



- (51) International Patent Classification:  
G06F 1/26 (2006.01) G06F 1/20 (2006.01)  
G06F 1/32 (2006.01)
- (21) International Application Number:  
PCT/US2015/038125
- (22) International Filing Date:  
26 June 2015 (26.06.2015)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
14/461,039 15 August 2014 (15.08.2014) US
- (71) Applicant: INTEL CORPORATION [US/US]; 2200 Mission College Boulevard, Santa Clara, California 95054 (US).
- (72) Inventors: ROSENZWEIG, Nir; Haemek 128 PO 96, 23800 Givat Ella, Z (IL). RAJWAN,, Doron; 13 Odem st.,

75357 Rishon Lezion, M (IL). SHAPIRA, Dorit; Hameyasdim 16a St., P.O Box 450, 30300 Atlit, HA (IL). SHULMAN, Nadav; Hasayfan 19 st., POB 447, 40600 Tel Mond, M (IL). ZIV, Tomer; #5 Pnina & Moshe st., 75357 Rishon Lezion, M (IL).

(74) Agents: EPSTEIN, Samuel I. et al.; Trop, Pruner & Hu, P.C., 1616 S. Voss Rd., Ste. 750, Houston, Texas 77057-2631 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

[Continued on next page]

(54) Title: BALANCED CONTROL OF PROCESSOR TEMPERATURE

(57) Abstract: In an embodiment, a processor includes a plurality of cores and a plurality of temperature sensors, where each core is proximate to at least one temperature sensor. The processor also includes a power control unit (PCU) including temperature logic to receive temperature data that includes a corresponding temperature value from each of the temperature sensors. Responsive to an indication that a highest temperature value of the temperature data exceeds a threshold, the temperature logic is to adjust a plurality of domain frequencies according to a determined policy that is based on instruction execution characteristics of at least two of the plurality of cores. Each domain frequency is associated with a corresponding domain that includes at least one of the plurality of cores and each domain frequency is independently adjustable. Other embodiments are described and claimed.

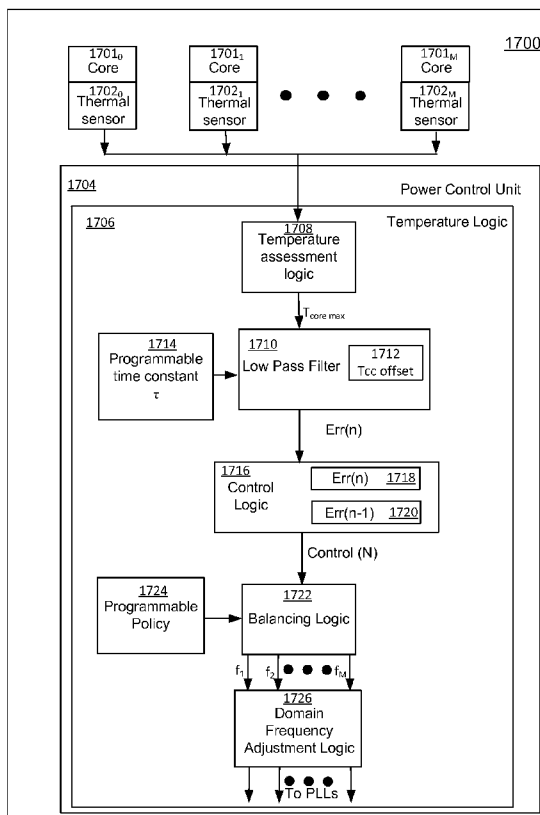


FIG. 17

WO 2016/025095 A1

**(84) Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

**Declarations under Rule 4.17:**

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

**Published:**

- *with international search report (Art. 21(3))*

## Balanced Control Of Processor Temperature

### Technical Field

[0001] Embodiments relate to balanced control of processor temperature.

### Background

[0002] 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).

### Brief Description of the Drawings

[0003] FIG. 1 is a block diagram of a portion of a system in accordance with an embodiment of the present invention.

[0004] FIG. 2 is a block diagram of a processor in accordance with an embodiment of the present invention.

[0005] FIG. 3 is a block diagram of a multi-domain processor in accordance with another embodiment of the present invention.

[0006] FIG. 4 is an embodiment of a processor including multiple cores.

[0007] FIG. 5 is a block diagram of a micro-architecture of a processor core in accordance with one embodiment of the present invention.

[0008] FIG. 6 is a block diagram of a micro-architecture of a processor core in accordance with another embodiment.

[0009] FIG. 7 is a block diagram of a micro-architecture of a processor core in accordance with yet another embodiment.

[0010] FIG. 8 is a block diagram of a micro-architecture of a processor core in accordance with a still further embodiment.

[0011] FIG. 9 is a block diagram of a processor in accordance with another embodiment of the present invention.

[0012] FIG. 10 is a block diagram of a representative SoC in accordance with an embodiment of the present invention.

[0013] FIG. 11 is a block diagram of another example SoC in accordance with an embodiment of the present invention.

[0014] FIG. 12 is a block diagram of an example system with which embodiments can be used.

[0015] FIG. 13 is a block diagram of another example system with which embodiments may be used.

[0016] FIG. 14 is a block diagram of a representative computer system.

[0017] FIG. 15 is a block diagram of a system in accordance with an embodiment of the present invention.

[0018] FIG. 16 is a block diagram of a system according to embodiments of the present invention.

[0019] FIG. 17 is a block diagram of a processor, according to an embodiment of the present invention.

[0020] FIG. 18 is a flow diagram of a method according to an embodiment of the present invention.

[0021] FIG. 19 is a block diagram of a system in accordance with an embodiment of the present invention.

#### Detailed Description

[0022] 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.

[0023] A processor such as a multi-core processor generates heat while in operation. In order for the processor to operate effectively, the heat generated can be managed. Management of processor heat can reduce overheating that can result in malfunction of the processor or reduction in efficiency.

[0024] In some embodiments, heat produced by a processor may be managed by throttling, e.g., reduction of throughput of the processor, such as by reduction of a clock frequency of a clock domain ("domain" herein) within the processor that may include an interconnect, one or more processing entities, e.g., one or more cores, a graphics processing unit (GPU), etc. One approach to processor heat management is to identify a "hot spot," e.g., a core (or GPU) within the processor that is operating at a high temperature, and to reduce the clock frequency of the domain of the hotspot. Reduction of the clock frequency (also "frequency" herein) throttles the core, e.g., reduces instruction execution rate of the core and results in a reduction of heat generated by the hot spot and a lower operating temperature of the processor.

[0025] According to embodiments of the present invention, throttling may be accomplished according to a balancing (also "balanced" herein) policy, in which the frequencies of a plurality of domains are adjusted responsive to detection of a high temperature within the processor. Throttling of more than one processing entity provides a more global solution to heat reduction in the processor, so that throttling of a particular processing unit identified to be a hot spot is less severe, e.g., a clock frequency (also domain frequency herein) of the hot spot may be reduced to a smaller degree than if it were the only clock frequency to be reduced, while other clock frequencies are also adjusted to reduce an overall heat generated by the processor.

[0026] A balancing policy may specify weighted adjustment of the clock frequencies, which may result in reduction of heat generated by the processor while lessening detrimental side effects (e.g., reduction of overall throughput of the processor) as compared with the side effects associated with reduction of only a single clock frequency to achieve the same reduction of heat generated.

[0027] The balancing policy may be determined based at least in part on execution characteristics of each of the execution units. For example, one or more cores may execute instructions rapidly (high throughput rate) and outputs of the one or more cores may be input to a graphical processing unit whose throughput rate may be less than the throughput rate of the one or more cores. The balancing policy may dictate that, responsive to detection of a hot spot, the cores are to be throttled to a greater extent than the GPU, e.g., clock frequencies of the cores are to be reduced to a greater extent than reduction of the clock frequency of the GPU so that input to the GPU from other cores does not “back up” (e.g., form a queue to be input to the GPU).

[0028] In another example, if the GPU has a faster throughput rate than any of the cores that feed the GPU, a different balancing policy may be applied. In such a case, to avoid “starvation” of the GPU the balancing policy may dictate that the clock frequency of the GPU is to be reduced by a smaller factor than reduction of the clock frequency of one or more of the cores feeding the GPU. In some embodiments, an interconnect frequency may be set approximately equal to the frequency of the cores feeding the GPU. In other embodiments, the balancing policy may indicate that the interconnect frequency is to have a specified functional relationship to another domain frequency. Yet other embodiments may include any combination of the above-described examples. Note that embodiments are not limited to those described herein.

[0029] 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 to processor 110.

[0030] As seen, processor 110 may be a single die processor including multiple cores  $120_a - 120_n$ . In addition, each core may be associated with an integrated voltage regulator (IVR)  $125_a - 125_n$  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.

[0031] Still referring to FIG. 1, additional components may be present within the processor including an input/output interface 132, another interface 134, and an integrated memory controller 136. As seen, each of these components may be powered by another integrated voltage regulator 125<sub>x</sub>. In one embodiment, interface 132 may be in accordance with the Intel® Quick Path Interconnect (QPI) protocol, 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 be in accordance with a Peripheral Component Interconnect Express (PCIe™) specification, e.g., the PCI Express™ Specification Base Specification version 2.0 (published January 17, 2007).

[0032] 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 to cause the voltage regulator to generate the appropriate regulated voltage. PCU 138 also provides control information to IVRs 125 via another digital interface 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 management power management source or system software).

[0033] 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.

[0034] Note that the power management techniques described herein may be independent of and complementary to an operating system (OS)-based mechanism, such as the Advanced Configuration and Platform Interface (ACPI) standard (e.g., Rev. 3.0b, published October 10, 2006). According to ACPI, a processor can operate at various performance states or levels, so-called P-states, namely from P0 to PN. 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 addition, according to ACPI, a processor can operate at various power states or levels. With regard to power states, ACPI specifies 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).

[0035] 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, CA, 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).

[0036] 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. Although described with these particular examples, understand that many other power management techniques may be used in particular embodiments.

[0037] 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<sub>a</sub> – 210<sub>n</sub>. 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. The various cores may be coupled via an interconnect 215 to a 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.

[0038] 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.

[0039] 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 310<sub>0</sub>-310<sub>n</sub>, 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.

[0040] 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) 340<sub>0</sub> - 340<sub>n</sub>. 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.

[0041] 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 logic to perform the power management techniques described herein.

[0042] 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  $380_0 - 380_n$  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 interfaces in accordance with an Intel® Quick Path Interconnect (QPI) protocol 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.

[0043] Referring to FIG. 4, an embodiment of a processor including multiple cores is illustrated in accordance with embodiments of the present invention. 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.

[0044] 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.

[0045] 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.

[0046] 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.

[0047] 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, ILTB 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 415, execution unit(s) 440, and portions of out-of-order unit 435 are potentially fully shared.

[0048] 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.

[0049] Core 401 further includes decode module 425 coupled to fetch unit 420 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.

[0050] 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.

[0051] 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.

[0052] Lower level data cache and data translation 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.

[0053] 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.

[0054] In the depicted configuration, processor 400 also includes bus interface module 405 and a power controller 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.

[0055] 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.

[0056] 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.

[0057] 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.

[0058] 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. Extended register file 535 may provide storage for vector-sized units, e.g., 256 or 512 bits per register.

[0059] 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.

[0060] 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.

[0061] 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.

[0062] 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 620 for storage and delivery to a given execution pipeline. A microcode ROM 618 is coupled to instruction decoder 615.

[0063] A floating point pipeline 630 includes a floating point 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.

[0064] An integer pipeline 640 also may be provided. In the embodiment shown, pipeline 640 includes an integer 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 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 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.

[0065] A memory execution scheduler 650 may schedule memory operations for execution in an address generation unit 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.

[0066] 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.

[0067] 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).

[0068] 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, CA. 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.

[0069] 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 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.

[0070] 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 815, which 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.

[0071] 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.

[0072] 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.

[0073] 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).

[0074] 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, CA, an ARM-based design from ARM Holdings, Ltd. or licensee thereof or a MIPS-based design from MIPS Technologies, Inc. of Sunnyvale, CA, 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.

[0075] In the high level view shown in FIG. 9, processor 900 includes a plurality of core units 910<sub>0</sub>-910<sub>n</sub>. 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 instructions 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 (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.

[0076] 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).

[0077] 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.

[0078] 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.

[0079] 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.

[0080] 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 according to a variety of communication protocols such as PCIe™, GPIO, USB, I<sup>2</sup>C, 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.

[0081] 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.

[0082] As seen in FIG. 10, SoC 1000 includes a first core domain 1010 having a plurality of first cores 1012<sub>0</sub> – 1012<sub>3</sub>. 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 1022<sub>0</sub> – 1022<sub>3</sub>. 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.

[0083] 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.

[0084] 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, 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).

[0085] 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.

[0086] 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 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.

[0087] 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.

[0088] 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 techniques such as Bluetooth™, IEEE 802.11, and so forth.

[0089] 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.

[0090] 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.

[0091] 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.

[0092] 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.

[0093] 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.

[0094] 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.

[0095] 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.

[0096] 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, such as according to a Bluetooth™ standard or an IEEE 802.11 standard such as IEEE 802.11 a/b/g/n can also be realized.

[0097] 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.

[0098] 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.

[0099] 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.

[0100] 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, including an IEEE 802.11 protocol, a Bluetooth™ protocol or any other wireless protocol.

[0101] 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.

[0102] Referring now to FIG. 14, a block diagram of a representative computer system such as notebook, Ultrabook™ 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 1400 is implemented as a SoC.

[0103] 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.

[0104] 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.

[0105] 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 I<sup>2</sup>C 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 I<sup>2</sup>C interconnect as touch screen 1425.

[0106] 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 I<sup>2</sup>C 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.

[0107] 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 in accordance with the Trusted Computing Group (TCG) TPM Specification Version 1.2, dated Oct. 2, 2003, may also couple to processor 1410 via this LPC interconnect.

[0108] 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.

[0109] 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 in accordance with a given IEEE 802.11 standard can be realized, while via Bluetooth unit 1452, short range communications via a Bluetooth protocol can occur. These units may communicate with processor 1410 via, e.g., a USB link or a universal asynchronous receiver transmitter (UART) link. Or these units may couple to processor 1410 via an interconnect according to a PCIe™ protocol or another such protocol such as a serial data input/output (SDIO) standard.

[0110] 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 USB protocol such as a USB 2.0 or 3.0 link, or a UART or I<sup>2</sup>C protocol.

[0111] An integrated camera module 1454 can be incorporated in the lid. 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.

[0112] 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.

[0113] 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.

[0114] 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.

[0115] FIG. 16 is a block diagram of a system 1600, according to embodiments of the present invention. The system 1600 includes a processor 1602, a power supply 1640, and a dynamic random access memory 1614.

[0116] The processor 1602 includes a plurality of cores  $1604_0 - 1604_N$ , a graphics processing unit (GPU) 1610, interconnect 1606, cache memory 1608, and power control unit (PCU) 1620. The PCU 1620 includes temperature logic 1630 to adjust two or more clock frequencies of a plurality of clock frequencies according to a programmable balancing policy 1632 and based on temperature data received from a plurality of temperature sensors  $1660_0 - 1660_N$  and 1662, where each clock frequency of the plurality of clock frequencies is generated by a corresponding phase locked loop (PLL)  $1650_0 - 1650_N$ , 1652, and 1654. In some embodiments, there are two or more temperature sensors to sense temperature of a single core, as there may be temperature variations within a core. Each of the temperature sensors may generate data, e.g., on a continual basis, and the temperature logic 1630 may receive the temperature data periodically and determine whether the clock frequencies are to be adjusted based on the temperature data received. In other embodiments (not shown) a plurality of the cores may be within the same domain (e.g., operate according to the same clock frequency determined by the same PLL), and adjustment of a particular frequency can affect throughput of the plurality of cores within the same domain.

[0117] FIG. 17 is a block diagram of a processor 1700, according to an embodiment of the present invention. The processor 1700 includes one or more execution units, e.g., cores and/or graphical processing units  $1701_0 - 1701_M$ , and a power control unit (PCU) 1704. The processor 1700 may include other components that are not shown, e.g., cache memory, uncore, etc. Thermal sensors  $1702_0 - 1702_M$  measure temperatures of the execution units. In an embodiment, each thermal sensor is located proximate to a corresponding execution unit (e.g., in physical contact with the corresponding execution unit) so as to measure a temperature of the corresponding execution unit. In other embodiments, two or more thermal sensors may be employed to measure temperatures of different portions of a single execution unit.

[0118] Temperature data from the thermal sensors 1702<sub>0</sub> – 1702<sub>M</sub> may be transmitted to the PCU 1704 and may be received by temperature logic 1706. Included in the temperature logic 1706 is temperature assessment logic 1708 that is to determine a largest temperature  $T_{\text{unit max}}$  of the temperature data received from the thermal sensors 1702<sub>0</sub> – 1702<sub>M</sub>.  $T_{\text{unit max}}$  is associated with an execution unit having a highest temperature of the plurality of execution units 1701<sub>0</sub> – 1701<sub>M</sub>.

[0119]  $T_{\text{unit max}}$  may be sent to a low pass filter 1710 that determines a temperature error signal  $\text{Err}(n)$  for an  $n^{\text{th}}$  time interval  $\Delta t$  (e.g., total elapsed time =  $n * \Delta t$ ). The low pass filter 1710 may have stored a configurable offset,  $T_{\text{cc offset}}$  1712 that is used in a calculation of the error signal  $\text{Err}(n)$ . In an embodiment, the low pass filter 1710 calculates  $\text{Err}(n+1)$  as

$$\text{Err}[n+1] = \alpha * \text{Err}[n] + (1 - \alpha) * (T_{\text{jmax}} - T_{\text{cc offset}} - T_{\text{unit max}})$$

$$\alpha = e^{-\Delta t/\tau}$$

where  $T_{\text{jmax}}$  is a maximum operable temperature of an execution unit, and  $\tau$  is a programmable time constant 1714.  $\text{Err}[0]$  typically has a value of 0.  $T_{\text{jmax}}$  represents a maximum transistor junction temperature value, which may be determined prior to operation and stored in the processor, e.g., in a fuse array.

