



(51) International Patent Classification:

G06F 1/32 (2006.01) G06F 15/16 (2006.01)
G06F 12/08 (2006.01)

(21) International Application Number:

PCT/US2013/054629

(22) International Filing Date:

13 August 2013 (13.08.2013)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

13/594,410 24 August 2012 (24.08.2012) US

(71) Applicant: **ADVANCED MICRO DEVICES, INC.**
[CA/CA]; One Commerce Valley Drive East, Markham,
ON L3T 7N6 (CA).

(72) Inventors; and

(71) Applicants: **KITCHIN, Paul** [US/US]; 4708 Trail Crest
Dr., Austin, TX 78735 (US). **WALKER, William, L.**
[US/US]; 2343 Timber Creek Drive, Fort Collins, CO
80528 (US). **KOMMRUSCH, Steven, J.** [US/US]; 6213
Eagle Ridge Ct., Fort Collins, CO 80525 (US).

(74) Agent: **EICH, Raymund, F.**; Williams, Morgan & Amer-
son, P.C., 10333 Richmond, Suite 1100, Houston, TX
77042 (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, IS, JP, KE, KG, KN, KP, KR,
KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME,
MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ,
OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA,
SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM,
TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM,
ZW.

(84) Designated States (unless otherwise indicated, for every
kind of regional protection available): ARIPO (BW, GH,
GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, 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).

Published:

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

(54) Title: POWER MANAGEMENT OF MULTIPLE COMPUTE UNITS SHARING A CACHE




FIGURE 1

(57) Abstract: We report methods, integrated circuit devices, and fabrication processes relating to power management transitions of multiple compute units sharing a cache. One method includes indicating that a first compute unit of a plurality of compute units of an integrated circuit device is attempting to enter a low power state, determining if the first compute unit is the only compute unit of the plurality in a normal power state, and in response to determining the first compute unit is the only compute unit in the normal power state: saving a state of a shared cache unit of the integrated circuit device, flushing at least a portion of a cache of the shared cache unit, repeating the flushing until either a second compute unit exits the low power state or the cache is completely flushed, and permitting the first compute unit to enter the low power state.

POWER MANAGEMENT OF MULTIPLE COMPUTE UNITS SHARING A CACHE**BACKGROUND****1. TECHNICAL FIELD**

Generally, the disclosed embodiments relate to integrated circuits, and, more particularly, to power management of multiple compute units sharing memory, such as cache memory.

2. DESCRIPTION OF THE RELATED ART

A computer system comprising two or more compute units (e.g., cores of a central processing unit (CPU)) can place those compute units into a lower power state when they are not needed to perform user- or system-requested operations. Placing unneeded compute units into a lower power state may reduce power consumption and heat generation by the computer system, thereby reducing operating expenses of the computer system and extending the service life of the computer system or components thereof. It is common for a computer system to contain a central power management unit (PMU) to orchestrate the low power transitions for each of the compute units within the system. Typically, the PMU can make requests directly to each compute unit to independently power down and power up.

At times, the compute units may share a common memory, such as a cache memory. When a compute unit is directed to power down, one issue to be addressed is the problem of shutting down any associated cache memory that may be shared with another compute unit. Commonly, when a compute unit is directed to power down, the compute unit will save off its architectural state to some memory retention area, flush its caches of all modified data (i.e., complete any writing of modified data from dirty cache locations to main memory and evict the modified data from the cache), and then signal its low power readiness to the PMU. At this point, the PMU will turn off power to that compute unit. When the PMU requires the compute unit to power up (e.g., exit a lower power state or enter a normal power state) to

service a process, the PMU will turn on power to the compute unit, and the compute unit will restore its architectural state from the memory retention area and start servicing the process.

A shared cache unit (SCU) is sometimes used within e.g. a CPU system so that all the included compute units can share cache resources. If an existing CPU system design were to add a SCU to its architecture, it can be beneficial for the PMU to directly interact with the SCU via one interface and for that SCU to interact with all the CPU cores directly. However, with such a topology, a legacy PMU might not “understand” the extra level of hierarchy containing the SCU, thereby possibly leading to incomplete or ineffectual power management. Reengineering a PMU/SCU/multiple cores system to understand a hierarchy containing a SCU would be a relatively complicated task which the person of ordinary skill in the art would wish to avoid. Further, engineering other elements of the computer system to aid the PMU in power management of an SCU/multiple cores system runs the risk of requiring a large number of interactions of possibly long latency, which could undesirably slow down power management transitions.

