen, processor 110 may be a single die processor including multiple cores 120a-120n. In addition, each core may be associated with an integrated voltage regulator (IVR) 125a-125n which receives the primary regulated voltage and generates an operating voltage to be provided to one or more agents of the processor associated with the IVR. Accordingly, an IVR implementation may be provided to allow for fine-grained control of voltage and thus power and performance of each individual core. As such, each core can operate at an independent voltage and frequency, enabling great flexibility and affording wide opportunities for balancing power consumption with performance. In some embodiments, the use of multiple IVRs enables the grouping of components into separate power planes, such that power is regulated and supplied by the IVR to only those components in the group. During power management, a given power plane of one IVR may be powered down or off when the processor is placed into a certain low power state, while another power plane of another IVR remains active, or fully powered. Similarly, cores 120 may include or be associated

with independent clock generation circuitry such as one or more phase lock loops (PLLs) to control operating frequency of each core 120 independently.

Still referring to FIG. 1, additional components may be present within the processor including an input/output interface (IF) 132, another interface 134, and an integrated memory controller (IMC) 136. As seen, each of these components may be powered by another integrated voltage regulator 125. In one embodiment, interface 132 may enable operation for an Intel® Quick Path Interconnect (QPI) interconnect, which provides for point-to-point (PtP) links in a cache coherent protocol that includes multiple layers including a physical layer, a link layer and a protocol layer. In turn, interface 134 may communicate via a Peripheral Component Interconnect Express (PCIe™) protocol.

Also shown is a power control unit (PCU) 138, which may include hardware, software and/or firmware to perform power management operations with regard to processor 110. As seen, PCU 138 provides control information to external voltage regulator 160 via a digital interface 162 to cause the voltage regulator to generate the appropriate regulated voltage. PCU 138 also provides control information to IVRs 125 via another digital interface 163 to control the operating voltage generated (or to cause a corresponding IVR to be disabled in a low power mode). In various embodiments, PCU 138 may include a variety of power management logic units to perform hardware-based power management. Such power management may be wholly processor controlled (e.g., by various processor hardware, and which may be triggered by workload and/or power, thermal or other processor constraints) and/or the power management may be performed responsive to external sources (such as a platform or power management source or system software). As described herein, in embodiments, PCU 138 may be configured to perform power management based at least in part on concurrency and/or overlap of execution of different processing engines of processor 110.

While not shown for ease of illustration, understand that additional components may be present within processor 110 such as uncore logic, and other components such as internal memories, e.g., one or more levels of a cache memory hierarchy and so forth. Furthermore, while shown in the implementation of FIG. 1 with an integrated voltage regulator, embodiments are not so limited. For example, other regulated voltages may be provided to on-chip resources from external voltage regulator 160 or one or more additional external sources of regulated voltages.

Note that the power management techniques described herein may be independent of and complementary to an operating system (OS)-based power management (OSPM) mechanism. According to one example OSPM technique, a processor can operate at various performance states or levels, so-called P-states, namely from P0 to PN, each associated with an operating frequency (and operating voltage). In general, the P1 performance state may correspond to the highest guaranteed performance state that can be requested by an OS. In addition to this P1 state, the OS can further request a higher performance state, namely a P0 state. This P0 state may thus be an opportunistic or turbo mode state in which, when power and/or thermal budget is available, processor hardware can configure the processor or at least portions thereof to operate at a higher than guaranteed frequency. In many implementations a processor can include multiple so-called bin frequencies above the P1 guaranteed maximum frequency, exceeding to a maximum peak frequency of the particular processor, as fused or otherwise written into the processor during manufacture. In

5

addition, according to one OSPM mechanism, a processor can operate at various power states or levels. With regard to power states, an OSPM mechanism may specify different power consumption states, generally referred to as C-states, C0, C1 to Cn states. When a core is active, it runs at a C0 state, and when the core is idle it may be placed in a core low power state, also called a core non-zero C-state (e.g., C1-C6 states), with each C-state being at a lower power consumption level (such that C6 is a deeper low power state than C1, and so forth).

Understand that many different types of power management techniques may be used individually or in combination in different embodiments. As representative examples, a power controller may control the processor to be power managed by some form of dynamic voltage frequency scaling (DVFS) in which an operating voltage and/or operating frequency of one or more cores or other processor logic may be dynamically controlled to reduce power consumption in certain situations. In an example, DVFS may be performed using Enhanced Intel SpeedStep™ technology available from Intel Corporation, Santa Clara, Calif., to provide optimal performance at a lowest power consumption level. In another example, DVFS may be performed using Intel TurboBoost™ technology to enable one or more cores or other compute engines to operate at a higher than guaranteed operating frequency based on conditions (e.g., workload and availability).

Another power management technique that may be used in certain examples is dynamic swapping of workloads between different compute engines. For example, the processor may include asymmetric cores or other processing engines that operate at different power consumption levels, such that in a power constrained situation, one or more workloads can be dynamically switched to execute on a lower power core or other compute engine. Another exemplary power management technique is hardware duty cycling (HDC), which may cause cores and/or other compute engines to be periodically enabled and disabled according to a duty cycle, such that one or more cores may be made inactive during an inactive period of the duty cycle and made active during an active period of the duty cycle.

Power management techniques also may be used when constraints exist in an operating environment. For example, when a power and/or thermal constraint is encountered, power may be reduced by reducing operating frequency and/or voltage. Other power management techniques include throttling instruction execution rate or limiting scheduling of instructions. Still further, it is possible for instructions of a given instruction set architecture to include express or implicit direction as to power management operations. Although described with these particular examples, understand that many other power management techniques may be used in particular embodiments.

Embodiments can be implemented in processors for various markets including server processors, desktop processors, mobile processors and so forth. Referring now to FIG. 2, shown is a block diagram of a processor in accordance with an embodiment of the present invention. As shown in FIG. 2, processor 200 may be a multicore processor including a plurality of cores 210_a-210_n. In one embodiment, each such core may be of an independent power domain and can be configured to enter and exit active states and/or maximum performance states based on workload. One or more cores 210 may be heterogeneous to the other cores, e.g., having different micro-architectures, instruction set architectures, pipeline depths, power and performance capabilities. The various cores may be coupled via an interconnect 215 to a

6

system agent or uncore 220 that includes various components. As seen, the uncore 220 may include a shared cache 230 which may be a last level cache. In addition, the uncore may include an integrated memory controller 240 to communicate with a system memory (not shown in FIG. 2), e.g., via a memory bus. Uncore 220 also includes various interfaces 250 and a power control unit 255, which may include logic to perform the power management techniques described herein. To this end, power control unit 255 may include a P-state control logic 256 which may be adapted to control P-state of one or more cores 210 based at least in part on concurrency and/or overlap of workload execution on multiple cores. Such logic may further be adapted to perform P-state control based at least in part on whether a workload is user interactive (or not).

In addition, by interfaces 250a-250n, connection can be made to various off-chip components such as peripheral devices, mass storage and so forth. While shown with this particular implementation in the embodiment of FIG. 2, the scope of the present invention is not limited in this regard.

Referring now to FIG. 3, shown is a block diagram of a multi-domain processor in accordance with another embodiment of the present invention. As shown in the embodiment of FIG. 3, processor 300 includes multiple domains. Specifically, a core domain 310 can include a plurality of cores 310a-310n, a graphics domain 320 can include one or more graphics engines, and a system agent domain 350 may further be present. In some embodiments, system agent domain 350 may execute at an independent frequency than the core domain and may remain powered on at all times to handle power control events and power management such that domains 310 and 320 can be controlled to dynamically enter into and exit high power and low power states. Each of domains 310 and 320 may operate at different voltage and/or power. Note that while only shown with three domains, understand the scope of the present invention is not limited in this regard and additional domains can be present in other embodiments. For example, multiple core domains may be present each including at least one core.

In general, each core 310 may further include low level caches in addition to various execution units and additional processing elements. In turn, the various cores may be coupled to each other and to a shared cache memory formed of a plurality of units of a last level cache (LLC) 340a-340n. In various embodiments, LLC 340 may be shared amongst the cores and the graphics engine, as well as various media processing circuitry. As seen, a ring interconnect 330 thus couples the cores together, and provides interconnection between the cores, graphics domain 320 and system agent circuitry 350. In one embodiment, interconnect 330 can be part of the core domain. However in other embodiments the ring interconnect can be of its own domain.

As further seen, system agent domain 350 may include display controller 352 which may provide control of and an interface to an associated display. As further seen, system agent domain 350 may include a power control unit 355 which can include P-state control logic 356 to perform the power management techniques described herein, include P-state control based on concurrency, overlap, and/or workload characterization.

As further seen in FIG. 3, processor 300 can further include an integrated memory controller (IMC) 370 that can provide for an interface to a system memory, such as a dynamic random access memory (DRAM). Multiple interfaces 380a-380n may be present to enable interconnection between the processor and other circuitry. For example, in one embodiment at least one direct media interface (DMI)

interface may be provided as well as one or more PCIe™ interfaces. Still further, to provide for communications between other agents such as additional processors or other circuitry, one or more QPI interfaces may also be provided. Although shown at this high level in the embodiment of FIG. 3, understand the scope of the present invention is not limited in this regard.

Referring to FIG. 4, an embodiment of a processor including multiple cores is illustrated. Processor 400 includes any processor or processing device, such as a microprocessor, an embedded processor, a digital signal processor (DSP), a network processor, a handheld processor, an application processor, a co-processor, a system on a chip (SoC), or other device to execute code. Processor 400, in one embodiment, includes at least two cores—cores 401 and 402, which may include asymmetric cores or symmetric cores (the illustrated embodiment). However, processor 400 may include any number of processing elements that may be symmetric or asymmetric.

