



- (51) **International Patent Classification:**
G06F 9/46 (2006.01) G06F 12/08 (2006.01)
G06F 9/30 (2006.01)
- (21) **International Application Number:**
PCT/US2012/030383
- (22) **International Filing Date:**
23 March 2012 (23.03.2012)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
61/467,939 25 March 2011 (25.03.2011) US
- (71) **Applicant (for all designated States except US):** **SOFT MACHINES, INC.** [US/US]; 3211 Scott Boulevard, Suite 202, Santa Clara, CA 95054 (US).
- (72) **Inventor; and**
- (75) **Inventor/Applicant (for US only):** **ABDALLAH, Mo-hammad** [US/US]; 3868 Suncrest Avenue, San Jose, CA 95132 (US).
- (74) **Agent:** **BARNES, Glenn, D.**; Murabito Hao & Barnes LLP, Two North Market Street, Third Floor, San Jose, CA 95113 (US).

- (81) **Designated States (unless otherwise indicated, for every kind of national protection available):** AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) **Designated States (unless otherwise indicated, for every kind of regional protection available):** ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report (Rule 48.2(g))

(54) **Title:** REGISTER FILE SEGMENTS FOR SUPPORTING CODE BLOCK EXECUTION BY USING VIRTUAL CORES INSTANTIATED BY PARTITIONABLE ENGINES

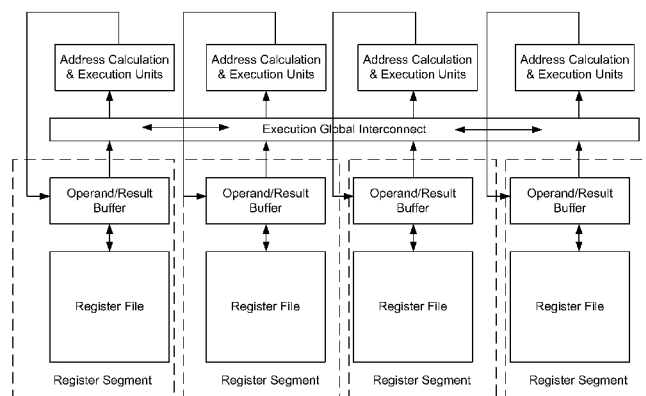


FIG. 6

(57) **Abstract:** A system for executing instructions using a plurality of register file segments for a processor. The system includes a global front end scheduler for receiving an incoming instruction sequence, wherein the global front end scheduler partitions the incoming instruction sequence into a plurality of code blocks of instructions and generates a plurality of inheritance vectors describing interdependencies between instructions of the code blocks. The system further includes a plurality of virtual cores of the processor coupled to receive code blocks allocated by the global front end scheduler, wherein each virtual core comprises a respective subset of resources of a plurality of partitionable engines, wherein the code blocks are executed by using the partitionable engines in accordance with a virtual core mode and in accordance with the respective inheritance vectors. A plurality register file segments are coupled to the partitionable engines for providing data storage.

WO 2012/135041 A2

REGISTER FILE SEGMENTS FOR SUPPORTING CODE BLOCK EXECUTION BY USING VIRTUAL CORES INSTANTIATED BY PARTITIONABLE ENGINES

This application claims the benefit co-pending commonly assigned US Provisional Patent Application serial number 61/467,939, titled "REGISTER FILE SEGMENTS FOR SUPPORTING CODE BLOCK EXECUTION BY USING VIRTUAL CORES INSTANTIATED BY PARTITIONABLE ENGINES" by Mohammad A. Abdallah, filed on March 25, 2011, and which is incorporated herein in its entirety.

CROSS REFERENCE TO RELATED APPLICATION

This application is related to co-pending commonly assigned US Patent Application serial number 2009/0113170, titled "APPARATUS AND METHOD FOR PROCESSING AN INSTRUCTION MATRIX SPECIFYING PARALLEL IN DEPENDENT OPERATIONS" by Mohammad A. Abdallah, filed on April 12, 2007, and which is incorporated herein in its entirety.

This application is related to co-pending commonly assigned US Patent Application serial number 2010/0161948, titled "APPARATUS AND METHOD FOR PROCESSING COMPLEX INSTRUCTION FORMATS IN A MULTITHREADED ARCHITECTURE SUPPORTING VARIOUS CONTEXT SWITCH MODES AND VIRTUALIZATION SCHEMES" by Mohammad A. Abdallah, filed on November 14, 2007, and which is incorporated herein in its entirety.

FIELD OF THE INVENTION

[001] The present invention is generally related to digital computer systems, more particularly, to a system and method for selecting instructions comprising an instruction sequence.

BACKGROUND OF THE INVENTION

[002] Processors are required to handle multiple tasks that are either dependent or totally independent. The internal state of such processors usually consists of registers that

might hold different values at each particular instant of program execution. At each instant of program execution, the internal state image is called the architecture state of the processor.

[003] When code execution is switched to run another function (e.g., another thread, process or program), then the state of the machine/processor has to be saved so that the new function can utilize the internal registers to build its new state. Once the new function is terminated then its state can be discarded and the state of the previous context will be restored and execution resumes. Such a switch process is called a context switch and usually includes 10's or hundreds of cycles especially with modern architectures that employ large number of registers (e.g., 64, 128, 256) and/or out of order execution.

[004] In thread-aware hardware architectures, it is normal for the hardware to support multiple context states for a limited number of hardware-supported threads. In this case, the hardware duplicates all architecture state elements for each supported thread. This eliminates the need for context switch when executing a new thread. However, this still has multiple draw backs, namely the area, power and complexity of duplicating all architecture state elements (i.e., registers) for each additional thread supported in hardware. In addition, if the number of software threads exceeds the number of explicitly supported hardware threads, then the context switch must still be performed.

[005] This becomes common as parallelism is needed on a fine granularity basis requiring a large number of threads. The hardware thread-aware architectures with duplicate context-state hardware storage do not help non-threaded software code and only reduces the number of context switches for software that is threaded. However, those threads are usually constructed for coarse grain parallelism, and result in heavy software overhead for initiating and synchronizing, leaving fine grain parallelism, such as function calls and loops parallel execution, without efficient threading initiations/auto generation. Such described overheads are accompanied with the difficulty of auto parallelization of such codes using state of the art compiler or user parallelization techniques for non-explicitly/easily parallelized/threaded software codes.

SUMMARY OF THE INVENTION

[006] In one embodiment, the present invention is implemented as a system for executing instructions using a plurality of register file segments for a processor. The system includes a global front end scheduler for receiving an incoming instruction sequence, wherein the global front end scheduler partitions the incoming instruction sequence into a plurality of code blocks of instructions and generates a plurality of inheritance vectors describing interdependencies between instructions of the code blocks. The system further includes a plurality of virtual cores of the processor coupled to receive code blocks allocated by the global front end scheduler, wherein each virtual core comprises a respective subset of resources of a plurality of partitionable engines, wherein the code blocks are executed by using the partitionable engines in accordance with a virtual core mode and in accordance with the respective inheritance vectors. A plurality register file segments are coupled to the partitionable engines for providing data storage.

[007] Other embodiments of the present invention utilize a common scheduler, a common register file and a common memory subsystem to implement fragmented address spaces for multiple partitionable engines of processor. The partitionable engines can be used to implement a plurality of virtual cores. Fragmentation enables the scaling of microprocessor performance by allowing additional virtual cores to cooperatively execute instruction sequences. The fragmentation hierarchy can be the same across each cache hierarchy (e.g., L1 cache, L2 cache, and the common register file). The fragmentation hierarchy can divide the address space into fragments using address bits, where the address bits are used such that the fragments are above cache line boundaries and below page boundaries. Each fragment can be configured to utilize a multiport bank structure for storage.

[008] The foregoing is a summary and thus contains, by necessity, simplifications, generalizations and omissions of detail; consequently, those skilled in the art will appreciate that the summary is illustrative only and is not intended to be in any way limiting. Other aspects, inventive features, and advantages of the present invention, as defined solely by the claims, will become apparent in the non-limiting detailed description set forth below.

BRIEF DESCRIPTION OF THE DRAWINGS

[009] The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements.

[010] Figure 1A shows an overview of the manner in which the global front end generates code blocks and inheritance vectors to support the execution of code sequences on their respective partitionable engines.

[011] Figure 1B shows an overview diagram of partitionable engines and their components, including segmented scheduler and register files, global interconnects and a fragmented memory subsystem for a multicore processor in accordance with one embodiment of the present invention.

[012] Figure 2 shows a scheduler flow diagram in accordance with one embodiment of the present invention.

[013] Figure 3 shows a diagram of exemplary hardware circuits that shows a segmented register file storing operands and results with an interconnect in accordance with one embodiment of the present invention.

[014] Figure 4 shows a diagram depicting a global front end Fetch & scheduler in accordance with one embodiment of the present invention.

[015] Figure 5 shows an alternative implementation of the distribution of instructions across many virtual cores in accordance with one embodiment of the present invention.

[016] Figure 6 shows a plurality of register segments with a corresponding plurality of register files and operand & result buffers in accordance with one embodiment of the present invention.

[017] Figure 7 shows a more detailed diagram of a fragmented memory subsystem for a multicore processor in accordance with one embodiment of the present invention.

[018] Figure 8 shows a diagram depicting how bits of an address can be used by address generation to enumerate fragments in accordance with one embodiment of the present invention.

[019] Figure 9 shows a diagram of how loads and stores are handled by embodiments of the present invention.

[020] Figure 10 shows the manner which fragments can be split into two or more domains in accordance with one embodiment of the present invention.

[021] Figure 11 shows a mode of operation of the processor wherein Virtual cores are configured as physical cores that correspond to logical cores in executing applications in accordance with one embodiment of the present invention.

[022] Figure 12 shows a mode of operation of the processor wherein virtual cores are configured as soft cores that correspond to logical cores in executing applications in accordance with one embodiment of the present invention.

[023] Figure 13 shows a mode of operation of the processor wherein the virtual cores are configured as soft cores that correspond to a single logical core in executing applications in accordance with one embodiment of the present invention.

[024] Figure 14 shows an exemplary implementation of fragment segmentation being used to support logical core and virtual core functionality in accordance with one embodiment of the present invention.

[025] Figure 15 shows a fragment memory of an exemplary four fragment processor implementing a many physicals to many logicals mode in accordance with one embodiment of the present invention.

[026] Figure 16 shows a fragment memory of an exemplary four fragment processor implementing a many physicals to many logicals mode in accordance with an alternative embodiment of the present invention.

[027] Figure 17 shows a fragment memory of an exemplary four fragment processor implementing a many soft cores to many logical cores mode in accordance with one embodiment of the present invention.

[028] Figure 18 shows a fragment memory of an exemplary four fragment processor implementing a many soft cores to one logical core mode in accordance with one embodiment of the present invention.

[029] Figure 19 shows address calculation and execution units, operand/result buffers, threaded register files, and common partition schedulers of an exemplary four fragment processor implementing a physicals to logicals mode in accordance with one embodiment of the present invention.

[030] Figure 20 shows an alternative implementation for the address calculation and execution units, operand/result buffers, threaded register files, and common partition schedulers of an exemplary four fragment processor to implement a many physicals to many logicals mode in accordance with one embodiment of the present invention.

[031] Figure 21 shows address calculation and execution units, register files, and common partition schedulers of an exemplary four fragment processor implementing a many soft cores to many logicals mode in accordance with one embodiment of the present invention.

[032] Figure 22 shows address calculation and execution units, register files, and common partition schedulers of an exemplary four fragment processor implementing a many soft cores to one logical core mode in accordance with one embodiment of the present invention.

[033] Figure 23 shows a diagram of an exemplary microprocessor pipeline in accordance with one embodiment of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

[034] Although the present invention has been described in connection with one embodiment, the invention is not intended to be limited to the specific forms set forth herein. On the contrary, it is intended to cover such alternatives, modifications, and equivalents as can be reasonably included within the scope of the invention as defined by the appended claims.

[035] In the following detailed description, numerous specific details such as specific method orders, structures, elements, and connections have been set forth. It is to be understood however that these and other specific details need not be utilized to practice embodiments of the present invention. In other circumstances, well-known structures,

elements, or connections have been omitted, or have not been described in particular detail in order to avoid unnecessarily obscuring this description.

[036] References within the specification to "one embodiment" or "an embodiment" are intended to indicate that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the present invention. The appearance of the phrase "in one embodiment" in various places within the specification are not necessarily all referring to the same embodiment, nor are separate or alternative embodiments mutually exclusive of other embodiments. Moreover, various features are described which may be exhibited by some embodiments and not by others. Similarly, various requirements are described which may be requirements for some embodiments but not other embodiments.

[037] Some portions of the detailed descriptions, which follow, are presented in terms of procedures, steps, logic blocks, processing, and other symbolic representations of operations on data bits within a computer memory. These descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. A procedure, computer executed step, logic block, process, etc., is here, and generally, conceived to be a self-consistent sequence of steps or instructions leading to a desired result. The steps are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals of a computer readable storage medium and are capable of being stored, transferred, combined, compared, and otherwise manipulated in a computer system. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

[038] It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the following discussions, it is appreciated that throughout the present invention, discussions utilizing terms such as "processing" or "accessing" or "writing" or "storing" or "replicating" or the like, refer to the action and processes of a computer system, or similar electronic computing device that manipulates and transforms data represented as physical (electronic)

quantities within the computer system's registers and memories and other computer readable media into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

[039] Embodiments of the present invention utilize a common global front end scheduler, a plurality of segmented register files, and a memory subsystem to implement fragmented address spaces for multiple cores of a multicore processor. In one embodiment, fragmentation enables the scaling of microprocessor performance by allowing additional virtual cores (e.g., soft cores) to cooperatively execute instruction sequences comprising one or more threads. The fragmentation hierarchy is the same across each cache hierarchy (e.g., L1 cache, L2 cache, and the common register file). The fragmentation hierarchy divides the address space into fragments using address bits, where the address bits are used such that the fragments are identified by bits that are above cache line boundaries and below page boundaries. Each fragment is configured to utilize a multiport bank structure for storage. Embodiments of the present invention are further described in the Figures 1A and 1B below.

[040] Figure 1A shows an overview diagram of a processor in accordance with one embodiment of the present invention. As depicted in Figure 1A, the processor includes a global front end fetch and scheduler 10 and a plurality of partitionable engines 11-14.

[041] Figure 1A shows an overview of the manner in which the global front end generates code blocks and inheritance vectors to support the execution of code sequences on their respective partitionable engines. Each of the code sequences 20-23 can belong to the same logical core/thread or to different logical cores/threads, depending upon the particular virtual core execution mode. The global front end fetch and scheduler will process the code sequences 20-23 to generate code blocks and inheritance vectors. These code blocks and inheritance vectors are allocated to the particular partitionable engines 11-14 as shown.

[042] The partitionable engines implement virtual cores, in accordance with a selected mode. A partitionable engine includes a segment, a fragment and a number of execution units. The resources within the partitionable engines can be used to implement virtual cores that have multiple modes. As provisioned by the virtual core mode, one soft core, or many soft cores, can be implemented to support one logical core/thread. In the Figure 1A embodiment, depending on the selected mode, the virtual cores can support one logical core/thread or four logical cores/threads. In an embodiment where the virtual cores

support four logical cores/threads, the resources of each virtual core are spread across each of the partitionable engines. In an embodiment where the virtual cores support one logical core/thread, the resources of all the engines are dedicated to that core/thread. The engines are partitioned such that each engine provides a subset of the resources that comprise each virtual core. In other words, a virtual core will comprise a subset of the resources of each of the engines 11-14. Communication between the resources of each of the engines 11-14 is provided by a global interconnection structure 30 in order to facilitate this process. Alternatively, the engines 11-14 can be used to implement a physical mode where the resources of the engines 11-14 are dedicated to support the execution of a dedicated core/thread. In this manner, the soft cores implemented by the engines comprise virtual cores that have resources spread across each of the engines. The virtual core execution modes are further described in the following figures below.

[043] It should be noted that in a conventional core implementation, only resources within one core/engine is solely allocated to one logical thread/core. In contrast, in embodiments of the present invention, the resources of any engine/core can be partitioned to, collectively with other engine/core partitions, instantiate a virtual core that is allocated to one logical thread/core. Additionally, embodiments of the present invention can implement multiple virtual execution modes in which those same engines can be partitioned to support many dedicated cores/threads, many dynamically allocated cores/threads, or an embodiment where all of the resources of all engines support the execution of a single core/thread. These embodiments are further described in the descriptions below.

[044] Figure 1B shows an overview diagram of partitionable engines and their components, including segmented scheduler and register files, global interconnects and a fragmented memory subsystem for a multicore processor in accordance with one embodiment of the present invention. As depicted in Figure 1, four fragments 101-104 are shown. The fragmentation hierarchy is the same across each cache hierarchy (e.g., L1 cache, L2 cache, and the load store buffer). Data can be exchanged between each of the L1 caches, each of the L2 caches and each of the load store buffers via the memory global interconnect 110a.

[045] The memory global interconnect comprises a routing matrix that allows a plurality of cores (e.g., the address calculation and execution units 121-124) to access data that may be stored at any point in the fragmented cache hierarchy (e.g., L1 cache, load store

buffer and L2 cache). Figure 1 also depicts the manner whereby each of the fragments 101-104 can be accessed by address calculation and execution units 121-124 via the memory global interconnect 110a.

[046] The execution global interconnect 110b similarly comprises a routing matrix allows the plurality of cores (e.g., the address calculation and execution units 121-124) to access data that may be stored at any of the segmented register files. Thus, the cores have access to data stored in any of the fragments and 2 data stored in any of the segments via the memory global interconnect 110a or the execution global interconnect 110b. Additionally, it should be noted that in one embodiment, another global interconnect exists between each of the common partition fetch and schedulers. This is shown by the horizontal arrows between and connecting each common partition fetch and scheduler.

[047] Figure 1B further shows a global front end Fetch & scheduler 150 which has a view of the entire machine and which manages the utilization of the register files segments and the fragmented memory subsystem. Address generation comprises the basis for fragment definition. The global front end Fetch & scheduler functions by allocating instruction sequences to each segment's partition scheduler. The common partition scheduler then dispatches those instruction sequences for execution on the address calculation and execution units 121-124.

