

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
22 March 2007 (22.03.2007)

PCT

(10) International Publication Number
WO 2007/031696 A1

(51) International Patent Classification:
G06F 12/08 (2006.01) G06F 9/38 (2006.01)

Green Ginger, Teh Green, Little Thurlow, Suffolk CB9 7JH (GB).

(21) International Application Number:
PCT/GB2005/003531

(74) Agents: HORNER, David, Richard et al.; D Young & Co, 12 Holborn, London EC1N 2DY (GB).

(22) International Filing Date:
13 September 2005 (13.09.2005)

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

(25) Filing Language: English

(26) Publication Language: English

(71) Applicant (for all designated States except US): ARM LIMITED [GB/GB]; 110 Fulbourn Road, Cherry Hinton, Cambridge CB1 9NJ (GB).

(72) Inventors; and


(75) Inventors/Applicants (for US only): GHOSH, Mrinmoy [IN/GB]; 31 Speedwell Close, Cambridge CB1 9YS (GB). ÖZER, Emre [TR/GB]; 43 Fulbourn Old Drift Road, Cambridge CB1 9ND (GB). BILES, Stuard, David [GB/GB];

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI,

[Continued on next page]

(54) Title: CACHE MISS DETECTION IN A DATA PROCESSING APPARATUS

(57) Abstract: A data processing apparatus and method are provided for detecting cache misses. The data processing apparatus has processing logic for executing a plurality of program threads, and a cache for storing data values for access by the processing logic. When access to a data value is required while executing a first program thread, the processing logic issues an access request specifying an address in memory associated with that data value, and the cache is responsive to the address to perform a lookup procedure to determine whether the data value is stored in the cache. Indication logic is provided which in response to an address portion of the address provides an indication as to whether the data value is stored in the cache, this indication being produced before a result of the lookup procedure is available, and the indication logic only issuing an indication that the data value is not stored in the cache if that indication is guaranteed to be correct. Control logic is then provided which, if the indication indicates that the data value is not stored in the cache, uses that indication to control a process having an effect on a program thread other than the first program thread.



WO 2007/031696 A1



FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT,
RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA,
GN, GQ, GW, ML, MR, NE, SN, TD, TG).

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

Published:

— *with international search report*

CACHE MISS DETECTION IN A DATA PROCESSING APPARATUS

FIELD OF THE INVENTION

The present invention relates to the use of cache miss detection techniques in a data processing apparatus.

BACKGROUND OF THE INVENTION

A data processing apparatus will typically include processing logic for executing sequences of instructions in order to perform processing operations on data items. The instructions and data items required by the processing logic will generally be stored in memory, and due to the long latency typically incurred when accessing memory, it is known to provide one or more levels of cache within the data processing apparatus for storing some of the instructions and data items required by the processing logic to allow a quicker access to those instructions and data items. For the purposes of the following description, the instructions and data items will collectively be referred to as data values, and accordingly when referring to a cache storing data values, that cache may be storing either instructions, data items to be processed by those instructions, or both instructions and data items. Further, the term data value is used herein to refer to a single instruction or data item, or alternatively to refer to a block of instructions or data items, as for example is the case when referring to a linefill process to the cache.

Significant latencies can also be incurred when cache misses occur within a cache. In the article entitled "Just Say No: Benefits of Early Cache Miss Determination" by G Memik et al, Proceedings of the Ninth International Symposium on High Performance Computer Architecture, 2003, a number of techniques are described for reducing the data access times and power consumption in a data processing apparatus with multi-level caches. In particular, the article describes a piece of logic called a "Mostly No Machine" (MNM) which, using the information about blocks placed into and replaced from caches, can quickly determine whether an access at any cache level will result in a cache miss. The accesses that are identified to miss are then aborted at that cache level. Since the MNM structures used to recognise misses are significantly smaller than the cache structures, data access time and power consumption is reduced.

The article entitled "Bloom Filtering Cache Misses for Accurate Data Speculation and Prefetching" by J Peir et al, Proceedings of the Sixteenth International Conference of

Supercomputing, Pages 189 to 198, 2002, describes a particular form of logic used to detect whether an access to a cache will cause a cache miss to occur, this particular logic being referred to as a Bloom filter. In particular, the paper uses a Bloom filter to identify cache misses early in the pipeline of the processor. This early identification of cache misses is then used to allow the processor to more accurately schedule instructions that are dependent on load instructions that are identified as resulting in a cache miss, and to more precisely prefetch data into the cache. Dependent instructions are those which require as a source operand the data produced by the instruction from which they depend, in this example the data being loaded by a load instruction.

10 The article entitled "Fetch Halting on Critical Load Misses" by N Mehta et al, Proceedings of the 22nd International Conference on Computer Design, 2004, also makes use of the Bloom filtering technique described in the above article by Peir et al. In particular, this article describes an approach where software profiling is used to identify instructions which are long latency instructions having many output dependencies, such
15 instructions being referred to as "critical" instructions. In particular, the software profiling techniques are used to identify load instructions that will be critical instructions. For any such load instructions, when those instructions are encountered in the processor pipeline, the Bloom filter technique is used to detect whether the cache lookup based on that load instruction will cause a cache miss to occur, and if so a fetch halting technique is
20 invoked. The fetch halting technique suspends instruction fetching during the period when the processor is stalled by the critical load instruction, which allows a power saving to be achieved in the issue logic of the processor.

From the above discussion, it will be appreciated that indication logic has been developed which can be used to provide an early indication of a cache miss when
25 accessing data, with that indication then being used to abort a cache access, and optionally also to perform particular scheduling or power saving activities in situations where there are dependent instructions, i.e. instructions that require the data being accessed.

SUMMARY OF THE INVENTION

30 Viewed from a first aspect, the present invention provides a data processing apparatus comprising: processing logic operable to execute a plurality of program threads; a cache operable to store data values for access by the processing logic, the

processing logic being operable when access to a data value is required whilst executing a first program thread to issue an access request specifying an address in memory associated with that data value, and the cache being operable in response to the address to perform a lookup procedure to determine whether the data value is stored in the cache;
5 indication logic operable in response to an address portion of the address to provide an indication as to whether the data value is stored in the cache, said indication being produced before a result of the lookup procedure is available, and the indication logic being operable to only issue an indication that the data value is not stored in the cache if that indication is guaranteed to be correct; and control logic operable, if the indication
10 indicates that the data value is not stored in the cache, to use said indication to control a process having an effect on a program thread other than the first program thread.

In accordance with the present invention, indication logic is used to provide an indication as to whether the data value is stored in the cache, with this indication being produced before a result of the lookup procedure is available. Further, the indication
15 logic is of the type which will only issue an indication that the data value is not stored in the cache if that indication is guaranteed to be correct. It should be noted that this does not necessarily mean that the converse will always be correct. Hence, whilst if the indication indicates that the data value is not stored in the cache, it will be the case that the cache lookup would result in a cache miss, if instead the indication does not identify
20 that the data value is not stored in the cache, it does not necessarily mean that the data value will be in the cache, and it is still the case that a cache miss may occur. Hence, the indication can produce a safe indication for a cache miss, but will not necessarily produce a safe indication for a cache hit. An example of a technique that can be used to implement the indication logic is the earlier-mentioned Bloom filter technique.

25 The address portion used by the indication logic may be the entire address of the access request or some subset of the address bits.

In accordance with the present invention, the processing logic is operable to execute a plurality of program threads, and control logic is provided which, if the indication produced by the indication logic for an access request of a first thread indicates
30 that the data value is not stored in the cache, uses that indication to control a process having an effect on a program thread other than the first program thread.

Hence, in accordance with the present invention, the early indication of a cache miss produced by the indication logic in relation to an access request of a particular program thread can be used by the control logic to influence the handling of other program threads, which can lead to significant performance improvements in such multi-
5 threaded systems.

In one embodiment, the data processing apparatus further comprises additional control logic operable in response to said indication to control a process having an effect on the first program thread. In one particular embodiment, this additional control logic is lookup control logic operable in response to said indication to abort the lookup procedure
10 if the indication indicates that the data value is not stored in the cache. Hence, in such embodiments, the data processing apparatus makes multiple uses of the indication produced by the indication logic. In particular, the lookup control logic will abort the lookup procedure if the indication indicates that the data value is not stored in the cache, which can result in significant power savings due to the fact that the indication is
15 produced before the result of the lookup procedure is available. In some embodiments, the indication will be produced before the lookup procedure has even been initiated, whilst in other embodiments the cache may have received the access request and have begun the lookup procedure, but power savings can still be achieved by abortion of the lookup procedure prior to completion. In addition to this use of the indication, if the
20 indication indicates that the data value is not stored in the cache, the earlier-mentioned control logic uses that indication to have an effect on a program thread other than the first program thread. This two-fold approach can produce both power savings and significant performance improvements within the multi-threaded system.

The data processing apparatus can be arranged to handle the multiple program
25 threads in a variety of ways. In one embodiment, the control logic is operable, if the indication indicates that the data value is not stored in the cache, to cause the processing logic to perform a context switch operation in order to switch execution between the first program thread and another program thread. In such embodiments, the early detection of the cache miss when handling the access request of the current thread can be used to
30 invoke an early initiation of a context switch operation, thereby improving the performance of the data processing apparatus by early switching to another program thread in situations where the first program thread is going to incur a latency from the

detected cache miss. An example of a data processing apparatus where such an implementation will produce significant performance benefits is in a coarse-grained multi-threading apparatus (also known as “block multi-threading”), where the entirety of the processing resources of the processing logic are switched to the other program thread by the context switch operation.

In another embodiment, the processing logic comprises a number of shared resources, each shared resource being allocatable to one of the program threads, and the control logic is operable, if the indication indicates that the data value is not stored in the cache, to cause the processing logic to perform a reallocation of at least one shared resource. In such embodiments the shared resources can be shared amongst the program threads, and the way in which these resources are allocated amongst the threads can be varied. A processor operating in such a manner is often referred to as a simultaneous multi-threading (SMT) processor. In accordance with this embodiment of the present invention, the control logic is arranged, if the indication indicates that the data value is not stored in the cache, to cause the processing logic to perform a reallocation of at least one shared resource. Hence, if the indication logic determines for an access request of a particular thread that a cache miss is going to occur, that information can be used by the control logic to cause a reallocation of at least one of the shared resources, and hence for example this can be used to free-up that shared resource for use by one or more of the other threads. Since the indication produced by the indication logic occurs before the result of the lookup procedure is available, this reallocation of resources can occur at a relatively early stage. This can lead to significant performance improvements through improved resource utilisation.

In another embodiment, the processing logic comprises a plurality of processors each sharing access to the cache, and each being operable when access to a data value is required to issue an access request specifying an address in memory associated with that data value; the plurality of program threads are distributed across the plurality of processors; the indication logic is operable in response to at least some of the access requests to provide said indication; and the control logic comprises arbitration logic operable to perform an arbitration process to arbitrate access to the cache between the plurality of processors, and being arranged to use the indications provided by the indication logic when performing said arbitration process.

In accordance with this embodiment, a plurality of processors share access to the cache and the plurality of program threads are distributed across the plurality of processors. Hence, for example, certain program threads may be executed by one processor, whilst other program threads are executed by another processor. Arbitration logic is used to arbitrate access to the cache between the plurality of processors, and the indications generated by the indication logic are routed to the arbitration logic for use when performing the arbitration process. There are a number of ways in which these indications could be used by the arbitration logic. However, as an example, an indication that a particular access request will result in a cache miss in the cache can be used to cause the arbitration logic to ignore that particular access request on the basis that it will be aborted, or at the very least to reduce the priority associated with that access request. As a result, the input of the indications from the indication logic into the arbitration logic can significantly reduce contention for access to the cache, and thereby increase the available bandwidth for access to the cache.