In one embodiment, a processing element refers to hardware or logic to support a software thread. Examples of hardware processing elements include: a thread unit, a thread slot, a thread, a process unit, a context, a context unit, a logical processor, a hardware thread, a core, and/or any other element, which is capable of holding a state for a processor, such as an execution state or architectural state. In other words, a processing element, in one embodiment, refers to any hardware capable of being independently associated with code, such as a software thread, operating system, application, or other code. A physical processor typically refers to an integrated circuit, which potentially includes any number of other processing elements, such as cores or hardware threads.

A core often refers to logic located on an integrated circuit capable of maintaining an independent architectural state, wherein each independently maintained architectural state is associated with at least some dedicated execution resources. In contrast to cores, a hardware thread typically refers to any logic located on an integrated circuit capable of maintaining an independent architectural state, wherein the independently maintained architectural states share access to execution resources. As can be seen, when certain resources are shared and others are dedicated to an architectural state, the line between the nomenclature of a hardware thread and core overlaps. Yet often, a core and a hardware thread are viewed by an operating system as individual logical processors, where the operating system is able to individually schedule operations on each logical processor.

Physical processor 400, as illustrated in FIG. 4, includes two cores, cores 401 and 402. Here, cores 401 and 402 are considered symmetric cores, i.e., cores with the same configurations, functional units, and/or logic. In another embodiment, core 401 includes an out-of-order processor core, while core 402 includes an in-order processor core. However, cores 401 and 402 may be individually selected from any type of core, such as a native core, a software managed core, a core adapted to execute a native instruction set architecture (ISA), a core adapted to execute a translated ISA, a co-designed core, or other known core. Yet to further the discussion, the functional units illustrated in core 401 are described in further detail below, as the units in core 402 operate in a similar manner.

As depicted, core 401 includes two hardware threads 401a and 401b, which may also be referred to as hardware thread slots 401a and 401b. Therefore, software entities, such as an operating system, in one embodiment potentially view processor 400 as four separate processors, i.e., four

logical processors or processing elements capable of executing four software threads concurrently. As alluded to above, a first thread is associated with architecture state registers 401a, a second thread is associated with architecture state registers 401b, a third thread may be associated with architecture state registers 402a, and a fourth thread may be associated with architecture state registers 402b. Here, each of the architecture state registers (401a, 401b, 402a, and 402b) may be referred to as processing elements, thread slots, or thread units, as described above. As illustrated, architecture state registers 401a are replicated in architecture state registers 401b, so individual architecture states/contexts are capable of being stored for logical processor 401a and logical processor 401b. In core 401, other smaller resources, such as instruction pointers and renaming logic in allocator and renamer block 430 may also be replicated for threads 401a and 401b. Some resources, such as re-order buffers in reorder/retirement unit 435, branch target buffer and instruction translation lookaside buffer (BTB and I-TLB) 420, load/store buffers, and queues may be shared through partitioning. Other resources, such as general purpose internal registers, page-table base register(s), low-level data-cache and data-TLB 450, execution unit(s) 440, and portions of out-of-order unit 435 are potentially fully shared.

Processor 400 often includes other resources, which may be fully shared, shared through partitioning, or dedicated by/to processing elements. In FIG. 4, an embodiment of a purely exemplary processor with illustrative logical units/resources of a processor is illustrated. Note that a processor may include, or omit, any of these functional units, as well as include any other known functional units, logic, or firmware not depicted. As illustrated, core 401 includes a simplified, representative out-of-order (OOO) processor core. But an in-order processor may be utilized in different embodiments. The OOO core includes a branch target buffer 420 to predict branches to be executed/taken and an instruction-translation buffer (I-TLB) 420 to store address translation entries for instructions.

Core 401 further includes decode module 425 coupled to a fetch unit to decode fetched elements. Fetch logic, in one embodiment, includes individual sequencers associated with thread slots 401a, 401b, respectively. Usually core 401 is associated with a first ISA, which defines/specifies instructions executable on processor 400. Often machine code instructions that are part of the first ISA include a portion of the instruction (referred to as an opcode), which references/specifies an instruction or operation to be performed. Decode logic 425 includes circuitry that recognizes these instructions from their opcodes and passes the decoded instructions on in the pipeline for processing as defined by the first ISA. For example, decoders 425, in one embodiment, include logic designed or adapted to recognize specific instructions, such as transactional instruction. As a result of the recognition by decoders 425, the architecture or core 401 takes specific, predefined actions to perform tasks associated with the appropriate instruction. It is important to note that any of the tasks, blocks, operations, and methods described herein may be performed in response to a single or multiple instructions; some of which may be new or old instructions.

In one example, allocator and renamer block 430 includes an allocator to reserve resources, such as register files to store instruction processing results. However, threads 401a and 401b are potentially capable of out-of-order execution, where allocator and renamer block 430 also reserves other resources, such as reorder buffers to track instruction results. Unit 430 may also include a register renamer to rename

program/instruction reference registers to other registers internal to processor **400**. Reorder/retirement unit **435** includes components, such as the reorder buffers mentioned above, load buffers, and store buffers, to support out-of-order execution and later in-order retirement of instructions executed out-of-order.

Scheduler and execution unit(s) block **440**, in one embodiment, includes a scheduler unit to schedule instructions/operation on execution units. For example, a floating point instruction is scheduled on a port of an execution unit that has an available floating point execution unit. Register files associated with the execution units are also included to store information instruction processing results. Exemplary execution units include a floating point execution unit, an integer execution unit, a jump execution unit, a load execution unit, a store execution unit, and other known execution units.

Lower level data cache and data translation lookaside buffer (D-TLB) **450** are coupled to execution unit(s) **440**. The data cache is to store recently used/operated on elements, such as data operands, which are potentially held in memory coherency states. The D-TLB is to store recent virtual/linear to physical address translations. As a specific example, a processor may include a page table structure to break physical memory into a plurality of virtual pages.

Here, cores **401** and **402** share access to higher-level or further-out cache **410**, which is to cache recently fetched elements. Note that higher-level or further-out refers to cache levels increasing or getting further away from the execution unit(s). In one embodiment, higher-level cache **410** is a last-level data cache—last cache in the memory hierarchy on processor **400**—such as a second or third level data cache. However, higher level cache **410** is not so limited, as it may be associated with or includes an instruction cache. A trace cache—a type of instruction cache—instead may be coupled after decoder **425** to store recently decoded traces.

In the depicted configuration, processor **400** also includes bus interface module **405** and a power control unit **460**, which may perform power management in accordance with an embodiment of the present invention. In this scenario, bus interface **405** is to communicate with devices external to processor **400**, such as system memory and other components.

A memory controller **470** may interface with other devices such as one or many memories. In an example, bus interface **405** includes a ring interconnect with a memory controller for interfacing with a memory and a graphics controller for interfacing with a graphics processor. In an SoC environment, even more devices, such as a network interface, coprocessors, memory, graphics processor, and any other known computer devices/interface may be integrated on a single die or integrated circuit to provide small form factor with high functionality and low power consumption.

Referring now to FIG. 5, shown is a block diagram of a micro-architecture of a processor core in accordance with one embodiment of the present invention. As shown in FIG. 5, processor core **500** may be a multi-stage pipelined out-of-order processor. Core **500** may operate at various voltages based on a received operating voltage, which may be received from an integrated voltage regulator or external voltage regulator.

As seen in FIG. 5, core **500** includes front end units **510**, which may be used to fetch instructions to be executed and prepare them for use later in the processor pipeline. For example, front end units **510** may include a fetch unit **501**,

an instruction cache **503**, and an instruction decoder **505**. In some implementations, front end units **510** may further include a trace cache, along with microcode storage as well as a micro-operation storage. Fetch unit **501** may fetch macro-instructions, e.g., from memory or instruction cache **503**, and feed them to instruction decoder **505** to decode them into primitives, i.e., micro-operations for execution by the processor.

Coupled between front end units **510** and execution units **520** is an out-of-order (OOO) engine **515** that may be used to receive the micro-instructions and prepare them for execution. More specifically OOO engine **515** may include various buffers to re-order micro-instruction flow and allocate various resources needed for execution, as well as to provide renaming of logical registers onto storage locations within various register files such as register file **530** and extended register file **535**. Register file **530** may include separate register files for integer and floating point operations. For purposes of configuration, control, and additional operations, a set of machine specific registers (MSRs) **538** may also be present and accessible to various logic within core **500** (and external to the core). For example, power limit information may be stored in one or more MSR and be dynamically updated as described herein.

Various resources may be present in execution units **520**, including, for example, various integer, floating point, and single instruction multiple data (SIMD) logic units, among other specialized hardware. For example, such execution units may include one or more arithmetic logic units (ALUs) **522** and one or more vector execution units **524**, among other such execution units.

Results from the execution units may be provided to retirement logic, namely a reorder buffer (ROB) **540**. More specifically, ROB **540** may include various arrays and logic to receive information associated with instructions that are executed. This information is then examined by ROB **540** to determine whether the instructions can be validly retired and result data committed to the architectural state of the processor, or whether one or more exceptions occurred that prevent a proper retirement of the instructions. Of course, ROB **540** may handle other operations associated with retirement.

As shown in FIG. 5, ROB **540** is coupled to a cache **550** which, in one embodiment may be a low level cache (e.g., an L1 cache) although the scope of the present invention is not limited in this regard. Also, execution units **520** can be directly coupled to cache **550**. From cache **550**, data communication may occur with higher level caches, system memory and so forth. While shown with this high level in the embodiment of FIG. 5, understand the scope of the present invention is not limited in this regard. For example, while the implementation of FIG. 5 is with regard to an out-of-order machine such as of an Intel® x86 instruction set architecture (ISA), the scope of the present invention is not limited in this regard. That is, other embodiments may be implemented in an in-order processor, a reduced instruction set computing (RISC) processor such as an ARM-based processor, or a processor of another type of ISA that can emulate instructions and operations of a different ISA via an emulation engine and associated logic circuitry.