SUMMARY OF EMBODIMENTS

The apparatuses, systems, and methods in accordance with the embodiments disclosed herein may facilitate power management of multiple compute units sharing an SCU by a communication process between a compute unit and a shared cache unit. Mechanisms controlling and implementing such a process may be formed within a microcircuit by any means, such as by growing or deposition.

Some embodiments provide an integrated circuit device that includes a plurality of compute units, a shared cache unit for caching data usable by two or more of the plurality of compute units, a power management unit for requesting each of the plurality of compute units to enter or exit a particular state (e.g., a low power state, a normal power state, etc.), and a

memory for storing at least a state of the shared cache unit. In these embodiments, the integrated circuit device is configured to: indicate that a first compute unit of a plurality of compute units of an integrated circuit device is attempting to enter the low power state, determine if the first compute unit is the only compute unit of the plurality in the normal power state, permit the first compute unit to enter the low power state, in response to determining the first compute unit is not the only compute unit in the normal power state, and in response to determining the first compute unit is the only compute unit in the normal power state: save a state, such as a configuration register state, of a shared cache unit of the integrated circuit device, flush at least a portion of a cache of the shared cache unit, repeat the flushing until either a second compute unit exits the low power state or the cache is completely flushed, and permit the first compute unit to enter the low power state.

Some embodiments provide a method that includes indicating that a first compute unit of a plurality of compute units of an integrated circuit device is attempting to enter a low power state, determining if the first compute unit is the only compute unit of the plurality in a normal power state, permitting the first compute unit to enter the low power state, in response to determining the first compute unit is not the only compute unit in the normal power state, and in response to determining the first compute unit is the only compute unit in the normal power state: saving a state of a shared cache unit of the integrated circuit device, flushing at least a portion of a cache of the shared cache unit, repeating the flushing until either a second compute unit exits the low power state or the cache is completely flushed, and permitting the first compute unit to enter the low power state.

Some embodiments provide a method that includes indicating that a first compute unit of a plurality of compute units of an integrated circuit device is attempting to exit a low power state, determining if all other compute units of the plurality are in the low power state, permitting the first compute unit to exit the low power state, in response to determining at

least one other compute unit is in a normal power state, and in response to determining all other compute units are in the low power state: restoring a state of a shared cache unit of the integrated circuit device, and permitting the first compute unit to exit the low power state.

The embodiments described herein may be used in any type of integrated circuit that uses multiple compute units, a shared cache unit, and a power management unit. One example is a general purpose microprocessor.

BRIEF DESCRIPTION OF THE FIGURES

The disclosed subject matter will hereafter be described with reference to the accompanying drawings, wherein like reference numerals denote like elements, and:

Figure 1 is a schematic diagram of an exemplary microcircuit design in accordance with some embodiments.

Figure 2 is a schematic diagram of a central processing unit depicted in Figure 1, in accordance with some embodiments.

Figure 3A provides a representation of a silicon die/chip that includes one or more circuits as shown in Figure 2, in accordance with some embodiments.

Figure 3B provides a representation of a silicon wafer which includes one or more dies/chips that may be produced in a fabrication facility, in accordance with some embodiments.

Figure 4 is a flowchart of a method relating to powering up or down a shared cache unit, in accordance with some embodiments.

Figure 5 is a flowchart of a method relating to powering up a shared cache unit, in accordance with some embodiments.

Figure 6 is a flowchart of a method relating to entry of a compute unit into a low power state, in accordance with some embodiments.

Figure 7 is a flowchart of a method relating to entry of a compute unit into a normal power state, in accordance with some embodiments.

While the disclosed subject matter is susceptible to various modifications and alternative forms, specific embodiments thereof have been shown by way of example in the drawings and are herein described in detail. It should be understood, however, that the description herein of specific embodiments is not intended to limit the disclosed subject matter to the particular forms disclosed, but on the contrary, the intention is to cover all modifications, equivalents, and alternatives falling within the spirit and scope of the disclosed subject matter as defined by the appended claims.