In such embodiments, it will be appreciated that the cache can take a variety of forms. For example, in one embodiment, the cache may be a multi-ported cache consisting of multiple cache banks, and in such instances the arbitration logic can be replicated for each cache bank.

In addition, whilst the indication logic may be a single piece of logic used to provide an indication for access requests issued by any of the plurality of processors, in alternative embodiments the indication logic comprises a plurality of indication units, each indication unit being associated with one of said processors and being operable in response to the address portion of an access request issued by the associated processor to provide said indication as to whether the data value is stored in the cache; the control logic being operable to use the indications provided by each of the indication units when performing said arbitration process.

The indication logic can take a variety of forms. In one embodiment, the indication logic comprises: counter logic having a plurality of counter entries, each counter entry containing a count value; vector logic having a vector entry for each counter entry in the counter logic, each vector entry containing a value which is set when the count value in the corresponding counter entry changes from a zero value to a non-zero value, and which is cleared when the count value in the corresponding counter

entry changes from a non-zero value to a zero value; index generation logic operable to generate from an address portion of an address at least one index, each index identifying a counter entry and associated vector entry; the index generation logic being operable whenever a data value is stored in, or removed from, the cache to generate from the address portion of the associated address said at least one index, and to cause the count value in each identified counter entry to be incremented if the data value is being stored in the cache or decremented if the data value is being removed from the cache; the index generation logic further being operable for at least some access requests to generate from the address portion of the associated address said at least one index, and to cause the vector logic to generate an output signal based on the value in each identified vector entry, the output signal indicating if the data value of the access request is not stored in the cache.

The actual incrementing or decrementing of the relevant count value(s) can take place at a variety of times during the storage or removal process. For example, the relevant count value(s) can be incremented at the time of allocating a data value to the cache or at some time later when the actual data value is stored as a result of the linefill procedure. Similarly, the relevant count value(s) can be decremented at the time when a data value has been chosen for eviction or at some time later when the data value is overwritten as a part of the linefill process following the eviction.

In accordance with this embodiment, separate counter logic and vector logic are provided within the indication logic, the counter logic being updated based on data values being stored in, or removed from, the cache. In particular, index generation logic generates from an address portion of an address at least one index, with each index identifying a counter entry. The count values in those counter entries are then incremented if data is being stored in the cache or decremented if data is being removed from the cache. The corresponding vector entries in the vector logic can then be queried for access requests using that address portion, in order to generate an output signal which indicates if the data value of the access request is not stored in the cache.

In one embodiment, each index is generated from the address portion using a hash function, and in some embodiments multiple hash functions can be used to generate multiple indexes from a single address portion. However, in one embodiment, a single

hash function is used, resulting in the generation of a single index for each address portion.

In one embodiment, a memory hierarchy exists within the data processing apparatus having a number of memory levels, the cache being provided at one of said
5 memory levels, and when processing access requests memory levels higher in the hierarchy being accessed before memory levels lower in the hierarchy; and whilst the vector logic and the counter logic are both associated with the cache, the vector logic is accessible separately to the counter logic and from a different level in the memory
10 hierarchy to that memory level from which the counter logic is accessible. Through such an approach the vector logic can be located within a different part of the data processing apparatus to that in which the counter logic is located. Whilst in one embodiment the index generation logic can be shared between the counter logic and the vector logic, in an alternative embodiment the index generation logic is replicated for
15 both the counter logic and the vector logic, which assists in facilitating the placement of the vector logic at a different location within the data processing apparatus to that in which the counter logic is located.

It should be noted that to generate the output signal indicating if the data value of an access request is not stored in the cache, only the vector logic needs to be accessed. The vector logic is typically significantly smaller than the counter logic, and hence by
20 arranging for the vector logic to be accessible separately to the counter logic and from a different level in the memory hierarchy, this enables the output signal to be generated relatively quickly, and with relatively low power. Furthermore, it is typically the case that the counter logic will be updated more frequently than the vector logic, since a vector entry in the vector logic only needs updating when the count value in the corresponding
25 counter entry changes from a zero value to a non-zero value, or from a non-zero value to a zero value. Hence, by enabling the vector logic to be accessed separately to the counter logic and from a different level in the memory hierarchy, this enables the counter logic to be run at a lower frequency than that at which the vector logic is run, thus producing power savings. Running of the counter logic at a lower frequency is acceptable since the
30 operation of the counter logic is not time critical.

Viewed from a second aspect, the present invention provides a data processing apparatus comprising: processing logic; a cache operable to store data values for access

by the processing logic, the processing logic being operable when access to a data value is required to issue an access request specifying an address in memory associated with that data value, and the cache being operable in response to the address to perform a lookup procedure to determine whether the data value is stored in the cache; counter logic
5 having a plurality of counter entries, each counter entry containing a count value; vector logic having a vector entry for each counter entry in the counter logic, each vector entry containing a value which is set when the count value in the corresponding counter entry changes from a zero value to a non-zero value, and which is cleared when the count value in the corresponding counter entry changes from a non-zero value to a zero value; index
10 generation logic operable to generate from an address portion of an address at least one index, each index identifying a counter entry and associated vector entry; the index generation logic being operable whenever a data value is stored in, or removed from, the cache to generate from the address portion of the associated address said at least one index, and to cause the count value in each identified counter entry to be incremented if
15 the data value is being stored in the cache or decremented if the data value is being removed from the cache; the index generation logic further being operable for at least some access requests to generate from the address portion of the associated address said at least one index, and to cause the vector logic to generate an output signal based on the value in each identified vector entry, the output signal indicating if the data value of the
20 access request is not stored in the cache; wherein a memory hierarchy exists within the data processing apparatus having a number of memory levels, the cache being provided at one of said memory levels, and when processing access requests memory levels higher in the hierarchy being accessed before memory levels lower in the hierarchy; and whilst the vector logic and the counter logic are both associated with the cache, the vector logic is
25 accessible separately to the counter logic and from a different level in the memory hierarchy to that memory level from which the counter logic is accessible.

In accordance with this aspect of the present invention, a novel form of indication logic is provided consisting of counter logic, vector logic and index generation logic. The vector logic is accessible separately to the counter logic and
30 from a different level in the memory hierarchy, and when generating an output signal indicating if the data value of the access request is not stored in the cache, only the vector logic needs to be accessed. As mentioned earlier, the output signal from the

vector logic can hence be generated quickly and with low power, and the counter logic can be run at a lower frequency than the vector logic, thereby producing power savings. Running of the counter logic at a lower frequency is acceptable since the operation of the counter logic is not time critical.

5 By such an approach, the indication logic can implement the functionality of a bloom filter type technique, by producing a safe indication of a cache miss, but in a more power efficient and flexible manner. In addition to the power and speed savings mentioned above, the provision of the vector logic separately to the counter logic enables the vector logic to be located in a different part of the data processing
10 apparatus to the counter logic, which can give rise to a number of benefits as will be discussed later.

In one embodiment the index generation logic comprises first index generation logic associated with the counter logic and second index generation logic associated with the vector logic; whenever a data value is stored in, or removed from, the cache the first
15 index generation logic being operable to generate from the address portion of the associated address said at least one index identifying one or more counter entries in the counter logic; and for at least some access requests the second index generation logic being operable to generate from the address portion of the associated address said at least one index identifying one or more vector entries in the vector logic.

20 The use of separate index generation logic for the counter logic and for the vector logic provides increased flexibility in the separation of the counter logic and the vector logic.

In one embodiment, the data processing apparatus further comprises lookup control logic operable in response to said output signal from the vector logic to abort
25 the lookup procedure if the output signal indicates that the data value is not stored in the cache. Preferably, the vector logic can be accessed such that the output signal is produced before the result of the lookup procedure is available, and hence abortion of the lookup procedure if the output signal indicates that the data value is not stored in the cache can result in significant power savings.

30 The output signal generated by the vector logic can be used in a variety of ways. However, in one embodiment, the processing logic is operable to execute a sequence of instructions in order; and if the output signal from the vector logic

indicates that the data value is not stored in the cache, the processing logic is operable to perform an energy management process to reduce energy consumption of the processing logic whilst waiting for the data value to be accessed. Since in such embodiments the instructions are executed in order, then a cache miss is likely to stall
5 the processing logic for a significant period of time, and by accessing the vector logic a fast and efficient indication of a cache miss can be obtained which enables an early invocation of the energy management process to reduce energy consumption within the processing logic. The energy management process may involve a number of different measures being put in place, and hence for example may involve running the
10 processing logic at a lower processing speed, or fully or partially shutting down elements of the processing logic until such time as the data value has been accessed.

In one embodiment, the data processing apparatus further comprises: an additional cache, if on issuance of the access request by the processing logic a lookup is performed in the additional cache resulting in a miss in the additional cache, the
15 cache being operable to perform said lookup procedure; and if the output signal from the vector logic indicates that the data value is not stored in the cache, at least one of the cache and additional cache being operable to perform an energy management process to reduce energy consumption whilst waiting for the data value to be accessed. Hence, in such embodiments the energy consumption in one or both of the cache and
20 the additional cache can be reduced dependent on the output signal produced by the vector logic. Given the provision of the vector logic separately to the counter logic, the output signal can be generated in a particularly quick and efficient manner, hence enabling the power savings achieved by such an early invocation of the energy management process to be significant. The energy management process applied to the
25 cache(s) can take a variety of forms. However, as an example, the caches can be put in a “drowsy” mode of operation, where the state of the cache, and hence the data stored therein, is retained, but the caches are not accessible whilst in the drowsy mode. This can for example be achieved by reducing the voltage supplied to the caches. To access the cache again, it will typically be necessary to exit the drowsy mode first.

30 In one embodiment employing an additional cache, the cache and additional cache exhibit inclusive behaviour, and if the output signal from the vector logic indicates that the data value is not stored in the cache the lookup in the additional

cache is aborted. Hence, in such embodiments, the cache and additional cache exhibit inclusive behaviour, and hence any data stored in the additional cache will be replicated in the cache. Accordingly, an indication of a miss in the cache will also indicate that a miss will occur in the additional cache. By provision of vector logic
5 separate to the counter logic it is possible to produce the output signal from the vector logic at an early enough stage to enable the cache lookup in the additional cache to be aborted, in addition to any abortion of the lookup in the cache. Hence, significant energy savings can be realised by avoiding lookups in multiple caches. It will be appreciated that this approach can be used when there are more than two levels of
10 cache in the system assuming inclusive behaviour is exhibited between the caches, thereby potentially enabling multiple cache lookups to be aborted, resulting in significant power savings.

In one embodiment, the vector logic is accessible prior to the access request being processed by the memory level in which the cache is provided. Due to the
15 provision of the separate vector logic and counter logic, it is possible in such embodiments to locate the vector logic so that it is accessible before the access request is processed by the memory level in which the cache is provided. This allows a particularly early generation of the output signal and hence enables any actions taken in dependence on that output signal to be initiated at a particularly early stage, thereby
20 increasing the benefits to be realised in terms of power saving, performance improvement, etc.