Referring now to FIG. 6, shown is a block diagram of a micro-architecture of a processor core in accordance with another embodiment. In the embodiment of FIG. 6, core **600** may be a low power core of a different micro-architecture, such as an Intel® Atom™-based processor having a relatively limited pipeline depth designed to reduce power consumption. As seen, core **600** includes an instruction

cache **610** coupled to provide instructions to an instruction decoder **615**. A branch predictor **605** may be coupled to instruction cache **610**. Note that instruction cache **610** may further be coupled to another level of a cache memory, such as an L2 cache (not shown for ease of illustration in FIG. 6). In turn, instruction decoder **615** provides decoded instructions to an issue queue (IQ) **620** for storage and delivery to a given execution pipeline. A microcode ROM **618** is coupled to instruction decoder **615**.

A floating point pipeline **630** includes a floating point (FP) register file **632** which may include a plurality of architectural registers of a given bit with such as 128, 256 or 512 bits. Pipeline **630** includes a floating point scheduler **634** to schedule instructions for execution on one of multiple execution units of the pipeline. In the embodiment shown, such execution units include an ALU **635**, a shuffle unit **636**, and a floating point adder **638**. In turn, results generated in these execution units may be provided back to buffers and/or registers of register file **632**. Of course understand while shown with these few example execution units, additional or different floating point execution units may be present in another embodiment.

An integer pipeline **640** also may be provided. In the embodiment shown, pipeline **640** includes an integer (INT) register file **642** which may include a plurality of architectural registers of a given bit with such as 128 or 256 bits. Pipeline **640** includes an integer execution (IE) scheduler **644** to schedule instructions for execution on one of multiple execution units of the pipeline. In the embodiment shown, such execution units include an ALU **645**, a shifter unit **646**, and a jump execution unit (JEU) **648**. In turn, results generated in these execution units may be provided back to buffers and/or registers of register file **642**. Of course understand while shown with these few example execution units, additional or different integer execution units may be present in another embodiment.

A memory execution (ME) scheduler **650** may schedule memory operations for execution in an address generation unit (AGU) **652**, which is also coupled to a TLB **654**. As seen, these structures may couple to a data cache **660**, which may be a L0 and/or L1 data cache that in turn couples to additional levels of a cache memory hierarchy, including an L2 cache memory.

To provide support for out-of-order execution, an allocator/renamer **670** may be provided, in addition to a reorder buffer **680**, which is configured to reorder instructions executed out of order for retirement in order. Although shown with this particular pipeline architecture in the illustration of FIG. 6, understand that many variations and alternatives are possible.

Note that in a processor having asymmetric cores, such as in accordance with the micro-architectures of FIGS. 5 and 6, workloads may be dynamically swapped between the cores for power management reasons, as these cores, although having different pipeline designs and depths, may be of the same or related ISA. Such dynamic core swapping may be performed in a manner transparent to a user application (and possibly kernel also).

Referring to FIG. 7, shown is a block diagram of a micro-architecture of a processor core in accordance with yet another embodiment. As illustrated in FIG. 7, a core **700** may include a multi-staged in-order pipeline to execute at very low power consumption levels. As one such example, processor **700** may have a micro-architecture in accordance with an ARM Cortex A53 design available from ARM Holdings, LTD., Sunnyvale, Calif. In an implementation, an 8-stage pipeline may be provided that is configured to

execute both 32-bit and 64-bit code. Core **700** includes a fetch unit **710** that is configured to fetch instructions and provide them to a decode unit **715**, which may decode the instructions, e.g., macro-instructions of a given ISA such as an ARMv8 ISA. Note further that a queue **730** may couple to decode unit **715** to store decoded instructions. Decoded instructions are provided to an issue logic **725**, where the decoded instructions may be issued to a given one of multiple execution units.

With further reference to FIG. 7, issue logic **725** may issue instructions to one of multiple execution units. In the embodiment shown, these execution units include an integer unit **735**, a multiply unit **740**, a floating point/vector unit **750**, a dual issue unit **760**, and a load/store unit **770**. The results of these different execution units may be provided to a writeback (WB) unit **780**. Understand that while a single writeback unit is shown for ease of illustration, in some implementations separate writeback units may be associated with each of the execution units. Furthermore, understand that while each of the units and logic shown in FIG. 7 is represented at a high level, a particular implementation may include more or different structures. A processor designed using one or more cores having a pipeline as in FIG. 7 may be implemented in many different end products, extending from mobile devices to server systems.

Referring to FIG. 8, shown is a block diagram of a micro-architecture of a processor core in accordance with a still further embodiment. As illustrated in FIG. 8, a core **800** may include a multi-stage multi-issue out-of-order pipeline to execute at very high performance levels (which may occur at higher power consumption levels than core **700** of FIG. 7). As one such example, processor **800** may have a microarchitecture in accordance with an ARM Cortex A57 design. In an implementation, a 15 (or greater)-stage pipeline may be provided that is configured to execute both 32-bit and 64-bit code. In addition, the pipeline may provide for 3 (or greater)-wide and 3 (or greater)-issue operation. Core **800** includes a fetch unit **810** that is configured to fetch instructions and provide them to a decoder/renamer/dispatcher unit **815** coupled to a cache **820**. Unit **815** may decode the instructions, e.g., macro-instructions of an ARMv8 instruction set architecture, rename register references within the instructions, and dispatch the instructions (eventually) to a selected execution unit. Decoded instructions may be stored in a queue **825**. Note that while a single queue structure is shown for ease of illustration in FIG. 8, understand that separate queues may be provided for each of the multiple different types of execution units.

Also shown in FIG. 8 is an issue logic **830** from which decoded instructions stored in queue **825** may be issued to a selected execution unit. Issue logic **830** also may be implemented in a particular embodiment with a separate issue logic for each of the multiple different types of execution units to which issue logic **830** couples.

Decoded instructions may be issued to a given one of multiple execution units. In the embodiment shown, these execution units include one or more integer units **835**, a multiply unit **840**, a floating point/vector unit **850**, a branch unit **860**, and a load/store unit **870**. In an embodiment, floating point/vector unit **850** may be configured to handle SIMD or vector data of 128 or 256 bits. Still further, floating point/vector execution unit **850** may perform IEEE-754 double precision floating-point operations. The results of these different execution units may be provided to a writeback unit **880**. Note that in some implementations separate writeback units may be associated with each of the execution units. Furthermore, understand that while each of the

units and logic shown in FIG. 8 is represented at a high level, a particular implementation may include more or different structures.

Note that in a processor having asymmetric cores, such as in accordance with the micro-architectures of FIGS. 7 and 8, workloads may be dynamically swapped for power management reasons, as these cores, although having different pipeline designs and depths, may be of the same or related ISA. Such dynamic core swapping may be performed in a manner transparent to a user application (and possibly kernel also).

A processor designed using one or more cores having pipelines as in any one or more of FIGS. 5-8 may be implemented in many different end products, extending from mobile devices to server systems. Referring now to FIG. 9, shown is a block diagram of a processor in accordance with another embodiment of the present invention. In the embodiment of FIG. 9, processor 900 may be a SoC including multiple domains, each of which may be controlled to operate at an independent operating voltage and operating frequency. As a specific illustrative example, processor 900 may be an Intel® Architecture Core™-based processor such as an i3, i5, i7 or another such processor available from Intel Corporation. However, other low power processors such as available from Advanced Micro Devices, Inc. (AMD) of Sunnyvale, Calif., an ARM-based design from ARM Holdings, Ltd. or licensee thereof or a MIPS-based design from MIPS Technologies, Inc. of Sunnyvale, Calif., or their licensees or adopters may instead be present in other embodiments such as an Apple A7 processor, a Qualcomm Snapdragon processor, or Texas Instruments OMAP processor. Such SoC may be used in a low power system such as a smartphone, tablet computer, phablet computer, Ultrabook™ computer or other portable computing device, which may incorporate a heterogeneous system architecture having a heterogeneous system architecture-based processor design.

In the high level view shown in FIG. 9, processor 900 includes a plurality of core units 910a-910n. Each core unit may include one or more processor cores, one or more cache memories and other circuitry. Each core unit 910 may support one or more instruction sets (e.g., an x86 instruction set (with some extensions that have been added with newer versions); a MIPS instruction set; an ARM instruction set (with optional additional extensions such as NEON)) or other instruction set or combinations thereof. Note that some of the core units may be heterogeneous resources (e.g., of a different design). In addition, each such core may be coupled to a cache memory (not shown) which in an embodiment may be a shared level two (L2) cache memory. A non-volatile storage 930 may be used to store various program and other data. For example, this storage may be used to store at least portions of microcode, boot information such as a BIOS, other system software or so forth.

Each core unit 910 may also include an interface such as a bus interface unit to enable interconnection to additional circuitry of the processor. In an embodiment, each core unit 910 couples to a coherent fabric that may act as a primary cache coherent on-die interconnect that in turn couples to a memory controller 935. In turn, memory controller 935 controls communications with a memory such as a DRAM (not shown for ease of illustration in FIG. 9).

In addition to core units, additional processing engines are present within the processor, including at least one graphics unit 920 which may include one or more graphics processing units (GPUs) to perform graphics processing as well as to possibly execute general purpose operations on the graphics

processor (so-called GPGPU operation). In addition, at least one image signal processor 925 may be present. Signal processor 925 may be configured to process incoming image data received from one or more capture devices, either internal to the SoC or off-chip.