DETAILED DESCRIPTION

Embodiments provide for facilitated power management of multiple compute units sharing an SCU. Because communications between a compute unit and an associated SCU typically have long latency, various embodiments of facilitated power management allow a reduced set of communications between each compute unit and the SCU. Thereby, power state transitions may be performed at an acceptably high speed. Various embodiments also substantially or completely hide SCU hierarchy from the PMU, thereby reducing the design changes for a computer system incorporating an SCU shared by multiple cores.

Turning now to Figure 1, a block diagram representation of a computer system comprising a processor, in accordance with some embodiments, is illustrated. Modern computer systems may exist in a variety of forms, such as telephones, tablet computers, desktop computers, laptop computers, servers, smart televisions, or other consumer electronic devices. The processor unit 110 may comprise one or more central processing units (CPUs) 140, each of which may comprise various compute units (CUs) 135, e.g. first CU 135a, second CU 135b, through Nth CU 135c). Each of the CPUs 140 may also comprise an internal cache 130 that provides memory resources for the CUs 135. The processor unit 110

may also comprise a shared cache unit (SCU) 152, which provides a shared cache 151 of memory resources shared among the various CPUs 140, compute units 135 of one or more CPUs 140, and/or graphics processing units (GPUs) 125 of a graphics card 120.

The SCU 152 may comprise a shared cache unit handshake register (HR) 153, which may facilitate communications between the SCU 152 and the computer units 135, CPUs 145, and/or the GPUs 125. Utilizing the shared cache unit HR 153, the compute units 135, CPUs 140, and/or GPUs 125 may write data into the SCU 152, and read data from the SCU 152, when performing various power management operations, as exemplified in more detail below. In some embodiments, each of the compute units 135, CPUs 140, and/or GPUs 125 may write the following values to the shared cache unit handshake register 153: request entry of low power state, request entry of low power state and flush of SCU cache, and request exit of low power state. The SCU 152 may write to, and the compute units 135, CPUs 140, and/or GPUs 125 read from, the shared cache unit handshake register 153 the following responses: entry of low power state permitted, entry of low power state forbidden, exit of low power state permitted, exit of low power state forbidden.

The computer system 100 may also comprise a northbridge 145. Among its various components, the northbridge 145 may comprise a power management unit (PMU) 132 that may regulate the amount of power consumed by the compute units 135, CPUs 140, GPUs 125, and/or the SCU 152. Particularly, in response to changes in demand for the compute units 135, CPUs 140, and/or GPUs 125, the PMU 132 may request each of the plurality of compute units 135, CPUs 140, and/or GPUs 125 to enter a low power state, exit the low power state, enter a normal power state, or exit the normal power state.

In some embodiments, the PMU 132 directly interacts solely with the SCU 152. Power management directives intended for the compute units 135, CPUs 140, and/or GPUs

125 are then relayed by the SCU 152 to the target unit(s) 135, 140, and/or 125. In some embodiments, the PMU 132 may interact with other portion(s) of the computer system.

The computer system 100 may also comprise a DRAM 155. The DRAM 155 may be configured to store one or more states of one or more components of the computer system 100. Particularly, the DRAM 155 may be configured to store one or more states of the SCU 152, one or more states of the compute units 135, one or more states of one or more CPUs 140, and/or one or more states of one or more GPUs 125. For example, the DRAM 155 may be configured to store a configuration register state of the SCU 152. In another example, the DRAM 155 may be configured to store an architectural state of a compute unit 135a, 135b, or 135c.

The computer system 100 may as a routine matter comprise other known units and/or components, e.g., one or more I/O interfaces 131, a southbridge 150, a data storage unit 160, display unit(s) 170, input device(s) 180, output device(s) 185, and/or peripheral devices 190, among others.

The computer system 100 may also comprise one or more data channels 195 for communication between one or more of the components described above.

Turning now to Figure 2 and Figure 3A, in some embodiments, the processor unit 110 may reside on a silicon die/chip 340. The silicon die/chip 340 may be housed on a motherboard or other structure of the computer system 100. In some embodiments, there may be more than one processor unit 110 on each silicon die/chip 340. Some embodiments of the processor unit 110 may be used in a wide variety of electronic devices.