[048] It should be noted that in one embodiment, the functionality of the common partition fetch and schedulers can be incorporated into the global front end scheduler 150. In such an embodiment, the segments would not include respective common partition fetch and schedulers, and there would be no need for an interconnect between them.

[049] Additionally, it should be noted that the partitionable engines shown in Figure 1A can be nested in a hierarchal way. In such an embodiment, a first level partitionable engine would include a local front end fetch and scheduler and multiple secondary partitionable engines connected to it.

[050] Figure 2 shows a scheduler flow diagram in accordance with one embodiment of the present invention. As depicted in Figure 2, a bucket buffer is shown that includes Speculative Thread Bucket-pointers, Bucket Sources and destinations lists. The scheduler and execution buckets include a bucket dispatch selector and the virtual register match and

read, including the possibility of a register hierarchy and a register cache. The Back end is where executed buckets are logged and exception ordering is enforced before retirement. The register hierarchy/cache also serves as an intermediate storage for the executed bucket results until they are non-speculative and can update the architecture state. The following discloses one possible implementation of the front end, the dispatch stage and the backend where executed buckets are logged.

[051] Figure 2 shows the manner in which the concept scales from a bucket buffer managing a small number of closely coupled threads into hardware circuits that manage multiple bucket buffers and threads. Those circuits that can be expanded to process larger numbers of threads that might have less close interaction are describe as a global front end (e.g., global front end scheduler 150 shown in Figure 1).

[052] The process starts by fetching a new thread Matrix/bucket/block, and then the new thread bucket is assigned into a vacant bucket slot in the bucket buffer. Each of the thread allocation pointers in the thread allocation pointer array 852 composes an interval of buckets that the thread is allowed physically to place its blocks/buckets of instruction in it. Each of those threads keeps allocating buckets into the bucket buffer array inside its corresponding interval of contiguous space in round-robin fashion. The buckets/blocks inside each thread space get assigned a new number 852 that is incremented each time a new bucket/block gets assigned. For each valid source in the bucket 850. The valid sources for each bucket has a valid read bit "Rv" indicating that this source is needed for the instructions inside this bucket. By the same convention, each destination register that is to be written back by instructions in this bucket has a valid bit "Wv" set in the bucket and it has a field in a destination inheritance vector 853. When a new bucket is to be fetched into the bucket buffer it inherits the destination inheritance vector from the previous allocated bucket pointed at by the thread bucket allocation pointer 852. The inheritance vector is copied from the previously allocated bucket and then it overwrites those valid destination fields that correspond to the registers which will be updated by those bucket instructions. The valid destinations will be labeled with the current bucket number, while the invalid destinations are copied from the corresponding inheritance vector inside the bucket. Then the thread bucket pointer is updated for the new fetched bucket by incrementing its pointer (it wraps around within its interval).

[053] In the bucket dispatch and execute stage, whenever a bucket is executed without any exception handling, then the bucket execution flag (containing the bucket number) 854 is set and broadcasted throughout the bucket buffer and is latched/monitored within each bucket that has a source with that bucket number as a source. It is also possible to pass other related information along the bucket number, such as information about virtual registers locations. When all the execution flags of the sources buckets are set within a bucket, then that bucket ready bit 855 is set and the bucket is ready to be dispatched and executed. When the bucket executes without any exception and it is ready to update the architecture state in the sequential order of the program, then it retires the bucket and the retirement thread pointer 857 is incremented to the next bucket in the array. The retired bucket location can be assigned to a new bucket.

[054] Those closely related threads can all coexist inside the Matrix/bucket/block buffer; each thread will occupy an interval of consecutive buckets that belongs to that thread. The allocation pointer of that thread moves inside this interval of buckets in a round robin fashion fetching new instruction buckets and allocating them inside the thread interval in the described round robin fashion. With such interval sectioning, the whole bucket buffer is divided dynamically with different or equal interval lengths of buckets.

[055] The concept of inheritance vector is introduced here for the instruction bucket as well as for the thread. Each instruction matrix/block/bucket writes into particular registers among the architectural registers. Each new bucket upon allocation stage updates this inheritance vector writing the thread and bucket number of its own into this vector leaving the fields for the registers that it does not write into un-updated. This bucket inheritance vector B_iv 856 is forwarded from each bucket to the next one in program order. In FIG. 8.b each matrix writes its own number into the architecture destination registers if the instructions in that matrix write into those registers, otherwise it inherits the value from the B_iv of the previous bucket in that thread.

[056] Figure 3 shows a diagram of exemplary hardware circuits that shows a segmented register file storing operands and results with an interconnect in accordance with one embodiment of the present invention. Figure 3 shows an operand result buffer coupled via the execution global interconnect to a plurality of execution units.

[057] Figure 4 shows a diagram depicting a global front end scheduler in accordance with one embodiment of the present invention. The global front end scheduler is configured to process larger numbers of threads that might have less close interaction (e.g., global front in scheduler 150 shown in Figure 1). This diagram shows how a sequence of instructions from one logical core is distributed across many virtual cores. This process would be repeated for each logical core present in the machine. It should be noted that the Figure 4 “Engine” comprises the components of a virtual core, where the register file is explicitly depicted to show aspects of inter virtual core communication at the register file level.

[058] For example, as depicted in Figure 4, the global front end scheduler can process a thread header 902 but does not need to process the actual instructions within the thread to enforce the dependency check across those distant threads. The header of the thread and the sub headers of its buckets contain only information about the architecture registers that those threads and buckets writes into (destination registers of those instructions) no need in those headers to include actual instructions or sources of those instructions. Actually it is enough to list those destination registers or a bit vector where each individual bit is set for each register that is a destination for an instruction. The header does not need to be physically placed as a header for the instructions; it can be any formatted packet or compact representation or of the destination registers of the instructions within the threads, which may or may not be stored with the rest of the instructions information.

[059] This global front-end fetches only the headers of the threads/blocks in program order and generates dynamic thread and/or bucket inheritance vectors 901 (Tiv and/or Biv). Each time a new thread is allocated, those inheritance vectors are forwarded by keeping the old fields that the current thread bucket will not write to or update as shown 903. Those inheritance vectors are distributed to a large number of engines/cores or processors 904 each of which might include a local front-end and a fetch unit (which will fetch and store the actual instructions produce the dependency vector for each bucket) and a local matrix/block/bucket buffer with local register files 905. The local front-ends then fetch the actual instructions and use the information from the inheritance vectors obtained from the global front end to fill the dependency information for the instruction sources of the instructions that are brought into those engines for execution. Figure 3 illustrates a global front-end implementation and the way it disseminates the inheritance vectors to the different engines 904 using only concise information about the instructions (e.g., which is just the registers that those instructions write

into). Other information that is of help to place in the header is that information about a change in the control path within or across the threads. A global branch predictor can be used to predict the flow of control across those threads. So such headers can include the branching destinations and offsets. In addition to the branch predictor to determine control flow, the hardware/compiler can decide to dispatch independent threads across the 2 control paths of a branch. In such case it will later merge the execution of those 2 paths using the inheritance vector. Figure 3 also shows the forwarding process when a header of a new thread is fetched by the global front, thread 2 (906) for example will update the corresponding inheritance vector 901 that is forwarded to it resulting in vector 910 where registers 1, 2, 3, 4, 6, 0 and 7 are updated with T2 labels. Note that in 910 register 5 was not written by T2 buckets and thus its label was inherited from a previous inheritance vector.

[060] One interesting observation is that the register files allow cross communication among the cores/engines. An early request (to reduce the access latency) of the registers that are needed from cross engines can be placed as soon as the instruction buckets of the thread are fetched and allocated in the local bucket buffer at that time the source dependency information is populated such that cross engine threads references can be issued at that time probably long before the actual instructions are dispatched for execution. In any case, the instruction will not be dispatched until the cross referenced source is forwarded and arrived. This cross referenced source can be stored in the local multi-threaded register file or register cache. Although this cross referenced source can be stored in a buffer similar to the load store buffer (it can reuse the load store buffer physical storage and dependency check mechanisms but as a register load instead of memory load). Many topologies can be used to connect the register files across the engines/cores, which may be a ring topology or cross bar topology or mesh routed interconnect.

[061] The following discussion can illustrate how register file segmentation can be used inside an engine and also across engines. When the bucket is dispatched, its sources are sent to both (simultaneously or sequentially) to the register file and the register cache. If the register file is physically unified and has direct support for threading, then the operand is read directly from the corresponding thread register section. If the register file is a virtual register, including a physically segmented register file that uses tags, then a tag match has to be done as a part of the virtual register read. If the tag matches, then the read happens from the segmented register file.

[062] Disclosed is register architecture that supports software threads, hardware generated threads, VLIW execution, SIMD & MIMD execution as well as emulation of out-of-order super-scalar execution. Although it is physically segmented, it looks as a unified architecture resource. This segmented register is part of the virtual register file which might include a register hierarchy and a register cache as well as mechanisms to store and check register tags. The tag access can be eliminated if we use a location based scheme that takes advantage of the dependency inheritance vector. The scheme works such that when the executed bucket number is broadcasted during dispatch stage all the sources of subsequent instructions perform a CAM (content addressable match) which compares their sources buckets with the just dispatched/executed bucket to set the ready flag for that source. Here the physical location of where that bucket executed can also be propagated along with the register number so that any ambiguity is resolved.

[063] For example, consider an implementation where there are 4 register file segments, each containing 16 registers. For example upon dispatching a bucket # x to section 2 the bucket number x is broadcasted to the bucket buffer and also the segment #2 is broadcasted with it such that all sources that have a dependency on bucket x will record that it wrote all its registers in segment 2. When time comes to dispatch those instructions, they know that they need to read their register from segment 2 and not any other segment, even though the same register number exists in the other segments. This also applies to the register cache to avoid using tags. We can extend this concept to the global front end where in addition to the thread information the inheritance vector can specify in which engine the instruction bucket writing to this register was allocated.

[064] Figure 5 shows an alternative implementation of the distribution of instructions across many virtual cores in accordance with one embodiment of the present invention. Figure 5 shows a runtime optimizer scheduler 550 functioning by distributing inheritance vectors encode segments to the virtual cores. In one embodiment, the optimizer looks at a number of code blocks of instructions and reschedules instructions across all of the code blocks to create code segments and inheritance vectors. The goal of the Optimizer would be to maximize the efficiency of execution of the overlapped execution of code segments on their respective virtual cores.

[065] Figure 6 shows a plurality of register segments with a corresponding plurality of register files and operand result buffers in accordance with one embodiment of the present invention. As depicted in Figure 6, and execution global interconnect connects each register segment to a plurality of address calculation and execution units.

[066] The register segments Figure 6 can be used to implement one of 3 execution modes: either by being grouped together by the compiler/programmer to form a MIMD super instruction matrix, or each matrix can be executed independently in a threaded mode where separate threads execute simultaneously on each of the 4 hardware sections. The last execution mode possible is the ability to dynamically execute 4 different instruction matrices from a single thread using a hardware dependency check to insure no dependency exists between those different matrices that execute simultaneously on the 4 different hardware sections

[067] The register files in Figure 6 may be alternately configured depending upon the execution mode. In one mode, the register files are viewed as either an MIMD sectioned register file serving a MIMD width of 4 sections or they serve as 4 individual register files, each serving a separate thread. The register files can also support a dynamic execution mode where the 4 sections are one unified register file where data written to any register in a particular section is accessible by all units in the other sections. Switching between those modes can be seamless as different executing modes can alternate between individual thread baseline instruction matrix and MIMD super instruction matrix threads.

[068] In a multithread execution mode, each register file and its execution unit that executes a thread is totally independent of other register files and their threads. This is similar to each thread having its own register state. However, dependency between those threads can be specified. Each matrix that belongs to a thread will execute in the execution unit of that thread's register file. If only one thread or non-threaded single program is executed on the hardware, then the following method is used to allow parallel matrices belonging to that single thread/program to be able to access the results written into the registers in the other sections. The way this is done is by allowing any matrix writing results into any one of the 4 register file sections to generate copies of those registers in the other register file sections. Physically this is done by extending the write ports of each section into the remaining sections. However, this is not scalable, as we cannot build an efficient register

file with each memory cell having as many as 4 times the write ports as needed for one section alone. We present a mechanism where the register file is built such that it will not be impacted with such single thread register-broadcast extension.

[069] It should be noted that additional aspects regarding register segments as used in embodiments of the present invention can be found in US Patent Application serial number 2010/0161948, titled "APPARATUS AND METHOD FOR PROCESSING COMPLEX INSTRUCTION FORMATS IN A MULTITHREADED ARCHITECTURE SUPPORTING VARIOUS CONTEXT SWITCH MODES AND VIRTUALIZATION SCHEMES" by Mohammad A. Abdallah, filed on November 14, 2007.

[070] Figure 7 shows a more detailed diagram of a fragmented memory subsystem for a multicore processor in accordance with one embodiment of the present invention. Figure 7 shows a comprehensive scheme and implementation of the synchronization scheme among threads and/or among loads and stores in general. The scheme describes a preferred method for synchronization and disambiguation of memory references across load/store architectures and/or across memory references and/or threads' memory accesses. In Figure 2, we show multiple segments of register files (address and or data registers), execution units, address calculation units, and fragments of level 1 caches and/or load store buffers and level 2 caches and address register interconnects 1200 and address calculation unit interconnects 1201. Those fragmented elements could be constructed within one core/processor by fragmenting and distributing its centralized resources into several engines or they can be constructed from elements of different cores/processors in multi-core/multi-processor configurations. One of those fragments 1211 is shown in the figure as fragment number 1; the fragments can be scaled to a large number (in general to N fragments as shown in the figure).

[071] This mechanism also serves also as a coherency scheme for the memory architecture among those engines/cores/processors. This scheme starts by an address request from one of the address calculation units in one fragment/core/processor. For example, assume the address is requested by fragment 1 (1211). It can obtain and calculate its address using address registers that belong to its own fragment and or from registers across other fragments using the address interconnect bus 1200. After calculating the address it creates the reference address of either 32-bit address or 64-bit address that is used to access caches and

memory. This address is usually fragmented into a tag field and a set and line fields. This particular fragment/engine/core will store the address into its load store buffer and/or L1 and/or L2 address arrays 1202, at the same time it will create a compressed version of the tag (with smaller number of bits than the original tag field of the address) by using a compression technique.

[072] More the different fragments/engines/cores/processors will use the set field or a subset of the set field as an index to identify which fragment/core/processor the address is maintained in. This indexing of the fragments by the address set field bits ensures exclusiveness of ownership of the address in a particular fragment/core/engine even though the memory data that corresponds to that address can live in another or multiple other fragments/engines/cores/processors. Even though the address CAM/tag arrays 1202/1206 are shown in each fragment to be coupled with the data arrays 1207, they might be only coupled in physical proximity of placement and layout or even by the fact that both belongs to a particular engine/core/processor, but there is no relation between addresses kept in the address arrays and the data in the data arrays inside one fragment.

[073] Figure 8 shows a diagram depicting how bits of an address can be used by address generation to enumerate fragments in accordance with one embodiment of the present invention. In the present embodiment, fragments are defined by the address bits that are above page boundaries and that are below cache line boundaries, as depicted in Figure 8. The present invention advantageously stays above the page boundaries to avoid causing TLB misses during the translation from the virtual addresses to physical addresses. The process stays below the cache line boundary in order to have complete cache lines in order to correctly fit within the hardware cache hierarchy. For example, in a system that employs 64 byte cache lines, the fragment boundary would avoid the last six address bits. In comparison, a system that employs 32 byte cache lines, the fragment boundary would avoid the last five bits. Once defined, the fragment hierarchy is the same across all cache hierarchies of the processor.

[074] Figure 9 shows a diagram of how loads and stores are handled by embodiments of the present invention. As depicted in Figure 9, each fragment is associated with its load store buffer and store retirement buffer. For any given fragment, loads and stores that designate an address range associated with that fragment or another fragment are

sent to that fragment's load store buffer for processing. It should be noted that they may arrive out of order as the cores execute instructions out of order. Within each core, the core has access to not only its own register file but each of the other cores' register files.

[075] Embodiments of the present invention implement a distributed load store ordering system. The system is distributed across multiple fragments. Within a fragment, local data dependency checking is performed by that fragment. This is because the fragment only loads and stores within the store retirement buffer of that particular fragment. This limits the need of having to look to other fragments to maintain data coherency. In this manner, data dependencies within a fragment are locally enforced.

[076] With respect to data consistency, the store dispatch gate enforces store retirement in accordance with strict in-program order memory consistency rules. Stores arrive out of order at the load store buffers. Loads arrive out of order also at the load store buffers. Concurrently, the out of order loads and stores are forwarded to the store retirement buffers for processing. It should be noted that although stores are retired in order within a given fragment, as they go to the store dispatch gate they can be out of order from the multiple fragments. The store dispatch gate enforces a policy that ensures that even though stores may reside across store retirement buffers out of order, and even though the buffers may forward stores to the store dispatch gate out of order with respect to other buffers' stores, the dispatch gate ensures that they are forwarded to fragment memory strictly in order. This is because the store dispatch gate has a global view of stores retiring, and only allows stores to leave to the global visible side of the memory in order across all the fragments, e.g., globally. In this manner, the store dispatch gate functions as a global observer to make sure that stores ultimately return to memory in order, across all fragments.

[077] Figure 10 shows the manner which fragments can be split into two or more domains in accordance with one embodiment of the present invention. Figure 10 shows the manner in which a single fragment can be split into multiple domains. Domain splitting can be implemented via the address generation process. Domain splitting changes the manner in which load store checks have to be done within a fragment, since in this case they only have to be done per domain, as opposed to across the entire fragment. Domain splitting also is advantageous in that it can enable single ported memory to behave like multiport memory, where the single port is accessed per different domain.

[078] Figure 11 shows a mode of operation of the processor wherein the hardware resources of the partitionable engines are used to function like logical cores in executing applications in accordance with one embodiment of the present invention. In this embodiment, the hardware resources of the engines of the virtual cores are configured as physical cores. In the Figure 11 mode, each physical core of is configured to function as a logical core. Multithreaded applications and multithreaded functionality is up to the threaded programmability of the software of the application.