Other accelerators also may be present. In the illustration of FIG. 9, a video coder 950 may perform coding operations including encoding and decoding for video information, e.g., providing hardware acceleration support for high definition video content. A display controller 955 further may be provided to accelerate display operations including providing support for internal and external displays of a system. In addition, a security processor 945 may be present to perform security operations such as secure boot operations, various cryptography operations and so forth.

Each of the units may have its power consumption controlled via a power manager 940, which may include control logic to perform the various power management techniques described herein.

In some embodiments, SoC 900 may further include a non-coherent fabric coupled to the coherent fabric to which various peripheral devices may couple. One or more interfaces 960a-960d enable communication with one or more off-chip devices. Such communications may be via a variety of communication protocols such as PCIe™, GPIO, USB, I²C, UART, MIPI, SDIO, DDR, SPI, HDMI, among other types of communication protocols. Although shown at this high level in the embodiment of FIG. 9, understand the scope of the present invention is not limited in this regard.

Referring now to FIG. 10, shown is a block diagram of a representative SoC. In the embodiment shown, SoC 1000 may be a multi-core SoC configured for low power operation to be optimized for incorporation into a smartphone or other low power device such as a tablet computer or other portable computing device. As an example, SoC 1000 may be implemented using asymmetric or different types of cores, such as combinations of higher power and/or low power cores, e.g., out-of-order cores and in-order cores. In different embodiments, these cores may be based on an Intel® Architecture™ core design or an ARM architecture design. In yet other embodiments, a mix of Intel and ARM cores may be implemented in a given SoC.

As seen in FIG. 10, SoC 1000 includes a first core domain 1010 having a plurality of first cores 1012a-1012d. In an example, these cores may be low power cores such as in-order cores. In one embodiment these first cores may be implemented as ARM Cortex A53 cores. In turn, these cores couple to a cache memory 1015 of core domain 1010. In addition, SoC 1000 includes a second core domain 1020. In the illustration of FIG. 10, second core domain 1020 has a plurality of second cores 1022a-1022d. In an example, these cores may be higher power-consuming cores than first cores 1012. In an embodiment, the second cores may be out-of-order cores, which may be implemented as ARM Cortex A57 cores. In turn, these cores couple to a cache memory 1025 of core domain 1020. Note that while the example shown in FIG. 10 includes 4 cores in each domain, understand that more or fewer cores may be present in a given domain in other examples.

With further reference to FIG. 10, a graphics domain 1030 also is provided, which may include one or more graphics processing units (GPUs) configured to independently execute graphics workloads, e.g., provided by one or more cores of core domains 1010 and 1020. As an example, GPU domain 1030 may be used to provide display support for a variety of screen sizes, in addition to providing graphics and display rendering operations.

As seen, the various domains couple to a coherent interconnect **1040**, which in an embodiment may be a cache coherent interconnect fabric that in turn couples to an integrated memory controller **1050**. Coherent interconnect **1040** may include a shared cache memory, such as an L3 cache, in some examples. In an embodiment, memory controller **1050** may be a direct memory controller to provide for multiple channels of communication with an off-chip memory, such as multiple channels of a DRAM (not shown for ease of illustration in FIG. **10**).

In different examples, the number of the core domains may vary. For example, for a low power SoC suitable for incorporation into a mobile computing device, a limited number of core domains such as shown in FIG. **10** may be present. Still further, in such low power SoCs, core domain **1020** including higher power cores may have fewer numbers of such cores. For example, in one implementation two cores **1022** may be provided to enable operation at reduced power consumption levels. In addition, the different core domains may also be coupled to an interrupt controller to enable dynamic swapping of workloads between the different domains.

In yet other embodiments, a greater number of core domains, as well as additional optional IP logic may be present, in that an SoC can be scaled to higher performance (and power) levels for incorporation into other computing devices, such as desktops, servers, high performance computing systems, base stations and so forth. As one such example, 4 core domains each having a given number of out-of-order cores may be provided. Still further, in addition to optional GPU support (which as an example may take the form of a GPGPU), one or more accelerators to provide optimized hardware support for particular functions (e.g. web serving, network processing, switching or so forth) also may be provided. In addition, an input/output interface may be present to couple such accelerators to off-chip components.

Referring now to FIG. **11**, shown is a block diagram of another example SoC. In the embodiment of FIG. **11**, SoC **1100** may include various circuitry to enable high performance for multimedia applications, communications and other functions. As such, SoC **1100** is suitable for incorporation into a wide variety of portable and other devices, such as smartphones, tablet computers, smart TVs and so forth. In the example shown, SoC **1100** includes a central processor unit (CPU) domain **1110**. In an embodiment, a plurality of individual processor cores may be present in CPU domain **1110**. As one example, CPU domain **1110** may be a quad core processor having 4 multithreaded cores. Such processors may be homogeneous or heterogeneous processors, e.g., a mix of low power and high power processor cores.

In turn, a GPU domain **1120** is provided to perform advanced graphics processing in one or more GPUs to handle graphics and compute APIs. A DSP unit **1130** may provide one or more low power DSPs for handling low-power multimedia applications such as music playback, audio/video and so forth, in addition to advanced calculations that may occur during execution of multimedia instructions. In turn, a communication unit **1140** may include various components to provide connectivity via various wireless protocols, such as cellular communications (including 3G/4G LTE), wireless local area protocols such as Bluetooth™, IEEE 802.11, and so forth.

Still further, a multimedia processor **1150** may be used to perform capture and playback of high definition video and audio content, including processing of user gestures. A sensor unit **1160** may include a plurality of sensors and/or a sensor controller to interface to various off-chip sensors

present in a given platform. An image signal processor **1170** may be provided with one or more separate ISPs to perform image processing with regard to captured content from one or more cameras of a platform, including still and video cameras.

A display processor **1180** may provide support for connection to a high definition display of a given pixel density, including the ability to wirelessly communicate content for playback on such display. Still further, a location unit **1190** may include a GPS receiver with support for multiple GPS constellations to provide applications highly accurate positioning information obtained using as such GPS receiver. Understand that while shown with this particular set of components in the example of FIG. **11**, many variations and alternatives are possible.

Referring now to FIG. **12**, shown is a block diagram of an example system with which embodiments can be used. As seen, system **1200** may be a smartphone or other wireless communicator. A baseband processor **1205** is configured to perform various signal processing with regard to communication signals to be transmitted from or received by the system. In turn, baseband processor **1205** is coupled to an application processor **1210**, which may be a main CPU of the system to execute an OS and other system software, in addition to user applications such as many well-known social media and multimedia apps. Application processor **1210** may further be configured to perform a variety of other computing operations for the device.

In turn, application processor **1210** can couple to a user interface/display **1220**, e.g., a touch screen display. In addition, application processor **1210** may couple to a memory system including a non-volatile memory, namely a flash memory **1230** and a system memory, namely a dynamic random access memory (DRAM) **1235**. As further seen, application processor **1210** further couples to a capture device **1240** such as one or more image capture devices that can record video and/or still images.

Still referring to FIG. **12**, a universal integrated circuit card (UICC) **1240** comprising a subscriber identity module and possibly a secure storage and cryptoprocessor is also coupled to application processor **1210**. System **1200** may further include a security processor **1250** that may couple to application processor **1210**. A plurality of sensors **1225** may couple to application processor **1210** to enable input of a variety of sensed information such as accelerometer and other environmental information. An audio output device **1295** may provide an interface to output sound, e.g., in the form of voice communications, played or streaming audio data and so forth.

As further illustrated, a near field communication (NFC) contactless interface **1260** is provided that communicates in a NFC near field via an NFC antenna **1265**. While separate antennae are shown in FIG. **12**, understand that in some implementations one antenna or a different set of antennae may be provided to enable various wireless functionality.

A power management integrated circuit (PMIC) **1215** couples to application processor **1210** to perform platform level power management. To this end, PMIC **1215** may issue power management requests to application processor **1210** to enter certain low power states as desired. Furthermore, based on platform constraints, PMIC **1215** may also control the power level of other components of system **1200**.

To enable communications to be transmitted and received, various circuitry may be coupled between baseband processor **1205** and an antenna **1290**. Specifically, a radio frequency (RF) transceiver **1270** and a wireless local area network (WLAN) transceiver **1275** may be present. In

general, RF transceiver **1270** may be used to receive and transmit wireless data and calls according to a given wireless communication protocol such as 3G or 4G wireless communication protocol such as in accordance with a code division multiple access (CDMA), global system for mobile communication (GSM), long term evolution (LTE) or other protocol. In addition a GPS sensor **1280** may be present. Other wireless communications such as receipt or transmission of radio signals, e.g., AM/FM and other signals may also be provided. In addition, via WLAN transceiver **1275**, local wireless communications can also be realized.

Referring now to FIG. **13**, shown is a block diagram of another example system with which embodiments may be used. In the illustration of FIG. **13**, system **1300** may be mobile low-power system such as a tablet computer, 2:1 tablet, phablet or other convertible or standalone tablet system. As illustrated, a SoC **1310** is present and may be configured to operate as an application processor for the device.

A variety of devices may couple to SoC **1310**. In the illustration shown, a memory subsystem includes a flash memory **1340** and a DRAM **1345** coupled to SoC **1310**. In addition, a touch panel **1320** is coupled to the SoC **1310** to provide display capability and user input via touch, including provision of a virtual keyboard on a display of touch panel **1320**. To provide wired network connectivity, SoC **1310** couples to an Ethernet interface **1330**. A peripheral hub **1325** is coupled to SoC **1310** to enable interfacing with various peripheral devices, such as may be coupled to system **1300** by any of various ports or other connectors.

In addition to internal power management circuitry and functionality within SoC **1310**, a PMIC **1380** is coupled to SoC **1310** to provide platform-based power management, e.g., based on whether the system is powered by a battery **1390** or AC power via an AC adapter **1395**. In addition to this power source-based power management, PMIC **1380** may further perform platform power management activities based on environmental and usage conditions. Still further, PMIC **1380** may communicate control and status information to SoC **1310** to cause various power management actions within SoC **1310**.

