

RECONFIGURABLE CACHE ARCHITECTURE AND METHODS FOR CACHE COHERENCY

CROSS-REFERENCE TO RELATED APPLICATIONS

[001] This application claims the benefit of US Provisional Application No. 62/540,854 filed on August 03, 2017, the contents of which are hereby incorporated by reference.

TECHNICAL FIELD

[002] The disclosure generally relates to memory architectures, and more specifically to embedded computing architectures and configurable computing architectures.

BACKGROUND

[003] In a shared memory multi-core processor with a separate cache memory for each processor, it is possible to have many copies of shared data: one copy in the main memory and one in the local cache of each processor that requested a copy of the data. When one of the data copies is changed, the other copies must reflect that change.

[004] Cache coherence is the uniformity of shared resource data that requires multiple local caches. When clients (e.g., processor cores) in a system maintain local caches of a common memory resource, problems may arise with incoherent data, e.g., the local caches have different values of a single address location.

[005] An example conventional architecture 100 for implementing cache coherence is shown in Fig. 1. Each processor core 110-1 through 110-M (hereinafter referred to individually as a processor core 110 and collectively as processor cores 110 for simplicity purposes) is associated with a corresponding local cache 120-1 through 120-M (hereinafter referred to individually as a local cache 120 and collectively as local caches 120 for simplicity purposes). All core processors 110 and their corresponding local caches 120 access a shared memory 130.

[006] As the memory 130 is shared by the multiple processor cores 110 (and their respective local caches 120), when accessing the shared memory 130, a processor core (e.g., the core 110-1) generally needs to copy a data block from the shared memory 130 to its own cache (e.g., the cache 120-1) in order to accelerate data access. When multiple processor cores 110 access the shared memory 130, a copy of the data block in the shared memory 130 exists in the local caches 120 of all such processor cores 110. To maintain coherence of the copies, a cache coherence mechanism (CCM) 125 is required to manage data sharing.

[007] Specifically, when performing a write (or store) operation on a shared data block or a copy of the shared data block, a write invalidate operation is sent to a processor core 110 that stores a copy of the shared data block, to avoid a data incoherence problem. To maintain cache coherence, the mechanism 125 records a cache status of a data block (or a data block interval). The cache status of the data block (or the data block interval) may include an access type and a sharer of the data block (or the data block interval).

[008] The cache coherence mechanism 125 utilized in conventional architectures operates in a pipeline fashion. As such, a large portion of the processing time is spent on moving data from one area of the memory 130 to the local cache(s) 120, and from one local cache 120 to another. In addition, the conventional architecture of caching as shown Fig. 1 is static by nature and therefore, certain inefficiencies occur as the static pipeline operation does not absolutely fit every use-case.

[009] The limitation of a shared memory resource can also be solved using a reconfigurable cache architecture. Typically, such architectures support dynamic cache partitioning at the hardware level. A reconfigurable cache architecture is typically designed to allow core processors to dynamically allocate cache resource while guaranteeing strict cache isolation among the real-time tasks.

[0010] Reconfigurable cache architectures mainly target for power reduction by using direct addressing mapping. However, such architectures do not improve the latency of memory access.

[0011] Thus, it would be advantageous to provide a processing architecture that overcomes the deficiencies noted above.

SUMMARY

[0012] A summary of several example embodiments of the disclosure follows. This summary is provided for the convenience of the reader to provide a basic understanding of such embodiments and does not wholly define the breadth of the disclosure. This summary is not an extensive overview of all contemplated embodiments, and is intended to neither identify key or critical elements of all embodiments nor to delineate the scope of any or all aspects. Its sole purpose is to present some concepts of one or more embodiments in a simplified form as a prelude to the more detailed description that is presented later. For convenience, the term “some embodiments” may be used herein to refer to a single embodiment or multiple embodiments of the disclosure.

[0013] Some embodiments disclosed herein include a method for cache coherency in a reconfigurable cache architecture. The method comprises receiving a memory access command, wherein the memory access command includes at least an address of a memory to access; determining at least one access parameter based on the memory access command; and determining a target cache bin for serving the memory access command based in part on the at least one access parameter and the address.

[0014] Some embodiments disclosed herein include a reconfigurable cache architecture, comprising: a memory; and a plurality of cache nodes coupled to the memory, wherein each cache node is partitioned to a plurality of cache bins, wherein access to any cache bin of the plurality of cache bins is determined based on an access parameter.