[0120] The temperature error signal  $\text{Err}[N]$  for an  $N^{\text{th}}$  time interval may be output to control logic 1716, which determines a value of a control signal based on successive values of  $\text{Err}[n]$ . In an embodiment, the control logic includes a proportional-integral-derivative controller (PID) controller to determine the control signal.

[0121] In an embodiment, the control signal for the  $N^{\text{th}}$  time interval is calculated as:

$$\text{Control}(N) = K_p * \text{Err}[N] + K_i * (\sum \text{Err}[n]) + K_d * (\text{Err}[N] - \text{Err}[N-1]),$$

where  $\sum \text{Err}[n]$  is a summation of all values of  $\text{Err}[n]$ ,  $n=1$  to  $N$ .  $K_p$ ,  $K_i$ , and  $K_d$  are constants determined, e.g., through laboratory tests.

[0122] The control signal  $\text{Control}(N)$  is sent to balancing logic 1722 that is to determine an adjusted domain frequency for each domain, and each adjusted domain frequency can be sent to a corresponding PLL that generates the

corresponding domain frequency. The balancing logic determines each adjusted domain frequency  $f_1, f_2, \dots, f_M$  based on the value of *Control* ( $N$ ) and based on a programmable policy 1724. For example, in an embodiment an adjusted domain frequency  $f_x$  can be determined for each domain as follows:

$$f_x = (\text{Policy}_{\text{min}} + \text{Policy}_{\text{domain } X}) * \text{Control}$$

where  $\text{Policy}_{\text{min}}$  is a minimum clock frequency multiplier (determined by the policy) for all domains, and  $\text{Policy}_{\text{domain } X}$  is a preset parameter for each domain  $X$ , e.g., determined by the policy. The balancing logic 1722 sends the adjusted domain frequency values  $f_1, f_2, \dots, f_M$  to domain frequency adjustment logic 1726 that may send an adjusted domain frequency command to each corresponding phase locked loop (PLL).

[0123] Responsive to detection of a hot spot, typically the adjusted frequencies  $f_x$  of at least some of the domain frequencies are reduced from their prior values. As an example, in one policy a first plurality of cores may feed data into a GPU, and each of the first plurality of cores may have approximately the same instruction execution rate and are part of a first domain with a first domain frequency. The GPU may have a slower execution rate than each core of the first plurality of cores and the GPU is within a second domain with a second domain frequency. The policy in place may specify reduction of the first domain frequency to a greater extent than reduction of the second domain frequency, which may have a minimal effect on overall throughput while reducing heat generated by the processor, thus reducing sensed temperatures including the temperature of the hotspot.

[0124] Other policies may specify a distinct frequency multiplier ( $\text{Policy}_{\text{domain } X}$ ) for each of three or more domains e.g., each core may be in its own domain with its own operating frequency. Criteria to determine a balancing policy may include throughput rates, thermal characteristics of each domain, synergistic effects among processing units, interdependences between processing units (cores and/or GPUs), etc. In some embodiments, the balancing policy may specify a relationship between domain frequencies. For example, one balancing policy may specify that an interconnect frequency of an interconnect domain is to be at least as large as a domain frequency of a domain that includes cache memory. In another embodiment, the balancing

policy may specify that the interconnect frequency of the interconnect domain is to be at a fixed ratio of another domain frequency, e.g., a domain that includes one or more cores. Other embodiments may specify a heuristic determination of frequency ratios. Note that the balancing policy is not limited to any of the examples provided above, and may include any combination thereof or other balancing policy specifications.

[0125] Thus, using the monitored temperature of each execution unit over time, adjusted values of each domain frequency may be determined responsive to detection of one or more hotspots, e.g., execution unit with an elevated temperature. In an embodiment, two or more of the domain frequencies  $f_1, f_2, \dots, f_M$  are changed responsive to detection of a localized elevated temperature within the processor, e.g., a measured temperature that exceeds a determined value. Adjustment of several domain frequencies can reduce heat generated by the processor 1700 while avoiding severe throttling of a single hot spot. Frequency reduction of several domains may enhance overall processing performance as compared with throttling only the hot spot.

[0126] FIG. 18 is a flow diagram of a method 1800 according to an embodiment of the present invention. At block 1802, temperature logic within a power control unit (PCU) of a processor receives temperature data from a plurality of temperature sensors. The temperature sensors may monitor temperature of each of a plurality of execution units (e.g., cores, GPUs, etc.). In some embodiments, there may be one temperature sensor per core. In other embodiments, one or more of the cores may be monitored by several temperature sensors. Continuing to block 1804, a largest temperature  $T_{\text{unit max}}$  is determined from the temperature data received. Advancing to block 1806, a low pass filter is applied to the largest temperature  $T_{\text{unit max}}$ . The low pass filter may have an adjustable time constant, and the low pass filter may smooth out a response to temperature spikes. The low pass filter determines a temperature error function value  $\text{Err}(n)$  for a current time interval  $n$ .

[0127] Proceeding to block 1808, control logic within the temperature logic of the PCU determines a control function value based on the temperature error function value. Continuing to block 1810, balancing logic within the PCU may receive the

control function value and may determine domain frequency changes based on the control function value, according to a balancing policy that is programmable. The balancing policy that is programmed may be determined based on, e.g., throughput characteristics of each of the execution units, interdependences between execution units, thermal characteristics of one or more of the execution units, and other factors. Advancing to block 1812, domain frequency adjustment logic within the PCU may send adjustment commands (e.g., to PLLs) to adjust a plurality of clock frequencies according to the determined domain frequency changes that are based on the control function value and the balancing policy. The method 1800 can be repeated, as temperatures may be continually monitored and temperature information may be received from the temperature sensors by the PCU temperature logic, e.g., periodically.

[0128] Note that the adjustment of multiple domain frequencies according to a programmable policy to reduce heat generated in a processor, as described herein, may be independent of and complementary to an operating system (OS)-based mechanism, such as the Advanced Configuration and Platform Interface (ACPI) standard (e.g., Rev. 3.0b, published October 10, 2006). According to ACPI, a processor can operate at various performance states or levels, so-called P-states, namely from  $P_0$  to  $P_N$ . In general, the  $P_1$  performance state may correspond to the highest guaranteed performance state that can be requested by an OS. In addition to this  $P_1$  state, the OS can further request a higher performance state, namely a  $P_0$  state. This  $P_0$  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  $P_1$  guaranteed maximum frequency, exceeding to a maximum peak frequency of the particular processor, as fused or otherwise written into the processor during manufacture. In addition, according to ACPI, a processor can operate at various power states or levels. With regard to power states, ACPI specifies different power consumption states, generally referred to as C-states,  $C_0$ ,  $C_1$  to  $C_n$  states. When a core is active, it runs at a  $C_0$  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.,

C<sub>1</sub>-C<sub>6</sub> states), with each C-state being at a lower power consumption level (such that C<sub>6</sub> is a deeper low power state than C<sub>1</sub>, and so forth).

[0129] Embodiments can be incorporated into other types of systems including mobile devices such as a cellular telephone. Referring now to FIG. 19, shown is a block diagram of a system in accordance with another embodiment of the present invention. As shown in FIG. 19, system 1900 may be a mobile device and may include various components. As shown in the high level view of FIG. 19, an applications processor 1910, which may be a central processing unit of the device, is in communication with various components, including a storage 1915. Storage 1915, in various embodiments, may include both program and data storage portions.

[0130] Applications processor 1910 may further be coupled to an input/output system 1920, which in various embodiments may include a display and one or more input devices such as a touch keypad, which itself can appear on the display when processed. The system 1900 may also include an integrated sensor hub (ISH) 1960 that may receive data from one or more sensors 1970. The applications processor 1910 may include one or more cores and may optionally include a graphics processing unit. The applications processor 1910 includes a power control unit (PCU) 1912 that includes temperature logic 1914 to, responsive to an indication of a hot spot within the applications processor 1910, adjust a plurality of domain frequencies of the processor 1910 according to a programmable balancing policy, in accordance with embodiments of the present invention.

[0131] Applications processor 1910 also may couple to baseband processor 1930, which may condition signals such as voice and data communications for output, as well as conditioning incoming telephone and other signals. As seen, baseband processor 1930 couples to transceiver 1940, which may enable both receive and transmit capabilities. In turn, transceiver 1940 may be in communication with an antenna 1950, e.g., any type of antenna capable of transmitting and receiving voice and data signals via one or more communication protocols such as via a wireless wide area network (e.g., a 3G or 4G network) and/or a wireless local area network, such as a BLUETOOTH™ or so-called WI-FI™ network in accordance with an Institute of Electrical and Electronics Engineers 802.11 standard. As seen, system

1900 may further include a rechargeable power supply 1925 having a rechargeable battery to enable operation in a mobile environment. While shown with this particular implementation in the embodiment of FIG. 19, the scope of the present invention is not limited in this regard.

[0132] The following examples pertain to further embodiments.

[0133] In a first example, a processor includes a plurality of cores, a plurality of temperature sensors, where each core is proximate to at least one temperature sensor, and a power control unit (PCU). The PCU includes temperature logic to receive temperature data including a corresponding temperature value from each of the temperature sensors. Responsive to an indication that a highest temperature value of the temperature data exceeds a threshold, the PCU is to adjust a plurality of domain frequencies according to a determined policy that is based on a comparison of instruction execution characteristics of at least two of the plurality of cores, where each domain frequency is associated with a corresponding domain that includes at least one of the plurality of cores and each domain frequency is independently adjustable.

[0134] A second example includes the features of the first example. Additionally, the temperature logic includes temperature assessment logic to identify the highest temperature value of the temperature values received from the temperature sensors.

[0135] A third example includes the features of the second example. The temperature logic further includes a low pass filter to determine a temperature error signal based on a comparison of the highest temperature value to a temperature offset value.

[0136] A fourth example includes the features of the third example, and additionally the temperature logic includes control logic to determine a control value based at least in part on the temperature error signal. Adjustment of the domain frequencies is based at least in part upon the control value.

[0137] A fifth example includes the features of the fourth example. Additionally, the control logic includes proportional integral derivative (PID) logic to calculate the

control value according to a PID calculation that is based upon the temperature error signal.

[0138] A sixth example includes the features of the fourth example. The temperature logic further includes balancing logic to determine each of a plurality of adjusted domain frequencies by multiplication of the control value by a corresponding weight factor of the determined policy.

[0139] A seventh example includes the features of any one of the first through the sixth examples. Additionally, the determined policy includes a plurality of weight factors, each domain having a corresponding weight factor. The weight factor of each domain is based at least in part on a corresponding instruction throughput rate of at least one core in the domain.

[0140] In an eighth example, a processor includes a plurality of execution units and a plurality of temperature sensors, each temperature sensor to provide a corresponding temperature value of a corresponding location within the processor. The processor also includes a power control unit (PCU) including temperature logic to receive the corresponding temperature value from each of the temperature sensors. Responsive to an indication that a particular temperature value received from one of the temperature sensors exceeds a threshold, the PCU is to adjust at least two domain frequencies according to a determined policy based at least in part on instruction execution characteristics of at least two of the plurality of execution units, where each domain frequency is associated with at least one of the execution units.

[0141] In a ninth example that includes the elements of the eighth example, adjustment of the at least two domain frequencies includes reduction of at least one of the at least two domain frequencies responsive to the indication.

[0142] In a tenth example that includes the elements of the eighth example, the determined policy is programmable.

[0143] In an 11<sup>th</sup> example that includes the elements of the eighth example, the temperature logic includes domain frequency adjustment logic to adjust the at least two domain frequencies according to a temperature error function that is based at

least in part on a difference between the particular temperature value and the threshold.

[0144] In a 12<sup>th</sup> example that includes the elements of the 11<sup>th</sup> example, the particular temperature value is a largest of the temperature values received from the plurality of temperature sensors.

[0145] In a 13<sup>th</sup> example that includes the elements of the 11<sup>th</sup> example, the temperature logic includes a low pass filter to determine a value of the temperature error function. The threshold is associated with a maximum operating temperature of the plurality of execution units.

[0146] In a 14<sup>th</sup> example that includes the elements of the 11<sup>th</sup> example, the temperature error function value is further based on a selectable time constant.

[0147] In a 15<sup>th</sup> example that includes the elements of any one of the 8<sup>th</sup> through the 14<sup>th</sup> examples, the determined policy includes a plurality of weight factors, each domain frequency associated with a respective weight factor. Each of the at least two domain frequencies is to be adjusted at least in part according to the respective weight factor responsive to the indication that the particular temperature value exceeds the threshold.

[0148] A 16<sup>th</sup> example is a computer readable medium storing executable instructions that when executed by a machine, cause the machine to monitor a plurality of temperatures, where each temperature is associated with a corresponding instruction execution unit within a processor that includes a plurality of domains, each domain having a corresponding domain frequency and each domain including at least one instruction execution unit, and to adjust a plurality of the domain frequencies based on the plurality of monitored temperatures and according to a balancing policy that is based at least in part on instruction throughput characteristics of the instruction execution units.

[0149] In a 17<sup>th</sup> example that includes the elements of the 16<sup>th</sup> example, the balancing policy is programmable.

[0150] An 18<sup>th</sup> example that includes the elements of the 16<sup>th</sup> example further includes instructions to determine a highest temperature of the plurality of

temperatures, to apply a low pass filter to the highest temperature to determine a temperature error value based on the highest temperature and upon a programmable time constant, and to determine a control value based on the temperature error value, and determine adjusted frequencies based on the control value.

[0151] In a 19<sup>th</sup> example that includes the elements of the 18<sup>th</sup> example, the temperature error value is determined further based on a temperature offset.

[0152] In a 20<sup>th</sup> example that includes the elements of any one of the 16<sup>th</sup> to the 19<sup>th</sup> examples, the balancing policy includes a plurality of weight factors, each weight factor associated with a corresponding domain. Additionally, the computer readable medium includes instructions to determine the adjusted frequency of each domain by multiplication of the control value by the corresponding weight factor of the corresponding domain.

[0153] In a 21<sup>th</sup> example, a method includes monitoring a plurality of temperatures, where each temperature is associated with a corresponding instruction execution unit within a processor that includes a plurality of domains, each domain having a corresponding domain frequency and each domain including at least one instruction execution unit. The method further includes adjusting a plurality of the domain frequencies based on the plurality of monitored temperatures and according to a balancing policy that is based at least in part on instruction throughput characteristics of the instruction execution units.

[0154] In a 22<sup>th</sup> example that includes the elements of the 21<sup>th</sup> example, the balancing policy is programmable.

[0155] In a 23<sup>th</sup> example that includes the elements of the 22<sup>th</sup> example, the method includes determining a highest temperature of the plurality of temperatures, applying a low pass filter to the highest temperature to determine a temperature error value based on the highest temperature and upon a programmable time constant, and determine a control value based on the temperature error value, and determining adjusted frequencies based on the control value.

[0156] In a 24<sup>th</sup> example that includes the elements of the 23<sup>th</sup> example, the temperature error value is determined further based on a temperature offset.

[0157] In a 25<sup>th</sup> example that includes the elements of the 21<sup>th</sup> example, the balancing policy includes a plurality of weight factors, each weight factor associated with a corresponding domain, and the method includes determining the adjusted frequency of each domain by multiplication of the control value by the corresponding weight factor of the corresponding domain.

[0158] A 26<sup>th</sup> example includes apparatus to perform the method of any of the 21<sup>st</sup> to the 25<sup>th</sup> examples.

[0159] A 27<sup>th</sup> example includes apparatus including means for performing the method of any of the 21<sup>st</sup> to the 25<sup>th</sup> examples.

[0160] Although the preceding 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, 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, Internet protocol devices, digital cameras, personal digital assistants (PDAs), and handheld PCs. Embedded applications typically include a microcontroller, a digital signal processor (DSP), network computers (NetPC), set-top boxes, network hubs, wide area network (WAN) switches, or any other system that can perform the functions and operations taught below. 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. 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.

[0161] 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.

[0162] 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. 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.

[0163] 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 1. A processor comprising:  
2 a plurality of cores;  
3 a plurality of temperature sensors, wherein each core is proximate to at least  
4 one temperature sensor; and  
5 a power control unit (PCU) including temperature logic to receive temperature  
6 data including a corresponding temperature value from each of the temperature  
7 sensors, and responsive to an indication that a highest temperature value of the  
8 temperature data exceeds a threshold, to adjust a plurality of domain frequencies  
9 according to a determined policy that is based on a comparison of instruction  
10 execution characteristics of at least two of the plurality of cores, wherein each  
11 domain frequency is associated with a corresponding domain that includes at least of  
12 one of the plurality of cores and each domain frequency is independently adjustable.
  
- 1 2. The processor of claim 1, wherein the temperature logic includes temperature  
2 assessment logic to identify the highest temperature value of the temperature values  
3 received from the temperature sensors.
  
- 1 3. The processor of claim 2, wherein the temperature logic further includes a low  
2 pass filter to determine a temperature error signal based on a comparison of the  
3 highest temperature value to a temperature offset value.
  