Still referring to FIG. **13**, to provide for wireless capabilities, a WLAN unit **1350** is coupled to SoC **1310** and in turn to an antenna **1355**. In various implementations, WLAN unit **1350** may provide for communication according to one or more wireless protocols.

As further illustrated, a plurality of sensors **1360** may couple to SoC **1310**. These sensors may include various accelerometer, environmental and other sensors, including user gesture sensors. Finally, an audio codec **1365** is coupled to SoC **1310** to provide an interface to an audio output device **1370**. Of course understand that while shown with this particular implementation in FIG. **13**, many variations and alternatives are possible.

Referring now to FIG. **14**, shown is a block diagram of a representative computer system such as notebook, Ultra-book™ or other small form factor system. A processor **1410**, in one embodiment, includes a microprocessor, multi-core processor, multithreaded processor, an ultra low voltage processor, an embedded processor, or other known processing element. In the illustrated implementation, processor **1410** acts as a main processing unit and central hub for communication with many of the various components of the system **1400**. As one example, processor **1410** is implemented as a SoC.

Processor **1410**, in one embodiment, communicates with a system memory **1415**. As an illustrative example, the

system memory **1415** is implemented via multiple memory devices or modules to provide for a given amount of system memory.

To provide for persistent storage of information such as data, applications, one or more operating systems and so forth, a mass storage **1420** may also couple to processor **1410**. In various embodiments, to enable a thinner and lighter system design as well as to improve system responsiveness, this mass storage may be implemented via a SSD or the mass storage may primarily be implemented using a hard disk drive (HDD) with a smaller amount of SSD storage to act as a SSD cache to enable non-volatile storage of context state and other such information during power down events so that a fast power up can occur on re-initiation of system activities. Also shown in FIG. **14**, a flash device **1422** may be coupled to processor **1410**, e.g., via a serial peripheral interface (SPI). This flash device may provide for non-volatile storage of system software, including a basic input/output software (BIOS) as well as other firmware of the system.

Various input/output (I/O) devices may be present within system **1400**. Specifically shown in the embodiment of FIG. **14** is a display **1424** which may be a high definition LCD or LED panel that further provides for a touch screen **1425**. In one embodiment, display **1424** may be coupled to processor **1410** via a display interconnect that can be implemented as a high performance graphics interconnect. Touch screen **1425** may be coupled to processor **1410** via another interconnect, which in an embodiment can be an I²C interconnect. As further shown in FIG. **14**, in addition to touch screen **1425**, user input by way of touch can also occur via a touch pad **1430** which may be configured within the chassis and may also be coupled to the same I²C interconnect as touch screen **1425**.

For perceptual computing and other purposes, various sensors may be present within the system and may be coupled to processor **1410** in different manners. Certain inertial and environmental sensors may couple to processor **1410** through a sensor hub **1440**, e.g., via an I²C interconnect. In the embodiment shown in FIG. **14**, these sensors may include an accelerometer **1441**, an ambient light sensor (ALS) **1442**, a compass **1443** and a gyroscope **1444**. Other environmental sensors may include one or more thermal sensors **1446** which in some embodiments couple to processor **1410** via a system management bus (SMBus) bus.

Also seen in FIG. **14**, various peripheral devices may couple to processor **1410** via a low pin count (LPC) interconnect. In the embodiment shown, various components can be coupled through an embedded controller **1435**. Such components can include a keyboard **1436** (e.g., coupled via a PS2 interface), a fan **1437**, and a thermal sensor **1439**. In some embodiments, touch pad **1430** may also couple to EC **1435** via a PS2 interface. In addition, a security processor such as a trusted platform module (TPM) **1438** may also couple to processor **1410** via this LPC interconnect.

System **1400** can communicate with external devices in a variety of manners, including wirelessly. In the embodiment shown in FIG. **14**, various wireless modules, each of which can correspond to a radio configured for a particular wireless communication protocol, are present. One manner for wireless communication in a short range such as a near field may be via a NFC unit **1445** which may communicate, in one embodiment with processor **1410** via an SMBus. Note that via this NFC unit **1445**, devices in close proximity to each other can communicate.

As further seen in FIG. **14**, additional wireless units can include other short range wireless engines including a

WLAN unit **1450** and a Bluetooth™ unit **1452**. Using WLAN unit **1450**, Wi-Fi™ communications can be realized, while via Bluetooth™ unit **1452**, short range Bluetooth™ communications can occur. These units may communicate with processor **1410** via a given link.

In addition, wireless wide area communications, e.g., according to a cellular or other wireless wide area protocol, can occur via a WWAN unit **1456** which in turn may couple to a subscriber identity module (SIM) **1457**. In addition, to enable receipt and use of location information, a GPS module **1455** may also be present. Note that in the embodiment shown in FIG. **14**, WWAN unit **1456** and an integrated capture device such as a camera module **1454** may communicate via a given link.

To provide for audio inputs and outputs, an audio processor can be implemented via a digital signal processor (DSP) **1460**, which may couple to processor **1410** via a high definition audio (HDA) link. Similarly, DSP **1460** may communicate with an integrated coder/decoder (CODEC) and amplifier **1462** that in turn may couple to output speakers **1463** which may be implemented within the chassis. Similarly, amplifier and CODEC **1462** can be coupled to receive audio inputs from a microphone **1465** which in an embodiment can be implemented via dual array microphones (such as a digital microphone array) to provide for high quality audio inputs to enable voice-activated control of various operations within the system. Note also that audio outputs can be provided from amplifier/CODEC **1462** to a headphone jack **1464**. Although shown with these particular components in the embodiment of FIG. **14**, understand the scope of the present invention is not limited in this regard.

Embodiments may be implemented in many different system types. Referring now to FIG. **15**, shown is a block diagram of a system in accordance with an embodiment of the present invention. As shown in FIG. **15**, multiprocessor system **1500** is a point-to-point interconnect system, and includes a first processor **1570** and a second processor **1580** coupled via a point-to-point interconnect **1550**. As shown in FIG. **15**, each of processors **1570** and **1580** may be multicore processors, including first and second processor cores (i.e., processor cores **1574a** and **1574b** and processor cores **1584a** and **1584b**), although potentially many more cores may be present in the processors. Each of the processors can include a PCU or other power management logic to perform processor-based power management as described herein, including P-state control based at least in part on one or more of workload concurrency, overlap, and/or characterization.

Still referring to FIG. **15**, first processor **1570** further includes a memory controller hub (MCH) **1572** and point-to-point (P-P) interfaces **1576** and **1578**. Similarly, second processor **1580** includes a MCH **1582** and P-P interfaces **1586** and **1588**. As shown in FIG. **15**, MCH's **1572** and **1582** couple the processors to respective memories, namely a memory **1532** and a memory **1534**, which may be portions of system memory (e.g., DRAM) locally attached to the respective processors. First processor **1570** and second processor **1580** may be coupled to a chipset **1590** via P-P interconnects **1562** and **1564**, respectively. As shown in FIG. **15**, chipset **1590** includes P-P interfaces **1594** and **1598**.

Furthermore, chipset **1590** includes an interface **1592** to couple chipset **1590** with a high performance graphics engine **1538**, by a P-P interconnect **1539**. In turn, chipset **1590** may be coupled to a first bus **1516** via an interface **1596**. As shown in FIG. **15**, various input/output (I/O) devices **1514** may be coupled to first bus **1516**, along with a bus bridge **1518** which couples first bus **1516** to a second

bus **1520**. Various devices may be coupled to second bus **1520** including, for example, a keyboard/mouse **1522**, communication devices **1526** and a data storage unit **1528** such as a disk drive or other mass storage device which may include code **1530**, in one embodiment. Further, an audio I/O **1524** may be coupled to second bus **1520**. Embodiments can be incorporated into other types of systems including mobile devices such as a smart cellular telephone, tablet computer, netbook, Ultrabook™, or so forth.

One or more aspects of at least one embodiment may be implemented by representative code stored on a machine-readable medium which represents and/or defines logic within an integrated circuit such as a processor. For example, the machine-readable medium may include instructions which represent various logic within the processor. When read by a machine, the instructions may cause the machine to fabricate the logic to perform the techniques described herein. Such representations, known as "IP cores," are reusable units of logic for an integrated circuit that may be stored on a tangible, machine-readable medium as a hardware model that describes the structure of the integrated circuit. The hardware model may be supplied to various customers or manufacturing facilities, which load the hardware model on fabrication machines that manufacture the integrated circuit. The integrated circuit may be fabricated such that the circuit performs operations described in association with any of the embodiments described herein.

FIG. **16** is a block diagram illustrating an IP core development system **1600** that may be used to manufacture an integrated circuit to perform operations according to an embodiment. The IP core development system **1600** may be used to generate modular, re-usable designs that can be incorporated into a larger design or used to construct an entire integrated circuit (e.g., an SoC integrated circuit). A design facility **1630** can generate a software simulation **1610** of an IP core design in a high level programming language (e.g., C/C++). The software simulation **1610** can be used to design, test, and verify the behavior of the IP core. A register transfer level (RTL) design can then be created or synthesized from the simulation model **1600**. The RTL design **1615** is an abstraction of the behavior of the integrated circuit that models the flow of digital signals between hardware registers, including the associated logic performed using the modeled digital signals. In addition to an RTL design **1615**, lower-level designs at the logic level or transistor level may also be created, designed, or synthesized. Thus, the particular details of the initial design and simulation may vary.

