




- (51) **International Patent Classification:**  
*G06F 1/32* (2006.01)      *G06F 15/167* (2006.01)
- (21) **International Application Number:**  
PCT/GB2011/050319
- (22) **International Filing Date:**  
17 February 2011 (17.02.2011)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**  
12/659,235      1 March 2010 (01.03.2010)      US
- (71) **Applicant (for all designated States except US):** **ARM LIMITED** [—/GB]; 110 Fulbourn Road, Cherry Hinton, Cambridge CB1 9NJ (GB).
- (72) **Inventors; and**
- (75) **Inventors/Applicants (for US only):** **GREENHALGH, Peter Richard** [GB/GB]; 39 Argyle Street, Cambridge, Cambridgeshire CB1 3LR (GB). **GRISENTHWAITE, Richard Roy** [GB/GB]; 2 Connor's Close, Guilden Morden, Nr Royston, Cambridgeshire SG8 0PT (GB).

- (74) **Agent:** **HORNER, David**; D Young & Co LLP, 120 Holborn, London EC1N 2DY (GB).
- (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, RO, RS, RU, 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, 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).

[Continued on next page]

(54) **Title:** A DATA PROCESSING APPARATUS AND METHOD FOR SWITCHING A WORKLOAD BETWEEN FIRST AND SECOND PROCESSING CIRCUITRY



(57) **Abstract:** A data processing apparatus and method are provided for switching performance of a workload between two processing circuits. The data processing apparatus has first processing circuitry which is architecturally compatible with second processing circuitry, but with the first processing circuitry being micro-architecturally different from the second processing circuitry. At any point in time, a workload consisting of at least one application and at least one operating system for running that application is performed by one of the first processing circuitry and the second processing circuitry. A switch controller is responsive to a transfer stimulus to perform a handover operation to transfer performance of the workload from source processing circuitry to destination processing circuitry, with the source processing circuitry being one of the first and second processing circuitry and the destination processing circuitry being the other of the first and second processing circuitry. The switch controller is arranged, during the handover operation, to cause the source processing circuitry to make its current architectural state available to the destination processing circuitry, the current architectural state being that state not available from shared memory shared between the first and second processing circuitry at a time the handover operation is initiated, and that is necessary for the destination processing circuitry to successfully take over performance of the workload from the source processing circuitry. Further, the source processing

[Continued on next page]



WO 2011/107777 A1

**Published:**

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

— before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

---

circuitry and second processing circuitry implement an accelerated mechanism to make the current architectural state available to the destination processing circuitry without routing of the current architectural state via the shared memory. Since the accelerated mechanism is quick and energy efficient, it increases the number of situations in which it is energy efficient to make the switch from one processing circuitry to the other.

## A DATA PROCESSING APPARATUS AND METHOD FOR SWITCHING A WORKLOAD BETWEEN FIRST AND SECOND PROCESSING CIRCUITRY

### FIELD OF THE INVENTION

5           The present invention relates to a data processing apparatus and method for switching a workload between first and second processing circuitry, and in particular to a technique for performing such switching so as to improve energy efficiency of the data processing apparatus.

### BACKGROUND OF THE INVENTION

10           In modern data processing systems, the difference in performance demand between high intensity tasks such as games operation and low intensity tasks such as MP3 playback can exceed a ratio of 100:1. For a single processor to be used for all tasks, that processor would have to be high performance, but an axiom of processor micro-architecture is that high performance processors are less energy efficient than  
15 low performance processors. It is known to improve energy efficiency at the processor level using techniques such as Dynamic Voltage and Frequency Scaling (DVFS) or power gating to provide the processor with a range of performance levels and corresponding energy consumption characteristics. However, such techniques are generally becoming insufficient to allow a single processor to take on tasks with such  
20 diverging performance requirements.

          Accordingly, consideration has been given to using multi-core architectures to provide an energy efficient system for the performance of such diverse tasks. Whilst systems with multiple processor cores have been used for some time to increase performance, by allowing the different cores to operate in parallel on different tasks in  
25 order to increase throughput, analysis as to how such systems could be used to improve energy efficiency has been a relatively recent development.

          The article "Towards Better Performance Per Watt in Virtual Environments on Asymmetric Single-ISA Multi-Core Systems" by V Kumar et al, ACM SIGOPS Operating Systems Review, Volume 43, Issue 3 (July 2009), discusses Asymmetric  
30 Single Instruction Set Architecture (ASISA) multi-core systems, consisting of several cores exposing the same instruction set architecture (ISA) but differing in features, complexity, power consumption, and performance. In the paper, properties of virtualised workloads are studied to shed insight into how these workloads should be

scheduled on ASISA systems in order to improve performance and energy consumption. The paper identifies that certain tasks are more applicable to high frequency/performance micro-architectures (typically compute intensive tasks), while others are more suited to lower frequency/performance micro-architectures and as a side effect will consume less energy (typically input/output intensive tasks). Whilst such studies show how ASISA systems might be used to run diverse tasks in an energy efficient manner, it is still necessary to provide a mechanism for scheduling individual tasks to the most appropriate processors, and such scheduling management will typically place a significant burden on the operating system.

10           The article “Single-ISA Heterogeneous Multi-Core Architectures: The Potential for Processor Power Reduction” by R Kumar et al, Proceedings of the 36<sup>th</sup> International Symposium of Microarchitecture (MICRO-36’03) discusses a multi-core architecture where all cores execute the same instruction set, but have different capabilities and performance levels. At run time, system software evaluates the resource requirements of an application and chooses the core that can best meet these requirements while minimising energy consumption. As discussed in section 2 of that paper, during an application’s execution, the operating system software tries to match the application to the different cores, attempting to meet a defined objective function, for example a particular performance requirement. In section 2.3, it is noted that there is a cost to switching cores, which necessitates restriction of the granularity of switching. A particular example is then discussed where, if the operating system decides a switch is in order, it powers up the new core, triggers a cache flush to save all dirty cache data to a shared memory structure, and then signals the new core to start at a predefined operating system entry point. The old core can then be powered down, whilst the new core retrieves required data from memory. Such an approach is described in section 2.3 as allowing an application to be switched between cores by the operating system. The remainder of the paper then discusses how such switching may be performed dynamically within a multi-core setting with the aim of reducing energy consumption.

30           Whilst the above paper discusses the potential for single-ISA heterogeneous multi-core architectures to provide energy consumption reductions, it still requires the operating system to be provided with sufficient functionality to enable scheduling

5 decisions for individual applications to be made. The role of the operating system in this respect is made more complex when switching between processor instances with different architectural features. In this regard it should be noted that the Alpha cores EV4 to EV8 considered in the paper are not fully ISA compatible, as discussed for example in the fifth paragraph of section 2.2.

Further, the paper does not address the problem that there is a significant overhead involved in switching applications between cores, which can significantly reduce the benefits to be achieved from such switching.

#### SUMMARY OF THE INVENTION

10 Viewed from a first aspect the present invention provides a data processing apparatus comprising: first processing circuitry for performing data processing operations; second processing circuitry for performing data processing operations; the first processing circuitry being architecturally compatible with the second processing circuitry, such that a workload to be performed by the data processing apparatus can be  
15 performed on either the first processing circuitry or the second processing circuitry, said workload comprising at least one application and at least one operating system for running said at least one application; the first processing circuitry being micro-architecturally different from the second processing circuitry, such that performance of the first processing circuitry is different to performance of the second processing  
20 circuitry; the first and second processing circuitry being configured such that the workload is performed by one of the first processing circuitry and the second processing circuitry at any point in time; a switch controller, responsive to a transfer stimulus, to perform a handover operation to transfer performance of the workload from source processing circuitry to destination processing circuitry, the source  
25 processing circuitry being one of the first processing circuitry and the second processing circuitry, and the destination processing circuitry being the other of the first processing circuitry and the second processing circuitry; the switch controller being arranged, during the handover operation, to cause the source processing circuitry to make its current architectural state available to the destination processing circuitry, the  
30 current architectural state being that state not available from shared memory shared between the first and second processing circuitry at a time the handover operation is initiated, and that is necessary for the destination processing circuitry to successfully

take over performance of the workload from the source processing circuitry; the source processing circuitry and second processing circuitry arranged to implement an accelerated mechanism to make the current architectural state available to the destination processing circuitry without routing of the current architectural state via the shared memory.

In accordance with the present invention, a data processing apparatus is provided with first and second processing circuitry, which are architecturally compatible with each other, but micro-architecturally different. Due to the architectural compatibility of the first and second processing circuitry, a workload consisting not just of one or more applications, but also including at least one operating system for running those one or more applications, can be moved between the first and second processing circuitry. Further because the first and second processing circuitry are micro-architecturally different, the performance characteristics (and hence energy consumption characteristics) of the first and second processing circuitry differs.

In accordance with the present invention, at any point in time the workload is performed by one of the first and second processing circuits and a switch controller is responsive to a transfer stimulus to perform a handover operation to transfer performance of the workload between the processing circuits. Upon receipt of a transfer stimulus, whichever of the two processing circuits is currently performing the workload is considered to be the source processing circuitry, and the other is considered to be the destination processing circuitry. The switch controller responsible for performing the handover operation causes the source processing circuitry's current architectural state to be made available to the destination processing circuitry through the use of an accelerated mechanism without routing of the current architectural state via the shared memory. As used herein, the term "shared memory" refers to memory which can be directly accessed by both the first processing circuitry and the second processing circuitry, for example main memory coupled to both the first and second processing circuitry via an interconnect.

Hence, by such an approach, the source processing circuitry makes its current architectural state available to the destination processing circuitry without reference by the destination processing circuitry to the shared memory in order to obtain that

current architectural state. This results not only in a performance improvement during the transfer operation, but also a reduction in energy consumption associated with the transfer operation.

This addresses the problem in existing prior art, namely that irrespective of the manner in which a switch between different processing circuits takes place, there is a need to transfer in a fast and energy efficient manner the information required for that switch to be successful, in particular the earlier-mentioned current architectural state. It would be possible for all of the current architectural state to be written out to shared memory as part of the handover operation, so that it could then be read from shared memory by the destination processing circuitry. However, such a process would not only take a significant amount of time, but would also consume significant energy, which could dramatically offset the potential benefits that could be achieved by performing the switch.

Through use of the present invention, it is possible to ensure that the necessary architectural state that is not available in shared memory at the time the handover operation is initiated is made available to the destination processing circuitry in a quick and energy efficient manner, so that it can successfully takeover performance of the workload. Since the accelerated mechanism is quick and energy efficient, it increases the number of situations in which it is energy efficient to make the switch from one processing circuitry to the other.

For the purposes of the present invention, it is immaterial whether the operating system is involved in the switching process (either by generation of the transfer stimulus, or by forming at least part of the switch controller), or whether instead the switch controller is arranged to make the transfer transparent to the operating system. Whichever approach is taken, the accelerated mechanism of the present invention will give significant performance and energy savings in the transfer of the architectural state to the destination processing circuitry.

In one embodiment, the data processing apparatus further comprises: power control circuitry for independently controlling power provided to the first processing circuitry and the second processing circuitry; wherein prior to occurrence of the transfer stimulus the destination processing circuitry is in a power saving condition, and during the handover operation the power control circuitry causes the destination

processing circuitry to exit the power saving condition prior to the destination processing circuitry taking over performance of the workload. Through use of such power control circuitry, it is possible to reduce the energy consumed by any processing circuitry not currently performing the workload.

5           In one embodiment, following the handover operation, the power control circuitry causes the source processing circuitry to enter the power saving condition. This can occur immediately following the handover operation, or in alternative embodiments the source processing circuitry may be arranged to only enter the power saving condition after some predetermined period of time has elapsed, which can allow  
10 data still retained by the source processing circuitry to be made available to the destination processing circuitry in a more energy efficient and higher performance manner.

          In one embodiment, at least the source circuitry has an associated cache, the data processing apparatus further comprises snoop control circuitry, and the  
15 accelerated mechanism comprises transfer of the current architectural state to the destination processing circuitry through use of the source circuitry's associated cache and the snoop control circuitry.