- 1 4. The processor of claim 3, wherein the temperature logic includes control logic  
2 to determine a control value, based at least in part on the temperature error signal,  
3 wherein adjustment of the domain frequencies is based at least in part upon the  
4 control value.
  
- 1 5. The processor of claim 4, wherein the control logic includes proportional  
2 integral derivative (PID) logic to calculate the control value according to a PID  
3 calculation that is based upon the temperature error signal.

1 6. The processor of claim 4, wherein the temperature logic further comprises  
2 balancing logic to determine each of a plurality of adjusted domain frequencies by  
3 multiplication of the control value by a corresponding weight factor of the determined  
4 policy.

1 7. The processor of any one of claims 1 to 6, wherein the determined policy  
2 includes a plurality of weight factors, each domain having a corresponding weight  
3 factor, wherein the weight factor of each domain is based at least in part on a  
4 corresponding instruction throughput rate of at least one core in the domain.

1 8. A processor comprising:  
2 a plurality of execution units;  
3 a plurality of temperature sensors, each temperature sensor to provide a  
4 corresponding temperature value of a corresponding location within the processor;  
5 and  
6 a power control unit (PCU) including temperature logic to receive the  
7 corresponding temperature value from each of the temperature sensors, and  
8 responsive to  
9 an indication that a particular temperature value received from one of the  
10 temperature sensors exceeds a threshold, adjust at least two domain frequencies  
11 according to a determined policy based at least in part on instruction execution  
12 characteristics of at least two of the plurality of execution units, wherein each domain  
13 frequency is associated with at least one of the execution units.

1 9. The processor of claim 8, wherein adjustment of the at least two domain  
2 frequencies comprises reduction of at least one of the at least two domain  
3 frequencies responsive to the indication.

1 10. The processor of claim 8, wherein the temperature logic includes domain  
2 frequency adjustment logic to adjust the at least two domain frequencies according  
3 to a temperature error function that is based at least in part on a difference between  
4 the particular temperature value and the threshold.

1 11. The processor of claim 10, wherein the particular temperature value is a  
2 largest of the temperature values received from the plurality of temperature sensors.

1 12. The processor of claim 10, wherein the temperature logic includes a low pass  
2 filter to determine a value of the temperature error function, wherein the threshold is  
3 associated with a maximum operating temperature of the plurality of execution units.

1 13. The processor of claim 10, wherein the temperature error function value is  
2 further based on a selectable time constant.

1 14. The processor of any one of claims 8 to 13, wherein the determined policy  
2 includes a plurality of weight factors, each domain frequency associated with a  
3 respective weight factor, wherein each of the at least two domain frequencies is to  
4 be adjusted at least in part according to the respective weight factor responsive to  
5 the indication that the particular temperature value exceeds the threshold.

1 15. A computer readable medium storing executable instructions that when  
2 executed by a machine, cause the machine to:  
3 monitor a plurality of temperatures, wherein each temperature is associated  
4 with a corresponding instruction execution unit within a processor that includes a  
5 plurality of domains, each domain having a corresponding domain frequency and  
6 each domain including at least one instruction execution unit; and  
7 adjust a plurality of the domain frequencies based on the plurality of  
8 monitored temperatures and according to a balancing policy that is based at least in  
9 part on instruction throughput characteristics of the instruction execution units.

1 16. The computer readable medium of claim 15, wherein the balancing policy is  
2 programmable.

1 17. The computer readable medium of claim 15, further including instructions to:  
2 determine a highest temperature of the plurality of temperatures;

3           apply a low pass filter to the highest temperature to determine a temperature  
4 error value based on the highest temperature and upon a programmable time  
5 constant, and determine a control value based on the temperature error value; and  
6           determine adjusted frequencies based on the control value.

1 18.   The computer readable medium of claim 17, wherein the temperature error  
2 value is determined further based on a temperature offset.

1 19.   The computer readable medium of any one of claims 16 to 18, wherein the  
2 balancing policy includes a plurality of weight factors, each weight factor associated  
3 with a corresponding domain, and wherein the computer readable medium further  
4 includes instructions to determine the adjusted frequency of each domain by  
5 multiplication of the control value by the corresponding weight factor of the  
6 corresponding domain.

1 20.   A method comprising:  
2           monitoring a plurality of temperatures, wherein each temperature is  
3 associated with a corresponding instruction execution unit within a processor that  
4 includes a plurality of domains, each domain having a corresponding domain  
5 frequency and each domain including at least one instruction execution unit; and  
6           adjusting a plurality of the domain frequencies based on the plurality of  
7 monitored temperatures and according to a balancing policy that is based at least in  
8 part on instruction throughput characteristics of the instruction execution units.

1 21.   The method of claim 20, wherein the balancing policy is programmable.

1 22.   The method of claim 20, further comprising:  
2           determining a highest temperature of the plurality of temperatures;  
3 applying a low pass filter to the highest temperature to determine a temperature error  
4 value based on the highest temperature and upon a programmable time constant,  
5 and  
6           determining a control value based on the temperature error value; and

7           determining adjusted frequencies based on the control value.

1   23.    The method of claim 22, wherein the temperature error value is determined  
2   further based on a temperature offset.

1   24.    The method of claim 20, wherein the balancing policy includes a plurality of  
2   weight factors, each weight factor associated with a corresponding domain, and  
3   wherein the method includes determining the adjusted frequency of each domain by  
4   multiplication of the control value by the corresponding weight factor of the  
5   corresponding domain.

1   25.    Apparatus comprising means for performing the method of any one of claims  
2   20 to 24.

100

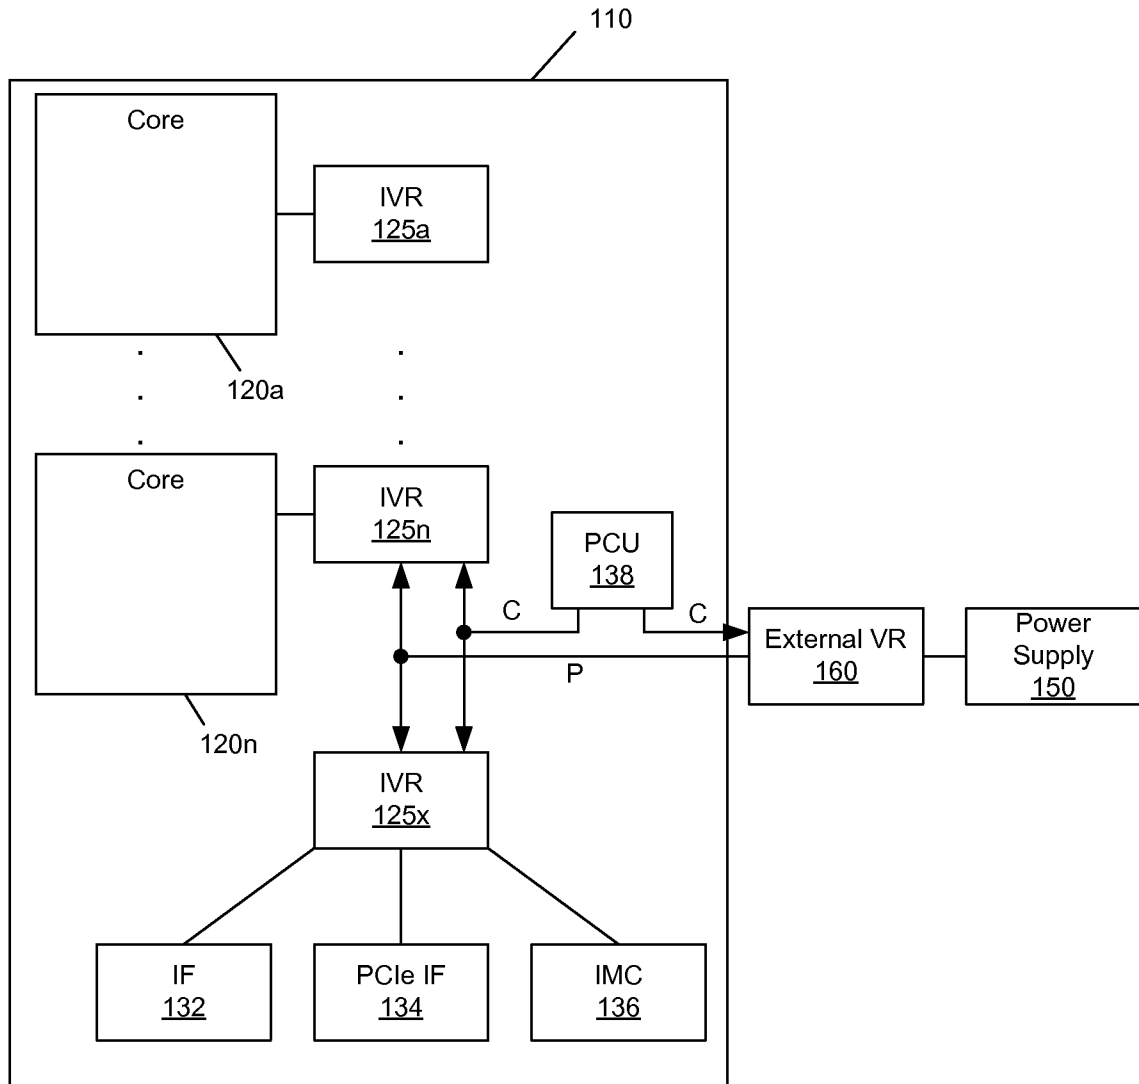


FIG. 1

200

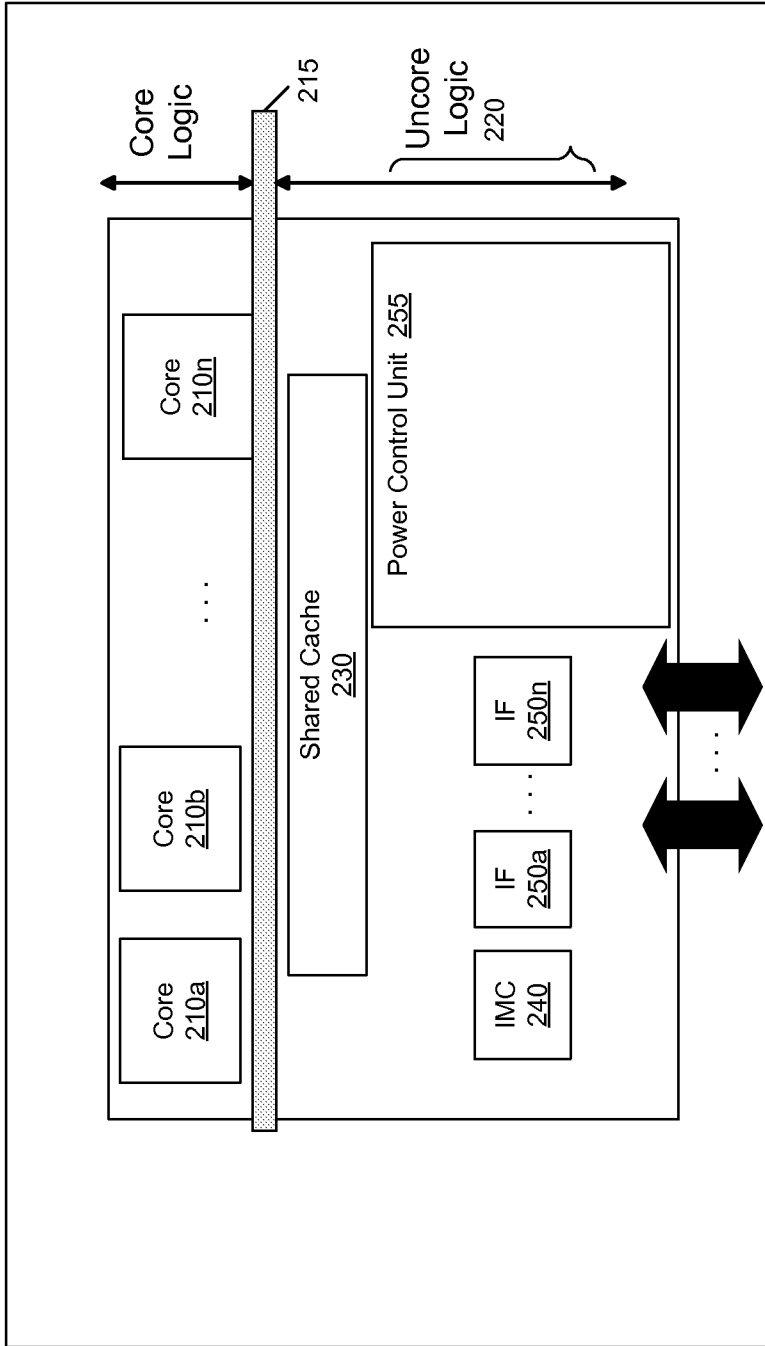


FIG. 2

300

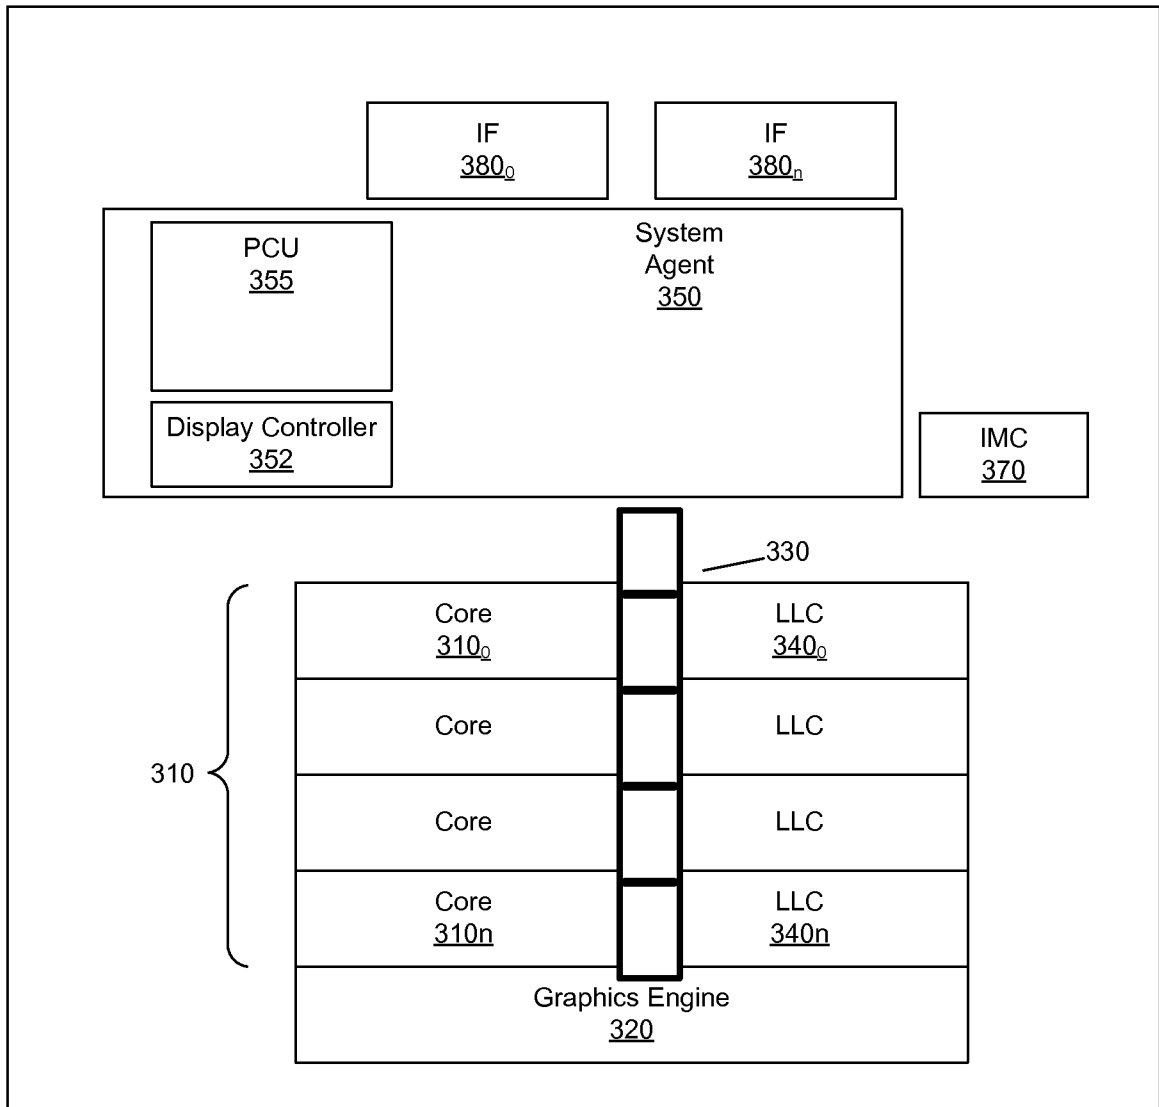
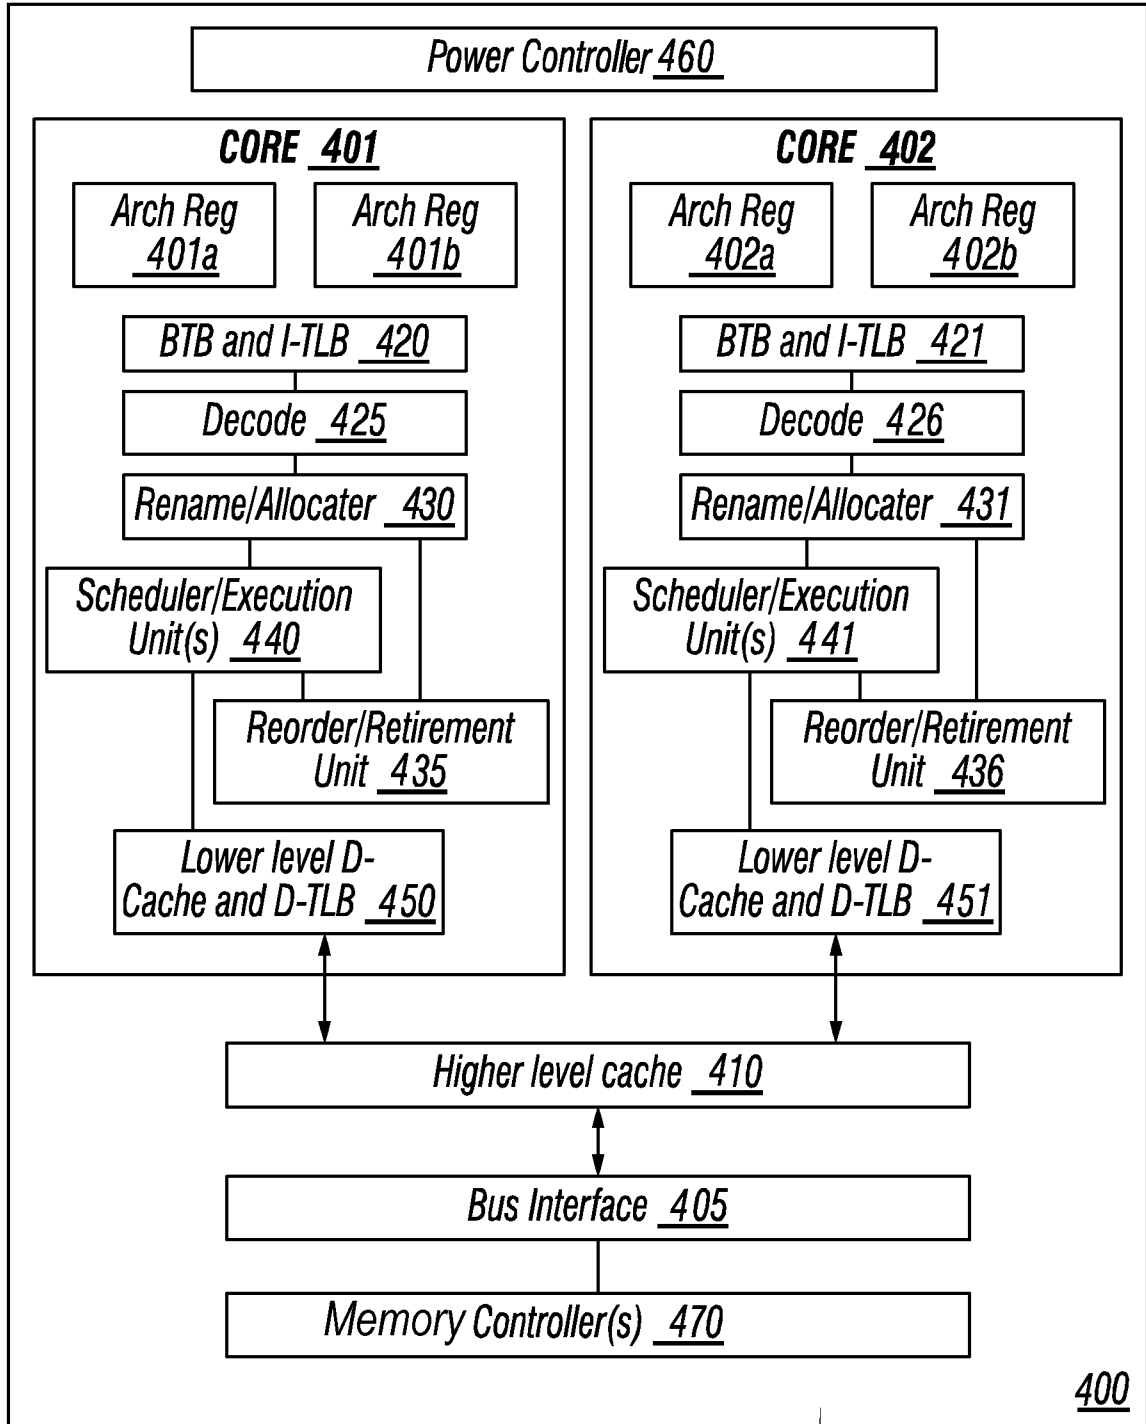


