



- (51) International Patent Classification: **G01R 31/28** (2006.01)
- (21) International Application Number: PCT/US2012/031464
- (22) International Filing Date: 30 March 2012 (30.03.2012)
- (25) Filing Language: English
- (26) Publication Language: English
- (71) Applicant (for all designated States except US): **INTEL CORPORATION** [US/US]; 2200 Mission College Boulevard, Santa Clara, California 95052 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **VARMA, Ankush** [IN/US]; 179 SE 49th Court, Hillsboro, Oregon 97123 (US). **SISTLA, Krishnakanth V.** [IN/US]; 6000 SW Fountain Grove Terrace, Beaverton, Oregon 97007 (US). **ROWLAND, Martin T.** [US/US]; 6570 SW Ecola Place, Beaverton, Oregon 97007 (US). **GARG, Vivek** [US/US]; 257 Oxborough Drive, Folsom, California 95630 (US). **BURNS, James S.** [US/US]; 19700 Alderbrook Way, Cupertino, California 95014 (US).
- (74) Agent: **ROZMAN, Mark J.**; 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, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— as to the identity of the inventor (Rule 4.17(i))

Published:

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

(54) Title: DYNAMICALLY MEASURING POWER CONSUMPTION IN A PROCESSOR

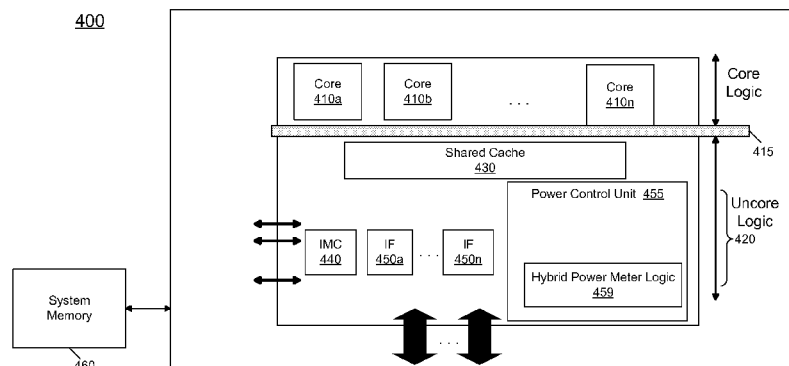


FIG. 6

(57) Abstract: In one embodiment, the present invention includes a processor having multiple cores to independently execute instructions, a first sensor to measure a first power consumption level of the processor based at least in part on events occurring on the cores, and a hybrid logic to combine the first power consumption level and a second power consumption level. Other embodiments are described and claimed.

WO 2013/147849 A1

DYNAMICALLY MEASURING POWER CONSUMPTION IN A PROCESSOR

Technical Field

[0001] The field relates to semiconductor devices for use in a variety of systems.

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).

[0003] Power management for integrated circuits such as processors (used in both server and client systems) depends on accurate measurements of estimates of current processor power consumption. Various components of a processor may have their voltage and frequency modulated to stay within specified power limits. Since exceeding a power constraint is undesirable, processors are tuned to always stay under the power limit. Errors in power measurement are accounted for as a guardband, resulting in reduced power being available for processor performance.

[0004] Consider a processor that has a 100 watt (W) power limit, also referred to as a thermal design power (TDP), and a +/-5% error in power measurement. Since the processor must stay below its power limit, it caps power when a power consumption

of 95W is measured, making 5W unavailable for use because it is reserved as a guardband. The size of this guardband is directly proportional to the amount of the error. These guardbands thus reduce available power. Further, inaccuracies of different types of power monitors can vary at low and high loads.

Brief Description of the Drawings

[0005] FIG. 1 is a block diagram of a system in accordance with one embodiment of the present invention.

[0006] FIG. 2 is a flow diagram of a method in accordance with an embodiment of the present invention.

[0007] FIG. 3 is a flow diagram of combining multiple power consumption level calculations in accordance with one embodiment of the present invention.

[0008] FIG. 4 is a flow diagram of a method for performing a correlation-based combination of multiple power sensors in accordance with an embodiment of the present invention.

[0009] FIG. 5 is a flow diagram of a method of performing power measurements in accordance with another embodiment of the present invention.

[0010] FIG. 6 is a block diagram of a processor in accordance with an embodiment of the present invention.

[0011] FIG. 7 is a block diagram of a multi-domain processor in accordance with another embodiment of the present invention.

[0012] FIG. 8 is a block diagram of a system in accordance with an embodiment of the present invention.

[0013] FIG. 9 is a block diagram of a multiprocessor system with a point-to-point (PtP) interconnect in accordance with one embodiment of the present invention.

[0014] FIG. 10 is a block diagram of a partially connected quad processor system in accordance with one embodiment of the present invention.

Detailed Description

[0015] Embodiments may be used to provide a power measurement for a processor that is highly accurate at all load levels of the processor. To this end, embodiments can obtain information both from a digital power meter and a voltage regulator-based current sensor to determine power consumption levels from such sensor information. More specifically, embodiments may generate a combined or hybrid power measurement based on sensor information obtained from these multiple power sensors. In this way, a more accurate power measurement can be obtained. Then based on this accurate power measurement, a greater power head room is realized. Thus based on this power information, a processor may be controlled to operate at higher operating frequencies and/or voltages to realize greater performance within a given power budget.

[0016] 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.