In one embodiment, the data processing apparatus further comprises an additional cache at a higher memory level than said cache, if on issuance of the access request by the processing logic a lookup is performed in the additional cache resulting in
25 a miss in the additional cache, said cache being operable to perform said lookup procedure; and the vector logic being accessed in response to a cache miss signal issued by the additional cache in the event of a cache miss in the additional cache, thereby causing the vector logic to generate said output signal. Hence, in such embodiments the cache miss signal issued by the additional cache is used to trigger the
30 vector logic to generate the output signal. This can enable the subsequent access to the cache to be aborted at an early stage, and can also in some embodiments enable power

management processes to be instigated within the processing logic and the various levels of cache earlier than would otherwise be possible.

In one embodiment having an additional cache at a higher memory level than said cache, said cache and said additional cache exhibit inclusive behaviour, and the vector logic is located within the processing logic and is accessed for every access request, thereby causing the vector logic to generate said output signal. Hence, for embodiments where the caches exhibit inclusive behaviour, the locating of the vector logic within the processing logic enables the output signal to be generated for every access request prior to that access request being issued to the caches within the memory hierarchy. Hence, this enables the cache lookups within those caches to be avoided in situations where the output signal indicates that there will be a cache miss in the cache.

In one embodiment, the processing logic comprises a plurality of processors, each processor being operable when access to a data value is required to issue an access request specifying an address in memory associated with that data value; the cache comprises a plurality of cache portions, each cache portion associated with one of said processors, and each cache portion being accessible by said plurality of processors; separate counter logic, vector logic and index generation logic being provided in association with each cache portion; the vector logic associated with at least one cache portion being accessed for access requests issued by a processor with which the at least one cache portion is not associated.

In such embodiments, the cache comprises of plurality of cache portions and separate counter logic, vector logic and index generation logic is provided in association with each cache portion. The cache portions may be separate cache banks within a cache, or instead may be physically separate. In accordance with this embodiment, the vector logic associated with at least one cache portion is arranged so that it can be accessed for access requests issued by a processor with which the at least one cache portion is not associated. Hence, this processor can get an early indication of a cache miss with a single and fast access to the relevant vector logic. Since in this embodiment each processor can access the cache portions associated with other processors, there is likely to be a significant contention for access to the cache portions, and by adopting the above technique, the number of access requests being

propagated to the cache can be reduced, thereby reducing such contention.

The way in which the vector logic is arranged to be accessible by multiple processors can be implemented in a variety of ways. However, in one embodiment the vector logic is replicated for each processor. This will result in each processor having
5 accessing to multiple instances of vector logic, and in such cases these multiple instances of the vector can be accessed simultaneously using the same index generation logic. Hence, with a single access using a single index generation logic, the processing logic will know which cache portions, including the cache portion associated with it, will result in a cache miss, and this information can be used to
10 reduce the number of access to the cache portions.

In one embodiment the cache has a plurality of potential locations where the data value the subject of the access request may reside, when performing the lookup procedure the cache being operable to perform a sequence of lookups, each lookup relating to a subset of said potential locations; and if the output signal from the vector
15 logic indicates that the data value is not in the cache, the sequence of lookups being terminated.

There are many examples of caches where a sequence of lookups needs to be performed in order to look in each of the potential locations where the data value may reside. For example, in a high way set associative cache, the cache may be unable to
20 look in all ways in the same cycle, and accordingly will need to perform a sequence of lookups in order to determine whether the data value is in the cache. By using the vector logic of embodiments of the present invention, the early indication produced by the vector logic can be used to terminate the sequence of lookups if the output signal indicates the data value is not in the cache, i.e. is not in any of the potential locations.
25 This can result in a saving of much of the cache energy that would otherwise be consumed in performing the sequence of lookups. As another example, caches that are accessed using virtual addresses (often referred to as virtual caches) may require a sequence of lookups to take place in order to perform a lookup in each of the potential locations where the data value may reside. This results from the possibility of multiple
30 virtual addresses mapping to the same physical address, often referred to as the synonym problem. Again, the early indication from the output signal can be used to terminate the sequence of lookups early if it is determined that the data value is not in

the cache.

Viewed from a third aspect, the present invention provides a method of operating a data processing apparatus having processing logic for executing a plurality of program threads, and a cache for storing data values for access by the processing logic, the method comprising the steps of: when access to a data value is required whilst the processing logic is executing a first program thread, issuing an access request specifying an address in memory associated with that data value; in response to the address, performing a lookup procedure in the cache to determine whether the data value is stored in the cache; in response to an address portion of the address, employing indication logic to provide an indication as to whether the data value is stored in the cache, said indication being produced before a result of the lookup procedure is available, and the indication logic only issuing an indication that the data value is not stored in the cache if that indication is guaranteed to be correct; and if the indication indicates that the data value is not stored in the cache, using said indication to control a process having an effect on a program thread other than the first program thread.

Viewed from a fourth aspect, the present invention provides a method of operating a data processing apparatus, the data processing apparatus comprising processing logic, a cache for storing data values for access by the processing logic, counter logic associated with the cache and having a plurality of counter entries, each counter entry containing a count value, and vector logic associated with the cache and having a vector entry for each counter entry in the counter logic, each vector entry containing a value which is set when the count value in the corresponding counter entry changes from a zero value to a non-zero value, and which is cleared when the count value in the corresponding counter entry changes from a non-zero value to a zero value, the method comprising the steps of: when access to a data value is required by the processing logic, issuing an access request specifying an address in memory associated with that data value; in response to the address, performing a lookup procedure in the cache to determine whether the data value is stored in the cache; whenever a data value is stored in, or removed from, the cache, generating from the address portion of the associated address at least one index, each index identifying a counter entry, and causing the count value in each identified counter entry to be incremented if the data value is being stored in the cache or decremented if the data value is being removed from the cache; for at

least some access requests, generating from the address portion of the associated address at least one index, each index identifying a vector entry, and causing the vector logic to generate an output signal based on the value in each identified vector entry, the output signal indicating if the data value of the access request is not stored in the cache; wherein
5 a memory hierarchy exists within the data processing apparatus having a number of memory levels, the cache being provided at one of said memory levels, and when processing access requests memory levels higher in the hierarchy being accessed before memory levels lower in the hierarchy; and said method further comprising the step of
10 accessing the vector logic separately to the counter logic and from a different level in the memory hierarchy to that memory level from which the counter logic is accessible.

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:

Figure 1 is a block diagram illustrating application of an early miss indication
15 technique in a coarse grained multithreaded processor in accordance with one embodiment;

Figure 2 is a block diagram illustrating application of the early cache miss indication technique in a simultaneous multi-threaded processor of one embodiment;

Figure 3 is a block diagram illustrating application of the early cache miss
20 indication technique in a chip multi-processor of one embodiment;

Figure 4 is a block diagram schematically illustrating the use of a segmented counting bloom filter in a data processing apparatus in accordance with one embodiment;

Figure 5 is a diagram schematically illustrating in more detail the segmented counting bloom filter of Figure 4;

Figure 6 is a block diagram illustrating in more detail the data processing
25 apparatus of Figure 4 in accordance with one embodiment;

Figure 7 is a block diagram illustrating in more detail the data processing apparatus of Figure 4 in accordance with an alternative embodiment;

Figure 8 is a diagram schematically illustrating the operation of the hash function
30 illustrated in Figure 5 in accordance with one embodiment;

Figure 9 is a diagram schematically illustrating the application of the segmented bloom filter of one embodiment in association with a virtually indexed physically tagged cache;

Figure 10 is a diagram schematically illustrating the application of the segmented bloom filter of one embodiment in association with a virtually indexed virtually tagged cache;

Figure 11 is a block diagram illustrating a cross-firing technique used with a segmented bloom filter in accordance with one embodiment; and

Figure 12 is a diagram schematically illustrating the cross-firing technique employed in a further embodiment.

DESCRIPTION OF EMBODIMENTS

Figure 1 is a block diagram of a data processing apparatus in accordance with one embodiment. In particular Figure 1 illustrates a coarse grained multi-threaded processor 10 which includes a processing pipeline consisting of fetch logic 70, decode and issue logic 75 and execution logic 80. At any point in time, this processing pipeline is used to execute a particular program thread. A context bank 30 is provided for maintaining the context for a number of different program threads, and for the particular program thread being handled by the processing pipeline, the various elements of the processing pipeline can obtain the relevant information pertaining to that thread from the context bank 30. Hence, the context bank 30 will include multiple program counter values 40, 42, 44 for the various program threads, multiple program status registers 50, 52, 54 for the various program threads, and multiple register banks 60, 62, 64 for storing the data required when executing instructions of the various program threads.

The fetch logic 70 can then access the relevant program counter value 40, 42, 44 for the particular program thread being executed, and can issue updated program counter values back to the context bank 30. Similarly, the decode and issue logic 75 can output to the context bank 30 a source register identifier and a thread identifier (TID), to cause the required register in one of the register banks 60, 62, 64 to be accessed and for the value in that register to be returned to the decode and issue logic for routing to the execution logic 80. The result value produced by the execution logic can then be routed back to the context bank 30 along with the relevant TID to cause the data value to be stored in the relevant register of the appropriate register bank 60, 62, 64.

The CPU 10 also includes thread switch logic 20 which is used to determine which thread should execute and to cause the appropriate context switch to occur when switching between program threads.

When the fetch logic 70 wishes to retrieve an instruction, it issues an access request identifying an instruction address to the level 1 instruction cache 90. If the instruction is found in that cache, then this is returned to the fetch logic, along with a control signal indicating that there has been a hit. However, if the instruction is not in the cache, then a miss signal is returned to the fetch logic, and the appropriate address for a linefill to the level 1 instruction cache 90 is output to the level 2 cache 105. If there is a hit in the level 2 cache, then the relevant line of instructions is returned to the level 1 instruction cache 90 along with a control signal indicating a hit in the level 2 cache. However, if there is a miss in the level 2 cache, then a miss signal is returned to the level one instruction cache 90, and the line address is propagated from the level 2 cache 105 to memory 110. This will ultimately result in the line of instructions being returned to the level 2 cache and propagated on to the fetch logic 70 via the level 1 instruction cache 90.

Similarly, when the execution logic 80 executes a load or a store instruction, the address of the data to be accessed will be output from the execution logic to the level 1 data cache 95. In the event of a store operation, this will also be accompanied by the write data to be stored to memory. The level 1 data cache 95 will issue a hit/miss control signal to the execution logic 80 indicating whether a cache hit or a cache miss has occurred in the level 1 data cache, and in the event of a load operation which has hit in the cache will also return the data to the execution logic 80. In the event of a cache miss in the level 1 data cache, the line address will be propagated onto the level 2 cache 105, and in the event of a hit the line of data values will be accessed in the level 2 cache. For a load this will cause the line of data values to be returned to the level 1 data cache 95 for storing therein. In the event of a miss in the level 2 cache 105, the line address will be propagated onto memory 110 to cause the line of data to be accessed in the memory.

In accordance with the embodiment described in Figure 1, cache miss indication logic 100 is provided which receives any line addresses output by either the instruction cache 90 or the data cache 95, and in response to that line address information determines whether a miss in the level 2 cache 105 will occur. The

indication produced by the cache miss indication logic 100 will be produced before the result of the lookup procedure in the level 2 cache 105 is available, and that cache miss indication will be routed to the thread switch logic 20 within the CPU 10. If the cache miss indication indicates that a miss will occur in the level 2 cache 105 then the thread
5 switch logic 20 is arranged to perform a context switch operation in order to switch execution between the current program thread and another program thread. As a result, it will issue a control signal to the context bank 30 to cause the context of the current thread to be saved, after which the pipeline processing logic will start executing the designated new thread.