          In accordance with this technique, the source processing circuitry's local cache is used to store the current architectural state that must be made available to the  
20 destination processing circuitry. That state is then marked as shareable, which allows that state to be snooped by the destination processing circuitry using the snoop control circuitry. Hence, in such embodiments, the first and second processing circuitry are made hardware cache coherent with one another, this reducing the amount of time, energy and hardware complexity involved in switching from the source processing  
25 circuitry to the destination processing circuitry.

          In one particular embodiment, the accelerated mechanism is a save and restore mechanism, which causes the source processing circuitry to store its current architectural state to its associated cache, and causes the destination processing circuitry to perform a restore operation via which the snoop control circuitry retrieves  
30 the current architectural state from the source processing circuitry's associated cache and provides that retrieved current architectural state to the destination processing circuitry. The save/store mechanism provides a particularly efficient technique for

saving the architectural state into the source circuitry's local cache, and for the destination processing circuitry to then retrieve that state.

Such an approach may be used irrespective of whether the destination processing circuitry has its own associated local cache or not. Whenever a request for an item of the architectural state is received by the snoop control circuitry, either  
5 directly from the destination processing circuitry, or from an associated local cache of the destination processing circuitry in the event of a cache miss, then it will determine that the required item of architectural state is stored in the local cache associated with the source circuitry and retrieve that data from the source circuitry's local cache for  
10 return to the destination processing circuitry (either directly or via the destination processing circuitry's associated cache if present).

In one particular embodiment, the destination processing circuitry does have an associated cache in which the transferred architectural state obtained by the snoop control circuitry is stored for reference by the destination processing circuitry.

15 However, the hardware cache coherency approach described above is not the only technique that could be used for providing the earlier-mentioned accelerated mechanism. For example, in an alternative embodiment, the accelerated mechanism comprises a dedicated bus between the source processing circuitry and the destination processing circuitry over which the source processing circuitry provides its current  
20 architectural state to the destination processing circuitry. Whilst such an approach will typically have a higher hardware cost than employing the cache coherency approach, it would provide an even faster way of performing the switching, which could be beneficial in certain implementations.

The transfer stimulus can be generated for a variety of reasons. However, in  
25 one embodiment, timing of the transfer stimulus is chosen so as to improve energy efficiency of the data processing apparatus. This can be achieved in a variety of ways. For example, the performance counters of the processing circuitry can be set up to count performance sensitive events (for example the number of instructions executed, or the number of load-store operations). Coupled with a cycle counter or a system  
30 timer, this allows identification that a highly compute intensive application is executing that may be better served by switching to the higher performance processing circuitry, identification of a large number of load-store operations indicating an IO

intensive application which may be better served on the energy efficient processing circuitry, etc. An alternative approach is for applications to be profiled and marked as 'big', 'little' or 'big/little', whereby the operating system can interface with the switch controller to move the workload accordingly (here the term "big" refers to a higher performance processing circuitry, and the term "little" refers to a more energy efficient processing circuitry).

The architectural state that is required for the destination processing circuitry to successfully take over performance of the workload from the source processing circuitry can take a variety of forms. However, in one embodiment, the architectural state comprises at least the current value of one or more special purpose registers of the source processing circuitry, including a program counter value. In addition to the program counter value, various other information may be stored within the special purpose registers. For example, other special purpose registers include processor status registers (e.g. the CPSR and SPSR in the ARM architecture) that hold control bits for processor mode, interrupt masking, execution state and flags. Other special purpose registers include architectural control (the CP15 system control register in the ARM architecture) that hold bits to alter data endianness, turn the MMU on or off, turn data/instruction caches on or off, etc. Other special purpose registers in CP15 store exception address and status information.

In one embodiment, the architectural state further comprises the current values stored in an architectural register file of the source processing circuitry. As will be understood by those skilled in the art, the architectural register file contains registers that will be referred to by the instructions executed whilst applications are running, those registers holding source operands for computations, and providing locations to which results of those computations are stored.

In one embodiment, at least one of the first processing circuitry and the second processing circuitry comprise a single processing unit. Further, in one embodiment, at least one of the first processing circuitry and the second processing circuitry comprise a cluster of processing units with the same micro-architecture. In one particular embodiment, the first processing circuitry may comprise a cluster of processing units with the same micro-architecture, whilst the second processing circuitry comprises a

single processing unit (with a different micro-architecture to the micro-architecture of the processing units within the cluster forming the first processing circuitry).

The power saving condition that the power control circuitry can selectively place the first and second processing circuits in can take a variety of forms. In one embodiment, the power saving condition is one of: a powered off condition; a partial/full data retention condition; a dormant condition; or an idle condition. Such conditions will be well understood by a person skilled in the art, and accordingly will not be discussed in more detail herein.

There are a number of ways in which the first and second processing circuits can be arranged to be micro-architecturally different. In one embodiment, the first processing circuitry and second processing circuitry are micro-architecturally different by having at least one of: different execution pipeline lengths; or different execution resources. Differences in pipeline length will typically result in differences in operating frequency, which in turn will have an effect on performance. Similarly, differences in execution resources will have an effect on throughput and hence performance. For example, a processing circuit having wider execution resources will enable more information to be processed at any particular point in time, improving throughput. In addition, or alternatively, one processing circuit may have more execution resources than the other, for example, more arithmetic logic units (ALUs), which again will improve throughput. As another example of different execution resources, an energy efficient processing circuit may be provided with a simple in-order pipeline, whilst a higher performance processing circuit may be provided with an out-of-order superscalar pipeline.

A further problem that can arise when using high performance processing circuits, for example running at GHz frequencies, is that such processors are approaching, and sometimes exceeding, the thermal limits that they were designed to operate within. Known techniques for seeking to address these problems can involve the processing circuit being put into a low-power condition to reduce heat output, which may include clock throttling and/or voltage reduction, and potentially even turning the processing circuit off completely for a period of time. However, when adopting the technique of embodiments of the present invention, it is possible to implement an alternative approach to avoid the thermal limits being exceeded. In

particular, in one embodiment, the source processing circuitry is higher performance than the destination processing circuitry, and the data processing apparatus further comprises thermal monitoring circuitry for monitoring a thermal output of the source processing circuitry, and for triggering said transfer stimulus when said thermal output reaches a predetermined level. In accordance with such techniques, the entire workload can be migrated from the higher performance processing circuitry to the lower performance processing circuitry, whereafter less heat will be generated, and the source processing circuitry will be allowed to cool down. Hence, the package containing the two processing circuits can cool while continued program execution can take place, albeit at lower throughput.

The data processing apparatus can be arranged in a variety of ways. However, in one embodiment the first processing circuitry and the second processing circuitry reside within a single integrated circuit.

Viewed from a second aspect, the present invention provides a data processing apparatus comprising: first processing means for performing data processing operations; second processing means for performing data processing operations; the first processing means being architecturally compatible with the second processing means, such that a workload to be performed by the data processing apparatus can be performed on either the first processing means or the second processing means, said workload comprising at least one application and at least one operating system for running said at least one application; the first processing means being micro-architecturally different from the second processing means, such that performance of the first processing means is different to performance of the second processing means; the first and second processing means being configured such that the workload is performed by one of the first processing means and the second processing means at any point in time; a transfer control means, responsive to a transfer stimulus, for performing a handover operation to transfer performance of the workload from source processing means to destination processing means, the source processing means being one of the first processing means and the second processing means, and the destination processing means being the other of the first processing means and the second processing means; the transfer control means, during the handover operation, for causing the source processing means to make its current architectural state available to

the destination processing means, the current architectural state being that state not available from shared memory means shared between the first and second processing means at a time the handover operation is initiated, and that is necessary for the destination processing means to successfully take over performance of the workload from the source processing means; the source processing means and second processing means for implementing an accelerated mechanism to make the current architectural state available to the destination processing means without routing of the current architectural state via the shared memory means.

Viewed from a third aspect the present invention provides a method of operating a data processing apparatus having first processing circuitry for performing data processing operations and second processing circuitry for performing data processing operations, the first processing circuitry being architecturally compatible with the second processing circuitry, such that a workload to be performed by the data processing apparatus can be performed on either the first processing circuitry or the second processing circuitry, said workload comprising at least one application and at least one operating system for running said at least one application, and the first processing circuitry being micro-architecturally different from the second processing circuitry, such that performance of the first processing circuitry is different to performance of the second processing circuitry, the method comprising the steps of: performing, at any point in time, the workload on one of the first processing circuitry and the second processing circuitry; performing, in response to a transfer stimulus, a handover operation to transfer performance of the workload from source processing circuitry to destination processing circuitry, the source processing circuitry being one of the first processing circuitry and the second processing circuitry, and the destination processing circuitry being the other of the first processing circuitry and the second processing circuitry; during the handover operation, causing the source processing circuitry to make its current architectural state available to the destination processing circuitry, the current architectural state being that state not available from shared memory shared between the first and second processing circuitry at a time the handover operation is initiated, and that is necessary for the destination processing circuitry to successfully take over performance of the workload from the source processing circuitry; and said step of making the current architectural state available to

the destination processing circuitry comprising the source processing circuitry and second processing circuitry implementing an accelerated mechanism to make the current architectural state available to the destination processing circuitry without routing of the current architectural state via the shared memory.

## 5 BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be described further, by way of example only, with reference to embodiments thereof as illustrated in the accompanying drawings, in which:

10 Figure 1 is a block diagram of a data processing system in accordance with one embodiment;

15 Figure 2 schematically illustrates the provision of a switch controller (also referred to herein as a workload transfer controller) in accordance with one embodiment to logically separate the workload being performed by the data processing apparatus from the particular hardware platform within the data processing apparatus being used to perform that workload;

Figure 3 is a diagram schematically illustrating the steps performed by both a source processor and a destination processor in response to a switching stimulus in order to transfer the workload from the source processor to the destination processor in accordance with one embodiment;

20 Figure 4A schematically illustrates the storing of the source processing circuitry's current architectural state into its associated cache during the save operation of Figure 3;

25 Figure 4B schematically illustrates the use of the snoop control unit to control the transfer of the source processing circuit's current architectural state to the destination processing circuit during the restore operation of Figure 3;

Figure 5 illustrates an alternative structure for providing an accelerated mechanism for transferring the current architectural state of the source processing circuitry to the destination processing circuitry during the transfer operation in accordance with one embodiment;

30 Figures 6A to 6I schematically illustrate the steps performed to transfer a workload from a source processing circuit to a destination processing circuit in accordance with one embodiment;

Figure 7 is a graph showing energy efficiency variation with performance, and illustrating how the various processor cores illustrated in Figure 1 are used at various points along that curve in accordance with one embodiment;

Figures 8A and 8B schematically illustrate a low performance processor pipeline and a high performance processor pipeline, respectively, as utilised in one embodiment; and

Figure 9 is a graph showing the variation in power consumed by the data processing system as performance of a processing workload is switched between a low performance, high energy efficiency, processing circuit and a high performance, low energy efficiency, processing circuit.

#### DESCRIPTION OF EMBODIMENTS

Figure 1 is a block diagram schematically illustrating a data processing system in accordance with one embodiment. As shown in Figure 1, the system contains two architecturally compatible processing circuit instances (the processing circuitry 0 10 and the processing circuitry 1 50), but with those different processing circuit instances having different micro-architectures. In particular, the processing circuitry 10 is arranged to operate with higher performance than the processing circuitry 50, but with the trade-off that the processing circuitry 10 will be less energy efficient than the processing circuitry 50. Examples of micro-architectural differences will be described in more detail below with reference to Figures 8A and 8B.

Each processing circuit may include a single processing unit (also referred to herein as a processor core), or alternatively at least one of the processing circuit instances may itself comprise a cluster of processing units with the same micro-architecture.

In the example illustrated in Figure 1, the processing circuit 10 includes two processor cores 15, 20 which are both architecturally and micro-architecturally identical. In contrast, the processing circuit 50 contains only a single processor core 55. In the following description, the processor cores 15, 20 will be referred to as “big” cores, whilst the processor core 55 will be referred to as a “little” core, since the processor cores 15, 20 will typically be more complex than the processor core 55 due to those cores being designed with performance in mind, whereas in contrast the

processor core 55 is typically significantly less complex due to being designed with energy efficiency in mind.