Turning now to Figure 3B, in accordance with some embodiments, and as described above, the processor unit 110 may be included on the silicon chip/die 340. The silicon chip/die 340 may contain one or more different configurations of the processor unit 110. The silicon chip/die 340 may be produced on a silicon wafer 330 in a fabrication facility (or

“fab”) 390. That is, the silicon wafer 330 and the silicon die/chip 340 may be referred to as the output, or product of, the fab 390. The silicon chip/die 340 may be used in electronic devices.

The circuits described herein may be formed on a semiconductor material by any known means in the art. Forming can be done, for example, by growing or deposition, or by any other means known in the art. Different kinds of hardware descriptive languages (HDL) may be used in the process of designing and manufacturing the microcircuit devices. Examples include VHDL and Verilog/Verilog-XL. In some embodiments, the HDL code (e.g., register transfer level (RTL) code/data) may be used to generate GDS data, GDSII data and the like. GDSII data, for example, is a descriptive file format and may be used in some embodiments to represent a three-dimensional model of a semiconductor product or device. Such models may be used by semiconductor manufacturing facilities to create semiconductor products and/or devices. The GDSII data may be stored as a database or other program storage structure. This data may also be stored on a computer readable storage device (e.g., data storage units, RAMs, compact discs, DVDs, solid state storage and the like) and, in some embodiments, may be used to configure a manufacturing facility (e.g., through the use of mask works) to create devices capable of embodying various aspects of some embodiments. As understood by one of ordinary skill in the art, this data may be programmed into a computer, processor, or controller, which may then control, in whole or part, the operation of a semiconductor manufacturing facility (or fab) to create semiconductor products and devices. In other words, some embodiments relate to a non-transitory computer-readable medium storing instructions executable by at least one processor to fabricate an integrated circuit. These tools may be used to construct the embodiments described herein.

Figure 4 presents a flowchart depicting a method 400 according to some embodiments. As illustrated in Figure 4, the method 400 may comprise: monitoring (at 410)

activity of one or more CPUs (or compute unit(s), GPU(s), etc. For convenience, the term “CPU” will be used in further description of this figure). For example, monitoring may comprise monitoring activity of at least two components each selected from central processing units (CPUs), graphics processing units (GPUs), or compute units of an integrated circuit device. Upon a determination (at 420) that activity of the one or more CPUs is such that a shared cache may be powered down (e.g., sufficiently low activity of the component(s) is occurring), a shared cache power-down process may be performed (at 430). A more detailed description of some embodiments of the shared cache power-down process is provided below. If the determination (at 420) is that the shared cache may not be powered down, flow may be returned to monitoring (at 410). When a determination is made (at 440) that the shared cache is required to be powered up (e.g., the shared cache is needed by at least one of the component(s)), a shared cache power-up process may be performed (at 450), with flow returning to monitoring (at 410). If the determination (at 440) is that the shared cache is not required to be powered up, then the powered-down mode may be continued (at 460), with flow returning to determining (at 440).

Figure 5 presents a flowchart depicting a method 500 according to some embodiments. As illustrated in Figure 5, the method 500 may comprise: receiving (at 510) a request from a power management unit (PMU) for a power transition (e.g., from normal power to low power) (which may be termed “receiving a power transition request”). Upon receiving the request, the computer system (at 520) may check with each first compute unit (or CPU, GPU, etc.) of a plurality of compute units of an integrated circuit device to determine if a low power state is possible (e.g., whether powering down a shared cache is possible). (For convenience, the term “compute unit” will be used in further description of this figure). If a power down is determined (at 530) to be possible, then checking (at 540) with each next compute unit may be performed to determine if a low power state is possible.

If the power down is determined (at 560) to be possible, then a determination is made whether there are any additional compute units left to check to see if low power state is possible (at 570). If there is a determination (at 570) that no additional compute units need to be checked regarding the possibility of a low power state (e.g., whether powering down a shared cache is possible), then a shared cache may be powered down (at 580). If a determination is made that additional compute units need to be checked, then flow returns to checking (at 540). If either of the determinations (at 530 or 560) find that power down is not possible, then an indication that the shared cache power-down cannot be implemented can be made (at 550), with flow returning to receipt (at 510) of power transition requests from the PMU.