10 The early cache miss indication produced by the cache miss indication logic 100 can hence be useful in starting the process of a context switch earlier than may otherwise be possible, thereby improving the performance of the processor by starting the next thread earlier.

The cache miss indication logic 100 can take a variety of forms, but is arranged
15 to provide a safe indication of a cache miss. Accordingly, if the cache miss indication indicates that there will be a cache miss in the level 2 cache, then this is guaranteed to be correct, as opposed to being a mere prediction. However, in one embodiment, if the cache miss indication is unable to determine that there will be a cache miss, this does not necessarily mean that there will not be a cache miss, and accordingly the cache
20 miss indication logic does not provide a safe indication of a cache hit.

In one embodiment, the cache miss indication logic 100 uses a bloom filter technique and maintains information which is updated each time a data value is stored in, or removed from, the level 2 cache 105, based on replacement and linefill information routed to the cache miss indication logic 100 from the level 2 cache 105.
25 More details of one particular embodiment of the cache miss indication logic 100 will be discussed in detail later.

Whilst the cache miss indication logic 100 has been shown in Figure 1 as being accessed based on the line addresses issued by the level 1 instruction cache 90 or data cache 95, it could instead be located so as to be accessed based on the instruction
30 address issued by the fetch logic 70 and/or the data address issued by the execution logic 80. This may for example be an appropriate location if the level 1 and level 2 caches employ inclusive behaviour, such that an indication of a miss in the level 2

cache will also indicate a miss in the level 1 cache. In such instances, the cache miss indication produced by the cache miss indication logic 100 can produce an earlier indication for routing to the thread switch logic 20.

Figure 2 is a block diagram of a data processing apparatus in accordance with an alternative embodiment of the present invention, where the CPU 200 takes the form of a simultaneous multi-threaded (SMT) processor. The core 200 consists of fetch, decode and rename logic 205, an issue queue 210, issue logic 220, execution logic 225, and retire logic 230. A number of these logic blocks have resources that are shared by multiple program threads. Examples of such shared resources are the fetch queue in the fetch, decode and rename logic 205, the issue queue 210 and a shared register file within the retire logic 230. For simplicity, figure 2 only illustrates the sharing of the issue queue 210 in detail for a two-way SMT processor.

As illustrated in Figure 2, the issue queue 210 has a partition 215 to divide entries of two threads for the two-way SMT. For an n-way SMT, the issue queue 210 would have n-1 partitions. The arrows illustrated next to the partition 215 in Figure 2 illustrate that the location of the partition 215 can be varied. In particular, the partition can be moved in either direction to increase or decrease entries allocated to each thread.

The instruction cache 90, data cache 95, level 2 cache 105, memory 110 and cache miss indication logic 100 operate in an analogous manner to that described earlier with reference to Figure 1. However, in the example of Figure 2, the cache miss indication from the cache miss indication logic 100 is routed to the fetch, decode and rename logic 205, the issue queue 210 and the retire logic 230 to control reallocation of the shared resources therein. In particular, if for a particular thread a cache miss indication is produced indicating that an access request of that thread will cause a cache miss in the level 2 cache 105, that indication can be used by these various logic elements to control repartitioning of the shared resources. Considering by way of example the issue queue 210, if the cache miss indication from the cache miss indication logic 100 indicates that an access for thread 0 is going to result in a level 2 cache miss 105, the issue queue 210 may use this information to slide the partition 215 so as to allocate more of the shared resource to thread 1. Similar

processes can also take place within the fetch, decode and rename logic 205 in relation to the fetch queue, and within the retire logic 230 in relation to the shared register file.

Whilst the cache miss indication logic 100 has been illustrated in Figure 2 as being responsive to the line address outputs from the level 1 instruction cache 90 and data cache 95, in an alternative embodiment the cache miss indication logic 100 may be located between the CPU 200 and the level 1 caches, and accessed based on the instruction address issued by the fetch, decode and rename logic 205 or the data address issued by the execution logic 225.

For either placement of the cache miss indication logic 100, the early cache miss indication produced by the cache miss indication logic 100 can be used to cause an early re-allocation of shared resources, thereby helping in freeing-up resources for use by other threads. This can lead to significant performance improvements within the SMT system.

Whilst not explicitly shown in Figure 1 or Figure 2, the cache miss indication produced by the cache miss indication logic 100 can also be routed to the level 2 cache to cause the associated access request to be aborted within the level 2 cache if the cache miss indication indicates that a cache miss will occur.

Figure 3 is a block diagram of an alternative embodiment of the present invention where the data processing apparatus has multiple processor cores 300, 310, each having their own level 1 caches 305, 315, respectively, but each sharing access to a level 2 cache. Such a system will be referred to herein as a chip multi-processor (CMP) system. Whilst only two cores 300, 310 are shown in Figure 3, the techniques illustrated can be used for a CMP system having any number of cores.

As shown in Figure 3, each core 300, 310 has an associated level 1 cache 305, 315 and in the event of a miss in the level 1 cache, an access request is propagated onto the level 2 cache via the arbiter logic 320. The job of the arbiter logic 320 is to arbitrate between multiple access requests seeking access to the level 2 cache.

For the purposes of illustration, the level 2 cache is shown as a multi-ported cache having multiple cache banks 350, 352, 354, 356, each cache bank having its associated tag portion and data portion. Within the arbiter logic 320, separate arbitration logic 340, 342, 344, 346 is provided for each cache bank. Routing logic 330, 335 within the arbiter 320 is arranged based on the received access requests to

route those access requests to the appropriate arbitration logic 340, 342, 344, 346, dependent on which cache bank 350, 352, 354, 356 is to be accessed.

Cache miss indication logic 325 is provided within the arbiter logic 320 which, based on the address information associated with each access request routed to the
5 arbiter 320, determines if a cache miss in the level 2 cache will occur, and outputs that cache miss indication to one or more of the arbitration logic units 340, 342, 344, 346. Each arbitration logic then takes that cache miss indication information into account when arbitrating access to the associated cache bank.

Hence, by way of example, if the arbitration logic 340 is seeking to arbitrate
10 between an access request issued by the core 300 and an access request issued by the core 310, and the cache miss indication logic 325 indicates that the access request from the core 300 will result in a cache miss in the level 2 cache bank 350, then the arbitration logic 340 will be arranged to grant access to the cache bank 350 to the access request issued by the core 310. This can be achieved in a variety of ways. For
15 example, the arbitration logic 340 may be arranged to merely ignore the access request issued by the core 300.

As shown in Figure 3, the cache miss indication logic 325 receives from the level 2 cache the replacement and linefill information used to maintain the necessary information within the cache miss indication logic required to produce the cache miss
20 indication.

Whilst in Figure 3 the cache miss indication logic 325 is shown residing within the arbiter 320, in an alternative embodiment separate cache miss indication logic can be provided in association with each core 300, 310, and in particular can be arranged to review any access requests output to the arbiter 320 from the associated core in
25 order to produce the corresponding cache miss indication. The cache miss indications output by each cache miss indication logic can then be routed to the relevant arbitration logic units 340, 342, 344, 346 within the arbiter 320.

The use of the cache miss indication logic 325 as shown in Figure 3 can significantly reduce the number of accesses to the level 2 cache, and for a particular
30 modelled example has been shown to reduce the accesses by about 30%. Hence, it can be seen that such an approach reduces contention for the level 2 cache, and as a result can increase the available bandwidth for accessing the level 2 cache.

The cache miss indication logic 100, 325 shown in Figures 1, 2 and 3 can be implemented in a variety of ways, but its main purpose is to produce the indication before a result of the lookup procedure in the level 2 cache is available, and furthermore the indication logic will only issue an indication that the data value is not stored in the cache if that indication is guaranteed to be correct, i.e. an indication of a cache miss is a safe indication, rather than merely a prediction. In one particular embodiment, the cache miss indication logic uses a Bloom filter technique.

Bloom Filters were named after Burton Bloom for his seminal paper entitled "Space/time trade-offs in hash coding with allowable errors", Communications of the ACM, Volume 13, Issue 4, July 1970. The purpose was to build memory efficient database applications. Bloom filters have found numerous uses in networking and database applications in the following articles:

A. Border and M. Mitzenmacher, "Network application of Bloom Filters: A Survey", in 40th Annual Allerton Conference on Communication, Control, and Computing, 2002;

S. Rhea and J. Kubiatowicz, "Probabilistic Location and Routing", IEEE INFOCOM'02, June 2002;

S. Dharmapurikar, P. Krishnamurthy, T. Sproull and J. Lockwood, "Deep Packet Inspection using Parallel Bloom Filters", IEEE Hot Interconnects 12, Stanford, CA, August 2003;

A. Kumar, J. Xu, J. Wang, O. Spatschek, L. Li, "Space-Code Bloom Filter for Efficient Per-Flow Traffic Measurement", Proc. IEEE INFOCOM, 2004;

F. Chang, W. Feng and K. Li, "Approximate Caches for Packet Classification", IEEE INFOCOM'04, Mar. 2004;

S. Cohen and Y. Matias, "Spectral Bloom Filters", Proceedings of the 2003 ACM SIGMOD International Conference on Management of Data, 2003; and

L. Fan, P. Cao, J. Almeida, and A. Broder, "Summary cache: A scalable wide-area Web cache sharing protocol," IEEE/ACM Transactions on Networking, vol. 8, no. 3, pp. 281–293, 2000.

For a generic Bloom filter, a given address portion in N bits is hashed into k hash values using k different random hash functions. The output of each hash function is an m -bit index value that addresses a Bloom filter bit vector of 2^m . Here, m is typically

much smaller than N . Each element of the Bloom filter bit vector contains only 1 bit that can be set. Initially, the Bloom filter bit vector is zero. Whenever an N -bit address is observed, it is hashed to the bit vector and the bit value hashed by each m -bit index is set.

5 When a query is to be made whether a given N -bit address has been observed before, the N -bit address is hashed using the same hash functions and the bit values are read from the locations indexed by the m -bit hash values. If at least one of the bit values is 0, this means that this address has *definitely not* been observed before. If all of the bit values are 1, then the address *may have been* observed but this cannot be
10 guaranteed. The latter case is also called a *false hit*.

 As the number of hash functions increases, the Bloom filter bit vector is polluted much faster. On the other hand, the probability of finding a zero during a query increases if more hash functions are used.

 Recently, Bloom filters have been used in the field of computer micro-
15 architecture. Sethumadhvan et al in the article "Scalable Hardware Memory Disambiguation for High ILP Processors", Proceedings of the 36th International Symposium for Microarchitecture pp.399-410, 2003, uses Bloom Filters for memory disambiguation to improve the scalability for load store queues. Roth in the article
20 "Store Vulnerability Window (SVW): Re-Execution Filtering for Enhanced Load Optimization", Proceedings of the 32nd International Symposium on Computer Architecture (ISCA-05), June 2005, uses a Bloom filter to reduce the number of load re-executions for load/store queue optimizations. Akkary et al in the article
 "Checkpoint Processing and Recovery: Towards Scalable Large Instruction Window Processors", Proceedings of the 36th International Symposium for Microarchitecture,
25 Dec., 2003, also uses a Bloom filter to detect the load-store conflicts in the store queue. Moshovos et al in the article "JETTY: Snoop filtering for reduced power in SMP servers", Proceedings of International Symposium on High Performance Computer Architecture (HPCA-7), Jan 2001, uses a Bloom filter to filter out cache coherence requests or snoops in SMP systems.