In Figure 1, each of the cores 15, 20, 55 is assumed to have its own associated local level 1 cache 25, 30, 60, respectively, which may be arranged as a unified cache for storing both instructions and data for reference by the associated core, or can be arranged with a Harvard architecture, providing distinct level 1 data and level 1 instruction caches. Whilst each of the cores is shown as having its own associated level 1 cache, this is not a requirement, and in alternative embodiments, one or more of the cores may have no local cache.

In the embodiment shown in Figure 1, the processing circuitry 10 also includes a level 2 cache 35 shared between the core 15 and the core 20, with a snoop control unit 40 being used to ensure cache coherency between the two level 1 caches 25, 30 and the level 2 cache 35. In one embodiment, the level 2 cache is arranged as an inclusive cache, and hence any data stored in either of the level 1 caches 25, 30 will also reside in the level 2 cache 35. As will be well understood by those skilled in the art, the purpose of the snoop control unit 40 is to ensure cache coherency between the various caches, so that it can be ensured that either core 15, 20 will always access the most up-to-date version of any data when it issues an access request. Hence, purely by way of example, if the core 15 issues an access request for data that does not reside in the associated level 1 cache 25, then the snoop control unit 40 intercepts the request as propagated on from the level 1 cache 25, and determines with reference to the level 1 cache 30 and/or the level 2 cache 35 whether that access request can be serviced from the contents of one of those other caches. Only if the data is not present in any of the caches is the access request then propagated on via the interconnect 70 to main memory 80, the main memory 80 being memory that is shared between both the processing circuitry 10 and the processing circuitry 50.

The snoop control unit 75 provided within the interconnect 70 operates in a similar manner to the snoop control unit 40, but in this instance seeks to maintain coherency between the cache structure provided within the processing circuitry 10 and the cache structure provided within the processing circuitry 50. In examples where the level 2 cache 35 is an inclusive cache, then the snoop control unit maintains hardware cache coherency between the level 2 cache 35 of the processing circuitry 10 and the

level 1 cache 60 of the processing circuitry 50. However, if the level 2 cache 35 is arranged as an exclusive level 2 cache, then the snoop control unit 75 will also snoop the data held in the level 1 caches 25, 30 in order to ensure cache coherency between the caches of the processing circuitry 10 and the cache 60 of the processing circuitry 50.

In accordance with one embodiment, only one of the processing circuitry 10 and the processing circuitry 50 will be actively processing a workload at any point in time. For the purposes of the present application, the workload can be considered to comprise at least one application and at least one operating system for running that at least one application, such as illustrated schematically by the reference numeral 100 in Figure 2. In this example, two applications 105, 110 are running under control of the operating system 115, and collectively the applications 105, 110 and the operating system 115 form the workload 100. The applications can be considered to exist at a user level, whilst the operating system exists at a privileged level, and collectively the workload formed by the applications and the operating system runs on a hardware platform 125 (representing the hardware level view). At any point in time that hardware platform will either be provided by the processing circuitry 10 or by the processing circuitry 50.

As shown in Figure 1, power control circuitry 65 is provided for selectively and independently providing power to the processing circuitry 10 and the processing circuitry 50. Prior to a transfer of the workload from one processing circuit to the other, only one of the processing circuits will typically be fully powered, i.e. the processing circuit currently performing the workload (the source processing circuitry), and the other processing circuit (the destination processing circuitry) will typically be in a power saving condition. When it is determined that the workload should be transferred from one processing circuit to the other, there will then be a period of time during the transfer operation where both processing circuits are in the powered on state, but at some point following the transfer operation, the source processing circuit from which the workload has been transferred will then be placed into the power saving condition.

The power saving condition can take a variety of forms, dependent on implementation, and hence for example may be one of a powered off condition, a

partial/full data retention condition, a dormant condition or an idle condition. Such conditions will be well understood by a person skilled in the art, and accordingly will not be discussed in more detail herein.

The aim of the described embodiments is to perform switching of the workload  
5 between the processing circuits depending on the required performance/energy level of the workload. Accordingly, when the workload involves the execution of one or more performance intensive tasks, such as execution of games applications, then the workload can be executed on the high performance processing circuit 10, either using one or both of the big cores 15, 20. However, in contrast, when the workload is only  
10 performing low performance intensity tasks, such as MP3 playback, then the entire workload can be transferred to the processing circuit 50, so as benefit from the energy efficiencies that can be realised from utilising the processing circuit 50.

To make best use of such switching capabilities, it is necessary to provide a mechanism that allows the switching to take place in a simple and efficient manner, so  
15 that the action of transferring the workload does not consume energy to an extent that will negate the benefits of switching, and also to ensure that the switching process is quick enough that it does not in itself degrade performance to any significant extent.

In one embodiment, such benefits are at least in part achieved by arranging the processing circuitry 10 to be architecturally compatible with the processing circuitry  
20 50. This ensures that the workload can be migrated from one processing circuitry to the other whilst ensuring correct operation. As a bare minimum, such architectural compatibility requires both processing circuits 10 and 50 to share the same instruction set architecture. However, in one embodiment, such architectural compatibility also entails a higher compatibility requirement so as to ensure that the two processing  
25 circuit instances are seen as identical from a programmer's view. In one embodiment, this involves use of the same architectural registers, and one or more special purpose registers storing data used by the operating system when executing applications. With such a level of architectural compatibility, it is then possible to mask from the operating system  
30 the transfer of the workload between processing circuits, so that the operating system is entirely unaware as to whether the workload is being executed on the processing circuitry 10 or on the processing circuitry 50.

In one embodiment, the handling of the transfer from one processing circuit to the other is managed by the switch controller 120 shown in Figure 2 (also referred to therein as a virtualiser and elsewhere herein as a workload transfer controller). The switch controller can be embodied by a mixture of hardware, firmware and/or software features, but in one embodiment includes software similar in nature to hypervisor software found in virtual machines to enable applications written in one native instruction set to be executed on a hardware platform adopting a different native instruction set. Due to the architectural compatibility between the two processing circuits 10, 50, the switch controller 120 can mask the transfer from the operating system 115 merely by masking one or more items of predetermined processor specific configuration information from the operating system. For example, the processor specific configuration information may include the contents of a CP15 processor ID register and CP15 cache type register.

In such an embodiment, the switch controller then merely needs to ensure that any current architectural state held by the source processing circuit at the time of the transfer, and that is not at the time the transfer is initiated already available from shared memory 80, is made available to the destination processing circuit in order to enable the destination circuit to be in a position to successfully take over performance of the workload. Using the earlier described example, such architectural state will typically comprise the current values stored in the architectural register file of the source processing circuitry, along with the current values of one or more special purpose registers of the source processing circuitry. Due to the architectural compatibility between the processing circuits 10, 50, if this current architectural state can be transferred from the source processing circuit to the destination processing circuit, the destination processing circuit will then be in a position to successfully take over performance of the workload from the source processing circuit.

Whilst architectural compatibility between the processing circuits 10, 50 facilitates transfer of the entire workload between the two processing circuits, in one embodiment the processing circuits 10, 50 are micro-architecturally different from each other, such that there are different performance characteristics, and hence energy consumption characteristics, associated with the two processing circuits. As discussed earlier, in one embodiment, the processing circuit 10 is a high performance, high

energy consumption, processing circuit, while the processing circuit 50 is a lower performance, lower energy consumption, processing circuit. The two processing circuits can be micro-architecturally different from each other in a number of respects, but typically will have at least one of different execution pipeline lengths, and/or  
5 different execution resources. Differences in pipeline length will typically result in differences in operating frequency, which in turn will have an effect on performance. Similarly, differences in execution resources will have an effect on throughput and hence performance. Hence, by way of example, the processing circuitry 10 may have wider execution resources and/or more execution resources, in order to improve  
10 throughput. Further, the pipelines within the processor cores 15, 20 may be arranged to perform out-of-order superscalar processing, whilst the simpler core 55 within the energy efficient processing circuit 50 may be arranged as an in-order pipeline. A further discussion of micro-architectural differences will be provided later with reference to Figures 8A and 8B.

15 The generation of a transfer stimulus to cause the switch controller 120 to instigate a handover operation to transfer the workload from one processing circuit to another can be triggered for a variety of reasons. For example, in one embodiment, applications may be profiled and marked as 'big', 'little' or 'big/little', whereby the operating system can interface with the switch controller to move the workload  
20 accordingly. Hence, by such an approach, the generation of the transfer stimulus can be mapped to particular combinations of applications being executed, to ensure that when high performance is required, the workload is executed on the high performance processing circuit 10, whereas when that performance is not required, the energy efficient processing circuit 50 is instead used. In other embodiments, algorithms could  
25 be executed to dynamically determine when to trigger a transfer of the workload from one processing circuit to the other based on one or more inputs. For example, the performance counters of the processing circuitry can be set up to count performance sensitive events (for example the number of instructions executed, or the number of load-store operations). Coupled with a cycle counter or a system timer, this allows  
30 identification that a highly compute intensive application is executing that may be better served by switching to the higher performance processing circuitry, or identification of a large number of load-store operations indicating an IO intensive

application which may be better served on the energy efficient processing circuitry, etc.

As a yet further example of when a transfer stimulus might be generated, the data processing system may include one or more thermal sensors 90 for monitoring the temperature of the data processing system during operation. It can be the case that modern high performance processing circuits, for example those running at GHz frequencies, sometimes reach, or exceed, the thermal limits that they were designed to operate within. By using such thermal sensors 90, it can be detected when such thermal limits are being reached, and under those conditions a transfer stimulus can be generated to trigger a transfer of the workload to a more energy efficient processing circuit in order to bring about an overall cooling of the data processing system. Hence, considering the example of Figure 1 where the processing circuit 10 is a high performance processing circuit and the processing circuit 50 is a lower performance processing circuit consuming less energy, migration of the workload from the processing circuit 10 to the processing circuit 50 when the thermal limits of the device are being reached will bring about a subsequent cooling of the device, whilst still allowing continued program execution to take place, albeit at lower throughput.

Whilst in Figure 1 two processing circuits 10, 50 are shown, it will be appreciated that the techniques of the above described embodiments can also be applied to systems incorporating more than two different processing circuits, allowing the data processing system to span a larger range of performance/energy levels. In such embodiments, each of the different processing circuits will be arranged to be architecturally compatible with each other to allow the ready migration of the entire workload between the processing circuits, but will also be micro-architecturally different to each other to allow choices to be made between the use of those processing circuits dependent on required performance/energy levels.

Figure 3 is a flow diagram illustrating the sequence of steps performed on both the source processor and the destination processor when the workload is transferred from the source processor to the destination processor upon receipt of a transfer stimulus. Such a transfer stimulus may be generated by the operating system 115 or the virtualiser 120 via a system firmware interface resulting in the detection of the switching stimulus at step 200 by the source processor (which will be running not only

the workload, but also the virtualiser software forming at least part of the switch controller 120). Receipt of the transfer stimulus (also referred to herein as the switching stimulus) at step 200 will cause the power controller 65 to initiate a power on and reset operation 205 on the destination processor. Following such power on and reset, the destination processor will invalidate its local cache at step 210, and then  
5 enable snooping at step 215. At this point, the destination processor will then signal to the source processor that it is ready for the transfer of the workload to take place, this signal causing the source processor to execute a save state operation at step 225. This save state operation will be discussed in more detail later with reference to Figure 4A,  
10 but in one embodiment involves the source processing circuitry storing to its local cache any of its current architectural state which is not available from shared memory at the time the handover operation is initiated, and that is necessary for the destination processor to successfully take over performance of the workload.

