



US 20110255542A1

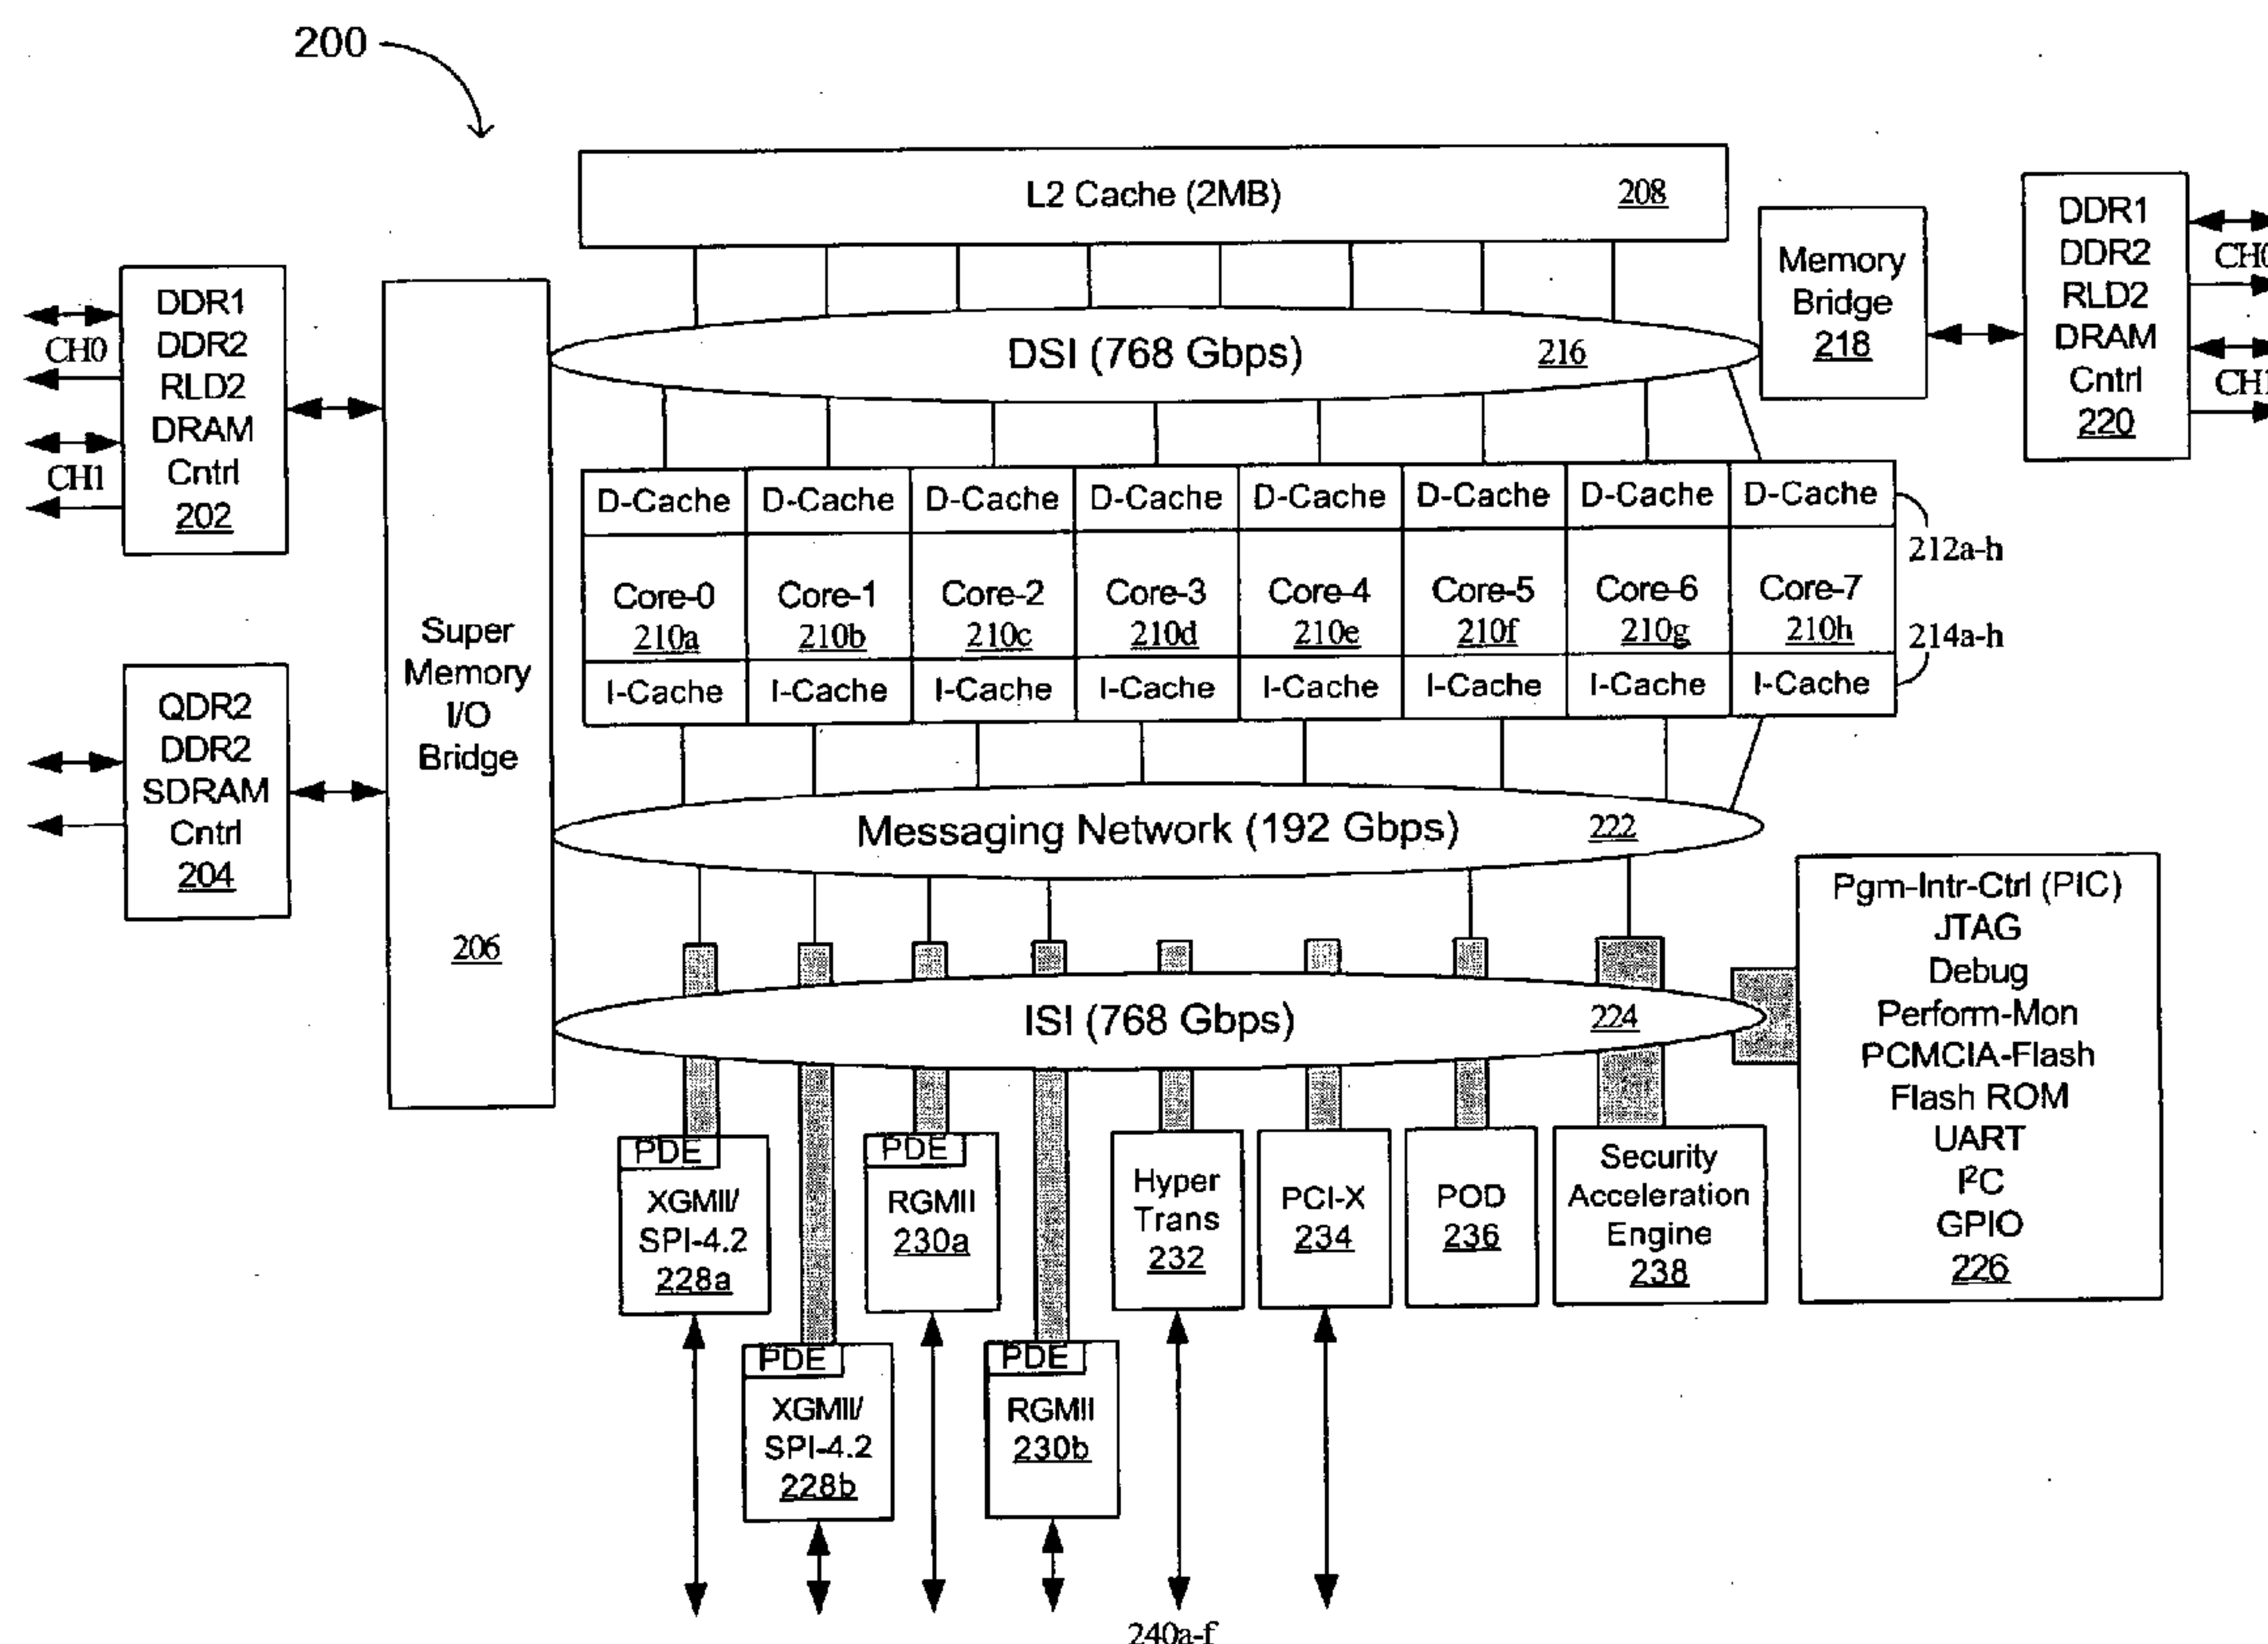
(19) **United States**(12) **Patent Application Publication**
HASS et al.(10) **Pub. No.: US 2011/0255542 A1**(43) **Pub. Date: Oct. 20, 2011**(54) **ADVANCED PROCESSOR WITH
MECHANISM FOR FAST PACKET QUEUING
OPERATIONS**

(60) Provisional application No. 60/490,236, filed on Jul. 25, 2003, provisional application No. 60/416,838, filed on Oct. 8, 2002.

Publication Classification(75) Inventors: **David T. HASS**, Santa Clara, CA (US); **Abbas RASHID**, Fremont, CA (US)(51) **Int. Cl.**
H04L 12/56 (2006.01)(52) **U.S. Cl.** **370/394**(73) Assignee: **NETLOGIC MICROSYSTEMS, INC.**(57) **ABSTRACT**(21) Appl. No.: **13/084,516**(22) Filed: **Apr. 11, 2011****Related U.S. Application Data**

(63) Continuation of application No. 10/930,455, filed on Aug. 31, 2004, now Pat. No. 7,924,828, which is a continuation-in-part of application No. 10/898,008, filed on Jul. 23, 2004, now Pat. No. 7,334,086, which is a continuation-in-part of application No. 10/682,579, filed on Oct. 8, 2003, now abandoned.

An advanced processor comprises a plurality of multi-threaded processor cores each having a data cache and instruction cache. A data switch interconnect is coupled to each of the processor cores and configured to pass information among the processor cores. A messaging network is coupled to each of the processor cores and a plurality of communication ports. In one aspect of an embodiment of the invention, the data switch interconnect is coupled to each of the processor cores by its respective data cache, and the messaging network is coupled to each of the processor cores by its respective message station. Advantages of the invention include the ability to provide high bandwidth communications between computer systems and memory in an efficient and cost-effective manner.



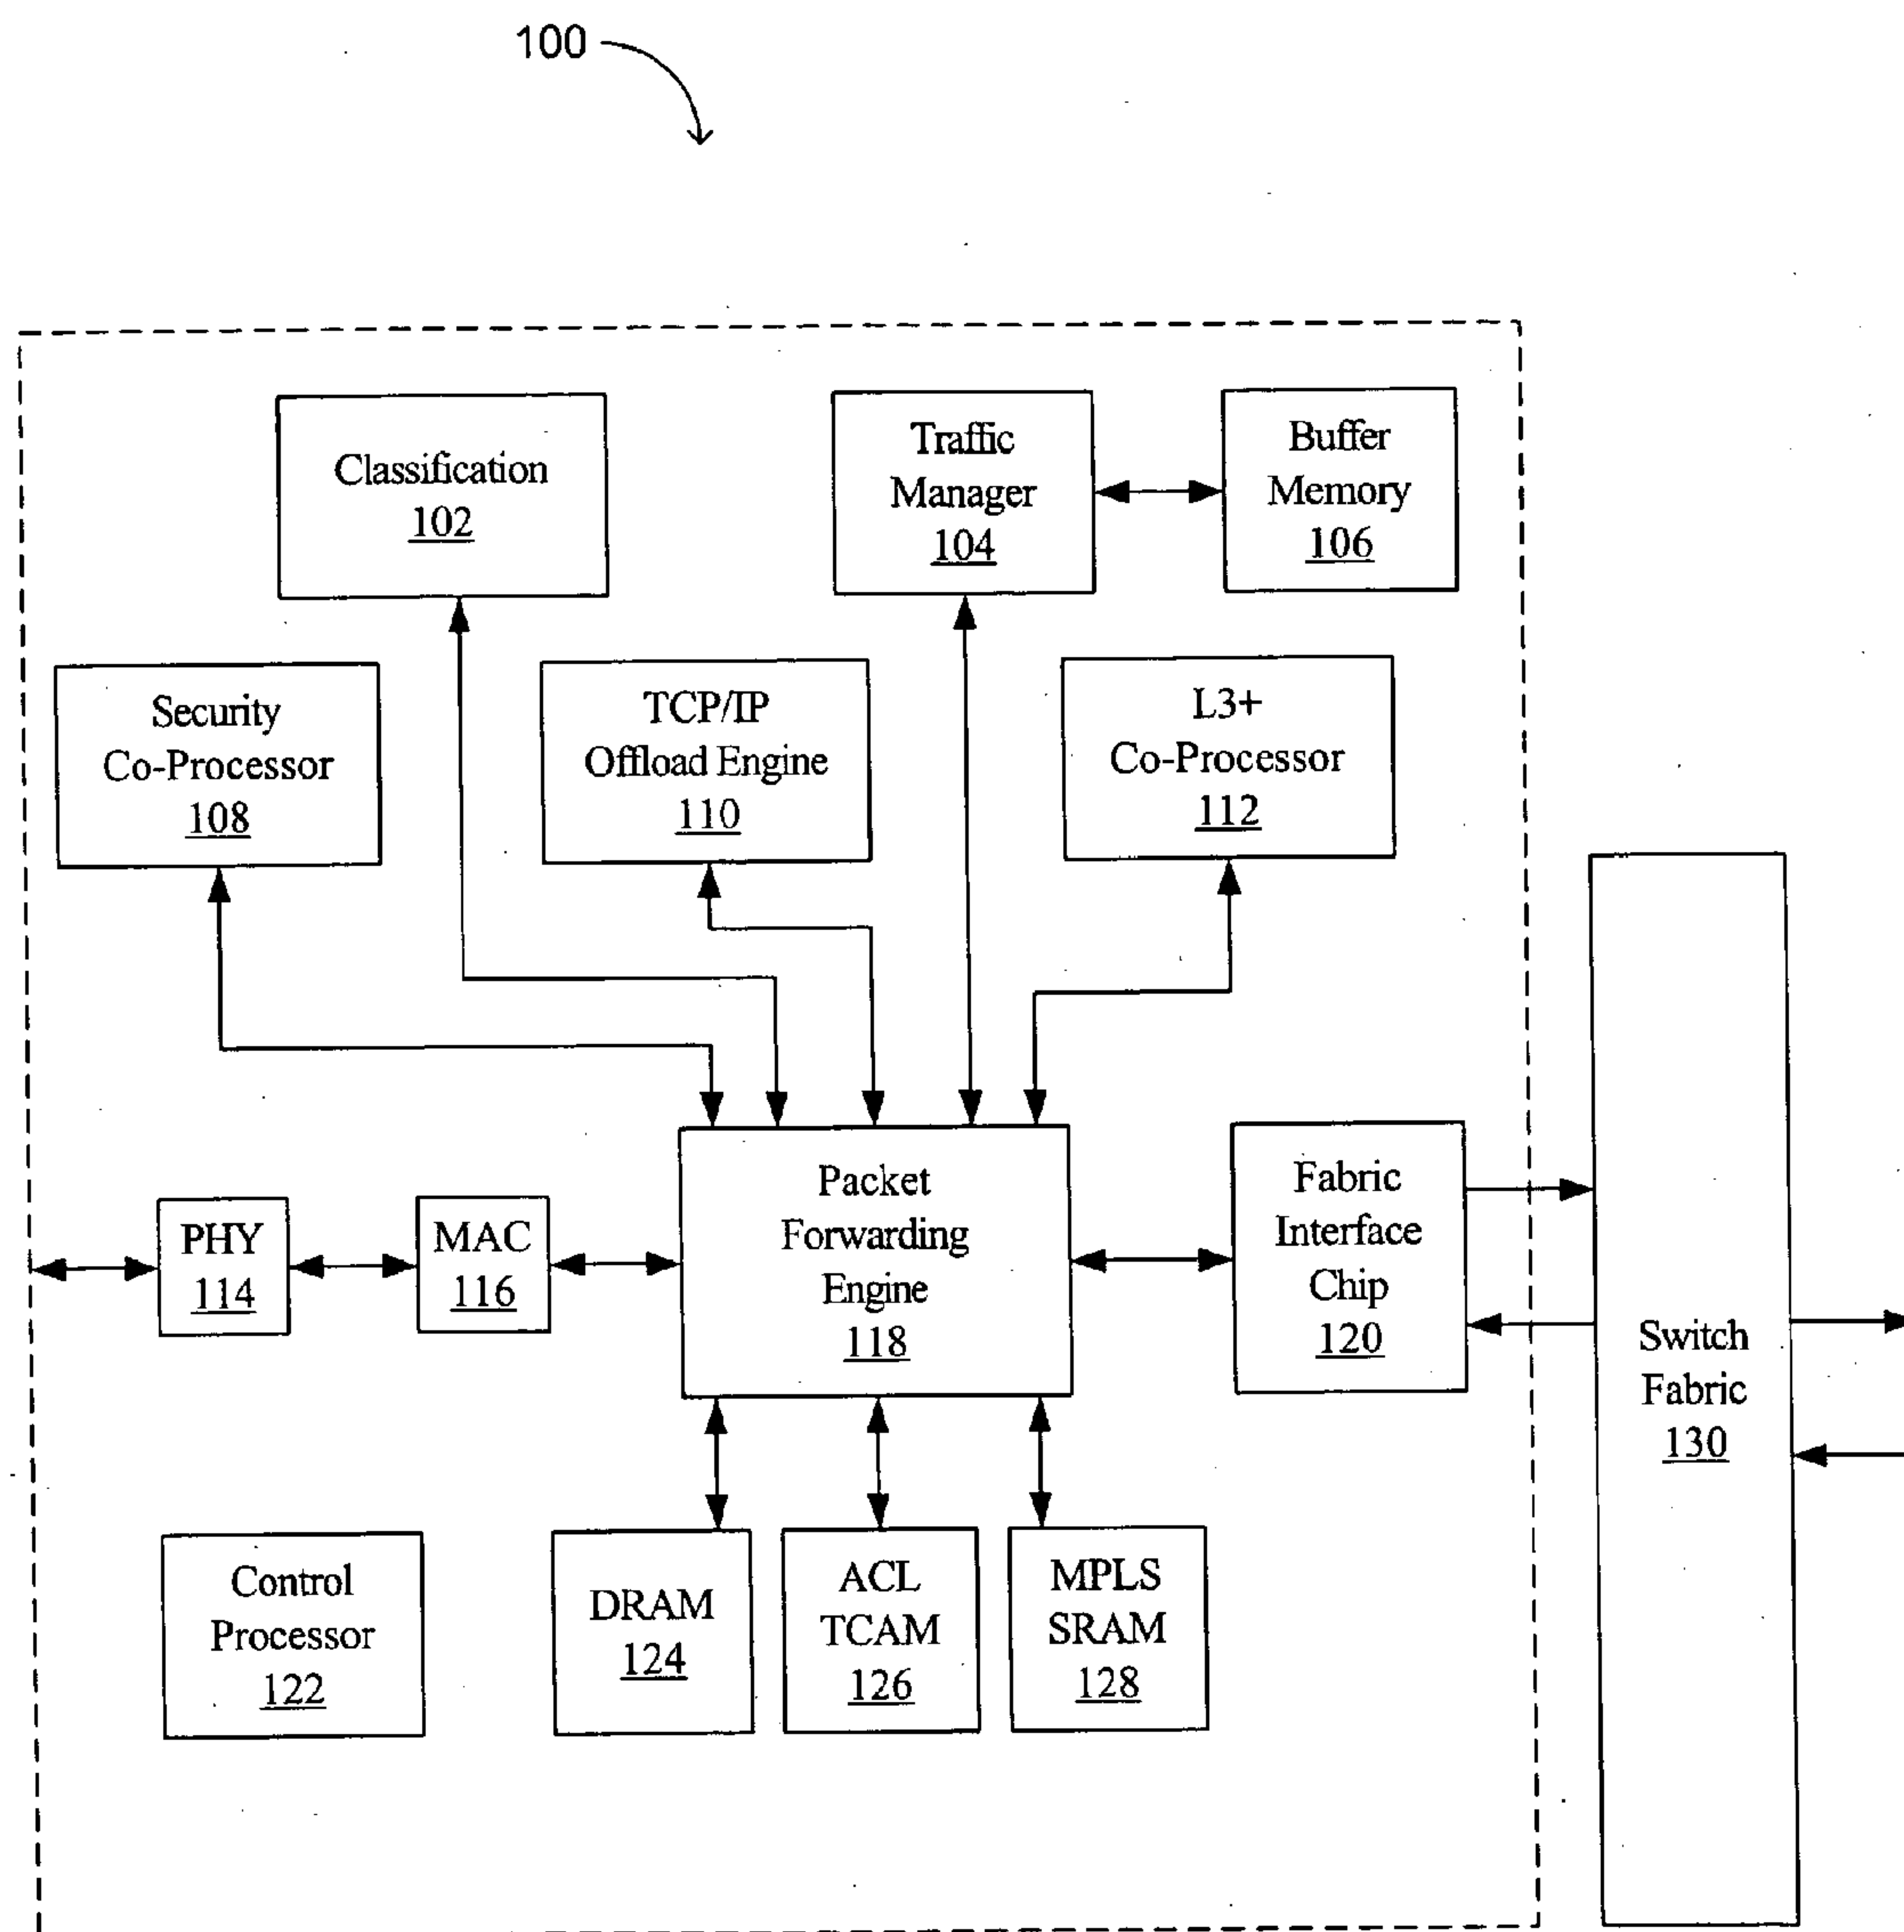


FIG. 1 (conventional)