30 In the embodiments described herein the Bloom filter technique is employed for the purpose of detecting cache misses. In one particular embodiment, a counting Bloom filter technique is used. The counting Bloom Filter has one counter

corresponding to each bit in the bit vector of a normal Bloom Filter. The counting Bloom Filter can be arranged to track addresses entering and leaving a cache. Whenever there is a line fill in the cache, the counters corresponding to the hashes for the address are incremented. In the case of a cache replacement, the corresponding
5 counters are decremented. If the counter is zero, the address has never been seen before. If the counter is greater than one, the address may have been encountered.

In one particular embodiment, a segmented counting Bloom filter design is used. In accordance with this novel segmented approach, the cache miss indication logic has counter logic and separate vector logic as illustrated schematically in Figure 5. The
10 counter logic 510 has a plurality of counter entries, with each counter entry containing a count value, and the vector logic 520 has a vector entry for each counter entry in the counter logic 510, each vector entry containing a value (preferably a single bit) which is set when the count value in the corresponding counter entry changes from a zero value to a non-zero value, and which is cleared when the count value in the
15 corresponding counter entry changes from a non-zero value to a zero value.

Whenever a data value is stored in, or removed from, the cache, a portion of the associated address is provided to the hash function 500 which generates an m-bit index identifying a particular counter in the counter logic 510. This counter value is then incremented if a data value is being stored in the cache, or is decremented if a data
20 value is being removed from the cache. For certain access requests, the cache miss indication logic can then be accessed by issuing the relevant address portion to the hash function 530 to cause an index to be generated identifying a particular vector entry in the vector logic 520. If the value in that vector entry is not set, then this indicates that the data value is not present in the cache, and accordingly a cache miss
25 indication can be generated indicating that the data value is not stored in the cache. If in contrast the vector entry is set, then the data value may or may not be in the cache.

Multiple hash functions can be provided so as to cause multiple counter entries or multiple vector entries to be accessed dependent on each address. However, in one embodiment, only a single hash function is used to form the index generation logic
30 used to generate the index into the either the counter logic or the vector logic, but this index generation logic is replicated separately for the counter logic and the vector logic, thereby facilitating placement of the vector logic in a part of the apparatus

separate to the counter logic. It has been found that having one hash function produces false hit rates (i.e. indications which are inconclusive) similar to the rates achieved when having more than one hash function. The number of bits “L” provided for each counter in the counter logic depends on the hash functions chosen. In the worst case, 5 if all cache lines map to the same counter, the bit-width of the counter must be at most $\log_2(\# \text{ of cache lines})$. One form of hash function that can be used in one embodiment involves bitwise XORing of adjacent bits. A schematic of this hash function which converts a 32 bit address to a m bit hash is shown in Figure 8.

As shown in Figure 8, each pair of m bits in the address is passed through an 10 XOR function 610, 620, 630 so as to ultimately reduce the address 600 to a single m -bit index 640.

Another simplistic hash function that may be chosen uses the lower $\log_2(N)$ bits of the block address (N is the size of the bloom filter). It can be proven with this hash function that the number of bits per counter is equal to $\log_2(\text{Associativity of cache})$. 15 Other simple hash functions could be chosen, such as $\text{Hash} = \text{Addr} \% \text{Prime}$, where Prime is the greatest prime number less than N .

As mentioned earlier, the segmented design provides separate counter logic 510 and bit vector logic 520. There are several benefits realised by such a segmented design. Firstly, to know the outcome of a query to the Bloom filter, only the bit vector 20 logic 520 needs to be accessed, and its size is smaller than the counter logic 510. Keeping the bit vectors separate hence enables faster and lower power accesses to the Bloom Filter. In addition, updates to the counter logic 510 are much more frequent than updates to the bit vector logic 520. Thus segmenting the counter logic and the bit vector logic allows the counter logic to be run at a much lower frequency than the bit 25 vector logic, which is acceptable since the operation of the counter logic is not time critical. In particular, the bit vector logic only needs updating if a particular counter entry in the counter logic changes from a non-zero value to a zero value, or from a zero value to a non-zero value.

Hence, the use of such a segmented design to form the cache miss indication 30 logic 100 used in the embodiments of Figures 1 to 3 leads to a particularly efficient approach for generating an early cache miss indication for use by the data processing apparatus embodiments illustrated in those figures.

Whilst the embodiments of Figures 1 to 3 all involve the use of the indication produced by the cache miss indication logic to control a process having an effect on a program thread other than a program thread whose access request caused the indication to be generated, it has also been found that the particular segmented counting Bloom filter technique described above can provide particularly beneficial results when employed within a data processing apparatus irrespective of whether multiple program threads are being executed or not. Figure 4 illustrates a general data processing apparatus in which the segmented counting Bloom filter 420 of one embodiment can be used. In this example, a CPU 400 is shown coupled to external memory 440 via two levels of cache, namely the level 1 (L1) cache 410, which may be a unified cache or may provide separate instruction and data caches, and the level 2 (L2) cache 430, which again may provide separate instruction and data caches, but more usually will be formed as a unified cache.

The increasing complexity and frequencies of present day microprocessors has led to energy becoming an important design constraint. Techniques to save energy may be applied at the circuit level or at the architecture level. The embodiment illustrated schematically with reference to Figure 4 describes an architecture level technique to save energy in the core and the caches of a microprocessor by early detection of a miss in the L2 Cache.

In one embodiment, energy in an in-order microprocessor can be saved without loss of performance. In an in-order processor, severe stalls in the core 400 may occur due to an L2 Cache miss on a load operation. This is because after a data access misses the L2 cache 430, it accesses the DRAM memory 440. Since a DRAM memory is much slower than caches, a memory access may take more than 100 cycles depending on the processor frequency before the data returns. A good way to save energy here would be to use the event of an L2 Cache miss as a trigger to perform an energy management process on the core. However such a trigger may only be obtained after it is known that the access would miss the L2 Cache. The time taken to access a L2 Cache depends on the size and whether the cache is on chip or off chip. A typical estimate of an L2 cache access time is more than 10 cycles.

However, in accordance with the embodiment used in Figure 4, the segmented counting Bloom filter 420 can be used to produce an early indication of the level 2

cache miss by snooping all accesses to the level 2 cache 430 from the level 1 cache 410. As discussed earlier, the Bloom filter 420 can provide an indication which will either identify that the data value associated with the address is definitely not present in the level 2 cache 430, or instead may identify that the data value the subject of the address may be in the level 2 cache. As shown schematically by the dotted lines in Figure 4, if the segmented counting Bloom filter 420 can determine that the data value is definitely not present in the level 2 cache, then energy can be saved by doing one or more of the following:

- i) performing an energy management process on the processor core 400;
- ii) turning the level 1 cache 410 and the level 2 cache 430 drowsy, i.e. putting them into a low power state preserving mode; and/or
- iii) preventing the level 2 cache 430 access at an early stage.

By early detection of the cache miss it is also possible to reduce the bus activity between the level 1 and level 2 caches and therefore reduce the bus power.

There will be some overhead incurred as a result of turning on and turning off the various logic units. However, this overhead would typically not be much of a concern because the turn-off period happens during a memory access that may take hundreds of cycles. Also, since it is known exactly when the data would return from memory, the structures which have been turned off may be turned on slowly in stages to save energy.

Figure 6 is a block diagram illustrating in more detail one particular embodiment of the apparatus of Figure 4. As can be seen from Figure 6, the CPU 400 can issue an access request to the level 1 cache, which in Figure 6 is shown as consisting of a separate instruction cache 500 and data cache 510. The bit vector segment 520 of the segmented counting Bloom filter 420 is located between the level 1 cache and the level 2 cache 430, and in the event of a miss in the level 1 cache, the bit vector segment is arranged to snoop the address output from the level 1 cache to the level 2 cache, and can signal in a single cycle if the level 2 cache does not have the cache line indicated by that address. This cache miss indication can then be routed via the OR gate function 540 to the CPU 400 to cause the CPU 400 to implement an energy management process, and similarly this cache miss indication can be routed via

the OR gate function 540 to the level 1 cache 500, 510 and the level 2 cache 430 to implement an energy management process, for example putting the caches into a drowsy mode of operation. The access to the level 2 cache can also be aborted.

As also shown in Figure 6, the counter segment of the segmented counting Bloom filter 420 can be considered as existing at the same level of the memory hierarchy as the level 2 cache 430, and is updated based on the linefill and eviction activities occurring between the level 2 cache 430 and the memory 440. Whenever a particular entry in the counter segment 530 changes from a non-zero value to a zero value, or from a zero value to a non-zero value, the relevant vector entry in the bit vector segment 520 is updated.

As also shown in Figure 6, an actual miss in the level 2 cache 430 can result in a signal being routed via the OR gate logic 540 to implement an energy management process in the CPU 400, and a corresponding energy management process in the caches.

Figure 7 illustrates an alternative design where the bit vector segment 520 resides within the CPU 400 itself, this design being useful for embodiments where the caches exhibit inclusive behaviour. Since the bit vector segment 520 is relatively small, it can be retained within a register or set of registers within the CPU 400. In this embodiment, the bit vector segment can monitor all accesses issued to the level 1 cache. If it is determined that such accesses will miss in the level 2 cache then it is also clear that the access will miss in the level 1 cache, if the inclusion property holds. Thus, in such embodiments, the access to both the level 1 cache and the level 2 cache can be aborted, providing further energy savings. The design of figure 7 can also be used for embodiments where the caches do not exhibit inclusive behaviour, since the cache miss indication logic will be used prior to the level 1 cache access, and in the event of a miss in the level 1 cache there is no need to propagate the request to the level 2 cache, thereby achieving power and performance benefits.

Hence, as is clear from Figures 6 and 7, by adopting the segmented approach, the bit vector segment 520 can reside between the level 1 and the level 2 caches, or even inside the core. Since the bit vector segment is much smaller than the counter segment 530, it provides a fast access time, whereas the larger counter segment 530 can be run at a lower frequency to save energy (which is acceptable since the operation

of the counter segment is not time critical). The only additional overhead of this segmented approach is the duplication of the hash function hardware used to access the counter segment 530 and the bit vector segment 520. However, as discussed earlier, in one embodiment a single hash function can be used in that hardware, which simplifies
5 the implementation and reduces the impact of the duplication of the hardware.

For some implementations of cache, there may be a plurality of potential locations where the requested data value the subject of an access request could reside, and not all of these potential locations may be searchable via a single lookup in the cache. In such instances, when performing the lookup procedure the cache may be
10 arranged to perform a sequence of lookups with each lookup relating to a subset of the potential locations. One particular example of a cache where this may occur is the so-called virtual cache, where the cache is accessed using a virtual address. Figures 9 and 10 illustrate how the segmented counting Bloom filter approach of one embodiment can be used in association with two such virtual caches, in particular Figure 9 showing
15 a virtually indexed, physically tagged cache, and Figure 10 illustrating a virtually indexed virtually tagged cache.

An advantage of using a virtual cache is that the cache access time can be improved, since it is unnecessary to translate the virtual address to a physical address using table lookaside buffer (TLB) logic 730 to translate the virtual address to a
20 physical address before the lookup takes place. However, one identified problem with virtual caches is the so-called synonym problem, which results from the possibility that multiple virtual addresses can map to the same physical address. The synonym problem occurs when the cache index bits use some of the bits from the virtual page number or superset bits for which the TLB address translation has not yet been
25 performed. In such cases, an address tag can be located in multiple cache lines, so the tag lookup may need to be performed in multiple cache lines for direct-mapped caches and in multiple cache sets for set-associative caches. The superset bits determine the number of synonym sets.