Following the save state operation 225, a switch state signal will be issued to  
15 the destination processor 230 indicating to the destination processor that it should now begin snooping the source processor in order to retrieve the required architectural state. This process takes place via a restore state operation 230 which will be discussed in more detail later with reference to Figure 4B, but which in one embodiment involves the destination processing circuitry initiating a sequence of accesses which are  
20 intercepted by the snoop control unit 75 within the interconnect 70, and which cause the cached copy of the architectural state in the source processor's local cache to be retrieved and returned to the destination processor.

Following step 230, the destination processor is then in a position to take over processing of the workload, and accordingly normal operation begins at step 235.

25 In one embodiment, once normal operation begins on the destination processor, the source processor's cache could be cleaned as indicated at step 250, in order to flush any dirty data to the shared memory 80, and then the source processor could be powered down at step 255. However, in one embodiment, to further improve the efficiency of the destination processor, the source processor is arranged to remain  
30 powered up for a period of time referred to in Figure 3 as the snooping period. During this time, at least one of the caches of the source circuit remains powered up, so that its contents can be snooped by the snoop control circuit 75 in response to access requests

issued by the destination processor. Following the transfer of the entire workload using the process described in Figure 3, it is expected that for at least an initial period of time after which the destination processor begins operation of the workload, some of the data required during the performance of the workload will reside in the source processor's cache. If the source processor had flushed its contents to memory, and been powered down, then the destination processor would during these early stages operate relatively inefficiently, since there would be a lot of cache misses in its local cache, and a lot of fetching of data from shared memory, resulting in a significant performance impact whilst the destination processor's cache is "warmed up", i.e. filled with data values required by the destination processor circuit to perform the operations specified by the workload. However, by leaving the source processor's cache powered up during the snooping period, the snoop control circuit 75 will be able to service a lot of these cache miss requests with reference to the source circuit's cache, yielding significant performance benefits when compared with the retrieval of that data from shared memory 80.

However, this performance benefit is only expected to last for a certain amount of time following the switch, after which the contents of the source processor's cache will become stale. Accordingly, at some point a snoop stop event will be generated to disable snooping at step 245, whereafter the source processor's cache will be cleaned at step 250, and then the source processor will be powered down at step 255. A discussion of the various scenarios under which the snoop stop event may be generated will be discussed in more detail later with reference to Figure 6G.

Figure 4A schematically illustrates the save operation performed at step 225 in Figure 3 in accordance with one embodiment. In particular, in one embodiment, the architectural state that needs to be stored from the source processing circuitry 300 to the local cache 330 consists of the contents of a register file 310 referenced by an arithmetic logic unit (ALU) 305 during the performance of data processing operations, along with the contents of various special purpose registers 320 identifying a variety of pieces of information required by the workload to successfully enable that workload to be taken over by the destination processing circuitry. The contents of the special purpose registers 320 will include for example a program counter value identifying a current instruction being executed, along with various other information. For example,

other special purpose registers include processor status registers (e.g. the CPSR and SPSR in the ARM architecture) that hold control bits for processor mode, interrupt masking, execution state and flags. Other special purpose registers include architectural control (the CP15 system control register in the ARM architecture) that hold bits to alter data endianness, turn the MMU on or off, turn data/instruction caches on or off, etc. Other special purpose registers in CP15 store exception address and status information.

As schematically illustrated in Figure 4A, the source processing circuit 300 will also typically hold some processor specific configuration information 315, but this information does not need saving to the cache 330, since it will not be applicable to the destination processing circuitry. The processor specific configuration information 315 is typically hard-coded in the source processing circuit 300 using logic constants, and may include, for example, the contents of the CP15 processor ID register (which will be different for each processing circuit) or the contents of the CP15 cache type register (which will depend on the configuration of the caches 25, 30, 60, for example indicating that the caches have different line lengths). When the operating system 115 requires a piece of processor specific configuration information 315, then unless the processor is already in hypervisor mode, an execution trap to hypervisor mode occurs. In response, the virtualiser 120 may in one embodiment indicate the value of the information requested, but in another embodiment will return a “virtual” value. In the case of the processor ID value, this virtual value can be chosen to be the same for both “big” and “little” processors, thereby causing the actual hardware configuration to be hidden from the operating system 115 by the virtualiser 120.

As illustrated schematically in Figure 4A, during the save operation, the contents of the register file 310 and of the special purpose registers 320 is stored by the source processing circuitry into the cache 330 to form a cached copy 335. This cached copy is then marked as shareable, which allows the destination processor to snoop this state via the snoop control unit 75.

The restore operation subsequently performed on the destination processor is then illustrated schematically in Figure 4B. In particular, the destination processing circuitry 350, which may or may not have its own local cache) will issue a request for a particular item of architectural state, with that request being intercepted by the snoop

control unit 75. The snoop control unit will then issue a snoop request to the source processing circuit's local cache 330 to determine whether that item of architectural state is present in the source's cache. Because of the steps taken during the save operation discussed in Figure 4A, a hit will be detected in the source's cache 330, resulting in that cached architectural state being returned via the snoop control unit 75 to the destination processing circuit 350. This process can be repeated iteratively until all of the items of architectural state have been retrieved via snooping of the source processing circuit's cache. Any processor specific configuration information relevant to destination processing circuit 350 is typically hard-coded in the destination processing circuit 350 as discussed earlier. Thus, once the restore operation has been completed, the destination processing circuitry then has all the information required to enable it to successfully take over handling of the workload.

Further, in one embodiment, regardless of whether the workload 100 is being performed by the "big" processing circuit 10 or "little" processing circuit 50, the virtualiser 120 provides the operating system 115 with virtual configuration information having the same values, and so the hardware differences between the "big" and "little" processing circuits 10, 50 are masked from the operating system 115 by the virtualiser 120. This means that the operating system 115 is unaware that the performance of the workload 100 has been transferred to a different hardware platform.

In accordance with the save and restore operations described with reference to Figures 4A and 4B, the various processor instances 10, 50 are arranged to be hardware cache coherent with one another in order to reduce the amount of time, energy and hardware complexity involved in transferring the architectural state from the source processor to the destination processor. The technique uses the source processor's local cache to store all of the state that must be transferred from the source processor to the destination processor and which is not available from shared memory at the time the transfer operation takes place. Because the state is marked as shareable within the source processor's cache, this allows the hardware cache coherent destination processor to snoop this state during the transfer operation. By using such a technique, it is possible to transfer the state between the processor instances without the need to save that state either to main memory or to a local memory mapped storage element.

This hence yields significant performance and energy consumption benefits, increasing the variety of situations in which it would be appropriate to switch the workload in order to seek to realise energy consumption benefits.

However, whilst the technique of using cache coherence as described above provides one accelerated mechanism for making the current architectural state available to the destination processor without routing of the current architectural state via the shared memory, it is not the only way in which such an accelerated mechanism could be implemented. For example, Figure 5 illustrates an alternative mechanism where a dedicated bus 380 is provided between the source processing circuitry 300 and the destination processing circuitry 350 in order to allow the architectural state to be transferred during the handover operation. Hence, in such embodiments, the save and restore operations 225, 230 of Figure 3 are replaced with an alternative transfer mechanism utilising the dedicated bus 380. Whilst such an approach will typically have a higher hardware cost than employing the cache coherency approach (the cache coherency approach typically making use of hardware already in place within the data processing system), it would provide an even faster way of performing the switching, which could be beneficial in certain implementations.

Figures 6A to 6I schematically illustrate a series of steps that are performed in order to transfer performance of a workload from the source processing circuitry 300 to the destination processing circuitry 350. The source processing circuitry 300 is whichever of the processing circuits 10, 50 is performing the workload before the transfer, with the destination processing circuitry being the other of the processing circuits 10, 50.

Figure 6A shows the system in an initial state in which the source processing circuitry 300 is powered by the power controller 65 and is performing the processing workload 100, while the destination processing circuitry 350 is in the power saving condition. In this embodiment, the power saving condition is a power off condition, but as mentioned above other types of power saving condition may also be used. The workload 100, including applications 105, 110 and an operating system 115 for running the applications 105, 110, is abstracted from the hardware platform of the source processing circuitry 300 by the virtualiser 120. While performing the workload 100, the source processing circuitry 300 maintains architectural state 400, which may

comprise for example the contents of the register file 310 and special purpose registers 320 as shown in Figure 4A.

In Figure 6B, a transfer stimulus 430 is detected by the virtualiser 120. While the transfer stimulus 430 is shown in Figure 6B as an external event (e.g. detection of thermal runaway by the thermal sensor 90), the transfer stimulus 430 could also be an event triggered by the virtualiser 120 itself or by the operating system 115 (e.g. the operating system 115 could be configured to inform the virtualiser 120 when a particular type of application is to be processed). The virtualiser 120 responds to the transfer stimulus 430 by controlling the power controller 65 to supply power to the destination processing circuitry 350, in order to place the destination processing circuitry 350 in a powered state.

In Figure 6C, the destination processing circuitry 350 starts executing the virtualiser 120. The virtualiser 120 controls the destination processing circuitry 350 to invalidate its cache 420, in order to prevent processing errors caused by erroneous data values which may be present in the cache 420 on powering up the destination processing circuitry 350. While the destination cache 420 is being invalidated, the source processing circuitry 350 continues to perform the workload 100. When invalidation of the destination cache 420 is complete, the virtualiser 120 controls the destination processing circuitry 350 to signal to the source processing circuitry 300 that it is ready for the handover of the workload 100. By continuing processing of the workload 100 on the source processing circuitry 300 until the destination processing circuitry 350 is ready for the handover operation, the performance impact of the handover can be reduced.

At the next stage, shown in Figure 6D, the source processing circuitry 300 stops performing the workload 100. During this stage, neither the source processing circuitry 300 nor the destination processing circuitry 350 performs the workload 100. A copy of the architectural state 400 is transferred from the source processing circuitry 300 to the destination processing circuitry 350. For example, the architectural state 400 can be saved to the source cache 410 and restored to the destination processing circuitry 350 as shown in Figures 4A and 4B, or can be transferred over a dedicated bus as shown in Figure 5. The architectural state 400 contains all the state information

required for the destination processing circuitry 350 to perform the workload 100, other than the information already present in the shared memory 80.

Having transferred the architectural state 400 to the destination processing circuitry 350, the source processing circuitry 300 is placed in the power saving state by the power control circuitry 65 (see Figure 6E), with the exception that the source cache 410 remains powered. Meanwhile, the destination processing circuitry 350 begins performing the workload 100 using the transferred architectural state 400.

When the destination processing circuitry 350 begins processing the workload 100, the snooping period begins (see Figure 6F). During the snooping period, the snoop control unit 75 can snoop the data stored in the source cache 410 and retrieve the data on behalf of the destination processing circuitry 350. When the destination processing circuitry 350 requests data that is not present in the destination cache 420, the destination processing circuitry 350 requests data from the snoop control unit 75. The snoop control unit 75 then snoops the source cache 410, and if the snoop results in a cache hit then the snoop control unit 75 retrieves the snooped data from the source cache 410 and returns it to the destination processing circuitry 350 where the snooped data can be stored in the destination cache 420. On the other hand, if the snoop results in a cache miss in the source cache 410 then the requested data is fetched from the shared memory 80 and returned to the destination processing circuitry 350. Since accesses to data in the source cache 410 are faster and require less energy than accesses to shared memory 80, snooping the source cache 410 for a period improves processing performance and reduces energy consumption during an initial period following the handover of the workload 100 to the destination processing circuitry 350.

At the step shown in Figure 6G, the snoop control unit 75 detects a snoop stop event which indicates that it is no longer efficient to maintain the source cache 410 in the powered state. The snoop stop event triggers the end of the snooping period. The snoop stop event may be any one of a set of snoop stop events monitored by the snoop control circuitry 75. For example, the set of snoop stop events can include any one or more of the following events:

- a) when the percentage or fraction of snoop hits that result in a cache hit in the source cache 410 (i.e. a quantity proportional to number of snoop hits/number

- of total snoops) drops below a predetermined threshold level after the destination processing circuitry 350 has started performing the workload 100;
- b) when the number of transactions, or the number of transactions of a predetermined type (e.g. cacheable transactions), performed since the destination processing circuitry 350 began performing the workload 100 exceeds a predetermined threshold;
  - c) when the number of processing cycles elapsed since the destination processing circuitry 350 began performing the workload 100 exceeds a predetermined threshold;
  - d) when a particular region of the shared memory 80 is accessed for the first time since the destination processing circuitry 350 began performing the workload 100;
  - e) when a particular region of the shared memory 80, which was accessed for an initial period after the destination processing circuitry 350 began performing the workload 100, is not accessed for a predetermined number of cycles or a predetermined period of time;
  - f) when the destination processing circuitry 350 writes to a predetermined memory location for the first time since starting to perform the transferred workload 100.
- These snoop stop events can be detected using programmable counters in the coherent interconnect 70 that includes the snoop control unit 75. Other types of snoop stop event may also be included in the set of snoop stop events.