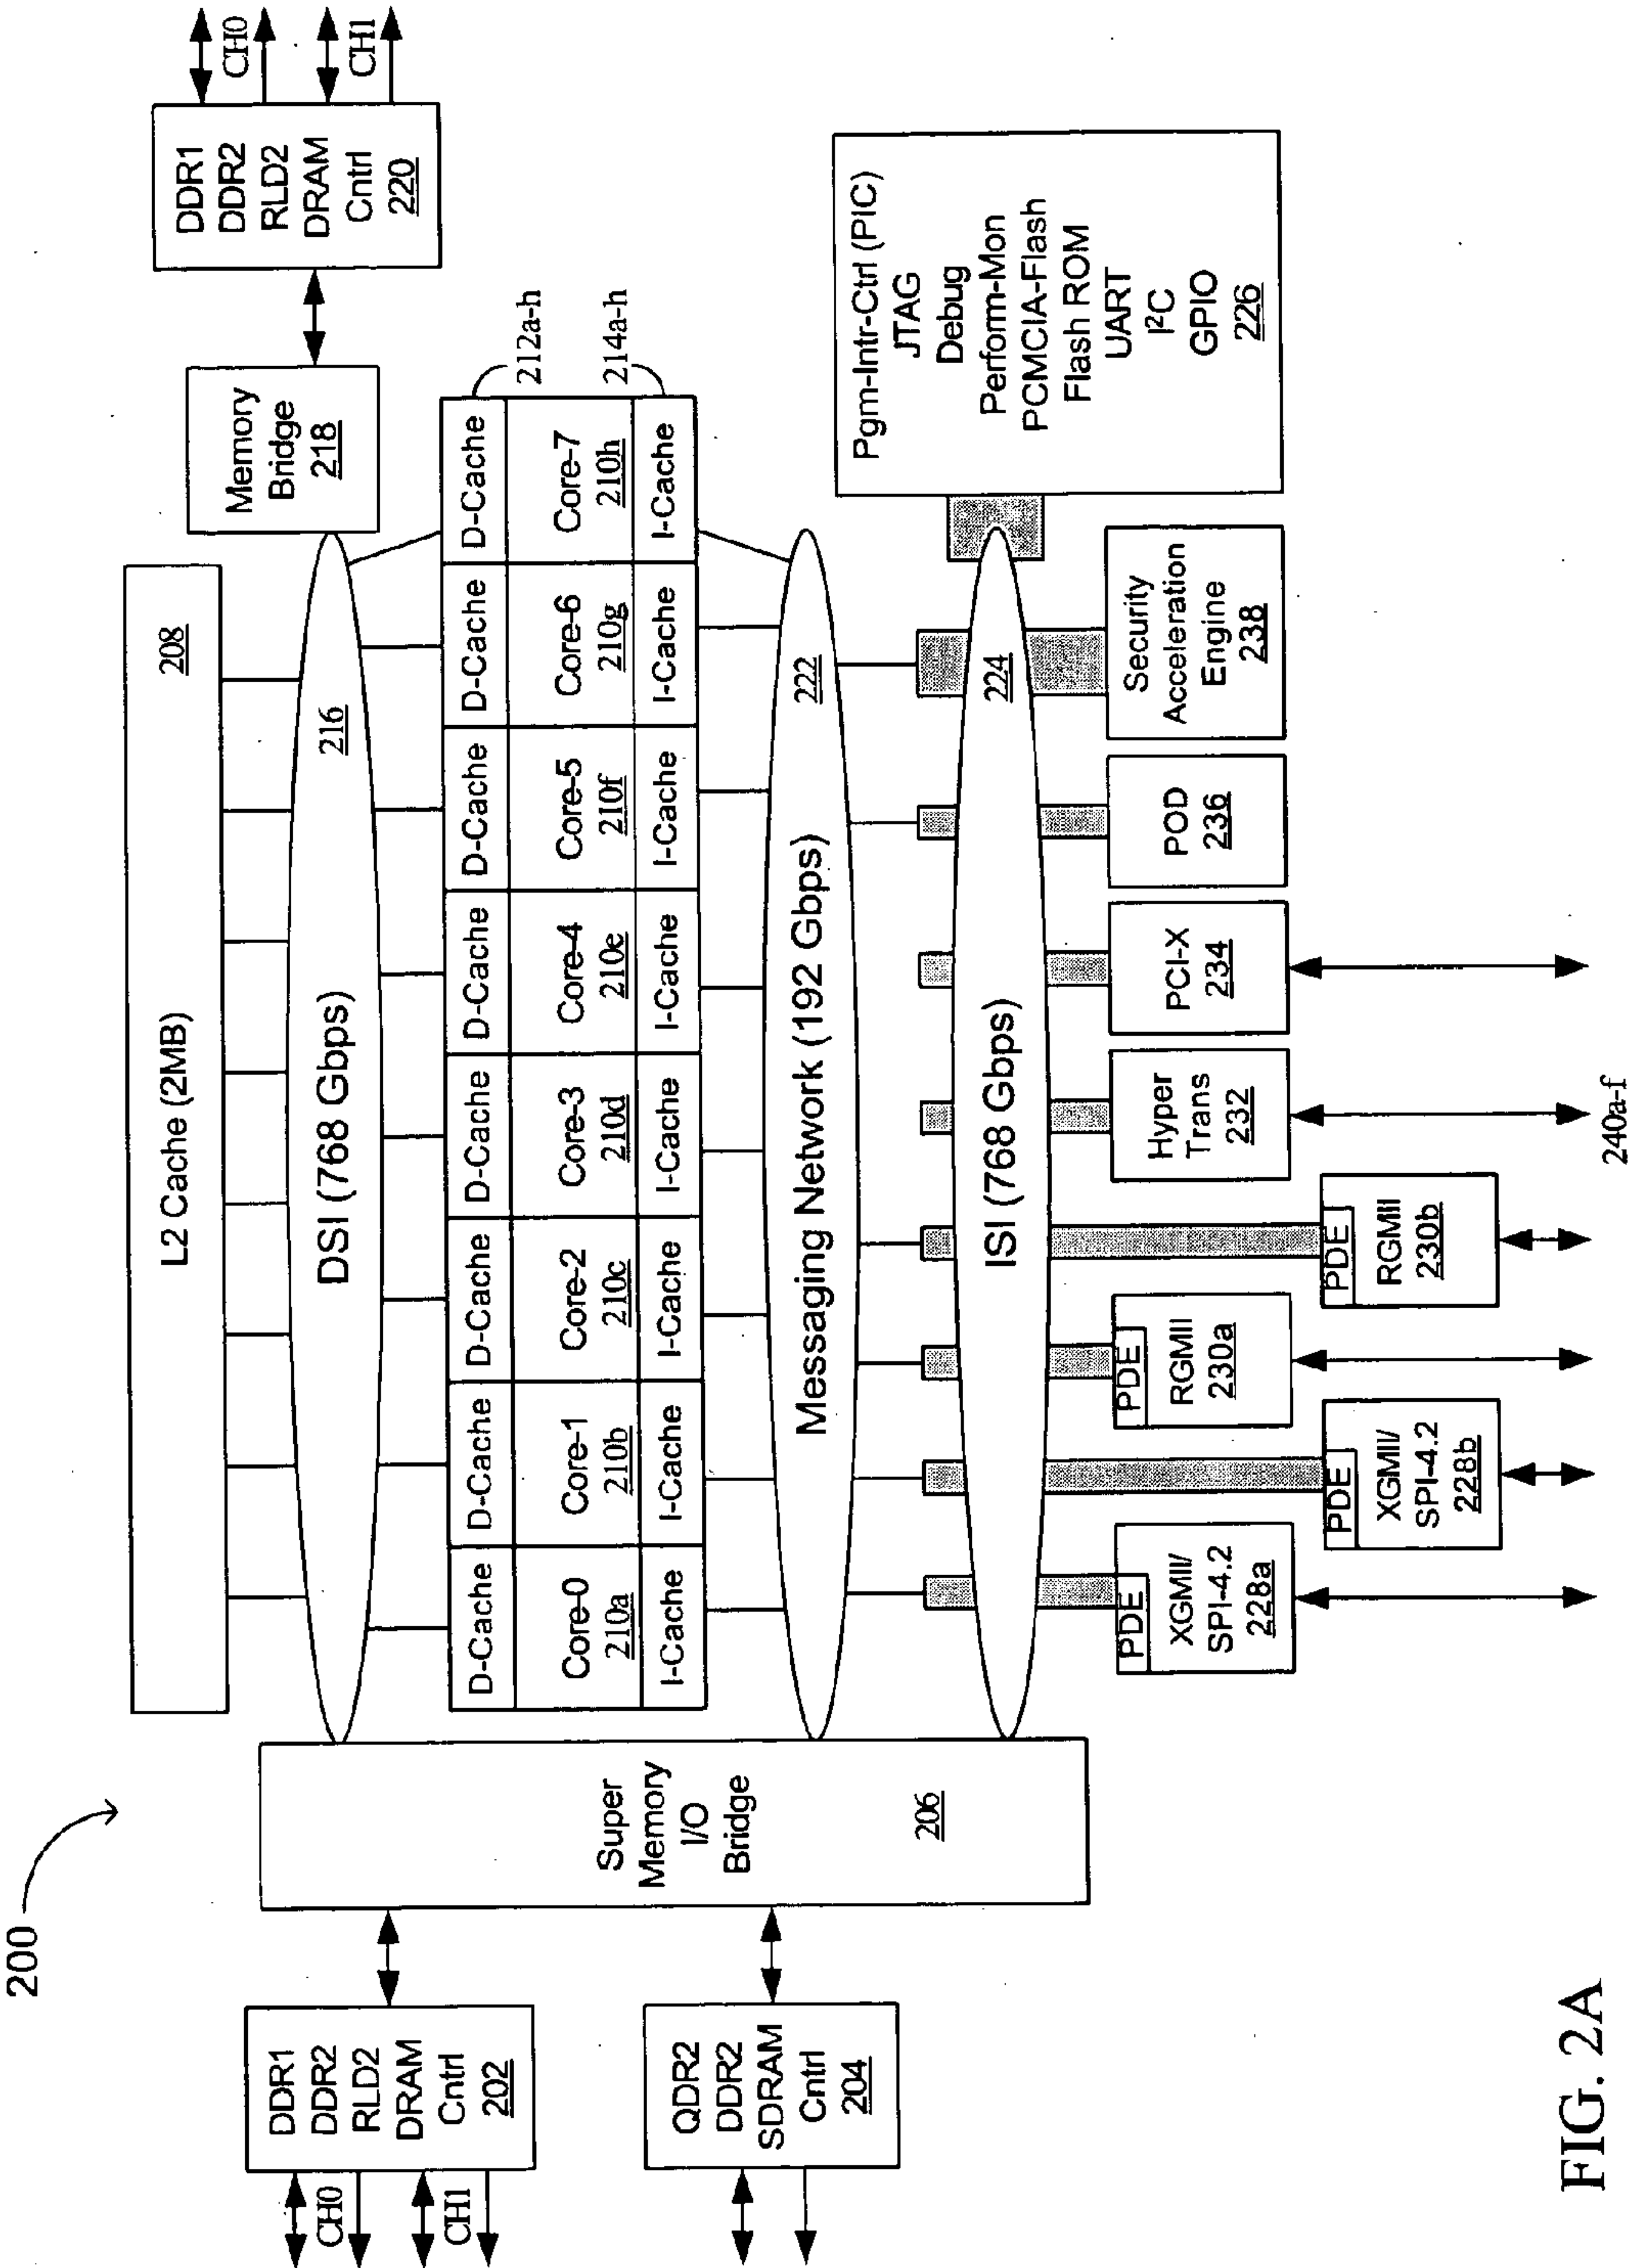


FIG. 2A

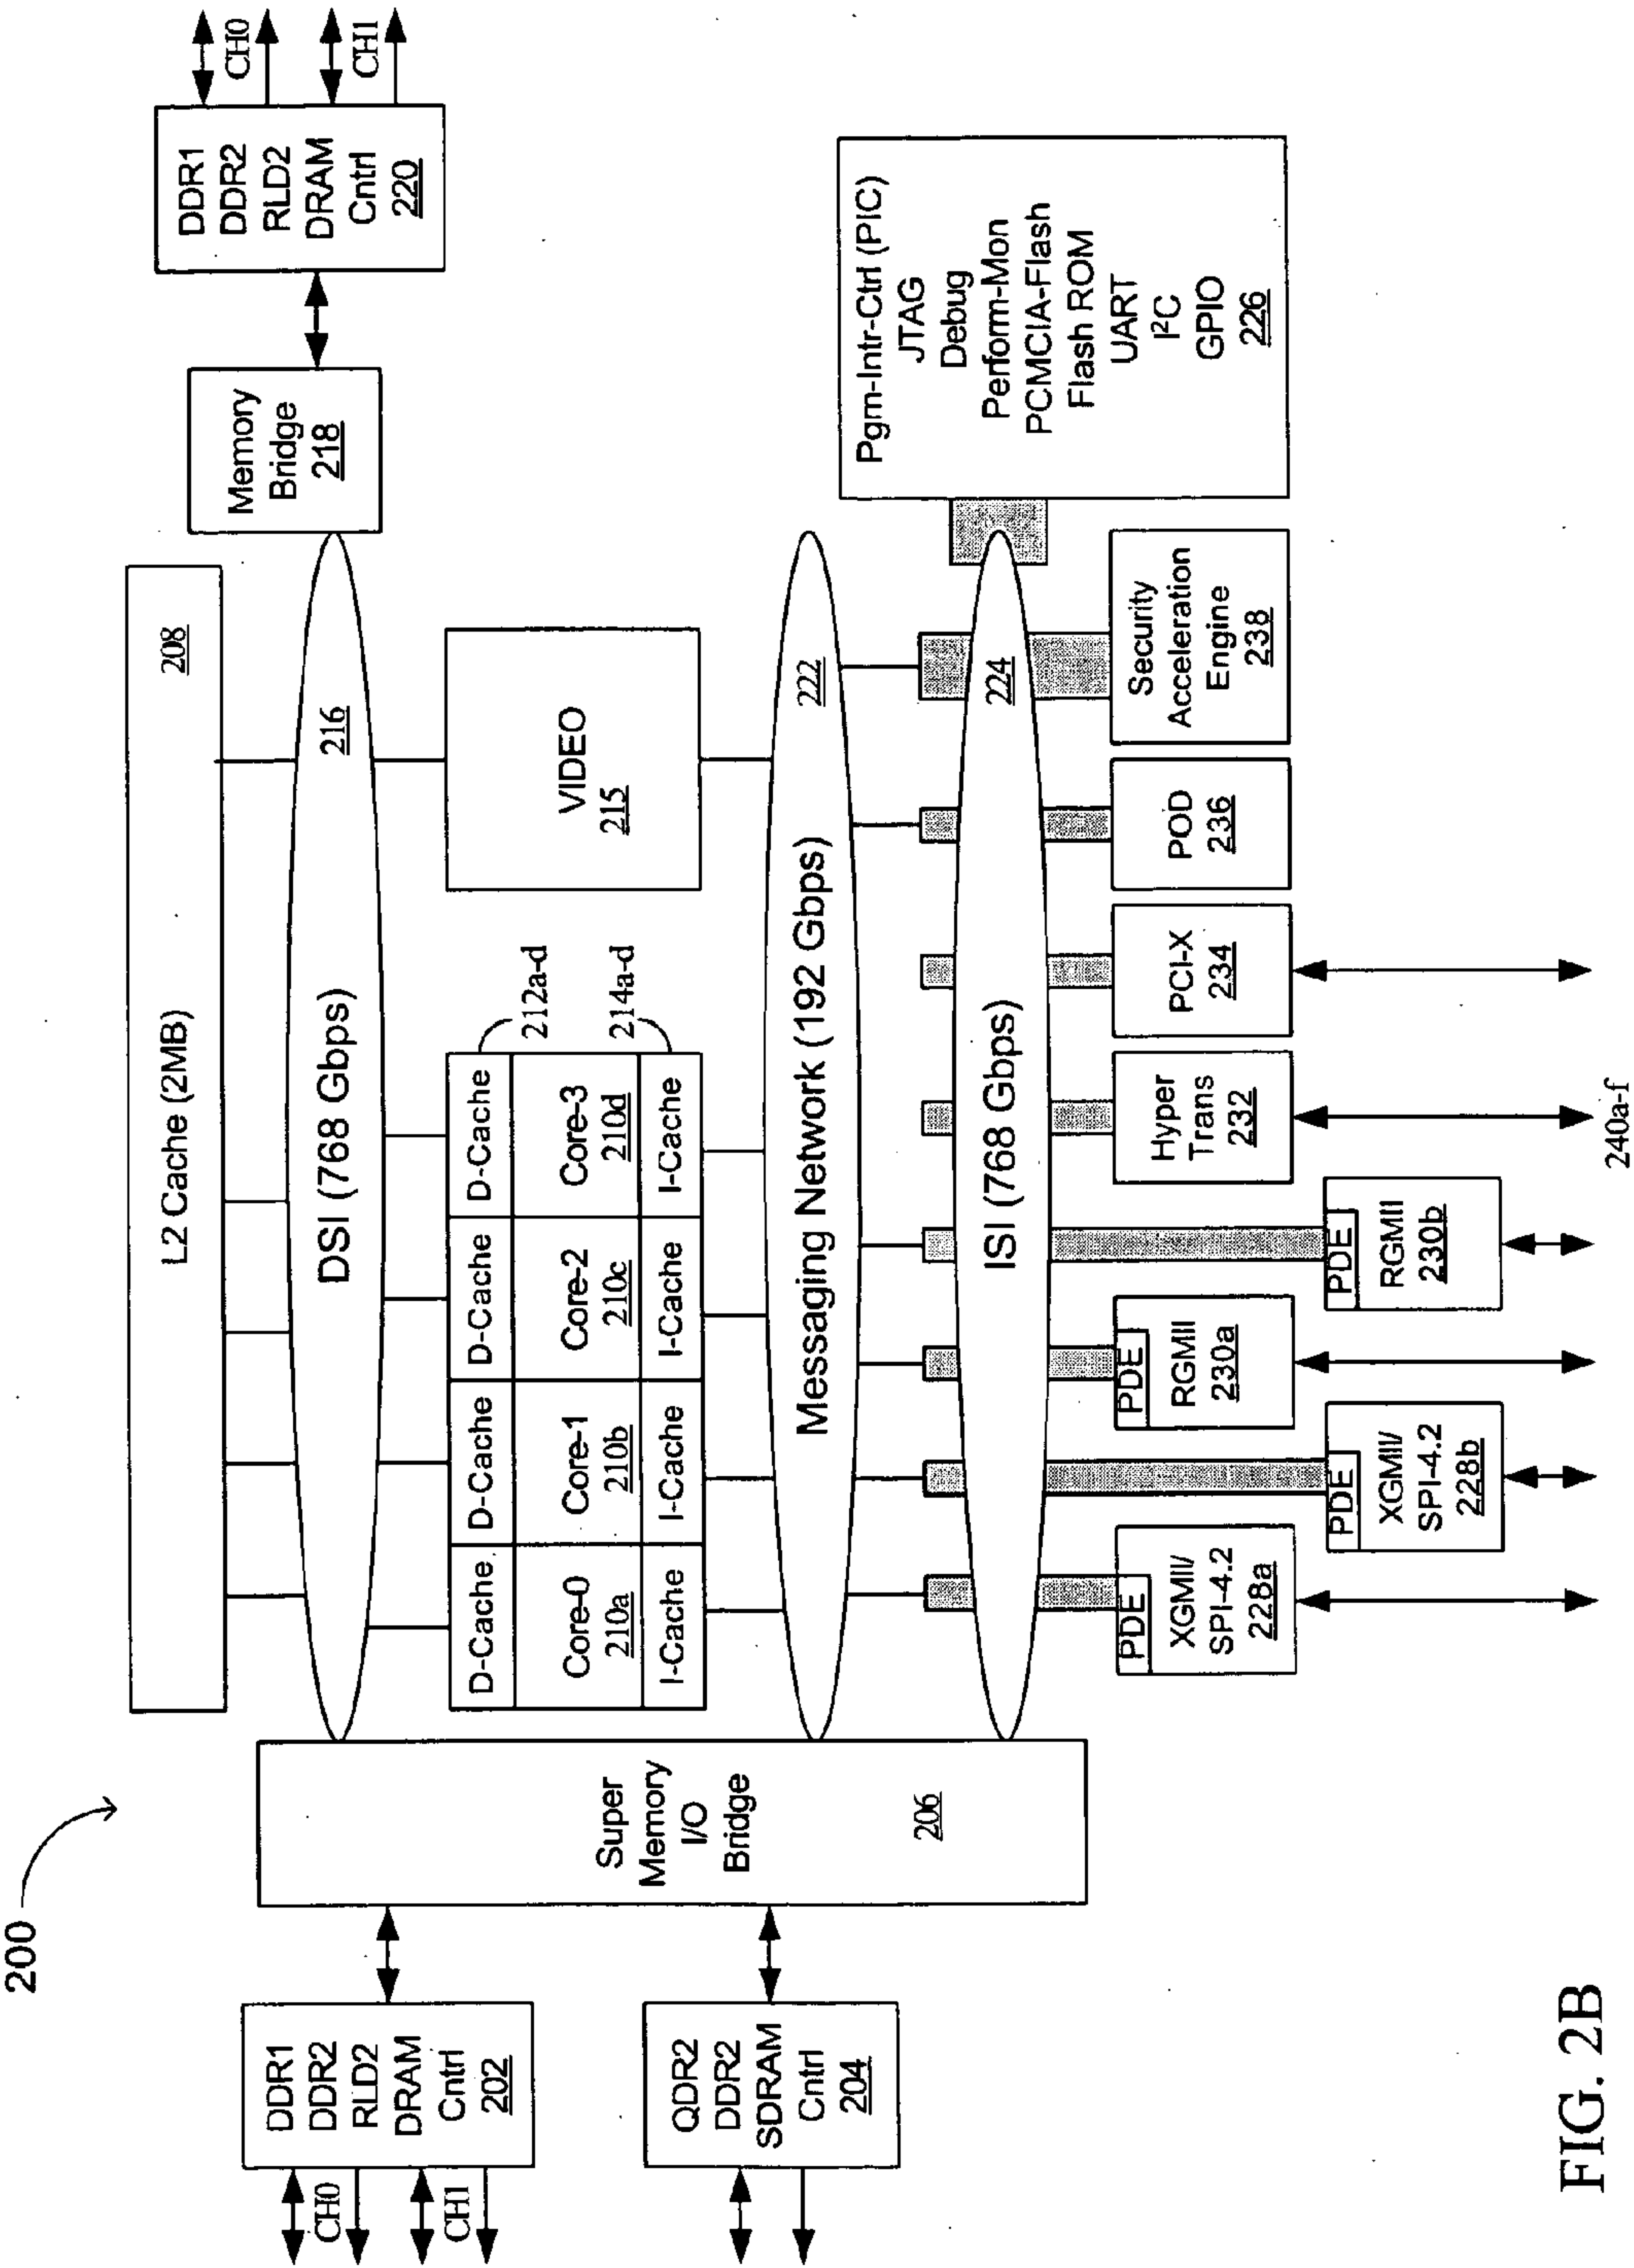


FIG. 2B

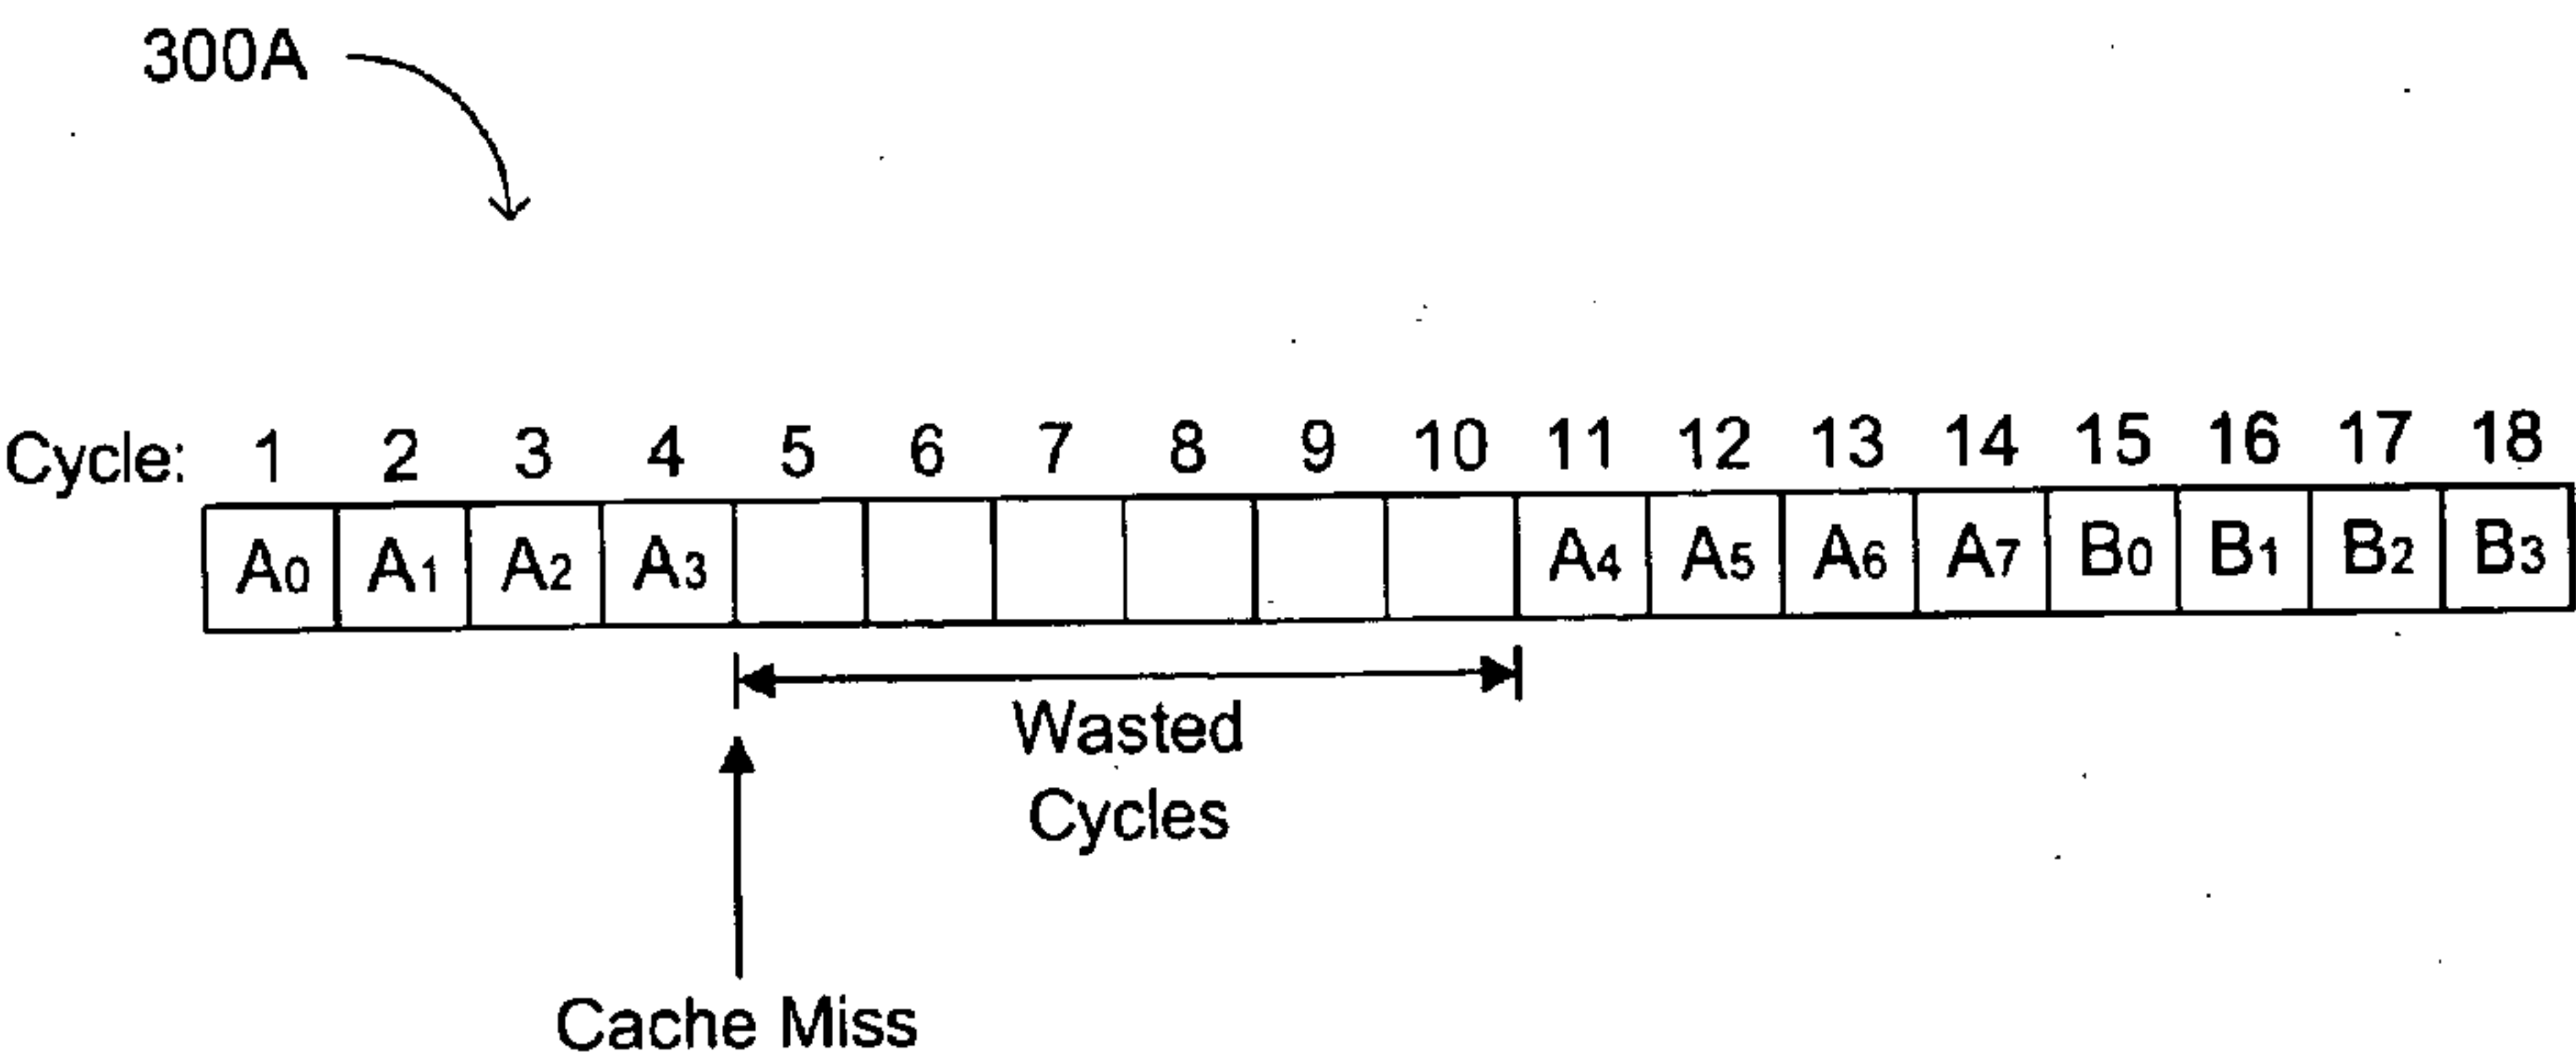


FIG. 3A (conventional)

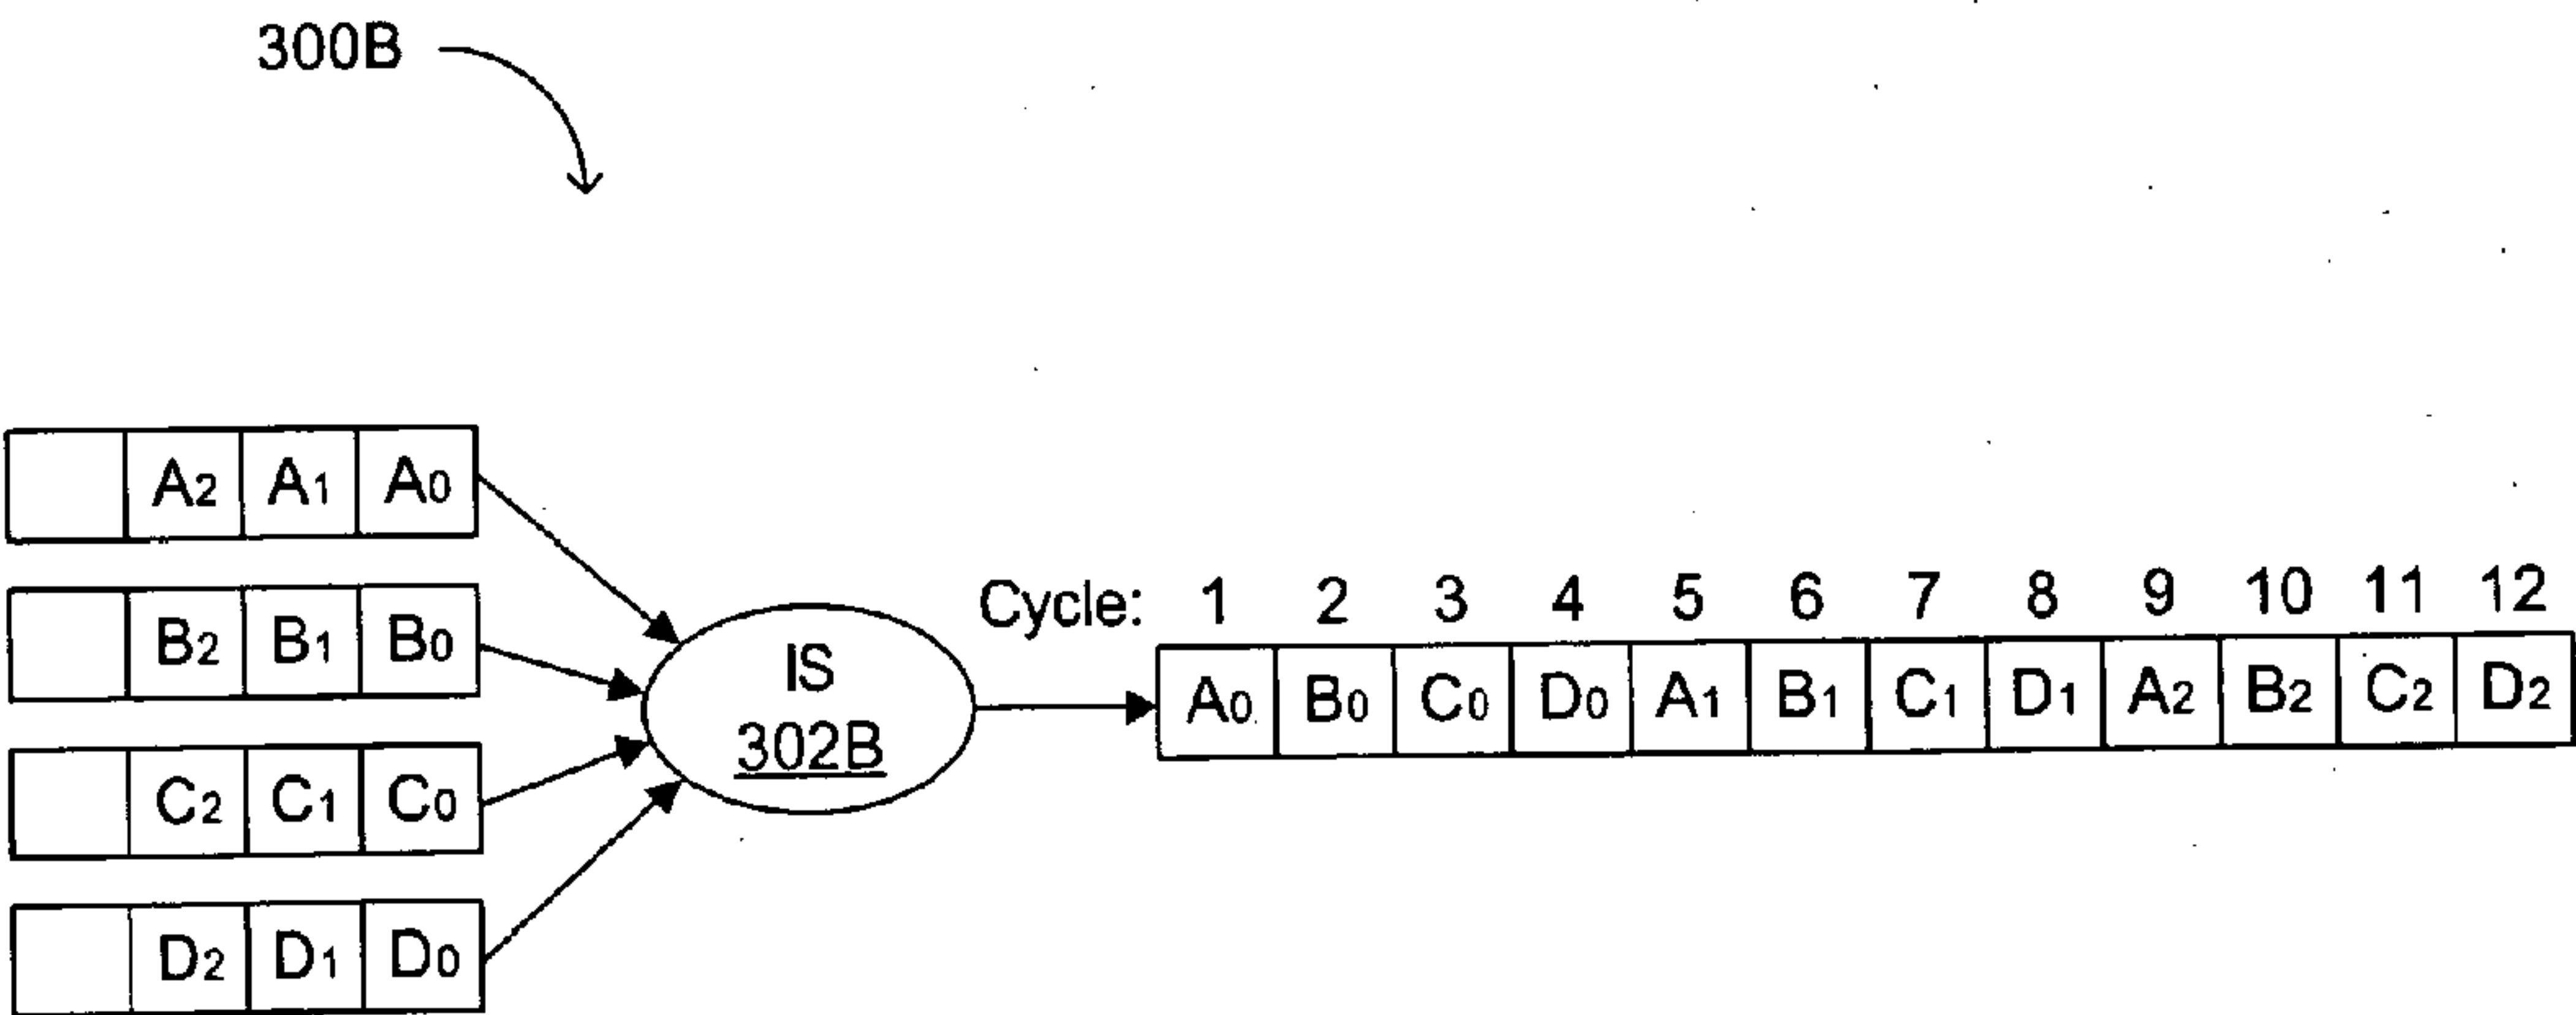


FIG. 3B (conventional)




FIG. 3C (conventional)