[079] Figure 12 shows a mode of operation of the processor wherein soft cores are used to function like logical cores in executing applications in accordance with one embodiment of the present invention. In this embodiment, the partitionable engines of virtual cores will support a plurality of soft cores. In the Figure 12 mode, each soft core is configured to function as a logical core. Multithreaded applications and multithreaded functionality is up to the threaded programmability of the software of the application.

[080] Figure 13 shows a mode of operation of the processor wherein the soft cores are used to function like a single logical core in executing applications in accordance with one embodiment of the present invention. In the Figure 13 mode, each soft core is configured to function as a single logical core. In such an implementation, a single threaded application has its instruction sequence divided up and allocated among the virtual cores where they are cooperatively executed to achieve high single threaded performance. In this manner, single threaded performance can scale with the addition of additional soft cores.

[081] A number of strategies can be used in choosing the operating mode of the processor. For a processor having a large number of engines (e.g., 8 engines, 12 engines, etc.), a number of soft cores can be configured to function as a single logical core, while the remaining cores can operate in the other modes. This attribute allows an intelligent partition of resources to ensure maximum utilization of the hardware and/or minimal wasted power consumption. For example, in one embodiment, cores (e.g., soft or logical) can be allocated on a per thread basis depending upon the type of application that is executing.

[082] Figure 14 shows an exemplary implementation of fragment segmentation being used to support logical core and virtual core functionality in accordance with one embodiment of the present invention. As described above, the fragment segmentation allows

the processor to be configured to support different virtual core execution modes, as described above.

[083] The global interconnect allows cores' threads to access any of the ports 1401. It should be noted that the term "thread" as used herein refers to either a representation of instruction sequences from different logical cores, instruction sequences from the same logical core, or some mixture of the two.

[084] The manner in which the threads utilize one of the ports 1401 to access the load store buffer is adjustable in accordance with the policies of the arbiters, as shown. Thus, a thread using any one of the ports 1401 can have a greater or lesser amount of access to the load store buffer via the ports 1402. The size of the allocation and the manner in which the allocation is managed is controlled by the arbiter. The arbiter can dynamically allocate access to the ports in accordance with demands of a particular thread.

[085] The load store buffer is configured to have a plurality of entries spread across the ports. Access to the load store buffer is controlled by the arbiter. In this way, the arbiter can dynamically allocate entries in the load store buffer to the different threads.

[086] Figure 14 also shows arbiters on the ports between load store buffer and the L1 cache. Thus, as with the load store buffer described above, a thread using any one of the ports 1403 can have a greater or lesser amount of access to the L1 cache via the ports 1404. The size of the allocation and the manner in which the allocation is managed is controlled by the arbiter. The arbiter can dynamically allocate access to the ports in accordance with demands of a particular thread.

[087] The L1 cache is configured to have a plurality of ways spread across the ports. Access to the L1 cache is controlled by the arbiter. In this way, the arbiter can dynamically allocate entries in the L1 cache to the different threads.

[088] In one embodiment, the arbiters are configured to function with a plurality of counters 1460 that are used for tracking functionality and a plurality of threshold limit registers 1450 that provide a limiting function. The limiting function specifies the maximum resource allocation percentage for a given thread. The tracking function tracks the actual resources allocated to a given thread at any given time. These tracking and limiting functionalities affect the allocation of the number of per thread entries, ways, or ports for the

load store buffer, L1 cache, L2 cache or the global interconnects. For example, the total number of entries in the load store buffer allocated for each thread can be dynamically checked against a variable threshold. This variable threshold can be updated in accordance with a given thread's forward progress. For example, in one embodiment, threads that are slowed down (e.g., large number or L2 misses, etc.) are quantified as making slow forward progress, and thus, their respective resource allocation thresholds are lowered, including the entries thresholds, the ways thresholds and the ports thresholds.

[089] Figure 14 also shows a shared L2 cache. In the present embodiment, the shared L2 cache has a fixed port arrangement without any arbitration between accesses coming from the L1 cache. Threads executing on the processor would all share access to the L2 cache and the resources of the L2 cache.

[090] Figure 15 shows a fragment memory of an exemplary four fragment processor implementing a many physicals to many logicals mode in accordance with one embodiment of the present invention.

[091] One example logical core and its relationship with the resources of the processor is shown by the shading on Figure 15. In the Figure 11 mode of operation, the many physical cores to many logical cores mode, wherein the physical cores are used to function like logical cores in executing applications, each logical core will be configured to have a fixed ratio of the resources of the load store buffer and the L1 cache. The ports can be specifically assigned to each thread or core. Entries in the load store buffer can be specifically reserved per thread or core. Ways within the L1 cache can be specifically reserved per thread or core. Multithreaded applications and multithreaded functionality is up to the threaded programmability of the software of the application. This is shown by the one logical core having an allocated port and an allocated portion of the store buffer and the L1 cache of each of the fragments. In this manner, the logical core comprises a fixed allocated slice of the resources of each fragment.

[092] In one embodiment, in the many physical cores to many logical cores mode, the four fragments can be partitioned in accordance with the number of ports (e.g., ports 1401) that access each fragment. For example, in an embodiment where there are 6 ports per fragment, the resources of each fragment, and hence the resources of each partition will engine, can be divided in such a way to support 6 physical cores across the 4 fragments and

the 4 partition double engines. Each partition can be allocated its own port. Similarly, the resources of the load store buffer and the L1 cache would be allocated in such a way to support 6 physical cores. For example, in an embodiment where the load store buffer has 48 entries, the 48 entries can be allocated such that there are 12 entries per physical core to support a mode where 4 physical cores are implemented, or they can be allocated such that there are 8 entries per physical core disappointment where 6 physical cores are implemented.

[093] Figure 16 shows a fragment memory of an exemplary four fragment processor implementing a many physicals to many logicals mode in accordance with an alternative embodiment of the present invention.

[094] As with Figure 15, the one example logical core and its relationship with the resources of the processor is shown by the shading on Figure 16. In the Figure 11 mode of operation, the many physical cores to many logical cores mode, an entire partition table engine is dedicated to support the execution of a single logical core. This is shown by the shading in Figure 16. The physical resource is engine is used to function like logical core in executing applications.

[095] Figure 17 shows a fragment memory of an exemplary four fragment processor implementing a many soft cores to many logical cores mode in accordance with one embodiment of the present invention.

[096] One example logical core and its relationship with the resources of the processor is shown by the shading on Figure 17. In the Figure 12 mode of operation, the many soft cores to many logicals mode, where virtual cores are used to function like logical cores in executing applications, the size of the allocation of the resources of the load store buffer and the manner in which the allocation is managed is controlled by the arbiter. The arbiter can dynamically allocate access to the ports in accordance with demands of a particular thread or core. Similarly, the size of the allocation of the resources of the L1 cache and the manner in which the allocation is managed is controlled by the arbiter. The arbiter can dynamically allocate access to the ports in accordance with demands of a particular thread or core. Thus, at any given instance the logical thread/core (e.g. shaded) can use different arbiters and different ports.

[097] In this manner, access to the resources of the load store buffer and access to the resources of the L1 cache can be more policy driven and can be more based on the needs of individuals threads or cores making forward progress. This is shown by the one logical core having a dynamically allocated port and a dynamically allocated portion of the store buffer and the L1 cache of each of the fragments. In this manner, the logical core comprises a non-fixed, dynamically allocated slice of the resources of each fragment.

[098] Figure 18 shows a fragment memory of an exemplary four fragment processor implementing a many soft cores to one logical core mode in accordance with one embodiment of the present invention.

[099] In the Figure 13 mode of operation, the many soft cores to one logical core mode, wherein the soft cores are used to function like a single logical core in executing applications, each of the soft cores is configured to function cooperatively with the other soft cores as a single logical core. A single thread or core has all the resources of the load store buffers and all of the resources of the L1 caches. In such an implementation, a single threaded application has its instruction sequence divided up and allocated among the soft cores where they are cooperatively executed to achieve high single threaded performance. In this manner, single threaded performance can scale with the addition of additional soft cores. This is shown in Figure 17 where the one example logical core and its relationship with the resources of the processor is shown by the shading of all of the resources of the processor.

[0100] Figure 19 shows address calculation and execution units, operand/result buffers, threaded register files, and common partition schedulers of an exemplary four fragment processor implementing a many physicals to many logicals mode in accordance with one embodiment of the present invention.

[0101] One example logical core and its relationship with the resources of the processor is shown by the shading on Figure 19. In the Figure 11 mode of operation, the many physical cores to many logical cores mode, wherein the physical cores are used to function like logical cores in executing applications, each logical core will be configured to have a fixed ratio of the resources of the address calculation units, operand/result buffers, threaded register files, and common partition scheduler. Multithreaded applications and multithreaded functionality is up to the threaded programmability of the software of the application. This is shown by the one logical core having an allocated address calculation and

execution unit, an allocated threaded register file and an allocated common partition scheduler. In this manner, the logical core comprises a fixed allocated segment. However, in one embodiment, in this mode of operation, the address calculation and execution units can still be shared (e.g., meaning each of the address calculation and execution units would be unshaded).

[0102] Figure 20 shows an alternative implementation for the address calculation and execution units, operand/result buffers, threaded register files, and common partition schedulers of an exemplary four fragment processor to implement a many physicals to many logicals mode in accordance with one embodiment of the present invention.

[0103] One example logical core and its relationship with the resources of the processor is shown by the shading on Figure 20. In the figure 20 embodiment however, the resources of a physical core is spread across each of the segments and each of the partitionable engines. This is shown by the one logical core having an allocated portion of the address calculation and execution units, an allocated portion of the threaded register files and an allocated portion of common partition schedulers across each of the segments. Additionally, figure 20 shows how the one logical core would have been allocated portion of the resources of each of the address calculation execution units. In this manner, the logical core comprises a fixed allocated portion of each of the segments.

[0104] Figure 21 shows address calculation and execution units, register files, and common partition schedulers of an exemplary four fragment processor implementing a many soft cores to many logical cores mode in accordance with one embodiment of the present invention.

[0105] One example logical core and its relationship with the resources of the processor is shown by the shading on Figure 21. In the Figure 12 mode of operation, the many soft cores to many logical cores mode, wherein the soft cores are used to function like logical cores in executing applications, each logical core will be configured to have a shared access to any one the address calculation units, and a dynamic allocated portion of the operand/result buffers, threaded register files, and common partition scheduler. Multithreaded applications and multithreaded functionality is up to the threaded programmability of the software of the application.

[0106] Figure 22 shows address calculation and execution units, register files, and common partition schedulers of an exemplary four fragment processor implementing a many soft cores to one logical core mode in accordance with one embodiment of the present invention.

[0107] One example logical core and its relationship with the resources of the processor is shown by the shading on Figure 22. In the Figure 13 mode of operation, the many soft cores to one logical core mode, wherein the soft cores are used to function like a single logical core in executing applications, each logical core will be configured to have a shared access to all of the address calculation units, and all of the operand/result buffers, threaded register files, and common partition schedulers. In such an implementation, a single threaded application has its instruction sequence divided up and allocated among the virtual cores where they are cooperatively executed to achieve high single threaded performance. In this manner, single threaded performance can scale with the addition of additional soft cores.

[0108] Figure 23 shows a diagram of an exemplary microprocessor pipeline 2300 in accordance with one embodiment of the present invention. The microprocessor pipeline 2300 includes a fetch module 2301 that implements the functionality of the process for identifying and extracting the instructions comprising an execution, as described above. In the Figure 23 embodiment, the fetch module is followed by a decode module 2302, an allocation module 2303, a dispatch module 2304, an execution module 2305 and a retirement modules 2306. It should be noted that the microprocessor pipeline 2300 is just one example of the pipeline that implements the functionality of embodiments of the present invention described above. One skilled in the art would recognize that other microprocessor pipelines can be implemented that include the functionality of the decode module described above.

[0109] The foregoing description, for the purpose of explanation, has been described with reference to specific embodiments. However, the illustrated discussions above are not intended to be exhaustive or to limit the invention to the precise forms disclosed. Many modifications and variations are possible in view of the above teachings. Embodiments were chosen and described in order to best explain the principles of the invention and its practical applications, to thereby enable others skilled in the art to best utilize the invention and various embodiments with various modifications as may be suited to the particular use contemplated.

CLAIMS

What is claimed is:

- 5 1. A system for executing instructions using a plurality of register file segments for a processor, comprising:
- a global front end scheduler for receiving an incoming instruction sequence, wherein the global front end scheduler partitions the incoming instruction sequence into a plurality of code blocks of instructions and generates a plurality of inheritance vectors describing
- 10 interdependencies between instructions of the code blocks;
- a plurality of virtual cores of the processor coupled to receive code blocks allocated by the global front end scheduler, wherein each virtual core comprises a respective subset of resources of a plurality of partitionable engines, wherein the code blocks are executed by using the partitionable engines in accordance with a virtual core mode and in accordance with
- 15 the respective inheritance vectors; and
- a plurality register file segments coupled to the partitionable engines for providing data storage.
2. The system of claim 1 wherein the plurality register file segments implement an
- 20 execution mode wherein a subset of physical resources of each register file segment are allocated to support execution of a single logical thread of a logical core.
3. The system of claim 2, each register file segment implements a portion of a
- 25 plurality of logical cores.
4. The system of claim 1, wherein the plurality of register file segments implement an execution mode wherein physical resources of each register file segment are dynamically allocated in accordance with an adjustable threshold to support execution of a single logical thread of a single logical core.
- 30 5. The system of claim 4, wherein the plurality of register file segments implement a portion of a plurality of logical cores.

6. The system of claim 1, wherein the plurality of register file segments implement an execution mode wherein the set of physical resources of each register file segment are allocated to support execution of a single logical thread.

5 7. The system of claim 1, wherein each register file segment further comprises a common partition scheduler, an operand and result buffer and a threaded register file.

8. The system of claim 1, wherein an interconnect links each of the plurality of register file segments to each of the plurality of execution units.

10

9. A processor for executing instructions using a plurality of register file segments, comprising:

a global front end scheduler for receiving an incoming instruction sequence, wherein the global front end scheduler allocates incoming the instruction sequence into a plurality of code segments of instructions, and generates a plurality of inheritance vectors describing interdependencies between instructions of the code segments;

15

a plurality execution units coupled to the global front scheduler for executing the code segments in accordance with respective inheritance vectors; and

a plurality register file segments coupled to the execution units for providing data storage.

20

10. The processor of claim 9 wherein the plurality register file segments implement an execution mode wherein a subset of physical resources of each register file segment are allocated to support execution of a single logical thread of a logical core.

25

11. The processor of claim 10, each register file segment implements a portion of a plurality of logical cores.

12. The processor of claim 9, wherein the plurality of register file segments implement an execution mode wherein physical resources of each register file segment are dynamically allocated in accordance with an adjustable threshold to support execution of a single logical thread of a single logical core.

30

13. The processor of claim 12, wherein the plurality of register file segments implement a portion of a plurality of logical cores.

5 14. The processor of claim 9, wherein the plurality of register file segments implement an execution mode wherein the set of physical resources of each register file segment are allocated to support execution of a single logical thread.

10 15. The processor of claim 9, wherein each register file segment further comprises a common partition scheduler, an operand and result buffer and a threaded register file.

16. The processor of claim 9, wherein an interconnect links each of the plurality of register file segments to each of the plurality of execution units.

15 17. A system for executing instructions using a plurality of register file segments for a processor, comprising:

a global front end scheduler for receiving an incoming instruction sequence, wherein the global front end scheduler partitions the incoming instruction sequence into a plurality of code blocks of instructions;

20 a plurality of virtual cores of the processor coupled to receive code blocks allocated by the global front end scheduler, wherein each virtual core comprises a respective subset of resources of a plurality of partitionable engines, wherein the code blocks are executed by using the partitionable engines in accordance with a virtual core mode; and

25 a plurality register file segments coupled to the partitionable engines for providing data storage.

18. The system of claim 17 wherein the plurality register file segments implement an execution mode wherein a subset of physical resources of each register file segment are allocated to support execution of a single logical thread of a logical core.

30 19. The system of claim 18, each register file segment implements a portion of a plurality of logical cores.

20. The system of claim 17, wherein the plurality of register file segments implement an execution mode wherein physical resources of each register file segment are dynamically allocated in accordance with an adjustable threshold to support execution of a single logical thread of a single logical core.

5

21. The system of claim 20, wherein the plurality of register file segments implement a portion of a plurality of logical cores.

22. The system of claim 17, wherein the plurality of register file segments implement an execution mode wherein the set of physical resources of each register file segment are allocated to support execution of a single logical thread.

10

23. The system of claim 17, wherein each register file segment further comprises a common partition scheduler, an operand and result buffer and a threaded register file.

15

24. The system of claim 17, wherein an interconnect links each of the plurality of register file segments to each of the plurality of execution units.