On detecting a snoop stop event, the snoop control unit 75 sends a snoop stop signal 440 to the source processor 300. The snoop control unit 75 stops snooping the source cache 410 and from now on responds to data access requests from the destination processing circuitry 350 by fetching the requested data from shared memory 80 and returning the fetched data to the destination processing circuitry 350, where the fetched data can be cached.

In Figure 6H, the source cache's control circuit is responsive to the snoop stop signal 440 to clean the cache 410 in order to save to the shared memory 80 any valid and dirty data values (i.e. whose cached value is more up-to-date than the corresponding value in shared memory 80).

In Figure 6I, the source cache 410 is then powered down by the power controller 65 so that the source processing circuitry 300 is entirely in the power saving state. The destination processing circuitry 350 continues to perform the workload 100. From the point of view of the operating system 115, the situation is now the same as in  
5 Figure 6A. The operating system 115 is not aware that execution of the workload has transferred from one processing circuit to another processing circuit. When another transfer stimulus occurs, then the same steps of Figures 6A to 6I can be used to switch performance of the workload back to the first processor (in this case which of the processing circuits 10, 50 are the “source processing circuitry” and “destination  
10 processing circuitry” will be reversed).

In the embodiment of Figures 6A to 6I, independent power control to the cache 410 and the source processing circuitry 300 is available so that the source processing circuitry 300, other than the source cache 410, can be powered down once the destination processing circuitry 350 has started performing the workload (see Figure  
15 6E), while only the cache 410 of the source processing circuitry 350 remains in the powered state (see Figures 6F to 6H). The source cache 410 is then powered down in Figure 6I. This approach can be useful to save energy, especially when the source processing circuitry 300 is the “big” processing circuit 10.

However, it is also possible to continue to power the entire source processing  
20 circuitry 300 during the snooping period, and to then place the source processing circuitry 300 as a whole in the power saving state at Figure 6I, following the end of the snooping period and the cleaning of the source cache 410. This may be more useful in the case where the source cache 410 is too deeply embedded with the source processor core to be able to be powered independently from the source processor core. This  
25 approach can also be more practical when the source processor is the “little” processing circuit 50, whose power consumption is insignificant in comparison to the “big” processing circuit 10, since once the “big” processing circuit 10 has started processing the transferred workload 100 then switching the “little” processing circuit 50, other than the cache 60, to the power saving state during the snooping period may  
30 have little effect on the overall power consumption of the system. This may mean that the extra hardware complexity of providing individual power control to the “little” processing circuit 50 and the “little” core’s cache 60 may not be justified.

In some situations, it may be known before the workload transfer that the data stored in the source cache 410 will not be needed by the destination processing circuitry 350 when it begins to perform the workload 100. For example, the source processing circuitry 300 may just have completed an application when the transfer occurs, and therefore the data in the source cache 410 at the time of the transfer relates to the completed application and not the application to be performed by the destination processing circuitry 350 after the transfer. In such a case, a snoop override controller can trigger the virtualiser 120 and snoop control circuitry 75 to override the snooping of the source cache 410 and to control the source processing circuit 300 to clean and power down the source cache 410 without waiting for a snoop stop event to signal the end of the snooping period. In this case, the technique of Figures 6A to 6I would jump from the step of Figure 6E straight to the step of Figure 6G, without the step of Figure 6F in which data is snooped from the source cache 410. Thus, if it is known in advance that the data in the source cache 410 will not be useful for the destination processing circuitry 350, power can be saved by placing the source cache 410 and source processing circuitry 300 in the power saving condition without waiting for a snoop stop event. The snoop override controller can be part of the virtualiser 120, or can be implemented as firmware executing on the source processing circuitry 300. The snoop override controller could also be implemented as a combination of elements, for example the operating system 115 could inform the virtualiser 120 when an application has finished, and the virtualiser 120 could then override snooping of the source cache 410 if a transfer occurs when an application has finished.

Figure 7 is a graph on which the line 600 illustrates how energy consumption varies with performance. For various portions of this graph, the data processing system can be arranged to utilise different combinations of the processor cores 15, 20, 55 illustrated in Figure 1 in order to seek to obtain the appropriate trade-off between performance and energy consumption. Hence, by way of example, when a number of very high performance tasks need to be executed, it is possible to run both of the big cores 15, 20 of the processing circuit 10 in order to achieve the desired performance. Optionally supply voltage variation techniques can be used to allow some variation in performance and energy consumption when utilising these two cores.

When the performance requirements drop to a level where the required performance can be achieved using only one of the big cores, then the tasks can be migrated on to just one of the big cores 15, 20, with the other core being powered down or put into some other power saving condition. Again supply voltage variation can be used to allow some variation between performance and energy consumption when using such a single big core. It should be noted that the transition from two big cores to one big core will not require a generation of a transfer stimulus, nor the use of the above described techniques for transferring workload, since in all instances it is the processing circuit 10 that is being utilised, and the processing circuit 50 will be in a power saving condition. However, as indicated by the dotted line 610 in Figure 7, when the performance drops to a level where the small core is able to achieve the required performance, then a transfer stimulus can be generated to trigger the earlier described mechanism for transferring the entire workload from the processing circuit 10 to the processing circuit 50, such that the entire workload is then run on the small core 55, with the processing circuit 10 being placed into a power saving condition. Again, supply voltage variation can be used to allow some variation in the performance and energy consumption of the small core 55.

Figures 8A and 8B respectively illustrate micro-architectural differences between a low performance processor pipeline 800 and a high performance processor pipeline 850 according to one embodiment. The low performance processor pipeline 800 of Figure 8A would be suitable for the little processing core 55 of Figure 1, whereas the high performance processor pipeline 850 of Figure 8B would be suitable for the big cores 15, 20.

The low performance processor pipeline 800 of Figure 8A comprises a fetch stage 810 for fetching instructions from memory 80, a decode stage 820 for decoding the fetched instructions, an issue stage 830 for issuing instructions for execution, and multiple execution pipelines including an integer pipeline 840 for performing integer operations, a MAC pipeline 842 for performing multiply accumulate operations, and a SIMD/FPU pipeline 844 for performing SIMD (single instruction, multiple data) operations or floating point operations. In the low performance processor pipeline 800, the issue stage 830 issues a single instruction at a time, and issues the instructions in the order in which the instructions are fetched.

The high performance processor pipeline 850 of Figure 8B comprises a fetch stage 860 for fetching instructions from memory 80, a decode stage 870 for decoding the fetched instructions, a rename stage 875 for renaming registers specified in the decoded instructions, a dispatch stage 880 for dispatching instructions for execution, and multiple execution pipelines including two integer pipelines 890, 892, a MAC pipeline 894, and two SIMD/FPU pipelines 896, 898. In the high performance processor pipeline 850, the dispatch stage 880 is a parallel issue stage which can issue multiple instructions to different ones of the pipelines 890, 892, 894, 896, 898 at once. The dispatch stage 880 can also issue the instructions out-of-order. Unlike in the low performance processor pipeline 800, the SIMD/FPU pipelines 896, 898 are variable length, which means that operations proceeding through the SIMD/FPU pipelines 896, 898 can be controlled to skip certain stages. An advantage of such an approach is that if multiple execution pipelines each have different resources, there is no need to artificially lengthen the shortest pipeline to make it the same length as the longest pipeline, but instead logic is required to deal with the out-of-order nature of the results produced by the different pipelines (for example to place everything back in order if a processing exception occurs).

The rename stage 875 is provided to map register specifiers, which are included in program instructions and identify particular architectural registers when viewed from a programmer's model point of view, to physical registers which are the actual registers of the hardware platform. The rename stage 875 enables a larger pool of physical registers to be provided by the microprocessor than are present within the programmer's model view of the microprocessor. This larger pool of physical registers is useful during out-of-order execution because it enables hazards such as write-after-write (WAW) hazards to be avoided by mapping the same architectural register specified in two or more different instructions to two or more different physical registers, so that the different instructions can be executed concurrently. For more details of register renaming techniques, the reader is referred to commonly owned US patent application US 2008/114966 and US patent 7,590,826.

The low-performance pipeline 800 and high-performance pipeline 850 are micro-architecturally different in a number of ways. The micro-architectural differences can include:

- a) the pipelines having different stages. For example, the high-performance pipeline 850 has a rename stage 875 which is not present in the low-performance pipeline 800.
- 5 b) the pipeline stages having different capabilities. For example, the issue stage 830 of the low-performance pipeline 800 is capable only of single issue of instructions, whereas the dispatch stage 880 of the high performance pipeline 850 can issue instructions in parallel. Parallel issue improves the processing throughput of the pipeline and so improves performance.
- 10 c) the pipeline stages having different lengths. For example, the decode stage 870 of the high-performance pipeline 850 may include three sub-stages whereas the decode stage 820 of the low-performance pipeline 800 may include only a single sub-stage. The longer a pipeline stage (the greater the number of sub-stages), the greater the number of instructions which can be in flight at the same time, and so greater the operating frequency at which the pipeline can
- 15 operate, which results in a higher level of performance.
- d) a different number of execution pipelines (e.g. the high-performance pipeline 850 has more execution pipelines than the low-performance pipeline 800). By providing more execution pipelines, more instructions can be processed in parallel and so performance is increased.
- 20 e) providing in-order execution (as in pipeline 800) or out-of-order execution (as in pipeline 850). When instructions can be executed out-of-order, then performance is improved since the execution of instructions can be dynamically scheduled to optimize performance. For example, in the low-performance in-order pipeline 800 a series of MAC instructions would need to
- 25 be executed one by one by the MAC pipeline 842 before a later instruction could be executed by one of the integer pipeline 840 and SIMD/floating point pipeline 844. In contrast, in the high-performance pipeline 850 then the MAC instructions could be executed by the MAC pipe 894, while (subject to any data hazards which cannot be resolved by renaming) a later instruction using a
- 30 different execution pipeline 890, 892, 896, 898 can be executed in parallel with the MAC instructions. This means that out-of-order execution can improve processing performance.

These, and other examples of, micro-architectural differences result in the pipeline 850 providing higher performance processing than the pipeline 800. On the other hand, the micro-architectural differences also make the pipeline 850 consume more energy than the pipeline 800. Thus, providing micro-architecturally different pipelines 800, 850 enables the processing of the workload to be optimised for either high performance (by using a “big” processing circuit 10 having the high-performance pipeline 850) or energy efficiency (by using a “little” processing circuit 50 having the low-performance pipeline 800).

Figure 9 shows a graph illustrating the variation in power consumption of the data processing system as performance of the workload 100 is switched between the big processing circuit 10 and the little processing circuit 50.

At point A of Figure 9, the workload 100 is being performed on the little processing circuitry 50 and so power consumption is low. At point B, a transfer stimulus occurs indicating that high-intensity processing is to be performed and so the performance of the workload is handed over to the big processing circuitry 10. The power consumption then rises and remains high at point C while the big processing circuitry 10 is performing the workload. At point D it is assumed that both big cores are operating in combination to process the workload. If however the performance requirements drop to a level where the workload can be handled by only one of the big cores, then the workload is migrated to only one of the big cores, and the other is powered down, as indicated by the drop in power to the level adjacent point E. However, at point E, another transfer stimulus occurs (indicating that a return to low-intensity processing is desired) to trigger a transfer of the performance of the workload back to the little processing circuitry 50.