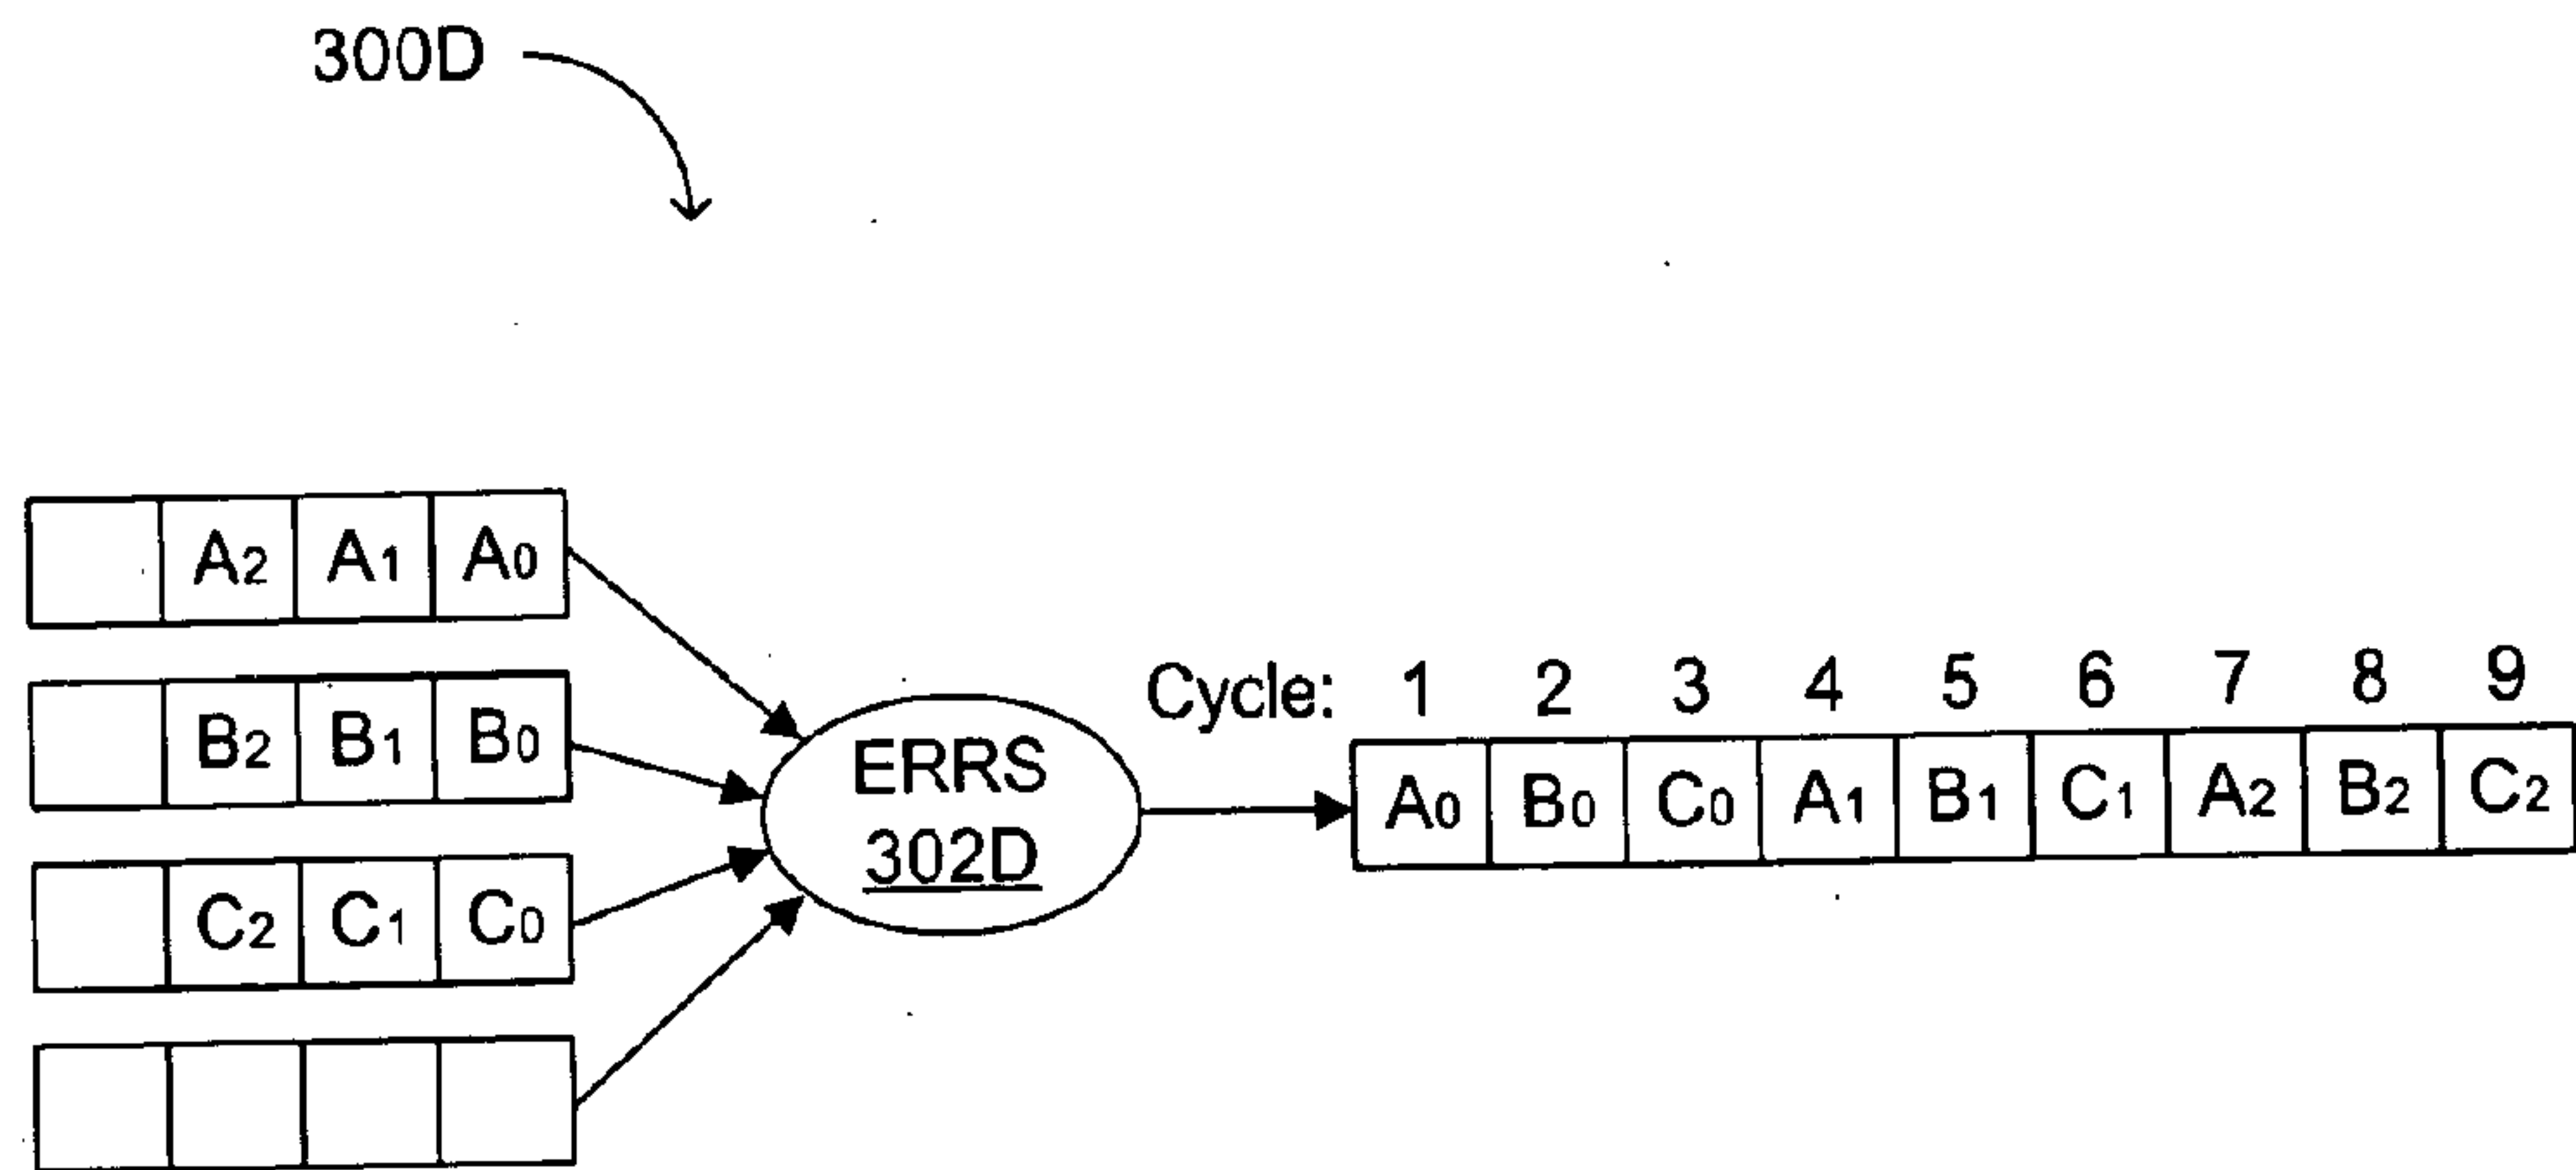


FIG. 3D

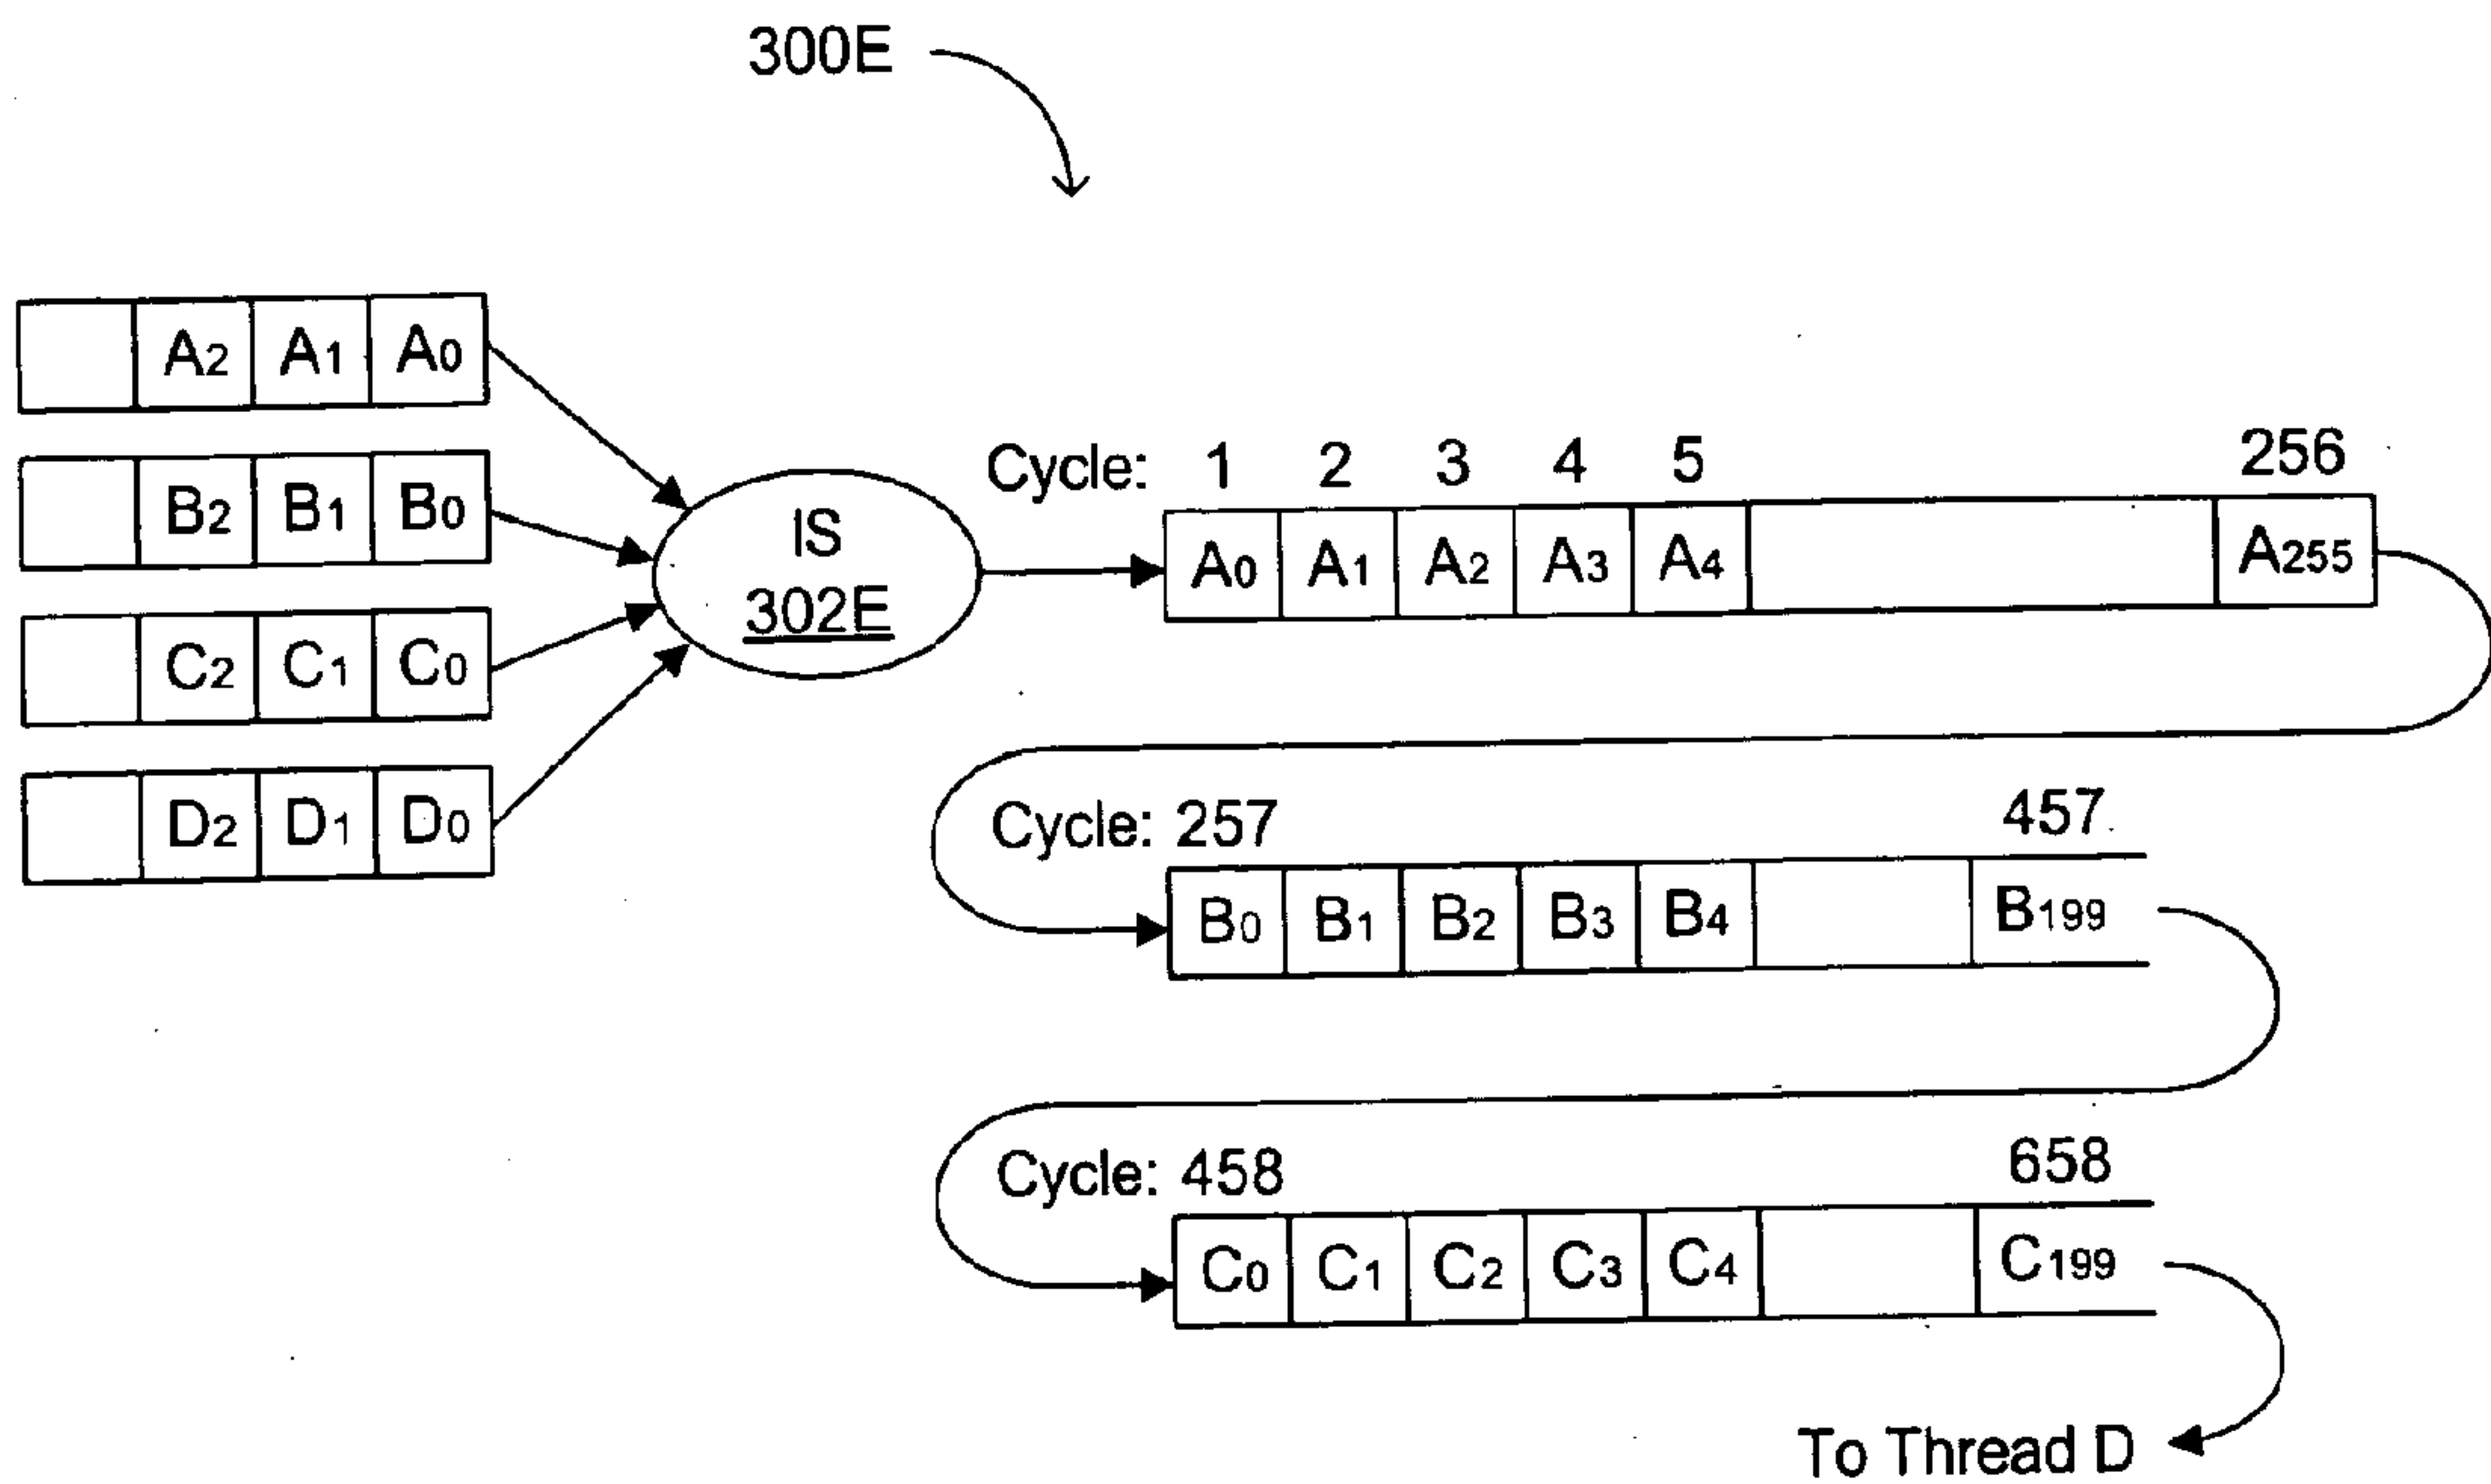


FIG. 3E

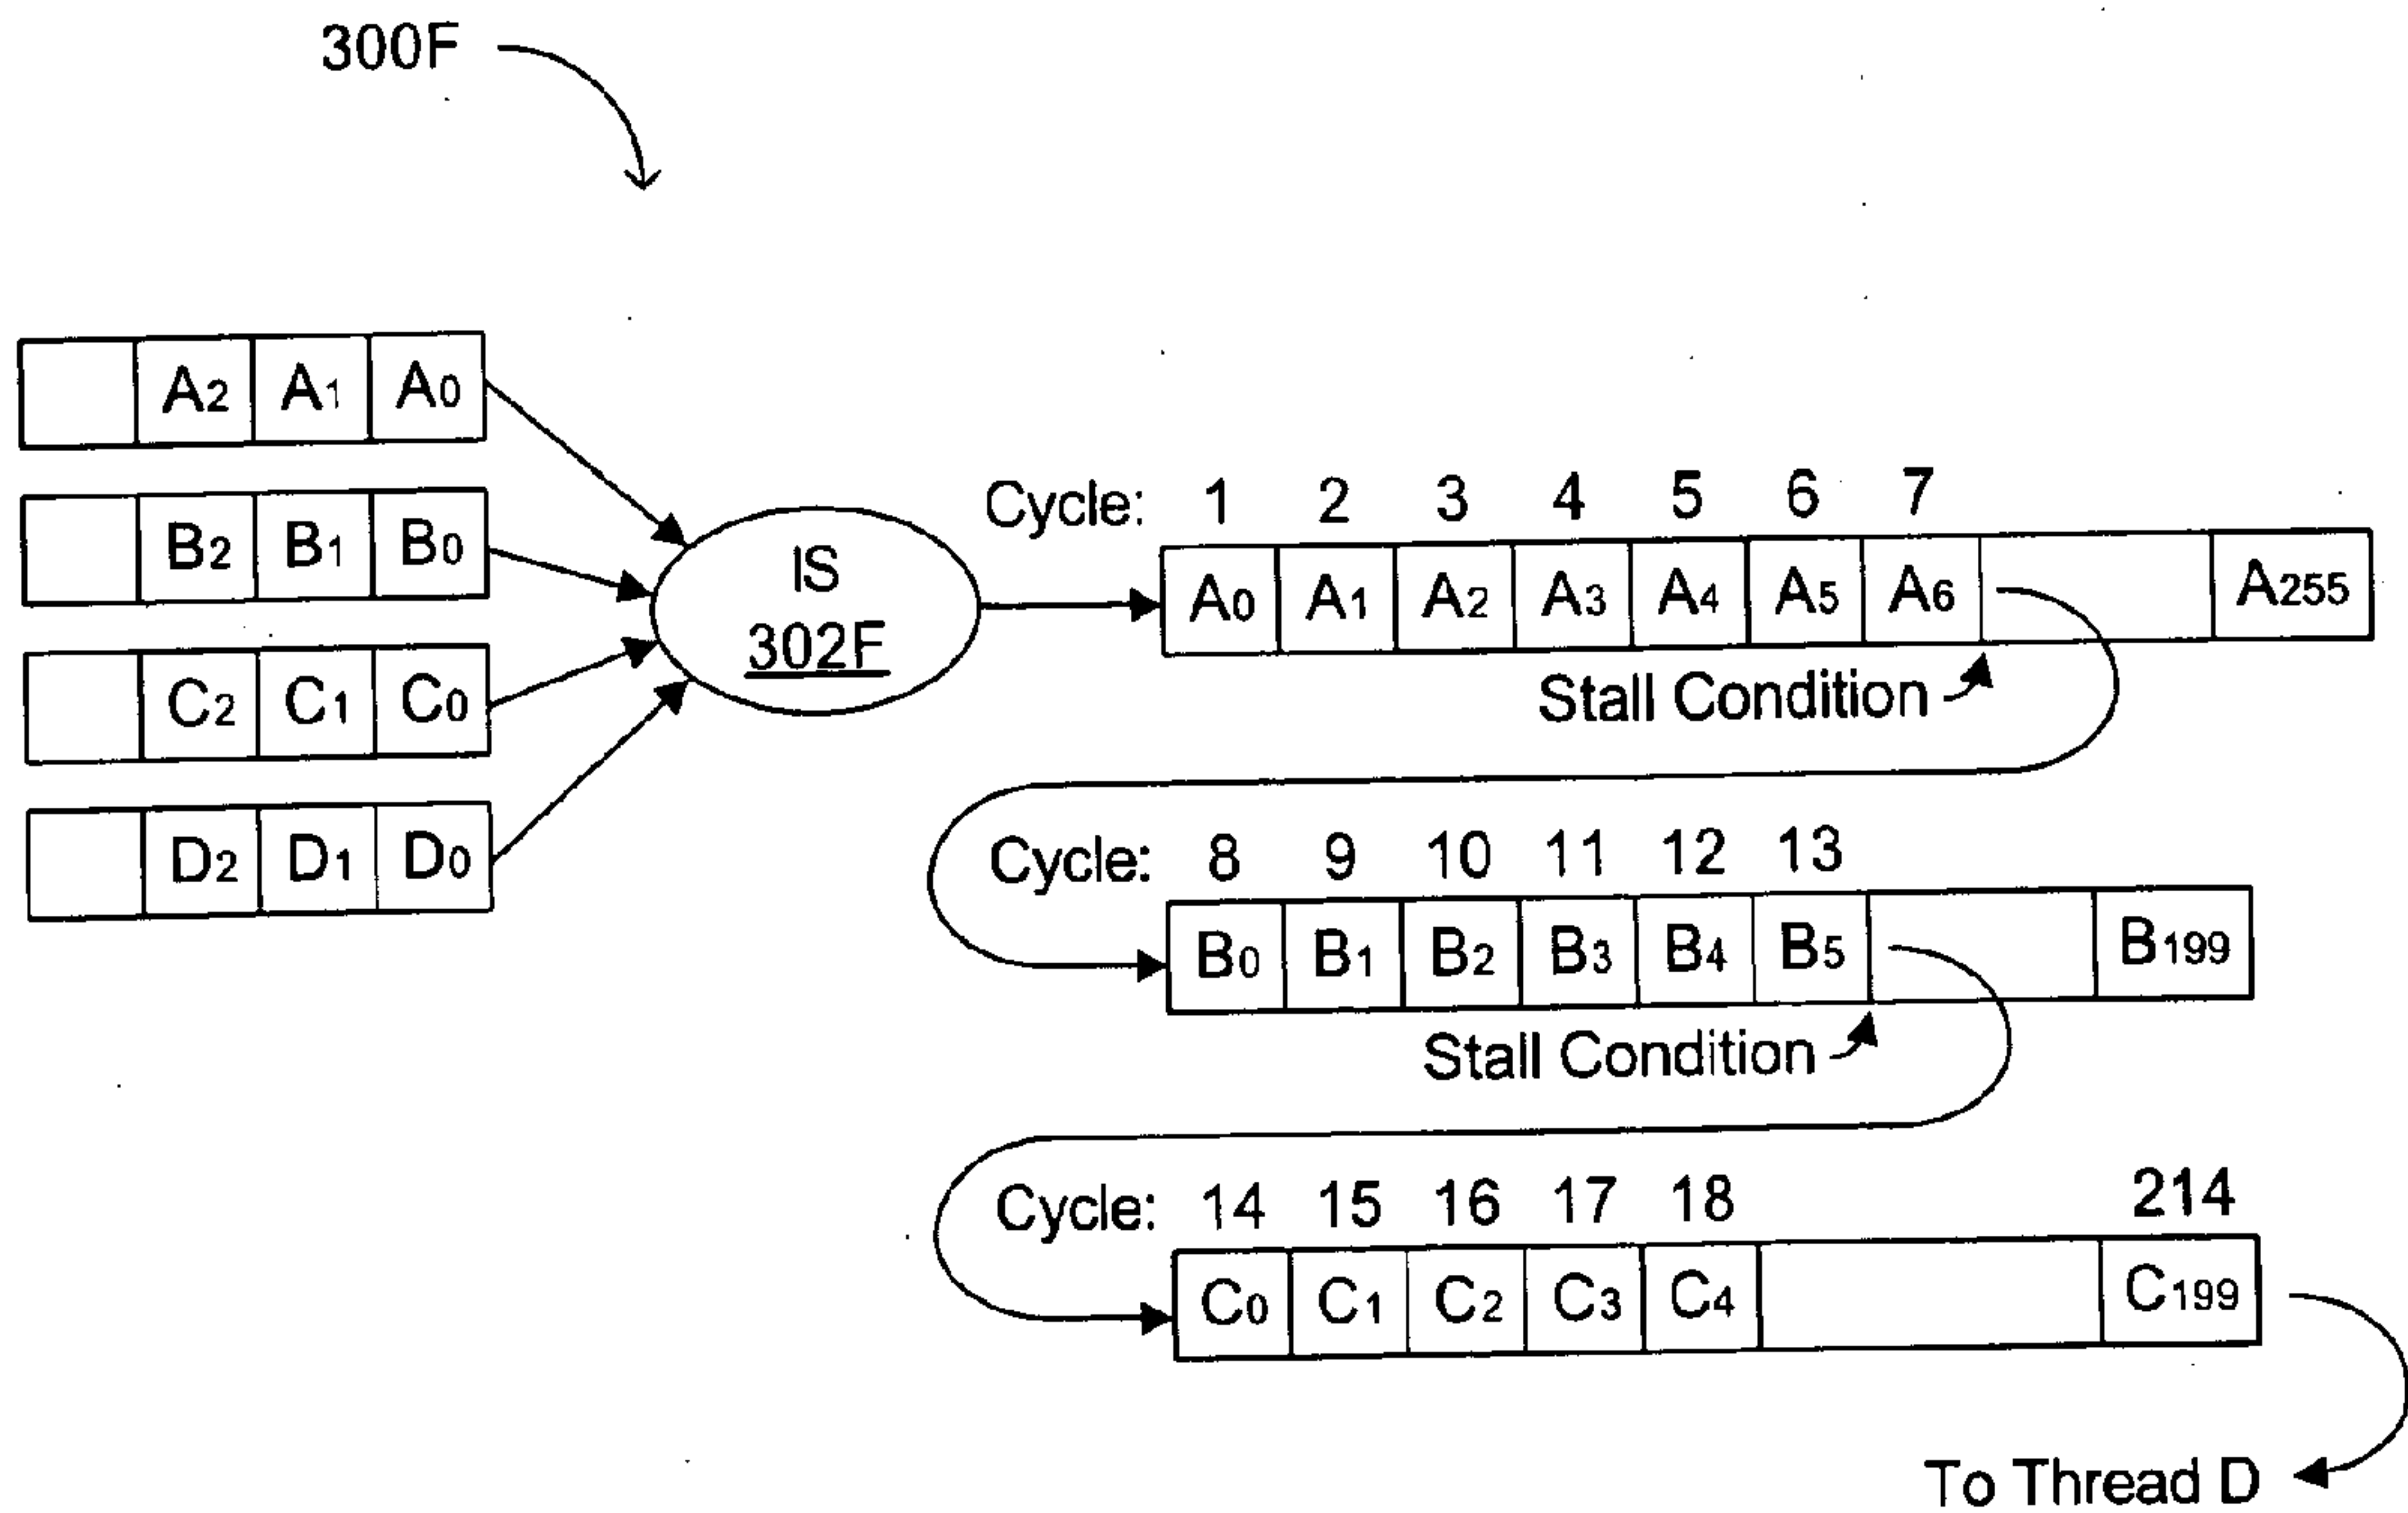


FIG. 3F