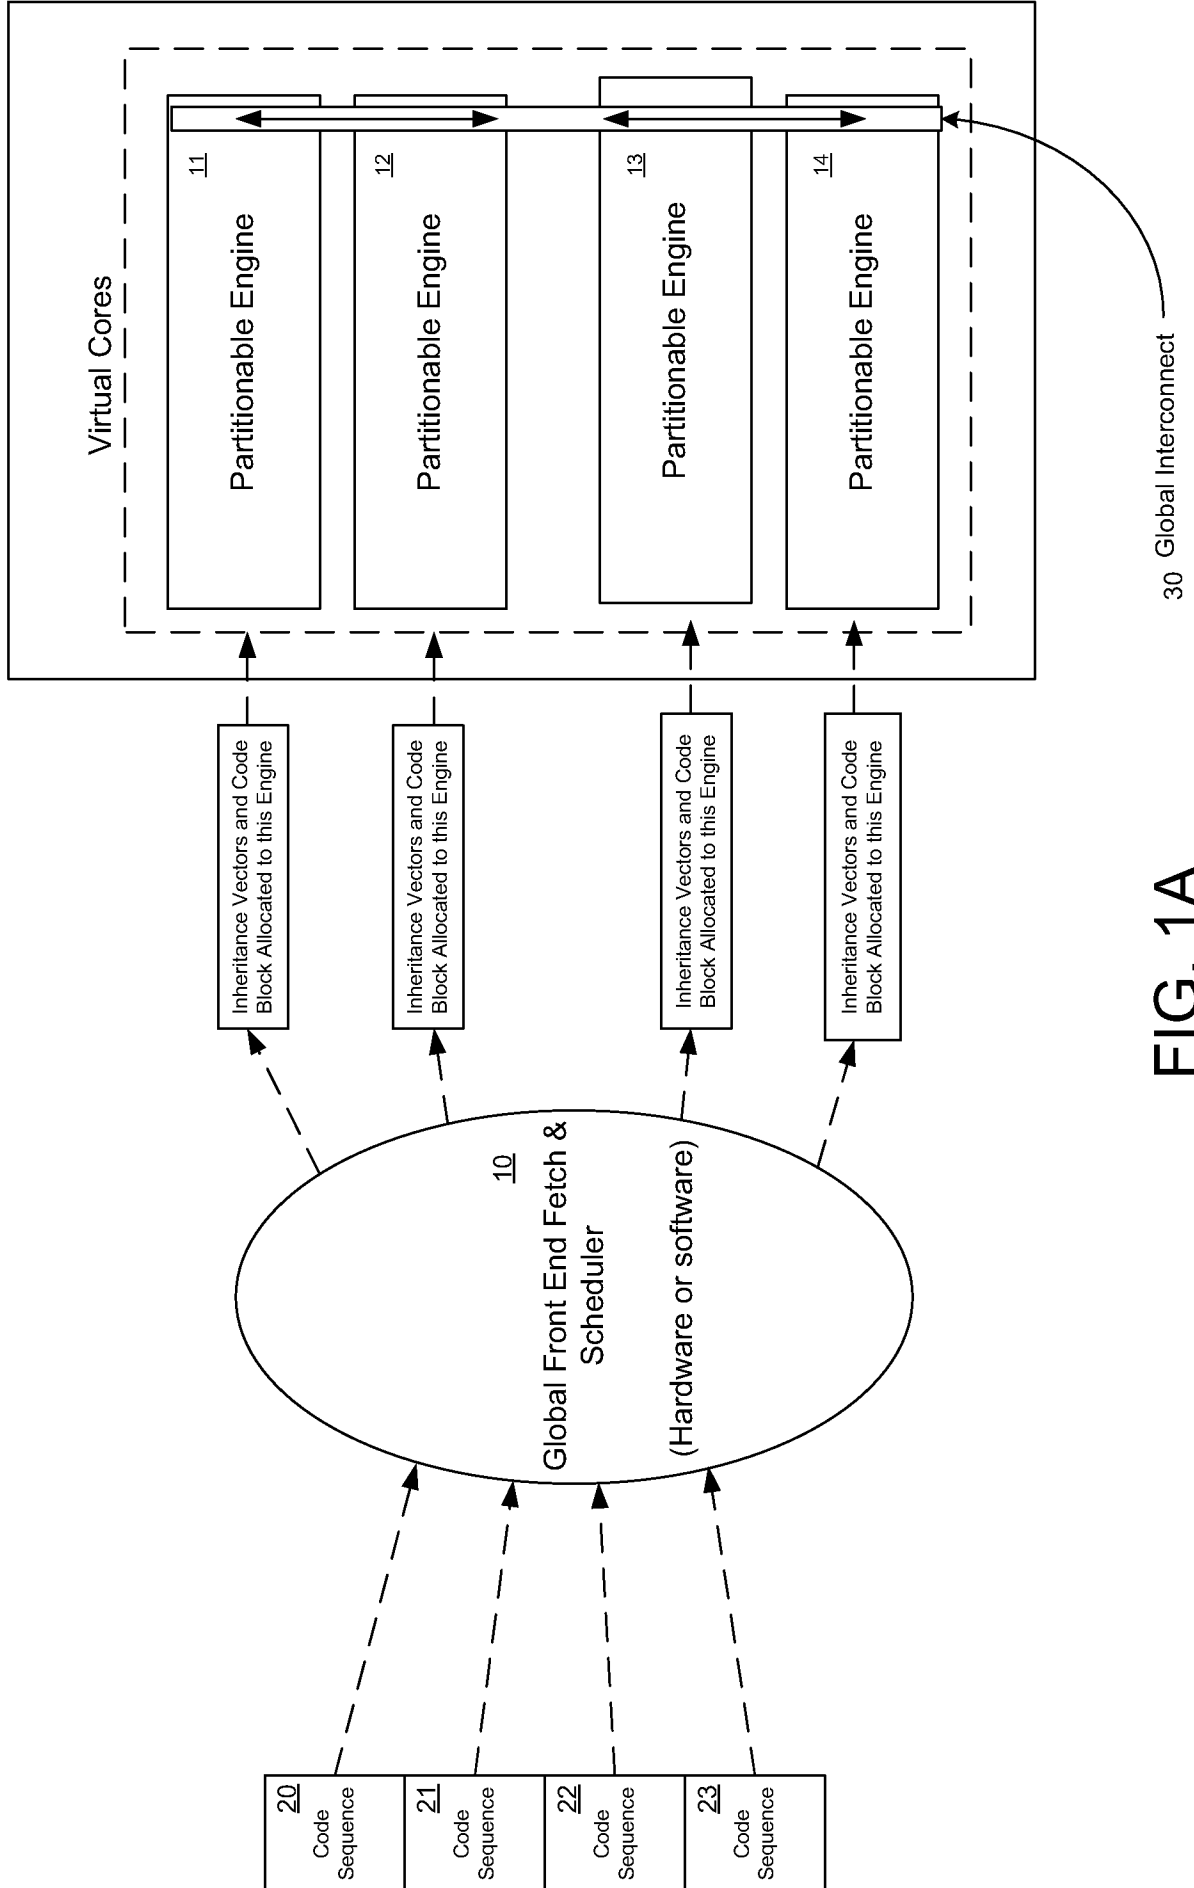


FIG. 1A

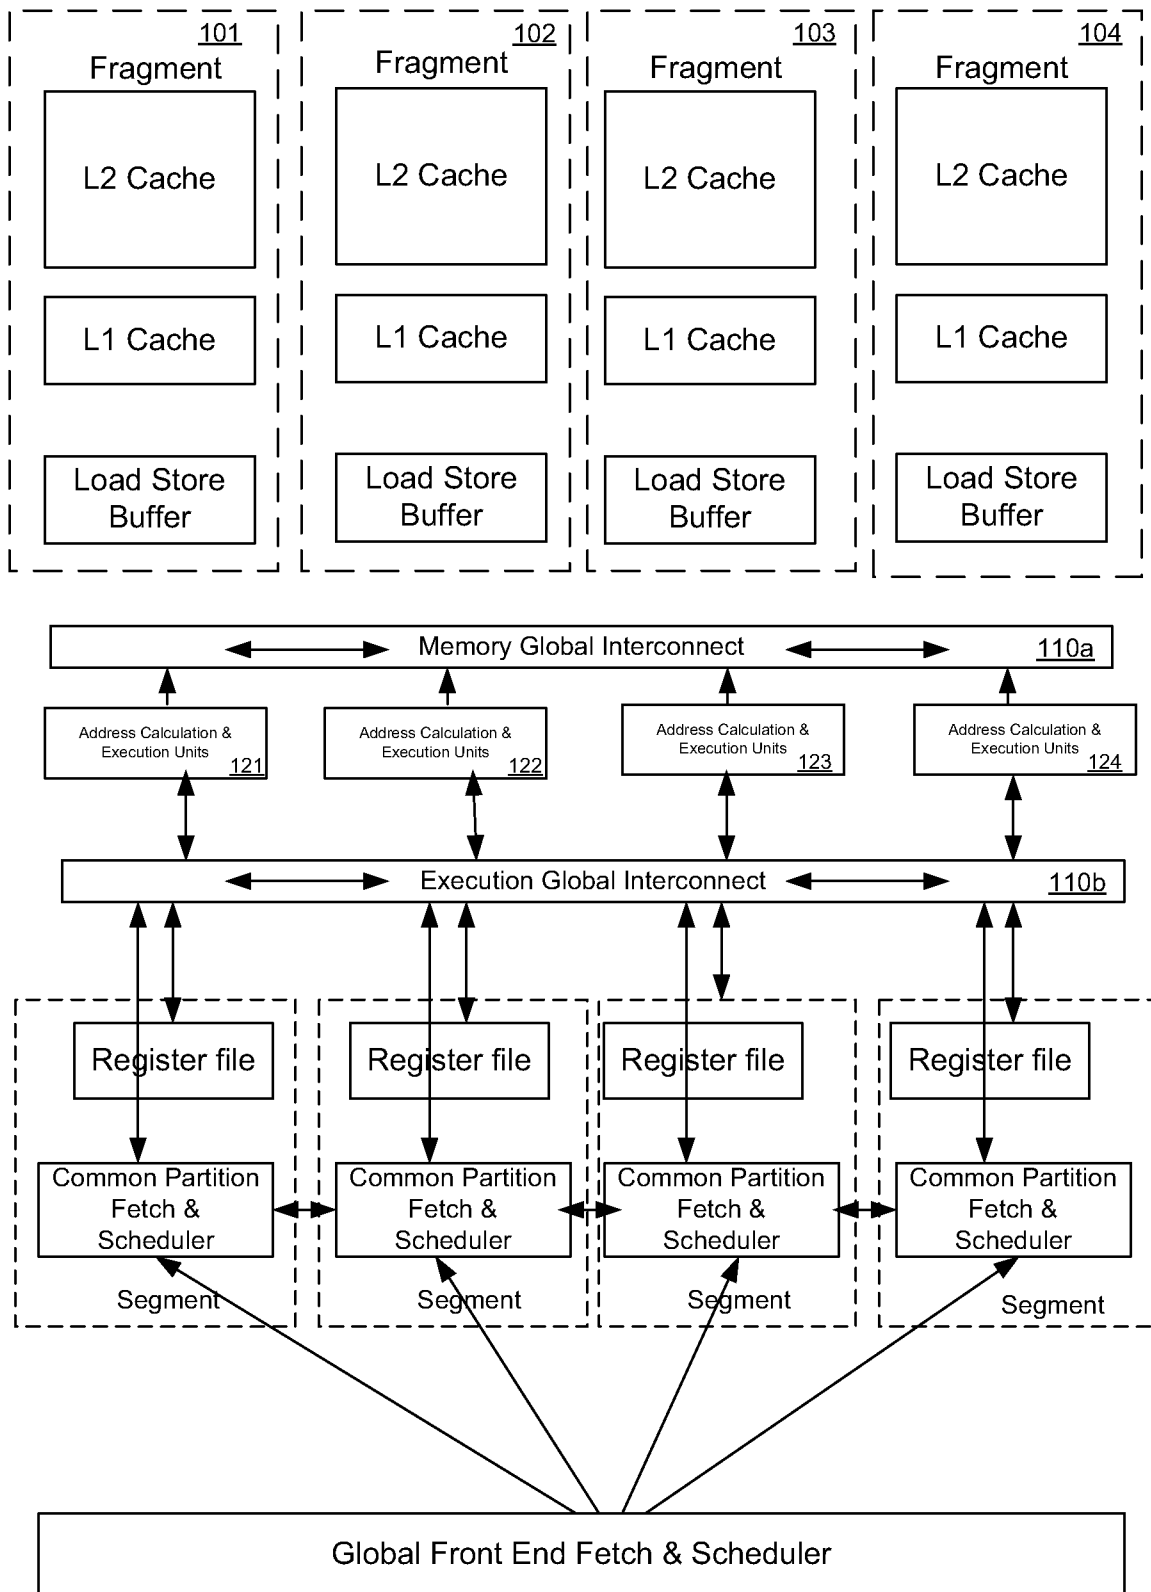


FIG. 1B




FIG. 2




FIG. 3

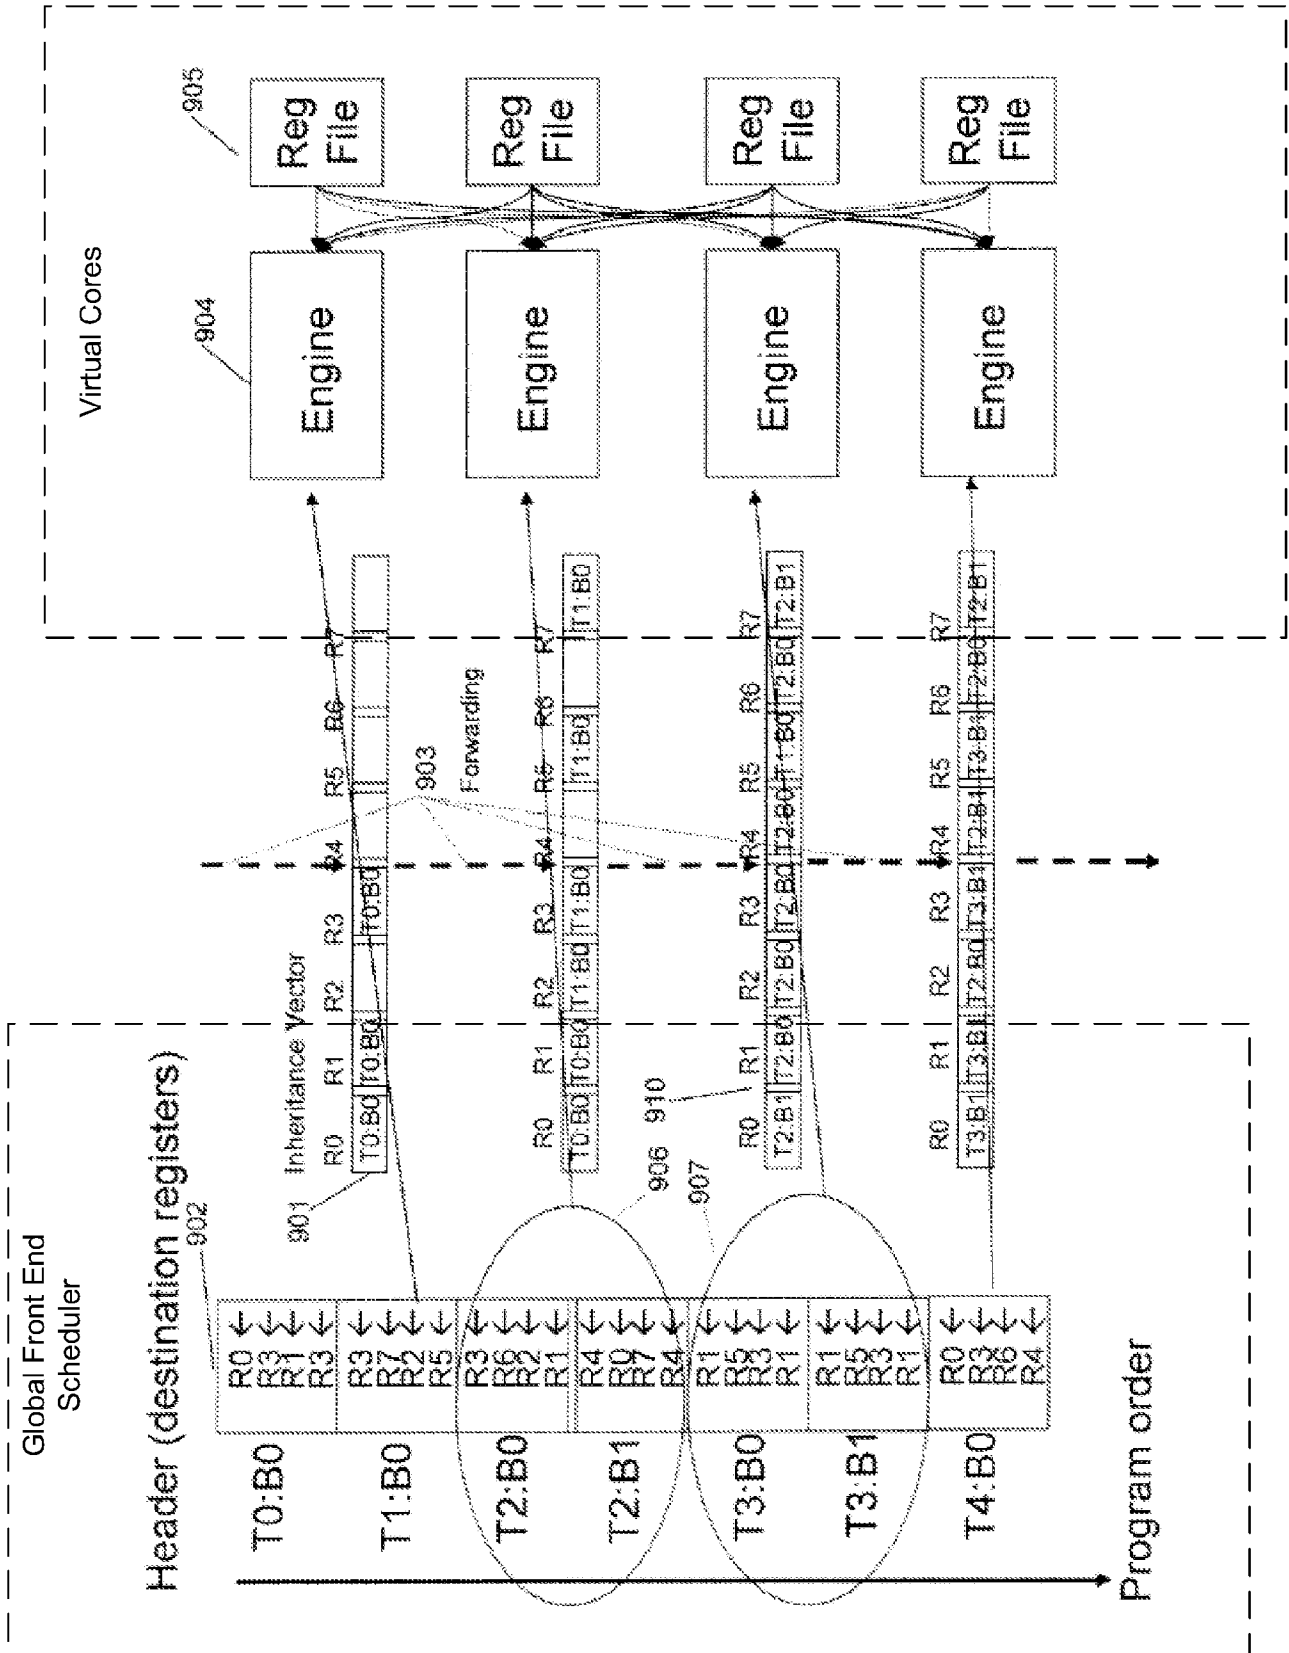


FIG. 4

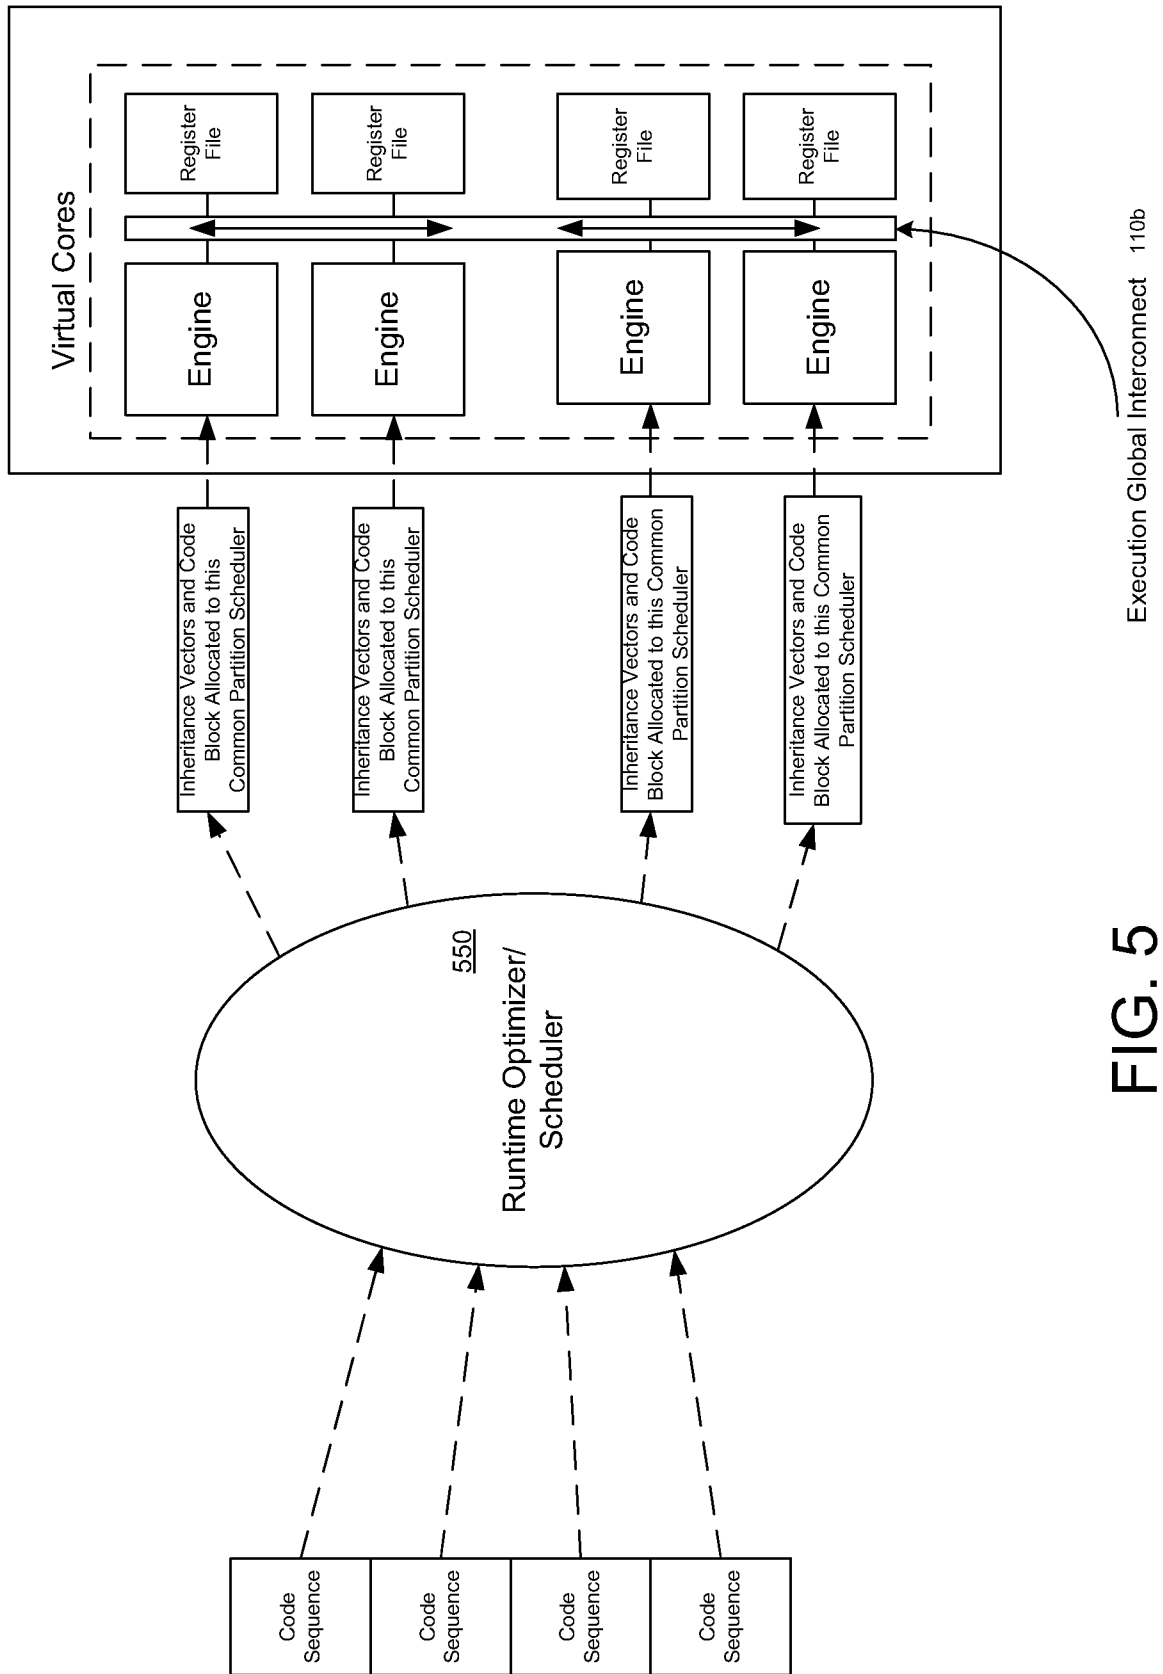


FIG. 5

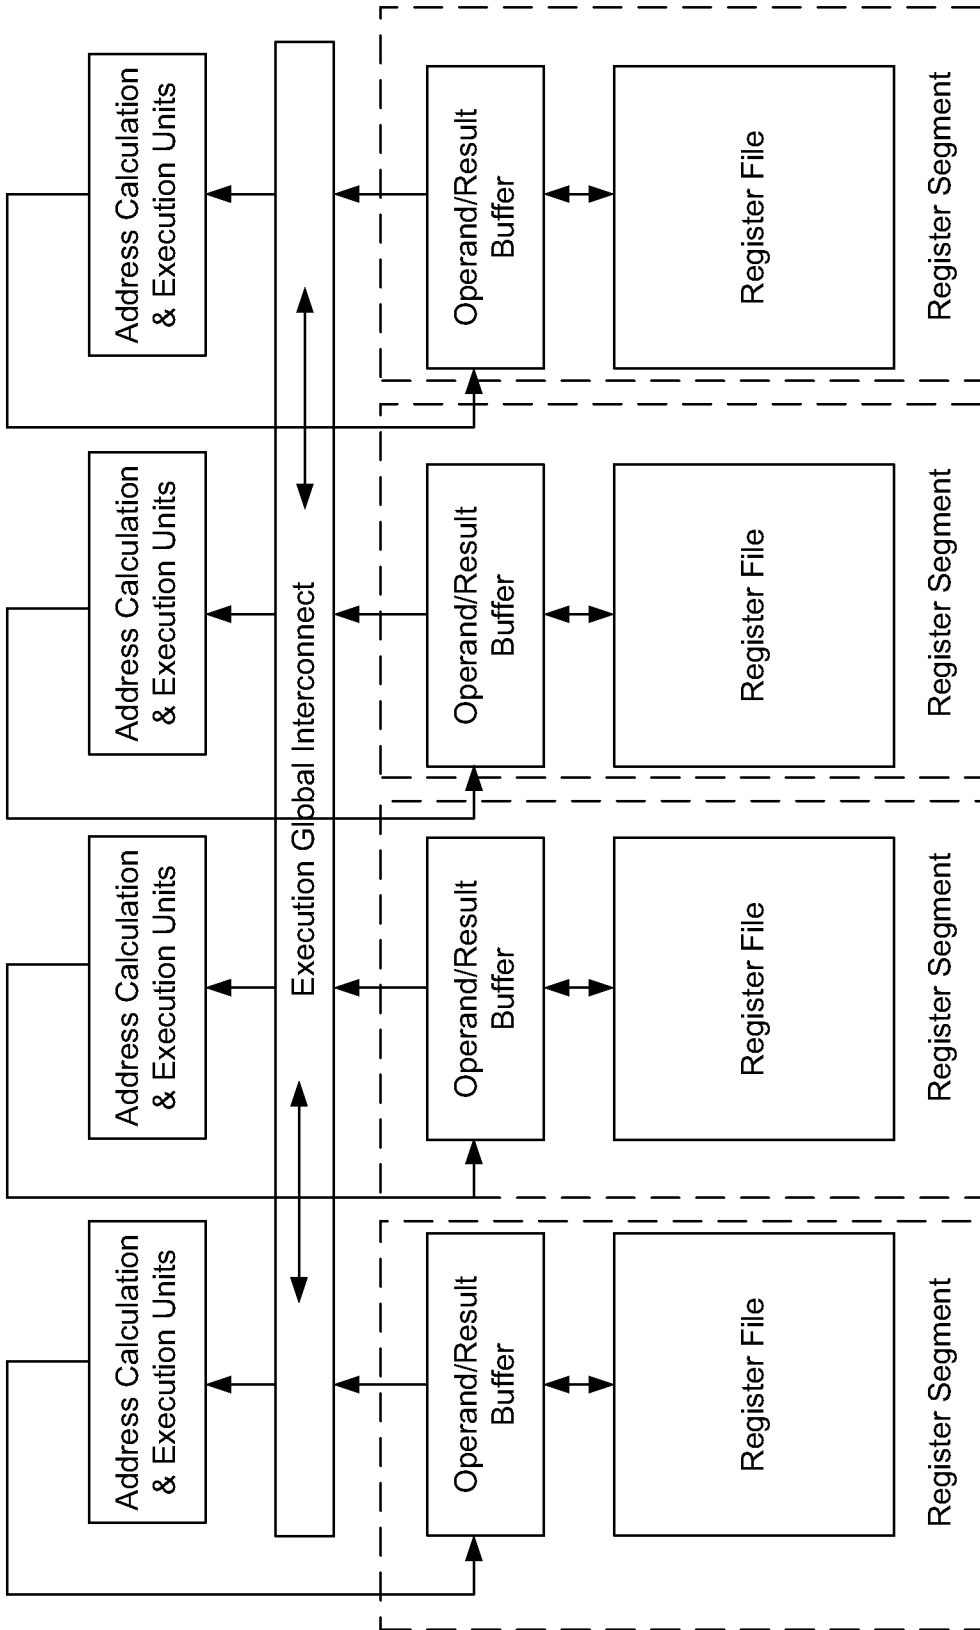