The RTL design **1615** or equivalent may be further synthesized by the design facility into a hardware model **1620**, which may be in a hardware description language (HDL), or some other representation of physical design data. The HDL may be further simulated or tested to verify the IP core design. The IP core design can be stored for delivery to a third party fabrication facility **1665** using non-volatile memory **1640** (e.g., hard disk, flash memory, or any non-volatile storage medium). Alternately, the IP core design may be transmitted (e.g., via the Internet) over a wired connection **1650** or wireless connection **1660**. The fabrication facility **1665** may then fabricate an integrated circuit that is based at least in part on the IP core design. The fabricated integrated circuit can be configured to perform operations in accordance with at least one embodiment described herein.

With hardware-controlled performance state updates, false positives may occur in certain periodic repetitive workloads such as frame-based algorithms, in which performance state may be increased (e.g., to a maximum turbo

mode frequency) every frame (or few frames), which can waste energy and performance. Embodiments may be used to discriminate between an interactive workload that gains from performance state increases and cyclic workloads that do not. Embodiments may also consider the impact of a multi-threaded workload in determining whether to increase performance state of a processor. In various embodiments, the power management techniques disclosed herein maximize responsiveness with minimal impact on throughput of workloads. More specifically, embodiments may improve the responsiveness of short burst user interactive action while minimizing the energy lost when executing non-user interactive workloads. Stated another way, a performance state of at least one engine may be increased when a processor executes a user interactive workload and to prevent such increase when the processor is not executing a user interactive workload.

In different embodiments, a technique may be used to identify a parallel workload and as a result of it, take an action to relax in some way the aggressiveness of a P-state control logic in order to save power. Note also that such techniques may be used in multi-threaded cases where the assumption is that power consumption is higher than in single thread cases. Assume for purposes of discussion a two core base processor with the following assumptions:

P=utilization percentage of each core (assume that both cores have the same utilization percent).

As such:

Two cores are active: P^2

Only one core is active: $2*(1-P)*P$

Zero cores are active: $(1-P)^2$

Sum of number of active cores: $2*P^2+1*2*(1-P)*P$

At least one core is active: $1-(1-P)^2$

Average number of active cores=(Sum of number of active cores)/(At least one core active).

Based on this information, a configuration table may be provided for incorporation in a processor. Referring now to Table 1, shown is a representative configuration table in accordance with an embodiment. Understand that while shown with these representation values based on above equations, other embodiments may provide a different configuration table derived from the same or different equations. This configuration table may be fused into a processor during manufacture, written into a non-volatile storage, e.g., as firmware-coded values, and/or input into a storage at processor reset.

TABLE 1

P(C0) for one core	2	1	0	any sum	C0 any	avg dc
0	0	0	1	0	0	1
0.05	0.00	0.10	0.90	0.10	0.10	1.03
0.1	0.01	0.18	0.81	0.20	0.19	1.05
0.15	0.02	0.26	0.72	0.30	0.28	1.08
0.2	0.04	0.32	0.64	0.40	0.36	1.11
0.25	0.06	0.38	0.56	0.50	0.44	1.14
0.3	0.09	0.42	0.49	0.60	0.51	1.18
0.35	0.12	0.46	0.42	0.70	0.58	1.21
0.4	0.16	0.48	0.36	0.80	0.64	1.25
0.45	0.20	0.50	0.30	0.90	0.70	1.29
0.5	0.25	0.50	0.25	1.00	0.75	1.33
0.55	0.30	0.50	0.20	1.10	0.80	1.38
0.6	0.36	0.48	0.16	1.20	0.84	1.43
0.65	0.42	0.46	0.12	1.30	0.88	1.48
0.7	0.49	0.42	0.09	1.40	0.91	1.54
0.75	0.56	0.38	0.06	1.50	0.94	1.60
0.8	0.64	0.32	0.04	1.60	0.96	1.67
0.85	0.72	0.26	0.02	1.70	0.98	1.74

TABLE 1-continued

P(C0) for one core	2	1	0	any sum	C0 any	avg dc
0.9	0.81	0.18	0.01	1.80	0.99	1.82
0.95	0.90	0.10	0.00	1.90	1.00	1.90
1	1.00	0.00	0.00	2.00	1.00	2.00

In various embodiments an average number of active cores can be determined using count values obtained from certain processor counters. More specifically, in one embodiment these counters include: a C0_ANY_SUM counter, which counts the amount of time that each core is active; and a C0_ANY_counter, which counts the amount of time that any core is active. The average number of active cores can be determined by computing an average over time of the hardware counters.

By using a rolling average technique (or another averaging technique), the average number of active cores can be determined over a given evaluation time window. Of course, other methods of averaging can also be used. In case that this average value is higher than a theoretical calculated value (also referred to herein as an estimated core overall active residency level), it can be assumed that the current active workload is parallel. If the average number of active cores is lower than this value, it can be assumed that there is no dependency between the execution of workloads on each core.

In an embodiment, P-state control logic is generally adapted to measure the average active time of each core over a defined time window. If this average measured value exceeds a defined high active residency threshold, the logic requests an increase of the P-state (e.g., increase in core frequency and/or voltage). If the average measured value falls below a defined low active residency threshold, the logic requests a decrement of the P-state.

In case that a workload is identified as parallel, embodiments may update one or more parameters used by the P-state control logic. More specifically, lower and higher P-state selection thresholds can be updated (e.g., increased) and/or the evaluation time window used to estimate the average core utilization can be increased. Increasing the thresholds may cause a lower P-state (e.g., lower operating frequency) because a very high value of core utilization may be present in order to increase P-state (and vice versa). This threshold control also may make the flow more relaxed compared to a case where the thresholds are lower. Note that the level of P-state selection thresholds change and time window change can be a function of actual power change compared to the current status. As an example, for a unit that executes at a high P-state, its thresholds and time window change can be higher as compared to a unit that executes only at lower P-state ranges. (Assume that a highest P-state consumes more power than a lower P-state). Note that when the time window is longer, changes of P-state will be less frequent, and the higher the threshold values used will relax the number of P-state changes. The control of change can also be a function of the number of the cores that the processor supports. Increasing the time window used to estimate the average core utilization helps to stable this estimation and minimize the effect of random active events.

Referring now to FIG. 17, shown is a flow diagram of a method in accordance with an embodiment of the present invention. More specifically, method 1700 of FIG. 17 may be performed by power control logic of a processor, which in an embodiment can be implemented as one or more

microcontrollers or other hardware-based logic. In some cases such logic may be implemented by appropriate combinations of hardware, software, and/or firmware.

As seen, method 1700 begins by predicting an average number of active cores (block 1710). Note that for purposes of illustration, the discussion of FIG. 17 is with regard to multiple cores of a multi-core processor over a given evaluation interval. Understand that in other cases, processing engines may be of other types than a core, such as a graphics processor or any other type of processing engine. In an embodiment, the average number of active core prediction may be made for an evaluation interval on the order of approximately between 10-100 milliseconds. Still further, in an embodiment such prediction may be made based on count values maintained in one or more counters as described herein. Next control passes to block 1720, where a maximum core active residency level can be determined. In an embodiment, this determination may identify the core having a highest active state residency level (generally referred to as a maximum utilization level), which can be expressed in terms of a utilization percentage.

Still with reference to FIG. 17, next at block 1730 an estimated core overall residency level can be determined. More specifically, such determination may be made as a function of the maximum core active residency level. In one embodiment, this maximum core active residency level may be used to access an entry in a configuration storage as in the example of Table 1 above, that provides a table of estimated core overall residency levels based on a given active residency level or utilization value. Thereafter, control passes to diamond 1740 to determine whether a sum of the estimated core overall residency level minus a threshold value is less than the average number of active cores. This threshold value may be used to provide a measure of hysteresis. Note that in other cases, this comparison between the estimated core overall residency level and the active number of active cores can be directly made without inclusion of a threshold (e.g., based on a function of the maximum C0 %) (or equally, the threshold may be applied to the average number of active cores instead of the estimated core overall residency level).

In any case, if the comparison indicates that the sum value is less than the average number of active cores, control passes to block 1750 where an adjustment may be made to one or more input parameters used by P-state control logic in determining an appropriate P-state. More specifically as shown in FIG. 17, at least one of one or more active residency thresholds and an evaluation time window for the P-state control logic may be adjusted. Note that if the sum is not less than the average number of active cores, it can be assumed that there is a dependency between the cores and the interval may be counted as a consumer producer case of software thread migration between the different cores.

As an example, to reduce aggressiveness of P-state increases when the multiple cores are executing a parallel workload, a high active residency threshold may be increased to a higher level (e.g., 90%), and a low active residency threshold can be increased (e.g., to approximately 70%). Still further, the evaluation time window may be increased. For example, in one embodiment this time window may be doubled from approximately 30 milliseconds to approximately 60 milliseconds. Such adjustments reduce the aggressiveness of performance state control, as with the longer evaluation interval and higher high active residency threshold, a determination of a performance state increase is less likely to occur, in the absence of a very active processor.

Understand while shown at this high level in the embodiment of FIG. 17, many variations and alternatives are possible.

The amount of threshold and time window change is also relative to the estimated number of active cores. As soon as this estimation is higher, the changes also may be increased. For example, up to a given value (e.g., up to 1.4), the amount of change can be zero, above this value (e.g., between 1.4 to 1.6), the changes be a linear function of increasing value, e.g., up to 20% for the low and high threshold and increasing the time window, e.g., by 30 milliseconds.

The amount of changes of the threshold and time windows can also be a function of the actual physical definition of the cores and the actual possible energy consumption, so that different values of threshold and time window change can be based on the possible maximum ratio energy consumption that the cores can consume.

In addition to the above-described methods, embodiments may further analyze activity overlap between multiple processing engines and control performance states accordingly. More specifically, embodiments may independently (and/or alternately) identify whether two or more processing engines are correlated or anti-correlated with each other. Based on such identification, it can be determined whether an increase in an operating frequency of one or more of the engines by way of a performance state increase will degrade or improve energy consumption. As discussed above, such engines may be cores, graphics, processors or other processing engines of a SoC or other processor.