FIG. 3G

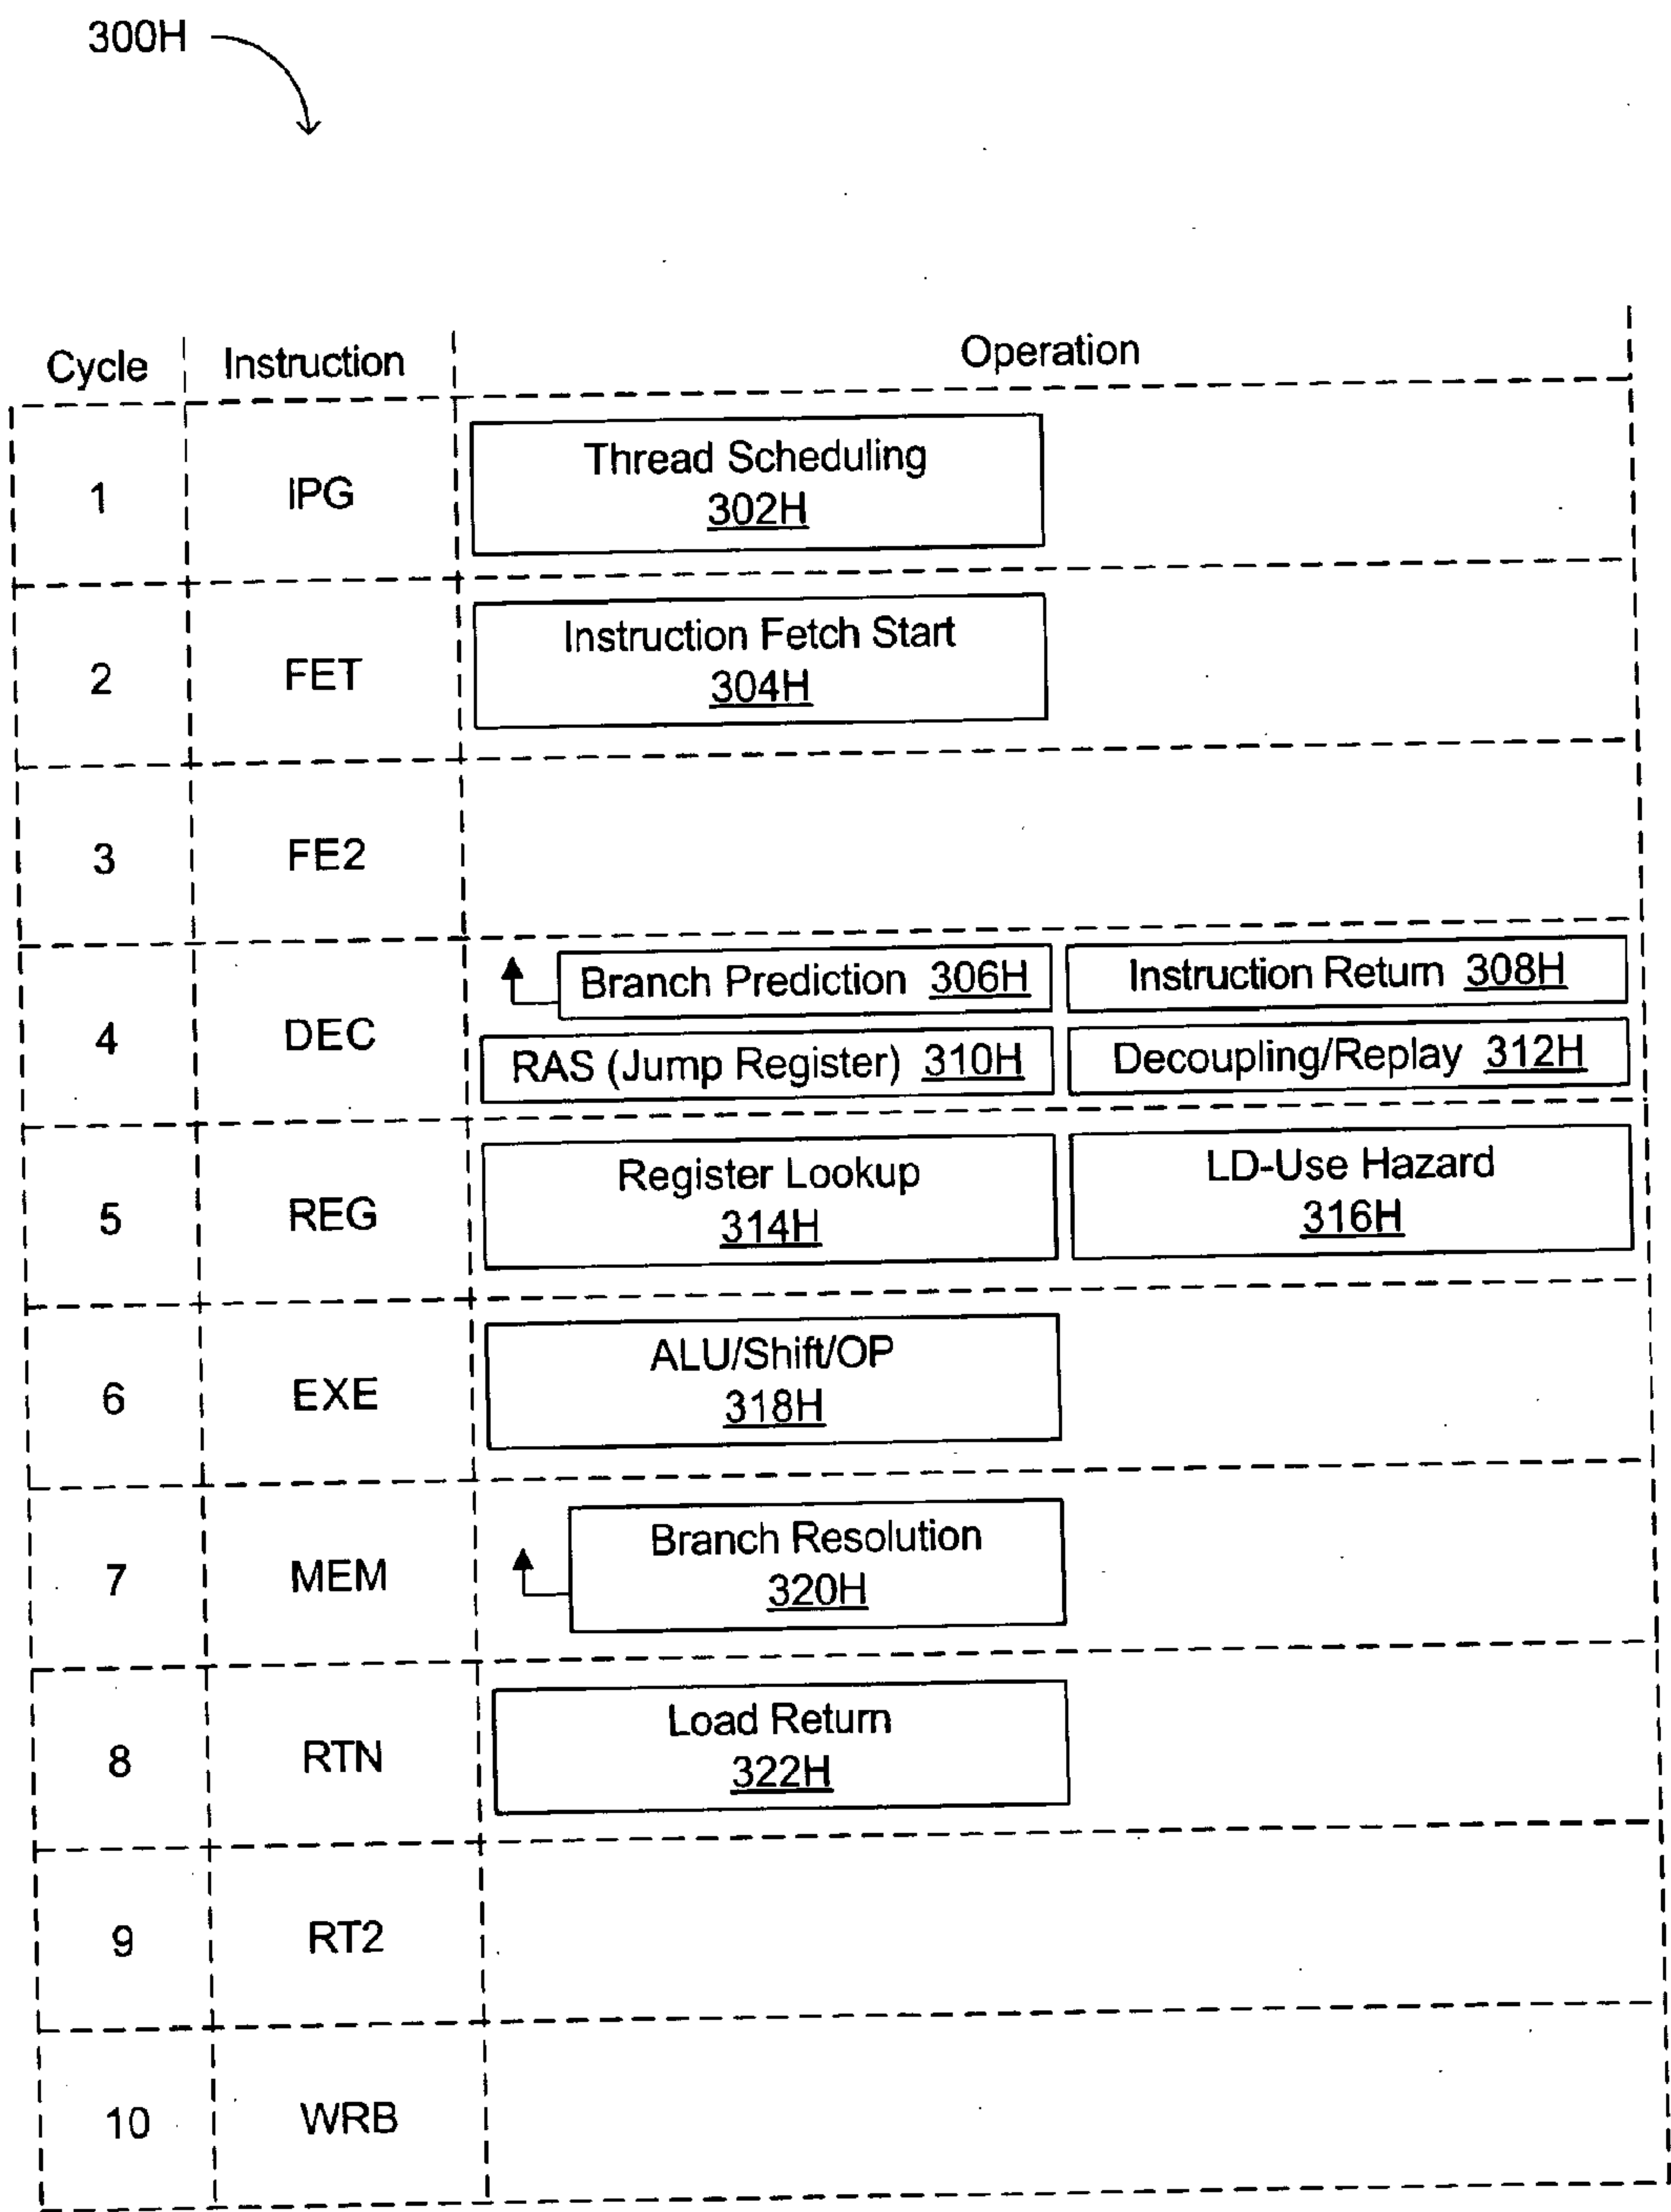


FIG. 3H

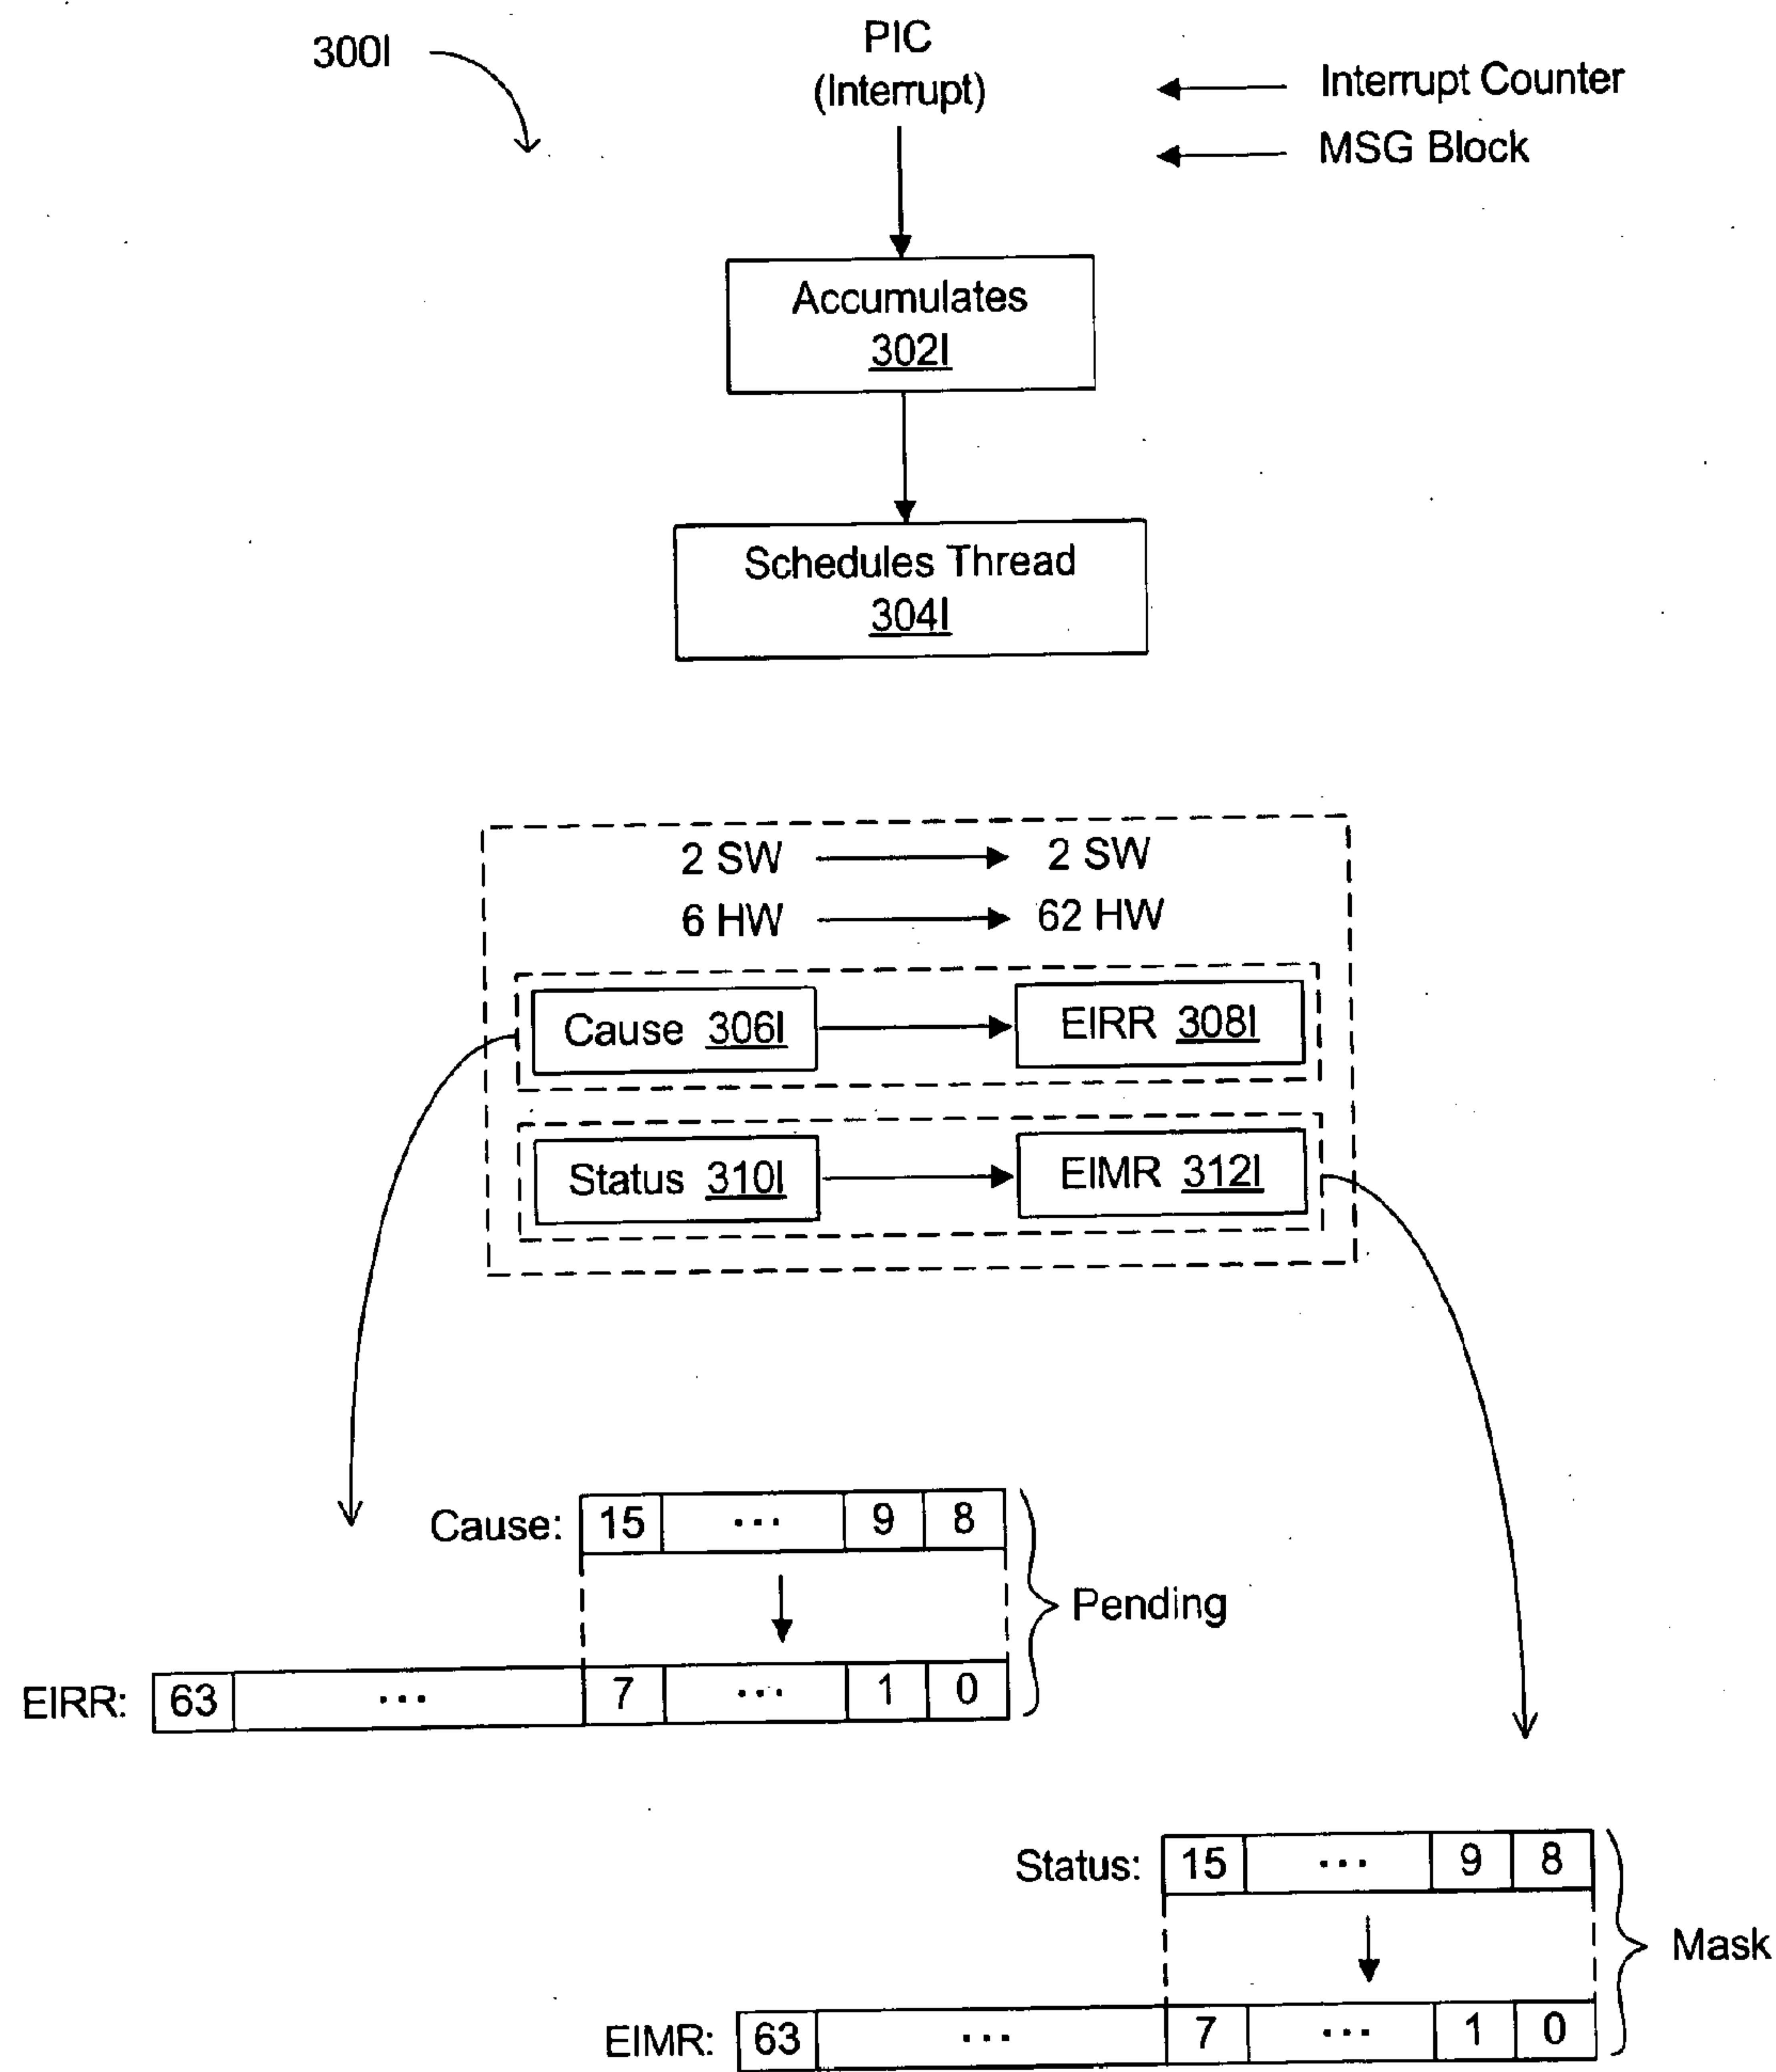


FIG. 3I

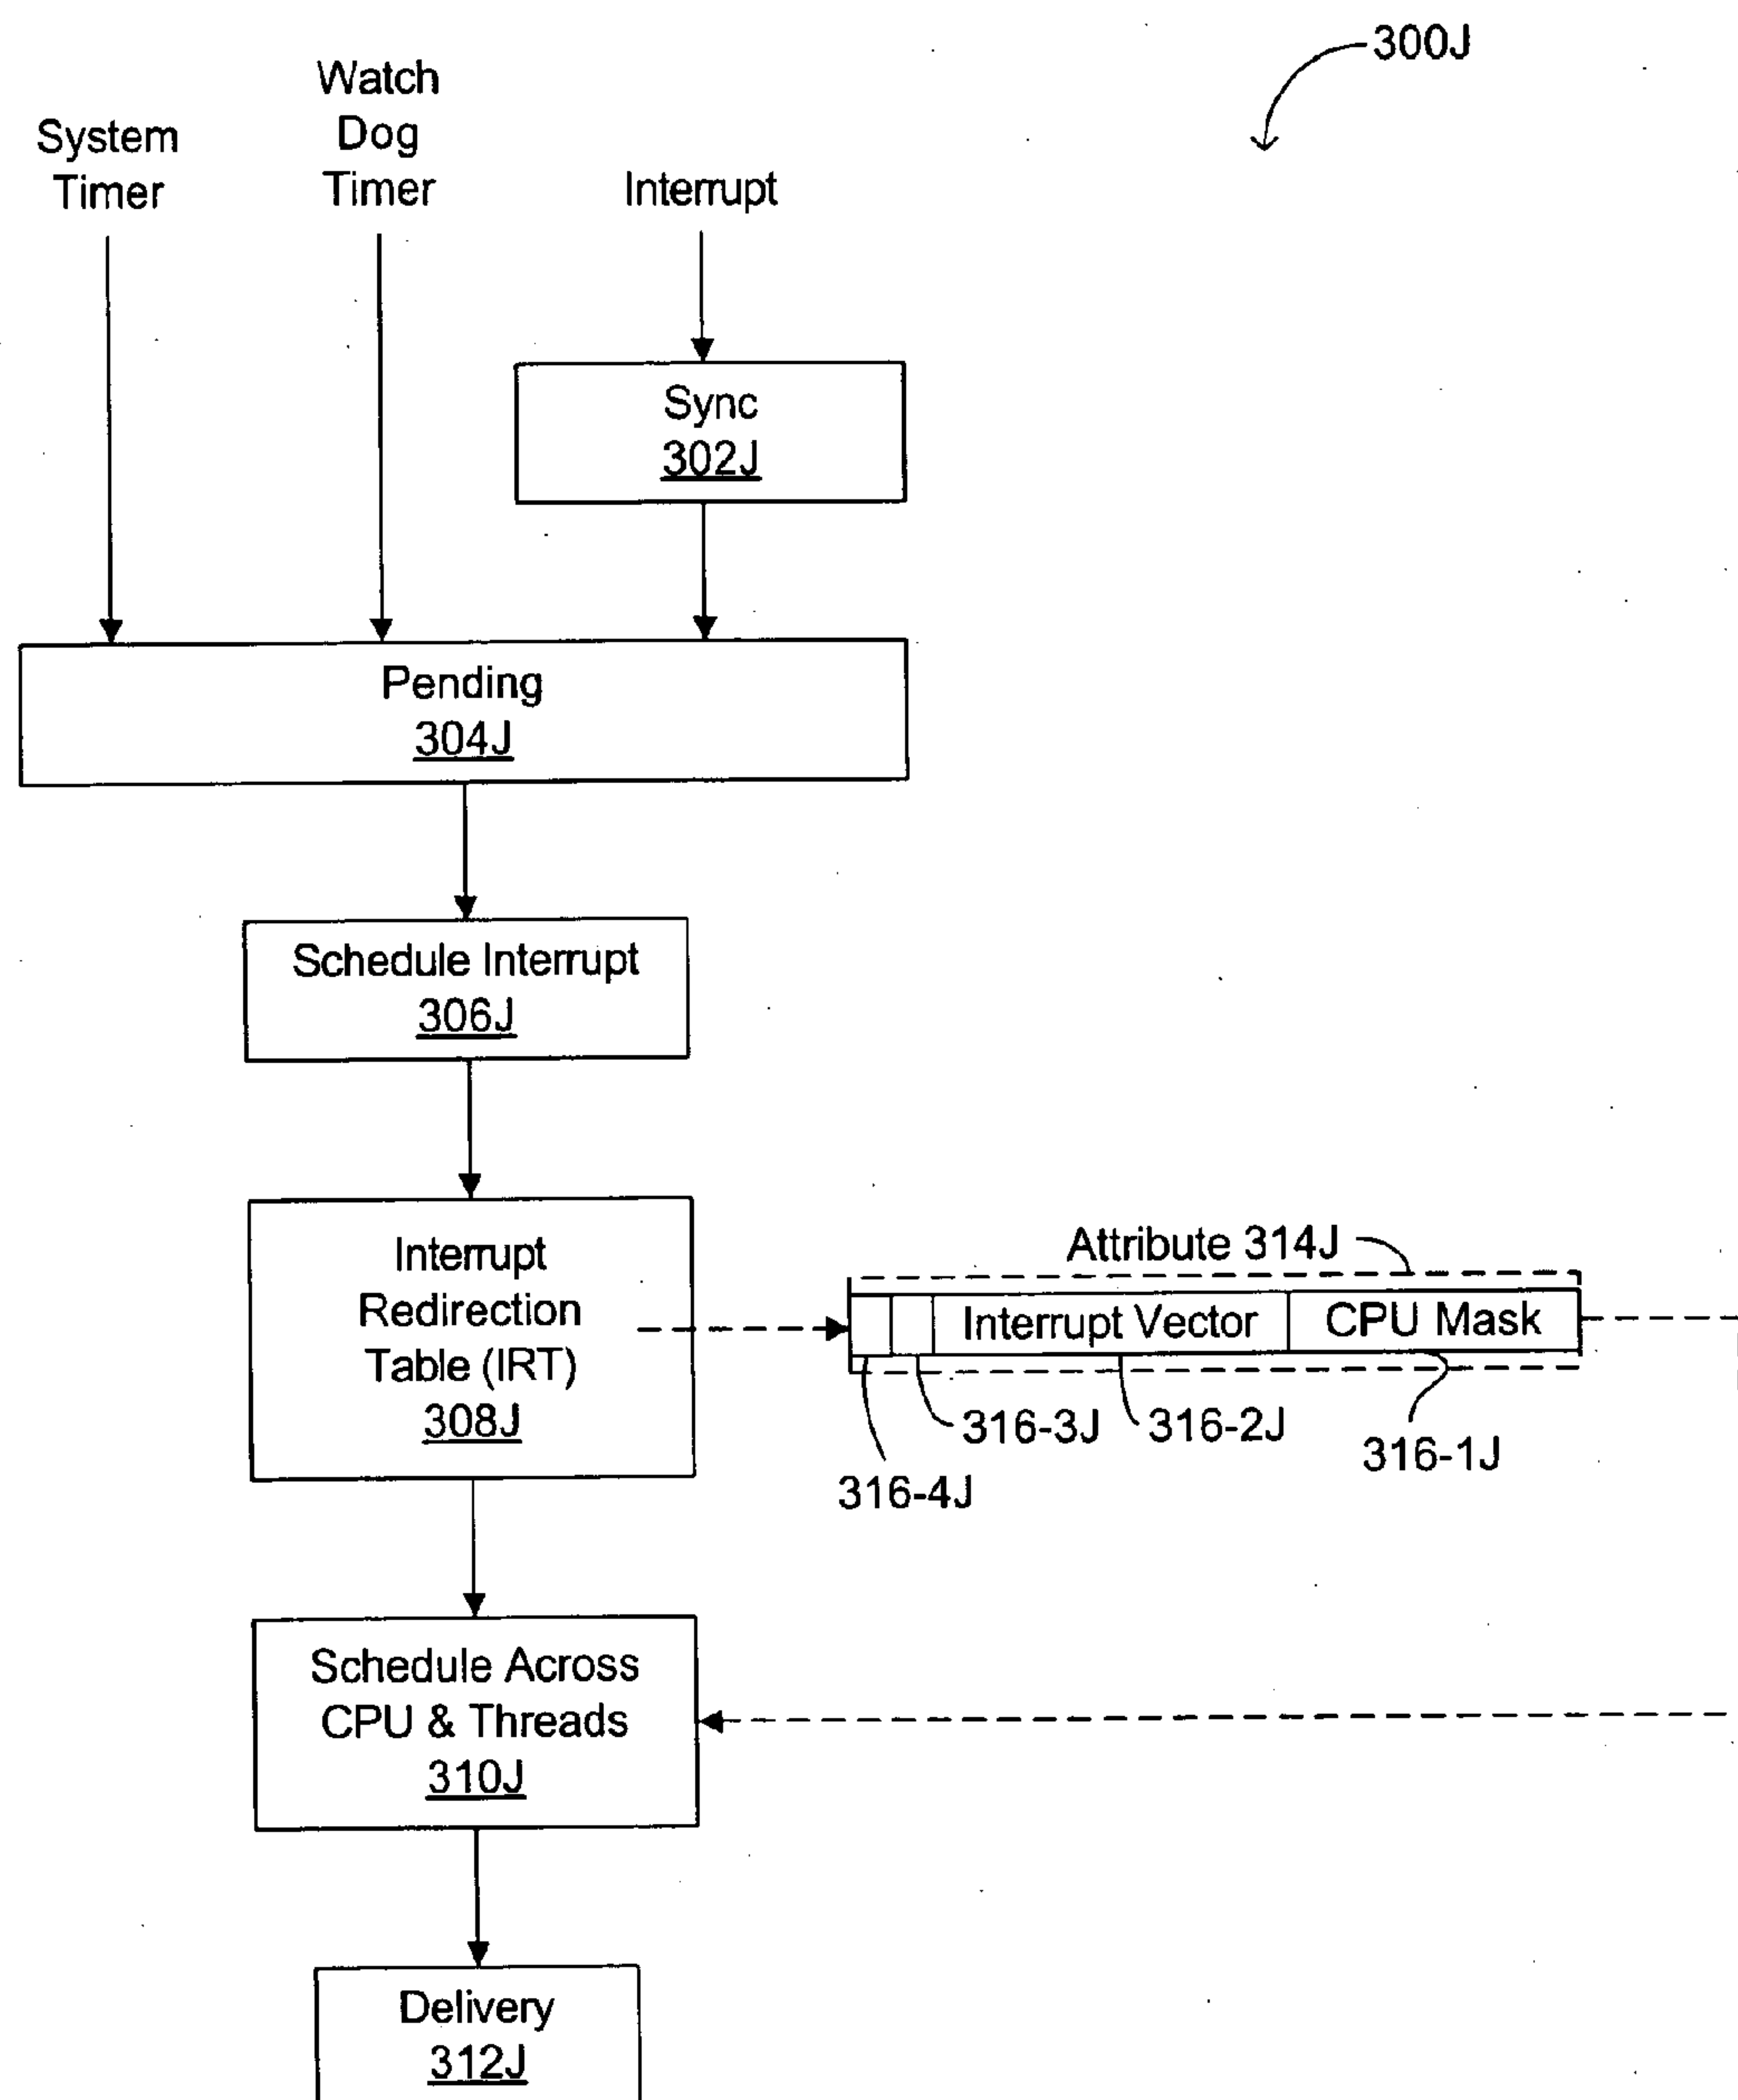