[0017] 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 individual voltage regulator $125_a - 125_n$. Accordingly, a fully integrated voltage regulator (FIVR) 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. As further seen, each core 120 can include at least one event counter 122_0-122_n to count certain events occurring on the core, such as instructions that use certain core circuitry (e.g., high power circuitry). This event information can be used to provide sensor information for a digital power meter in accordance with an embodiment of the present invention. Note that in some embodiments, the cores themselves may

contain a digital power meter to use this information to generate a power consumption level for the core, and communicate this information to a power control unit described below.

[0018] Also although only a single event counter per core is shown for ease of illustration, understand that the scope of the present invention is not limited in this regard. For example, in other implementations of a number of event counters can be present per core. Each counter can be configured to count a number of instructions executed of a given type in the corresponding core. For example, all instructions of an instruction set architecture (ISA) can be associated with one of these counters, where instructions of roughly the same power consumption level (e.g., due to the units of the core used for such instruction's execution) can be associated or bucketed with the same counter. In this way, a relatively accurate measure of actual power consumption based on the instructions being executed in the cores can be achieved. In one embodiment, each of these counters may be associated with a given weight value, generally corresponding to its relative power consumption level. Thus in one embodiment, a digital power meter (DPM) can operate based on counting events, assigning an energy weight to each event, and scaling for temperature. The rate at which various events occur is an indicator of the dynamic power consumption of the processor.

[0019] Further understand that instead of having an internal (to the core) digital power meter, instead the weighted count information from each of the event counters (of each core) can be provided to logic of a power control unit that can perform digital power measurements based on this information. In addition, this power controller-based digital power meter can further receive information from other portions of a processor such as uncore or system agent circuitry, interface circuitry, interconnect circuitry and so forth. Based on all such information, this digital power meter can generate a relatively accurate measure of power consumption. Note that for the system agent or uncore circuitry, the information may be associated with cache accesses. For interconnect circuitry, relative bandwidth may be used as an indication of power consumption. Similarly, for interface circuitry, a measure of the amount of data packets sent and received can be a good proxy for power

consumption. While such a digital power meter can be relatively accurate across a full load line of the processor, it can be very difficult to tune the meter appropriately for the various operations and events occurring within the processor. Accordingly, a hybrid power meter in accordance with an embodiment of the present invention can improve the accuracy without the need for more complex tuning of a digital power meter.

[0020] 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_x. 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).

[0021] Also shown is a power control unit (PCU) 138, which may include hardware, software and firmware to perform power management operations with regard to processor 110. In various embodiments, PCU 138 may include logic to perform digital power measurements, as described above. In addition, PCU 138 may include logic to perform hybrid power measurements in accordance with an embodiment of the present invention. Furthermore, PCU 138 may be coupled via a dedicated interface to external voltage regulator 160. In this way, PCU 138 can instruct the voltage regulator to provide a requested regulated voltage to the processor. In addition, voltage regulator 160 can provide information regarding its current delivery to the processor. In different implementations, voltage regulator 160 can store this information in a register of the voltage regulator that the PCU accesses. Or a current sensor, located either in voltage regulator 160 or on the path between voltage regulator 160 and PCU 138 can provide this information. This current information can be used by power meter logic of PCU 138 to generate a power consumption level based on this current delivery. Thus a voltage regulator-

based current sensor can directly measure the current supplied by voltage regulator 160 to the processor. When multiplied by the supply voltage, this provides a measurement of processor power consumption.

[0022] As will be described below, logic within PCU 138 can be used to both calculate power consumption levels in multiple manners, including as described above as well as possibly in other manners and then, using a hybrid power measurement logic in accordance with an embodiment of the present invention, determine a hybrid power consumption level based on a combination of these individual power consumption levels.

[0023] 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.

[0024] 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, 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 apparatus', 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, apparatus', 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.

[0025] Referring now to FIG. 2, shown is a flow diagram of a method in accordance with an embodiment of the present invention. As shown in FIG. 2, method 200 can be implemented by logic of a PCU. More specifically, method 200 can be implemented by hybrid power measurement logic of a PCU that can receive and process information from multiple power sensors.

[0026] As seen in FIG. 2, method 200 may begin by receiving sensor information from a digital power meter (block 210). As described above, in some embodiments at least portions of this digital power meter can be implemented within the PCU. From this sensor information, a first power consumption level of the processor can be calculated (block 220). For example, the logic can calculate a power consumption level based on the sensor information. As an example, this digital power meter can provide various information such as an indication of events occurring on various cores (e.g., instructions bucketed into different counters, which are weighted according to their relative power consumption) and other circuitry (e.g., as determined by cache activity, interconnect bandwidth, interface activity and so forth) of the processor. From this, a first power consumption level can be calculated. Note that in other embodiments, the first power consumption level can be received from the digital power meter itself.

[0027] Still referring to FIG. 2, similarly sensor information can be obtained from a current sensor such as a sensor that measures the current supplied by a voltage regulator to the processor (block 230). Different implementations of this current sensor may be present, from internal hardware of an external voltage regulator that provides power to the processor, to a current sensing circuit located in a path from the voltage regulator to the processor (and possibly within the processor itself).

From this sensor information, control passes to block 240 where a second power consumption level can be calculated using this current sensor information. In one embodiment, the calculation may be implemented by multiplying the measured current by the supply voltage and thus a power consumption level corresponding to: $P=IV$, can be determined.

[0028] As further shown in FIG. 2, control next passes to block 250 where the first and second power consumption levels can be combined to obtain a hybrid power consumption level. As will be described further below, different manners of combining these different power consumption levels can occur. The different manners can range from using both of the power levels, e.g., according to scaled values, selecting a particular one of the power consumption levels based on a load level of the processor, or other such combinations.