FIG. 6




FIG. 7




FIG. 8

10/24

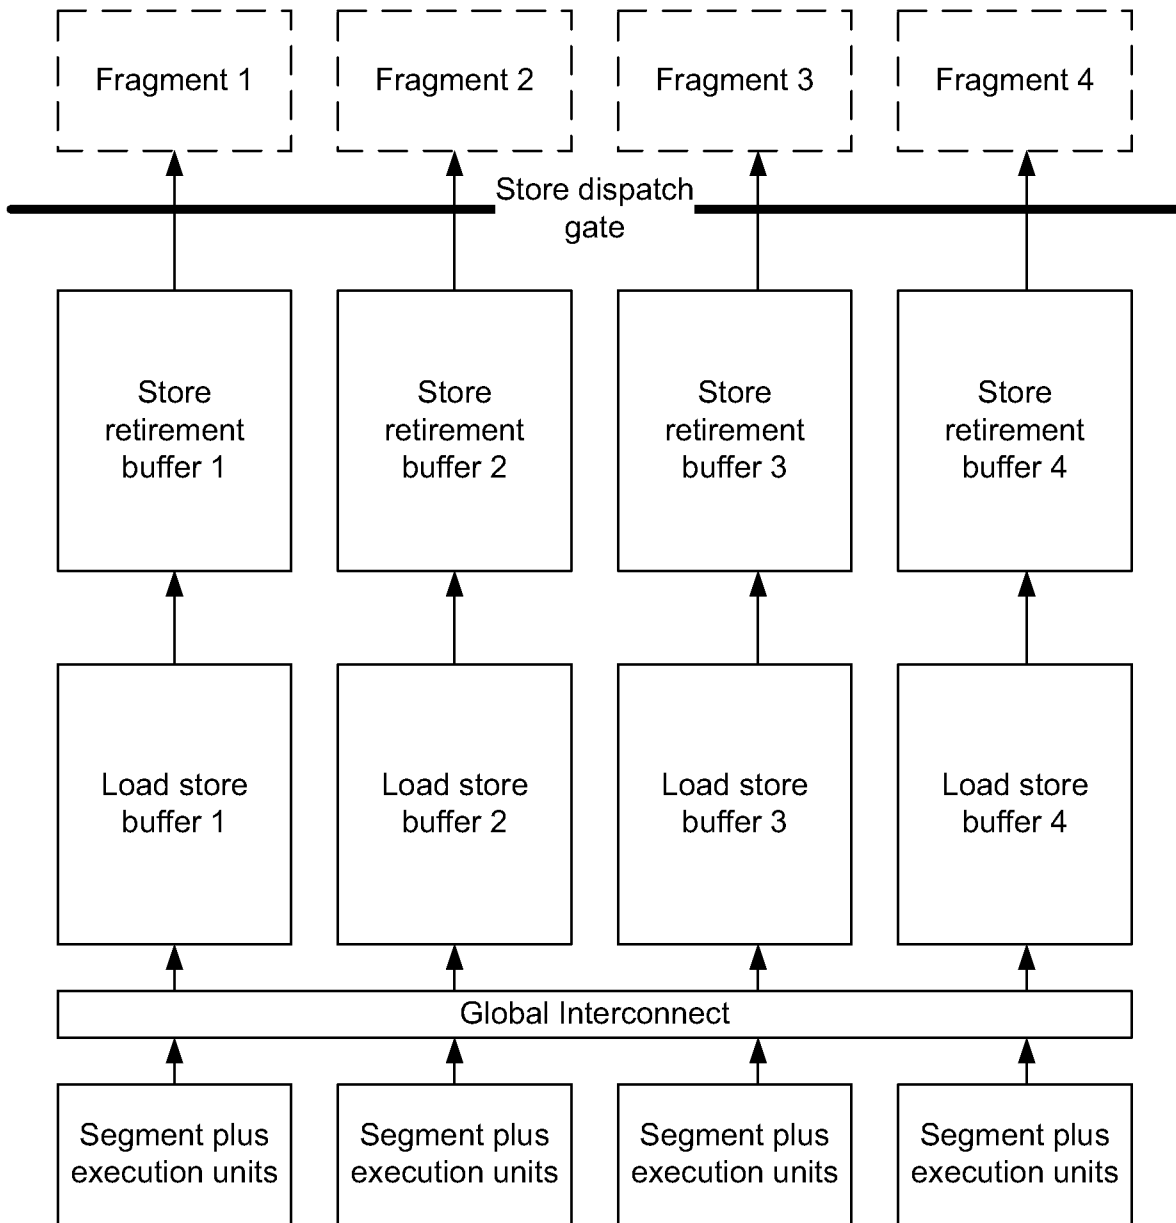


FIG. 9

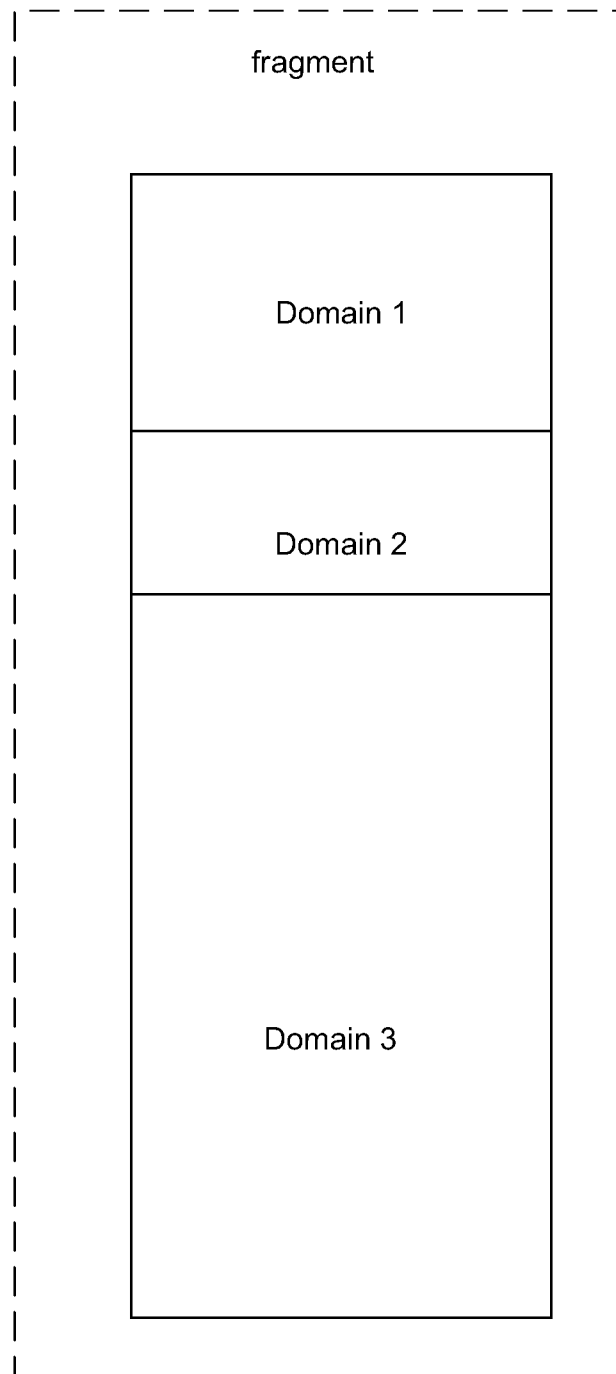


FIG. 10

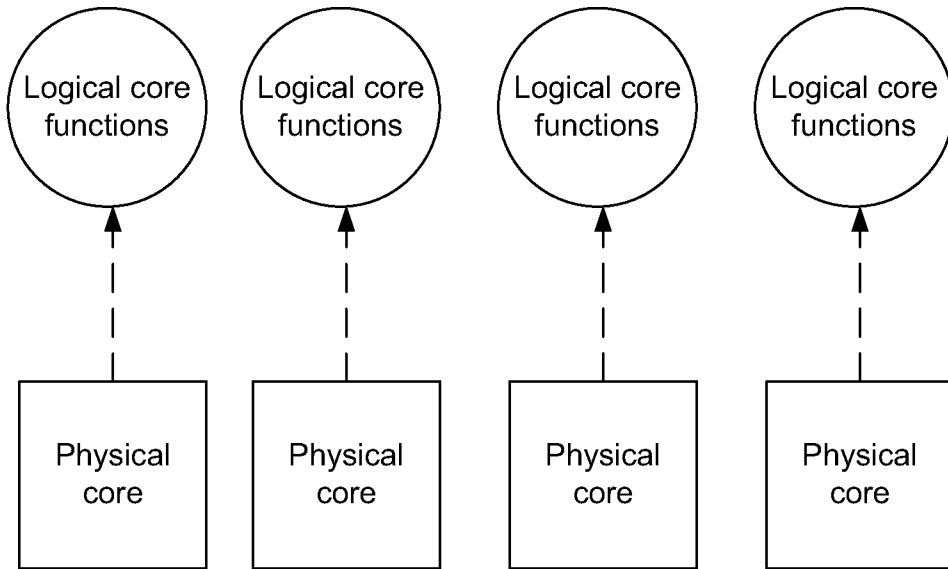


FIG. 11

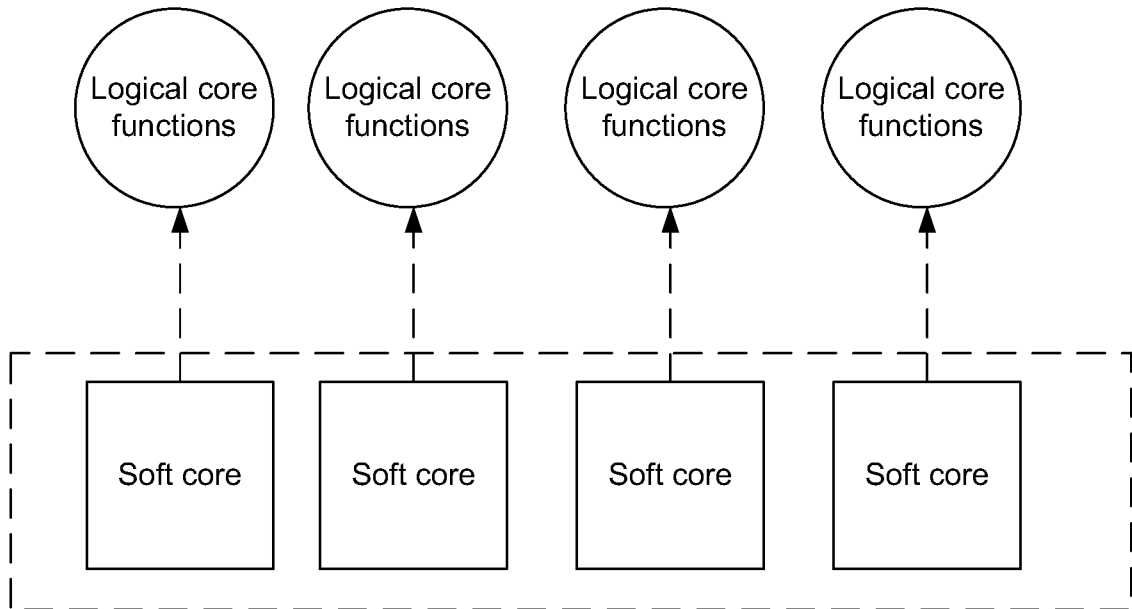


FIG. 12

14/24

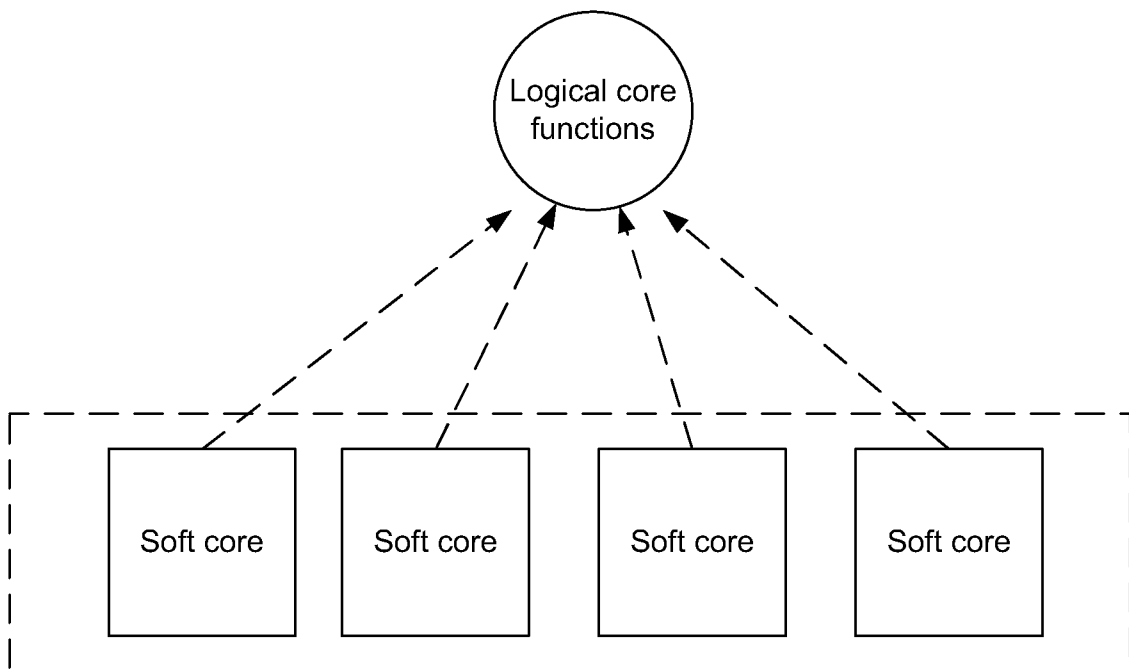