FIG. 3J

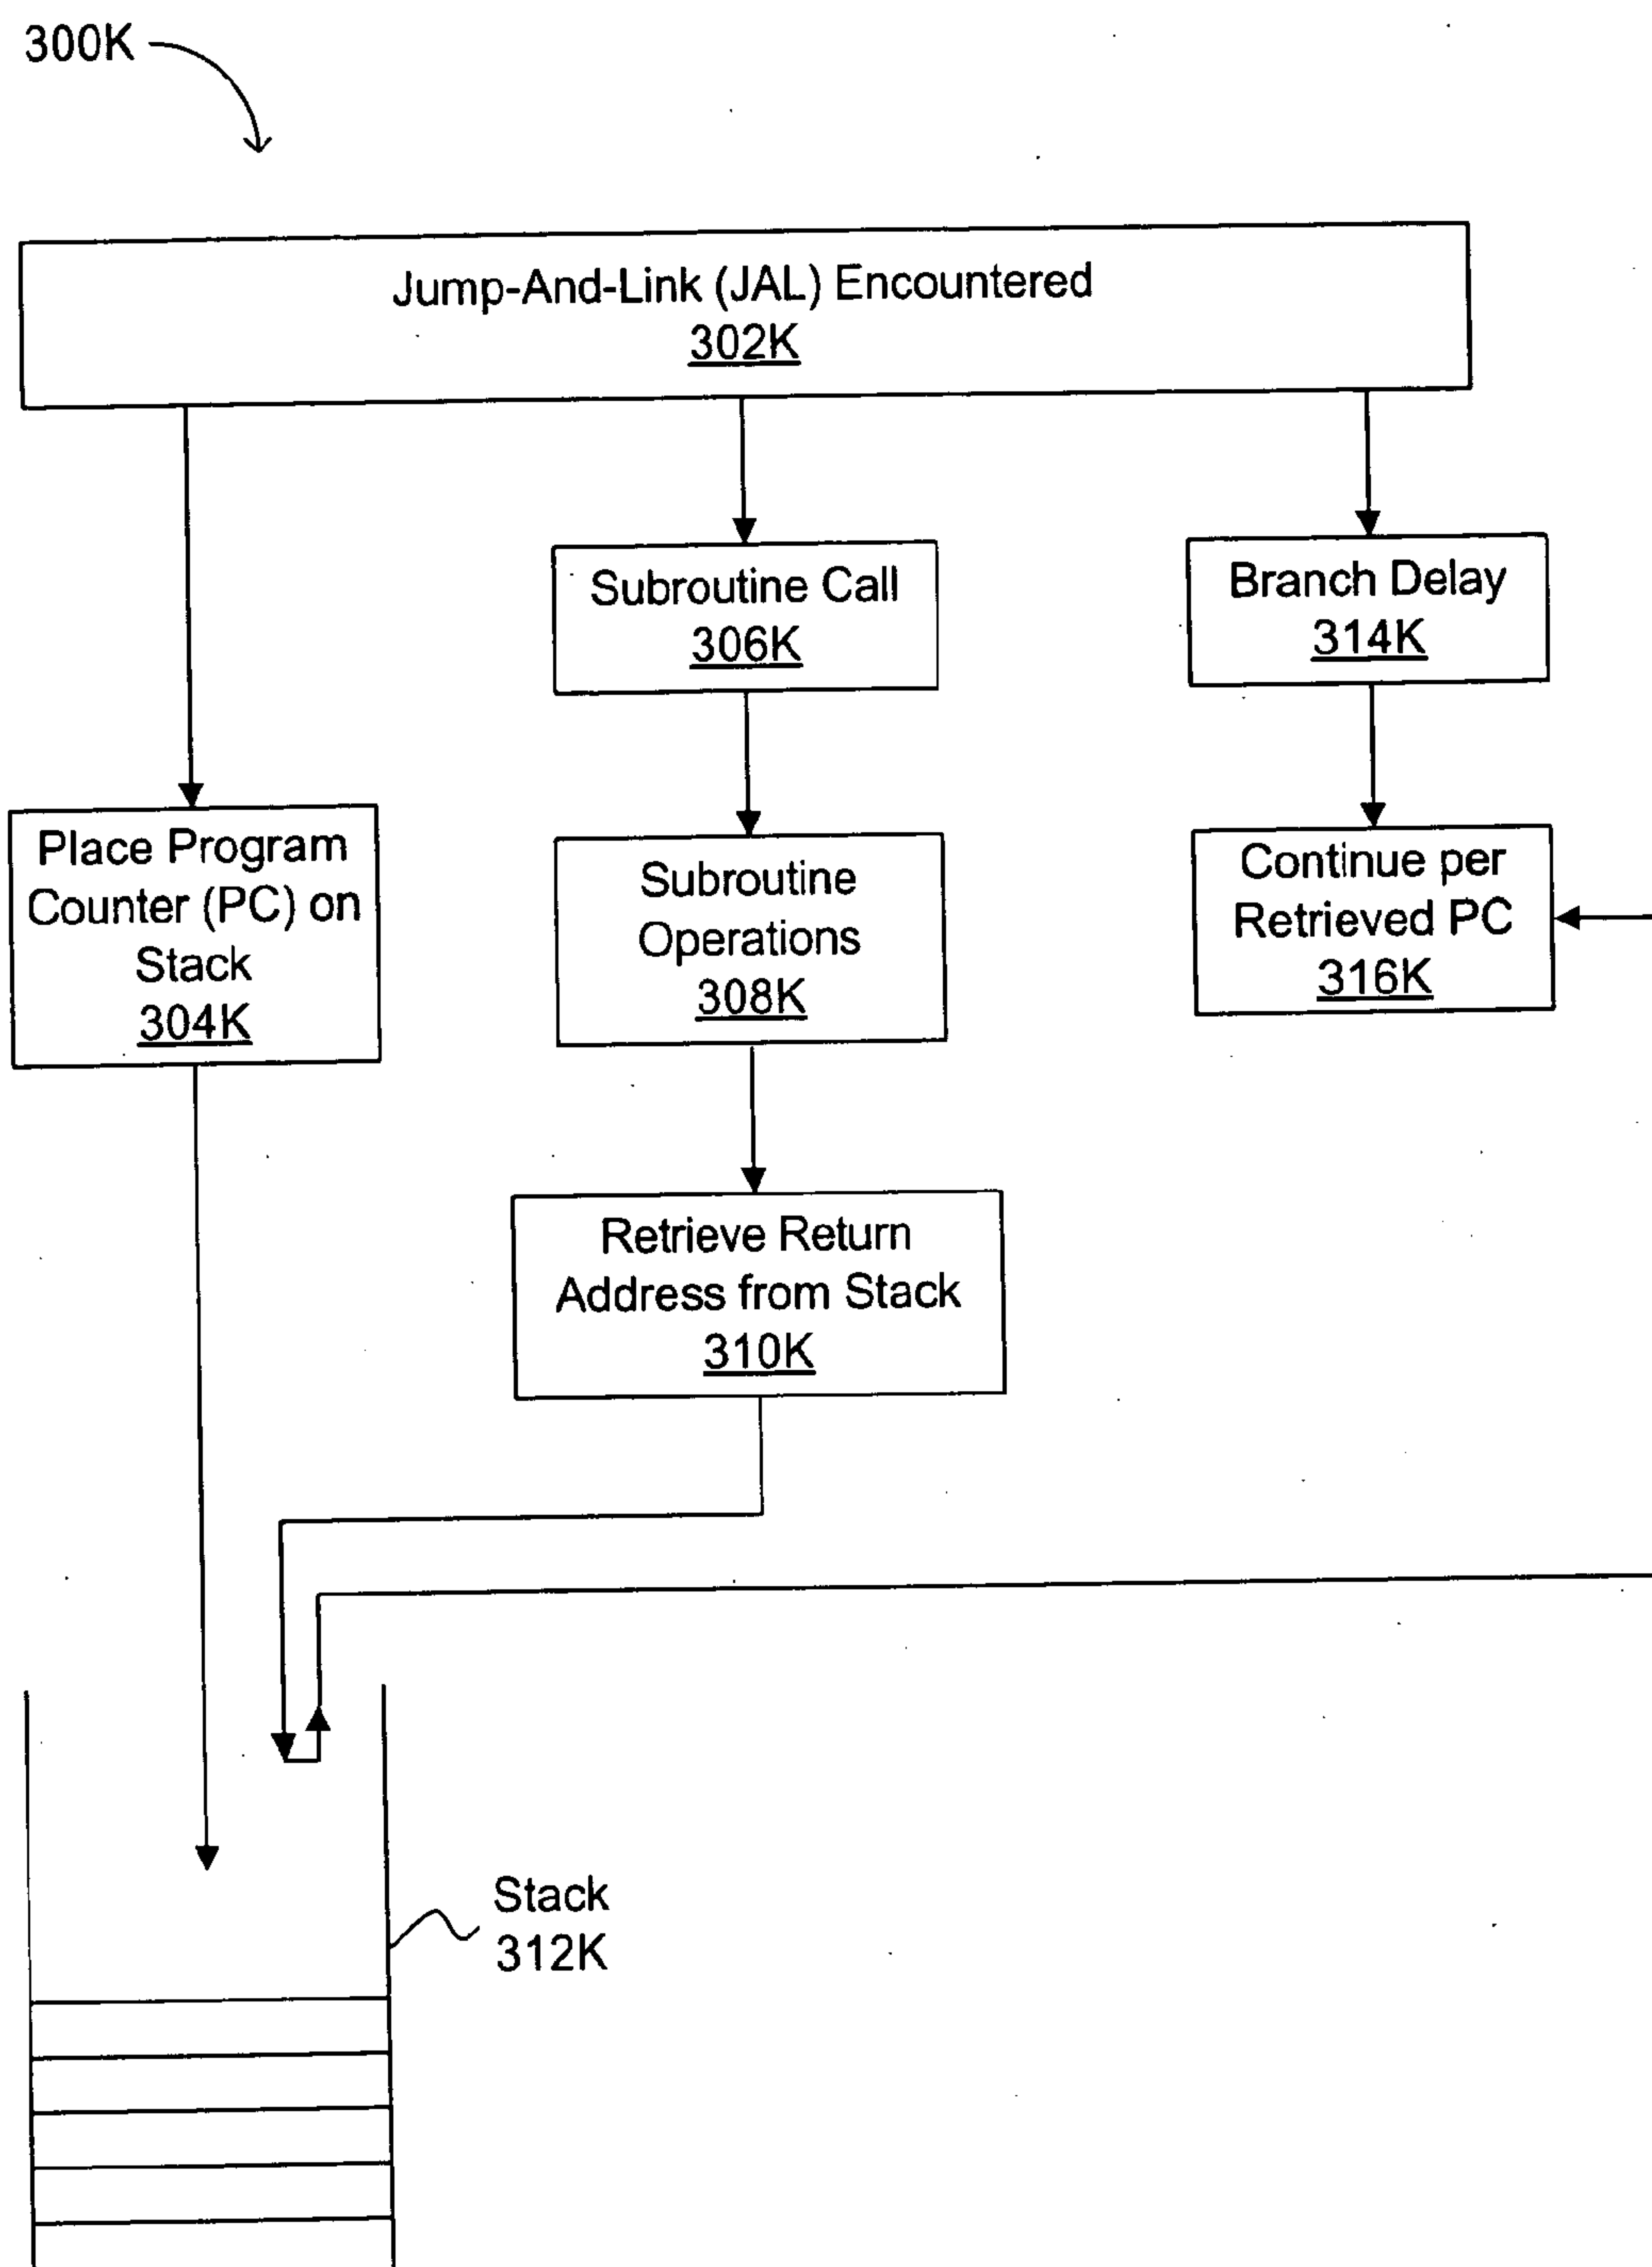


FIG. 3K




FIG. 4A




FIG. 4B

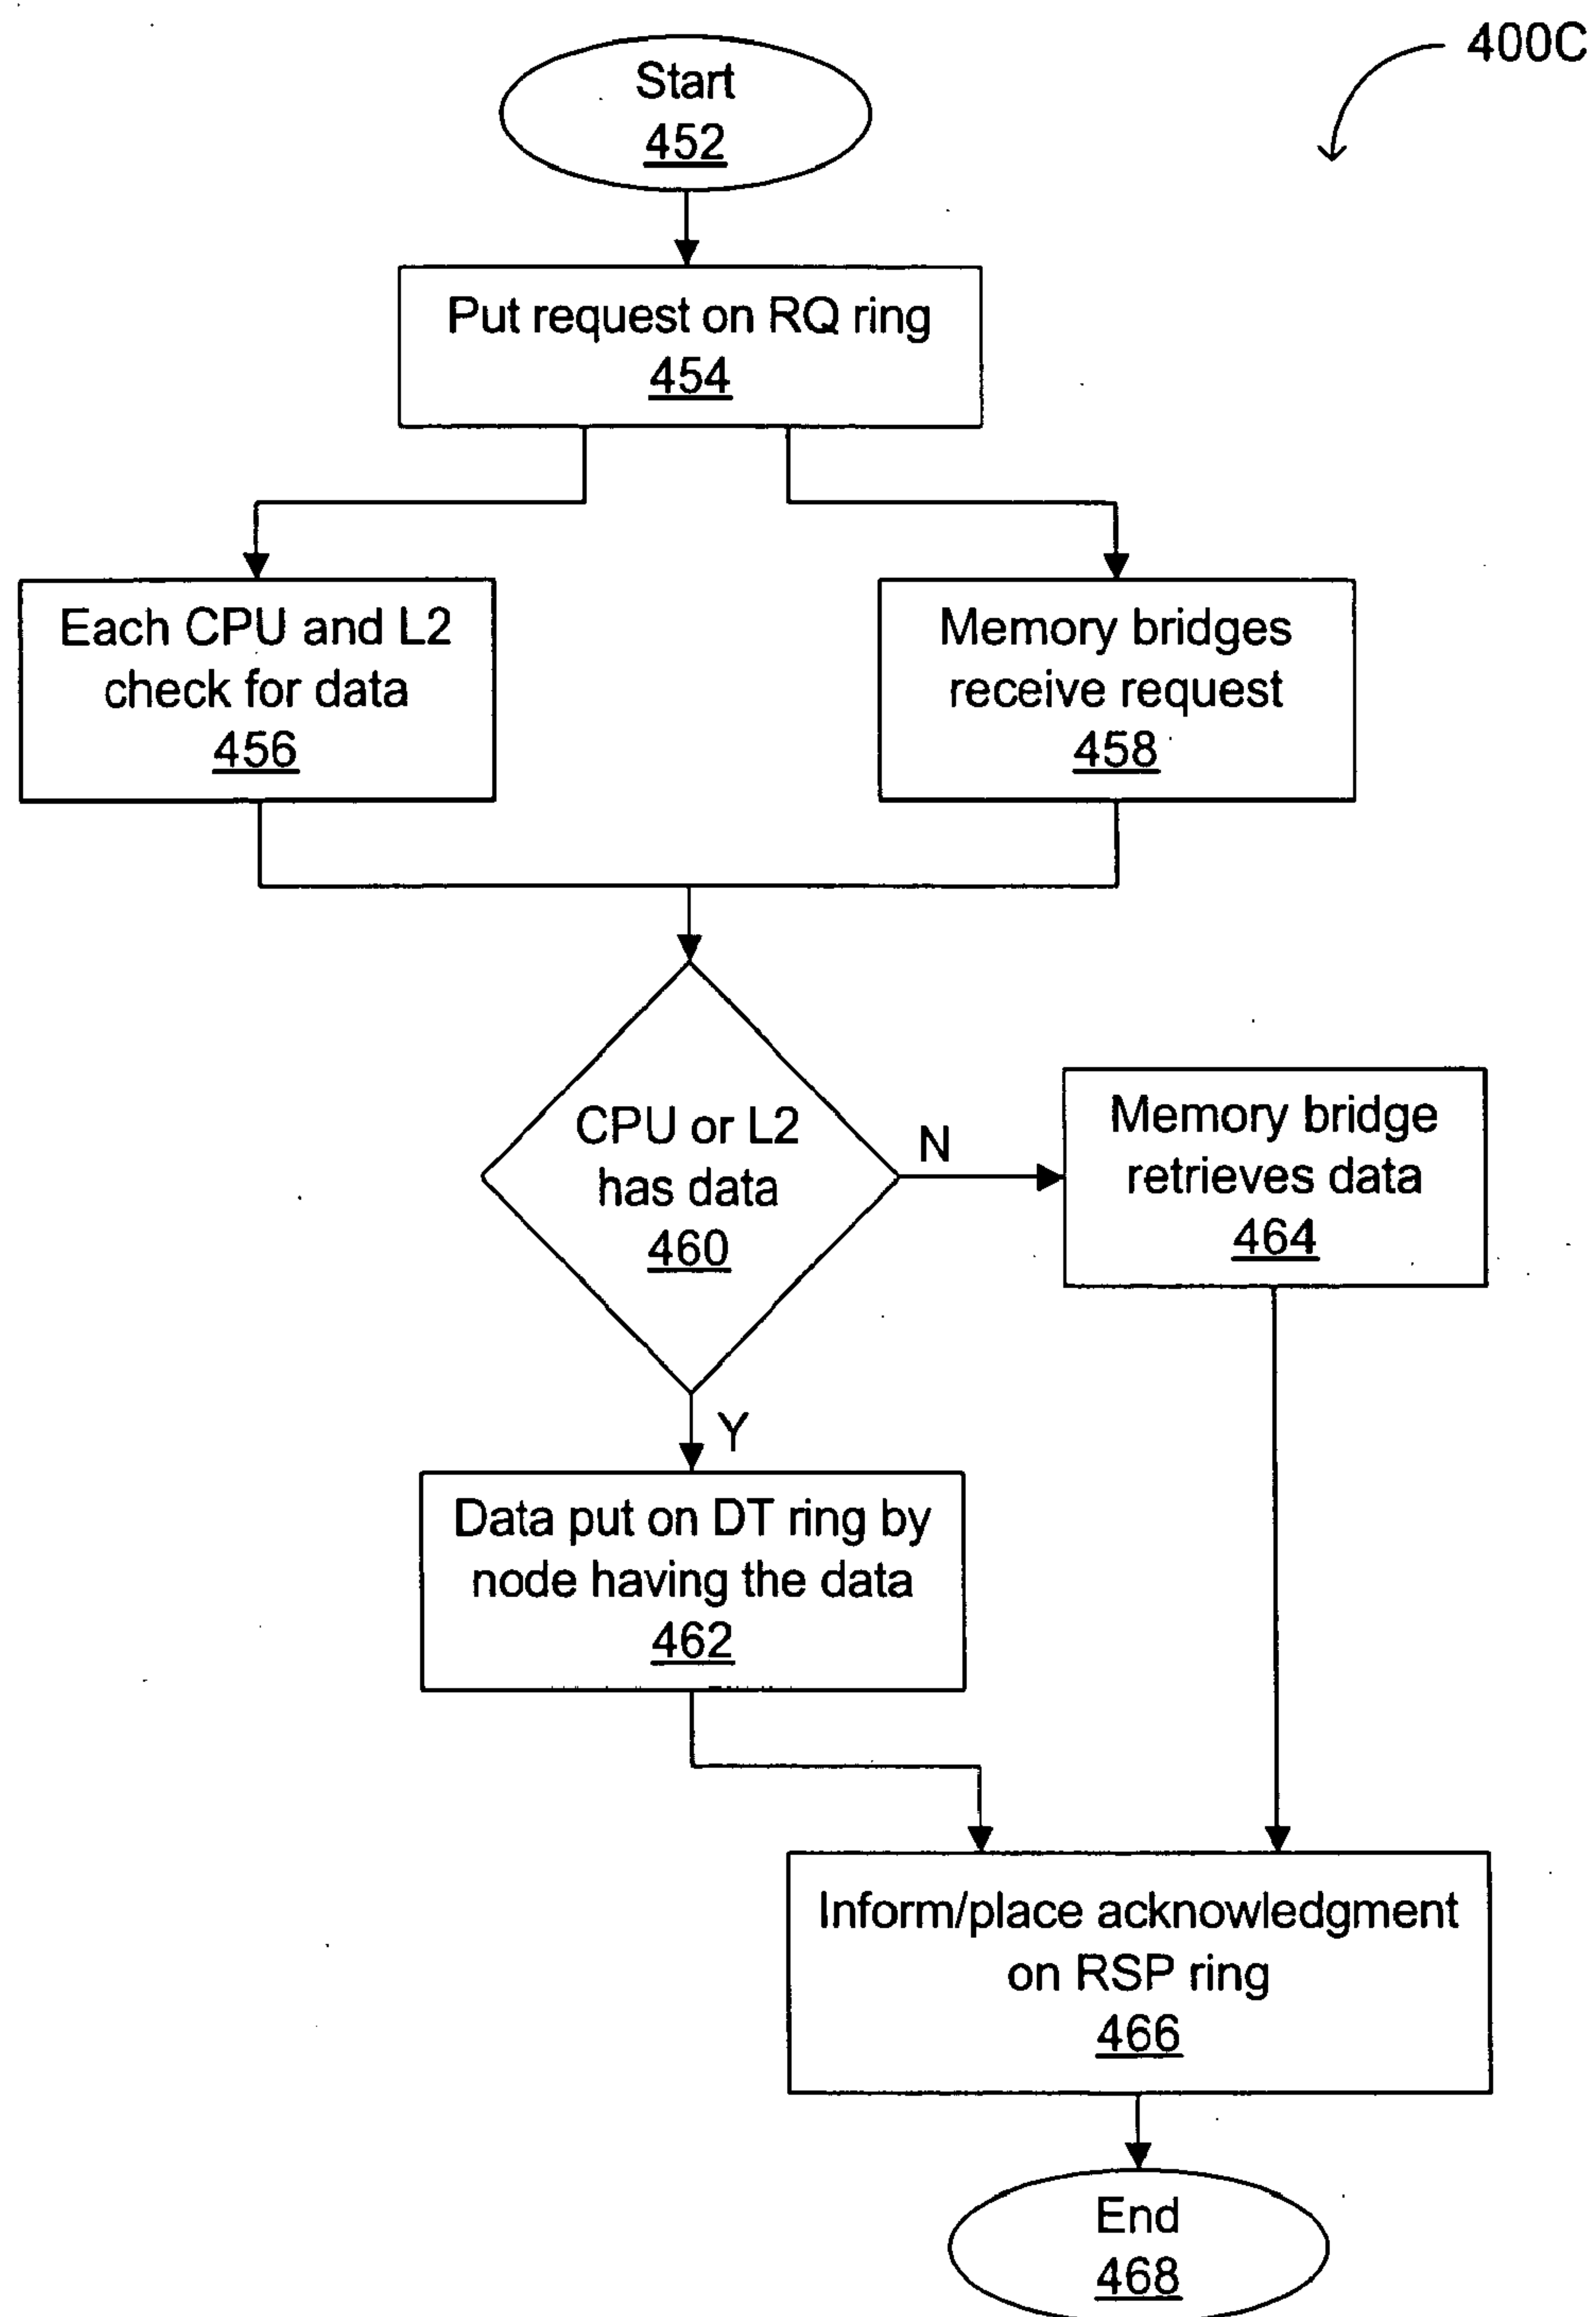


FIG. 4C