[0029] Finally, control passes to block 260 where an operating frequency and/or voltage of the processor can be controlled based on this hybrid power consumption level and a power limit for the processor. As an example, typical processors can have a thermal design power (TDP) that corresponds to a maximum power dissipation that the processor can output (that can be handled by a cooling system). Using this as a maximum value and understanding a present loading of the processor and thus a current power consumption level (namely this hybrid power consumption level), in addition to obtaining a more accurate power measurement regardless of where on the load line the processor is executing, it may be possible to increase the operating frequency and/or voltage to thus obtain greater performance while remaining within the power budget, namely the TDP. Although shown at this high level in the embodiment of FIG. 2, understand the scope of the present invention is not limited in this regard.

[0030] As discussed, in different embodiments different manners of combining the power consumption levels determined by the different sensors can occur. Referring now to FIG. 3, shown is a flow diagram of one embodiment of combining multiple power consumption level calculations in accordance with an embodiment of the present invention. As seen in FIG. 3, method 300 may begin by determining whether

the first power consumption level (which as discussed above is generated from sensor information received from a digital power meter) is less than a first threshold. If so, this is an indication that the processor is operating at a low load, meaning that relatively low processor utilization is occurring (e.g., as a relatively low number of instructions per cycle are being executed). At such low loads the digital power meter may provide a more accurate measure of processor power consumption than other techniques. Accordingly, control passes in FIG. 3 to block 315 where this first power consumption level can be used to control operating frequency and/or voltage.

[0031] If instead this first power consumption level is greater than the threshold, control passes to block 320 where instead the second power consumption level can be used to control operating frequency and/or voltage, as a current-based sensor may be more accurate at higher load levels. Although the scope of the present invention is not limited in this regard in some embodiments this threshold level may correspond to a load level of the processor between approximately 40 and 60% of a processor utilization, e.g., corresponding to roughly half of a TDP value of the processor.

[0032] Still other manners of combining power consumption levels determined by multiple power sensors can occur. As an example, a correlation-based combination can occur. In this way, one of the power consumption levels can be used to apply a correction factor to the other power consumption level. Assume for example given that a DPM may be accurate at low loads and a current sensor method is more accurate at higher loads, a correlation factor can be computed by executing a low power load and obtaining sensor information, and then executing a high power load and obtaining sensor information.

[0033] Referring now to FIG. 4, shown is a flow diagram of a method for performing a correlation-based combination of multiple power sensors in accordance with an embodiment of the present invention. As shown in FIG. 4, method 330 can begin by controlling a processor for a low power load (block 335). In one embodiment, this control operation can occur during a boot process in which a predetermined load is executed on the processor. During execution of this load,

control passes to block 340 where the first and second power consumption levels can be calculated using sensor information from the digital power meter and the current sensor. In other implementations, rather than executing a low power load, which can be a predetermined load during a boot process, instead the PCU can control the processor to be in a predetermined low power state, e.g., by controlling the cores and other circuitry of the processor accordingly to thus obtain sensor information while the processor is in this low load state. In any event, calculations can occur as discussed above in FIG. 2, in one embodiment. Next an offset can be determined and stored based on these power consumption levels (block 345). In one embodiment, this offset can be determined as follows:

Offset = First Power Consumption Level – Second Power Consumption Level.

In one embodiment, this offset corresponding to the difference between the power consumption levels can be stored in a power management storage, e.g., present in a PCU. Although described as being stored within a storage of the PCU, understand the scope of the present invention is not limited in this regard. For example, in other implementations, this offset value (in other calibration values discussed further below) can be stored in a non-volatile storage, e.g., a non-volatile storage associated with the basic input/output system (BIOS).

[0034] Still referring to FIG. 4, next control passes to block 350 where the processor can be controlled for a high power load. During execution of this high power load, which again can occur during a boot process and may correspond to a predetermined load, at least the second power consumption level can be calculated using the sensor information from the current sensor (block 355). Control next passes to block 360 where a slope can be determined and stored based on the offset and the second power consumption level (namely the power consumption level obtained from the current sensor). In one embodiment, the slope can be determined according to:

Slope = (Second Power Consumption Level – Offset)/Second Power Consumption Level. Note that this slope can also be stored in the power management storage.

[0035] These operations as discussed above in FIG. 4 can be performed during a boot process. Or they can be performed under control of the PCU, e.g., when a different temperature level of the processor is reached (e.g., when a processor temperature exceeds a previous temperature by a threshold amount), responsive to a user request, or at another interval. After obtaining of these slope and offset values, the values can be used during normal operation.

[0036] Referring now to FIG. 5, shown is a flow diagram of a method of performing power measurements in accordance with another embodiment of the present invention. As shown in FIG. 5, method 370 may be used to measure power consumption during normal operation using these stored values. As seen, method 370 may begin by receiving sensor information from a current sensor (block 375). Control next passes to block 380 where a second power consumption level can be calculated using this current sensor information. Control then passes to block 385 where a hybrid power consumption level can be determined using this second power consumption level, and the slope and offset. In one embodiment, this hybrid power consumption level can be determined according to:

Hybrid = Second Power Consumption Level x Slope + Offset.

[0037] Accordingly, a hybrid power consumption level can be determined and can be used at block 390 to control the operating frequency and/or voltage of the processor. Although shown at this high level in the embodiment of FIGS. 4 and 5, understand the scope of the present invention is not limited in this regard. For example, in another implementation, the calculations can be performed to enable applying a correction factor to the power consumption level determined according to the digital power meter, where the correction factors are obtained using the power consumption levels from the digital power meter and the current sensor. And furthermore understand that different equations for generating the correction values and the hybrid power consumption level can occur in other embodiments.