One of the basic approaches for handling synonyms in virtual caches is to keep
30 only a single copy in the cache. When there is a cache miss, the access request is sent to the lower level memory and all cache lines in every set in the superset are looked up for a potential synonym. For virtually-indexed physically-tagged caches, all physical

tags in the superset are looked up for a synonym, and the access request is aborted and the cache line is moved to the indexed set if a synonym is found. The situation gets worse in virtually-indexed virtually-tagged caches, where every virtual tag must be translated into a physical tag and then each must be compared with the tag of the missed cache address in order to detect a synonym. When a synonym is found, the memory request is aborted and the cache line is remapped or retagged and the line can also be moved to the indexed set.

Thus, detecting synonyms is very slow and consumes power due to a potentially large number of tag lookups. The lookup speed will drop and the power will increase as the cache size and associativity gets larger. Given this problem, the segmented Bloom filter approach of embodiments of the present invention can be used to filter out several tag lookups and accordingly save energy. In particular, if the segmented Bloom filter can identify that there will be a miss in the cache, then the lookups can be aborted at an early stage. Since the segmented Bloom filter is only looking at physical addresses, it can quickly detect synonyms and save a lot of tag lookup power. In particular, it can determine from only checking a single bit location in the bit vector logic if a synonym is not present, since if that lookup indicates that the data value is not in the cache, then there will also be no synonym in the cache. In particular, use of the segmented Bloom filter approach has the possibility of eliminating ($w \times n$) tag lookup comparisons, where w is the cache associativity and n is the number of supersets, in situations where the segmented Bloom filter detects that there is no synonym.

As shown in Figure 9, for virtually-indexed physically-tagged caches the bit vector logic 720 of the segmented Bloom filter keeps track of physical addresses translated by the TLB 730 rather than tracking virtual addresses. During a linefill, the counter logic 710 is accessed by the address of the linefill, which is the physical address sent by the lower level memory. Similarly, during a cache line eviction, the counter logic 720 is accessed by the physical index of the linefill and the physical tag of the evicted line since the tags in the cache are physical. When a cache miss occurs, the bit vector logic 720 is accessed by the physical address of the line, which is obtained from the TLB 730. If the output from the vector logic 720 indicates that there will be a cache miss, this means that there is no synonym in the cache 700 associated

with the physical address. Accordingly, a significant number of lookups can be avoided in such situations.

As shown in Figure 10, for virtually-indexed virtually-tagged caches, the counter logic 710 is accessed by the physical address of the linefill during the linefill process. However, for evictions the evicted address must be translated into the physical address by the TLB 730 before the counter logic 710 is accessed. The dirty evicted lines are already translated by the TLB before sending to the physical memory. Hence, the only extra TLB translations required are in association with clean evicted lines. When a cache miss occurs, the bit vector logic 720 is accessed by the physical address acquired from the TLB logic 730 in a similar way to that described earlier for the virtually-indexed physically tagged caches, and if the bit vector logic 720 indicates that there is a cache miss in the cache 750, then this means that there is no synonym in the cache associated with the physical address, and accordingly a significant number of lookups can be avoided.

Figure 11 illustrates another application of the segmented Bloom filter design in the context of a CMP system in which the level 2 cache is distributed amongst the various processor cores. This means that each processor core 800, 810 has a particular level 2 cache portion 820, 830, respectively, associated with it, but can also access the level 2 cache portions associated with other cores. Hence, by way of example, the core 800 has the level 2 cache portion 820 associated with it, but can also access the level 2 cache portion 830 associated with the core 810. Hence, the core 800 can allocate data within the level 2 cache portion 830 without replicating that data within its own level two cache portion 820. An analogous situation exists for the other core 810. In accordance with this model, victims from one level 2 cache can also be placed into other level 2 caches to increase the effective level 2 cache size.

Each level 2 cache portion has associated counter logic and vector logic forming a segmented Bloom filter design. Hence, the level 2 cache portion 820 has associated therewith the counter logic 840 and the bit vector logic 845 and similarly the level 2 cache portion 830 has associated therewith the counter logic 850 and the bit vector logic 855. However, as shown in Figure 11, each core has access to the bit vector segments associated with level 2 cache portions other than the level 2 cache portion associated with that core. Hence, as shown, the core 800 has access to the bit

vector logic 855 associated with the level 2 cache 830. Similarly, the core 810 has access to the bit vector logic 845 associated with the level 2 cache 820. The counter segments 840, 850 are updated in the usual manner, and are caused to update their associated bit vector segments as required when individual counter entries change
5 from a zero to a non-zero value or from a non-zero value to a zero value. A core can then detect misses in the level 2 cache portions other than its associated level 2 cache portion early by performing an appropriate lookup in the relevant bit vector logic.

Since each core can access the various level 2 caches other than its own associated level 2 cache portion via the bus arbiter 860, the traffic through the bus
10 arbiter 860 can potentially be very high. However, the bus congestion, arbiter queue size and the number of level 2 ports can be reduced by this cross-firing design illustrated in Figure 11, since if the output from the bit vector logic indicates that a cache miss will occur, the cache access can be aborted early. Hence, as a particular example, if the core 800 wishes to issue an access request to the level 2 cache 830 via
15 the bus arbiter 860, but the bit vector logic 855 issues an output signal indicating that a cache miss will occur, then that cache access can be aborted thereby reducing bus congestion.

Whilst in Figure 11 only two cores are shown, it will be appreciated that the technique described therein can be generalised for an n-core model as well.

20 In an alternative embodiment illustrated in Figure 12, each core can also have the bit vector segment associated with its own cache portion to further decrease access to the level 2 cache. Hence, core 800 can be provided with bit vector logic 870 associated with the level 2 cache 820, and core 810 can be provided with bit vector logic 880 associated with the level 2 cache 830. In this design, it should be noted that
25 only the bit vector logic is replicated, and in contrast the counter logic is shared between the cross-firing and the local bit vector logic units. The various bit vector logic units accessible via a particular core can be accessed simultaneously using the same index generation logic 890, 895. Hence, with a single access using the index generation logic, a particular core can identify which level 2 portions, including its
30 own, will result in a cache miss, thereby enabling an early abortion of those cache lookups.

It should be noted that the demand for level 2 bandwidth will increase if each core also has an n-way SMT design. Hence, this will further complicate the design of the bus arbiter 860 and the level 2 caches with multiple ports, thereby increasing their cost. This cost can be mitigated by using the cross-firing design of segmented Bloom filter illustrated schematically in Figures 11 and 12.

From the above description of various embodiments, it will be appreciated that the cache miss indication logic designs described herein enable a number of power saving and performance improvements to be realised. The segmented Bloom filter design enables lookups to the vector logic to be fast and consume low power. Further, the flexibility to place the vector logic at locations within the data processing apparatus separate to the associated counter logic produces further flexibility which can result in further efficiencies.

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 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:
processing logic operable to execute a plurality of program threads;
5 a cache operable to store data values for access by the processing logic, the processing logic being operable when access to a data value is required whilst executing a first program thread to issue an access request specifying an address in memory associated with that data value, and the cache being operable in response to the address to perform a lookup procedure to determine whether the data value is stored in the cache;
10 indication logic operable in response to an address portion of the address to provide an indication as to whether the data value is stored in the cache, said indication being produced before a result of the lookup procedure is available, and the indication logic being operable to only issue an indication that the data value is not stored in the cache if that indication is guaranteed to be correct; and
15 control logic operable, if the indication indicates that the data value is not stored in the cache, to use said indication to control a process having an effect on a program thread other than the first program thread.
2. A data processing apparatus as claimed in Claim 1, further comprising additional
20 control logic operable in response to said indication to control a process having an effect on the first program thread.
3. A data processing apparatus as claimed in Claim 1 or Claim 2, wherein:
the control logic is operable, if the indication indicates that the data value is not
25 stored in the cache, to cause the processing logic to perform a context switch operation in order to switch execution between the first program thread and another program thread.
4. A data processing apparatus as claimed in Claim 1 or Claim 2, wherein:
the processing logic comprises a number of shared resources, each shared
30 resource being allocatable to one of the program threads; and

the control logic is operable, if the indication indicates that the data value is not stored in the cache, to cause the processing logic to perform a reallocation of at least one shared resource.

- 5 5. A data processing apparatus as claimed in Claim 1 or Claim 2, wherein:
the processing logic comprises a plurality of processors each sharing access to the cache, and each being operable when access to a data value is required to issue an access request specifying an address in memory associated with that data value;
the plurality of program threads are distributed across the plurality of processors;
10 the indication logic is operable in response to at least some of the access requests to provide said indication; and
the control logic comprises arbitration logic operable to perform an arbitration process to arbitrate access to the cache between the plurality of processors, and being arranged to use the indications provided by the indication logic when performing said
15 arbitration process.
6. A data processing apparatus as claimed in Claim 5, wherein the indication logic comprises a plurality of indication units, each indication unit being associated with one of said processors and being operable in response to the address portion of an access request
20 issued by the associated processor to provide said indication as to whether the data value is stored in the cache;
the control logic being operable to use the indications provided by each of the indication units when performing said arbitration process.
- 25 7. A data processing apparatus as claimed in any preceding claim, wherein the indication logic comprises:
counter logic having a plurality of counter entries, each counter entry containing a count value;
vector logic having a vector entry for each counter entry in the counter logic,
30 each vector entry containing a value which is set when the count value in the corresponding counter entry changes from a zero value to a non-zero value, and which is

cleared when the count value in the corresponding counter entry changes from a non-zero value to a zero value;

index generation logic operable to generate from an address portion of an address at least one index, each index identifying a counter entry and associated vector entry;

5 the index generation logic being operable whenever a data value is stored in, or removed from, the cache to generate from the address portion of the associated address said at least one index, and to cause the count value in each identified counter entry to be incremented if the data value is being stored in the cache or decremented if the data value is being removed from the cache;

10 the index generation logic further being operable for at least some access requests to generate from the address portion of the associated address said at least one index, and to cause the vector logic to generate an output signal based on the value in each identified vector entry, the output signal indicating if the data value of the access request is not stored in the cache.

15

8. A data processing apparatus as claimed in Claim 7, wherein:

a memory hierarchy exists within the data processing apparatus having a number of memory levels, the cache being provided at one of said memory levels, and when processing access requests memory levels higher in the hierarchy being accessed before
20 memory levels lower in the hierarchy; and

whilst the vector logic and the counter logic are both associated with the cache, the vector logic is accessible separately to the counter logic and from a different level in the memory hierarchy to that memory level from which the counter logic is accessible.

25

9. A data processing apparatus comprising:

processing logic;

a cache operable to store data values for access by the processing logic, the processing logic being operable when access to a data value is required to issue an access
30 request specifying an address in memory associated with that data value, and the cache being operable in response to the address to perform a lookup procedure to determine whether the data value is stored in the cache;

counter logic having a plurality of counter entries, each counter entry containing a count value;

vector logic having a vector entry for each counter entry in the counter logic, each vector entry containing a value which is set when the count value in the
5 corresponding counter entry changes from a zero value to a non-zero value, and which is cleared when the count value in the corresponding counter entry changes from a non-zero value to a zero value;

index generation logic operable to generate from an address portion of an address
at least one index, each index identifying a counter entry and associated vector entry;

10 the index generation logic being operable whenever a data value is stored in, or removed from, the cache to generate from the address portion of the associated address said at least one index, and to cause the count value in each identified counter entry to be incremented if the data value is being stored in the cache or decremented if the data value is being removed from the cache;

15 the index generation logic further being operable for at least some access requests to generate from the address portion of the associated address said at least one index, and to cause the vector logic to generate an output signal based on the value in each identified vector entry, the output signal indicating if the data value of the access request is not stored in the cache;

20 wherein a memory hierarchy exists within the data processing apparatus having a number of memory levels, the cache being provided at one of said memory levels, and when processing access requests memory levels higher in the hierarchy being accessed before memory levels lower in the hierarchy; and

whilst the vector logic and the counter logic are both associated with the cache,
25 the vector logic is accessible separately to the counter logic and from a different level in the memory hierarchy to that memory level from which the counter logic is accessible.