Figure 6 presents a flowchart depicting a method 600 according to some embodiments. As illustrated in Figure 6, the method 600 may comprise a request to place the first compute unit (CU) of a plurality of compute units of an integrated circuit device, into a low power state (at 602). In some embodiments, this request may be made by the PMU 132. Upon such a request, information relating to the state of the first CU is saved (at 604). For example, the state of the first CU saved at 604 may be an architectural state. The cache associated with the first CU may be flushed (at 606). An indication (at 610) may be received by the SCU 152 that the first CU is attempting to enter the low power state. A determination (at 620) may be made whether the first compute unit is the only compute unit of the plurality in a normal power state. Upon a determination that the first CU is the only CU in a normal power state, the state (e.g., a configuration register state) of the shared cache unit 152 is saved (at 630) to memory, e.g. to DRAM 155. At least a portion of the cache associated with SCU 152 is flushed (at 640). In some embodiments, this flushing process may be repeated until either a second compute unit exits the low power state (as determined at 650) or the cache is completely flushed (as determined at 660). Upon a determination that the second CU has

exited a low power state (at 650), or a determination that the SCU cache is completely flushed (at 660), the first compute unit may be permitted to enter the low power state (at 670). That is, the first compute unit may be permitted to enter the low power state in response to determining at 620 the first compute unit is not the only compute unit in the normal power state; or, in response to determining at 650 and 660 the first compute unit is the only compute unit in the normal power state and the SCU cache is completely flushed.

The size of the portion of the cache of the SCU flushed at 640 may be chosen in view of a tradeoff between reducing the latency of the process (which argues for reducing the sizes of the portions) and reducing the overhead of write/flush/read cycles (which argues for increasing the sizes of the portions). Also, the first compute unit must be able to respond to interrupts, e.g., those arising when a second compute unit exits the low power state as determined at 650. In some embodiments, this may be another reason for reducing the size of the portions.

Also as depicted in Figure 6, upon permitting the first CU to enter into a low power state (at 670), the first compute unit (at 672) may report (e.g., to the PMU 132) that it is permitted to enter the low power state subsequent the permitting (at 670); and the first CU (at 674) may enter into the low power state.

In some embodiments, the SCU 152 may enter into a low power state if all compute units are in a low power state. In some embodiments, indicating (at 610) comprises writing to a shared cache unit handshake register 153.

In some embodiments, various steps shown in Figure 6 (e.g., 602, 604, 610, etc.) may be omitted from Figure 6 and remain within the spirit and scope of the embodiments.

Figure 7 presents a flowchart depicting a method 700 according to some embodiments. An indication may be made (at 710) that a first compute unit (CU) of a plurality of compute units of an integrated circuit device is attempting to exit a low power

state. The prompt to exit the low power state may be the result of one or may possible actions, such as a message from PMU 132. A determination (at 720) is made whether all other compute units of the plurality are in the low power state. If not all other compute units are in the low power state, the first compute unit is permitted (at 750) to exit the low power state. That is, if at least one other compute unit is in a normal power state, the first compute unit is permitted (at 750) to exit the low power state. This is allowed since if at least one compute unit is in the normal power state, the state of the SCU 152 is intact and/or has been previously restored, and the CU in the low power state can be brought up to the normal power state without needing to restore the state of the SCU 152.

In the event a determination (at 720) is made that all other compute units are in the low power state (i.e., not even a single CU is in a normal power state), at 730, a state (e.g., a configuration register state) of a shared cache unit (SCU) of the integrated circuit device is restored. This is performed because when all compute units are in a low power mode, the SCU 152 may also have been in a power/sleep mode. In some embodiments (such as embodiments wherein there is no need to restore a state, such as an architectural state, of the first compute unit), upon restoration of the state (e.g., the configuration register state) of the SCU 152 from memory (e.g., from DRAM 155), the first CU may be permitted to exit the low power state.

In some embodiments, upon the determination (at 720) that all other compute units are in the low power state, a state (e.g., an architectural state) of the first compute unit may be restored (at 740) from memory (e.g., from DRAM 155), after which the first compute unit may be permitted to exit the low power state.