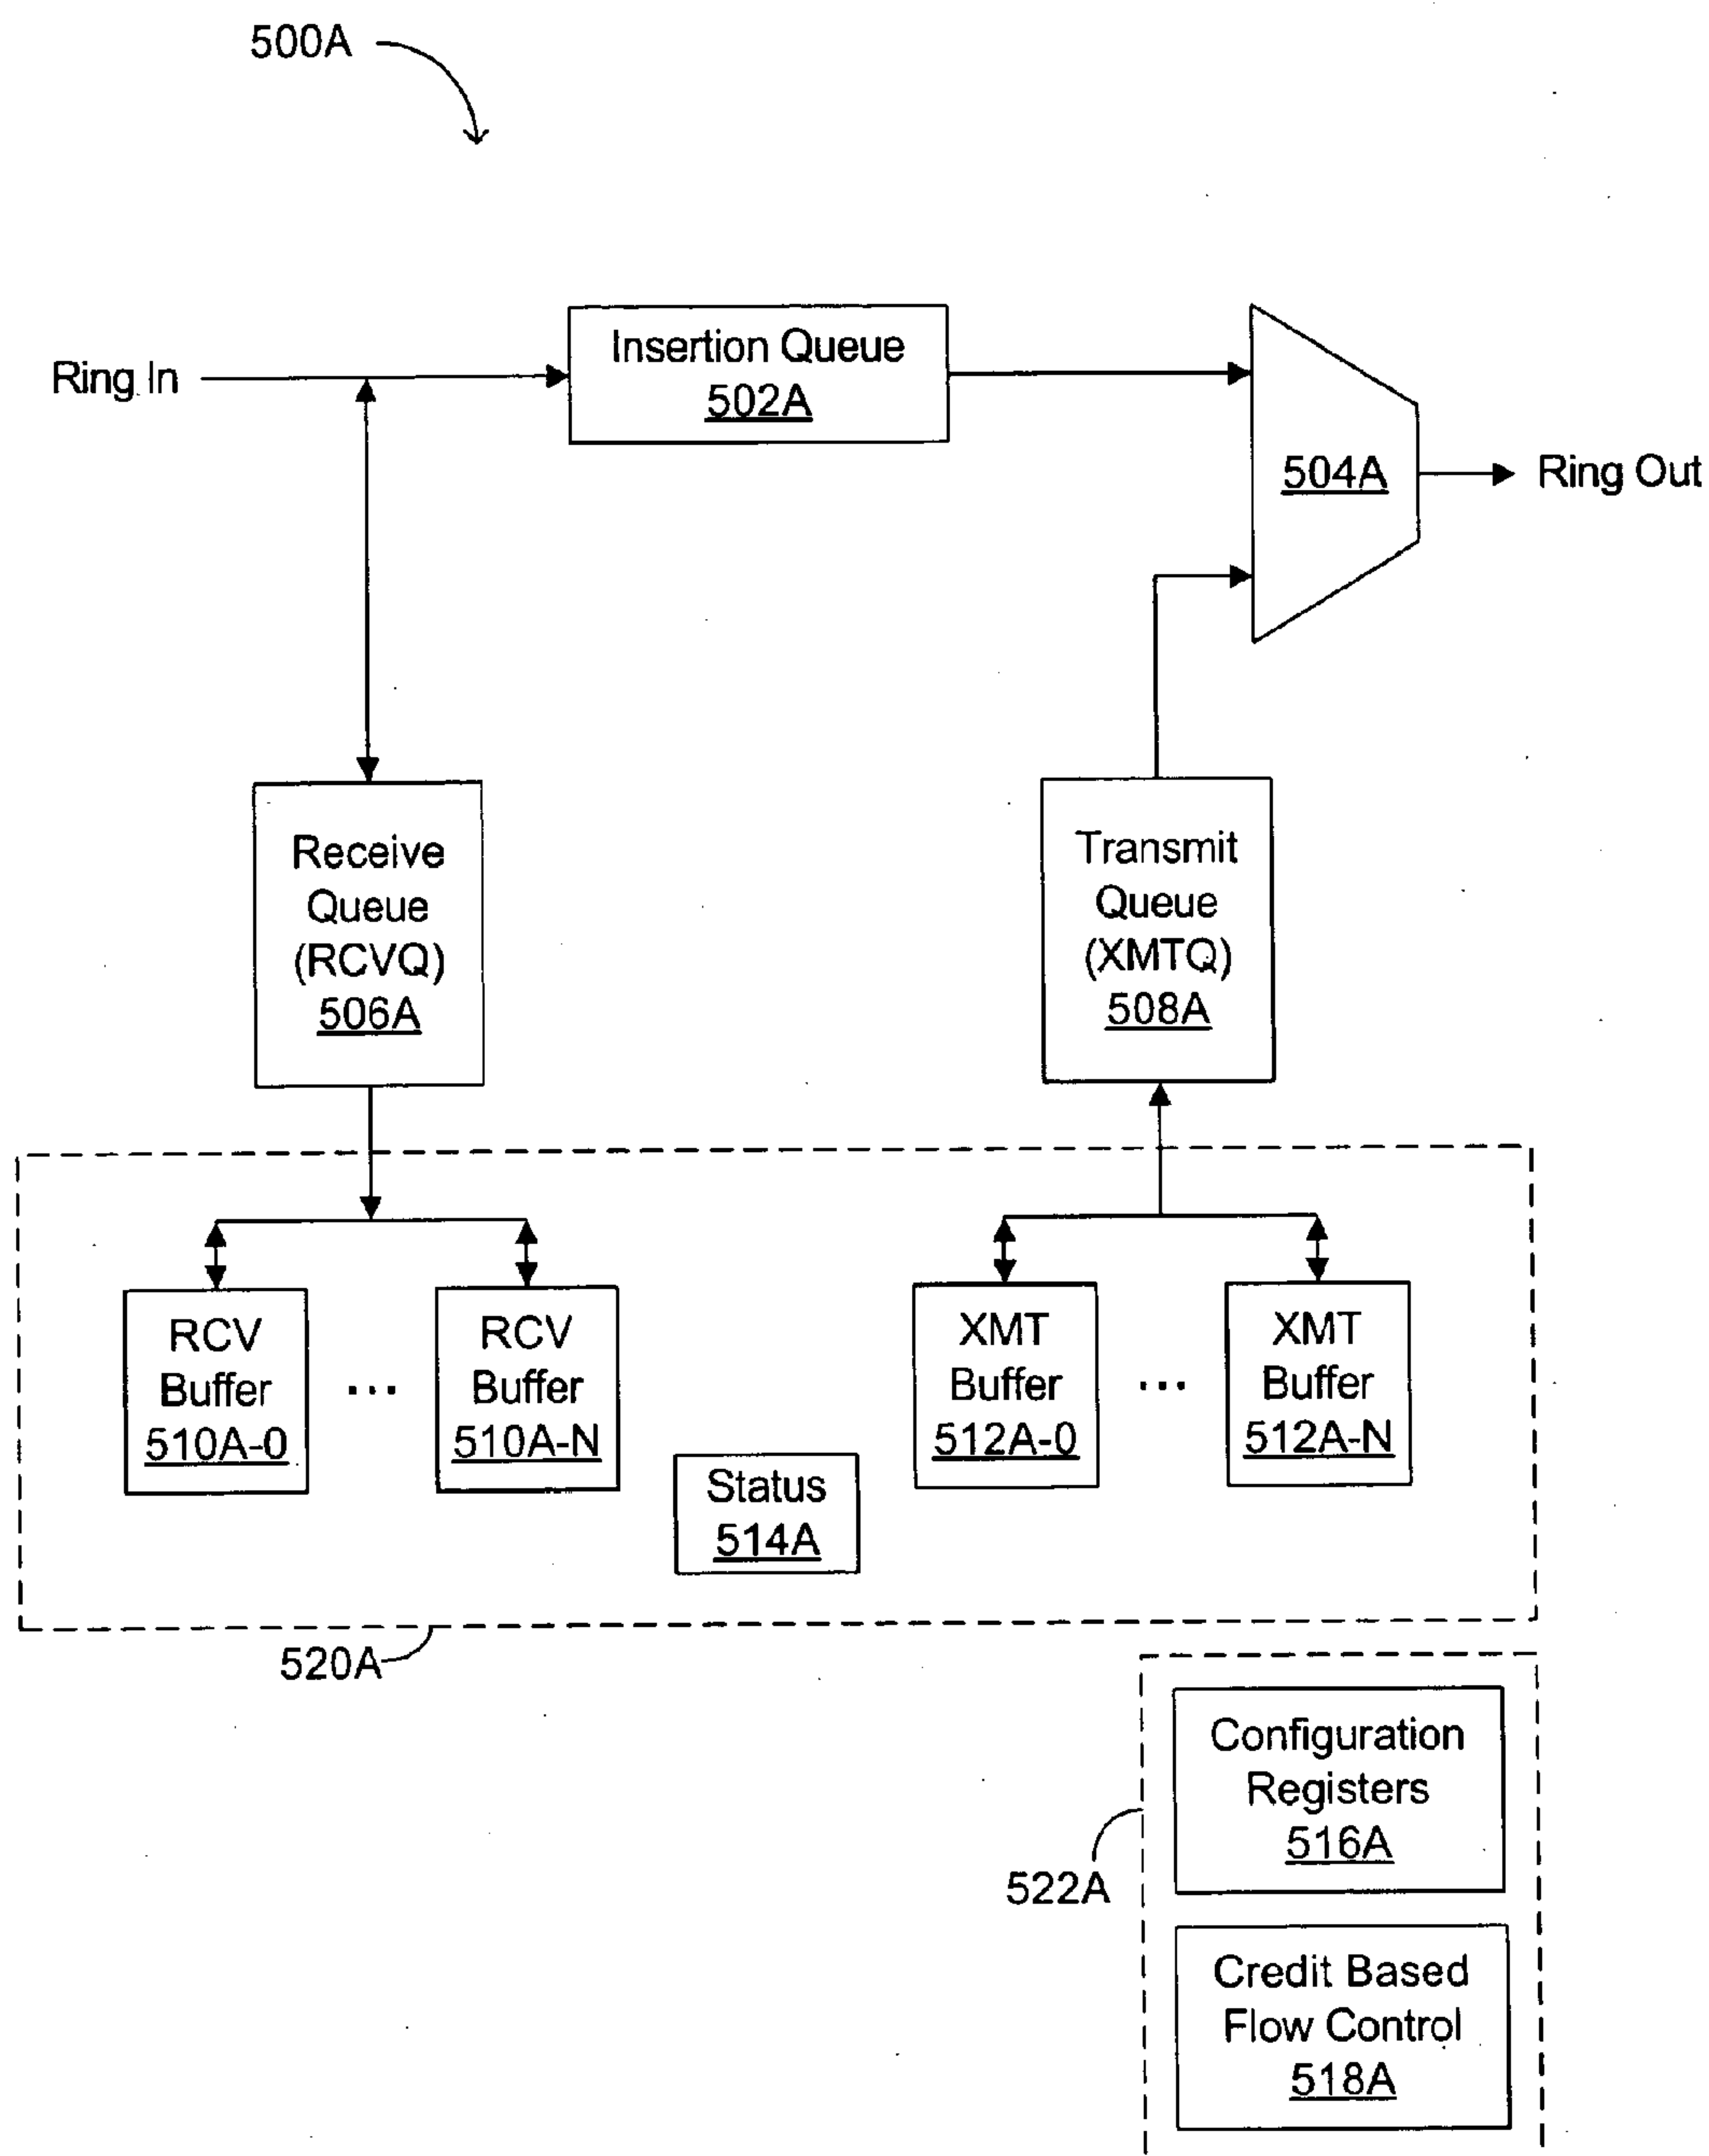
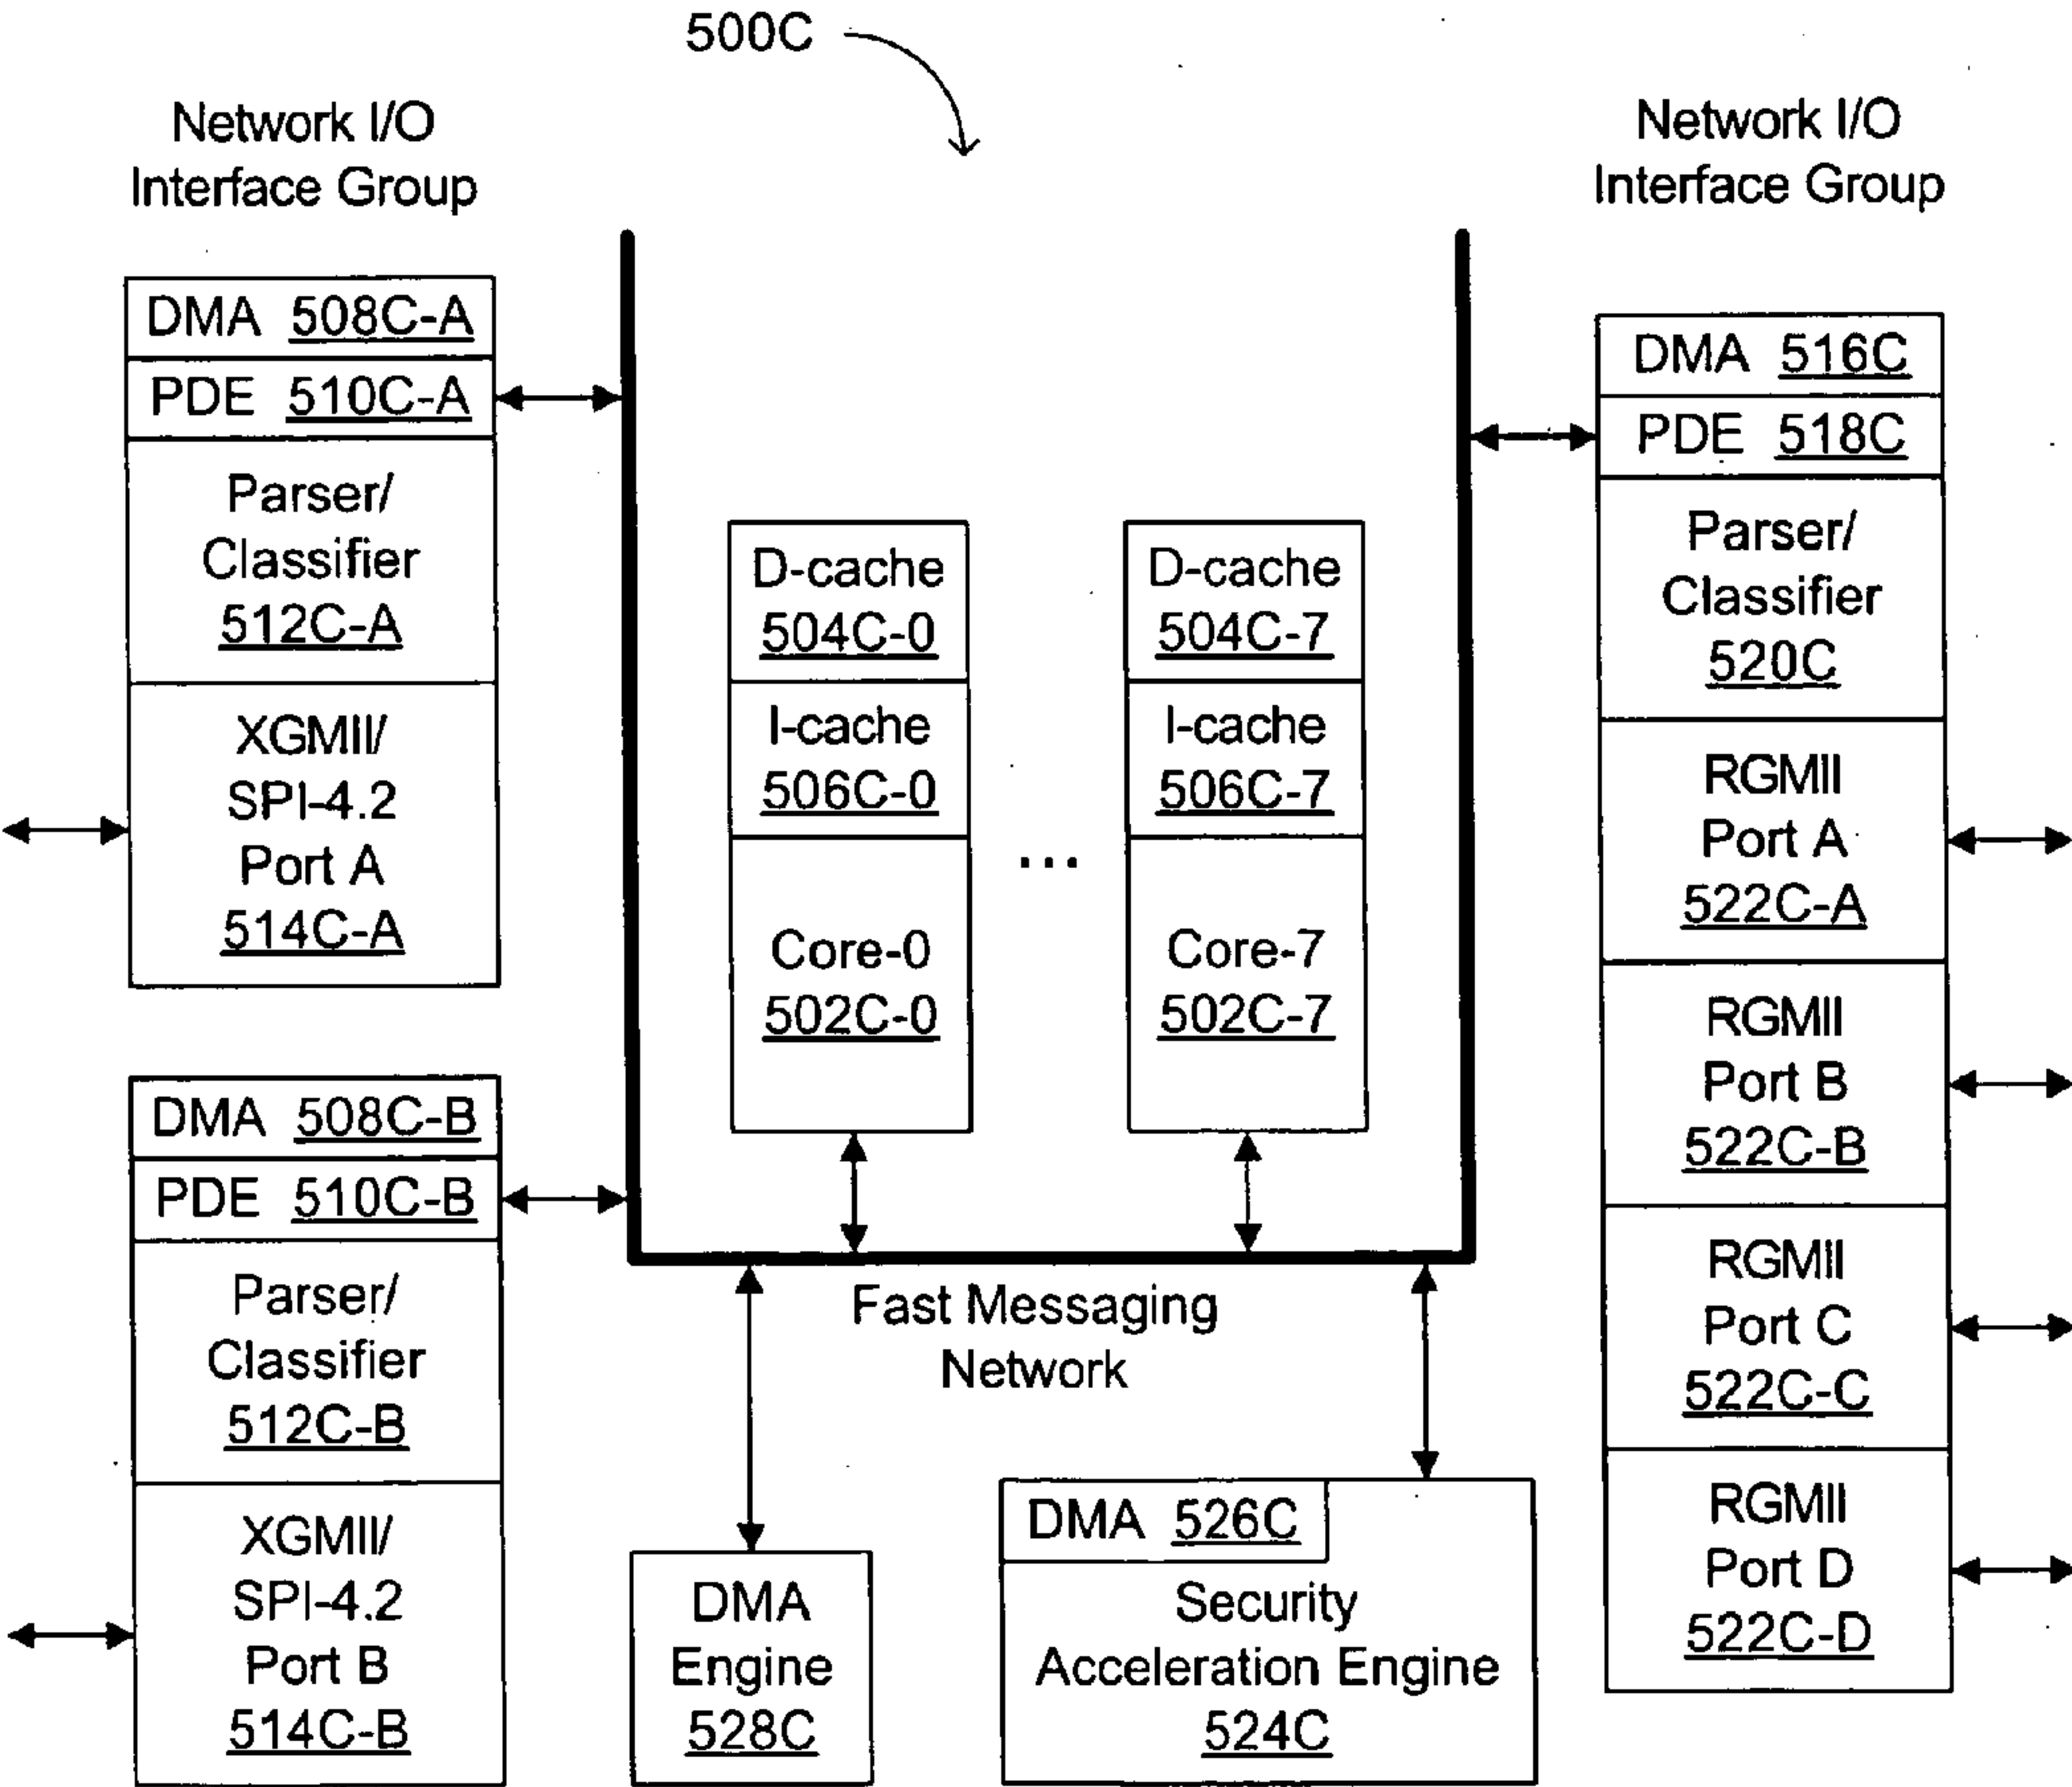



FIG. 5A



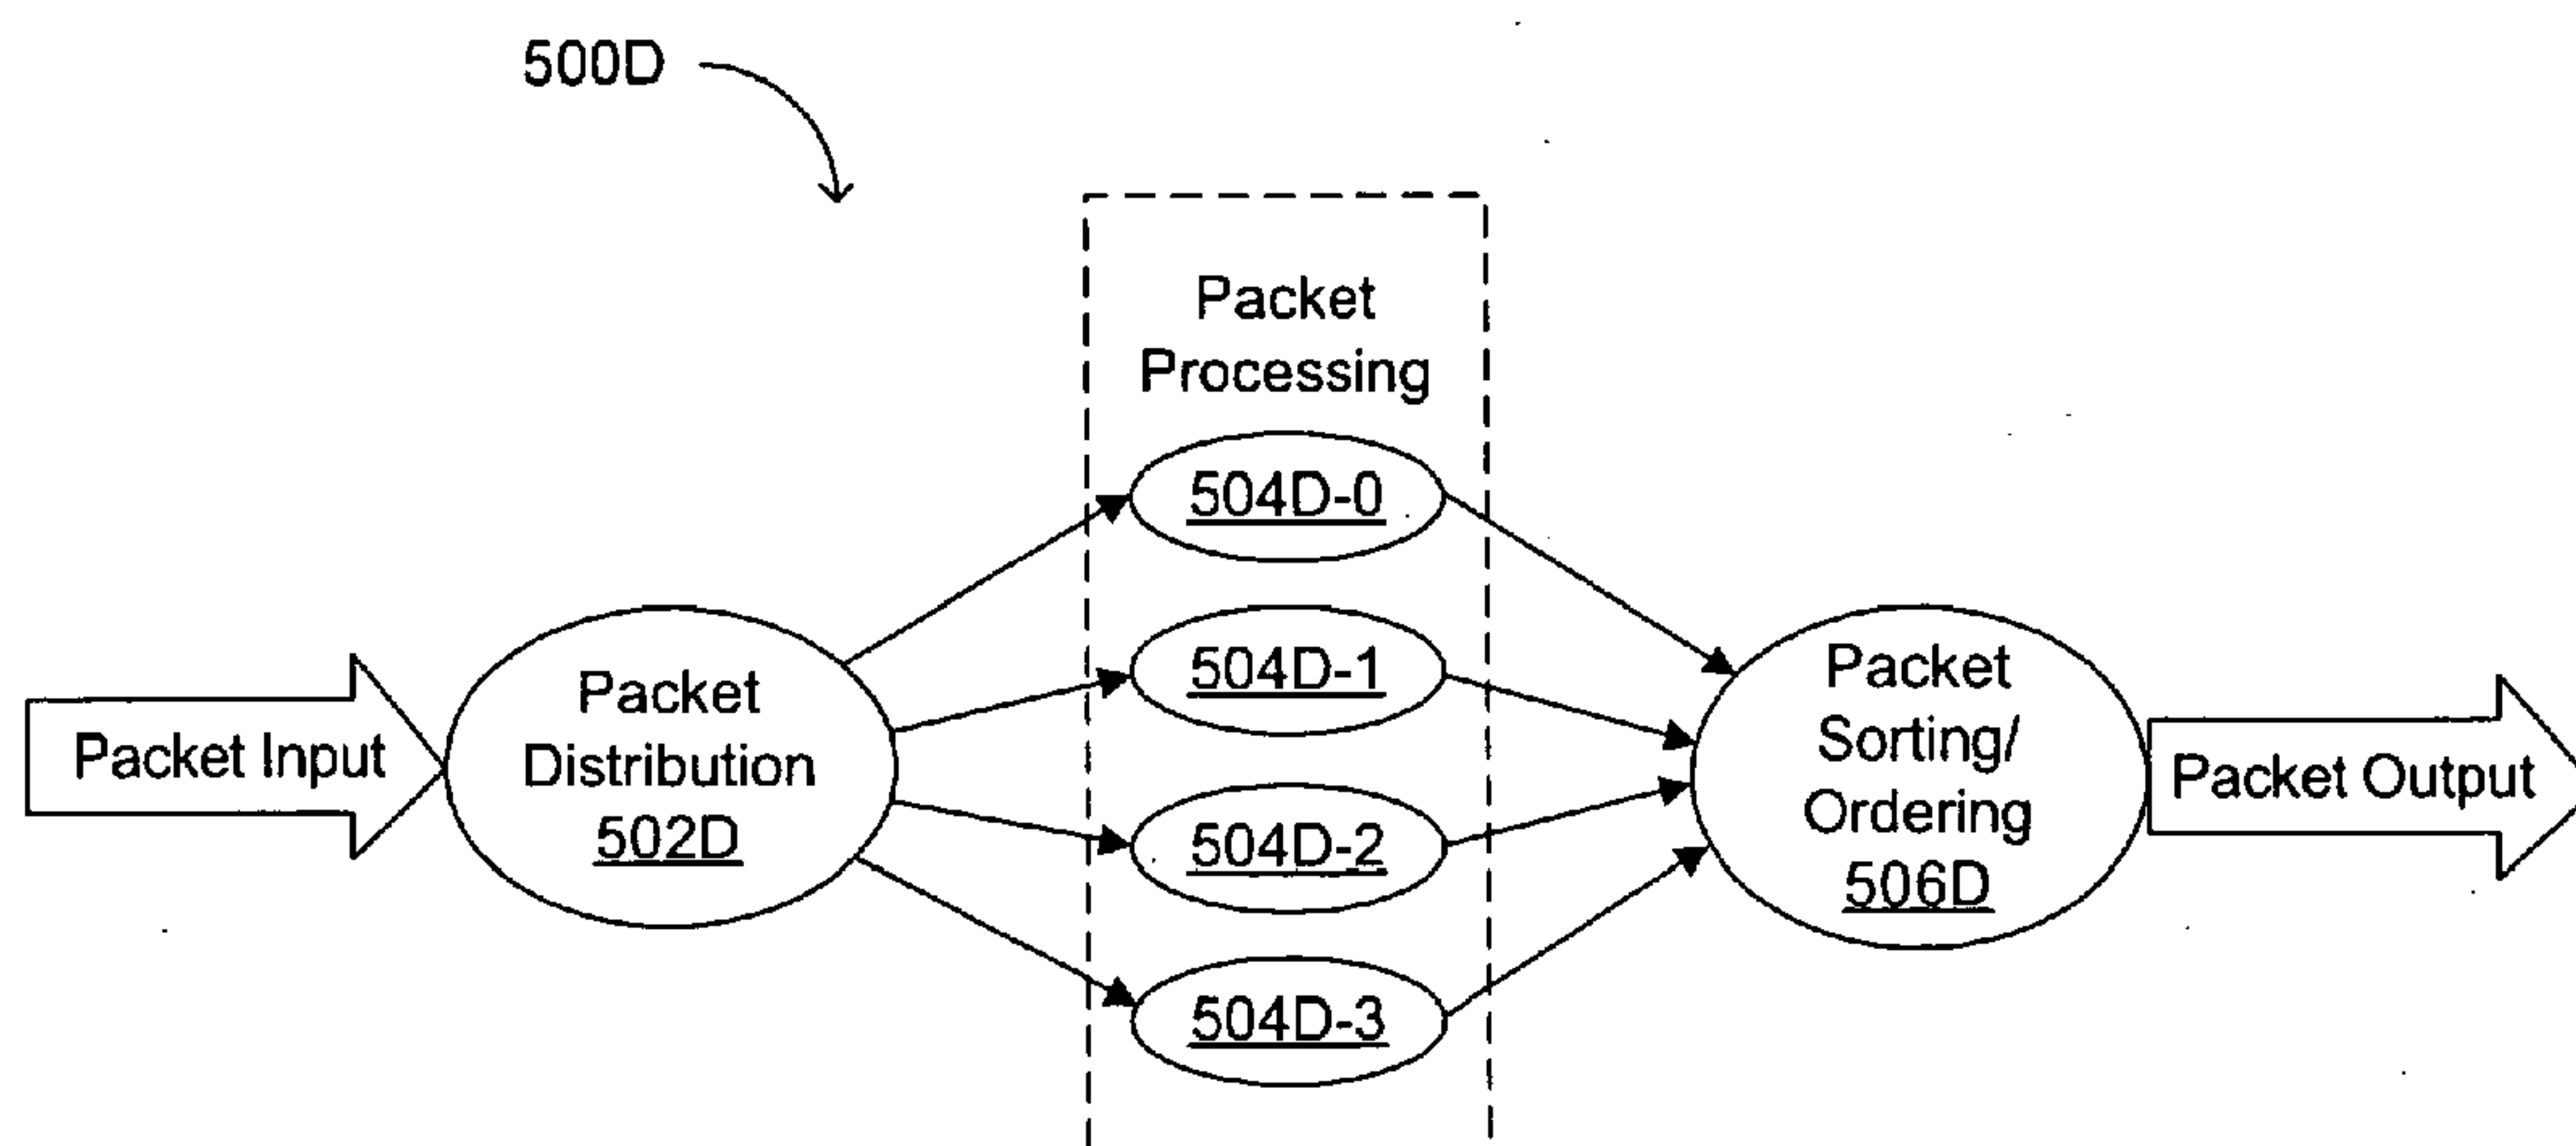


FIG. 5D (conventional)