[0038] Embodiments can be implemented in processors for various markets including server processors, desktop processors, mobile processors and so forth. Referring now to FIG. 6, shown is a block diagram of a processor in accordance with

an embodiment of the present invention. As shown in FIG. 6, processor 400 may be a multicore processor including a plurality of cores 410_a – 410_n. In one embodiment, each such core may be of an independent power domain and can be configured to enter and exit active states and/or maximum performance states based on workload. The various cores may be coupled via an interconnect 415 to a system agent or uncore 420 that includes various components. As seen, the uncore 420 may include a shared cache 430 which may be a last level cache. In addition, the uncore may include an integrated memory controller 440, various interfaces 450 and a power control unit 455.

[0039] In various embodiments, power control unit 455 may include a hybrid power meter logic 459 in accordance with an embodiment of the present invention. As described above, this power meter can receive sensor information from different power sensors, including a digital power meter and an analog or current-based sensor. Note that different implementations are possible, such as a hybrid power meter that receives sensor information from each of multiple pairs of such sensors, e.g., where each pair is associated with a given core. Or, digital power meters may be present in the individual cores (and other processor circuitry) and instead a single current sensor-based power meter may be present in the processor. Variations on these implementations are of course also possible.

[0040] Based on the information received from these various sensors, hybrid power meter logic 459 can combine the sensor information in a selected manner to obtain a very accurate measure of power consumption in the processor. In this way, processor 400 can be configured to operate with very little guardband from a maximum power consumption level, e.g., a TDP level of the processor.

[0041] With further reference to FIG. 6, processor 400 may communicate with a system memory 460, e.g., via a memory bus. In addition, by interfaces 450, 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. 6, the scope of the present invention is not limited in this regard.

[0042] Referring now to FIG. 7, 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. 7, processor 500 includes multiple domains. Specifically, a core domain 510 can include a plurality of cores 510_0 – 510_n , a graphics domain 520 can include one or more graphics engines, and a system agent domain 550 may further be present. In some embodiments, system agent domain 550 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 510 and 520 can be controlled to dynamically enter into and exit low power states. Each of domains 510 and 520 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.

[0043] In general, each core 510 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) 540_0 – 540_n . In various embodiments, LLC 540 may be shared amongst the cores and the graphics engine, as well as various media processing circuitry. As seen, a ring interconnect 530 thus couples the cores together, and provides interconnection between the cores, graphics domain 520 and system agent circuitry 550. In one embodiment, interconnect 530 can be part of the core domain. However in other embodiments the ring interconnect can be of its own domain.

[0044] As further seen, system agent domain 550 may include display controller 552 which may provide control of and an interface to an associated display. As further seen, system agent domain 550 may include a power control unit 555 which can include a hybrid power meter logic 559 in accordance with an embodiment of the present invention to dynamically and accurately measure power consumption in the processor to enable greater processor performance in view of greater available

processing power. In various embodiments, this logic may execute the algorithms described above in one or more of FIGS. 2-5.

[0045] As further seen in FIG. 7, processor 500 can further include an integrated memory controller (IMC) 570 that can provide for an interface to a system memory, such as a dynamic random access memory (DRAM). Multiple interfaces 580₀ – 580_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 Peripheral Component Interconnect Express (PCI Express™ (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. 7, understand the scope of the present invention is not limited in this regard.

[0046] Embodiments may be implemented in many different system types. Referring now to FIG. 8, shown is a block diagram of a system in accordance with an embodiment of the present invention. As shown in FIG. 8, multiprocessor system 600 is a point-to-point interconnect system, and includes a first processor 670 and a second processor 680 coupled via a point-to-point interconnect 650. As shown in FIG. 8, each of processors 670 and 680 may be multicore processors, including first and second processor cores (i.e., processor cores 674a and 674b and processor cores 684a and 684b), although potentially many more cores may be present in the processors. Each of the processors can include a PCU or other logic to perform hybrid power measurements, and corresponding operating frequency and/or voltage control, as described herein.

[0047] Still referring to FIG. 8, first processor 670 further includes a memory controller hub (MCH) 672 and point-to-point (P-P) interfaces 676 and 678. Similarly, second processor 680 includes a MCH 682 and P-P interfaces 686 and 688. As shown in FIG. 8, MCH's 672 and 682 couple the processors to respective memories, namely a memory 632 and a memory 634, which may be portions of system memory

(e.g., DRAM) locally attached to the respective processors. First processor 670 and second processor 680 may be coupled to a chipset 690 via P-P interconnects 652 and 654, respectively. As shown in FIG. 8, chipset 690 includes P-P interfaces 694 and 698.

[0048] Furthermore, chipset 690 includes an interface 692 to couple chipset 690 with a high performance graphics engine 638, by a P-P interconnect 639. In turn, chipset 690 may be coupled to a first bus 616 via an interface 696. As shown in FIG. 8, various input/output (I/O) devices 614 may be coupled to first bus 616, along with a bus bridge 618 which couples first bus 616 to a second bus 620. Various devices may be coupled to second bus 620 including, for example, a keyboard/mouse 622, communication devices 626 and a data storage unit 628 such as a disk drive or other mass storage device which may include code 630, in one embodiment. Further, an audio I/O 624 may be coupled to second bus 620. 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.