In some embodiments, indicating that the first CU is attempting to exit a low power state (at 710) comprises writing to a shared cache unit handshake register 153.

The methods illustrated in Figures 4-7 may be governed by instructions that are stored in a non-transitory computer readable storage medium and that are executed by at least one processor of the computer system 100. Each of the operations shown in Figures 4-7 may correspond to instructions stored in a non-transitory computer memory or computer readable storage medium. In various embodiments, the non-transitory computer readable storage medium includes a magnetic or optical disk storage device, solid state storage devices such as Flash memory, or other non-volatile memory device or devices. The computer readable instructions stored on the non-transitory computer readable storage medium may be in source code, assembly language code, object code, or other instruction format that is interpreted and/or executable by one or more processors.

The particular embodiments disclosed above are illustrative only, as the disclosed subject matter may be modified and practiced in different but equivalent manners apparent to those skilled in the art having the benefit of the teachings herein. Furthermore, no limitations are intended to the details of construction or design herein shown, other than as described in the claims below. It is therefore evident that the particular embodiments disclosed above may be altered or modified and all such variations are considered within the scope and spirit of the disclosed subject matter. Accordingly, the protection sought herein is as set forth in the claims below.

CLAIMS

We claim:

1. A method, comprising:
 - indicating that a first compute unit of a plurality of compute units of an integrated circuit device is attempting to enter a low power state;
 - determining if the first compute unit is the only compute unit of the plurality of compute units in a normal power state; and
 - when the first compute unit is the only compute unit in the normal power state:
 - saving a state of a shared cache unit of the integrated circuit device;
 - flushing a portion of a cache of the shared cache unit;
 - repeating the flushing until either a second compute unit exits the low power state or the cache is completely flushed; and
 - permitting the first compute unit to enter the low power state.
2. The method of claim 1, wherein prior to the indicating, the method further comprises:
 - requesting the first compute unit to enter the low power state;
 - saving a state of the first compute unit; and
 - flushing a cache of the first compute unit.
3. A method, comprising:
 - indicating that a first compute unit of a plurality of compute units of an integrated circuit device is attempting to exit a low power state;
 - determining if all other compute units of the plurality of compute units are in the low power state; and

when all other compute units are in the low power state:

restoring a state of a shared cache unit of the integrated circuit device; and

permitting the first compute unit to exit the low power state.

4. The method of claim 3, further comprising:

restoring a state of the first compute unit prior to permitting the first compute unit to exit the low power state.

5. An integrated circuit device, comprising:

a plurality of compute units;

a shared cache unit configured to cache data usable by two or more of the plurality of compute units;

a power management unit configured to request that a respective compute unit of the plurality of compute units enter a low power state, exit the low power state, enter a normal power state, or exit the normal power state; and,

a memory configured to store at least a state of the shared cache unit;

wherein a first compute unit of the plurality of compute units is configured to indicate that the first compute unit is attempting to enter the low power state, and, when the first compute unit is the only compute unit in the normal power state, to save a state of a shared cache unit to the memory, to flush a portion of a cache of the shared cache unit; and to repeat the flushing until either a second compute unit exits the low power state or the cache is completely flushed;

wherein the shared cache unit is further configured to determine if the first compute unit is the only compute unit of the plurality in the normal power state; and

wherein the power management unit is further configured to permit the first compute unit to enter the low power state upon completion of any required flushing by the first compute unit.

6. The integrated circuit device of claim 5, wherein the power management unit is further configured to, prior to the indicating, request the first compute unit to enter the low power state; and

wherein the first compute unit of the plurality of compute units is further configured to save a state of the first compute unit and flush a cache of the first compute unit subsequent the requesting and prior to the indicating.

7. An integrated circuit device, comprising:

a plurality of compute units;

a shared cache unit configured to cache data usable by two or more of the plurality of compute units;

a power management unit configured to request each of the plurality of compute units to enter a low power state, exit the low power state, enter a normal power state, or exit the normal power state; and,

a memory configured to store at least a state of the shared cache unit;

wherein a first compute unit is configured to indicate that the first compute unit is attempting to exit the low power state and, when all other compute units are in the low power state, to restore a state of the shared cache unit;

wherein the shared cache unit is configured to determine if all other compute units of the plurality are in the low power state; and

wherein the power management unit is configured to permit the first compute unit to exit the low power state subsequent restoring the state of the shared cache unit.