FIG. 13




FIG. 14

Virtual Cores Mode: Many Physical Cores to Many Logical Cores

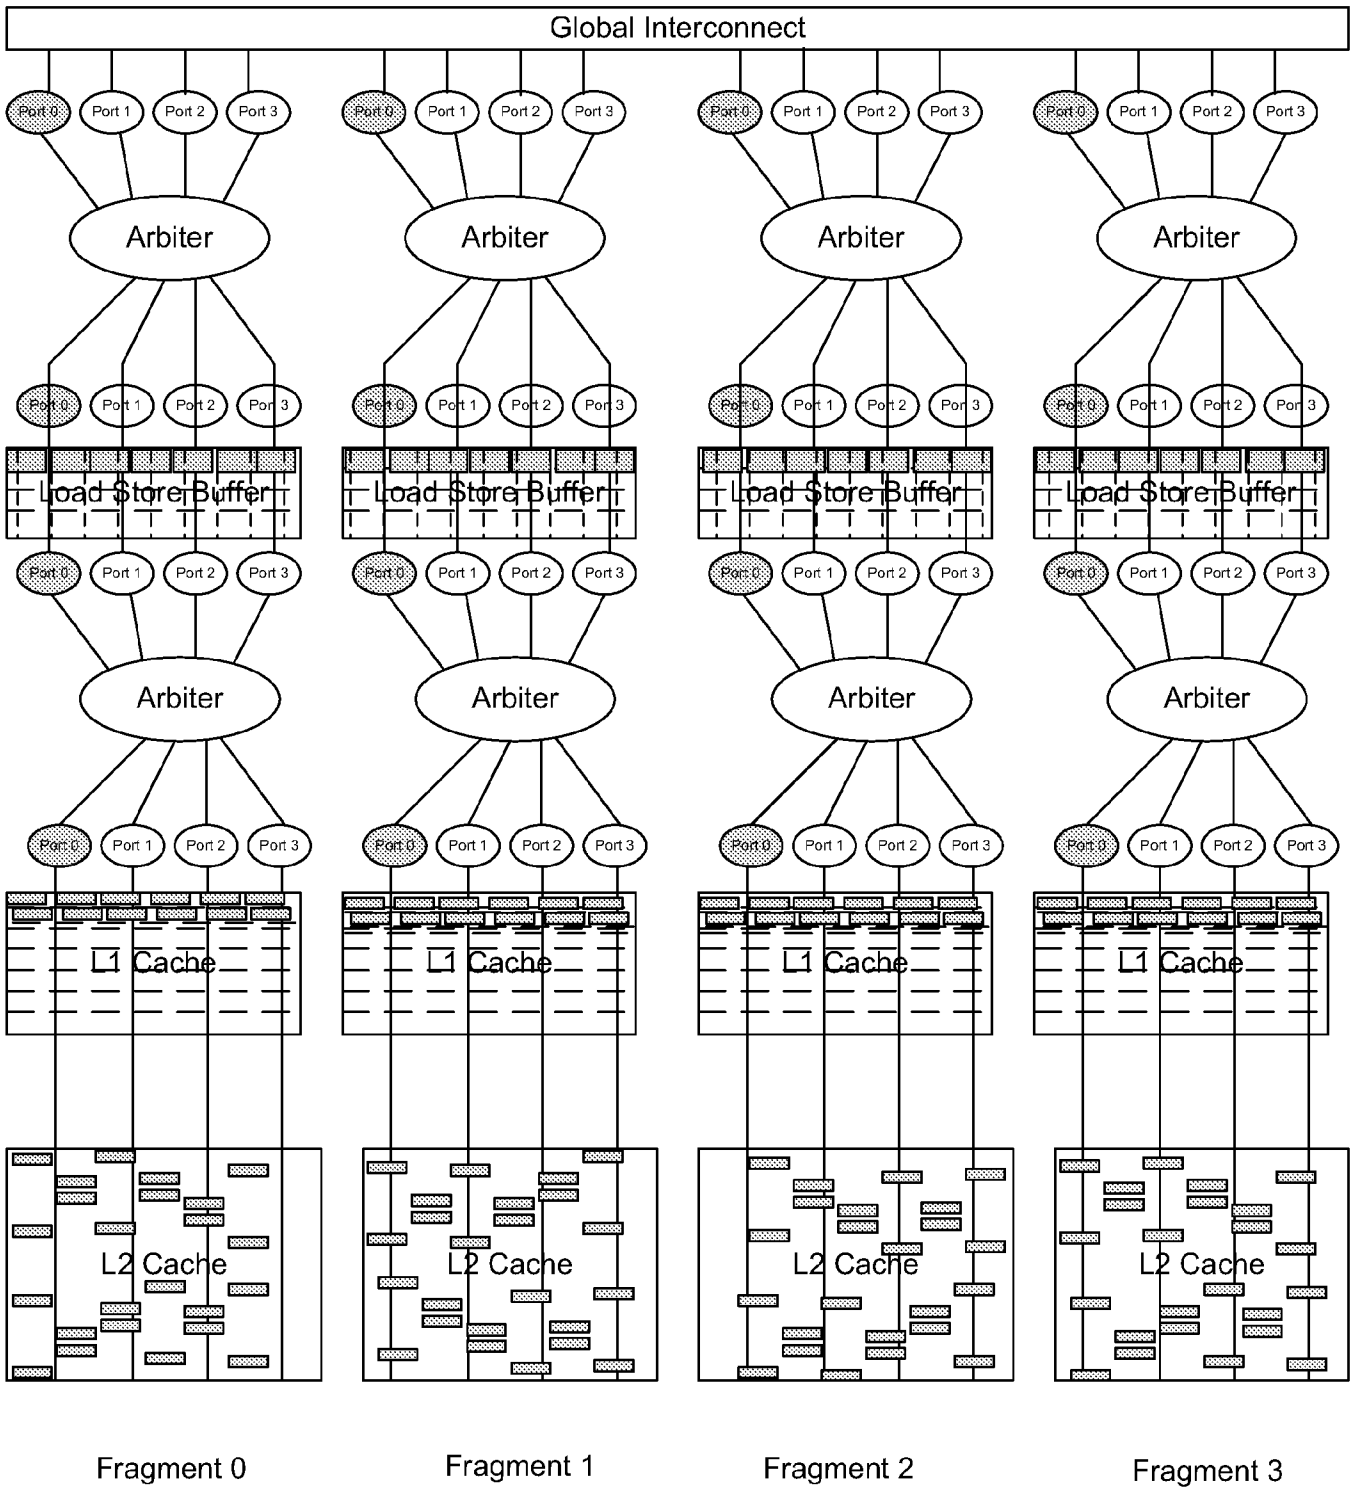


FIG. 15

Virtual Cores Mode: Many Physical Cores to Many Logical Cores




FIG. 16

Virtual Cores Mode: Many Soft Cores to Many Logical Cores