When the little processing circuitry 50 starts processing the processing workload, most of the big processing circuitry is in the power saving state, but the cache of the big processing circuitry 10 remains powered during the snooping period (point F in Figure 9) to enable the data in the cache to be retrieved for the little processing circuitry 50. Hence, the cache of the big processing circuitry 10 causes the power consumption at point F to be higher than at point A when only the little processing circuitry 50 was powered. At the end of the snooping period, the cache of

the big processing circuitry 10 is powered down and at point G power consumption returns to the low level when only the little processing circuitry 50 is active.

As mentioned above, in Figure 9 the power consumption is higher during the snooping period at point F than at point G due to the cache of the big processing circuitry 10 being powered during the snooping period. Although this increase in power consumption is indicated only following the big-to-little transition, following the little-to-big transition there may also be a snooping period, during which the data in the cache of the little processing circuitry 50 can be snooped on behalf of the big processing circuitry 10 by the snoop control unit 75. The snooping period for the little-to-big transition has not been indicated in Figure 9 because the power consumed by leaving the cache of the little processing circuitry 50 in a powered state during the snooping period is insignificant in comparison with the power consumed by the big processing circuitry 10 when performing the processing workload, and so the very small increase in power consumption due to the cache of the little processing circuitry 50 being powered is not visible in the graph of Figure 9.

The above described embodiments describe a system containing two or more architecturally compatible processor instances with micro-architectures optimised for energy efficiency or performance. The architectural state required by the operating system and applications can be switched between the processor instances depending on the required performance/energy level, in order to allow the entire workload to be switched between the processor instances. In one embodiment, only one of the processor instances is running the workload at any given time, with the other processing instance being in a power saving condition, or in the process of entering/exiting the power saving condition.

In one embodiment, the processor instances may be arranged to be hardware cache coherent with one another to reduce the amount of time, energy and hardware complexity involved in switching the architectural state from the source processor to the destination processor. This reduces the time to perform the switching operation, which increases the opportunities in which the techniques of embodiments can be used.

Such systems may be used in a variety of situations where energy efficiency is important for either battery life and/or thermal management, and the spread of

performance is such that a more energy efficient processor can be used for lower processing workloads while a higher performance processor can be used for higher processing workloads.

Because the two or more processing instances are architecturally compatible,  
5 from an application perspective the only difference between the two processors is the performance available. Through techniques of one embodiment, all architectural state required can be moved between the processors without needing to involve the operating system, such that it is then transparent to the operating system and the applications running on the operating system as to which processor that operating  
10 system and applications are running on.

When using architecturally compatible processor instances as described in the above embodiments, the total amount of architectural state that needs to be transferred can easily fit within a data cache, and since modern processing systems often implement cache coherence, then by storing the architectural state to be switched  
15 inside the data cache, the destination processor can rapidly snoop this state in an energy efficient way making use of existing circuit structures.

In one described embodiment, the switching mechanism is used to ensure thermal limits for the data processing system are not breached. In particular, when the thermal limits are about to be reached, the entire workload can be switched to a more  
20 energy efficient processor instance, allowing the overall system to cool while continued program execution takes place, albeit at a lower throughput.

Although a particular embodiment has been described herein, it will be appreciated that the invention is not limited thereto and that many modifications and additions thereto may be made within the scope of the invention. For example, various  
25 combinations of the features of the following dependent claims could be made with the features of the independent claims without departing from the scope of the present invention.

## CLAIMS

1. A data processing apparatus comprising:
  - 5 first processing circuitry for performing data processing operations;  
second processing circuitry for performing data processing operations;  
the first processing circuitry being architecturally compatible with the second processing circuitry, such that a workload to be performed by the data processing apparatus can be performed on either the first processing circuitry or the second processing circuitry, said workload comprising at least one application and at least one  
10 operating system for running said at least one application;  
the first processing circuitry being micro-architecturally different from the second processing circuitry, such that performance of the first processing circuitry is different to performance of the second processing circuitry;
  - 15 the first and second processing circuitry being configured such that the workload is performed by one of the first processing circuitry and the second processing circuitry at any point in time;  
a switch controller, responsive to a transfer stimulus, to perform a handover operation to transfer performance of the workload from source processing circuitry to  
20 destination processing circuitry, the source processing circuitry being one of the first processing circuitry and the second processing circuitry, and the destination processing circuitry being the other of the first processing circuitry and the second processing circuitry;
  - the switch controller being arranged, during the handover operation, to cause  
25 the source processing circuitry to make its current architectural state available to the destination processing circuitry, the current architectural state being that state not available from shared memory shared between the first and second processing circuitry at a time the handover operation is initiated, and that is necessary for the destination processing circuitry to successfully take over performance of the workload from the  
30 source processing circuitry;  
the source processing circuitry and second processing circuitry arranged to implement an accelerated mechanism to make the current architectural state available

to the destination processing circuitry without routing of the current architectural state via the shared memory.

2. A data processing apparatus as claimed in Claim 1, further comprising:  
5 power control circuitry for independently controlling power provided to the first processing circuitry and the second processing circuitry;  
wherein prior to occurrence of the transfer stimulus the destination processing circuitry is in a power saving condition, and during the handover operation the power control circuitry causes the destination processing circuitry to exit the power saving  
10 condition prior to the destination processing circuitry taking over performance of the workload.
3. A data processing apparatus as claimed in Claim 2, wherein following the handover operation the power control circuitry causes the source processing circuitry  
15 to enter the power saving condition.
4. A data processing apparatus as claimed in any preceding claim, wherein:  
at least the source circuitry has an associated cache;  
the data processing apparatus further comprises snoop control circuitry; and  
20 the accelerated mechanism comprises transfer of the current architectural state to the destination processing circuitry through use of the source circuitry's associated cache and the snoop control circuitry.
5. A data processing apparatus as claimed in Claim 4, wherein the accelerated  
25 mechanism is a save and restore mechanism, which causes the source processing circuitry to store its current architectural state to its associated cache, and causes the destination processing circuitry to perform a restore operation via which the snoop control circuitry retrieves the current architectural state from the source processing circuitry's associated cache and provides that retrieved current architectural state to the  
30 destination processing circuitry.

6. A data processing apparatus as claimed in Claim 4 or Claim 5, wherein the destination processing circuitry has an associated cache in which the transferred architectural state obtained by the snoop control circuitry is stored for reference by the destination processing circuitry.

5

7. A data processing apparatus as claimed in any preceding claim, wherein the accelerated mechanism comprises a dedicated bus between the source processing circuitry and the destination processing circuitry over which the source processing circuitry provides its current architectural state to the destination processing circuitry.

10

8. A data processing apparatus as claimed in any preceding claim, wherein timing of the transfer stimulus is chosen so as to improve energy efficiency of the data processing apparatus.

15

9. A data processing apparatus as claimed in any preceding claim, wherein said architectural state comprises at least the current value of one or more special purpose registers of the source processing circuitry, including a program counter value.

20

10. A data processing apparatus as claimed in Claim 9, wherein said architectural state further comprises the current values stored in an architectural register file of the source processing circuitry.

25

11. A data processing apparatus as claimed in any preceding claim, wherein at least one of the first processing circuitry and the second processing circuitry comprise a single processing unit.

30

12. A data processing apparatus as claimed in any preceding claim, wherein at least one of the first processing circuitry and the second processing circuitry comprise a cluster of processing units with the same microarchitecture.

13. A data processing apparatus as claimed in any preceding claim when dependent on Claim 2, wherein said power saving condition is one of:

a powered off condition;  
a partial/full data retention condition;  
a dormant condition; or  
an idle condition.

5

14. A data processing apparatus as claimed in any preceding claim wherein the first processing circuitry and second processing circuitry are micro-architecturally different by having at least one of:

different execution pipeline lengths; or

10

different execution resources.

15. A data processing apparatus as claimed in any preceding claim, wherein the source processing circuitry is higher performance than the destination processing circuitry, and the data processing apparatus further comprises:

15

thermal monitoring circuitry for monitoring a thermal output of the source processing circuitry, and for triggering said transfer stimulus when said thermal output reaches a predetermined level.

16. A data processing apparatus as claimed in any preceding claim, wherein the first processing circuitry and the second processing circuitry reside within a single integrated circuit.

20

17. A data processing apparatus comprising:

first processing means for performing data processing operations;

25

second processing means for performing data processing operations;

the first processing means being architecturally compatible with the second processing means, such that a workload to be performed by the data processing apparatus can be performed on either the first processing means or the second processing means, said workload comprising at least one application and at least one operating system for running said at least one application;

30

the first processing means being micro-architecturally different from the second processing means, such that performance of the first processing means is different to performance of the second processing means;

5 the first and second processing circuitry means being configured such that the workload is performed by one of the first processing means and the second processing means at any point in time;

a transfer control means, responsive to a transfer stimulus, for performing a handover operation to transfer performance of the workload from source processing means to destination processing means, the source processing means being one of the  
10 first processing means and the second processing means, and the destination processing means being the other of the first processing means and the second processing means;

the transfer control means, during the handover operation, for causing the source processing means to make its current architectural state available to the  
15 destination processing means, the current architectural state being that state not available from shared memory means shared between the first and second processing means at a time the handover operation is initiated, and that is necessary for the destination processing means to successfully take over performance of the workload from the source processing means;

20 the source processing means and second processing means for implementing an accelerated mechanism to make the current architectural state available to the destination processing means without routing of the current architectural state via the shared memory means.

25 18. A method of operating a data processing apparatus having first processing circuitry for performing data processing operations and second processing circuitry for performing data processing operations, the first processing circuitry being architecturally compatible with the second processing circuitry, such that a workload to be performed by the data processing apparatus can be performed on either the first  
30 processing circuitry or the second processing circuitry, said workload comprising at least one application and at least one operating system for running said at least one application, and the first processing circuitry being micro-architecturally different

from the second processing circuitry, such that performance of the first processing circuitry is different to performance of the second processing circuitry, the method comprising the steps of:

5 performing, at any point in time, the workload on one of the first processing circuitry and the second processing circuitry;

performing, in response to a transfer stimulus, a handover operation to transfer performance of the workload from source processing circuitry to destination processing circuitry, the source processing circuitry being one of the first processing circuitry and the second processing circuitry, and the destination processing circuitry  
10 being the other of the first processing circuitry and the second processing circuitry;

during the handover operation, causing the source processing circuitry to make its current architectural state available to the destination processing circuitry, the current architectural state being that state not available from shared memory shared between the first and second processing circuitry at a time the handover operation is  
15 initiated, and that is necessary for the destination processing circuitry to successfully take over performance of the workload from the source processing circuitry; and

said step of making the current architectural state available to the destination processing circuitry comprising the source processing circuitry and second processing circuitry implementing an accelerated mechanism to make the current architectural  
20 state available to the destination processing circuitry without routing of the current architectural state via the shared memory.