[0049] FIG. 9 is a block diagram of a system coupled with point-to-point (PtP) system interconnects in accordance with a given cache coherence protocol using QPI links as the system interconnect. In the embodiment shown, each processor 710 is coupled to two PtP links 725 and includes one instance of an integrated memory controller 715 that in turn is coupled to a corresponding local portion of a system memory 720. Each processor can perform hybrid power metering and corresponding power management policies as described herein. The processors are connected to an input/output hub (IOH) 730 using one link and the remaining link is used to connect the two processors.

[0050] Referring now to FIG. 10, shown is a block diagram of a system in accordance with another embodiment of the present invention. As shown in FIG. 10, system 800 may be a partially connected quad processor system in which each processor 810 (each of which may be multicore multi-domain processors) is coupled to each other processor via a PtP link and is coupled to a local portion of memory (e.g., dynamic random access memory (DRAM)) 820 via a memory interconnect

coupled to an integrated memory controller 815 of the corresponding processor. In the partially connected system of FIG. 10, note the presence of two IOHs 830 and 840 such that processors 810₀ and 810₁ are directly coupled to IOH 830 and similarly processors 810₂ and 810₃ are directly coupled to IOH 840.

[0051] In one aspect, a processor includes multiple cores to independently execute instructions, a first sensor to measure a first power consumption level of the processor based at least in part on events occurring on the cores, and a hybrid logic to combine the first power consumption level with a second power consumption level determined based on a dynamic current provided to the processor. To this end, a power controller included in or coupled to the processor may control at least one of an operating frequency and a voltage of the processor based on this combined consumption level and a power limit of the processor.

[0052] In another aspect, a method includes receiving, in a first logic of a processor, sensor information from a digital power meter of the processor, and calculating a first power consumption level of the processor using this sensor information; receiving, in the logic, sensor information from a current sensor configured to measure a current delivered by a voltage regulator coupled to the processor and calculating a second power consumption level of the processor using the sensor information from the current sensor; and combining, in the first logic, the first and second power consumption levels to obtain a hybrid power consumption level of the processor. From this information and a power limit of the processor, an operating frequency and/or voltage of the processor can be controlled.

[0053] Yet another aspect includes a system with a multicore processor and a system memory. The cores may each include one or more event counters to count events occurring on the core, a digital power meter to calculate a first power consumption level based on the count information, a second power meter to calculate a second power consumption level based on a current delivered to the processor from a voltage regulator, and a power controller including logic to generate a combined power consumption level of the processor using the first and second power consumption levels.

[0054] In another aspect, a processor means includes execution means each for independently executing instructions, sensor means for measuring a first power consumption level of the processor means based at least in part on events occurring on the execution means, and means for combining the first power consumption level and a second power consumption level of the processor means determined based on a dynamic current provided to the processor means. In turn an operating frequency and/or voltage of the processor means can be controlled via a controller means based on the combined first and second power consumption levels and a power limit of the processor means.

[0055] 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.

[0056] 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.

[0057] 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 to independently execute instructions;
3 a first sensor to measure a first power consumption level of the processor
4 based at least in part on a plurality of events occurring on the plurality of cores; and
5 a hybrid logic to combine the first power consumption level and a second
6 power consumption level of the processor determined based on a dynamic current
7 provided to the processor.

- 1 2. The processor of claim 1, further comprising a power controller to
2 control at least one of an operating frequency and a voltage of the processor based
3 on the combined first and second power consumption levels and a power limit of the
4 processor.

- 1 3. The processor of claim 2, wherein the power controller is to use the
2 first power consumption level to control the at least one of the operating frequency
3 and the voltage if the first power consumption level is less than a threshold, and
4 otherwise use the second power consumption level to control the at least one of the
5 operating frequency and the voltage.

- 1 4. The processor of claim 1, 2 or 3, wherein the first sensor comprises a
2 logic to receive counter information from a plurality of event counters associated with
3 each of the plurality of cores and to measure the first power consumption level based
4 on the counter information.

- 1 5. The processor of claim 1, 2 or 3, wherein the hybrid logic is to combine
2 the first power consumption level and the second power consumption level by:
3 generation of a first correction factor using the first power consumption level
4 and the second power consumption level; and
5 generation of a second correction factor using the second power consumption
6 level and the first correction factor.

1 6. The processor of claim 5, wherein the hybrid logic is to determine the
2 hybrid power consumption level using the second power consumption level and the
3 first and second correction factors.

1 7. The processor of claim 1, 2 or 3, wherein the hybrid logic is to execute
2 a low power load on the processor and calculate the first and second power
3 consumption levels using sensor information from the first sensor and the dynamic
4 current, respectively, and determine and store an offset based on the first and
5 second power consumption levels.

1 8. The processor of claim 7, wherein the hybrid logic is to execute a high
2 power load on the processor and calculate the second power consumption level
3 using the dynamic current, and determine and store a slope based on the offset and
4 the second power consumption level.

1 9. The processor of claim 8, wherein the hybrid logic is to receive the
2 dynamic current and calculate the second power consumption level using the
3 dynamic current, and determine the hybrid power consumption level using second
4 power consumption level, the slope and the offset.

1 10. A method comprising:
2 receiving, in a first logic of a processor, sensor information from a digital
3 power meter of the processor, and calculating a first power consumption level of the
4 processor using the sensor information from the digital power meter;
5 receiving, in the first logic, sensor information from a current sensor
6 configured to measure a current delivered by a voltage regulator coupled to the
7 processor, and calculating a second power consumption level of the processor using
8 the sensor information from the current sensor; and
9 combining, in the first logic, the first and second power consumption levels to
10 obtain a hybrid power consumption level of the processor.