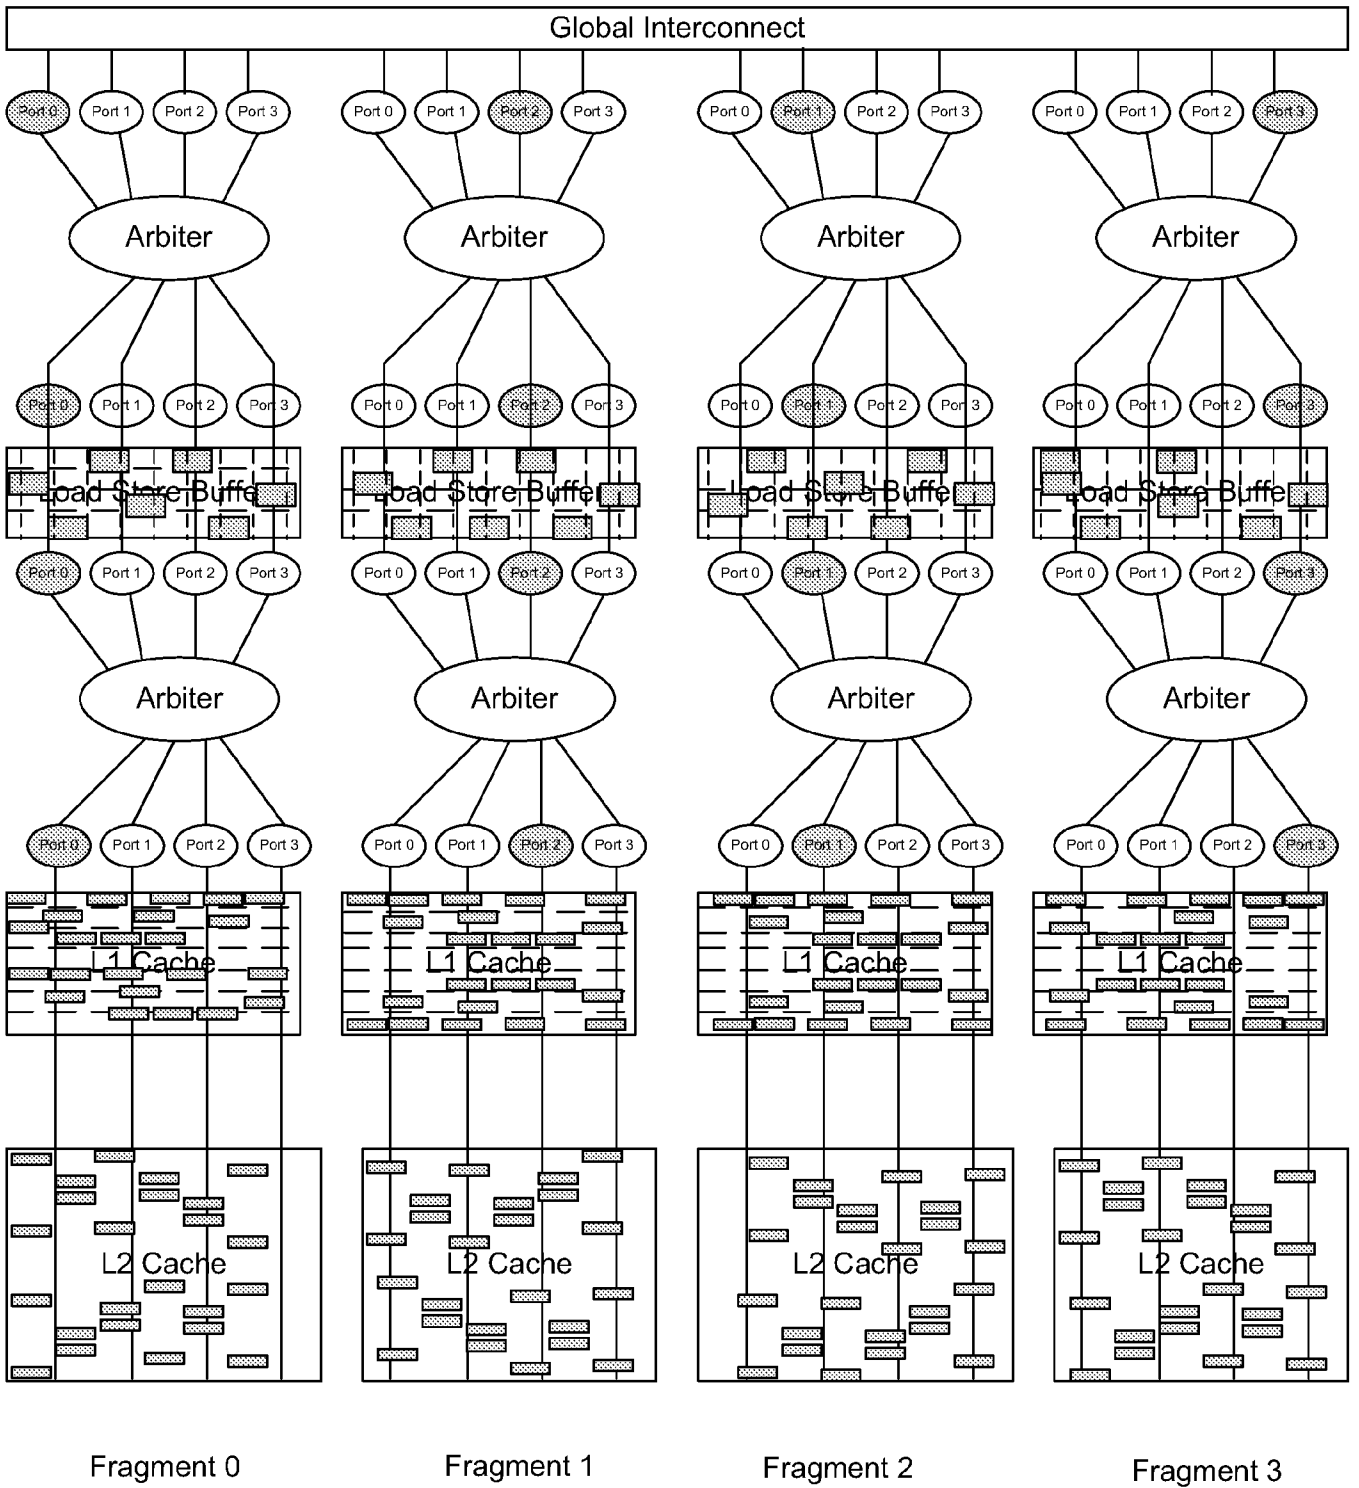


FIG. 17

Virtual Cores Mode: Many Soft Cores to One Logical Core

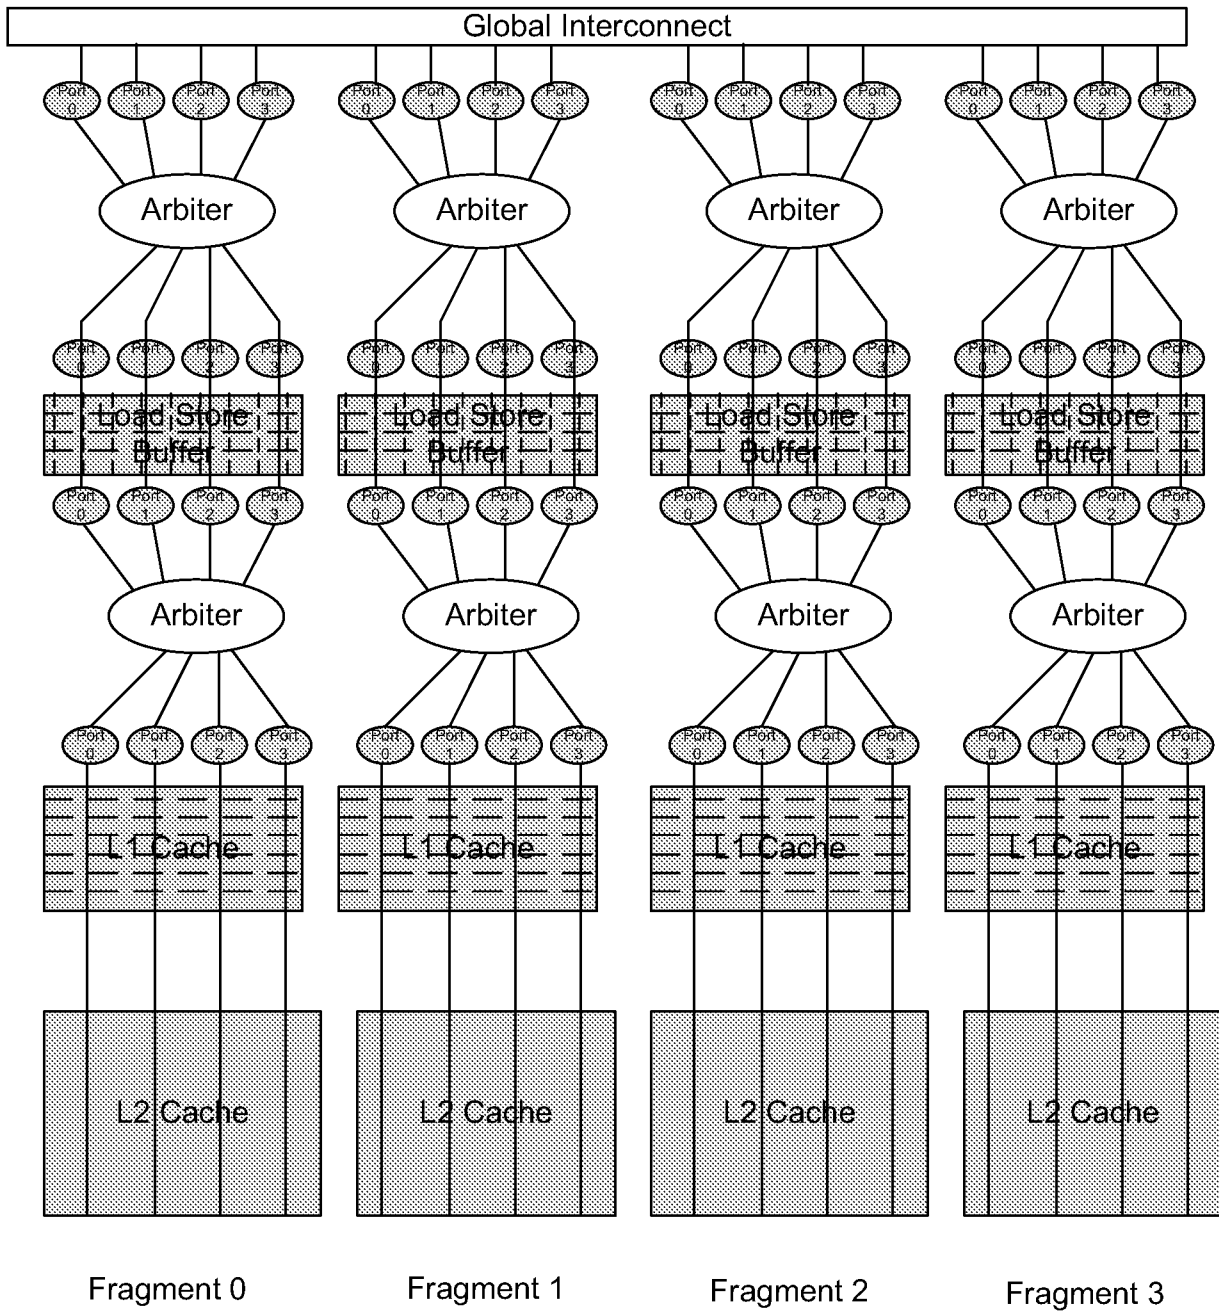


FIG. 18

Virtual Cores Mode:
Many Physical Cores
to Many Logical Cores

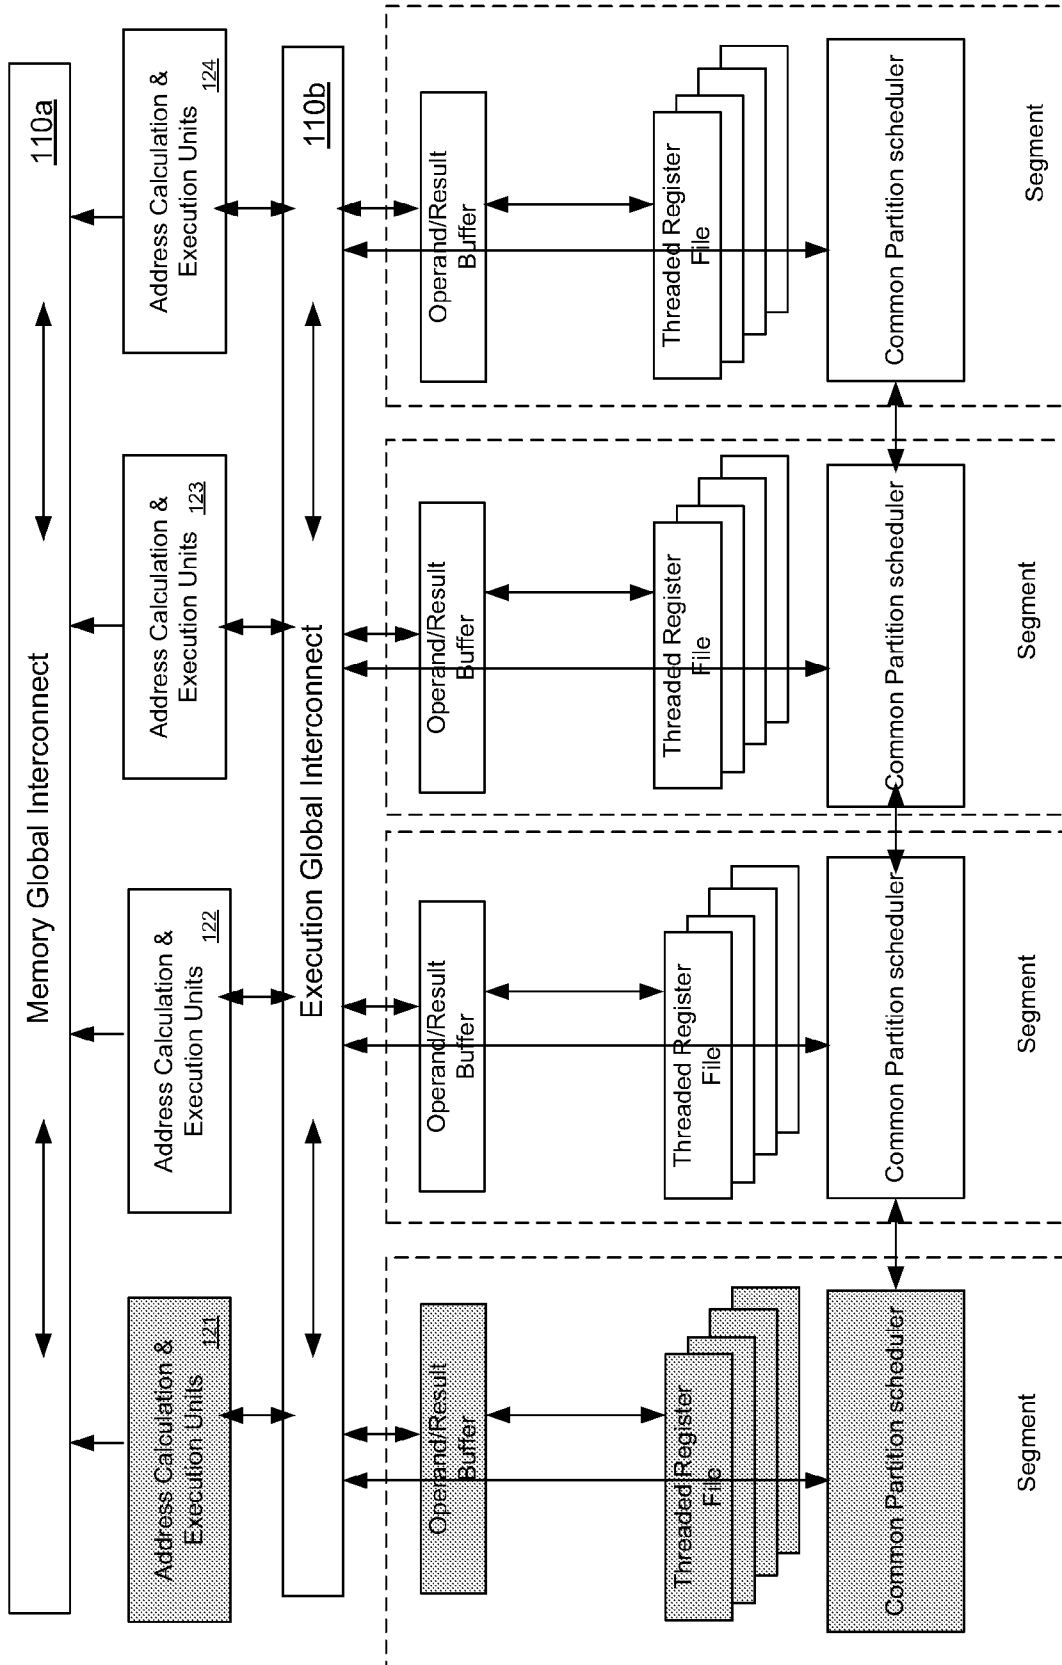


