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.