10. A data processing apparatus as claimed in Claim 9, wherein:

30 the index generation logic comprises first index generation logic associated with the counter logic and second index generation logic associated with the vector logic;

whenever a data value is stored in, or removed from, the cache the first index generation logic being operable to generate from the address portion of the associated address said at least one index identifying one or more counter entries in the counter logic; and

5 for at least some access requests the second index generation logic being operable to generate from the address portion of the associated address said at least one index identifying one or more vector entries in the vector logic.

11. A data processing apparatus as claimed in Claim 9 or Claim 10, further
10 comprising lookup control logic operable in response to said output signal from the vector logic to abort the lookup procedure if the output signal indicates that the data value is not stored in the cache.

12. A data processing apparatus as claimed in any of claims 9 to 11, wherein:
15 the processing logic is operable to execute a sequence of instructions in order; and
 if the output signal from the vector logic indicates that the data value is not stored in the cache, the processing logic is operable to perform an energy management process to reduce energy consumption of the processing logic whilst waiting for the data value to be accessed.

20 13. A data processing apparatus as claimed in Claim 12, further comprising:
 an additional cache, if on issuance of the access request by the processing logic a lookup is performed in the additional cache resulting in a miss in the additional cache, the cache being operable to perform said lookup procedure; and
25 if the output signal from the vector logic indicates that the data value is not stored in the cache, at least one of the cache and additional cache being operable to perform an energy management process to reduce energy consumption whilst waiting for the data value to be accessed.

30 14. A data processing apparatus as claimed in any of claims 9 to 11, further comprising:

an additional cache, if on issuance of the access request by the processing logic a lookup is performed in the additional cache resulting in a miss in the additional cache, the cache being operable to perform said lookup procedure, the cache and additional cache exhibiting inclusive behaviour; and

5 if the output signal from the vector logic indicates that the data value is not stored in the cache the lookup in the additional cache being aborted.

15. A data processing apparatus as claimed in any of claims 9 to 11, wherein:
the vector logic is accessible prior to the access request being processed by the
10 memory level in which the cache is provided.

16. A data processing apparatus as claimed in Claim 15, further comprising:
an additional cache at a higher memory level than said cache, if on issuance of the
access request by the processing logic a lookup is performed in the additional cache
15 resulting in a miss in the additional cache, said cache being operable to perform said
lookup procedure; and
the vector logic being accessed in response to a cache miss signal issued by the
additional cache in the event of a cache miss in the additional cache, thereby causing the
vector logic to generate said output signal.

20

17. A data processing apparatus as claimed in Claim 15, further comprising:
an additional cache at a higher memory level than said cache, if on issuance of the
access request by the processing logic a lookup is performed in the additional cache
resulting in a miss in the additional cache, said cache being operable to perform said
25 lookup procedure, said cache and said additional cache exhibiting inclusive behaviour;
and

the vector logic being located within the processing logic and being accessed for
every access request, thereby causing the vector logic to generate said output signal.

30 18. A data processing apparatus as claimed in any of claims 9 to 17, wherein:

the processing logic comprises a plurality of processors, each processor being operable when access to a data value is required to issue an access request specifying an address in memory associated with that data value;

5 the cache comprises a plurality of cache portions, each cache portion associated with one of said processors, and each cache portion being accessible by said plurality of processors;

separate counter logic, vector logic and index generation logic being provided in association with each cache portion;

10 the vector logic associated with at least one cache portion being accessed for access requests issued by a processor with which the at least one cache portion is not associated.

19. A data processing apparatus as claimed in Claim 18, wherein the vector logic is replicated for each processor.

15

20. A data processing apparatus as claimed in any of claims 9 to 19, wherein:

the cache has a plurality of potential locations where the data value the subject of the access request may reside, when performing the lookup procedure the cache being operable to perform a sequence of lookups, each lookup relating to a subset of said potential locations; and

20

if the output signal from the vector logic indicates that the data value is not in the cache, the sequence of lookups being terminated.

21. A method of operating a data processing apparatus having processing logic for executing a plurality of program threads, and a cache for storing data values for access by the processing logic, the method comprising the steps of:

25

when access to a data value is required whilst the processing logic is executing a first program thread, issuing an access request specifying an address in memory associated with that data value;

30

in response to the address, performing a lookup procedure in the cache to determine whether the data value is stored in the cache;

in response to an address portion of the address, employing indication logic to provide an indication as to whether the data value is stored in the cache, said indication being produced before a result of the lookup procedure is available, and the indication logic only issuing an indication that the data value is not stored in the cache if that
5 indication is guaranteed to be correct; and

if the indication indicates that the data value is not stored in the cache, using said indication to control a process having an effect on a program thread other than the first program thread.

10 22. A method of operating a data processing apparatus, the data processing apparatus comprising processing logic, a cache for storing data values for access by the processing logic, counter logic associated with the cache and having a plurality of counter entries, each counter entry containing a count value, and vector logic associated with the cache and having a vector entry for each counter entry in the counter logic, each vector entry
15 containing a value which is set when the count value in the corresponding counter entry changes from a zero value to a non-zero value, and which is cleared when the count value in the corresponding counter entry changes from a non-zero value to a zero value, the method comprising the steps of:

when access to a data value is required by the processing logic, issuing an access
20 request specifying an address in memory associated with that data value;

in response to the address, performing a lookup procedure in the cache to determine whether the data value is stored in the cache;

whenever a data value is stored in, or removed from, the cache, generating from the address portion of the associated address at least one index, each index identifying a
25 counter entry, and causing the count value in each identified counter entry to be incremented if the data value is being stored in the cache or decremented if the data value is being removed from the cache;

for at least some access requests, generating from the address portion of the associated address at least one index, each index identifying a vector entry, and causing
30 the vector logic to generate an output signal based on the value in each identified vector entry, the output signal indicating if the data value of the access request is not stored in the cache;

wherein a memory hierarchy exists within the data processing apparatus having a number of memory levels, the cache being provided at one of said memory levels, and when processing access requests memory levels higher in the hierarchy being accessed before memory levels lower in the hierarchy; and

5 said method further comprising the step of accessing the vector logic separately to the counter logic and from a different level in the memory hierarchy to that memory level from which the counter logic is accessible.