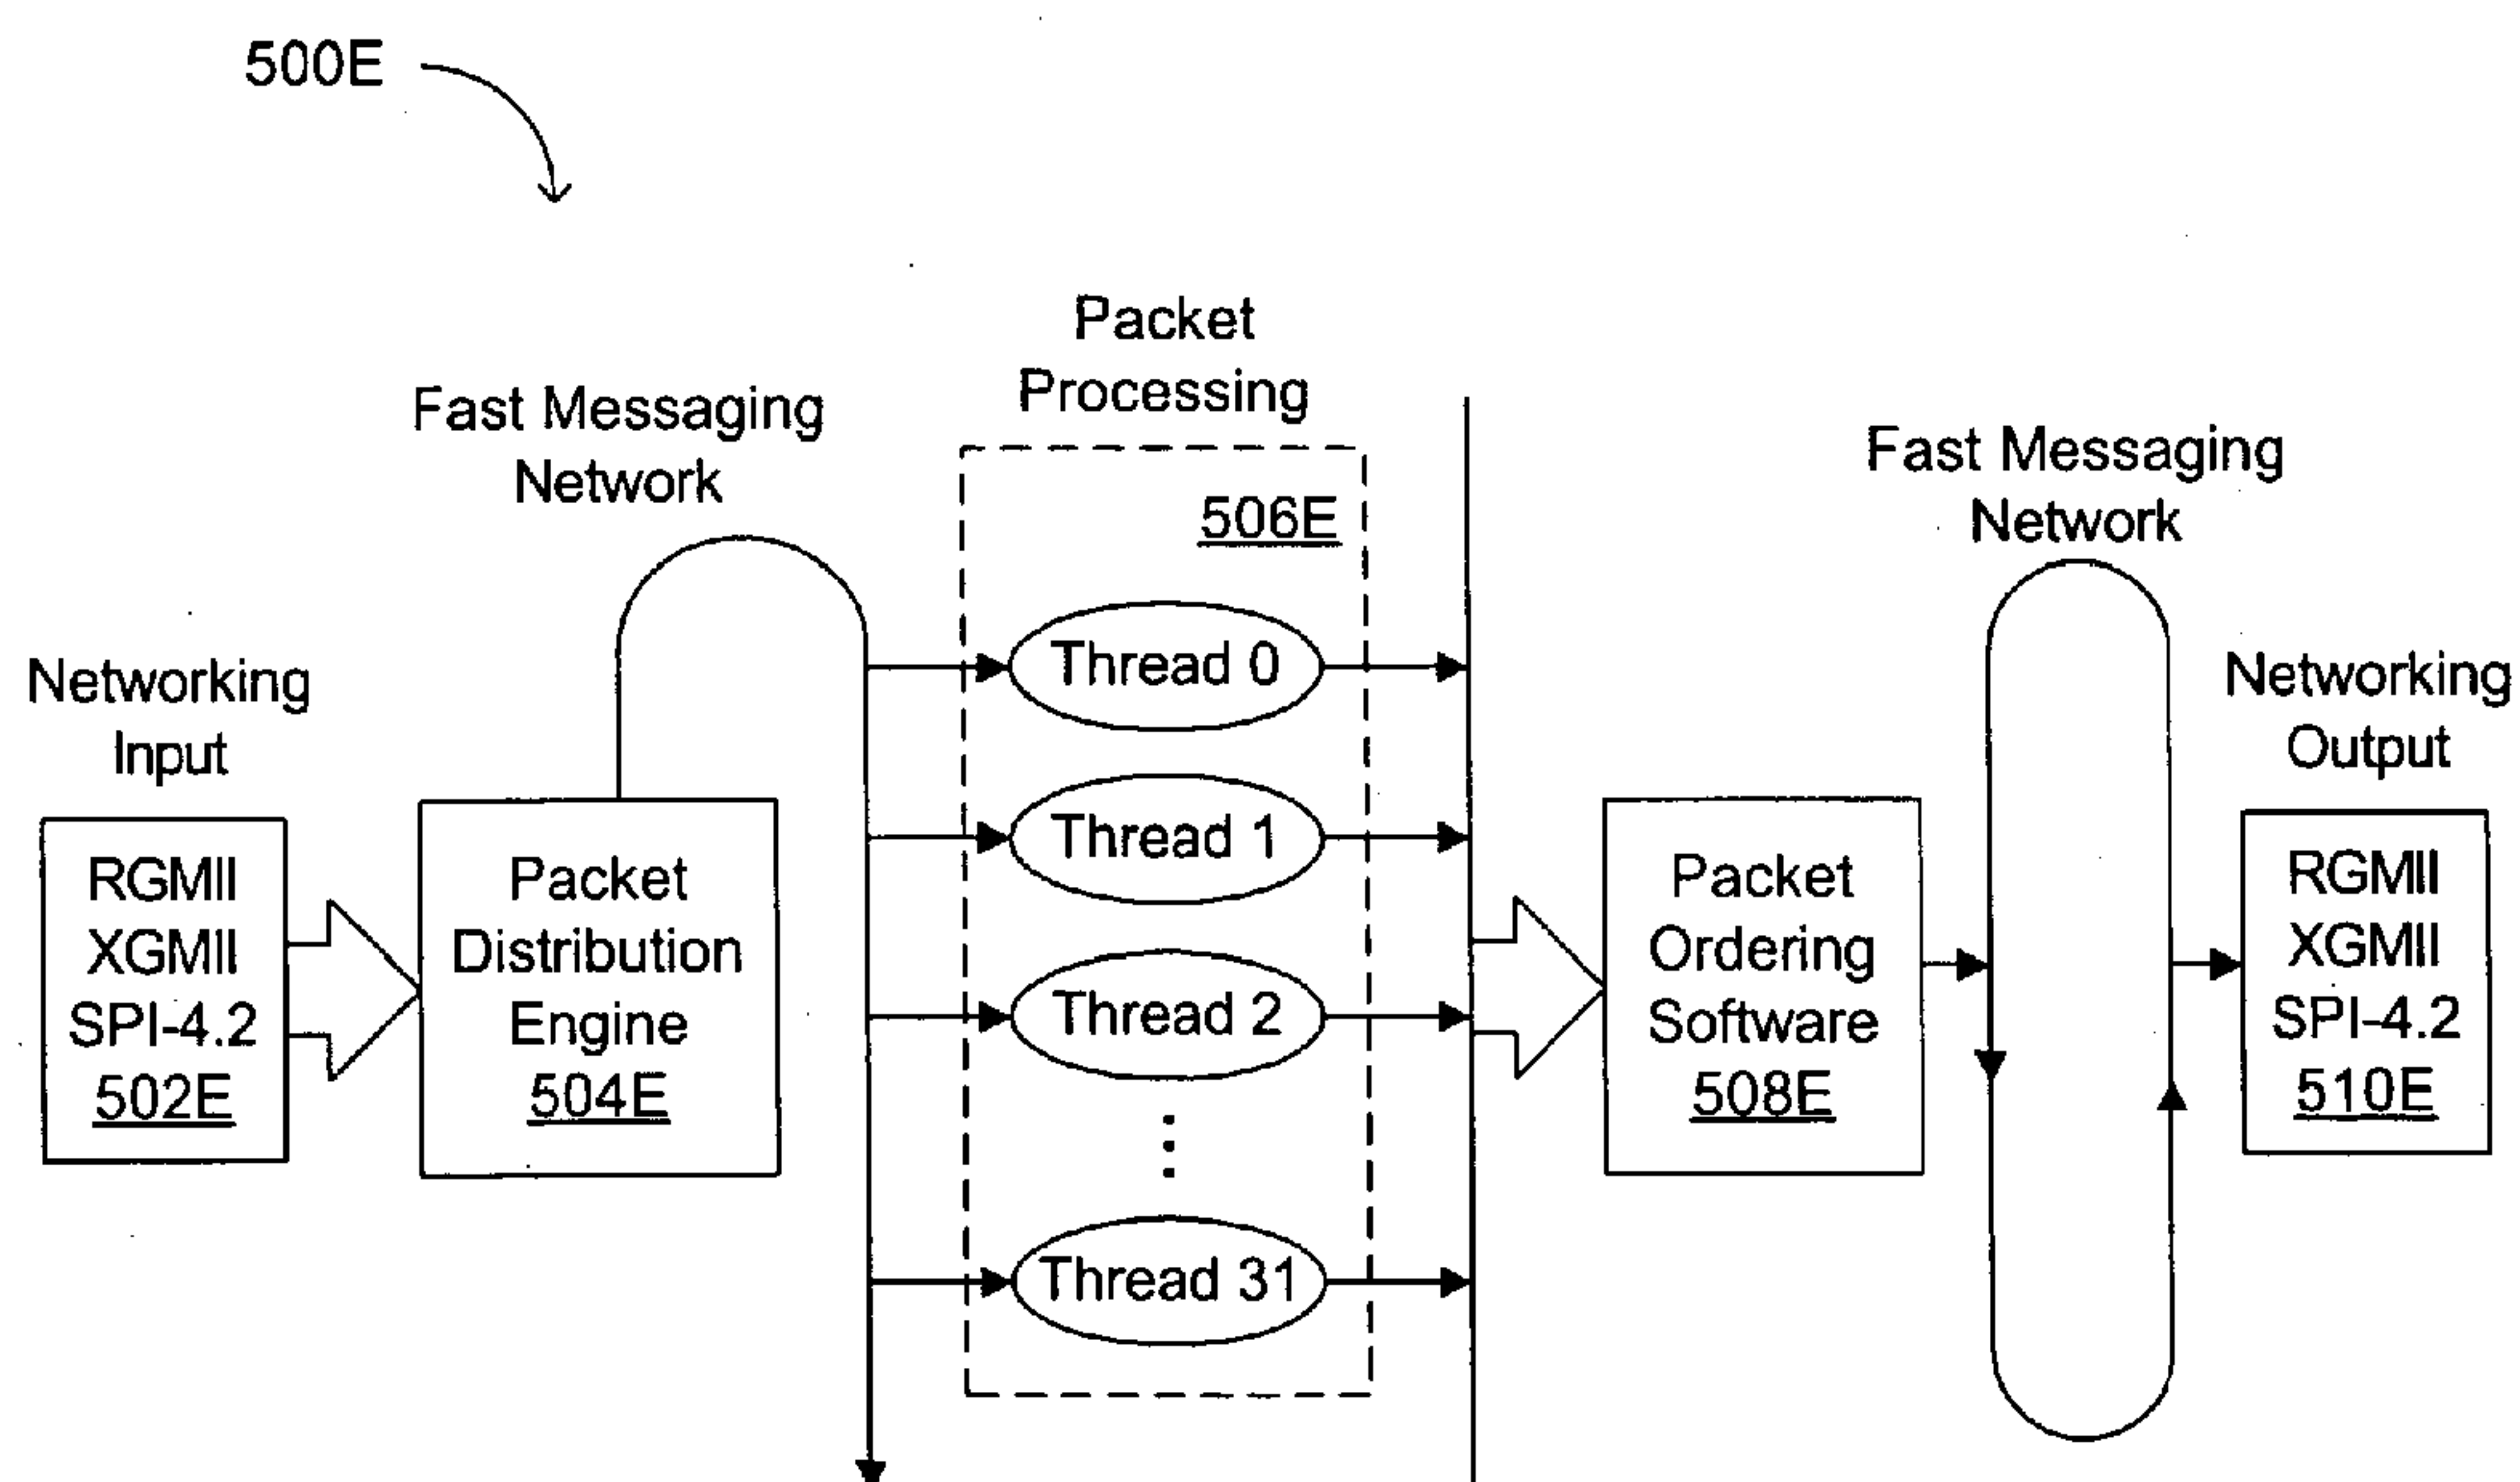


FIG. 5E

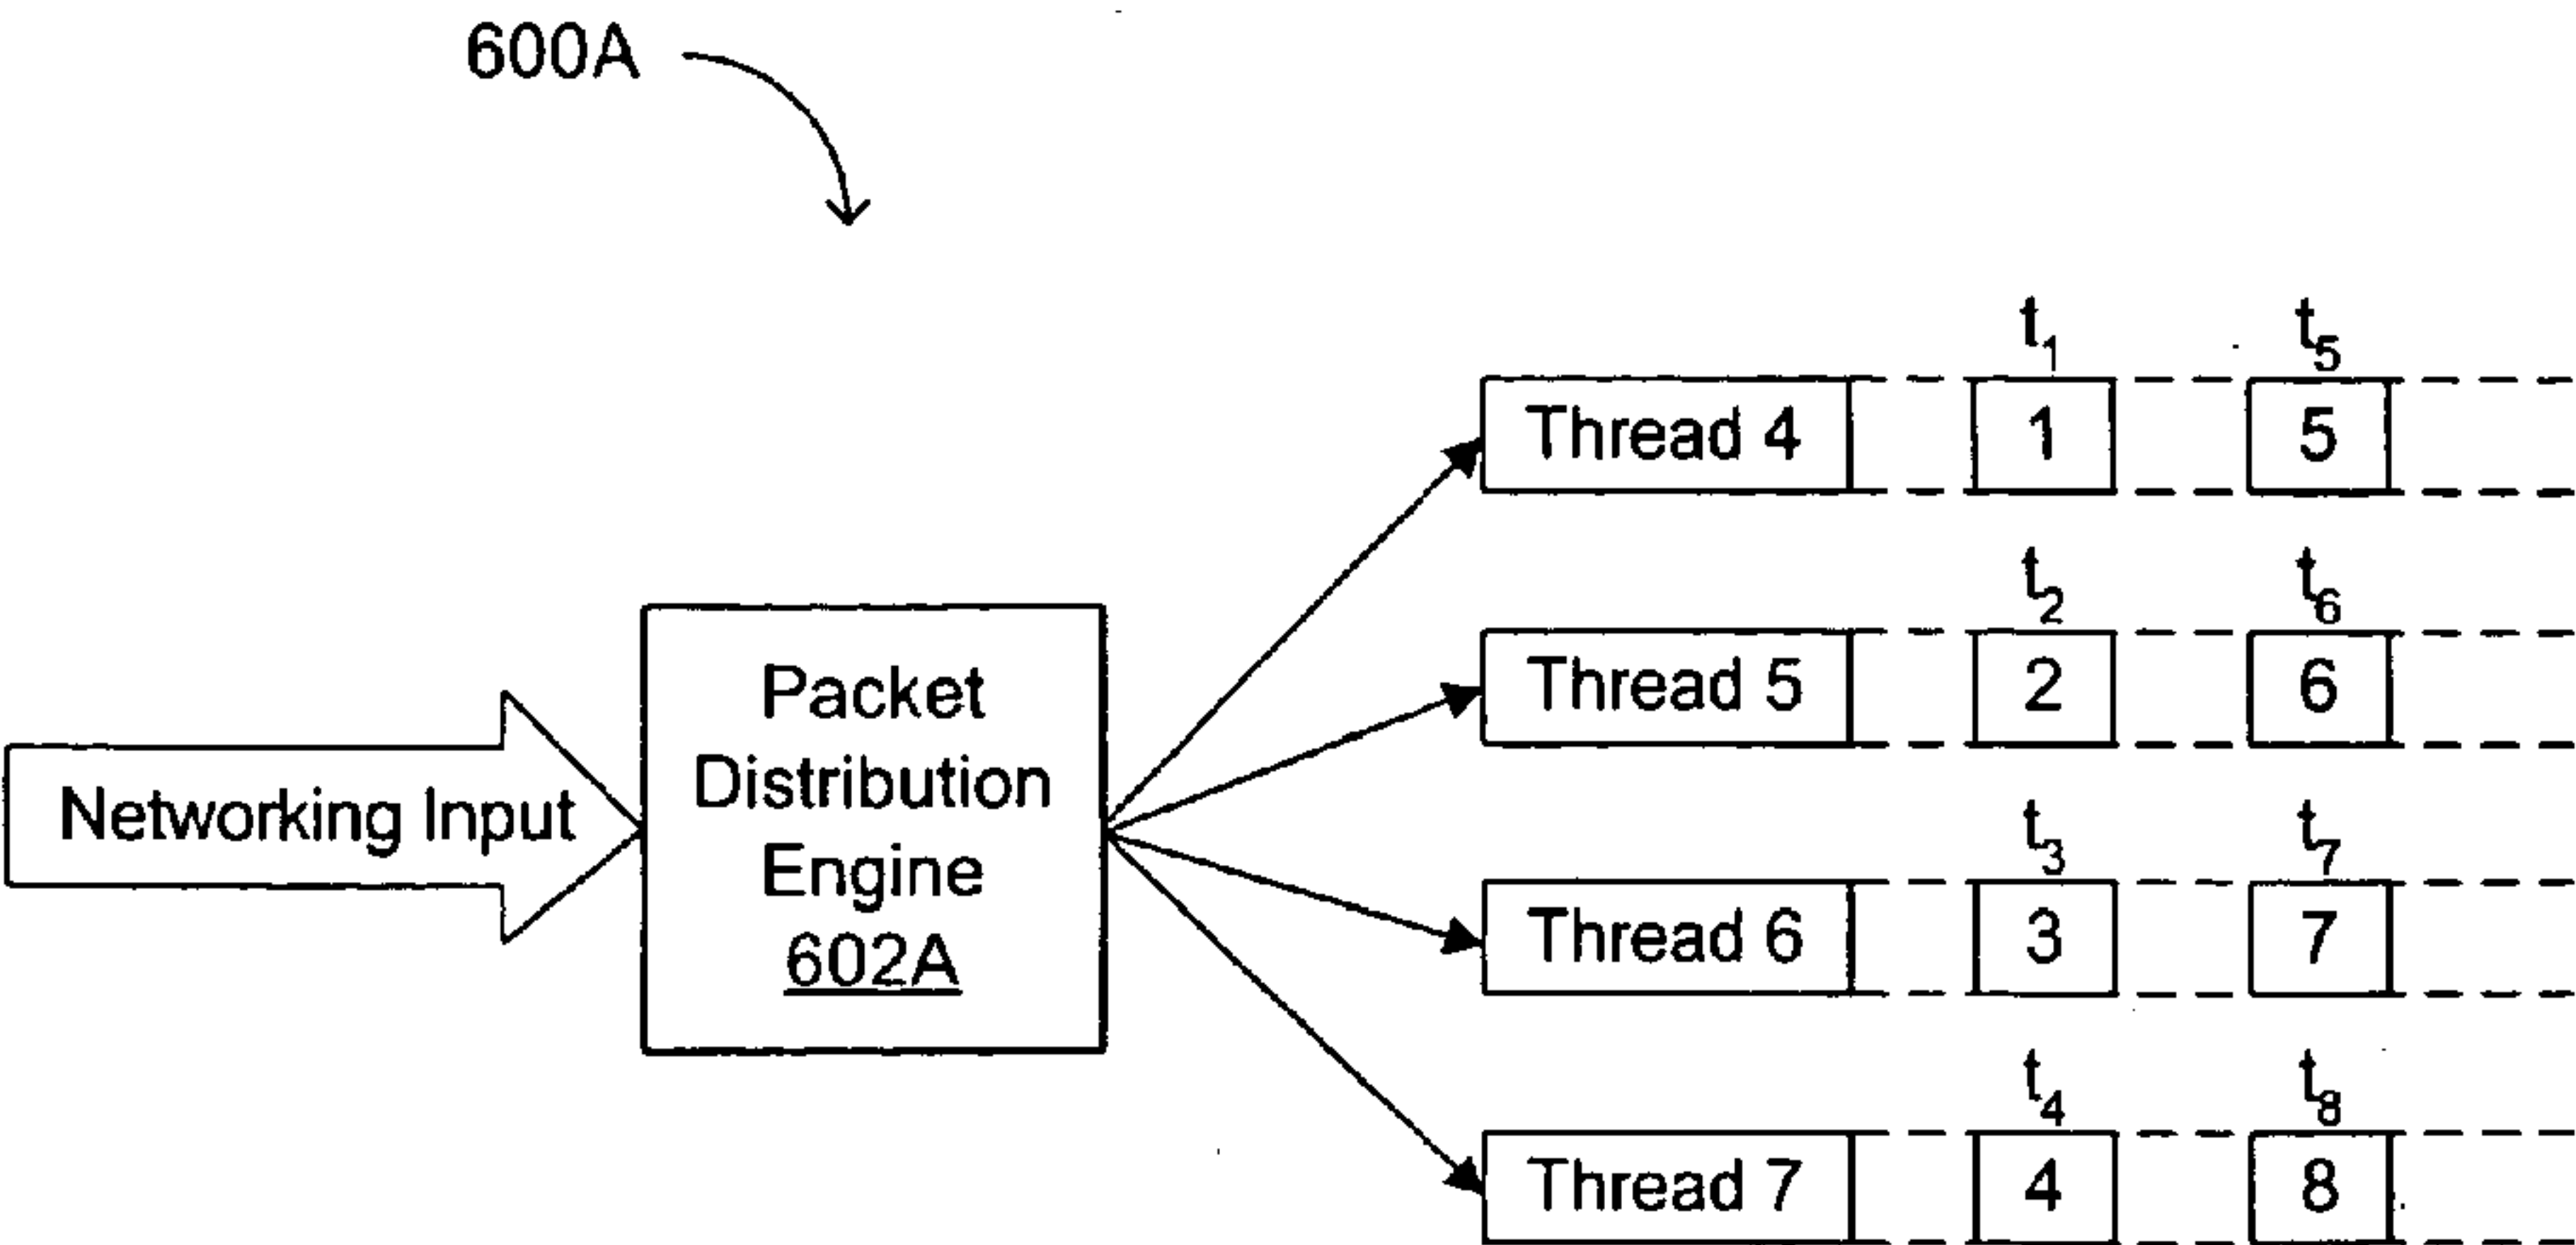


FIG. 6A

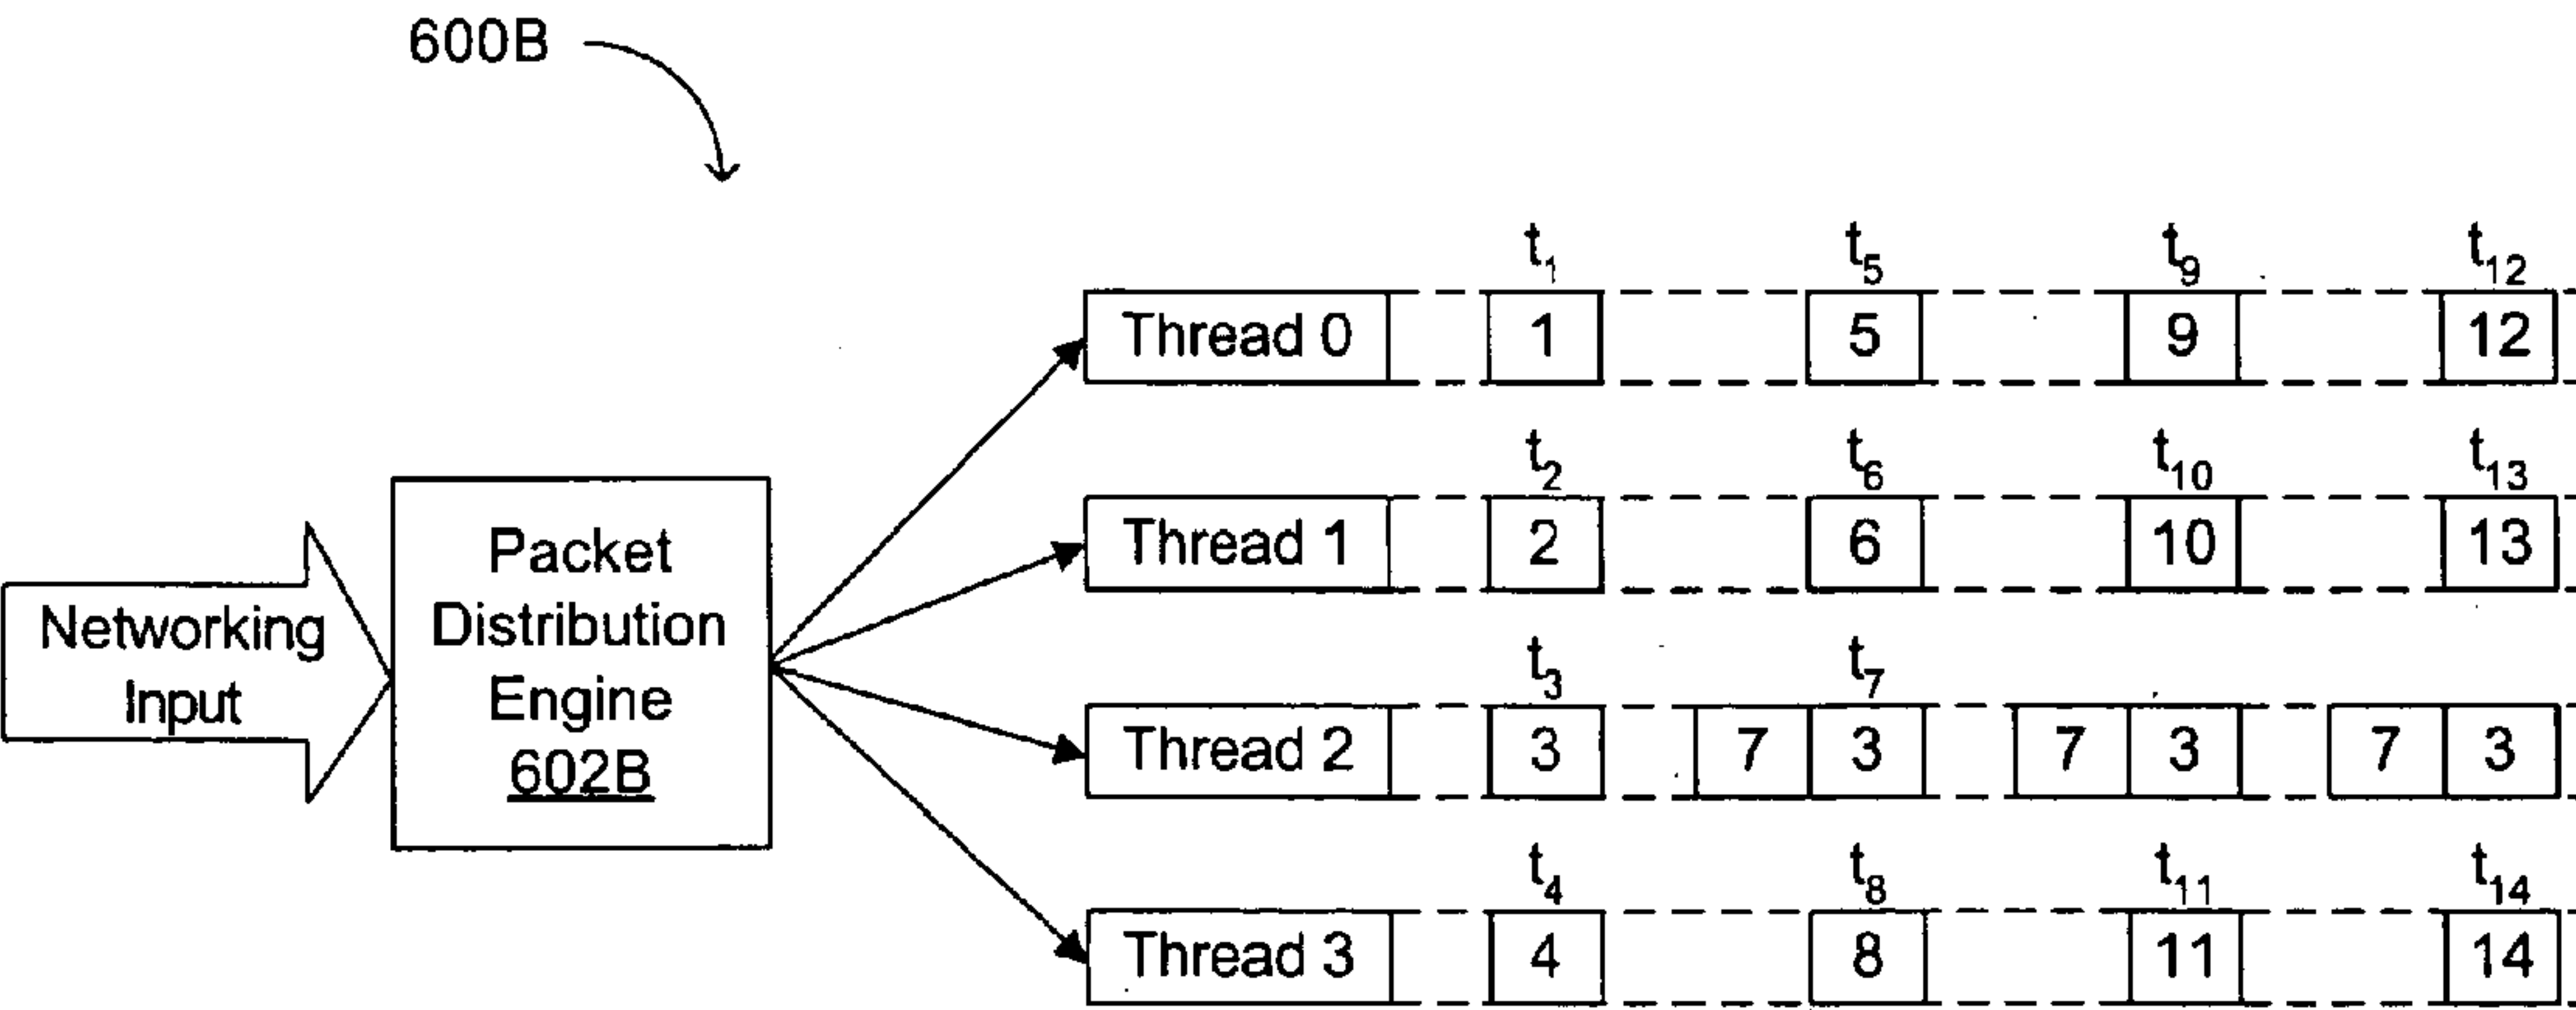


FIG. 6B

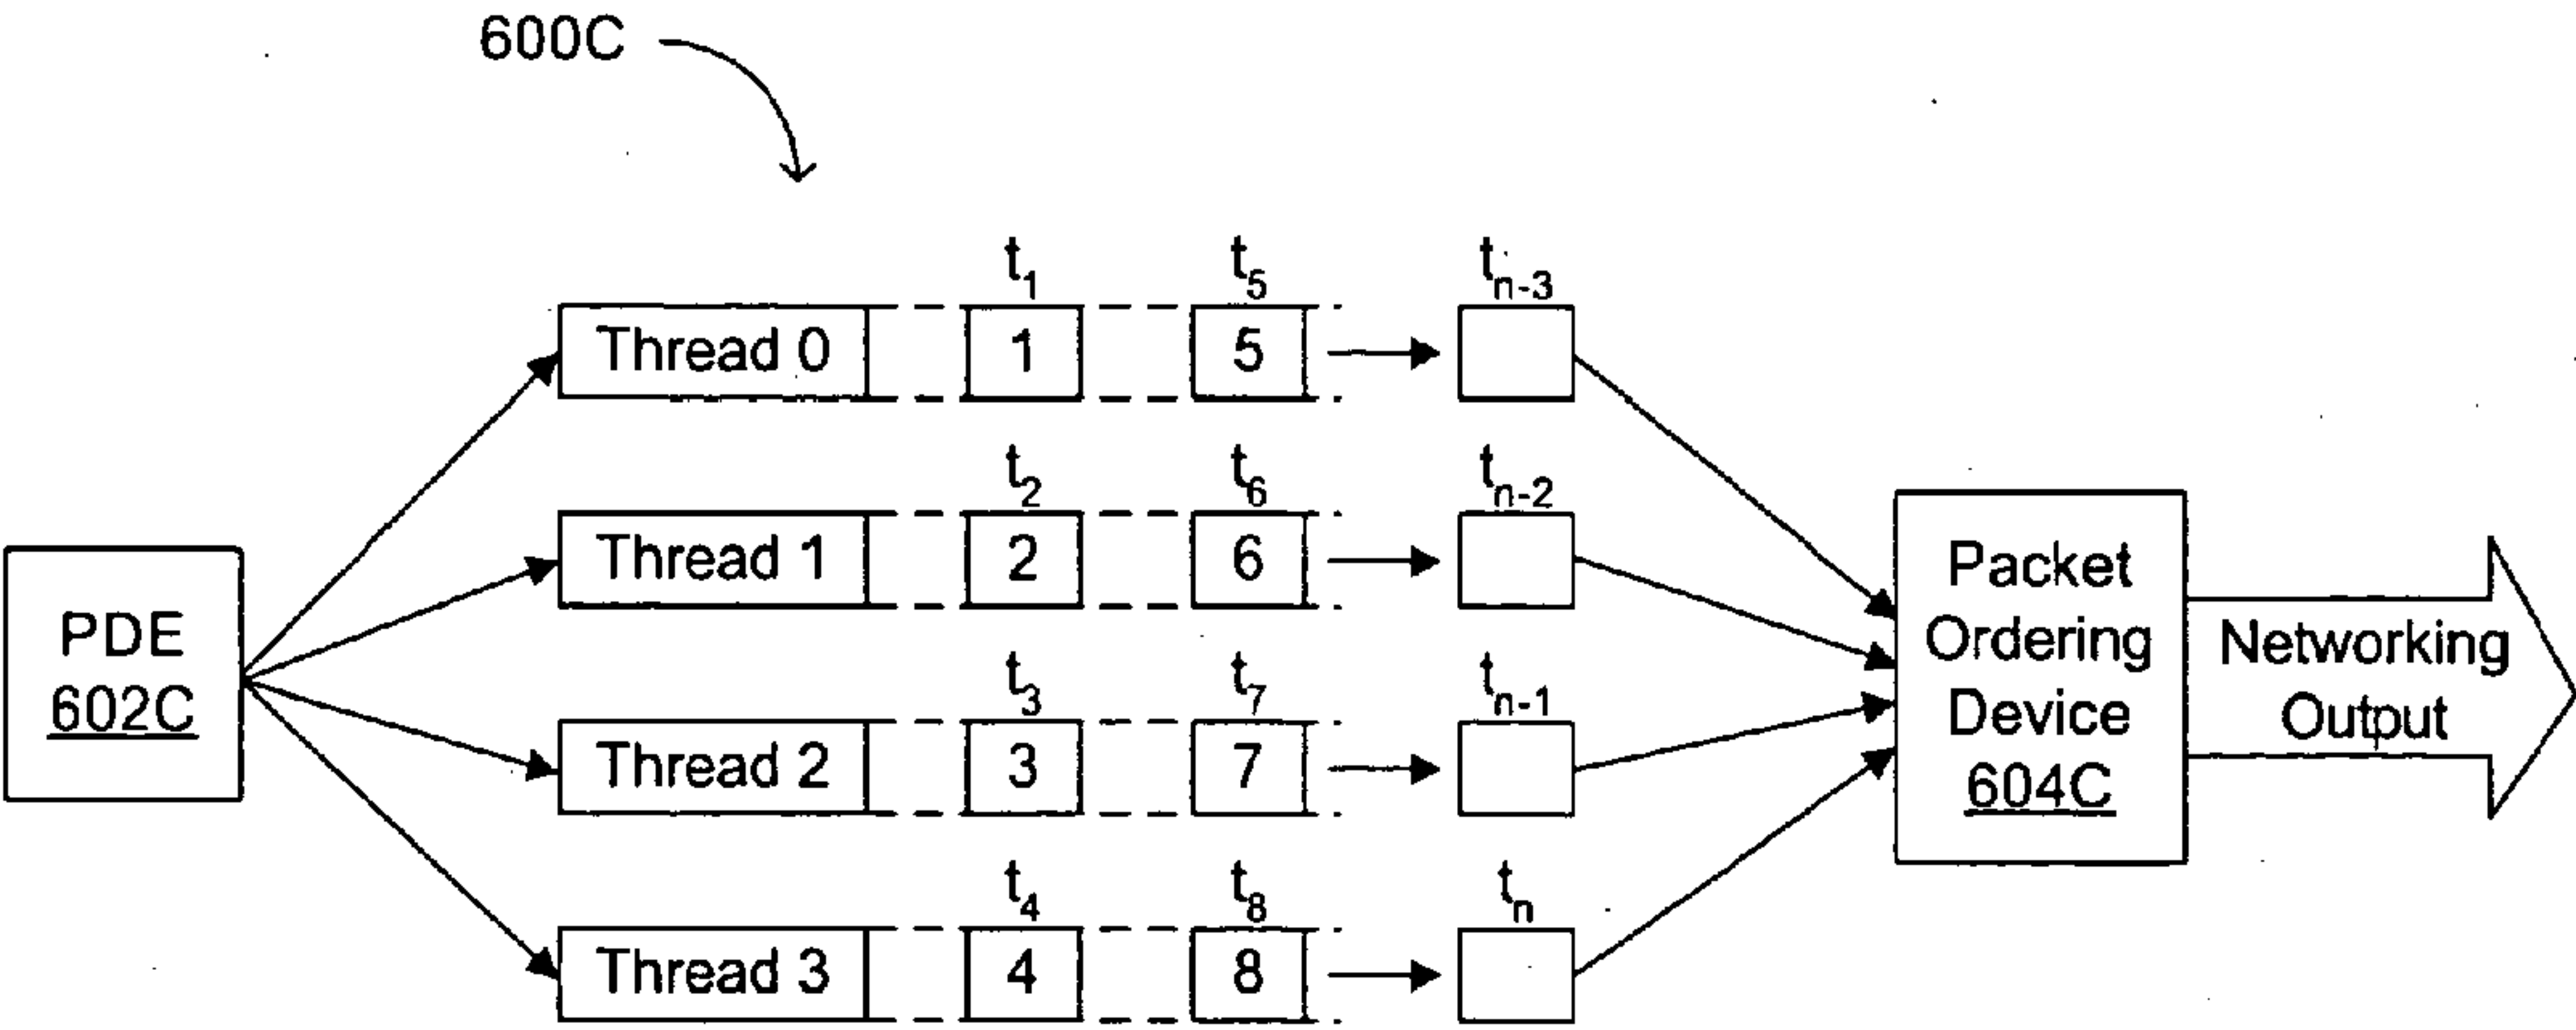


FIG. 6C




FIG. 6D

ADVANCED PROCESSOR WITH MECHANISM FOR FAST PACKET QUEUING OPERATIONS

RELATED APPLICATIONS

[0001] This application is a continuation in part of Ser. No. 10/898,008 filed Jul. 23, 2004 (RZMI-P102), which is a continuation in part of Ser. No. 10/682,579 filed Oct. 8, 2003 (RZMI-P101) claiming priority to Prov. No. 60/490,236 filed Jul. 25, 2003 (RZMI-P101P2) and Prov. No. 60/416,838 filed Oct. 8, 2002 (RZMI-P101P1), all incorporated herein by reference and all priorities claimed.

FIELD

[0002] The invention relates to the field of computers and telecommunications, and more particularly to an advanced processor for use in computers and telecommunications applications.

BACKGROUND

[0003] Modern computers and telecommunications systems provide great benefits including the ability to communicate information around the world. Conventional architectures for computers and telecommunications equipment include a large number of discrete circuits, which causes inefficiencies in both the processing capabilities and the communication speed.

[0004] For example, FIG. 1 depicts such a conventional line card employing a number of discrete chips and technologies. In FIG. 1, conventional line card **100** includes the following discrete components: Classification **102**, Traffic Manager **104**, Buffer Memory **106**, Security Co-Processor **108**, Transmission Control Protocol (TCP)/Internet Protocol (IP) Offload Engine **110**, L3+ Co-Processor **112**, Physical Layer Device (PHY) **114**, Media Access Control (MAC) **116**, Packet Forwarding Engine **118**, Fabric Interface Chip **120**, Control Processor **122**, Dynamic Random-Access Memory (DRAM) **124**, Access Control List (ACL) Ternary Content-Addressable Memory (TCAM) **126**, and Multiprotocol Label Switching (MPLS) Static Random-Access Memory (SRAM) **128**. The card further includes Switch Fabric **130**, which may connect with other cards and/or data.

[0005] Advances in processors and other components have improved the ability of telecommunications equipment to process, manipulate, store, retrieve and deliver information. Recently, engineers have begun to combine functions into integrated circuits to reduce the overall number of discrete integrated circuits, while still performing the required functions at equal or better levels of performance. This combination has been spurred by the ability to increase the number of transistors on a chip with new technology and the desire to reduce costs. Some of these combined integrated circuits have become so highly functional that they are often referred to as a System on a Chip (SoC). However, combining circuits and systems on a chip can become very complex and pose a number of engineering challenges. For example, hardware engineers want to ensure flexibility for future designs and software engineers want to ensure that their software will run on the chip and future designs as well.

[0006] The demand for sophisticated new networking and communications applications continues to grow in advanced switching and routing. In addition, solutions such as content-aware networking, highly integrated security, and new forms of storage management are beginning to migrate into flexible multi-service systems. Enabling technologies for these and

other next generation solutions must provide intelligence and high performance with the flexibility for rapid adaptation to new protocols and services.

[0007] Consequently, what is needed is an advanced processor that can take advantage of the new technologies while also providing high performance functionality. Additionally, this technology would be especially helpful if it included flexible modification ability.

SUMMARY

[0008] The present invention provides useful novel structures and techniques for overcoming the identified limitations, and provides an advanced processor that can take advantage of new technologies while also providing high performance functionality with flexible modification ability. The invention employs an advanced architecture System on a Chip (SoC) including modular components and communication structures to provide a high performance device.

[0009] An advanced processor comprises a plurality of multithreaded processor cores each having a data cache and instruction cache. A data switch interconnect is coupled to each of the processor cores and configured to pass information among the processor cores. A messaging network is coupled to each of the processor cores and a plurality of communication ports.

[0010] In one aspect of the invention, the data switch interconnect is coupled to each of the processor cores by its respective data cache, and the messaging network is coupled to each of the processor cores by its respective message station.

[0011] In one aspect of the invention, the advanced telecommunications processor further comprises a level 2 (L2) cache coupled to the data switch interconnect and configured to store information accessible to the processor cores.

[0012] In one aspect of the invention, the advanced telecommunications processor further comprises an interface switch interconnect coupled to the messaging network and the plurality of communication ports and configured to pass information among the messaging network and the communication ports.

[0013] In one aspect of the invention, the advanced telecommunications processor further comprises a memory bridge coupled to the data switch interconnect and at least one communication port, and that is configured to communicate with the data switch interconnect and the communication port.

[0014] In one aspect of the invention, the advanced telecommunications processor further comprises a super memory bridge coupled to the data switch interconnect, the interface switch interconnect and at least one communication port, and is configured to communicate with the data switch interconnect, the interface switch interconnect and the communication port.

[0015] Advantages of the invention include the ability to provide high bandwidth communications between computer systems and memory in an efficient and cost-effective manner.