BRIEF DESCRIPTION OF THE DRAWINGS

[0015] The subject matter disclosed herein is particularly pointed out and distinctly claimed in the claims at the conclusion of the specification. The foregoing and other objects, features and advantages of the invention will be

apparent from the following detailed description taken in conjunction with the accompanying drawings.

[0016] Figure 1 is a schematic diagram demonstrating a conventional cache coherence mechanism.

[0017] Figure 2 is a schematic diagram of a reconfigurable system architecture according to an embodiment.

[0018] Figures 3A-3C are schematic diagrams of a reconfigurable cache architecture according to an embodiment.

[0019] Figure 4 is a schematic diagram of I/O peripherals unit coupled to a reconfigurable cache architecture according to an embodiment.

[0020] Figure 5 is a flowchart illustrating a method for cache coherency in a reconfigurable cache according to an embodiment.

DETAILED DESCRIPTION

[0021] In general, statements made in the specification of the present application do not necessarily limit any of the various claimed embodiments. Moreover, some statements may apply to some inventive features but not to others. In general, unless otherwise indicated, singular elements may be in plural and vice versa with no loss of generality. In the drawings, like numerals refer to like parts through several views.

[0022] Fig. 2 illustrates an example schematic diagram of a processing architecture 200 demonstrating the operation of a reconfigurable cache in accordance with one embodiment.

[0023] In an embodiment, the processing architecture 200 includes a processing circuitry 210 coupled to a memory 220 via an interface or bus 240. An input/output (IO) and peripherals unit 230 is also connected to the interface or bus 240 to allow special functions, access to external elements, or both. The I/O and peripherals unit 230 may interface with a peripheral component interconnect (PCI) or PCI Express (PCIe) bus, co-processors, network controllers, and the like (not shown). It should be appreciated that PCIe bus enables connectivity to other peripheral devices.

[0024] The memory 220 is coupled to a plurality of cache nodes 225-1 through 225-n (hereinafter referred to individually as a cache node or collectively as cache nodes for simplicity purposes). Each cache node 225 is configured to store data processed by the processing circuitry 210 and to load data to the processing circuitry 210. Typically, access to the cache nodes 225 is performed through memory access commands, such as store (or write), load (or read). Each cache node 225 may be realized using high-speed static RAM (SRAM), dynamic RAM (DRAM), and the like. In an embodiment, each cache node 225 can be logically partitioned to a plurality of a cache bins (not shown in Fig. 2), as is discussed in detail herein below.

[0025] The processing circuitry 210 may be any processing device or computational device, such as, but not limited to, a central processing unit (CPU), a field-programmable gate array (FPGA), a graphics processing unit (GPU), a coarse-grained reconfigurable architecture (CGRA), an application-specific integrated circuit (ASIC), a quantum computer, and so on. Typically, the processing circuitry 210 is a multi-core processor. It should be noted that the processing architecture 200 can further support a plurality of processing devices 210, e.g., multiple CPUs, hybrid CPUs, and the like.

[0026] In an embodiment, the processing circuitry 210 may be realized as a reconfigurable processing architecture. Such an architecture may be realized as an array of logical elements and multiplexers (MUXs). The logical elements may include arithmetic logic units (ALUs) and functional units (FUs) configured to execute computing functions.

[0027] The processing circuitry 210 is configured to perform various processes to provide a configurable cache architecture which maintains cache coherency among the caches 225-1 through 225-n. As such, the configurable cache architecture is enabled without any additional dedicated hardware. The processing circuitry 210 providing the configurable cache also executes the main programs designed for the processing architecture 200. For example, the processing circuitry 210 may execute a computational machine learning process and run the cache coherency.

[0028] It should be appreciated that, by not using a dedicated hardware, low latency cache access and low power utilization by the processing architecture 200 is ensured. As such, the reconfigurable cache architecture, as disclosed herein, can be utilized to accelerate the operation of the processing circuitry 210 (e.g., a CPU, a FPGA, a GPU, an ASIC, etc.).

[0029] According to the disclosed embodiments, the cache coherency is achieved by determining the location of data in any of the nodes and their cache bins using a deterministic function computed over at least one access parameter. The access parameters are determined by the processing circuitry 210. An access parameter may include, for example, at least one of a unitary identification (ID) representing, a physical entity, and a logical entity. Examples for such entities include, a process ID, a thread ID, a core ID, a cache bit, a source instruction point, a memory port ID, the memory access address, or a combination thereof. The type of the access parameter may be assigned based on the type of memory being accessed. For example, bins of shared memory may be accessed through, for example, at least one cache bit, while bins of local memory can be accessed through at least one process ID. The type of access parameter may be determined during compilation or at runtime.