FIG. 3



400

FIG. 4

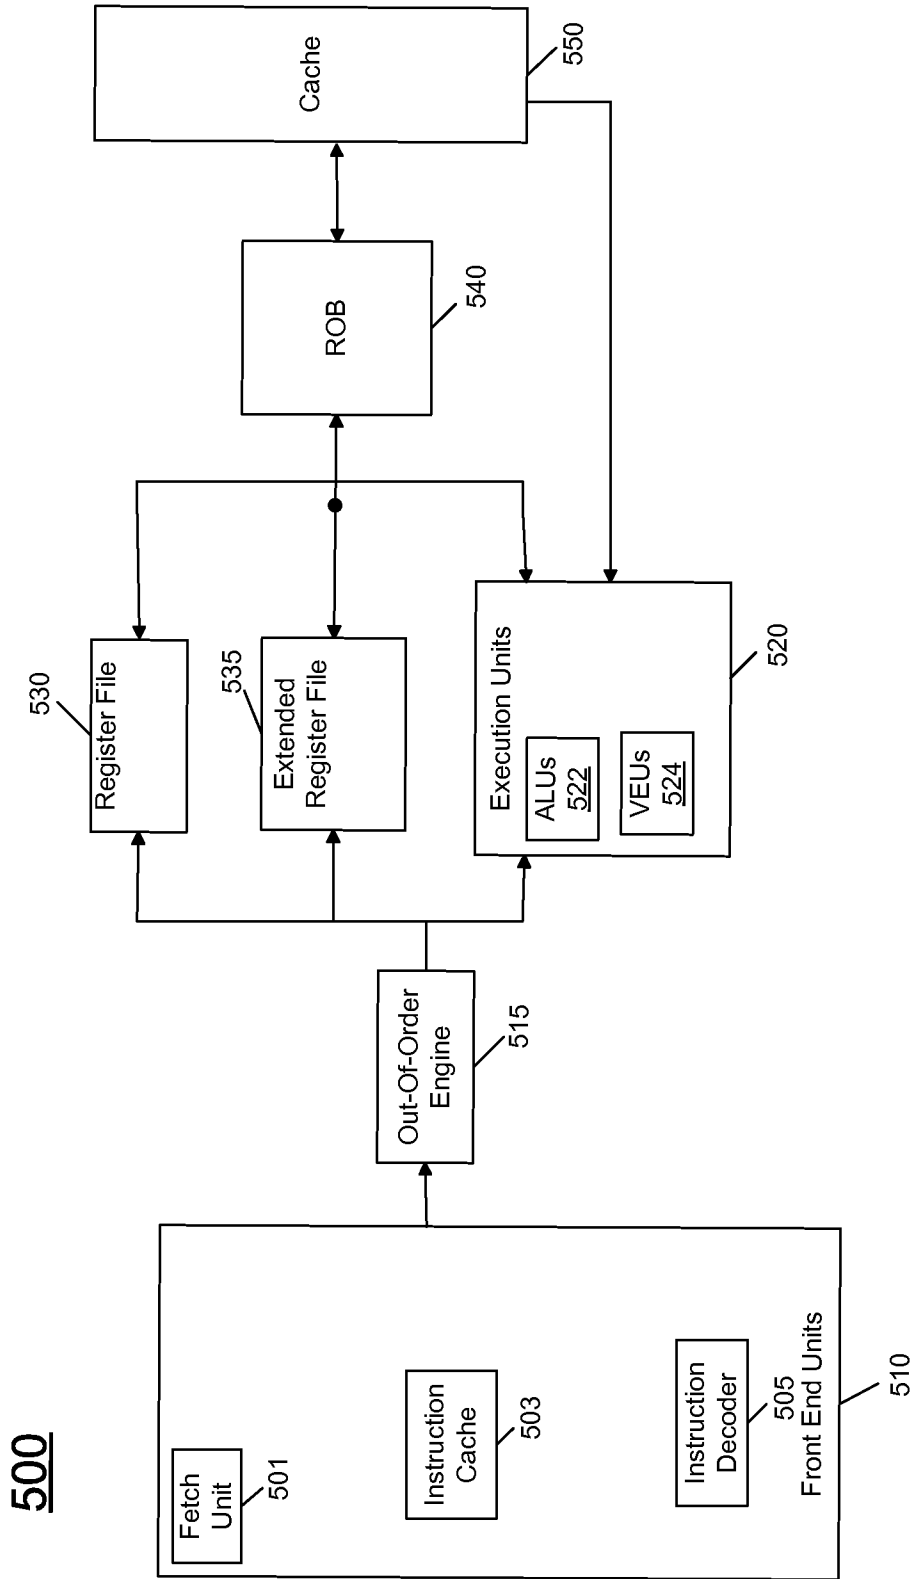


FIG. 5

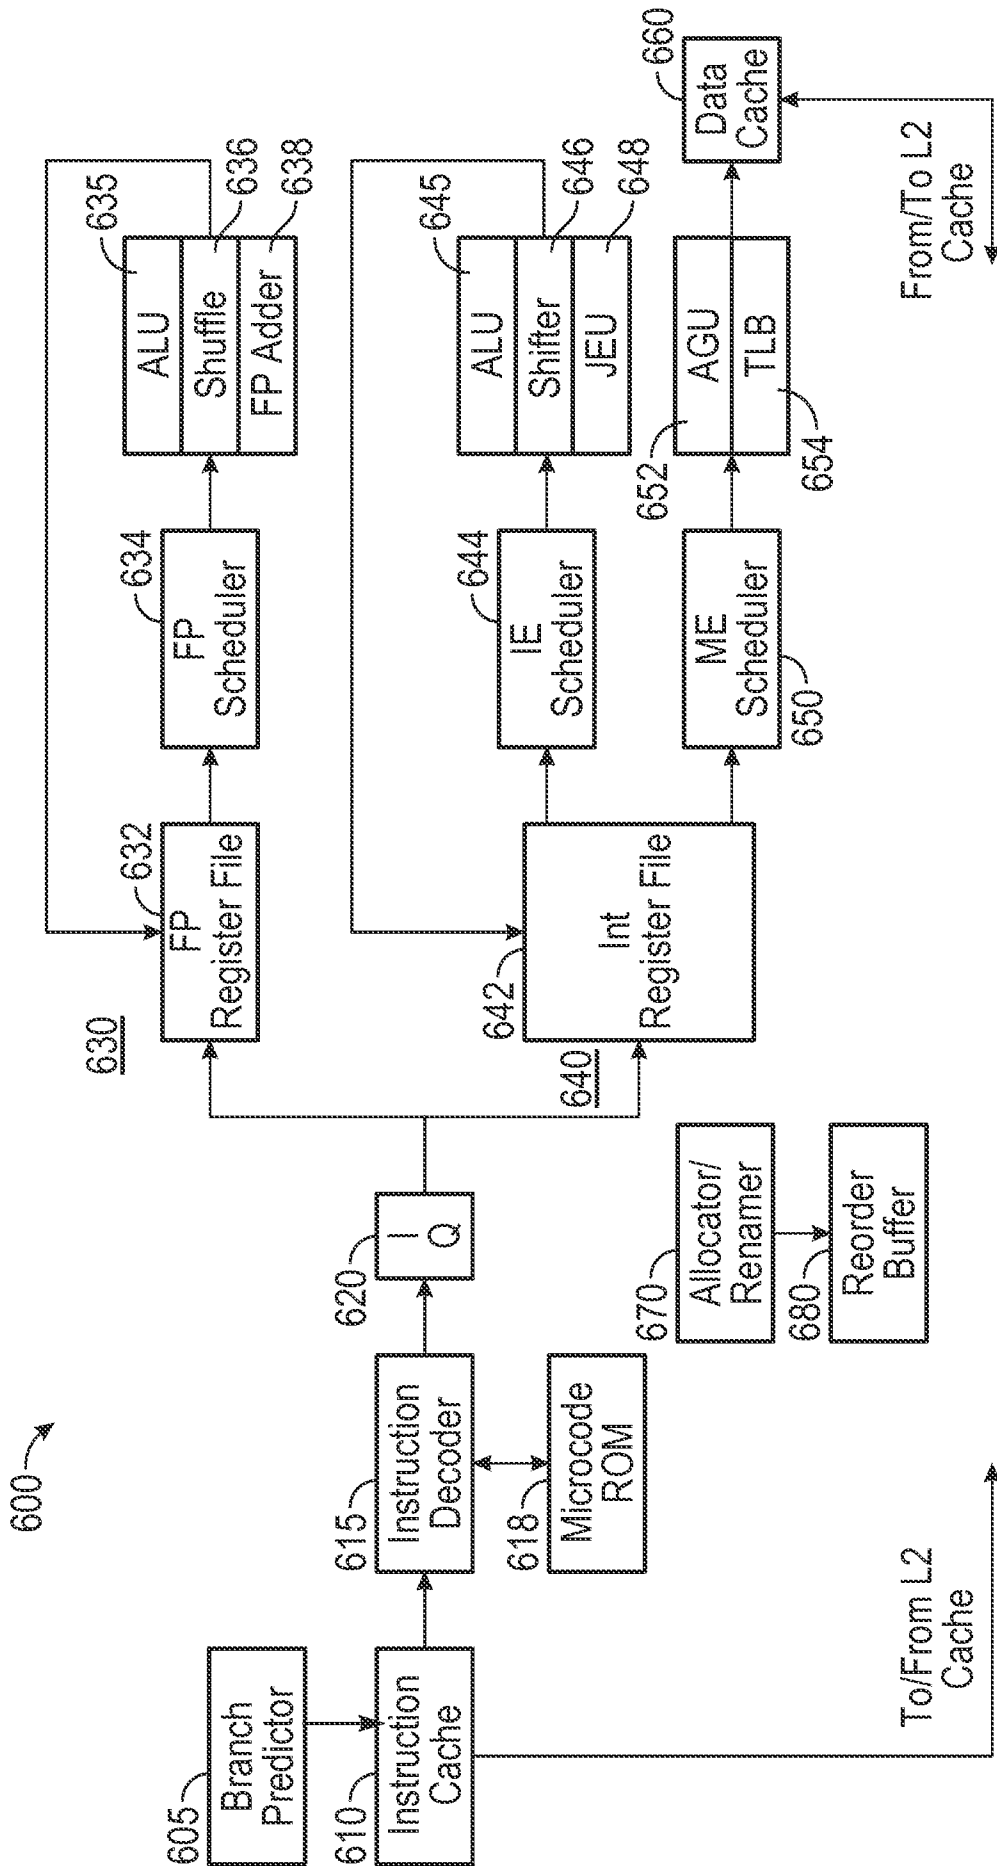


FIG. 6

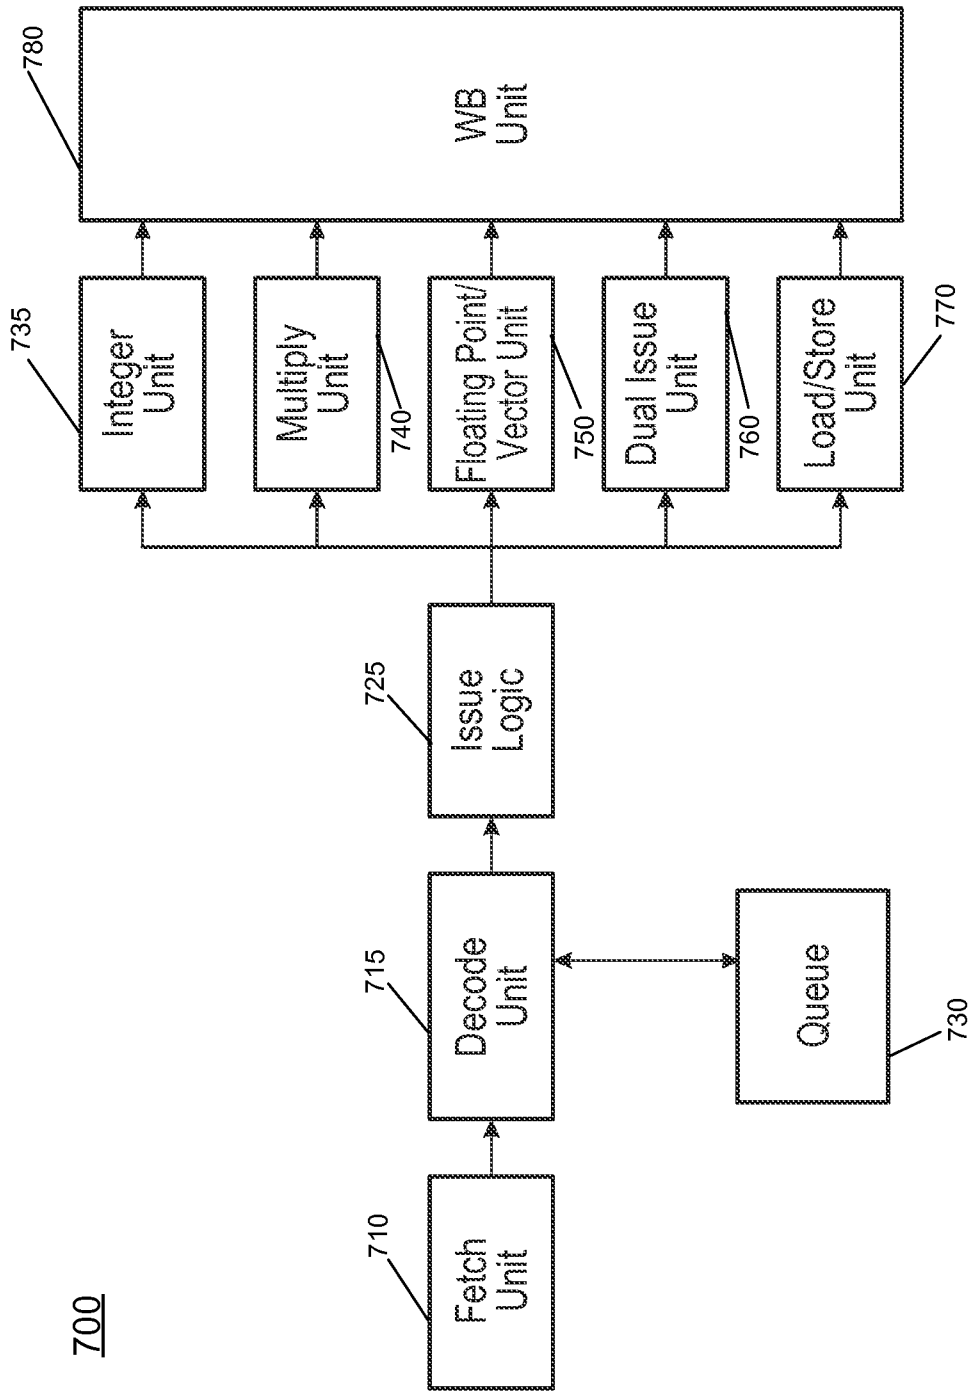


FIG. 7

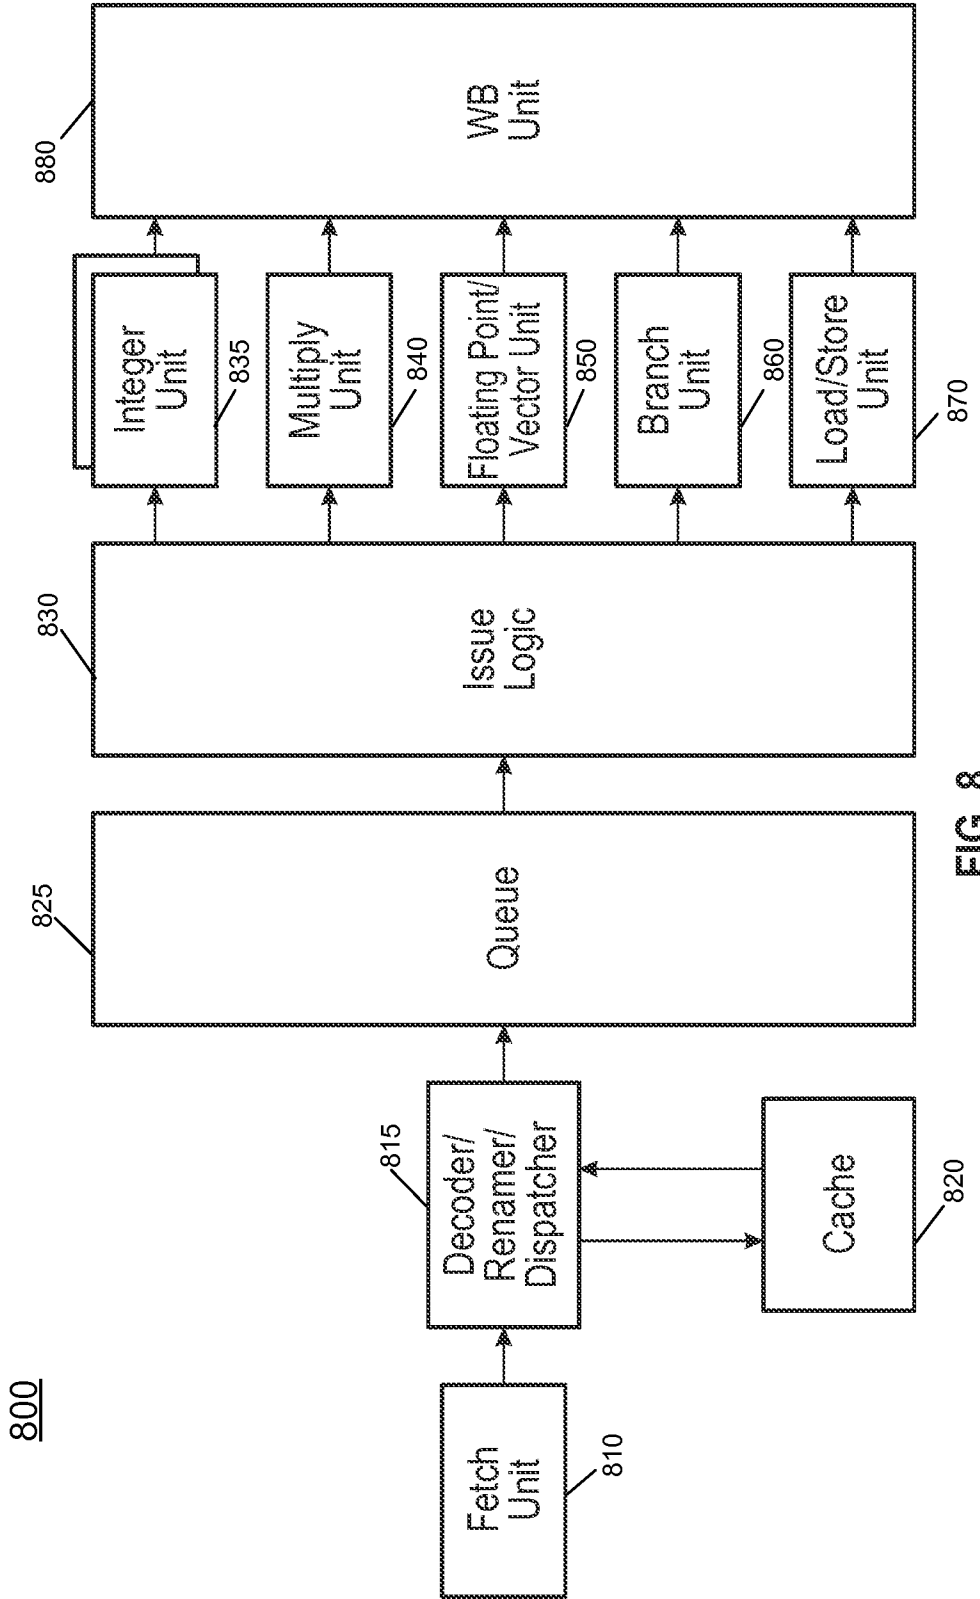


FIG. 8

900

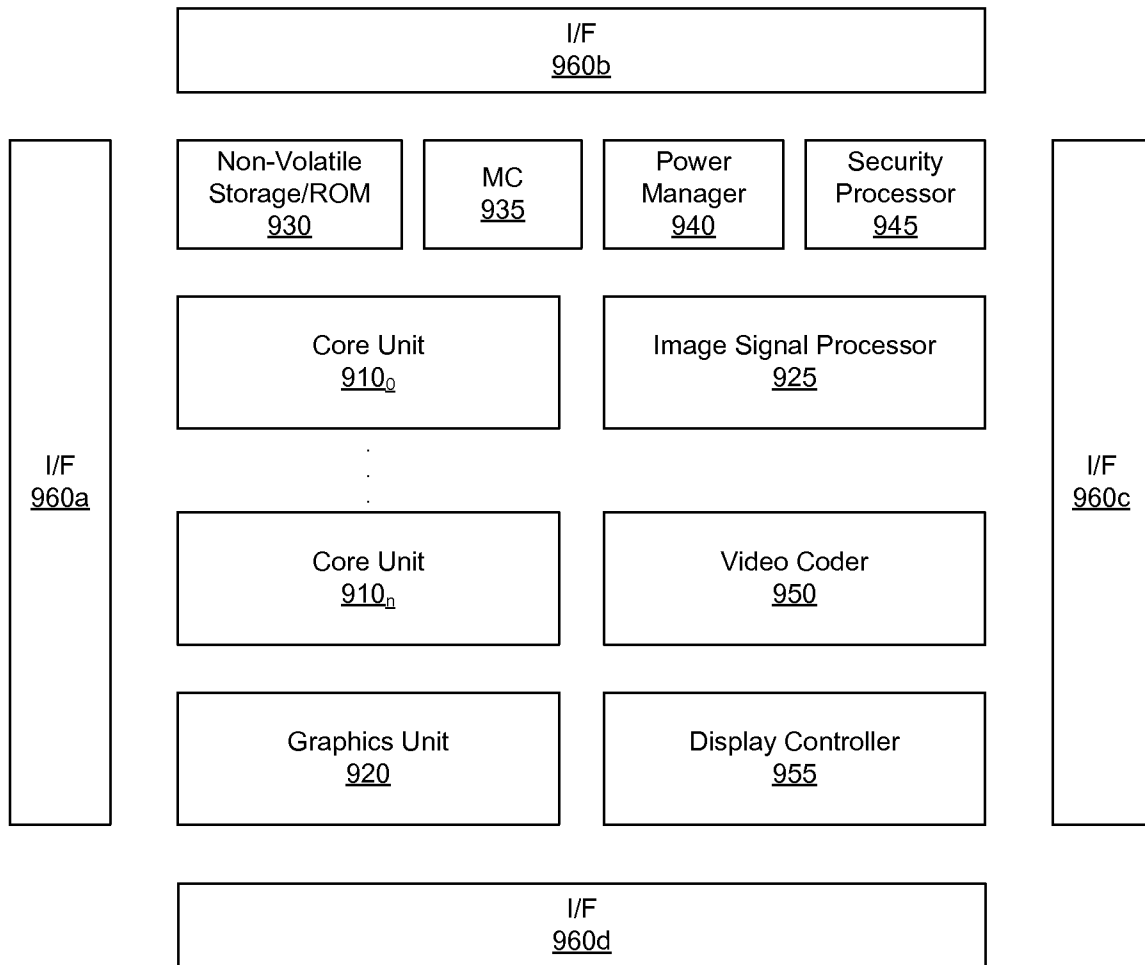


FIG. 9

1000

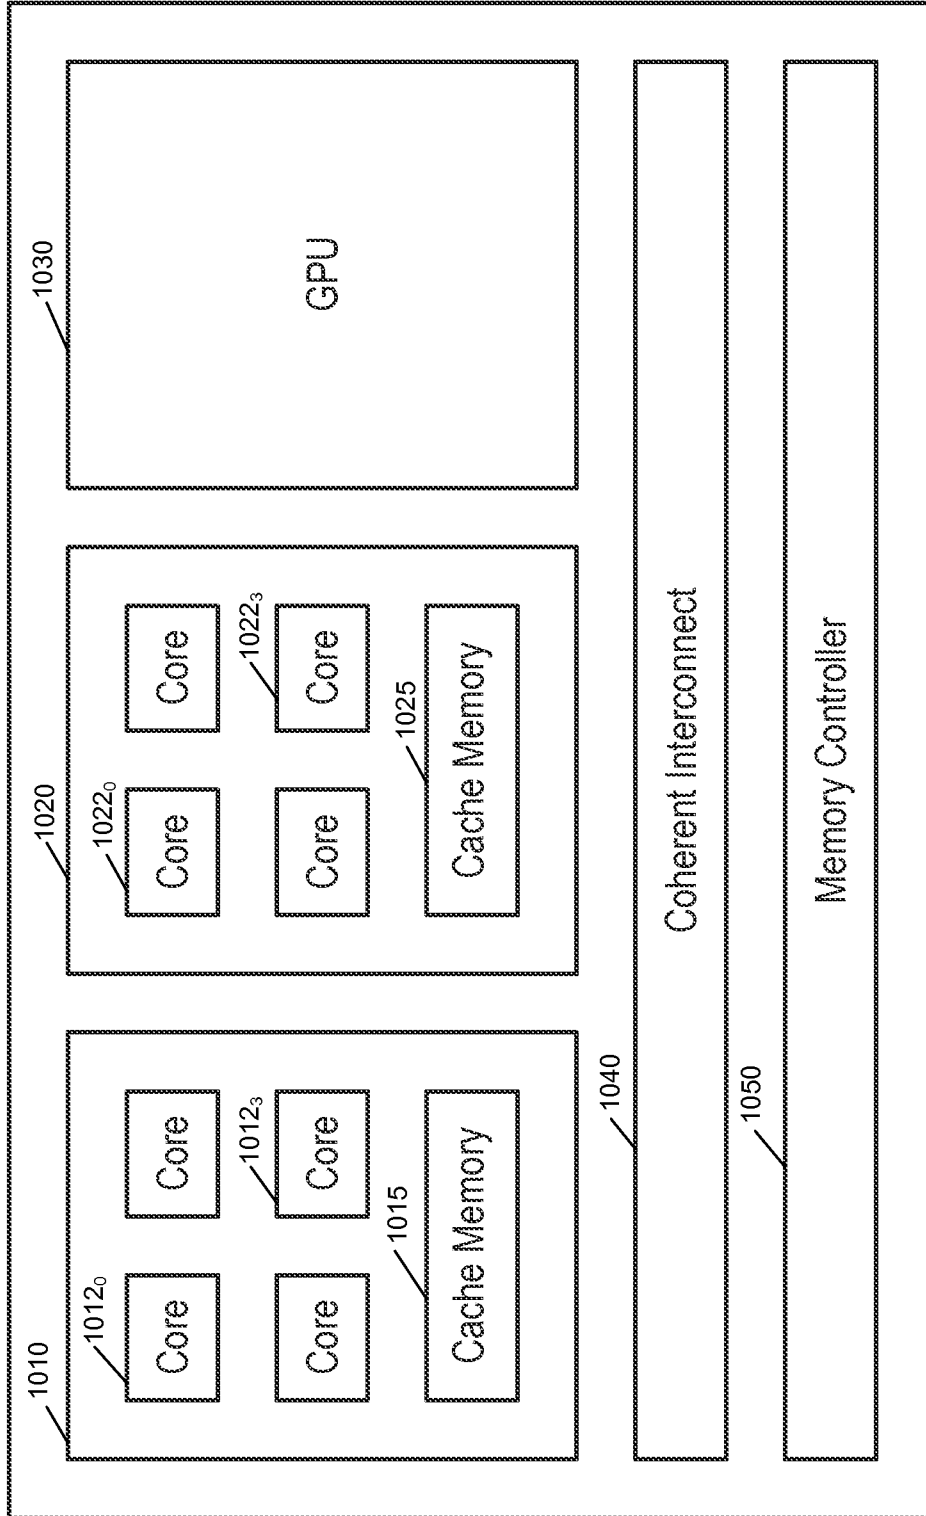


FIG. 10

1100

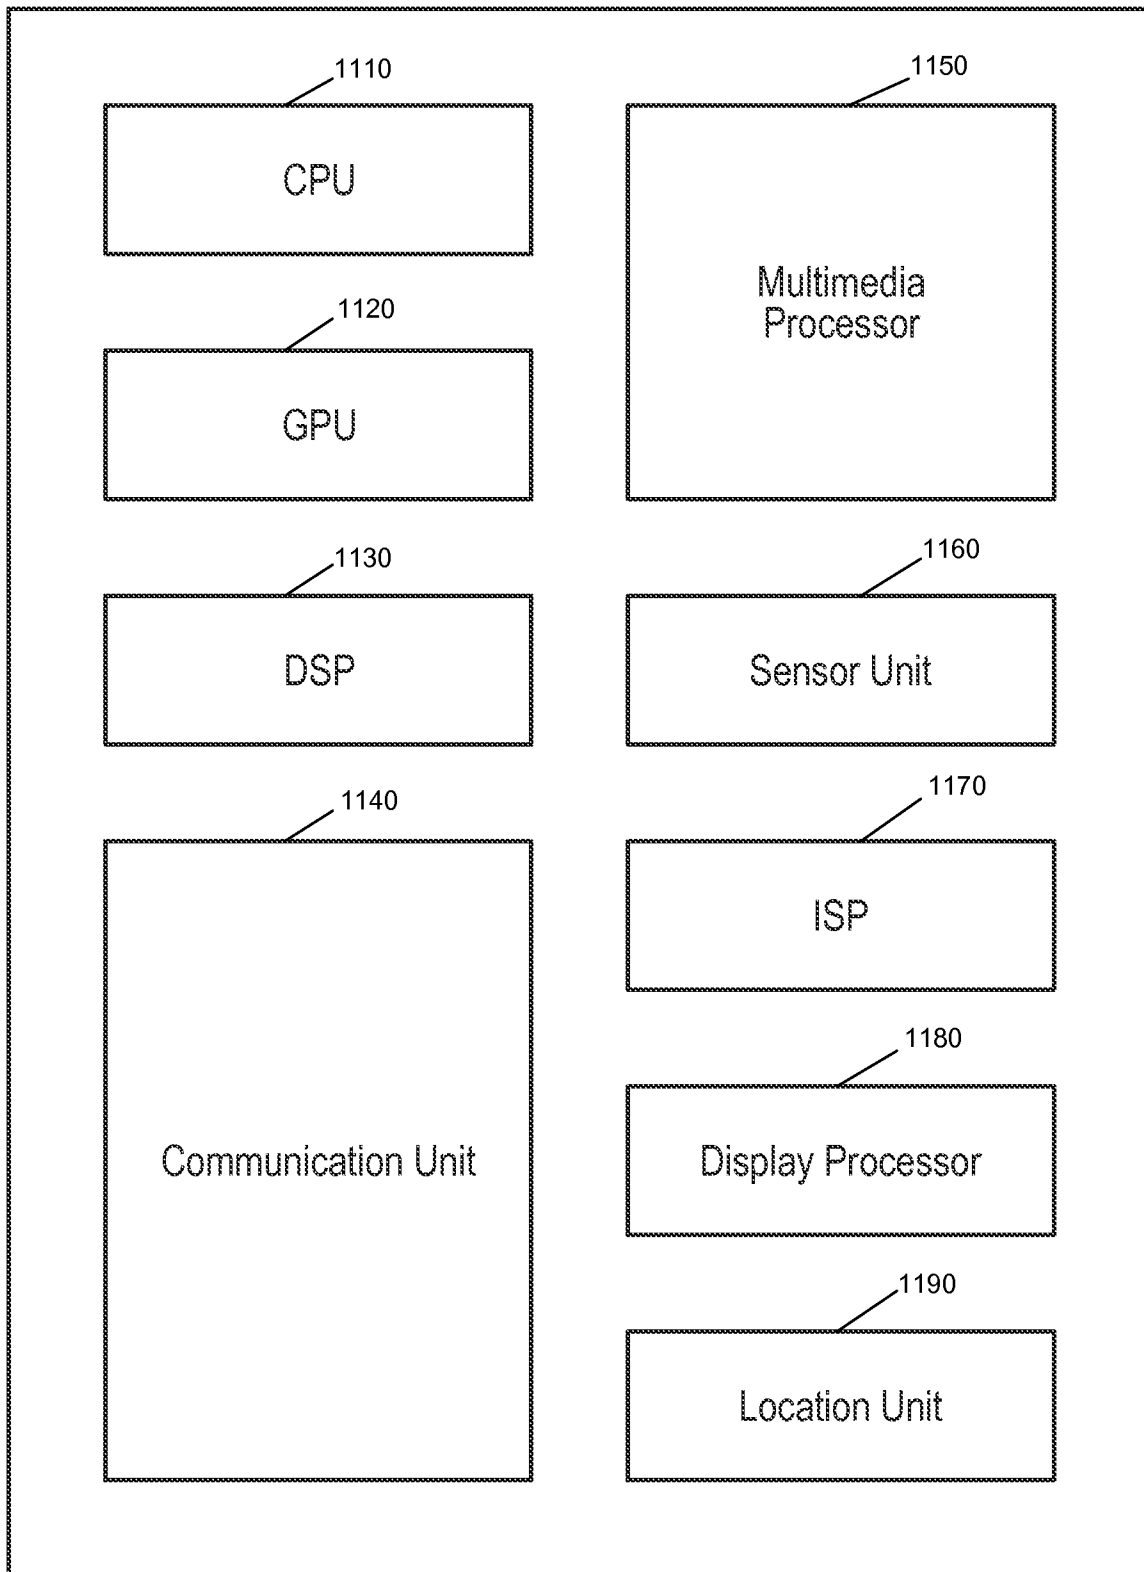


FIG. 11

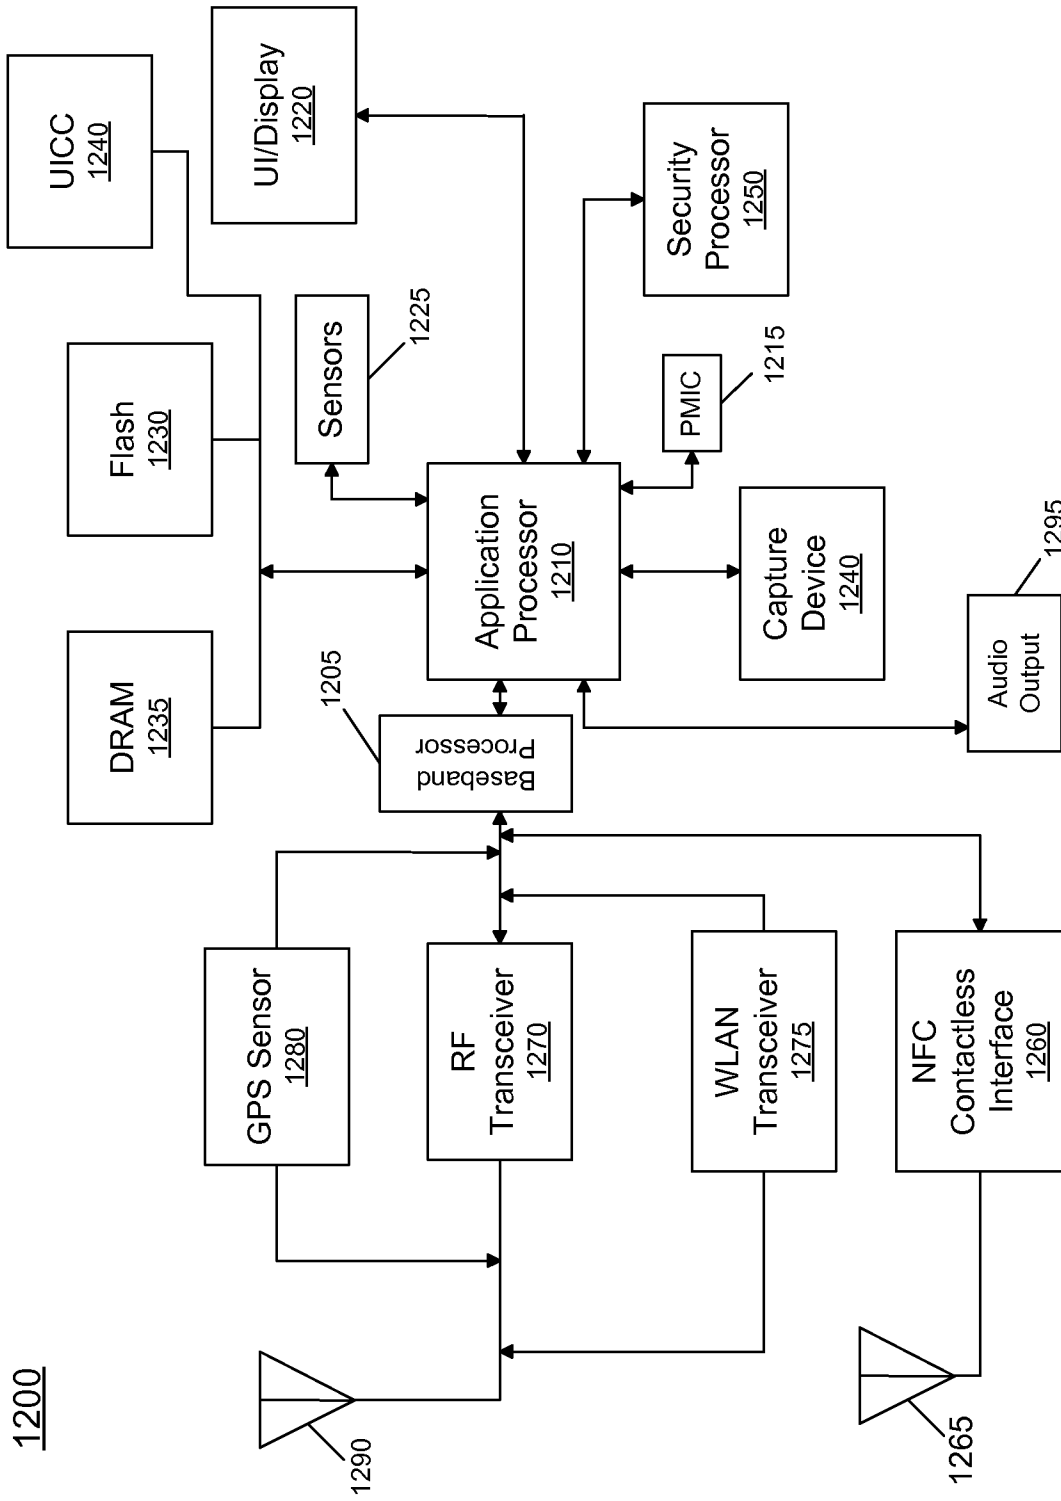


FIG. 12

1300

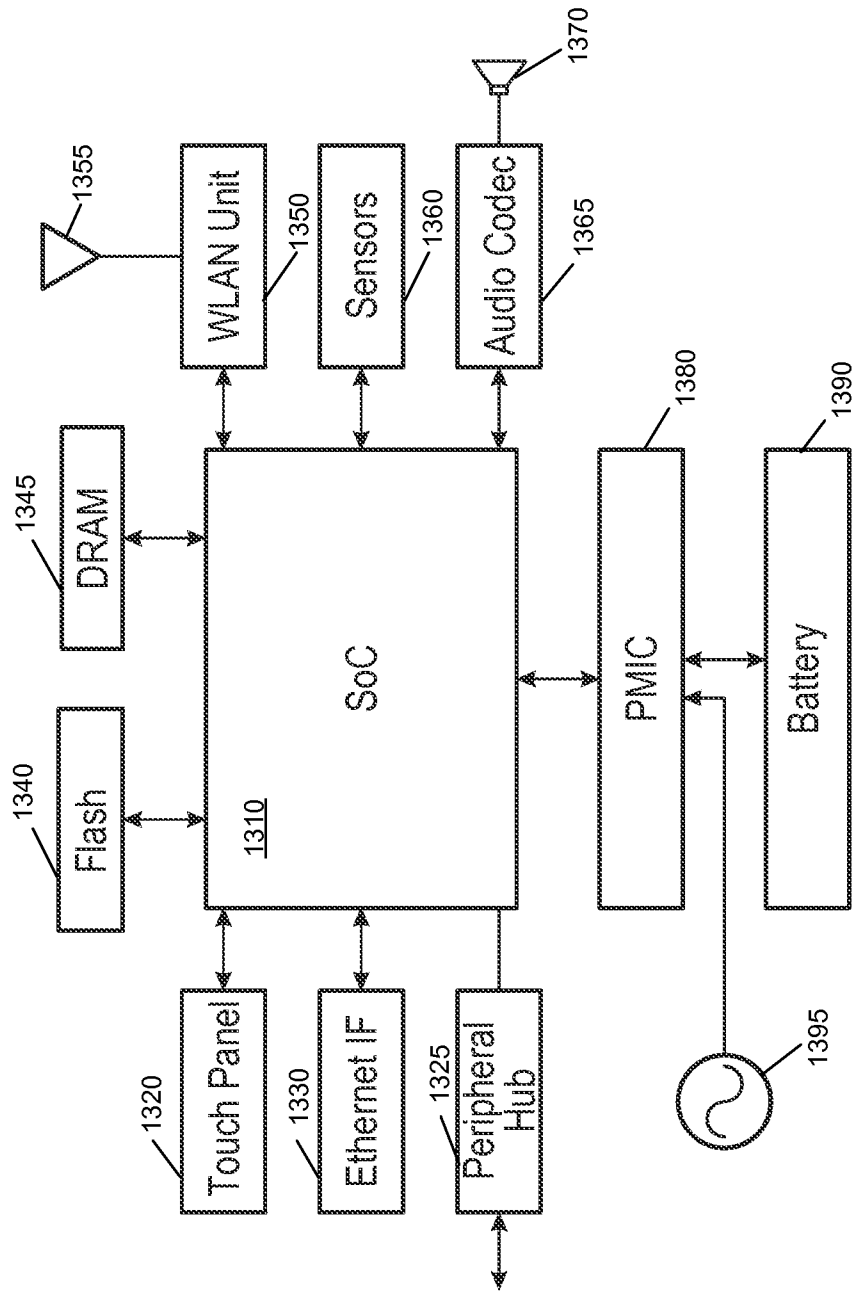


FIG. 13

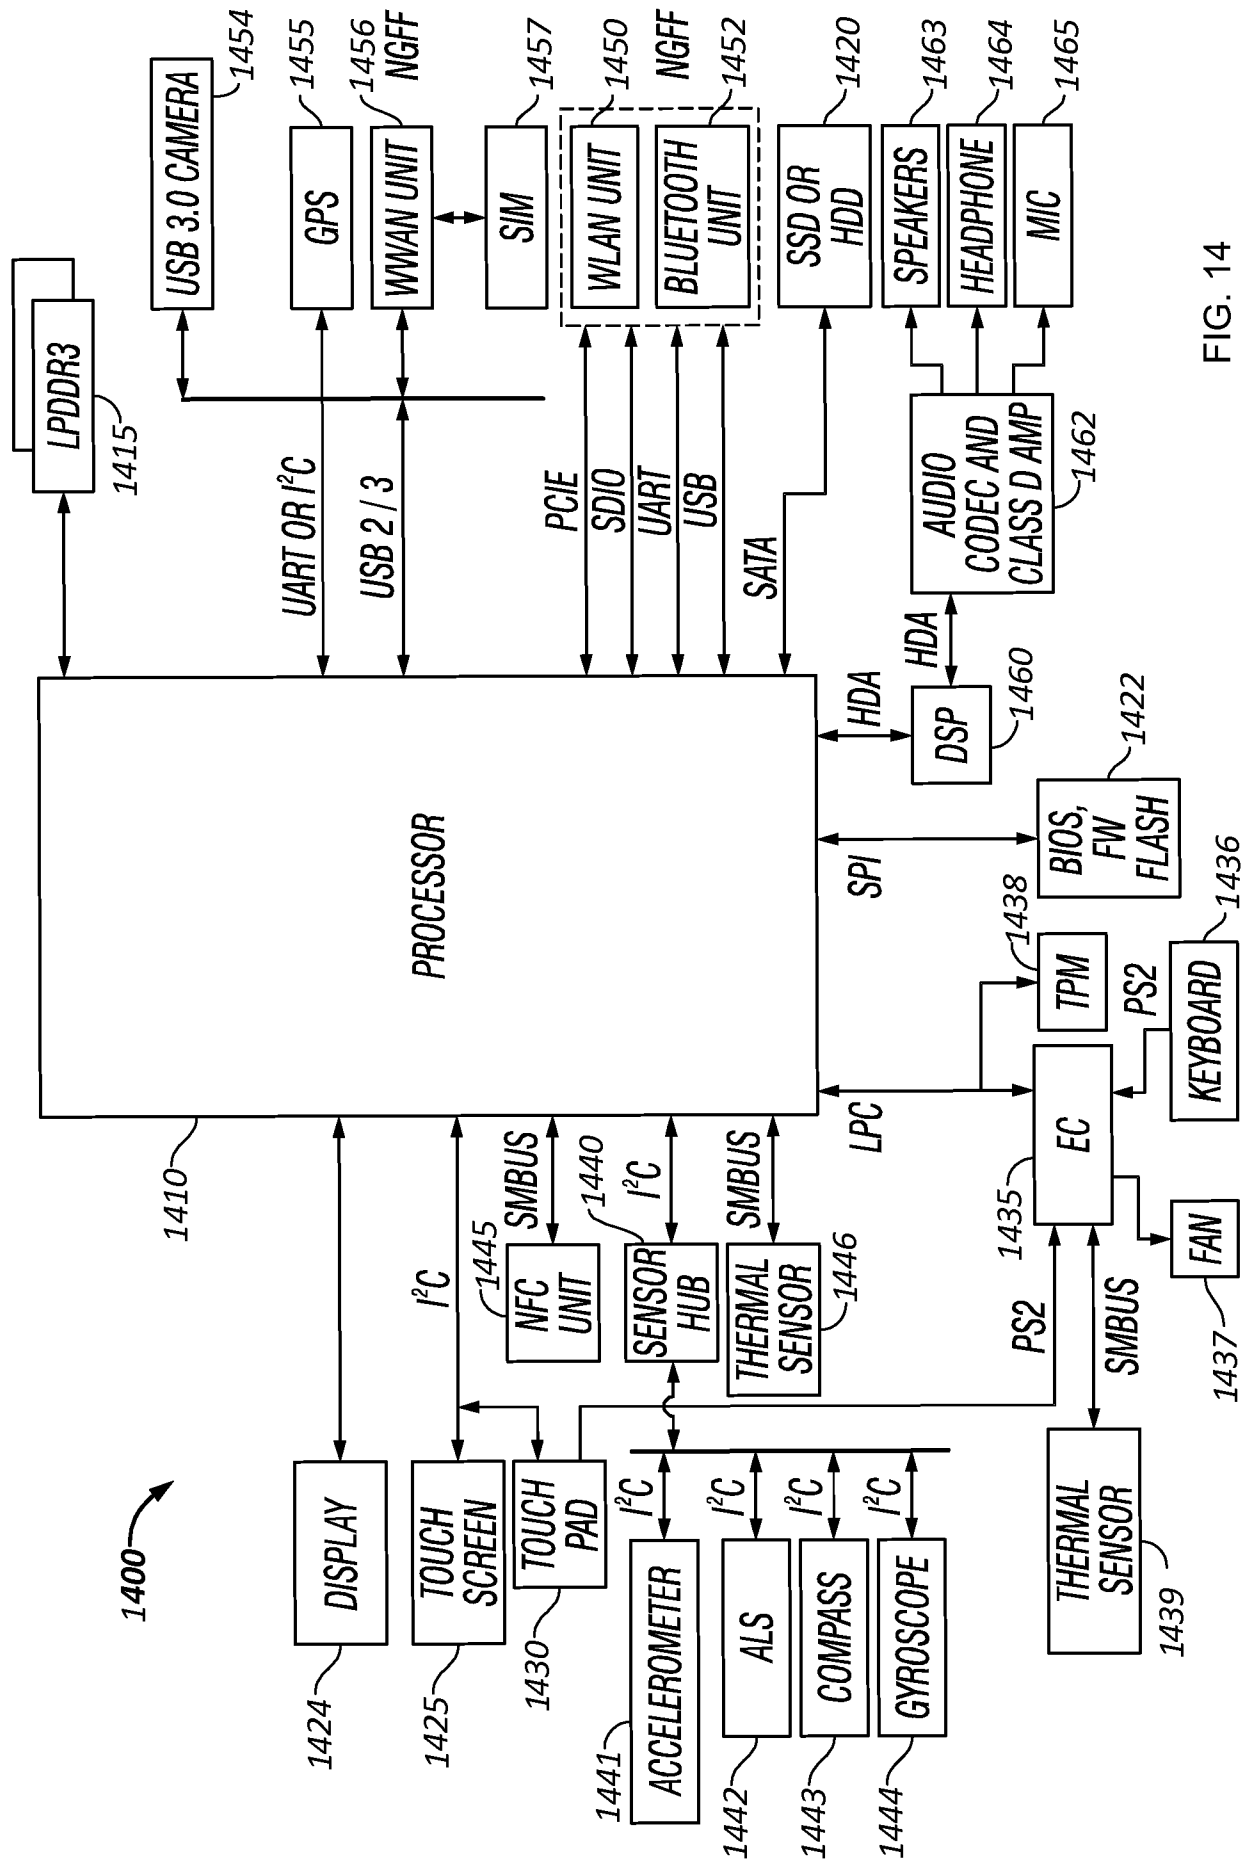


FIG. 14

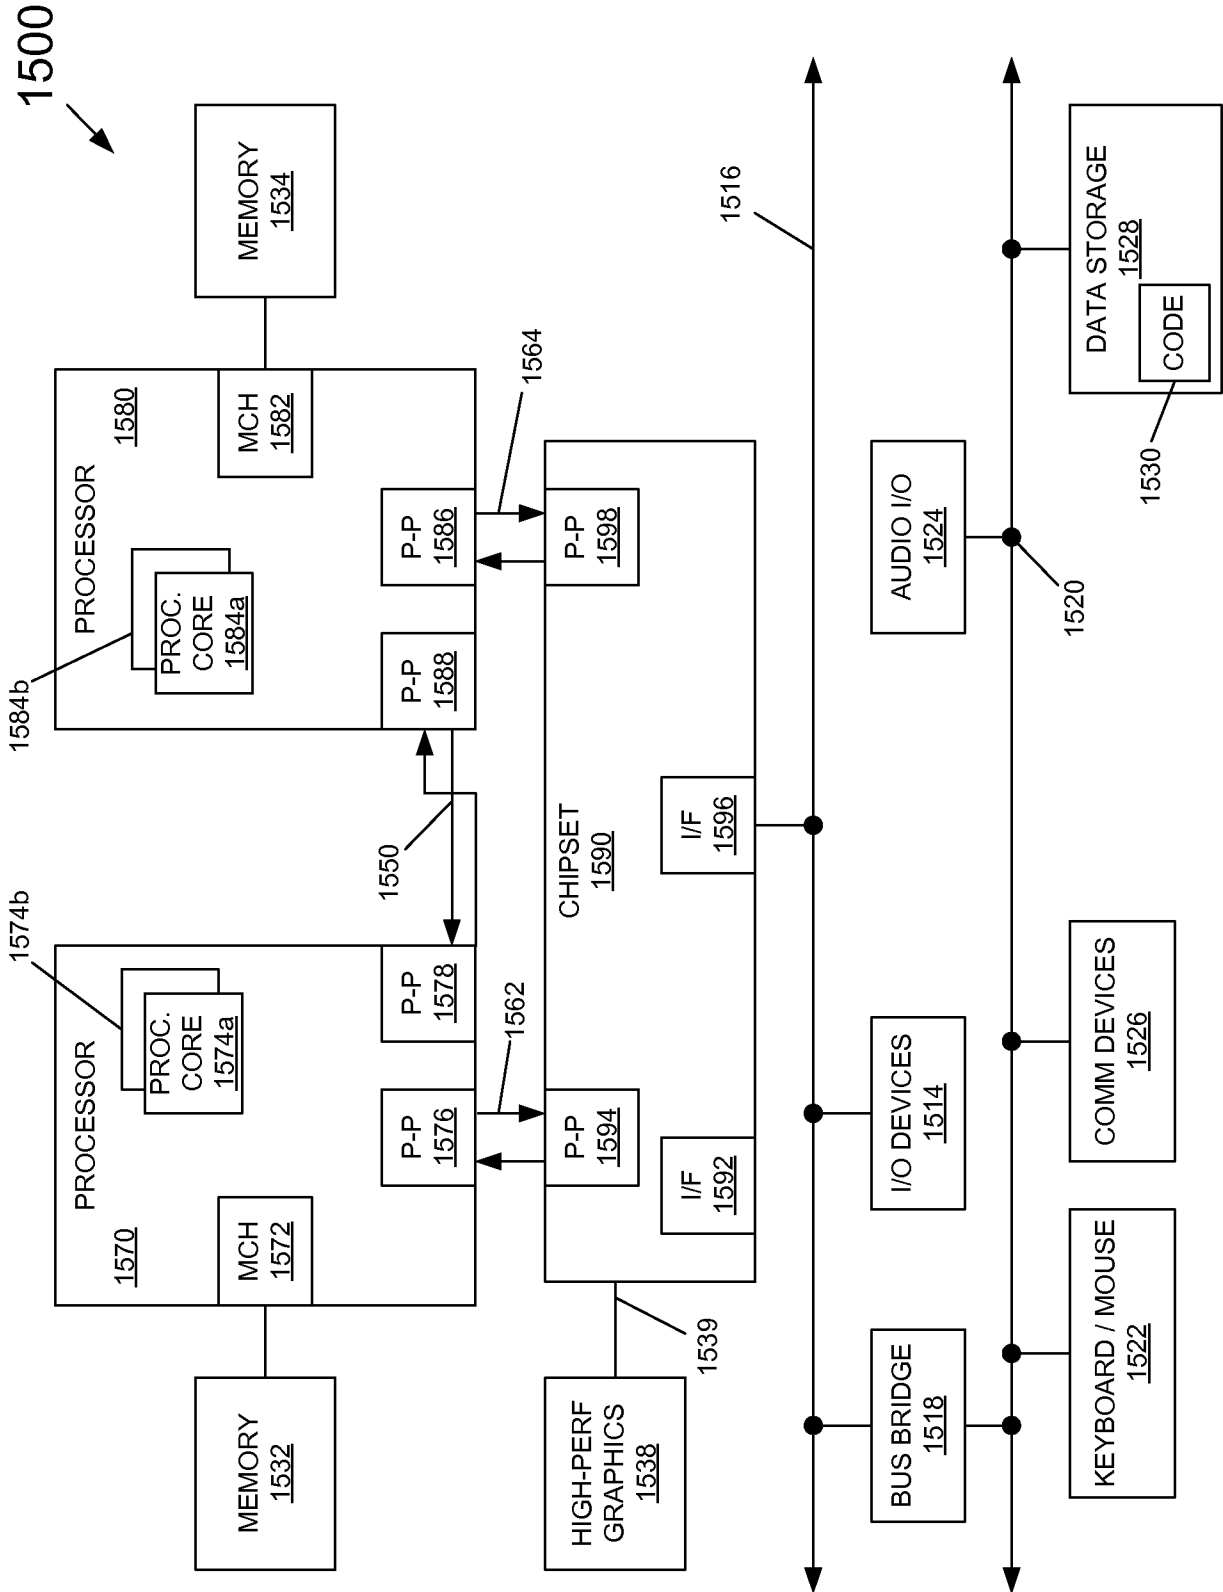


FIG. 15

16/19

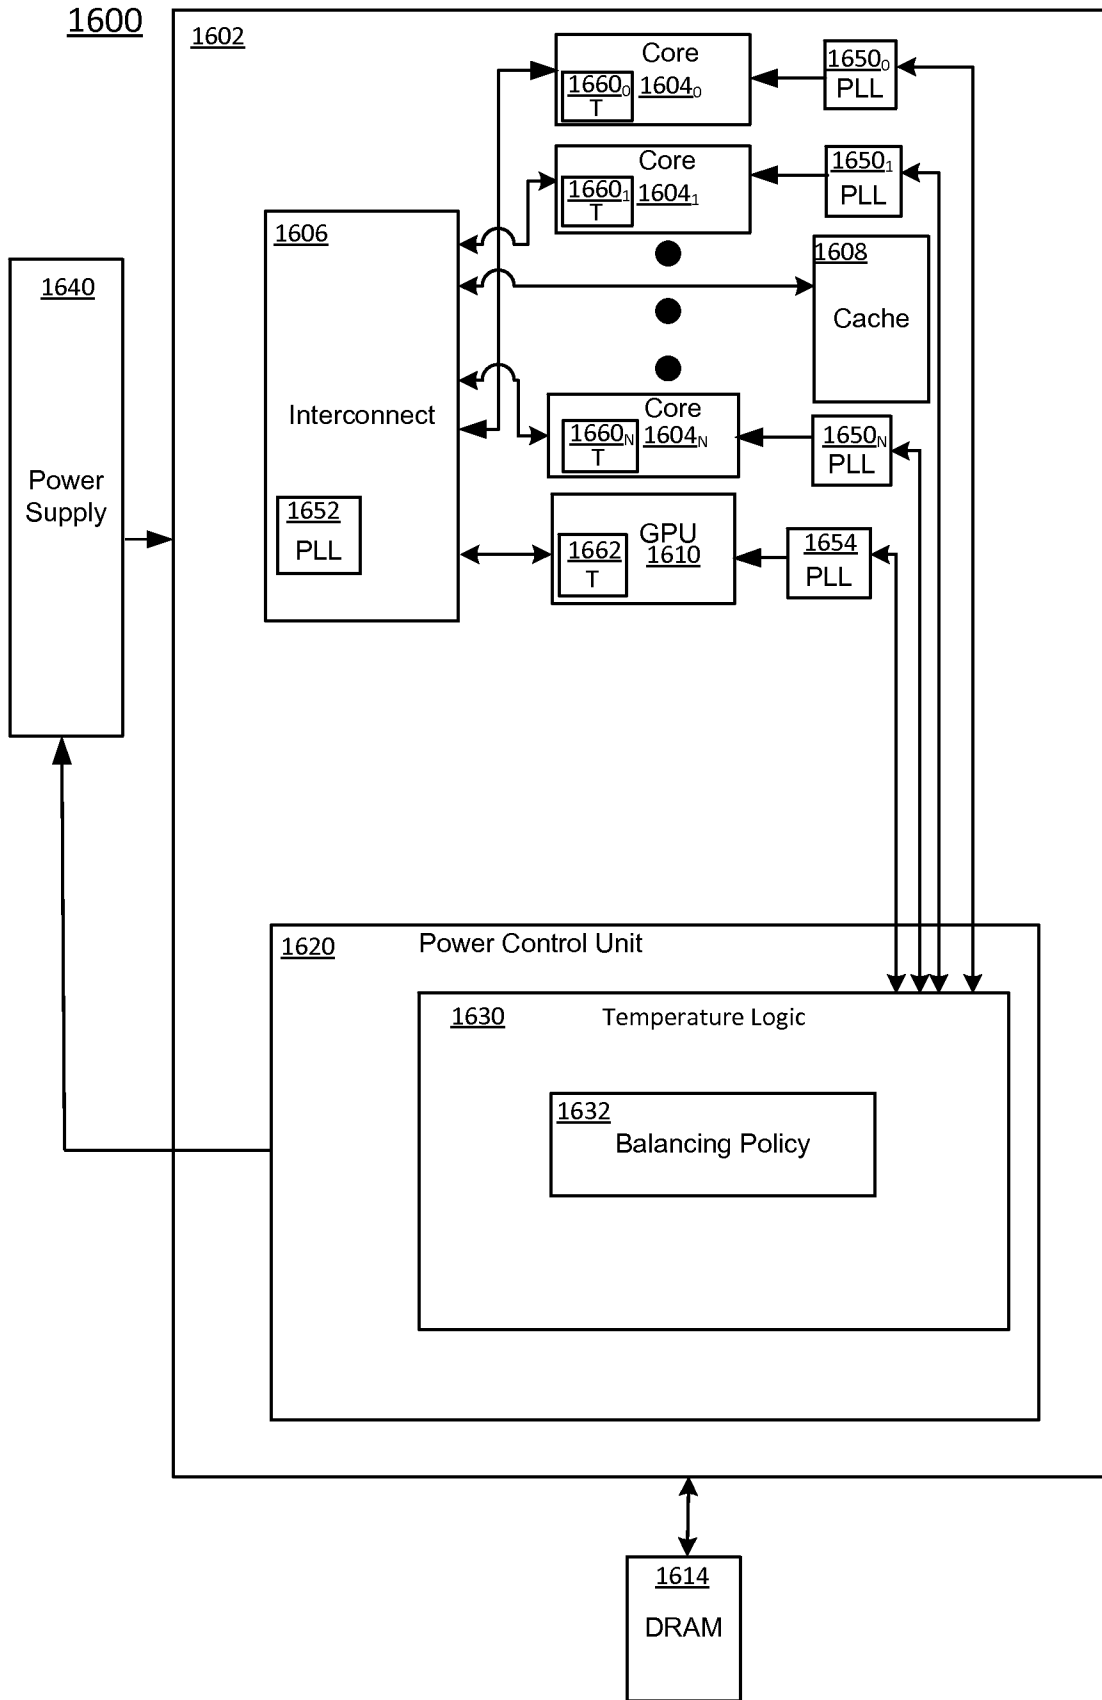


FIG. 16

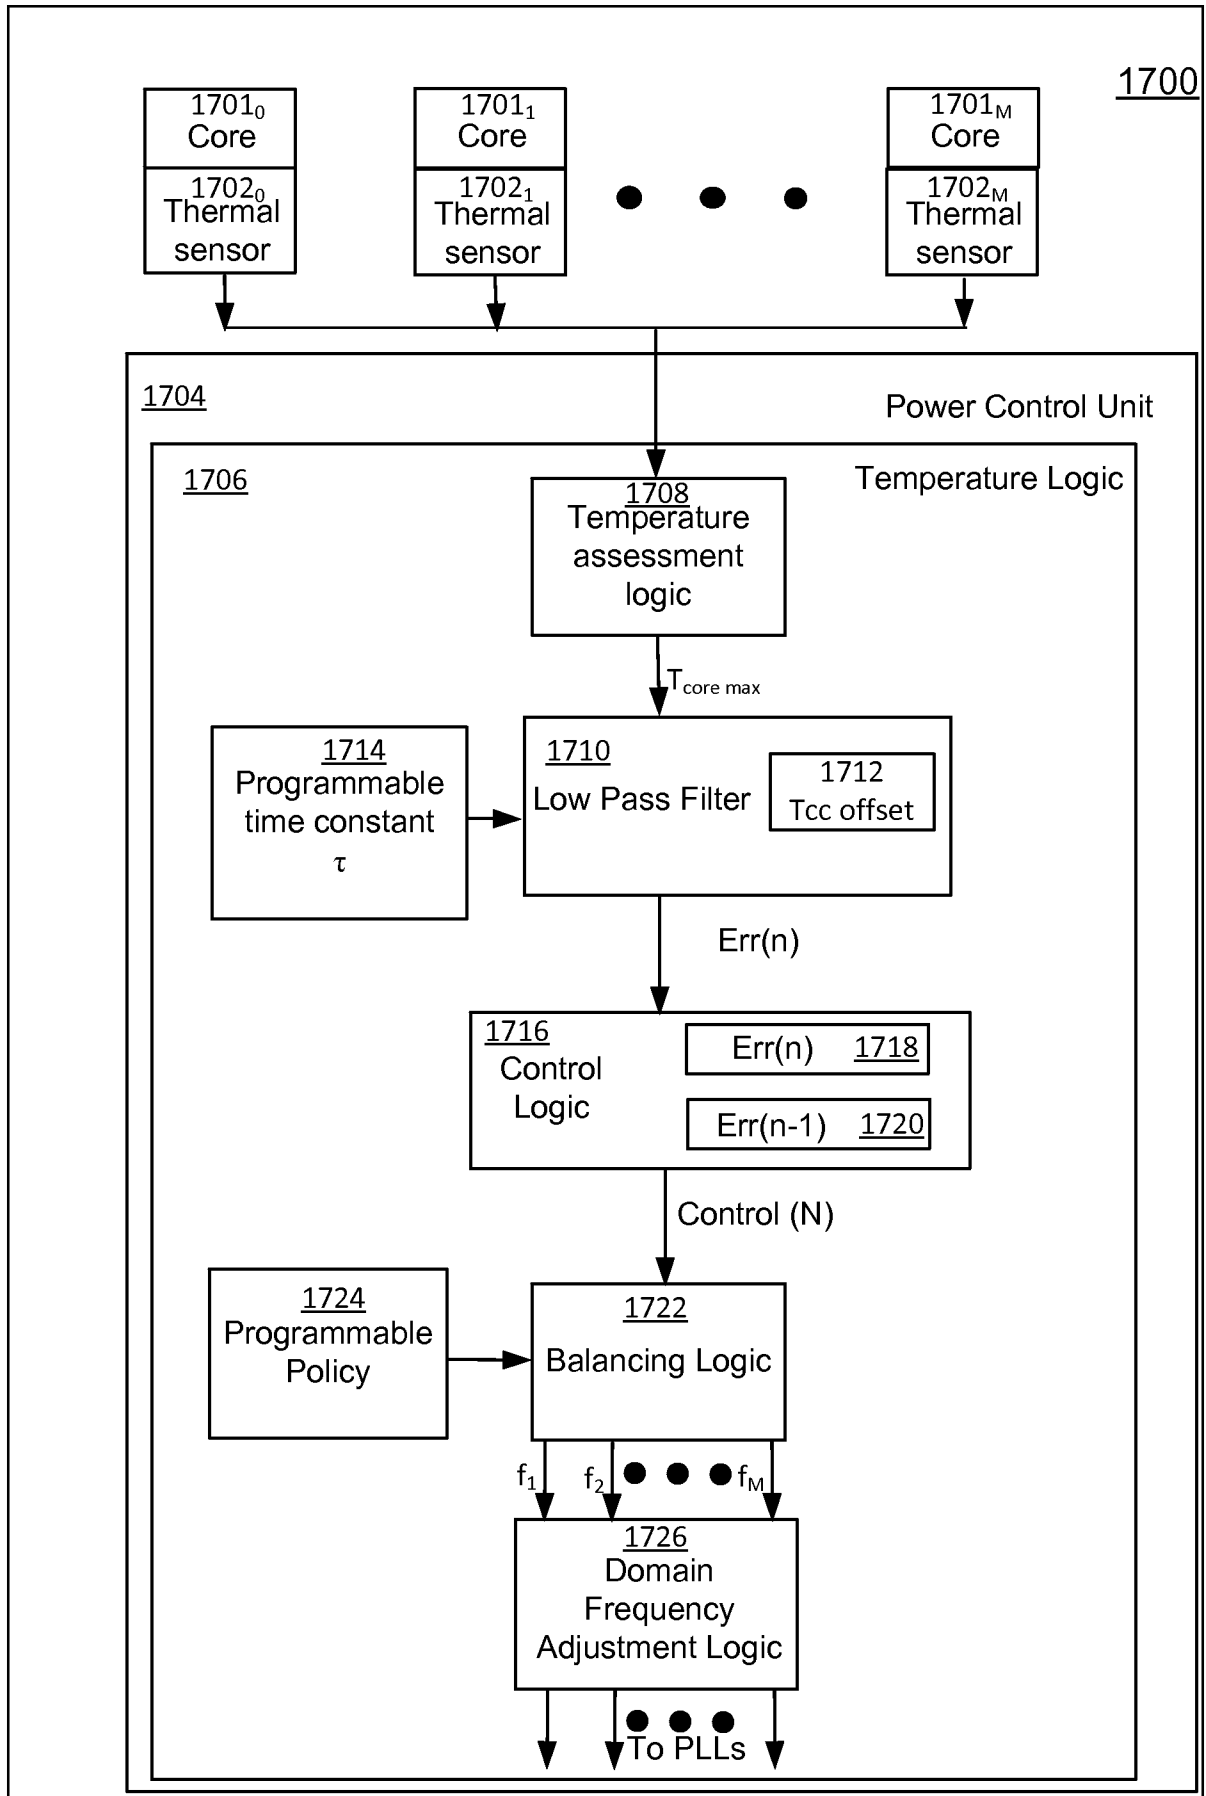


FIG. 17

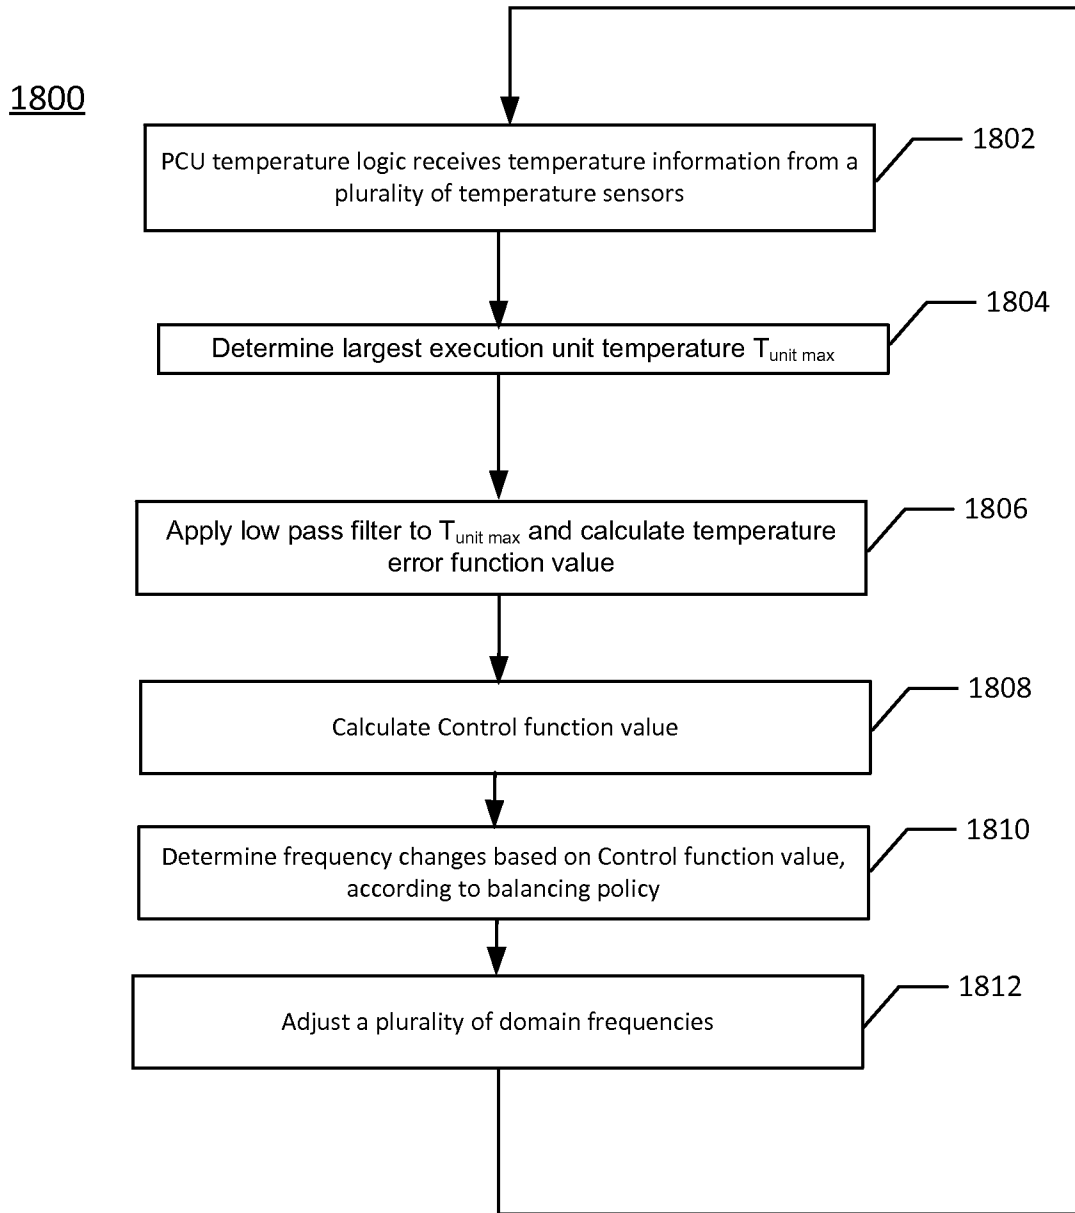


FIG. 18

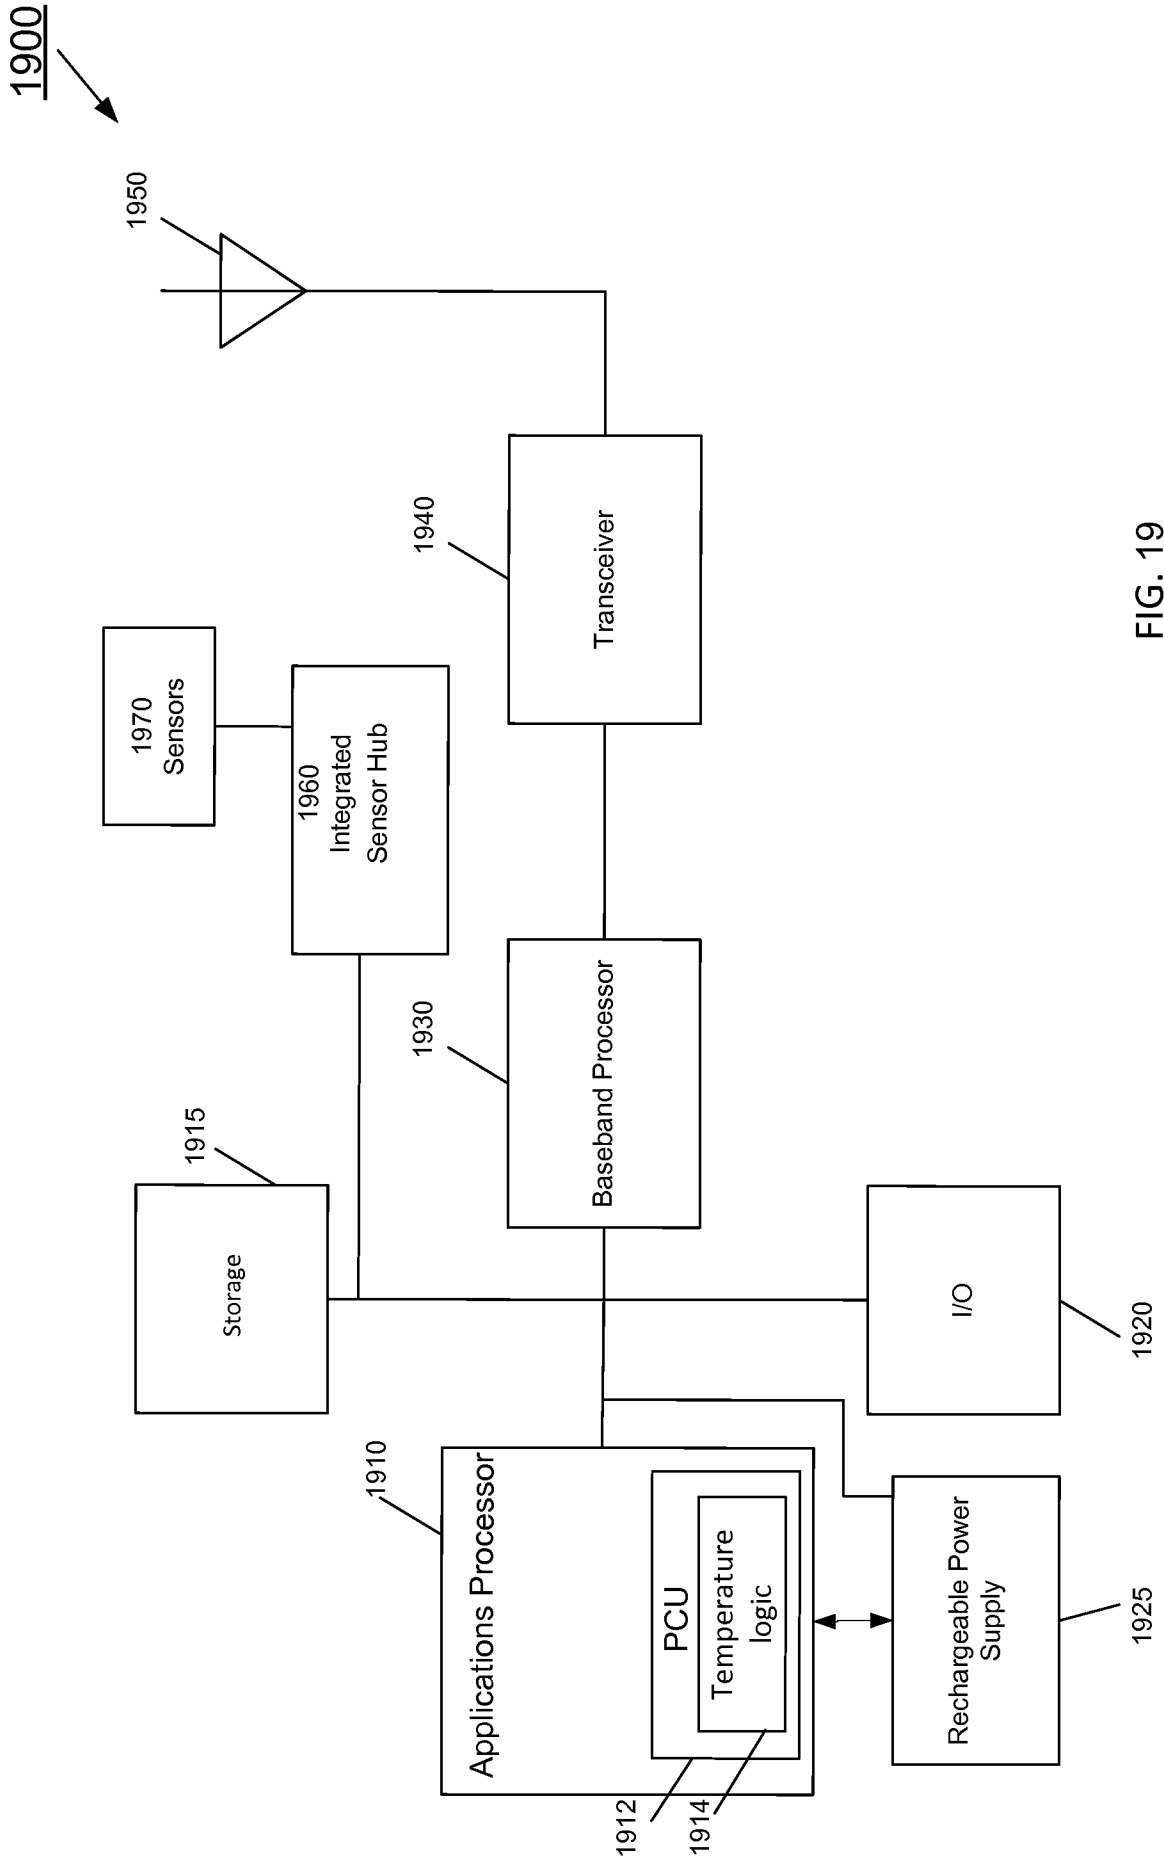


FIG. 19

## INTERNATIONAL SEARCH REPORT

International application No.  