1 11. The method of claim 10, further comprising controlling at least one of
2 an operating frequency and a voltage of the processor based on the hybrid power
3 consumption level and a power limit of the processor.

1 12. The method of claim 11, wherein combining the first and second
2 power consumption levels comprises:
3 if the first power consumption level is less than a threshold, using the first
4 power consumption level to control the at least one of the operating frequency and
5 the voltage; and
6 otherwise, using the second power consumption level to control the at least
7 one of the operating frequency and the voltage.

1 13. The method of claim 11, wherein combining the first and second power
2 consumption levels comprises:
3 generating a first correction factor using the first power consumption level and
4 the second power consumption level; and
5 generating a second correction factor using the second power consumption
6 level and the first correction factor.

1 14. The method of claim 13, further comprising determining the hybrid
2 power consumption level using the second power consumption level and the first and
3 second correction factors.

1 15. The method of claim 14, further comprising controlling at least one of
2 an operating frequency and a voltage of the processor based on the hybrid power
3 consumption level.

1 16. A system comprising:
2 a multicore processor including a plurality of cores to independently execute
3 instructions, each of the plurality of cores including at least one event counter to
4 count events occurring on the core, a digital power meter to calculate a first power
5 consumption level based on information from the event counters, a second power

6 meter to calculate a second power consumption level based on a current delivered to
7 the multicore processor from a voltage regulator, and a power controller including a
8 first logic to generate a combined power consumption level of the multicore
9 processor using the first and second power consumption levels;
10 the voltage regulator coupled to the multicore processor to provide a
11 regulated voltage to the multicore processor; and
12 a dynamic random access memory (DRAM) coupled to the multicore
13 processor.

1 17. The system of claim 16, wherein the first logic is to:
2 cause a low power load level to occur on the multicore processor and
3 calculation of the first and second power consumption levels during the low power
4 load level; and
5 determine and store an offset based on first and second power consumption
6 levels.

1 18. The system of claim 17, wherein the first logic is to:
2 cause a high power load level to occur on the multicore processor and
3 calculation of the second power consumption level during the high power load level;
4 and
5 determine and store a slope based on the offset and the second power
6 consumption level.

1 19. The system of claim 18, wherein the first logic is to receive the second
2 power consumption level during normal operation of the system and determine the
3 hybrid power consumption level using the second power consumption level, the
4 slope and the offset.

1 20. A communication device arranged to perform the method of claims 11,
2 12, 13, 14 or 15.

1 21. At least one machine readable medium comprising a plurality of
2 instructions that in response to being executed on a computing device, cause the
3 computing device to carry out a method according to any one of claims 11, 12, 13,
4 14 or 15.

1 22. An apparatus for processing instructions, configured to perform the
2 method of any one of claims 11, 12, 13, 14 or 15.

1 23. An apparatus comprising means for performing the method of claims
2 11, 12, 13, 14 or 15.

1 24. A processor means comprising:
2 a plurality of execution means each for independently executing instructions;
3 first sensor means for measuring a first power consumption level of the
4 processor means based at least in part on a plurality of events occurring on the
5 plurality of execution means; and
6 means for combining the first power consumption level and a second power
7 consumption level of the processor means determined based on a dynamic current
8 provided to the processor means.

1 25. The processor means of claim 24, further comprising controller means
2 for controlling at least one of an operating frequency and a voltage of the processor
3 means based on the combined first and second power consumption levels and a
4 power limit of the processor means.

1 26. The processor means of claim 24 or 25, further comprising regulator
2 means for providing a regulated voltage to the processor means, the regulator
3 means further providing a level of the dynamic current provided to the processor
4 means.