[0030] In an embodiment, the processing circuitry 210 is configured to receive a memory access command, to determine the access parameter, and to determine the target cache bin based on the access parameter and address designated in the memory access command. As a non-limiting example, a deterministic function, e.g., a hash function, a set of ternary content-addressable memory (TCAM) match rules, a combination thereof, and the like, is computed over the address and the access parameter is called to decide which cache bin of the cache nodes 225 maintains the data.

[0031] For example, a store command may be received at the processing circuitry 210 through the I/O and peripherals unit 230. Such a command may include a data block and a memory address in which to save the data block. The processing circuitry 210 is configured to determine if the command is associated with, for example, a particular process. If so, the process ID of the process is

used as an access parameter. A function computed over the address and process ID (serving as an access parameter) is used to determine the target cache bin for storing the data block. It should be noted that a thread-ID, a core-ID, a cache bit, and so on, can be used as an access parameter. For example, if the received stored command is associated with a particular thread, then a thread-ID will be utilized.

[0032] It should be appreciated that the system architecture 200 described hereinabove depicts a single computational device for the sake of simplicity, and that the architecture 200 can be equally implemented using a plurality of computational devices such as, e.g., CPUs, GPUs, combinations thereof, and so on.

[0033] In an embodiment, the processing circuitry 210 is configured to determine which of the cache nodes 225 should be partitioned, and is further configured to partition each node 225. That is, the processing circuitry 210 is configured to determine how many bins to partition the cache node 225, and the size of each partition. In an embodiment, the partitioning may be static, e.g., to a pre-defined number of bins having equal size. In another embodiment, the partitioning may be dynamic, where the allocation is based on the utilization of each cache bin. To this end, after each execution iteration, the utilization of each bin is measured, and based on the measured utilization, it is determined whether the bins' allocation should be modified. It should be noted that the measurement can be made after program termination or during runtime. For example, the size of popular bins may be increased, while the size of less popular bins is reduced. Further, the number of bins may be increased or decreased based on the measured utilization.

[0034] In certain embodiments, some cache nodes 225 may be statically partitioned, while other may be dynamically partitioned. It should be noted that, initially, the cache may be statistically partitioned, and as the program runs, the allocation of the bins may be dynamically modified.

[0035] In an embodiment, the cache address is divided among the cache bins. Each cache partition of the cache nodes 225 can be assigned a different logical

or physical entity. For example, the cache node 225-1 can be partitioned into two cache bins, with one cache bin dedicated to a first process and the other cache bin dedicated to a second process of a program. Alternatively, the cache bin can be assigned to processor cores of the processing circuitry 210. Other examples of entities that can be allocated cache bins include threads. A partitioning of a cache node to bins is further illustrated in Fig. 3.

[0036] It should be appreciated that this list is only illustrative and not exhaustive of the many types of logical entities and physical entities that can be assigned to cache bins. It should be further appreciated that a cache bin may be any portion of a cache node.

[0037] Figs. 3A–3C illustrate an example schematic diagram of a reconfigurable cache architecture 300 according to an embodiment. In the example illustrated in Figs. 3A-3C, a single cache node 225-n is shown being dynamically partitioned to a number of bins 310.

[0038] Specifically, as shown in Fig. 3A, the cache node 225-n is initially partitioned to 4 cache bins 310-1, 310-2, 310-3 and 310-4 having similar sizes. After a first execution iteration, during runtime or between runs, the partitioning of the node 225-n is changed to include 8 bins 310-1 through 310-8 having similar sizes (Fig. 3B). After another execution iteration, during runtime or between runs, the partitioning of the node 225-n changes to include 8 bins 310-1 through 310-8, but with different sizes. As shown in Fig. 3C, the memory allocated to bin 310-1 is different than bin 310-8.

[0039] According to an embodiment, the cache architecture 300 may be distributed over multiple physical nodes where each node is further divided into one or more logical bins. A processing circuitry of each physical node may access all or part of the cache nodes.

[0040] As shown in Figs. 3A through 3C, a deterministic hash function 225 is utilized to determine a target cache. The function 325 is computed by the processing circuitry 210. It should be appreciated that the reconfigurable cache architecture 300 enables a higher granularity of memory usage, therefore enhancing the system operation and improving runtime performance.

[0041] It should be further appreciated that the reconfigurable cache architecture 300 depicts a single cache node 225-n and a number of 4 or 8 bins 310 merely for the sake of simplicity. The architecture 300 would typically include a plurality of cache nodes that can be partitioned into any number of cache bins.