**PCT/US2015/038125****A. CLASSIFICATION OF SUBJECT MATTER****G06F 1/26(2006.01)i, G06F 1/32(2006.01)i, G06F 1/20(2006.01)i**

According to International Patent Classification (IPC) or to both national classification and IPC

**B. FIELDS SEARCHED**Minimum documentation searched (classification system followed by classification symbols)  
G06F 1/26; G06F 1/00; G06F 1/32; G05D 23/19; H01L 37/00; G06F 1/20Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched  
Korean utility models and applications for utility models  
Japanese utility models and applications for utility modelsElectronic data base consulted during the international search (name of data base and, where practicable, search terms used)  
eKOMPASS(KIPO internal) & Keywords: core, temperature sensor, power control unit, clock frequency, threshold, adjust**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2012-0110352 A1 (ALEXANDER BRANOVER et al.) 03 May 2012 See paragraphs [0023], [0037], [0051]-[0056]; claim 25; and figures 1-2.	1-25
A	US 2014-0068291 A1 (XIUTING C. MAN et al.) 06 March 2014 See paragraphs [0014]-[0016]; and figure 1.	1-25
A	US 2013-0228632 A1 (APPLE INC.) 05 September 2013 See paragraphs [0039]-[0040]; and figure 1.	1-25
A	WO 2013-048827 A2 (INTEL CORPORATION) 04 April 2013 See paragraphs [0023]-[0024]; and figure 2.	1-25
A	US 2014-0115351 A1 (AVINASH N. ANANTHAKRISHNAN et al.) 24 April 2014 See paragraphs [0047]-[0048]; and figure 6.	1-25

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

\* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

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

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

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

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