Here, multiple processing engines can be considered to be correlated when they tend to work together. In turn, multiple processing engines can be considered to be anti-correlated when they tend to work mostly apart (independently) from each other. If two engines are correlated, increasing a performance state of one engine (such as by increasing operating frequency) may not improve system energy consumption most of the time, since the other engine will block the system from entering into a particular, e.g., deep, low power state. If the engines are anti-correlated, increasing the performance state of one of the engines (such as by increasing operating frequency) will likely to result in additional sleep periods which could result in additional energy savings.

In an embodiment, the following variables may be used in determining whether to enable a performance state update to occur. More specifically, the following variables may be used (assume the presence of multiple processing engines E1 and E2 and an evaluation interval of T):

UTILIZATION(E_i,T): utilization of engine E_i over time interval T.

MEASURED_OVERLAP(E1,E2,T): a measured counter that represents the overlap time of engines E1 and E2 over time interval of T.

In an embodiment, this counter is incremented by one in a given clock cycle when both engines E1 and E2 are working together (both in an active state). In turn, RANDOM_OVERLAP(X,Y,T) is an estimation of the overlap time within time interval T between two engines, one of which has utilization X and the other has utilization Y (assume that X is the greater of the two utilizations). The overlap time is estimated assuming that the two engines are independent and unaware of each other.

While the UTILIZATION and MEASURED_OVERLAP metrics can be measured using the count values of counters present in the processor, the RANDOM_OVERLAP metric can be calculated, based on the assumption that there are two

25

independent engines, each of which works consecutively in the time frame of observation T.

Referring now to FIG. 18, shown is a block diagram illustrating activity of two processing engines X and Y during a time interval T. As seen, the two engines are active for different time durations, one of length X, the other of length Y, in the total time interval T.

In an embodiment, $RANDOM_OVERLAP(X,Y,T)$ can be calculated as $X*Y/T$.

In order to understand the correlation between the two engines, a comparison may be made between $RANDOM_OVERLAP(X,Y,T)$ and $MEASURED_OVERLAP(E1,E2,T)$. In an embodiment, if $RANDOM_OVERLAP(X,Y,T)$ is greater than $MEASURED_OVERLAP(E1,E2,T)$, the engines are assumed to be anti-correlated. If $RANDOM_OVERLAP(X,Y,T)$ is less than or equals $MEASURED_OVERLAP(E1,E2,T)$, the two engines are assumed to be correlated (note that a threshold may be applied to reduce hysteresis).

Referring now to FIG. 19, shown is a flow diagram of a method for controlling performance state within a processor in accordance with another embodiment of the present invention. Method 1800 may be performed by power control logic of a processor as described above. As illustrated, method 1800 begins by measuring an activity overlap between first and second engines of a processor (block 1810). Again, while analysis of two engines is discussed for ease of illustration understand that in other cases, more than two engines can be analyzed.

Control next passes to block 1820 where an expected overlap between the engines can be estimated. More specifically, this expected overlap may be referred to as a random overlap value. In different embodiments, different manners of estimating this expected overlap can occur. Thereafter at diamond 1830 it is determined whether the measured activity overlap is greater than the random overlap value. If so, no further operations occur for a particular evaluation interval. If it is determined that the measured activity overlap is not greater than the random overlap value, control passes to block 1840 where a performance state of at least a selected one of the processing engines can be increased. This is the case, as here it is likely that the engines are executing anti-correlated workloads such that an increased performance state for at least a selected one of the engines may enable a pending workload to be completed faster, likely enabling the processor to enter into a deeper low power state. Understand while shown at this high level in the embodiment of FIG. 19, many variations and alternatives are possible.

Referring now to FIG. 20, shown is a block diagram of a portion of a processor in accordance with an embodiment of the present invention. More specifically, the portion of processor 1900 shown in FIG. 20 includes various hardware to be used for performing the analysis and performance state control operations described herein. As seen, processor 1900 includes a counter storage area 1910. In the embodiment shown, storage area 1910 includes multiple independent counters, including corresponding core activity counters 1912₀-1912_n. In an embodiment, each such activity counter 1912 may be associated with a particular core or other processing engine (such as a graphics processor) and may be adapted to count clock cycles in which the corresponding engine is in an active state. In addition, an activity sum counter 1914 may be adapted to count clock cycles in which engines are active. As such, in an embodiment this counter may be incremented by one or more in each clock cycle depending on the number of active engines. For example, if

26

two engines are active in a particular clock cycle, this counter may be incremented by two. Finally, FIG. 20 shows that counter storage area 1910 includes an activity any counter 1916. In an embodiment, counter 1916 may be adapted to be incremented by one each clock cycle in which at least one engine is active. In addition, an overlap counter 1918 may be configured to measure the overlap duration between two engines. The definition is: the overlap counter between engine 1 and engine 2 is incremented by one every cycle that engine 1 and engine 2 are both working.

Note that the counts maintained by these counters may be used by one or more of an activity determination logic 1920, a P-state control logic 1930, and an engine correlation determination logic 1940. In an embodiment, activity determination logic 1920 and engine correlation determination logic 1940 may be adapted to perform analysis and control operations for P-state control logic 1930. As an example, activity determination logic 1920 may perform method 1700 of FIG. 17. In turn, engine correlation determination logic 1940 may be adapted to perform method 1800 of FIG. 19.

P-state control logic 1930 may receive, in addition to control inputs from these logics, a user input 1932, one or more thresholds from a threshold storage 1934, and an evaluation time window storage 1936. In various embodiments, user input 1932 may be a user-requested performance state. For example, based on a user request for a given mode of operation, e.g., a performance-biased mode, an energy efficient-biased mode and/or a balanced mode, an appropriate input to P-state control logic 1930 is provided. P-state control logic 1930 may select appropriate lower and upper activity threshold values from storage 1934, e.g., based on the determination made in activity determination logic 1920. In addition, an appropriate evaluation interval for making P-state determinations may be obtained from evaluation time window storage 1936, again based on activity determination logic 1920.

Based on all of this information, P-state control logic 1930 may determine an appropriate P-state for one or more cores or other processing engines for a next operation interval. For example, control logic 1930 may start with a baseline determination according to user input 1932 and then may fine tune this user request to generate an appropriate performance state. For example, in cases where it is determined that non-parallel or anti-correlated workloads are being executed, a performance state may be more aggressively increased. In contrast, in situations where a correlated or parallel workload is being executed, a user request can be clipped to a lower performance state, as it is likely that the energy efficiency of more rapidly completing a workload would not lead to a deeper power saving state.

FIG. 21 is a timing illustration showing utilization of two processing engines (engine0 and engine1), measured at time intervals of 50 msec. FIG. 22 is a timing illustration showing the measured overlap together with the calculated random overlap of the two engines for a given workload.

In various embodiments, performance state control tuning as described herein can be used in connection with a race to halt algorithm. The race to halt algorithm increases operating frequency of one or more processor cores of a core domain of a SoC or other processor, in order to increase SoC sleep time. Embodiments may accurately estimate whether the SoC will indeed go to sleep as the core domain finishes executing, or whether another compute engine (such as a graphics processor of a graphics domain) will block the SoC from sleeping. In such cases, a comparison of overlap between the two compute domains (core and graphics) can be performed. If the domains are positively correlated, no

increase in performance state for, e.g., the core domain is allowed, since the graphics domain is likely to block the SoC from sleeping (on completion of the core domain workload). If instead the domains are anti-correlated, a performance state of, e.g., the core domain may be increased when possible (such as where there are no processor constraints, and available power budget is present), because the probability to gain SoC sleep state is high.

The following examples pertain to further embodiments.

In one example, a processor comprises: a plurality of processing engines including a first processing engine and a second processing engine to independently execute instructions; and a power controller including a performance state control logic to control a performance state of at least one of the plurality of processing engines and a first logic to determine an average number of active processing engines over a first window, an estimated activity level of the processor for the first window, and adjust at least one of a window length at which the performance state control logic is to perform a performance state determination and at least one activity level threshold, based at least in part on a comparison of the estimated activity level and the average number of active processing engines.

In an example, the first logic is to increase the window length when a sum of the estimated activity level and a threshold value is less than the average number of active processing engines.

In an example, the increased window length is to cause the performance state control logic to control the performance state of the at least one of the plurality of processing engines at a reduced rate.

In an example, the first logic is to only adjust the at least one activity level threshold based at least in part on the comparison.

In an example, the first logic is to adjust a first activity level threshold and a second activity level threshold in a first direction, to cause the performance state control logic to increase the performance state of the at least one processing engine at a reduced rate.

In an example, the power state control logic is to increase the performance state of the at least one processing engine when an average core utilization value exceeds the first activity level threshold.

In an example, the first logic is to determine the estimated activity level according to a value stored in an entry of a configuration storage, the entry to be accessed based at least in part on a maximum active state residency of one of the plurality of processing engines.

In an example, the estimated activity level is less than the average number of active processing engines when the first processing engine and the second processing engine execute a parallel workload.

In an example, the estimated activity level is greater than the average number of active processing engines when the first processing engine and the second processing engine execute independent workloads.

In an example, the processor further comprises: a first counter to maintain a first count of clock cycles when the plurality of processing engines are active; and a second counter to maintain a second count of clock cycles when at least one of the plurality of processing engines is active.

In an example, the first logic is to determine the average number of active cores based at least in part on the first count and the second count.

In an example, the power state control logic is to receive a request from a user for a first performance state for the processor, and prevent operation of the processor at the first

performance state based on the comparison of the estimated activity level and the average number of active processing engines.