[0042] In an embodiment, a memory cache bin 310 may perform atomic memory access commands. Such commands may load, conditionally modify, and thereafter store the value of memory at a location, as a single operation. It is to be appreciated that when multiple atomic access commands are executed in parallel from multiple memory ports, and performed sequentially at the cache bin, they provide a coherent view to all memory ports.

[0043] Fig. 4 shows an example schematic diagram of a reconfigurable cache architecture 400 coupled to I/O peripherals (I/O P) 410 according to an embodiment. In this configuration, input/output (IO) and peripheral units 410-1 through 410-k (k is integer greater or equal to 1) may include a PCI bus, a PCI Express (PCIe), one or more co-processors, one or more network controllers, and the like.

[0044] As shown herein, the memory access commands are issued by the I/O peripherals 410. The processing circuitry 210 determines the target cache bin based in part on the received commands using a deterministic hash function 425.

[0045] In this configuration, any data or control signal (e.g., ack signal) received from the target cache bin is mapped to the I/O peripheral 410 that issued the received command. The mapping is performed by a mapping function 427 that can be implemented as a deterministic hash function, as a set of ternary content-addressable memory (TCAM) match rules, a combination thereof, and the like. It should be noted that the memory access is directed to the local caches 120 in order to perform the memory operation.

[0046] Fig. 5 shows an example flowchart 500 of a method for cache coherency in a reconfigurable cache architecture according to an embodiment. The reconfigurable cache architecture includes a plurality of cache nodes coupled to the memory, wherein each cache node is partitioned into a plurality of cache bins.

[0047] At S510, a memory access command is received. As mentioned above, the command may be to store (write) or load (read) data from the memory of a processing architecture. The command may be received via an interface such as, for example, the I/O peripherals unit 230. A received command includes at least a target address to which data is to be stored or from which data is to be loaded. In a store command, the data to be stored is also included in the received command. The memory address should be within the address boundaries determined during compilation of the code of the main program.

[0048] At S520, at least one access parameter is determined. As noted above, an access parameter may include a process ID, a thread ID, a cache bit, a storage pointer, a process core ID, and so on. In an embodiment, the determination includes determining a logical or physical entity that the received command is associated with. Examples for physical entities are discussed in detail above.

[0049] In an embodiment, if the received command is executed as part of a dedicated process or thread (both are considered logical entities), then the process-ID or thread-ID will be considered as the access parameter. In another embodiment, if the received command is executed on a dedicated processing core (considered a physical entity), then the core-ID will be considered as the access parameter. In yet another embodiment, if the received command is to access a shared memory (considered as a physical entity), then a cache bit will be considered as the access parameter.

[0050] In some embodiments, load/store attributes are determined. Such attributes include, for example, never cache certain values, always cache certain values, always check certain values, and so on. Furthermore, ordering of allocation, along with the access synchronization in the grid allows larger pipelines and higher throughput while simplifying mechanisms. Such attributes are advantageous for volatile memory as well as for locking mechanisms.

[0051] At S530, a target cache bin to access is determined. In an embodiment, the determination is performed using a deterministic function computed over the access parameter and the address designated in the received request. According

to another embodiment, the deterministic function is connected to the grid so that the determination is made using the same interfaces.

[0052] It should be noted that data is stored to, or loaded from, the target cache bin as determined by the deterministic function.

[0053] In an embodiment, S530 includes gathering the statistics about the target cache bin being accessed. For example, the number of the bin, the frequency of accessing the same bin, and the size of the data being written or read are determined. These gathered statistics can be utilized to dynamically change the partitions of the bins.

[0054] In S540, it is checked whether additional system calls have been received and if so, execution continues with S510; otherwise, execution terminates.

[0055] The embodiments disclosed herein can be implemented as hardware, firmware, software, or any combination thereof. The application program may be uploaded to, and executed by, a machine comprising any suitable architecture. Preferably, the machine is implemented on a computer platform having hardware such as one or more central processing units ("CPUs"), a memory, and input/output interfaces.

[0056] The computer platform may also include an operating system and microinstruction code. The various processes and functions described herein may be either part of the microinstruction code or part of the application program, or any combination thereof, which may be executed by a CPU, whether or not such computer or processor is explicitly shown.

[0057] In addition, various other peripheral units may be connected to the computer platform such as an additional data storage unit. Furthermore, a non-transitory computer readable medium is any computer readable medium except for a transitory propagating signal.

[0058] All examples and conditional language recited herein are intended for pedagogical purposes to aid the reader in understanding the principles of the disclosure and the concepts contributed by the inventor to furthering the art, and are to be construed as being without limitation to such specifically recited examples and conditions.