10

15




Fig. 1

2/12




Fig. 2




Fig. 3

4/12




Fig. 4




Fig. 5

6/12




Fig. 6

7/12




Fig. 7

8/12

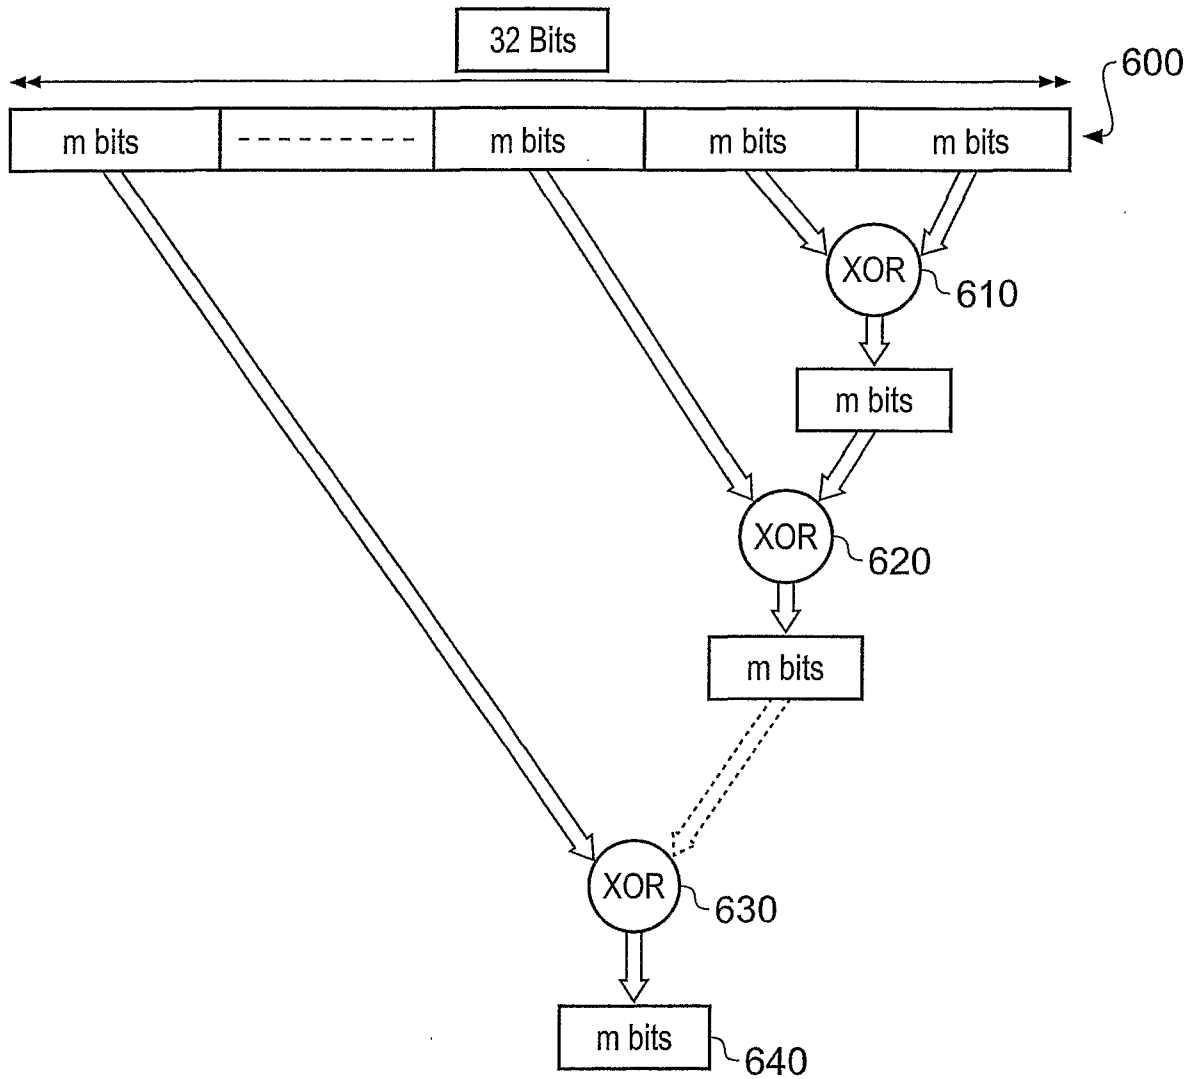


Fig. 8

9/12




Fig. 9

10/12




Fig. 10

11/12

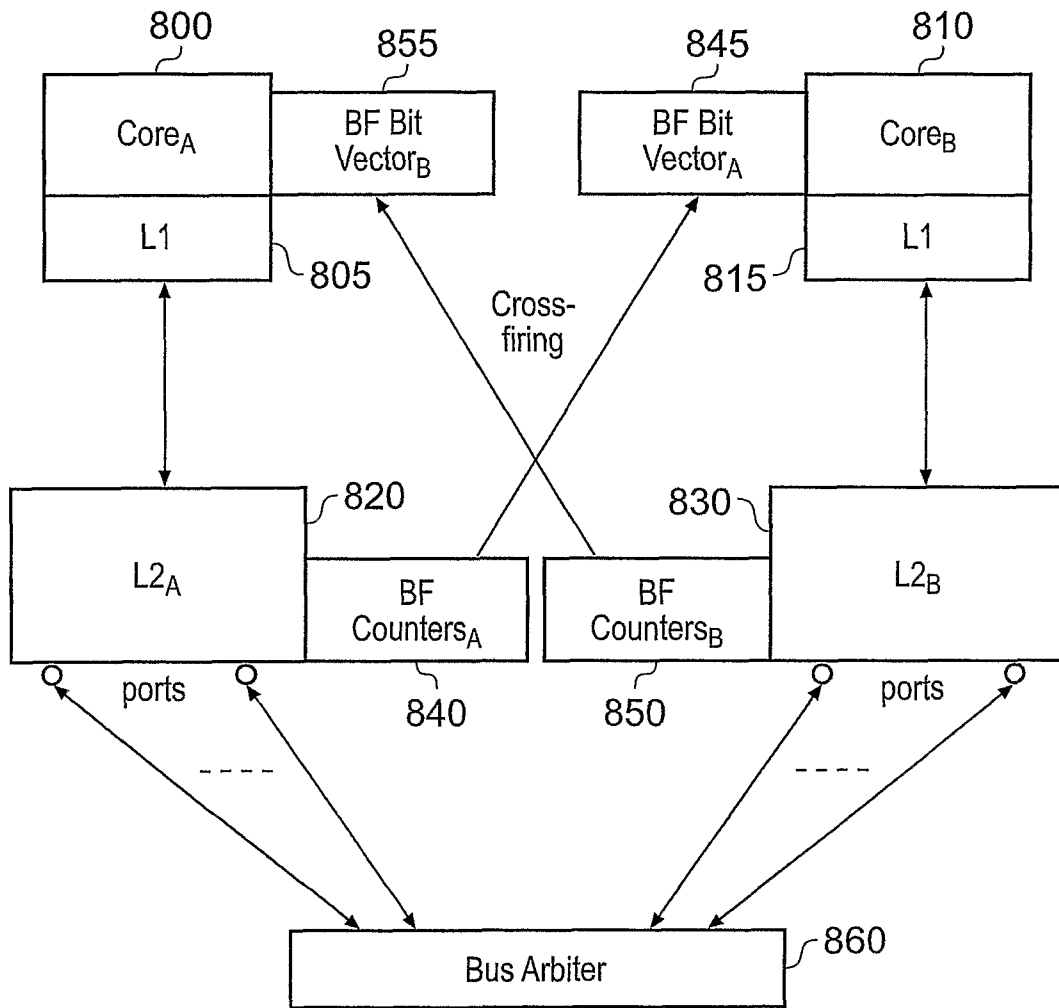


Fig. 11

12/12

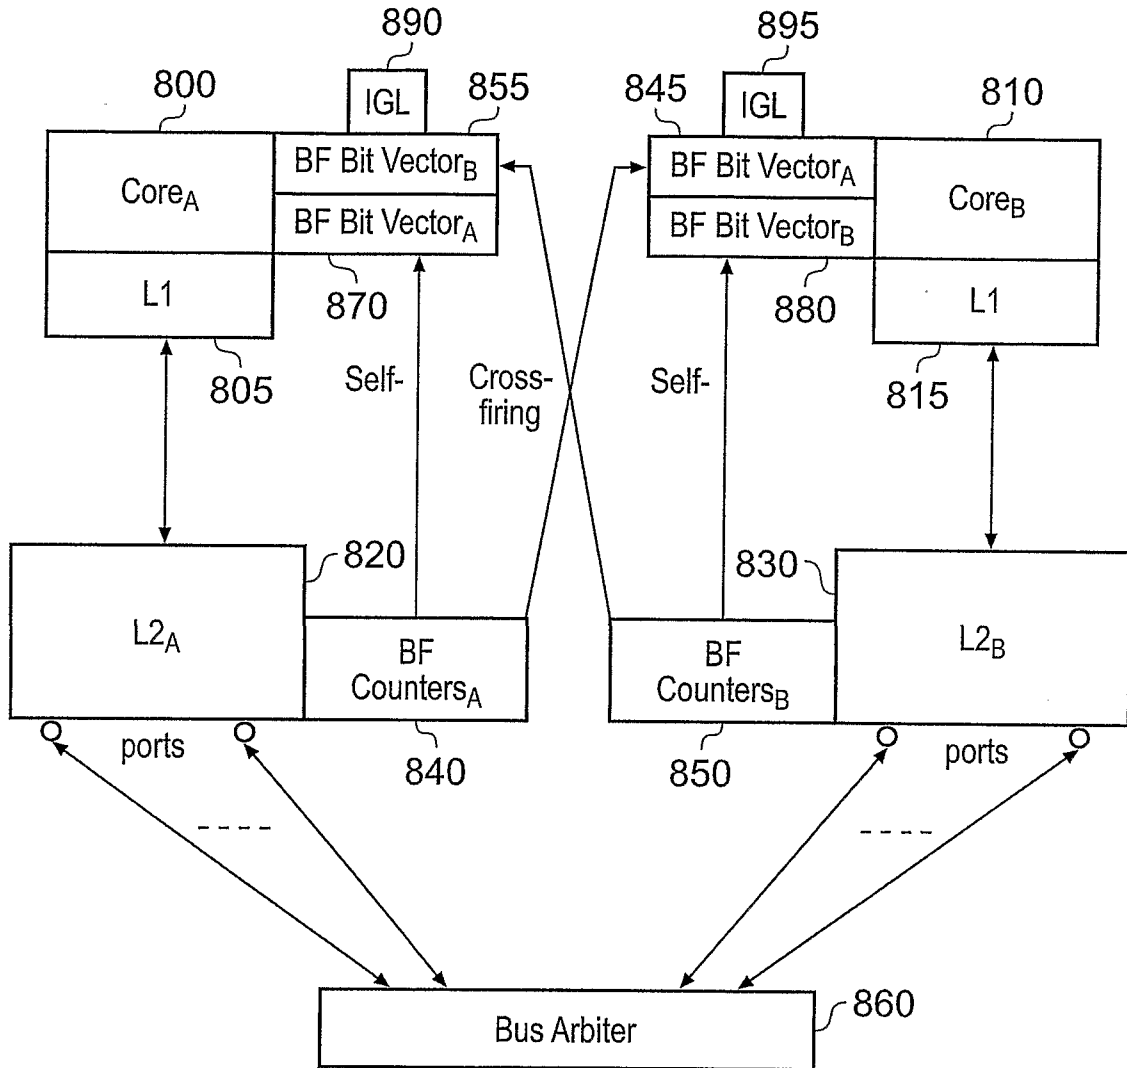


Fig. 12

INTERNATIONAL SEARCH REPORT

International application No
PCT/GB2005/003531

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F12/08 G06F9/38

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	PEIR J-K ET AL: "BLOOM FILTERING CACHE MISSES FOR ACCURATE DATA SPECULATION AND PREFETCHING" CONFERENCE PROCEEDINGS OF THE 2002 INTERNATIONAL CONFERENCE ON SUPERCOMPUTING. ICS'02. NEW YORK, NY, JUNE 22 - 26, 2002, ACM INTERNATIONAL CONFERENCE ON SUPERCOMPUTING, NEW YORK, NY : ACM, US, vol. CONF. 16, 22 June 2002 (2002-06-22), pages 189-198, XP001171515 ISBN: 1-58113-483-5 cited in the application	1-3,21
Y	abstract page 191, right-hand column, line 7 - line 23 ----- -/--	4,7,8

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier 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.
- *G* document member of the same patent family

Date of the actual completion of the international search

22 August 2006

Date of mailing of the international search report

29/08/2006

Name and mailing address of the ISA/
European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Nielsen, O

INTERNATIONAL SEARCH REPORT

 International application No
 PCT/GB2005/003531

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 6 907 520 B2 (PARADY BODO K) 14 June 2005 (2005-06-14) column 6, line 36 - line 50 column 9, line 20 - line 55; figures 1,3,5 -----	1-3,21
Y	US 6 272 520 B1 (SHARANGPANI HARSHVARDHAN ET AL) 7 August 2001 (2001-08-07) column 4, line 56 - column 5, line 5 -----	4
A	US 2005/060457 A1 (OLUKOTUN KUNLE A) 17 March 2005 (2005-03-17) paragraph [0010] - paragraph [0011] -----	5,6
Y	MEMIK G ET AL: "Just say no: benefits of early cache miss determination" HIGH-PERFORMANCE COMPUTER ARCHITECTURE, 2003. HPCA-9 2003. PROCEEDINGS. THE NINTH INTERNATIONAL SYMPOSIUM ON 8-12 FEB. 2003, PISCATAWAY, NJ, USA, IEEE, 8 February 2003 (2003-02-08), pages 307-316, XP010629523 ISBN: 0-7695-1871-0 cited in the application page 307, right-hand column; figures 1,8 page 309, left-hand column, line 32 - right-hand column, line 4 page 311, left-hand column, line 28 - right-hand column, line 31 -----	8-20,22
Y	MOSHOVOS A ET AL: "JETTY: filtering snoops for reduced energy consumption in SMP servers" HIGH-PERFORMANCE COMPUTER ARCHITECTURE, 2001. HPCA. THE SEVENTH INTERNATIONAL SYMPOSIUM ON MONTERREY, MEXICO 19-24 JAN. 2001, LOS ALAMITOS, CA, USA, IEEE COMPUT. SOC, US, 19 January 2001 (2001-01-19), pages 85-96, XP010531677 ISBN: 0-7695-1019-1 cited in the application page 90, left-hand column, line 16 - line 28; figure 3c -----	7-20,22

INTERNATIONAL SEARCH REPORT

International application No.
PCT/GB2005/003531

Box II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This International Search Report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. Claims Nos.:
because they relate to parts of the International Application that do not comply with the prescribed requirements to such an extent that no meaningful International Search can be carried out, specifically:

3. Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

see additional sheet

1. As all required additional search fees were timely paid by the applicant, this International Search Report covers all searchable claims.

2. As all searchable claims could be searched without effort justifying an additional fee, this Authority did not invite payment of any additional fee.

3. As only some of the required additional search fees were timely paid by the applicant, this International Search Report covers only those claims for which fees were paid, specifically claims Nos.:

4. No required additional search fees were timely paid by the applicant. Consequently, this International Search Report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

The additional search fees were accompanied by the applicant's protest.

No protest accompanied the payment of additional search fees.

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

This International Searching Authority found multiple (groups of) inventions in this international application, as follows:

1. claims: 1-6 21

program thread control

2. claims: 7-20 22

cache miss indication circuitry comprising separately
accessible vector logic

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No PCT/GB2005/003531

Patent document cited in search report	Publication date	Patent family member(s)	Publication date	
US 6907520	B2	14-06-2005	US 2002091915 A1	11-07-2002
US 6272520	B1	07-08-2001	NONE	
US 2005060457	A1	17-03-2005	EP 1656615 A2	17-05-2006
			WO 2005020079 A2	03-03-2005