FIG. 19

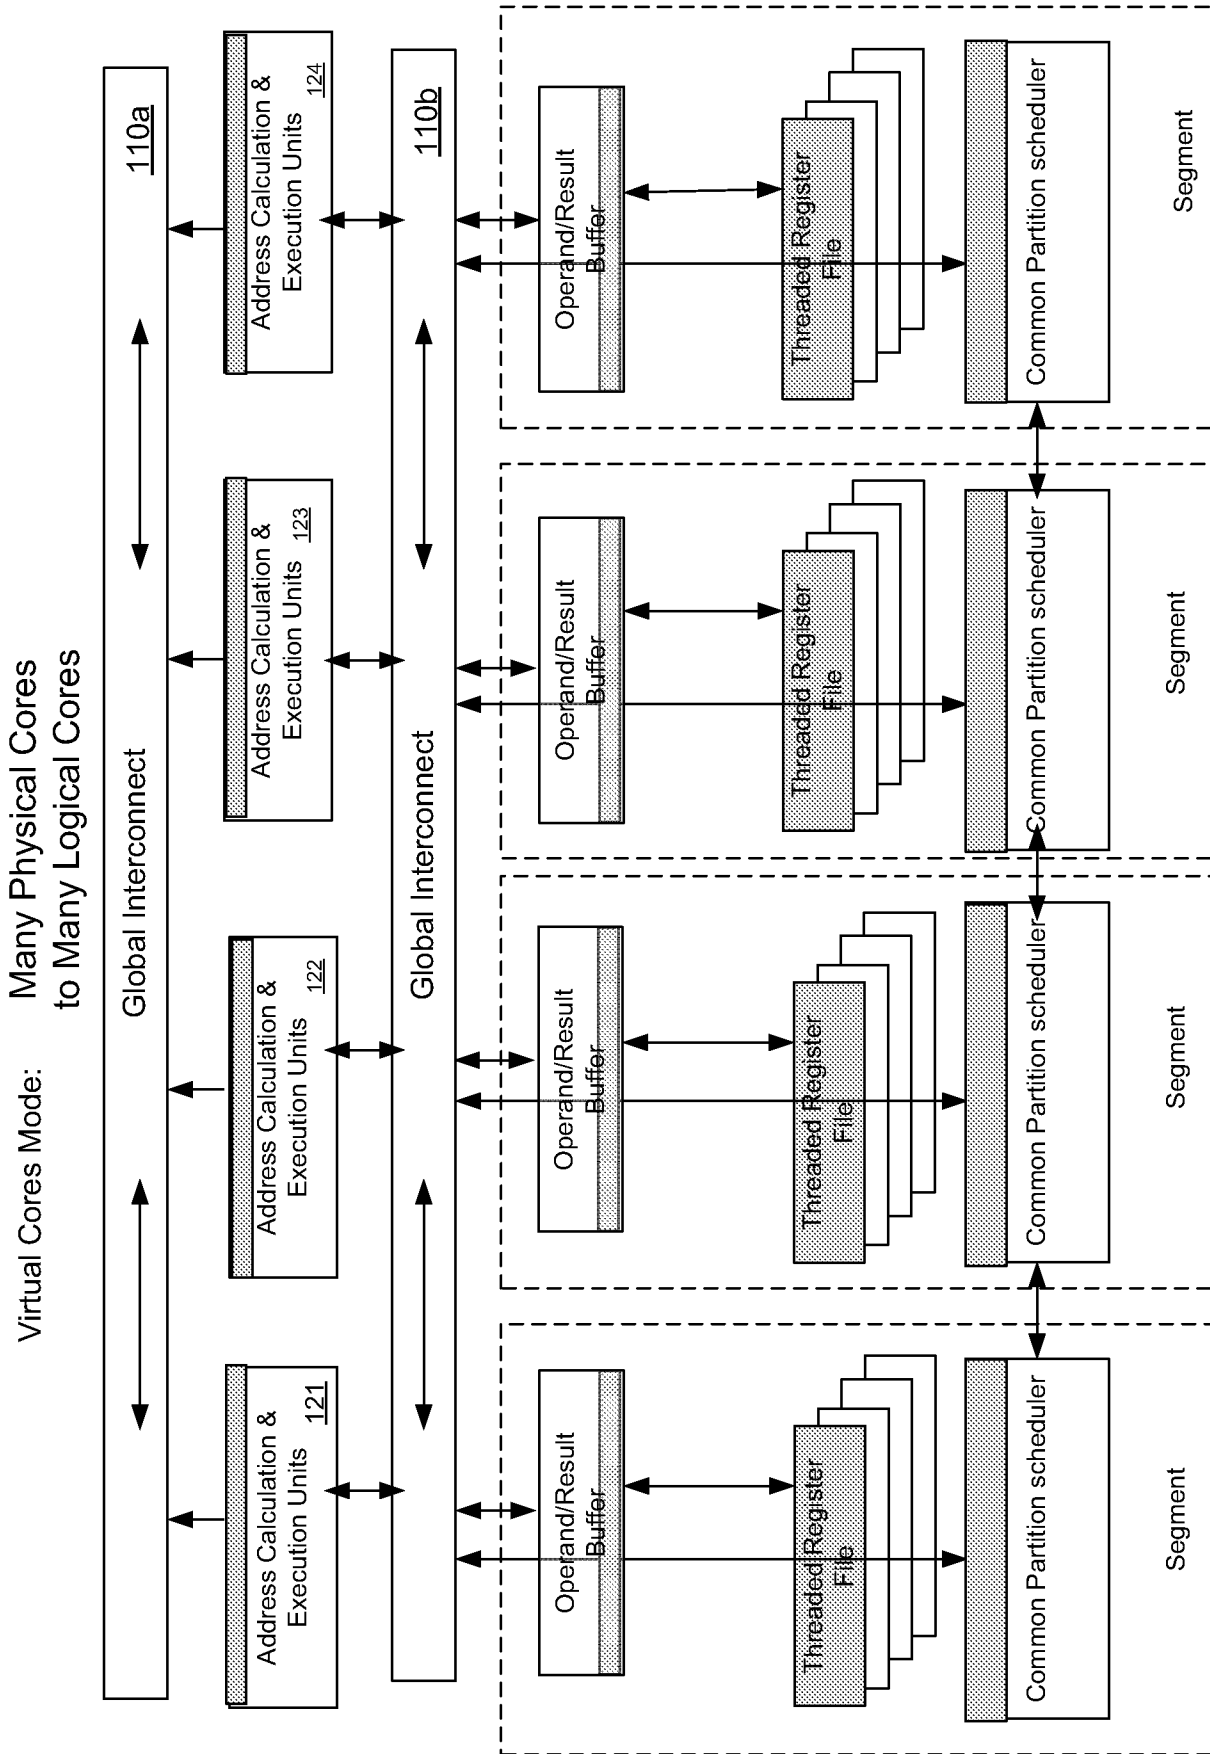


FIG. 20

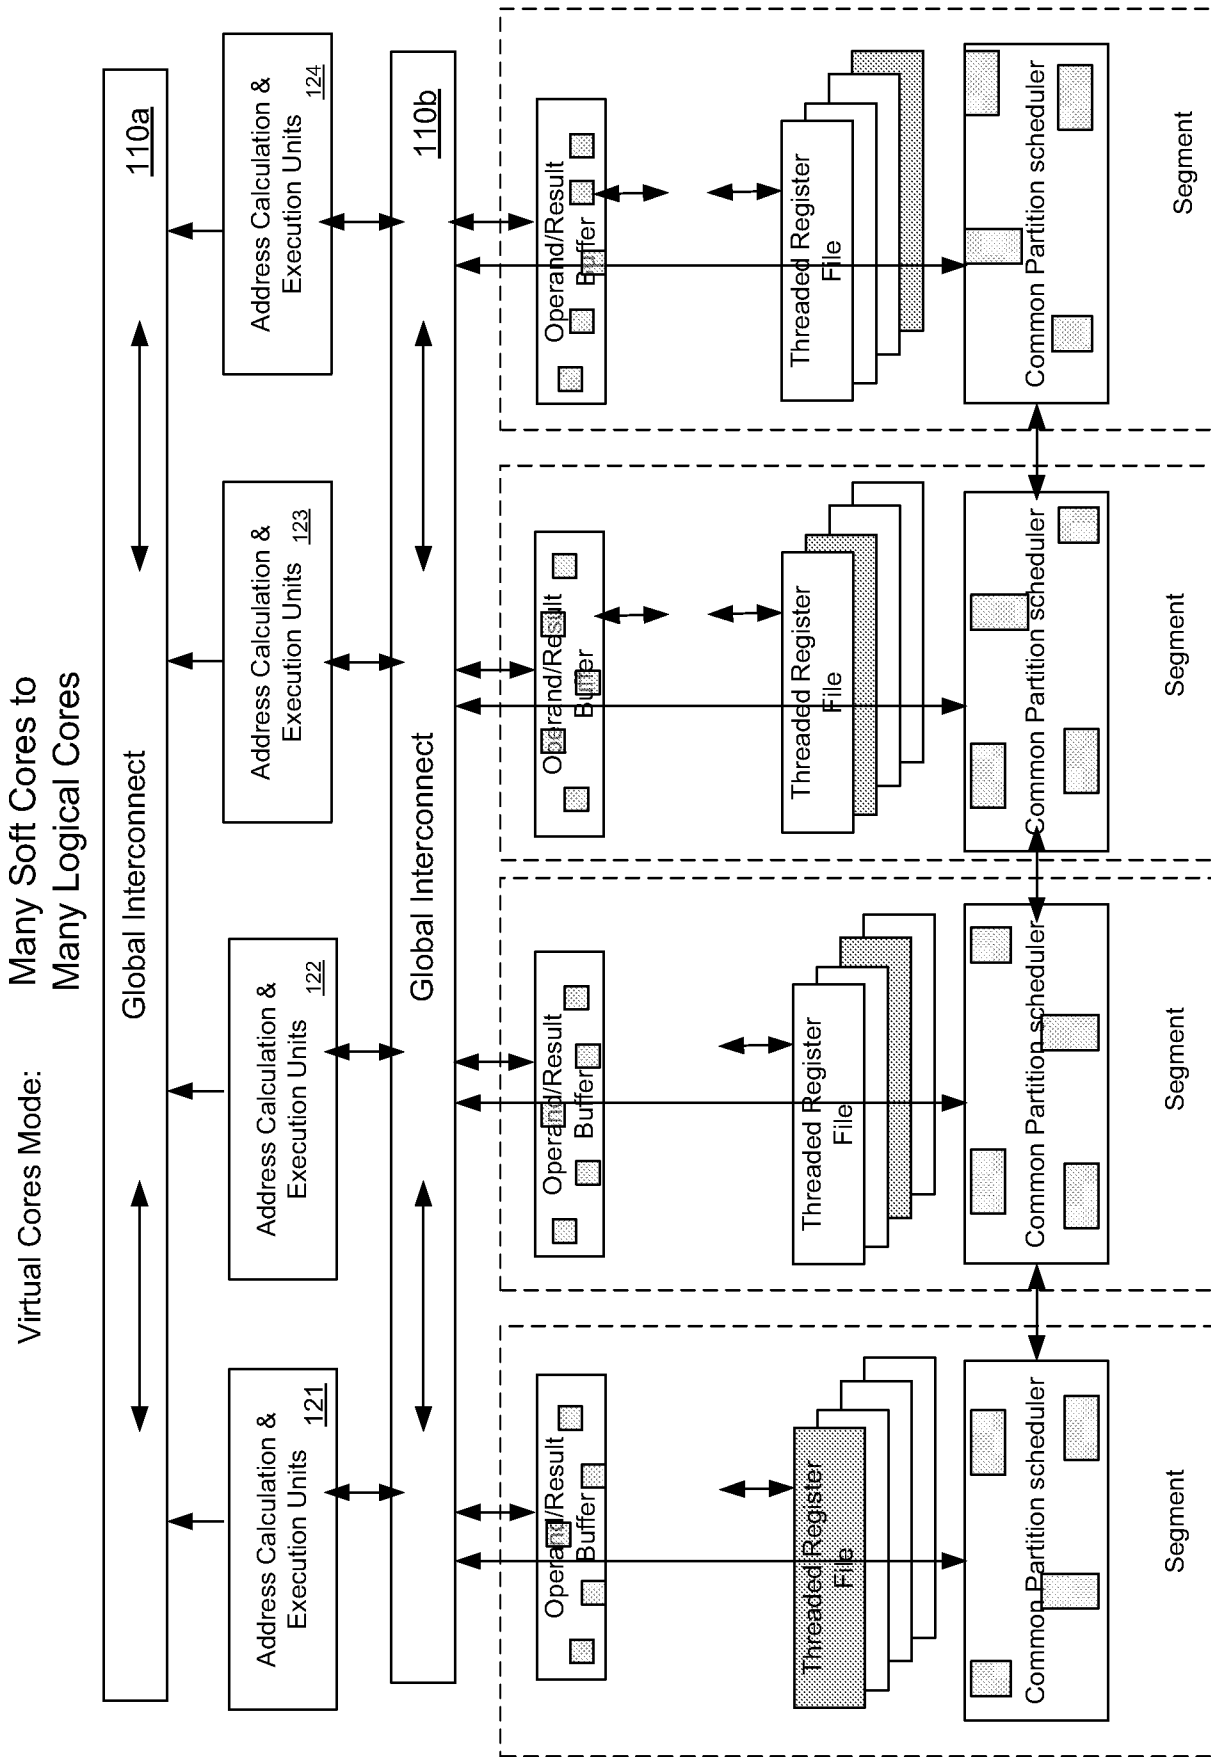


FIG. 21

Virtual Cores Mode:
Many Soft Cores to
One Logical Core

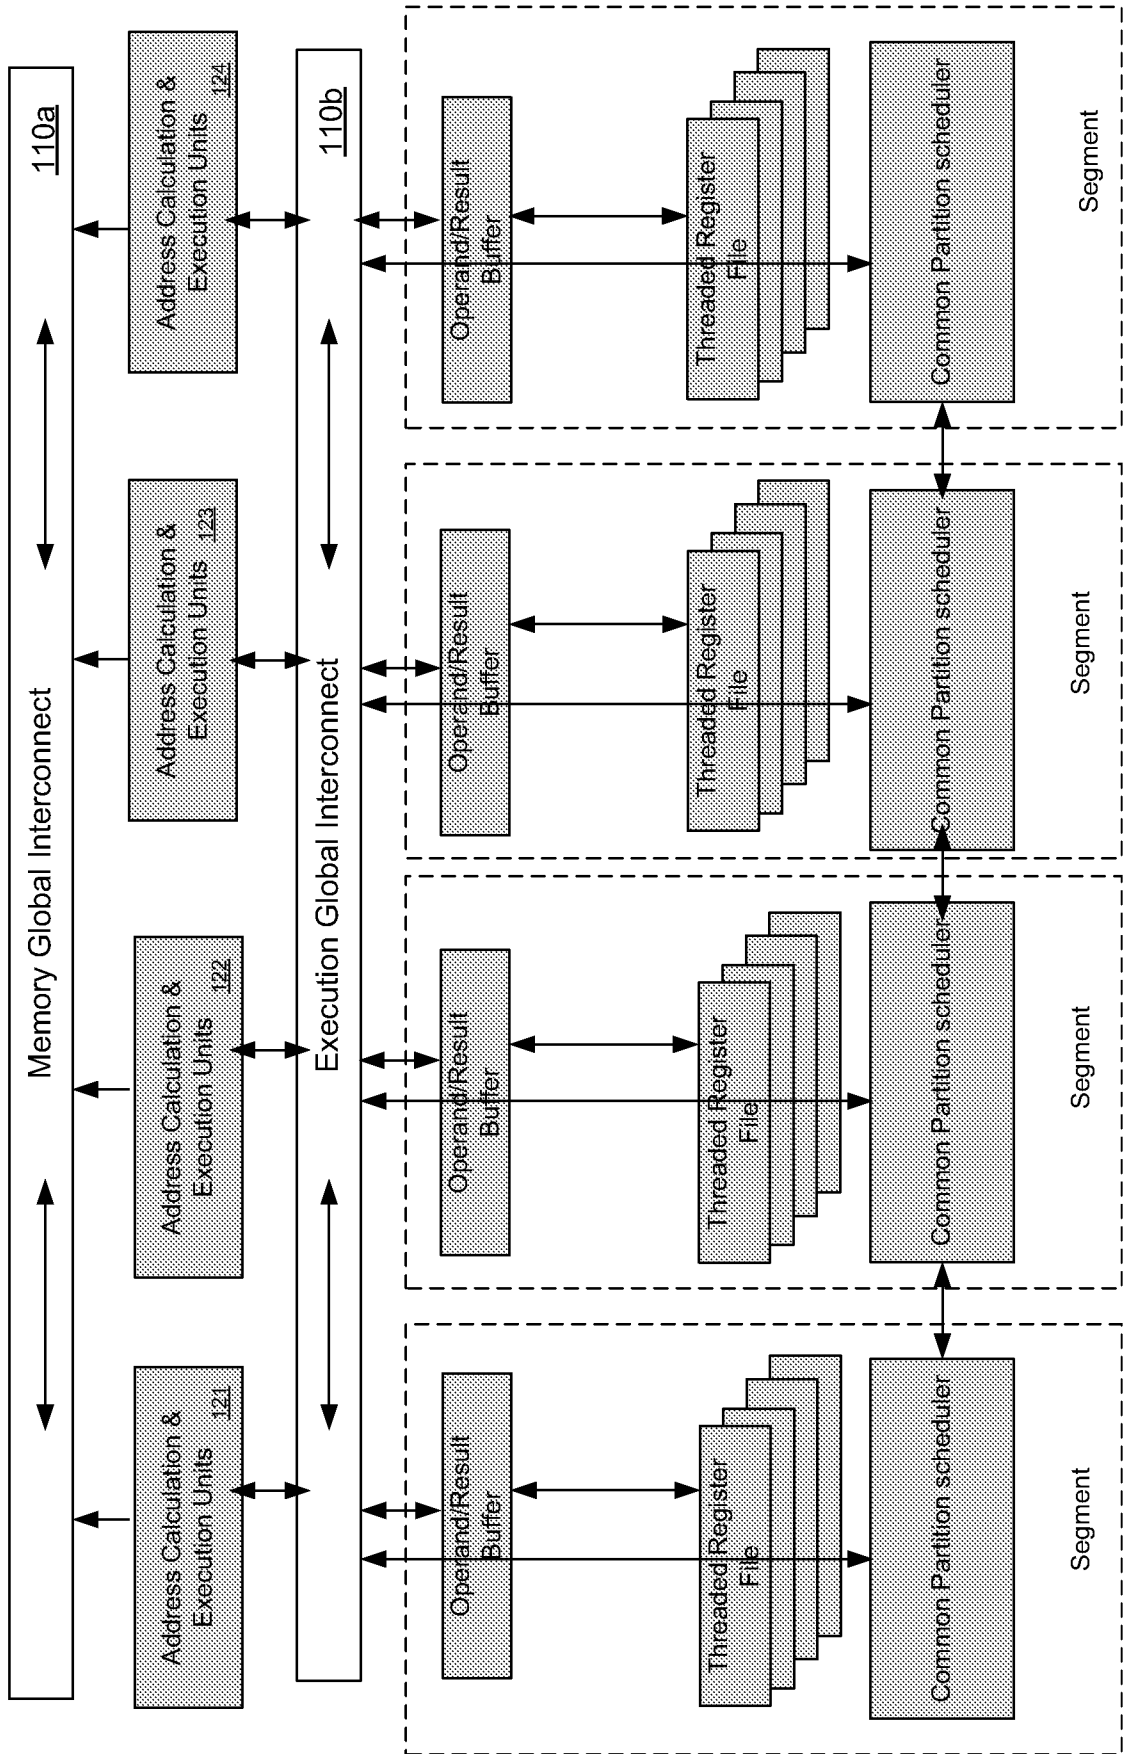


FIG. 22

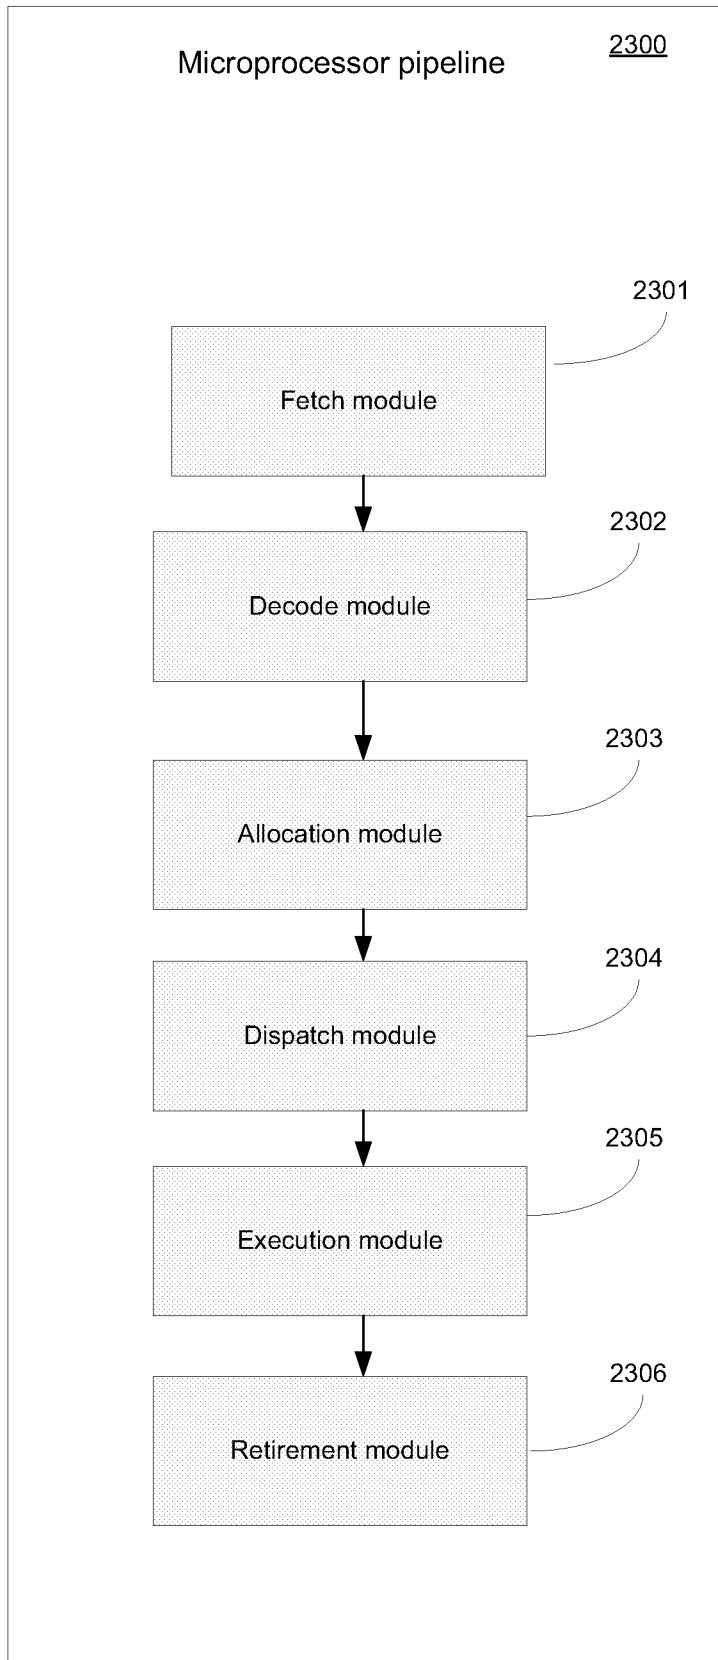


FIG. 23