WHAT IS CLAIMED IS:

1. A method for cache coherency in a reconfigurable cache architecture (300), comprising:

receiving a memory access command (S510), wherein the memory access command includes at least an address of a memory (220) to access;

determining at least one access parameter (S520) based on the memory access command; and

determining a target cache bin for serving the memory access command (S530) based in part on the at least one access parameter and the address;

wherein the reconfigurable cache architecture is distributed over a plurality of separate physical cache nodes (225-1...225-n), operating substantially independently and electrically coupled to the memory, wherein each cache node is partitionable into a plurality of cache bins (310-1...310-8), wherein the target cache bin is one of the plurality of cache bins, and wherein the partitioning of the plurality of cache bins within at least two of the plurality of cache nodes is reconfigurable; and

wherein the memory access command includes a unitary identification of any one of:

a physical entity, selected from a group consisting of a processing core, and a shared portion of the memory; and

a logical entity, selected from a group consisting of: a process, and a thread;

characterized by:

the at least one access parameter including at least one of: a process ID, a processing core ID, a thread ID, and a cache bit; and

the determination of the at least one access parameter further comprising at least one of:

setting the access parameter as a logical entity identifier, subject to determining that the memory access command is associated with a logical entity; and

setting the access parameter as a physical entity identifier subject to determining that the memory access command is associated with a physical entity.

2. The method of claim 1, wherein the target cache bin is at least a portion of at least one cache node.
3. The method of claim 1, wherein each cache bin is included in a respective cache node of the plurality of cache nodes.
4. The method of claim 1, further comprising:
dynamically partitioning each cache node into at least two cache bins based on utilization of the respective plurality of cache bins of the cache node.
5. The method of claim 4, further comprising:
initially partitioning each cache node into a predetermined number of cache bins (310-1, 310-2, 310-3, 310-4);
collecting statistics with respect to the usage of each cache bin; and
reconfiguring the initial partitioning (310-1...310-8) of each cache node based on the collected statistics.
6. The method of claim 5, wherein the reconfiguration of the partitioning of each cache node is performed after each execution iteration of a plurality of execution iterations.
7. The method of claim 5, further comprising:
dynamically allocating more cache storage to at least one of the cache bins (310-8).
8. The method of claim 1, further comprising:
determining at least one cache attribute, wherein the at least one cache attribute includes at least one of: a never cache certain value, an always cache certain value, and an always check certain value.

9. The method of claim 1, wherein determining the target cache bin for serving the memory access command further comprises:
 computing a deterministic function over the at least one access parameter and the address.
10. The method of claim 1, wherein the reconfigurable cache architecture is utilized to accelerate an execution of a program by a processing circuitry (210).
11. The method of claim 10, wherein the processing circuitry is any one of:
 a central processing unit, a field-programmable gate array, a graphics processing unit, a coarse-grained reconfigurable architecture, an application-specific integrated circuit, multi-core processor, and a quantum computer.
12. A non-transitory computer readable medium having stored thereon instructions which, when executed by at least one processing circuitry, cause the at least one processing circuitry to carry out the method of claim 1.
13. A system for executing a software program, comprising:
 a processing circuitry; and
 a reconfigurable cache connected to the processing circuitry, the reconfigurable cache comprising:
 a memory; and
 a plurality of separate physical cache nodes, operating substantially independently and electrically coupled to the memory, wherein each cache node is partitionable into a plurality of cache bins, wherein access to any cache bin of the plurality of cache bins is determined based on an access parameter, and wherein the partitioning of the plurality of cache bins within at least two of the plurality of cache nodes is reconfigurable;
 wherein the processing circuitry is adapted for cache coherency in a reconfigurable cache architecture, comprising:
 receiving a memory access command, wherein the memory access command includes at least an address of a memory to access;

determining at least one access parameter based on the memory access command; and

determining a target cache bin for serving the memory access command based in part on the at least one access parameter and the address; and

wherein the memory access command includes a unitary identification of any one of:

a physical entity, selected from a group consisting of a processing core, and a shared portion of the memory; and

a logical entity, selected from a group consisting of: a process, and a thread;

characterized by:

the at least one access parameter including at least one of: a process ID, a processing core ID, a thread ID, and a cache bit; and

the determination of the at least one access parameter further comprising at least one of:

setting the access parameter as a logical entity identifier, subject to determining that the memory access command is associated with a logical entity; and

setting the access parameter as a physical entity identifier subject to determining that the memory access command is associated with a physical entity.