FIG. 1




FIG. 2

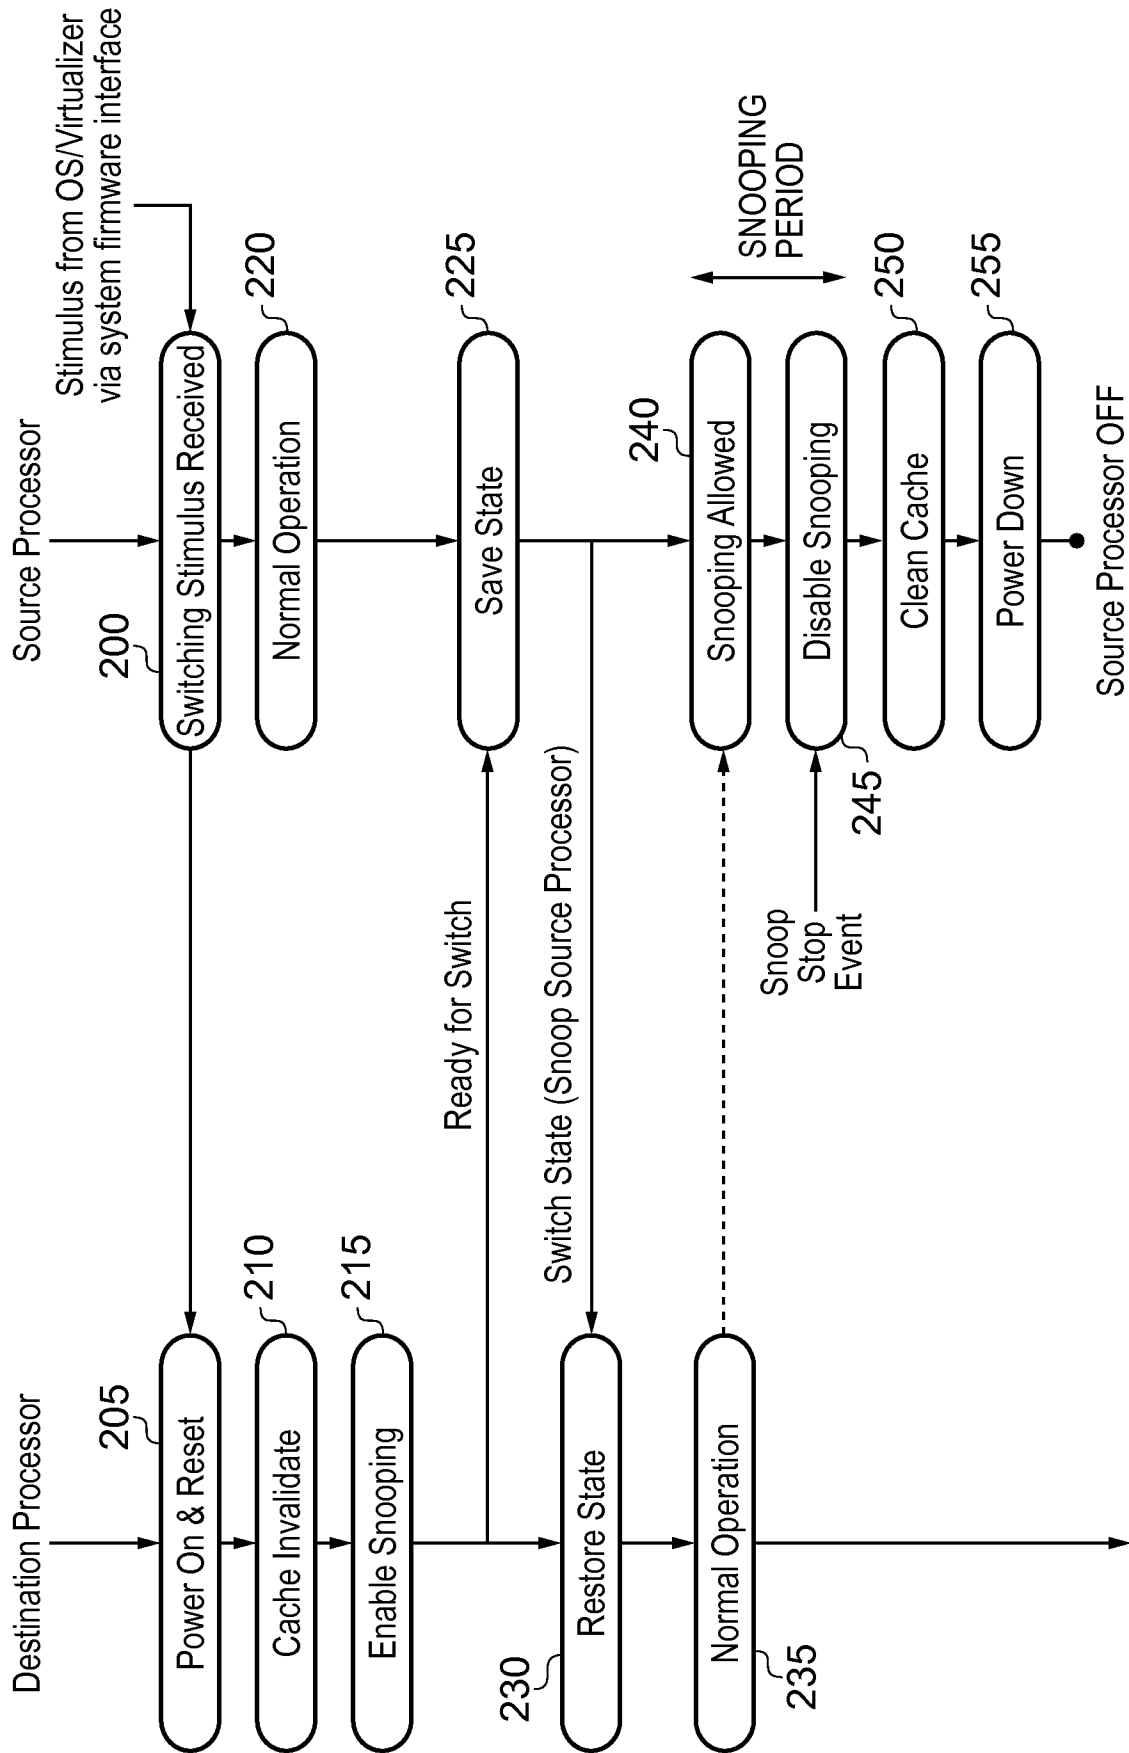


FIG. 3

4/14

SAVE OPERATION (DURING HANDOVER)

SOURCE PROCESSING CIRCUITRY




FIG. 4A

RESTORE OPERATION (DURING HANDOVER)