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

"&amp;" document member of the same patent family


Date of the actual completion of the international search

15 September 2015 (15.09.2015)

Date of mailing of the international search report

**16 September 2015 (16.09.2015)**

Name and mailing address of the ISA/KR

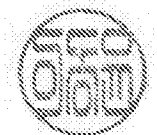

 International Application Division  
 Korean Intellectual Property Office  
 189 Cheongsu-ro, Seo-gu, Daejeon Metropolitan City, 35208,  
 Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

LEE, Dong Yun

Telephone No. +82-42-481-8734



**INTERNATIONAL SEARCH REPORT**

Information on patent family members

International application No.

**PCT/US2015/038125**

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2012-0110352 A1	03/05/2012	CN 103189814 A EP 2635948 A1 JP 2013-546070 A KR 10-2014-0001902 A US 8793512 B2 WO 2012-058202 A1	03/07/2013 11/09/2013 26/12/2013 07/01/2014 29/07/2014 03/05/2012
US 2014-0068291 A1	06/03/2014	CN 104508593 A GB 2520848 A US 2014-0068293 A1 US 9063727 B2 WO 2014-035542 A1	08/04/2015 03/06/2015 06/03/2014 23/06/2015 06/03/2014
US 2013-0228632 A1	05/09/2013	CN 104160357 A KR 10-2014-0136968 A TW 201351111 A WO 2013-130839 A1	19/11/2014 01/12/2014 16/12/2013 06/09/2013
WO 2013-048827 A2	04/04/2013	TW 201337526 A US 2013-0080804 A1 US 2013-0173941 A1 US 8954770 B2 WO 2013-048827 A3	16/09/2013 28/03/2013 04/07/2013 10/02/2015 23/05/2013
US 2014-0115351 A1	24/04/2014	DE 112012003701 T5 DE 202012008539 U1 TW 201333674 A TW I470409 B US 2013-0061064 A1 US 2013-0179704 A1 US 8769316 B2 US 8775833 B2 WO 2013-036497 A2 WO 2013-036497 A3	17/07/2014 12/12/2012 16/08/2013 21/01/2015 07/03/2013 11/07/2013 01/07/2014 08/07/2014 14/03/2013 10/05/2013