8. The integrated circuit device of claim 7, wherein the first compute unit is further configured to restore a state of the first compute unit from memory prior to permitting the first compute unit to exit the low power state.

9. A non-transitory computer-readable medium storing instructions executable by at least one processor to fabricate an integrated circuit, the integrated circuit comprising:

a plurality of compute units;

a shared cache unit configured to cache data usable by two or more of the plurality of compute units;

a power management unit configured to request that a respective compute unit of the plurality of compute units enter a low power state, exit the low power state, enter a normal power state, or exit the normal power state; and,

a memory configured to store at least a state of the shared cache unit;

wherein a first compute unit of the plurality of compute units is configured to indicate that the first compute unit is attempting to enter the low power state and, when the first compute unit is the only compute unit in the normal power state, to save a state of a shared cache unit to the memory, to flush a portion of a cache of the shared cache unit; and to repeat the flushing until either a second compute unit exits the low power state or the cache is completely flushed;

wherein the shared cache unit is further configured to determine if the first compute unit is the only compute unit of the plurality in the normal power state; and
wherein the power management unit is further configured to permit the first compute unit to enter the low power state upon completion of any required flushing by the first compute unit.

10. A non-transitory computer-readable medium storing instructions executable by at least one processor to fabricate an integrated circuit, the integrated circuit comprising:
a plurality of compute units;
a shared cache unit configured to cache data usable by two or more of the plurality of compute units;
a power management unit configured to request each of the plurality of compute units to enter a low power state, exit the low power state, enter a normal power state, or exit the normal power state; and,
a memory configured to store at least a state of the shared cache unit;
wherein a first compute unit of the plurality of compute units is configured to indicate that the first compute unit is attempting to exit the low power state and, when all other compute units are in the low power state, to restore a state of the shared cache unit;
wherein the shared cache unit is further configured to determine if all other compute units of the plurality are in the low power state; and
wherein the power management unit is further configured to permit the first compute unit to exit the low power state subsequent restoring the state of the shared cache unit.




FIGURE 1




FIGURE 3B




FIGURE 4




FIGURE 5




FIGURE 6

700




FIGURE 7

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2013/054629

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F1/32 G06F12/08 G06F15/16
ADD.

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06F

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

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

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 1 607 838 A2 (INTEL CORP [US]) 21 December 2005 (2005-12-21) abstract figure 1 figure 2A figure 3 figure 5 figure 7 paragraphs [0006] - [0009] paragraphs [0010] - [0026] paragraphs [0027] - [0028] paragraph [0029] paragraphs [0035] - [0042] ----- -/-	1-10



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

19 November 2013

Date of mailing of the international search report

06/12/2013

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

De Ceulaer, Bart

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2013/054629

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2008/077813 A1 (KELLER JAMES B [US] ET AL) 27 March 2008 (2008-03-27) paragraphs [0022] - [0027] figure 1 paragraph [0035] figure 2 paragraphs [0048] - [0051] figure 5 -----	1-10
A	"Hibernation (computing)", wikipedia.org 23 July 2012 (2012-07-23), XP002716587, Retrieved from the Internet: URL: http://en.wikipedia.org/w/index.php?title=Hibernation_%28computing%29&oldid=503801427 [retrieved on 2013-11-19] the whole document -----	1-10
A	"Effective Flush Cache command implementation", IP.COM JOURNAL, IP.COM INC., WEST HENRIETTA, NY, US, 4 September 2006 (2006-09-04), XP013115713, ISSN: 1533-0001 the whole document -----	1,3,5,7, 9,10

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2013/054629

Patent document cited in search report		Publication date		Patent family member(s)		Publication date
EP 1607838	A2	21-12-2005	CN	1710518 A		21-12-2005
			EP	1607838 A2		21-12-2005
			KR	20050120565 A		22-12-2005
			KR	20070029766 A		14-03-2007

US 2008077813	A1	27-03-2008	US	2008077813 A1		27-03-2008
			US	2010235670 A1		16-09-2010