FIG. 4B

6/14




FIG. 5

7/14




FIG. 6A




FIG. 6B

8/14

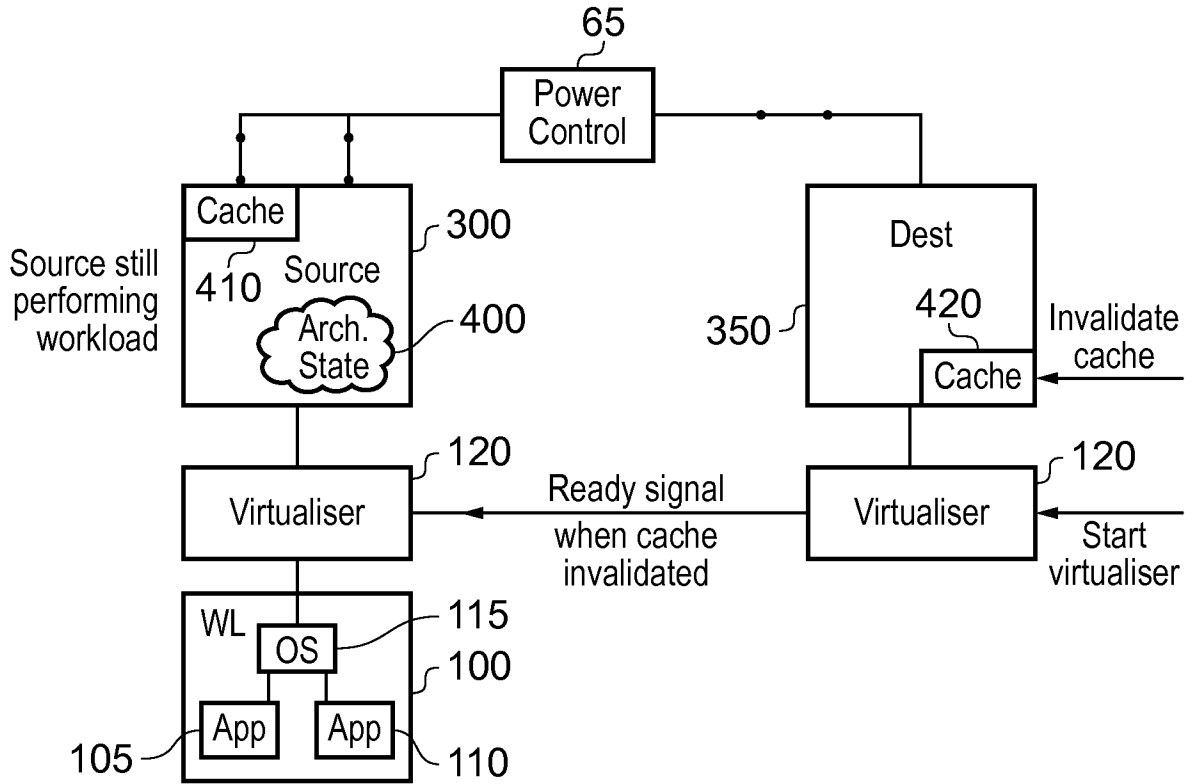


FIG. 6C




FIG. 6D

9/14

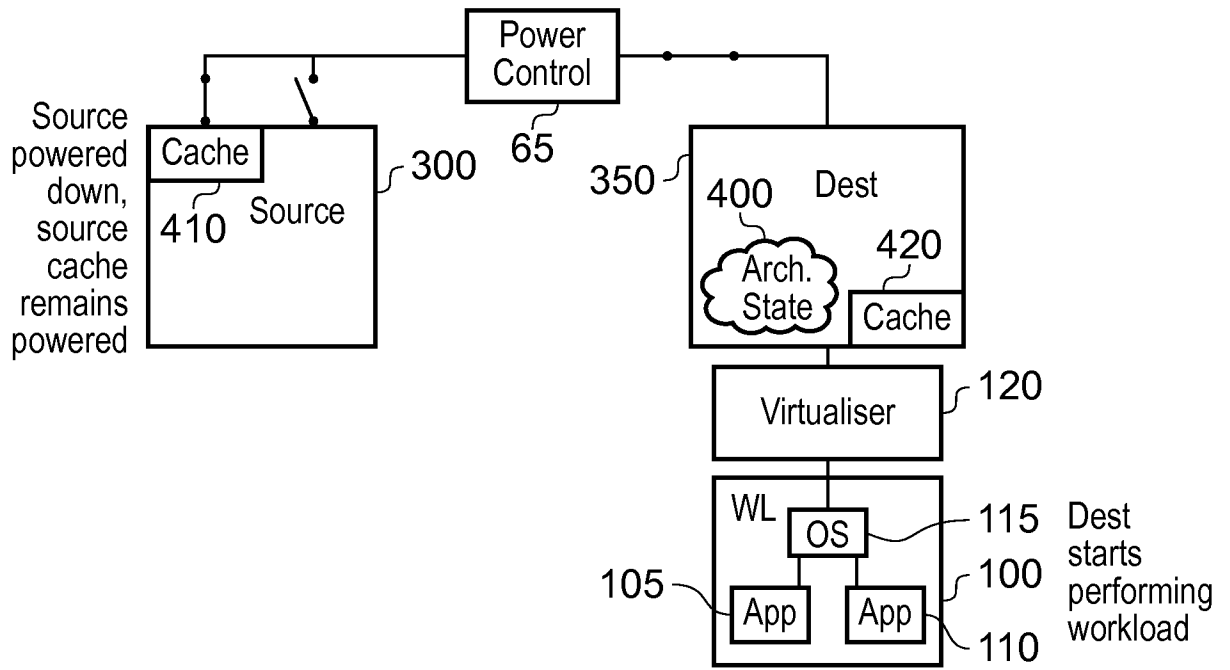


FIG. 6E

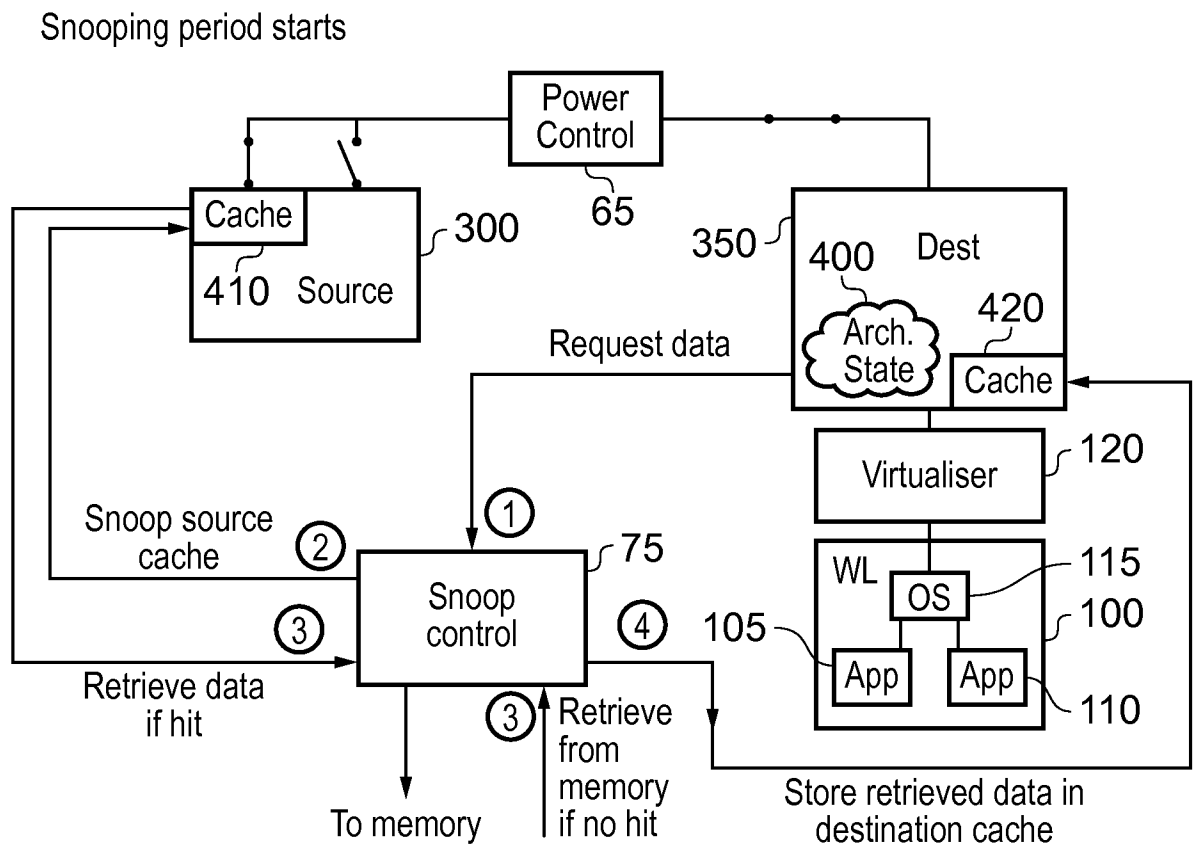


FIG. 6F

10/14

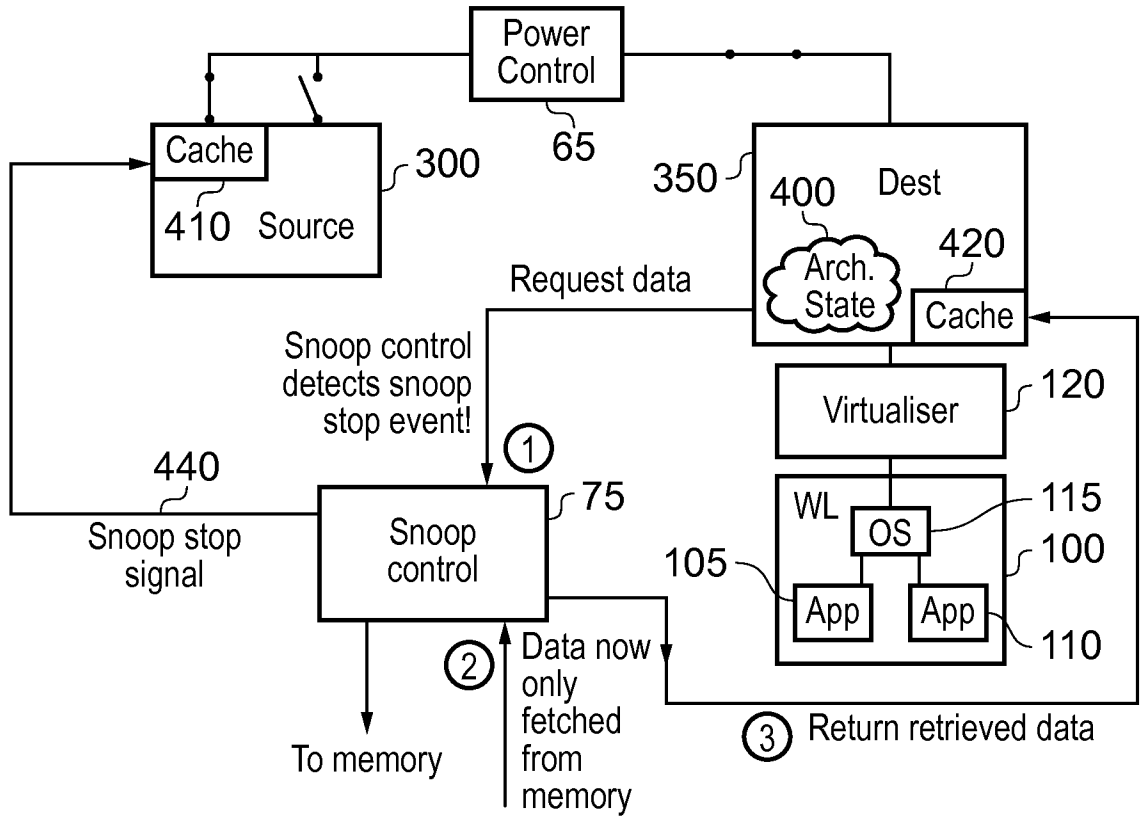


FIG. 6G

end of snooping period




FIG. 6H

11/14




FIG. 6I




FIG. 7




FIG. 8A




FIG. 8B

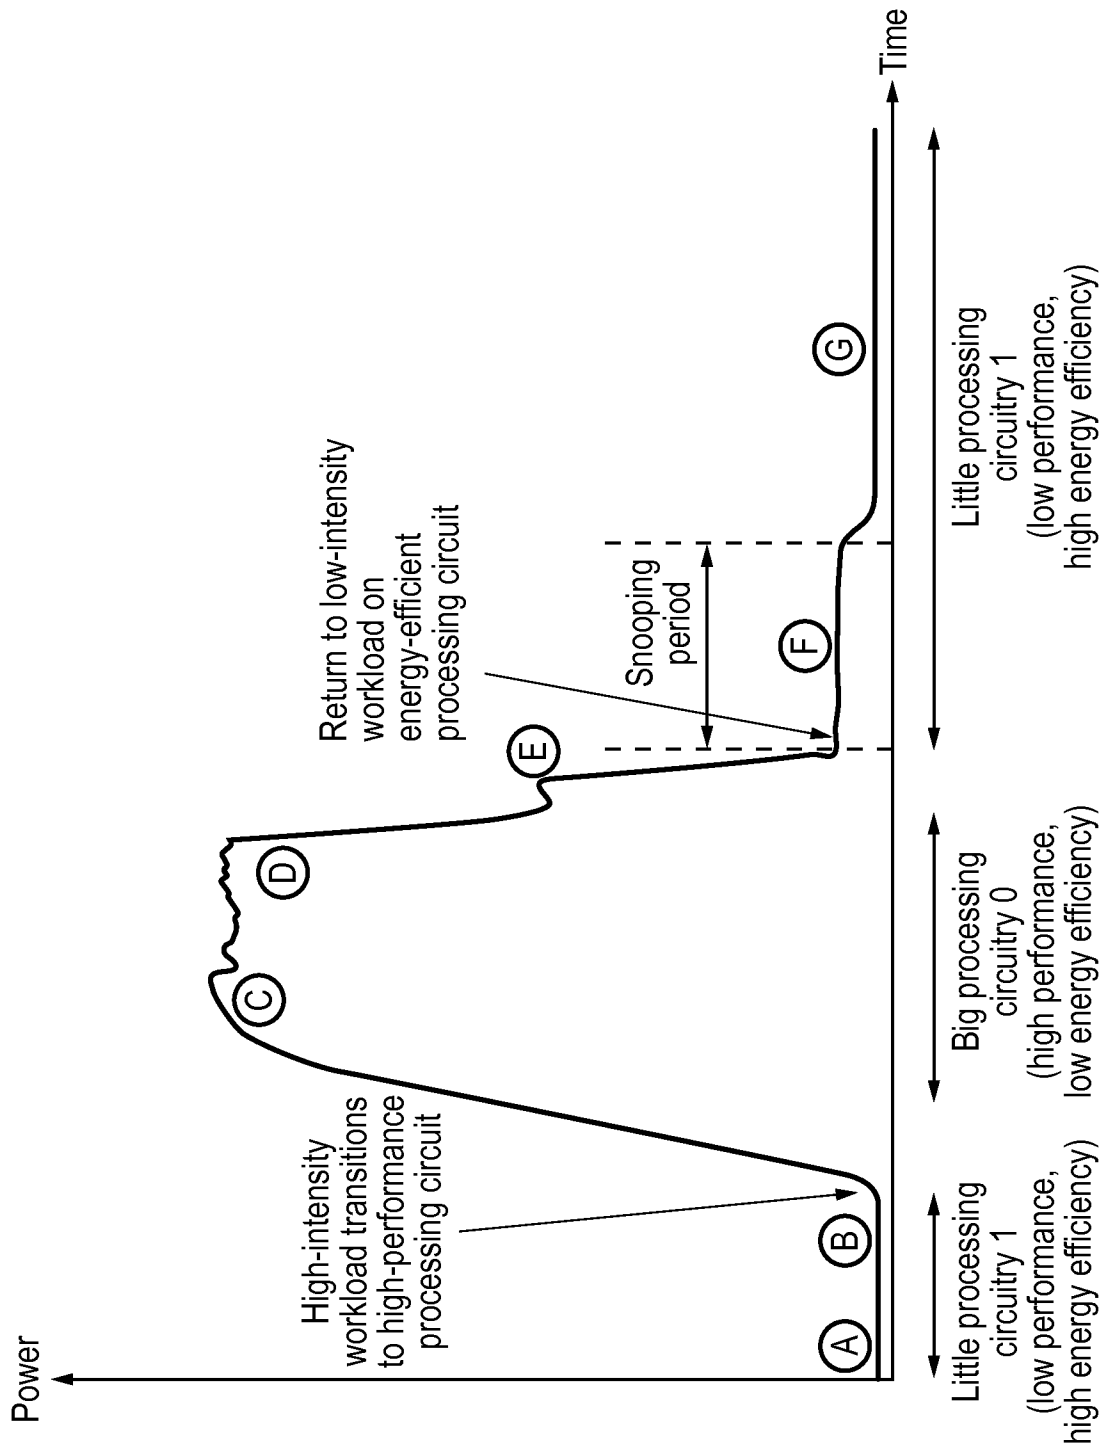


FIG. 9

## INTERNATIONAL SEARCH REPORT

International application No  
PCT/GB2011/050319

A. CLASSIFICATION OF SUBJECT MATTER INV. G06F1/32 G06F15/167 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 practical, search terms used) EPO-Internal		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 6 501 999 B1 (CAI ZHONG-NING [US]) 31 December 2002 (2002-12-31) the whole document	1-18
A	----- US 2006/064606 A1 (KIM SUHWAN [US] ET AL) 23 March 2006 (2006-03-23) the whole document	1-18
A	----- WO 2005/062180 A1 (APPLE COMPUTER [US]; ATHAS WILLIAM C [US]; MANSFIELD ROBERT L [US]; YO) 7 July 2005 (2005-07-07) the whole document	1-18
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input checked="" type="checkbox"/> 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 document 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  7 July 2011		Date of mailing of the international search report  25/07/2011
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  Moraiti, Marina

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/GB2011/050319

Patent document cited in search report	Publication date	Patent family member(s)	Publication date	
US 6501999	B1	31-12-2002	US 2003088800 A1	08-05-2003
US 2006064606	A1	23-03-2006	WO 2006034322 A2	30-03-2006
WO 2005062180	A1	07-07-2005	CN 1860446 A	08-11-2006
			EP 1695216 A1	30-08-2006
			US 2005132239 A1	16-06-2005