100

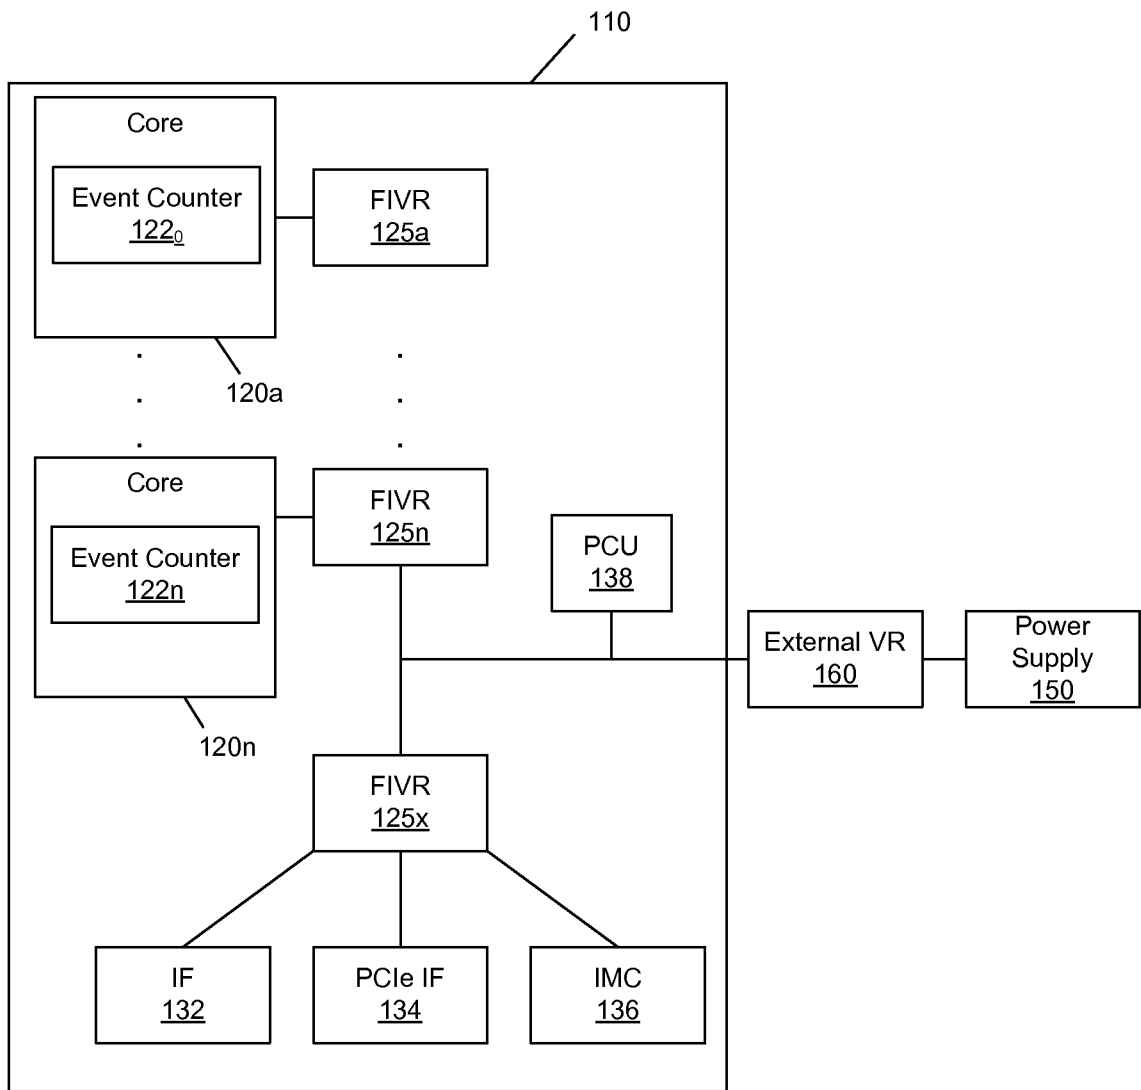


FIG. 1

2/10

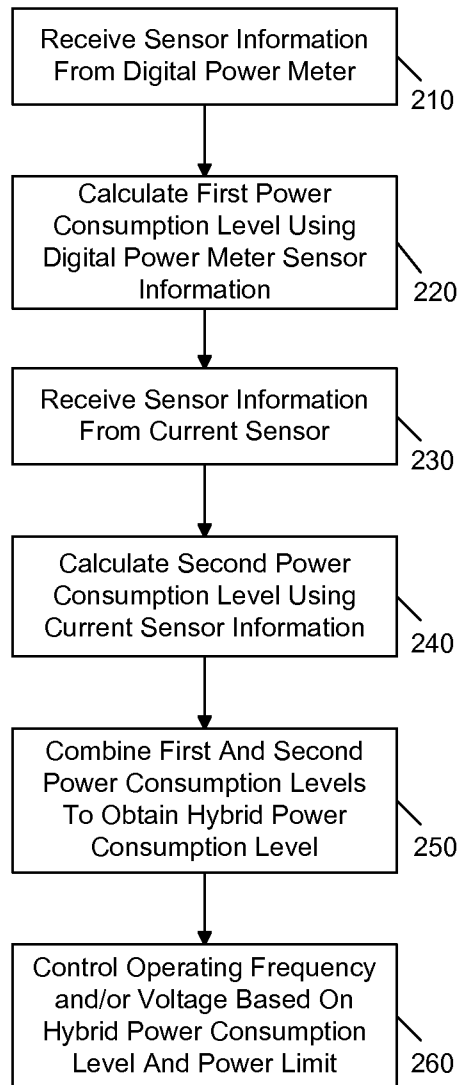
200

FIG. 2

300

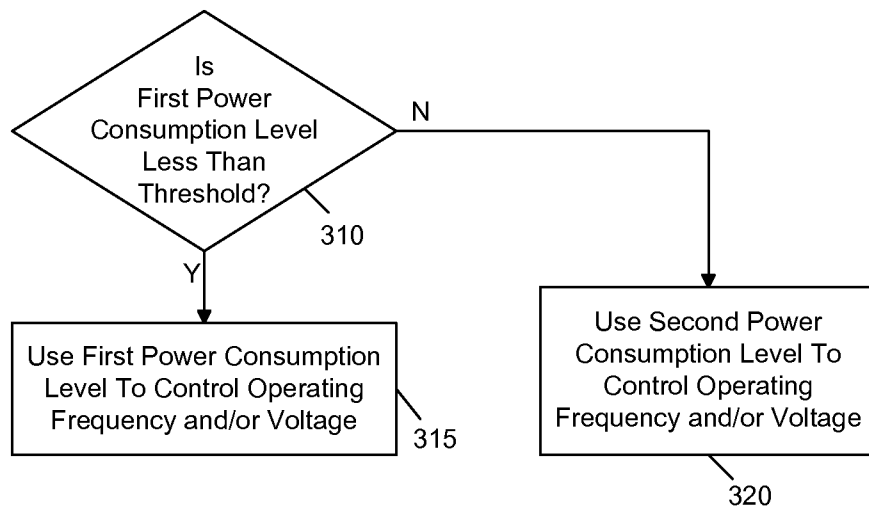


FIG. 3

4/10

330

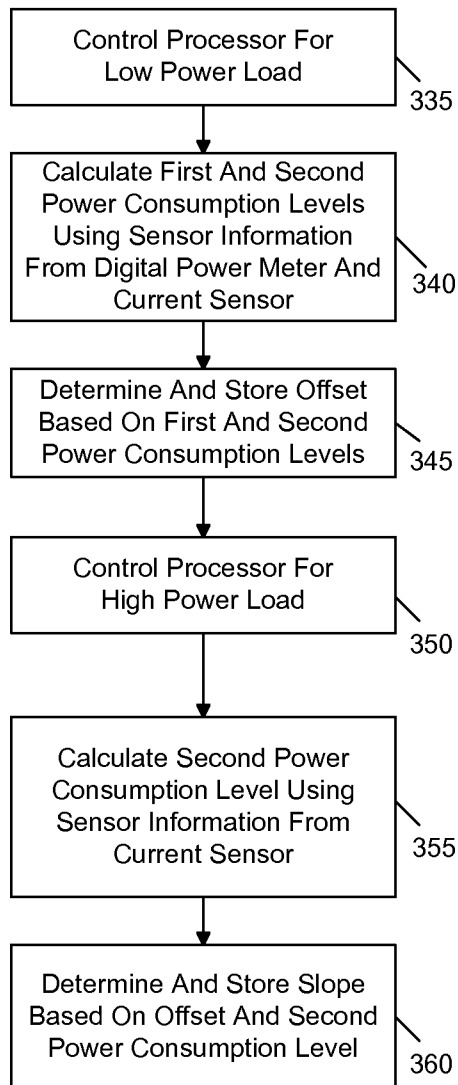


FIG. 4

370

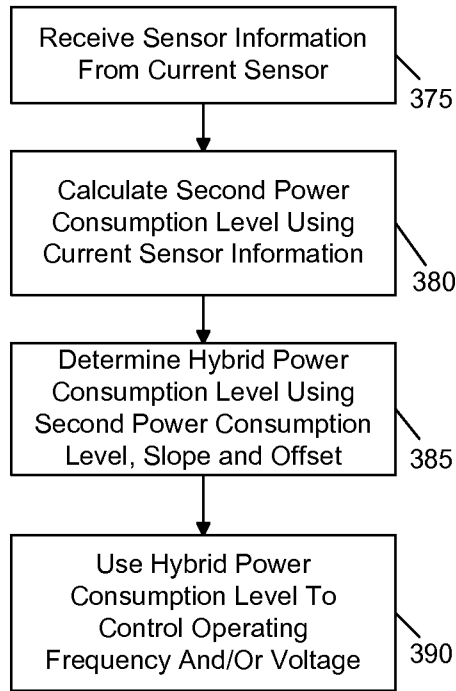


FIG. 5

400




FIG. 6

500




FIG. 7




FIG. 8

710




FIG. 9




FIG. 10

A. CLASSIFICATION OF SUBJECT MATTER**G01R 31/28(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)

G01R 31/28; F25B 19/02; G06F 1/26; F25D 23/12; G06F 1/00

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

eKOMPASS(KIPO internal) & Keywords: core, power, consumption, event, counter, voltage regulator, combine, logic, core, power consumption, event counter, voltage regulator, combine, logic

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2006-0090086 A1 (EFRAIM ROTEM et al.) 27 April 2006 See abstract, paragraphs [0013]-[0021], claim 1, and figures 1,2A,3.	1-26
A	US 2010-0115293 A1 (EFRAIM ROTEM et al.) 06 May 2010 See abstract, paragraphs [0017]-[0018], and figures 1,3.	1-26
A	US 2005-0210905 A1 (JAMES S. BURNS et al.) 29 September 2005 See abstract, paragraphs [0011]-[0017], claim 1, and figure 1.	1-26
A	US 2009-0006871 A1 (YEN-CHENG LIU et al.) 01 January 2009 See abstract, paragraphs [0042]-[0048], claim 1, and figure 4.	1-26