In an example, the power state control logic is to increase the performance state of the at least one processing engine when the processor is to execute a user interactive workload and to prevent the performance state increase when the processor is to execute a non-user interactive workload.

Note that the above processor can be implemented using various means.

In an example, the processor comprises a system on a chip (SoC) incorporated in a user equipment touch-enabled device.

In another example, a system comprises a display and a memory, and includes the processor of one or more of the above examples.

In another example, a method comprises: measuring an activity overlap between a first processing engine and a second processing engine of a processor during an evaluation interval; estimating an expected activity overlap between the first processing engine and the second processing engine; and responsive to the measured activity overlap exceeding the expected activity overlap, enabling a performance state of at least one of the first processing engine and the second processing engine to be increased.

In an example, the method further comprises responsive to the measured activity overlap being less than the expected activity overlap, preventing the performance state of the at least one of the first processing engine and the second processing engine from being increased.

In an example, the method further comprises measuring the activity overlap based at least in part on a count value of a counter, where the counter is updated when the first processing engine and the second processing engine are concurrently in an active state.

In an example, estimating the expected activity overlap comprises calculating the expected activity overlap based on a first utilization value of the first processing engine, a second utilization value of the second processing engine, and a duration of the evaluation interval.

In an example, the method further comprises preventing the first processing engine from operating at a higher performance state if the second processing engine is to be in an active state after the first processing engine completes a pending workload.

In another example, a computer readable medium including instructions is to perform the method of any of the above examples.

In another example, a computer readable medium including data is to be used by at least one machine to fabricate at least one integrated circuit to perform the method of any one of the above examples.

In another example, an apparatus comprises means for performing the method of any one of the above examples.

In another example, a system comprises: a processor including a first domain having a plurality of cores, a second domain having at least one graphics engine, and a power controller to control a performance state of the first domain based at least on part on a correlation of execution of a first workload on the first domain and execution of a second workload on the second domain; and a DRAM coupled to the processor.

In an example, the power controller is further to determine an average number of active processing engines of the first domain and the second domain over a first window, an estimated activity level of the processor for the first window, and adjust a window length at which the power controller is

to control the performance state, based at least in part on a comparison of the estimated activity level and the average number of active processing engines.

In an example, the power controller is to increase the window length when a sum of the estimated activity level and a threshold value is less than the average number of active processing engines.

In an example, the power controller is to receive a request from a user for a first performance state for the processor, and prevent operation of the processor at the first performance state based on the comparison of the estimated activity level and the average number of active processing engines.

In an example, the processor further comprises: a first counter to maintain a first count of clock cycles when the plurality of cores are active; and a second counter to maintain a second count of clock cycles when at least one of the plurality of cores is active.

In an example, the power controller is to determine the correlation by comparison of an activity overlap between the first domain and the second domain to an estimated activity overlap value.

Understand that various combinations of the above examples are possible.

Embodiments may be used in many different types of systems. For example, in one embodiment a communication device can be arranged to perform the various methods and techniques described herein. Of course, the scope of the present invention is not limited to a communication device, and instead other embodiments can be directed to other types of apparatus for processing instructions, or one or more machine readable media including instructions that in response to being executed on a computing device, cause the device to carry out one or more of the methods and techniques described herein.

Embodiments may be implemented in code and may be stored on a non-transitory storage medium having stored thereon instructions which can be used to program a system to perform the instructions. Embodiments also may be implemented in data and may be stored on a non-transitory storage medium, which if used by at least one machine, causes the at least one machine to fabricate at least one integrated circuit to perform one or more operations. The storage medium may include, but is not limited to, any type of disk including floppy disks, optical disks, solid state drives (SSDs), compact disk read-only memories (CD-ROMs), compact disk rewritables (CD-RWs), and magneto-optical disks, semiconductor devices such as read-only memories (ROMs), random access memories (RAMs) such as dynamic random access memories (DRAMs), static random access memories (SRAMs), erasable programmable read-only memories (EPROMs), flash memories, electrically erasable programmable read-only memories (EEPROMs), magnetic or optical cards, or any other type of media suitable for storing electronic instructions.

While the present invention has been described with respect to a limited number of embodiments, those skilled in the art will appreciate numerous modifications and variations therefrom. It is intended that the appended claims cover all such modifications and variations as fall within the true spirit and scope of this present invention.

What is claimed is:

1. A processor comprising:

a plurality of processing engines including a first processing engine and a second processing engine to independently execute instructions; and

a power controller including a performance state control logic to control a performance state of at least one of the plurality of processing engines, and a first logic to determine an average number of active processing engines over a first window, an estimated activity level of the processor for the first window, and adjust at least one of a window length at which the performance state control logic is to perform a performance state determination and at least one activity level threshold, based at least in part on a comparison of the estimated activity level and the average number of active processing engines,

wherein the first logic is to increase the window length when a sum of the estimated activity level and a threshold value is less than the average number of active processing engines.

2. The processor of claim 1, wherein the increased window length is to cause the performance state control logic to control the performance state of the at least one of the plurality of processing engines at a reduced rate.

3. The processor of claim 1, wherein the first logic is to only adjust the at least one activity level threshold based at least in part on the comparison.

4. The processor of claim 3, wherein the first logic is to adjust a first activity level threshold and a second activity level threshold in a first direction, to cause the performance state control logic to increase the performance state of the at least one processing engine at a reduced rate.

5. The processor of claim 4, wherein the power state control logic is to increase the performance state of the at least one processing engine when an average core utilization value exceeds the first activity level threshold.

6. The processor of claim 1, wherein the first logic is to determine the estimated activity level according to a value stored in an entry of a configuration storage, the entry to be accessed based at least in part on a maximum active state residency of one of the plurality of processing engines.

7. The processor of claim 1, wherein the estimated activity level is less than the average number of active processing engines when the first processing engine and the second processing engine execute a parallel workload.

8. The processor of claim 1, wherein the estimated activity level is greater than the average number of active processing engines when the first processing engine and the second processing engine execute independent workloads.

9. The processor of claim 1, wherein the processor further comprises:

a first counter to maintain a first count of clock cycles when the plurality of processing engines are active; and a second counter to maintain a second count of clock cycles when at least one of the plurality of processing engines is active.

10. The processor of claim 9, wherein the first logic is to determine the average number of active cores based at least in part on the first count and the second count.

11. The processor of claim 1, wherein the power state control logic is to receive a request from a user for a first performance state for the processor, and prevent operation of the processor at the first performance state based on the comparison of the estimated activity level and the average number of active processing engines.

12. The processor of claim 1, wherein the power state control logic is to increase the performance state of the at least one processing engine when the processor is to execute a user interactive workload and to prevent the performance state increase when the processor is to execute a non-user interactive workload.

31

13. A system comprising:
 a processor including a first domain having a plurality of
 cores, a second domain having at least one graphics
 engine, and a power controller to control a performance
 state of the first domain based at least in part on a
 correlation of execution of a first workload on the first
 domain and execution of a second workload on the
 second domain,

wherein the power controller is further to determine an
 average number of active processing engines of the first
 domain and the second domain over a first window, an
 estimated activity level of the processor for the first
 window, and adjust a window length at which the
 power controller is to control the performance state,
 based at least in part on a comparison of the estimated
 activity level and the average number of active pro-
 cessing engines, and increase the window length when
 a sum of the estimated activity level and a threshold
 value is less than the average number of active pro-
 cessing engines; and

a dynamic random access memory (DRAM) coupled to
 the processor.

14. The system of claim **13**, wherein the power controller
 is to receive a request from a user for a first performance
 state for the processor, and prevent operation of the proces-
 sor at the first performance state based on the comparison of
 the estimated activity level and the average number of active
 processing engines.

15. The system of claim **13**, wherein the processor further
 comprises: a first counter to maintain a first count of clock
 cycles when the plurality of cores are active; and a second
 counter to maintain a second count of clock cycles when at
 least one of the plurality of cores is active.

16. The system of claim **13**, wherein the power controller
 is to determine the correlation by comparison of an activity
 overlap between the first domain and the second domain to
 an estimated activity overlap value.

17. A processor comprising:

a plurality of processing engines including a first process-
 ing engine and a second processing engine to independ-
 ently execute instructions; and

a power controller including a performance state control
 logic to control a performance state of at least one of

32

the plurality of processing engines, and a first logic to
 determine an average number of active processing
 engines over a first window, an estimated activity level
 of the processor for the first window, and adjust at least
 one of a window length at which the performance state
 control logic is to perform a performance state deter-
 mination and at least one activity level threshold, based
 at least in part on a comparison of the estimated activity
 level and the average number of active processing
 engine, wherein the first logic is to determine the
 estimated activity level according to a value stored in
 an entry of a configuration storage, the entry to be
 accessed based at least in part on a maximum active
 state residency of one of the plurality of processing
 engines.

18. The processor of claim **17**, wherein the first logic is to
 increase the window length when a sum of the estimated
 activity level and a threshold value is less than the average
 number of active processing engines and wherein the
 increased window length is to cause the performance state
 control logic to control the performance state of the at least
 one of the plurality of processing engines at a reduced rate.

19. The processor of claim **17**, wherein the estimated
 activity level is less than the average number of active
 processing engines when the first processing engine and the
 second processing engine execute a parallel workload.

20. The processor of claim **17**, wherein the estimated
 activity level is greater than the average number of active
 processing engines when the first processing engine and the
 second processing engine execute independent workloads.

21. The processor of claim **17**, wherein the processor
 further comprises:

a first counter to maintain a first count of clock cycles
 when the plurality of processing engines are active; and
 a second counter to maintain a second count of clock
 cycles when at least one of the plurality of processing
 engines is active.

22. The processor of claim **21**, wherein the first logic is to
 determine the average number of active cores based at least
 in part on the first count and the second count.

* * * * *