 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 citation or other special reason (as specified)

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

14 NOVEMBER 2012 (14.11.2012)

Date of mailing of the international search report

16 NOVEMBER 2012 (16.11.2012)

Name and mailing address of the ISA/KR

Korean Intellectual Property Office
189 Cheongsu-ro, Seo-gu, Daejeon Metropolitan
City, 302-701, Republic of Korea

Facsimile No. 82-42-472-7140

Authorized officer

PARK, Keun Yong

Telephone No. 82-42-481-8508



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2012/031464

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2006-0090086 A1	27.04.2006	CN 101036103 A0	12.09.2007
		CN 101036103 B	26.05.2010
		TW 1309351A	01.05.2009
		TW 1309351B	01.05.2009
		US 7430672 B2	30.09.2008
		WO 2006-049690 A2	11.05.2006
		WO 2006-049690 A3	29.06.2006
		WO 2006-049690 A3	11.05.2006
		US 2010-0115293 A1	06.05.2010
TW 201032038 A	01.09.2010		
TW 201032038 A	01.09.2010		
US 2005-0210905 A1	29.09.2005	TW 1264515B	21.10.2006
		TW 264515 B	21.10.2006
		US 7062933 B2	20.06.2006
US 2009-0006871 A1	01.01.2009	CN 101334688 A	31.12.2008
		DE 112008001655 T5	22.07.2010
		GB 0921943 D0	03.02.2010
		GB 2463193 A	10.03.2010
		JP 2009-009577 A	15.01.2009
		RU 2010100908 A	20.07.2011
		US 7971074 B2	28.06.2011
		WO 2009-006014 A1	08.01.2009