BRIEF DESCRIPTION OF THE FIGURES

[0016] The invention is described with reference to the FIGS., in which:

[0017] FIG. 1 depicts a conventional line card;

[0018] FIG. 2A depicts an exemplary advanced processor according to an embodiment of the invention;

[0019] FIG. 2B depicts an exemplary advanced processor according to an alternate embodiment of the invention;

[0020] FIG. 3A depicts a conventional single-thread single-issue processing;

[0021] FIG. 3B depicts a conventional simple multi-threaded scheduling;

[0022] FIG. 3C depicts a conventional simple multi-threaded scheduling with a stalled thread;

[0023] FIG. 3D depicts an eager round-robin scheduling according to an embodiment of the invention;

[0024] FIG. 3E depicts a multithreaded fixed-cycle scheduling according to an embodiment of the invention;

[0025] FIG. 3F depicts a multithreaded fixed-cycle with eager round-robin scheduling according to an embodiment of the invention;

[0026] FIG. 3G depicts a core with associated interface units according to an embodiment of the invention;

[0027] FIG. 3H depicts an example pipeline of the processor according to embodiments of the invention;

[0028] FIG. 3I depicts a core interrupt flow operation within a processor according to an embodiment of the invention;

[0029] FIG. 3J depicts a programmable interrupt controller (PIC) operation according to an embodiment of the invention;

[0030] FIG. 3K depicts a return address stack (RAS) operation for multiple thread allocation according to an embodiment of the invention;

[0031] FIG. 4A depicts a data switch interconnect (DSI) ring arrangement according to an embodiment of the invention;

[0032] FIG. 4B depicts a DSI ring component according to an embodiment of the invention;

[0033] FIG. 4C depicts a flow diagram of an example data retrieval in the DSI according to an embodiment of the invention;

[0034] FIG. 5A depicts a fast messaging ring component according to an embodiment of the invention;

[0035] FIG. 5B depicts a message data structure for the system of FIG. 5A;

[0036] FIG. 5C depicts a conceptual view of various agents attached to the fast messaging network (FMN) according to an embodiment of the invention;

[0037] FIG. 5D depicts network traffic in a conventional processing system;

[0038] FIG. 5E depicts packet flow according to an embodiment of the invention;

[0039] FIG. 6A depicts a packet distribution engine (PDE) distributing packets evenly over four threads according to an embodiment of the invention;

[0040] FIG. 6B depicts a PDE distributing packets using a round-robin scheme according to an embodiment of the invention;

[0041] FIG. 6C depicts a packet ordering device (POD) placement during packet lifecycle according to an embodiment of the invention;

[0042] FIG. 6D depicts a POD outbound distribution according to an embodiment of the invention;

DETAILED DESCRIPTION

[0043] The invention is described with reference to specific architectures and protocols. Those skilled in the art will recognize that the description is for illustration and to provide the best mode of practicing the invention. The description is not meant to be limiting and references to telecommunications and other applications may be equally applicable to general computer applications, for example, server applications, dis-

tributed shared memory applications and so on. As described herein, reference is made to Ethernet Protocol, Internet Protocol, Hyper Transport Protocol and other protocols, but the invention may be applicable to other protocols as well. Moreover, reference is made to chips that contain integrated circuits while other hybrid or meta-circuits combining those described in chip form is anticipated. Additionally, reference is made to an exemplary MIPS architecture and instruction set, but other architectures and instruction sets can be used in the invention. Other architectures and instruction sets include, for example, x86, PowerPC, ARM and others.

A. ARCHITECTURE

[0044] The invention is designed to consolidate a number of the functions performed on the conventional line card of FIG. 1, and to enhance the line card functionality. In one embodiment, the invention is an integrated circuit that includes circuitry for performing many discrete functions. The integrated circuit design is tailored for communication processing. Accordingly, the processor design emphasizes memory intensive operations rather than computationally intensive operations. The processor design includes an internal network configured for high efficient memory access and threaded processing as described below.

[0045] FIG. 2A depicts an exemplary advanced processor (200) according to an embodiment of the invention. The advanced processor is an integrated circuit that can perform many of the functions previously tasked to specific integrated circuits. For example, the advanced processor includes a packet forwarding engine, a level 3 co-processor and a control processor. The processor can include other components, as desired. As shown herein, given the number of exemplary functional components, the power dissipation is approximately 20 watts in the exemplary embodiment. Of course, in other embodiments of the invention, the power dissipation may be more or less than about 20 watts.

[0046] The exemplary processor is designed as a network on a chip. This distributed processing architecture allows components to communicate with one another and not necessarily share a common clock rate. For example, one processor component could be clocked at a relatively high rate while another processor component is clocked at a relatively low rate. The network architecture further supports the ability to add other components in future designs by simply adding the component to the network. For example, if a future communication interface is desired, that interface can be laid out on the processor chip and coupled to the processor network. Then, future processors can be fabricated with the new communication interface.

[0047] The design philosophy is to create a processor that can be programmed using general purpose software tools and reusable components. Several exemplary features that support this design philosophy include: static gate design; low-risk custom memory design; flip-flop based design; design-for-testability including a full scan, memory built-in self-test (BIST), architecture redundancy and tester support features; reduced power consumption including clock gating; logic gating and memory banking; datapath and control separation including intelligently guided placement; and rapid feedback of physical implementation.

[0048] The software philosophy is to enable utilization of industry standard development tools and environment. The desire is to program the processing using general purpose software tools and reusable components. The industry stan-

standard tools and environment include familiar tools, such as gcc/gdb and the ability to develop in an environment chosen by the customer or programmer.

[0049] The desire is also to protect existing and future code investment by providing a hardware abstraction layer (HAL) definition. This enables relatively easy porting of existing applications and code compatibility with future chip generations.

[0050] Turning to the CPU core, the core is designed to be MIPS64 compliant and have a frequency target in the range of approximately 1.5 GHz+. Additional exemplary features supporting the architecture include: 4-way multithreaded single issue 10-stage pipeline; real time processing support including cache line locking and vectored interrupt support; 32 KB 4-way set associative instruction cache; 32 KB 4-way set associative data cache; and 128-entry translation-lookaside buffer (TLB).

[0051] One of the important aspects of the exemplary embodiment is the high-speed processor input/output (I/O), which is supported by: two XGMII/SPI-4 (e.g., boxes **228a** and **228b** of FIG. 2A); three 1 Gb MACs; one 16-bit Hyper-Transport (e.g., box **232**) that can scale to 800/1600 MHz memory, including one Flash portion (e.g., box **226** of FIG. 2A) and two Quad Data Rate (QDR2)/Double Data Rate (DDR2) SRAM portions; two 64-bit DDR2 channels that can scale to 400/800 MHz; and communication ports including 32-bit Peripheral Component Interconnect (PCI) (e.g., box **234** of FIG. 2A), Joint Test Access Group (JTAG) and Universal Asynchronous Receiver/Transmitter (UART) (e.g., box **226**).

[0052] Also included as part of the interface are two Reduced GMII (RGMII) (e.g., boxes **230a** and **230b** of FIG. 2A) ports. Further, Security Acceleration Engine (SAE) (e.g., box **238** of FIG. 2A) can use hardware-based acceleration for security functions, such as encryption, decryption, authentication, and key generation. Such features can help software deliver high performance security applications, such as IPsec and SSL.

[0053] The architecture philosophy for the CPU is to optimize for thread level parallelism (TLP) rather than instruction level parallelism (ILP) including networking workloads benefit from TLP architectures, and keeping it small.

[0054] The architecture allows for many CPU instantiations on a single chip, which in turn supports scalability. In general, super-scalar designs have minimal performance gains on memory bound problems. An aggressive branch prediction is typically unnecessary for this type of processor application and can even be wasteful.

[0055] The exemplary embodiment employs narrow pipelines because they typically have better frequency scalability. Consequently, memory latency is not as much of an issue as it would be in other types of processors, and in fact, any memory latencies can effectively be hidden by the multithreading, as described below.

[0056] Embodiments of the invention can optimize the memory subsystem with non-blocking loads, memory reordering at the CPU interface, and special instruction for semaphores and memory barriers.

[0057] In one aspect of the invention, the processor can acquire and release semantics added to load/stores. In another aspect of embodiments of the invention, the processor can employ special atomic incrementing for timer support.

[0058] As described above, the multithreaded CPUs offer benefits over conventional techniques. An exemplary

embodiment of the invention employs fine grained multithreading that can switch threads every clock and has 4 threads available for issue.

[0059] The multithreading aspect provides for the following advantages: usage of empty cycles caused by long latency operations; optimized for area versus performance trade-off; ideal for memory bound applications; enable optimal utilization of memory bandwidth; memory subsystem; cache coherency using MOSI (Modified, Own, Shared, Invalid) protocol; full map cache directory including reduced snoop bandwidth and increased scalability over broadcast snoop approach; large on-chip shared dual banked 2 MB L2 cache; error checking and correcting (ECC) protected caches and memory; 2 64-bit 400/800 DDR2 channels (e.g., 12.8 GByte/s peak bandwidth) security pipeline; support of on-chip standard security functions (e.g., AES, DES/3DES, SHA-1, MD5, and RSA); allowance of the chaining of functions (e.g., encrypt->sign) to reduce memory accesses; 4 Gbs of bandwidth per security pipeline, excluding RSA; on-chip switch interconnect; message passing mechanism for intra-chip communication; point-to-point connection between super-blocks to provide increased scalability over a shared bus approach; 16 byte full-duplex links for data messaging (e.g., 32 GB/s of bandwidth per link at 1 GHz); and credit-based flow control mechanism.

[0060] Some of the benefits of the multithreading technique used with the multiple processor cores include memory latency tolerance and fault tolerance.

[0061] FIG. 2B depicts an exemplary advanced processor according to an alternate embodiment of the invention. This embodiment is provided to show that the architecture can be modified to accommodate other components, for example, video processor **215**. In such a case, the video processor can communicate with the processor cores, communication networks (e.g. DSI and Messaging Network) and other components.

B. PROCESSOR CORES AND MULTI-THREADING

[0062] The exemplary advanced processor **200** of FIG. 2A includes a plurality of multithreaded processor cores **210a-h**. Each exemplary core includes an associated data cache **212a-h** and instruction cache **214a-h**. Data Switch Interconnect (DSI) **216** may be coupled to each of the processor cores **210a-h** and configured to pass data among the processor cores and between the L2 cache **208** and memory bridges **206**, **208** for main memory access. Additionally, a messaging network **222** may be coupled to each of the processor cores **210a-h** and a plurality of communication ports **240a-f**. While eight cores are depicted in FIG. 2A, a lesser or greater number of cores can be used in the invention. Likewise, in aspects of the invention, the cores can execute different software programs and routines, and even run different operating systems. The ability to run different software programs and operating systems on different cores within a single unified platform can be particularly useful where legacy software is desired to be run on one or more of the cores under an older operating system, and newer software is desired to be run on one or more other cores under a different operating system or systems. Similarly, as the exemplary processor permits multiple separate functions to be combined within a unified platform, the ability to run multiple different software and operating systems on

the cores means that the disparate software associated with the separate functions being combined can continue to be utilized.

[0063] The exemplary processor includes the multiple CPU cores **210a-h** capable of multithreaded operation. In the exemplary embodiment, there are eight 4-way multithreaded MIPS64-compatible CPUs, which are often referred to as processor cores. Embodiments of the invention can include 32 hardware contexts and the CPU cores may operate at over approximately 1.5 GHz. One aspect of the invention is the redundancy and fault tolerant nature of multiple CPU cores. So, for example, if one of the cores failed, the other cores would continue operation and the system would experience only slightly degraded overall performance. In one embodiment, a ninth processor core may be added to the architecture to ensure with a high degree of certainty that eight cores are functional.

[0064] The multithreaded core approach can allow software to more effectively use parallelism that is inherent in many packet processing applications. Most conventional processors use a single-issue, single-threaded architecture, but this has performance limitations in typical networking applications. In aspects of the invention, the multiple threads can execute different software programs and routines, and even run different operating systems. This ability, similar to that described above with respect to the cores, to run different software programs and operating systems on different threads within a single unified platform can be particularly useful where legacy software is desired to be run on one or more of the threads under an older operating system, and newer software is desired to be run on one or more other threads under a different operating system or systems. Similarly, as the exemplary processor permits multiple separate functions to be combined within a unified platform, the ability to run multiple different software and operating systems on the threads means that the disparate software associated with the separate functions being combined can continue to be utilized. Discussed below are some techniques used by the invention to improve performance in single and multithreaded applications.

[0065] Referring now to FIG. 3A, a conventional single-thread single-issue processing is shown and indicated by the general reference character **300A**. The cycle numbers are shown across the top of the blocks. “A” within the blocks can represent a first packet and “B” within the blocks can represent a next packet. The sub-numbers within the blocks can represent packet instructions and/or segments. The wasted cycles **5-10** after a cache miss, as shown, result from no other instructions being ready for execution. The system must essentially stall to accommodate the inherent memory latency and this is not desirable.

[0066] For many processors, performance is improved by executing more instructions per cycle, thus providing for instruction level parallelism (ILP). In this approach, more functional units are added to the architecture in order to execute multiple instructions per cycle. This approach is also known as a single-threaded, multiple-issue processor design. While offering some improvement over single-issue designs, performance typically continues to suffer due to the high-latency nature of packet processing applications in general. In particular, long-latency memory references usually result in similar inefficiency and increased overall capacity loss.

[0067] As an alternate approach, a multithreaded, single-issue architecture may be used. This approach takes advantage of, and more fully exploits, the packet level parallelism

commonly found in networking applications. In short, memory latencies can be effectively hidden by an appropriately designed multithreaded processor. Accordingly, in such a threaded design, when one thread becomes inactive while waiting for memory data to return, the other threads can continue to process instructions. This can maximize processor use by minimizing wasted cycles experienced by other simple multi-issue processors.

[0068] Referring now to FIG. 3B, a conventional simple multithreaded scheduling is shown and indicated by the general reference character **300B**. Instruction Scheduler (IS) **302B** can receive four threads: A, B, C, and D, as shown in the boxes to the left of IS **302B**. Each cycle can simply select a different packet instruction from each of the threads in “round-robin” fashion, as shown. This approach generally works well as long as every thread has an instruction available for issue. However, such a “regular” instruction issue pattern cannot typically be sustained in actual networking applications. Common factors, such as instruction cache miss, data cache miss, data use interlock, and non-availability of a hardware resource can stall the pipeline.

[0069] Referring now to FIG. 3C, a conventional simple multithreaded scheduling with a stalled thread is shown and indicated by the general reference character **300C**. Instruction Scheduler (IS) **302C** can receive four threads: A, B, and C, plus an empty “D” thread. As shown, conventional round-robin scheduling results in wasted cycles **4, 8, and 12**, the positions where instructions from the D thread would fall if available. In this example, the pipeline efficiency loss is 25% during the time period illustrated. An improvement over this approach that is designed to overcome such efficiency losses is the “eager” round-robin scheduling scheme.

[0070] Referring now to FIG. 3D, an eager round-robin scheduling according to an embodiment of the invention is shown and indicated by the general reference character **300D**. The threads and available instructions shown are the same as illustrated in FIG. 3C. However, in FIG. 3D, the threads can be received by an Eager Round-Robin Scheduler (ERRS) **302D**. The eager round-robin scheme can keep the pipeline full by issuing instructions from each thread in sequence as long as instructions are available for processing. When one thread is “sleeping” and does not issue an instruction, the scheduler can issue an instruction from the remaining three threads at a rate of one every three clock cycles, for example. Similarly, if two threads are inactive, the scheduler can issue an instruction from the two active threads at the rate of one every other clock cycle. A key advantage of this approach is the ability to run general applications, such as those not able to take full advantage of 4-way multithreading, at full speed. Other suitable approaches include multithreaded fixed-cycle scheduling.

[0071] Referring now to FIG. 3E, an exemplary multithreaded fixed-cycle scheduling is shown and indicated by the general reference character **300E**. Instruction Scheduler (IS) **302E** can receive instructions from four active threads: A, B, C, and D, as shown. In this programmable fixed-cycle scheduling, a fixed number of cycles can be provided to a given thread before switching to another thread. In the example illustrated, thread A issues 256 instructions, which may be the maximum allowed in the system, before any instructions are issued from thread B. Once thread B is started, it may issue 200 instructions before handing off the pipeline to thread C, and so on.

[0072] Referring now to FIG. 3F, an exemplary multi-threaded fixed-cycle with eager round-robin scheduling is shown and indicated by the general reference character 300F. Instruction Scheduler (IS) 302F can receive instructions from four active threads: A, B, C, and D, as shown. This approach may be used in order to maximize pipeline efficiency when a stall condition is encountered. For example, if thread A encounters a stall (e.g., a cache miss) before it has issued 256 instructions, the other threads may be used in a round-robin manner to “fill up” the potentially wasted cycles. In the example shown in FIG. 3F, a stall condition may occur while accessing the instructions for thread A after cycle 7, at which point the scheduler can switch to thread B for cycle 8. Similarly, another stall condition may occur while accessing the instructions for thread B after cycle 13, so the scheduler can then switch to thread C for cycle 14. In this example, no stalls occur during the accessing of instructions for thread C, so scheduling for thread C can continue though the programmed limit for the thread (e.g., 200), so that the last C thread instruction can be placed in the pipeline in cycle 214.

[0073] Referring now to FIG. 3G, a core with associated interface units according to an embodiment of the invention is shown and indicated by the general reference character 300G. Core 302G can include Instruction Fetch Unit (IFU) 304G, Instruction Cache Unit (ICU) 306G, Decoupling buffer 308G, Memory Management Unit (MMU) 310G, Instruction Execution Unit (IEU) 312G, and Load/Store Unit (LSU) 314. IFU 304G can interface with ICU 306G and IEU 312G can interface with LSU 314. ICU 306G can also interface with Switch Block (SWB)/Level 2 (L2) cache block 316G. LSU 314G, which can be a Level 1 (L1) data cache, can also interface with SWB/L2 316G. IEU 312G can interface with Message (MSG) Block 318G and, which can also interface with SWB 320G. Further, Register 322G for use in accordance with embodiments can include thread ID (TID), program counter (PC), and data fields.

[0074] According to embodiments of the invention, each MIPS architecture core may have a single physical pipeline, but may be configured to support multi-threading functions (i.e., four “virtual” cores). In a networking application, unlike a regular computational type of instruction scheme, threads are more likely to be waited on for memory accesses or other long latency operations. Thus, the scheduling approaches as discussed herein can be used to improve the overall efficiency of the system.

[0075] Referring now to FIG. 3H, an exemplary 10-stage (i.e., cycle) processor pipeline is shown and indicated by the general reference character 300H. In general operation, each instruction can proceed down the pipeline and may take 10-cycles or stages to execute. However, at any given point in time, there can be up to 10 different instructions populating each stage. Accordingly, the throughput for this example pipeline can be 1 instruction completing every cycle.

[0076] Viewing FIGS. 3G and 3H together, cycles 1-4 may represent the operation of IFU 304G, for example. In FIG. 3H, stage or cycle 1 (IPG stage) can include scheduling an instruction from among the different threads (Thread Scheduling 302H). Such thread scheduling can include round-robin, weighted round-robin (WRR), or eager round-robin, for example. Further, an Instruction Pointer (IP) may be generated in the IPG stage. An instruction fetch out of ICU 306G can occur in stages 2 (FET) and 3 (FE2) and can be initiated in Instruction Fetch Start 304H in stage 2. In stage 3, Branch Prediction 306H and/or Return Address Stack (RAS) (Jump

Register) 310H can be initiated and may complete in stage 4 (DEC). Also in stage 4, the fetched instruction can be returned (Instruction Return 308H). Next, instruction as well as other related information can be passed onto stage 5 and also put in Decoupling buffer 308G.

[0077] Stages 5-10 of the example pipeline operation of FIG. 3H can represent the operation of IEU 312G. In stage 5 (REG), the instruction may be decoded and any required register lookup (Register Lookup 314H) completed. Also in stage 5, hazard detection logic (LD-Use Hazard 316H) can determine whether a stall is needed. If a stall is needed, the hazard detection logic can send a signal to Decouple buffer 308G to replay the instruction (e.g., Decoupling/Replay 312H). However, if no such replay is signaled, the instruction may instead be taken out of Decoupling buffer 308G. Further, in some situations, such as where a hazard/dependency is due to a pending long-latency operation (e.g., a data-cache miss), the thread may not be replayed, but rather put to sleep. In stage 6 (EXE), the instruction may be “executed,” which may, for example, include an ALU/Shift and/or other operations (e.g., ALU/Shift/OP 318H). In stage 7 (MEM), a data memory operation can be initiated and an outcome of the branch can be resolved (Branch Resolution 320H). Further, the data memory lookup may extend to span stages 7, 8 (RTN), and 9 (RT2), and the load data can be returned (Load Return 322H) by stage 9 (RT2). In stage 10 (WRB), the instruction can be committed or retired and all associated registers can be finally updated for the particular instruction.

[0078] In general, the architecture is designed such that there are no stalls in the pipeline. This approach was taken for both ease of implementation as well as increased frequency of operation. However, there are some situations where a pipeline stall or stop is required. In such situations, Decoupling buffer 308G, which can be considered a functional part of IFU 304G, can allow for a restart or “replay” from a stop point instead of having to flush the entire pipeline and start the thread over to effect the stall. A signal can be provided by IFU 304G to Decoupling buffer 308G to indicate that a stall is needed, for example. In one embodiment, Decoupling buffer 308G can act as a queue for instructions whereby each instruction obtained by IFU 304G also goes to Decoupling buffer 308G. In such a queue, instructions may be scheduled out of order based on the particular thread scheduling, as discussed above. In the event of a signal to Decoupling buffer 308G that a stall is requested, those instructions after the “stop” point can be re-threaded. On the other hand, if no stall is requested, instructions can simply be taken out of the decoupling buffer and the pipeline continued. Accordingly, without a stall, Decoupling buffer 308G can behave essentially like a first-in first-out (FIFO) buffer. However, if one of several threads requests a stall, the others can proceed through the buffer and they may not be held up.

[0079] As another aspect of embodiments of the invention, a translation-lookaside buffer (TLB) can be managed as part of a memory management unit (MMU), such as MMU 310G of FIG. 3G. This can include separate, as well as common, TLB allocation across multiple threads. The 128-entry TLB can include a 64-entry joint main TLB and two 32-entry microTLBs, one each for the instruction and the data side. When a translation cannot be satisfied by accessing the relevant microTLB, a request may be sent to the main TLB. An interrupt or trap may occur if the main TLB also does not contain the desired entry.

[0080] In order to maintain compliance with the MIPS architecture, the main TLB can support paired entries (e.g., a pair of consecutive virtual pages mapped to different physical pages), variable page sizes (e.g., 4K to 256M), and software management via TLB read/write instructions. To support multiple threads, entries in the microTLB and in the main TLB may be tagged with the thread ID (TID) of the thread that installed them. Further, the main TLB can be operated in at least two modes. In a “partitioned” mode, each active thread may be allocated an exclusive subset or portion of the main TLB to install entries and, during translation, each thread only sees entries belonging to itself. In “global” mode, any thread may allocate entries in any portion of the main TLB and all entries may be visible to all threads. A “de-map” mechanism can be used during main TLB writes to ensure that overlapping translations are not introduced by different threads.

[0081] Entries in each microTLB can be allocated using a not-recently-used (NRU) algorithm, as one example. Regardless of the mode, threads may allocate entries in any part of the microTLB. However, translation in the microTLB may be affected by mode. In global mode, all microTLB entries may be visible to all threads, but in partitioned mode, each thread may only see its own entries. Further, because the main TLB can support a maximum of one translation per cycle, an arbitration mechanism may be used to ensure that microTLB “miss” requests from all threads are serviced fairly.

[0082] In a standard MIPS architecture, unmapped regions of the address space follow the convention that the physical address equals the virtual address. However, according to embodiments of the invention, this restriction is lifted and unmapped regions can undergo virtual-to-physical mappings through the microTLB/mainTLB hierarchy while operating in a “virtual-MIPS” mode. This approach allows a user to isolate unmapped regions of different threads from one another. As a byproduct of this approach, however, the normal MIPS convention that mainTLB entries containing an unmapped address in their virtual page number (VPN2) field can be considered invalid is violated. In one embodiment of the invention, this capability can be restored to the user whereby each entry in the mainTLB can include a special “master valid” bit that may only be visible to the user in the virtual MIPS-mode. For example, an invalid entry can be denoted by a master valid bit value of “0” and a valid entry can be denoted by a master valid bit value of “1.”

[0083] As another aspect of the invention, the system can support out-of-order load/store scheduling in an in-order pipeline. As an example implementation, there can be a user-programmable relaxed memory ordering model so as to maximize overall performance. In one embodiment, the ordering can be changed by user programming to go from a strongly ordered model to a weakly ordered model. The system can support four types: (i) Load-Load Re-ordering; (ii) Load-Store Re-ordering; (ii) Store-Store Re-ordering; and (iv) Store-Load Re-ordering. Each type of ordering can be independently relaxed by way of a bit vector in a register. If each type is set to the relaxed state, a weakly ordered model can be achieved.

[0084] Referring now to FIG. 3I, a core interrupt flow operation within a processor according to an embodiment of the invention is shown and indicated by the general reference character 300I. Programmable Interrupt Controller (PIC), as will be discussed in more detail below with reference to FIG. 3J, may provide an interrupt including Interrupt Counter and

MSG Block to Accumulates 302I. Accordingly, operation 300I can occur within any of the processors or cores of the overall system. Functional block Schedules Thread 304I can receive control interface from block 302I. Extensions to the MIPS architecture can be realized by shadow mappings that can include Cause 306I to EIRR 308I as well as Status 310I to EIMR 312I. The MIPS architecture generally only provides 2 bits for software interrupts and 6 bits for hardware interrupts for each of designated status and cause registers. This MIPS instruction architecture compatibility can be retained while providing extensions, according to embodiments of the invention.

[0085] As shown in more detail in FIG. 3I, a shadow mapping for Cause 306I to EIRR 308I for an interrupt pending can include bits 8-15 of the Cause 306I register mapping to bits 0-7 of EIRR 308I. Also, a software interrupt can remain within a core, as opposed to going through the PIC, and can be enacted by writing to bits 8 and/or 9 of Cause 306I. The remaining 6 bits of Cause 306I can be used for hardware interrupts. Similarly, a shadow mapping for Status 310I to EIMR 312I for a mask can include bits 8-15 of the Status 310I register mapping to bits 0-7 of EIMR 312I. Further, a software interrupt can be enacted by writing to bits 8 and/or 9 of Status 310I while the remaining 6 bits can be used for hardware interrupts. In this fashion, the register extensions according to embodiments of the invention can provide much more flexibility in dealing with interrupts. In one embodiment, interrupts can also be conveyed via the non-shadowed bits 8-63 of EIRR 308I and/or bits 8-63 of EIMR 312I.

[0086] Referring now to FIG. 3J, a PIC operation according to an embodiment of the invention is shown and indicated by the general reference character 300J. For example, flow 300J may be included in an implementation of box 226 of FIG. 2A. In FIG. 3J, Sync 302J can receive an interrupt indication and provide a control input to Pending 304J control block. Pending 304J, which can effectively act as an interrupt gateway, can also receive system timer and watch dog timer indications. Schedule Interrupt 306J can receive an input from Pending 304J. Interrupt Redirection Table (IRT) 308J can receive an input from Schedule Interrupt 306J.

[0087] Each interrupt and/or entry of IRT 308J can include associated attributes (e.g., Attribute 314J) for the interrupt, as shown. Attribute 314J can include CPU Mask 316-1J, Interrupt Vector 316-2J, as well as fields 316-3J and 316-4J, for examples. Interrupt Vector 316-2J can be a 6-bit field that designates a priority for the interrupt. In one embodiment, a lower number in Interrupt Vector 316-2J can indicate a higher priority for the associated interrupt via a mapping to EIRR 308I, as discussed above with reference to FIG. 3I. In FIG. 3J, Schedule across CPU & Threads 310J can receive an input from block 308J, such as information from Attribute 314J. In particular, CPU Mask 316-1J may be used to indicate to which of the CPUs or cores the interrupt is to be delivered. Delivery 312J block can receive an input from block 310J.

[0088] In addition to the PIC, each of 32 threads, for example, may contain a 64-bit interrupt vector. The PIC may receive interrupts or requests from agents and then deliver them to the appropriate thread. As one example implementation, this control may be software programmable. Accordingly, software control may elect to redirect all external type interrupts to one or more threads by programming the appropriate PIC control registers. Similarly, the PIC may receive an interrupt event or indication from the PCI-X interface (e.g., PCI-X 234 of FIG. 2A), which may in turn be redirected to a

specific thread of a processor core. Further, an interrupt redirection table (e.g., IRT **308J** of FIG. **3J**) may describe the identification of events (e.g., an interrupt indication) received by the PIC as well as information related to its direction to one or more “agents.” The events can be redirected to a specific core by using a core mask, which can be set by software to specify the vector number that may be used to deliver the event to a designated recipient. An advantage of this approach is that it allows the software to identify the source of the interrupt without polling.

[0089] In the case where multiple recipients are programmed for a given event or interrupt, the PIC scheduler can be programmed to use a global “round-robin” scheme or a per-interrupt-based local round-robin scheme for event delivery. For example, if threads **5**, **14**, and **27** are programmed to receive external interrupts, the PIC scheduler may deliver the first external interrupt to thread **5**, the next one to thread **14**, the next one to thread **27**, then return to thread **5** for the next interrupt, and so on.

[0090] In addition, the PIC also may allow any thread to interrupt any other thread (i.e., an inter-thread interrupt). This can be supported by performing a store (i.e., a write operation) to the PIC address space. The value that can be used for such a write operation can specify the interrupt vector and the target thread to be used by the PIC for the inter-thread interrupt. Software control can then use standard conventions to identify the inter-thread interrupts. As one example, a vector range may be reserved for this purpose.

[0091] As discussed above with reference to FIGS. **3G** and **3H**, each core can include a pipeline decoupling buffer (e.g., Decoupling **308G** of FIG. **3G**). In one aspect of embodiments of the invention, resource usage in an in-order pipeline with multiple threads can be maximized. Accordingly, the decoupling buffer is “thread aware” in that threads not requesting a stall can be allowed to flow through without stopping. In this fashion, the pipeline decoupling buffer can re-order previously scheduled threads. As discussed above, the thread scheduling can only occur at the beginning of a pipeline. Of course, re-ordering of instructions within a given thread is not normally performed by the decoupling buffer, but rather independent threads can incur no penalty because they can be allowed to effectively bypass the decoupling buffer while a stalled thread is held-up.

[0092] In one embodiment of the invention, a 3-cycle cache can be used in the core implementation. Such a 3-cycle cache can be an “off-the-shelf” cell library cache, as opposed to a specially-designed cache, in order to reduce system costs. As a result, there may be a gap of three cycles between the load and the use of a piece of data and/or an instruction. The decoupling buffer can effectively operate in and take advantage of this 3-cycle delay. For example, if there was only a single thread, a 3-cycle latency would be incurred. However, where four threads are accommodated, intervening slots can be taken up by the other threads. Further, branch prediction can also be supported. For branches correctly predicted, but not taken, there is no penalty. For branches correctly predicted and taken, there is a one-cycle “bubble” or penalty. For a missed prediction, there is a 5-cycle bubble, but such a penalty can be vastly reduced where four threads are operating because the bubbles can simply be taken up by the other threads. For example, instead of a 5-cycle bubble, each of the four threads can take up one so that only a single bubble penalty effectively remains.

[0093] As discussed above with reference to FIGS. **3D**, **3E**, and **3F**, instruction scheduling schemes according to embodiments of the invention can include eager round-robin scheduling (ERRS), fixed number of cycles per thread, and multi-threaded fixed-cycle with ERRS. Further, the particular mechanism for activating threads in the presence of conflicts can include the use of a scoreboard mechanism, which can track long latency operations, such as memory access, multiply, and/or divide operations.

[0094] Referring now to FIG. **3K**, a return address stack (RAS) operation for multiple thread allocation is shown and indicated by the general reference character **300K**. This operation can be implemented in IFU **304G** of FIG. **3G** and as also indicated in operation **310H** of FIG. **3H**, for example. Among the instructions supported in embodiments of the invention are: (i) a branch instruction where a prediction is whether it is taken or not taken and the target is known; (ii) a jump instruction where it is always taken and the target is known; and (iii) a jump register where it is always taken and the target is retrieved from a register and/or a stack having unknown contents.

[0095] In the example operation of FIG. **3K**, a Jump-And-Link (JAL) instruction can be encountered (**302K**) to initiate the operation. In response to the JAL, the program counter (PC) can be placed on the return address stack (RAS) (**304K**). An example RAS is shown as Stack **312K** and, in one embodiment, Stack **312K** is a first-in last-out (FILO) type of stack to accommodate nested subroutine calls. Substantially in parallel with placing the PC on Stack **312K**, a subroutine call can be made (**306K**). Various operations associated with the subroutine instructions can then occur (**308K**). Once the subroutine flow is complete, the return address can be retrieved from Stack **312K** (**310K**) and the main program can continue (**316K**) following any branch delay (**314K**).

[0096] For multiple thread operation, Stack **312K** can be partitioned so that entries are dynamically configured across a number of threads. The partitions can change to accommodate the number of active threads. Accordingly, if only one thread is in use, the entire set of entries allocated for Stack **312K** can be used for that thread. However, if multiple threads are active, the entries of Stack **312K** can be dynamically configured to accommodate the threads so as to utilize the available space of Stack **312K** efficiently.

[0097] In a conventional multiprocessor environment, interrupts are typically given to different CPUs for processing on a round-robin basis or by designation of a particular CPU for the handling of interrupts. However, in accordance with embodiments of the invention, PIC **226** of FIG. **2A**, with operation shown in more detail in FIG. **3J**, may have the ability to load balance and redirect interrupts across multiple CPUs/cores and threads in a multithreaded machine. As discussed above with reference to FIG. **3J**, IRT **308J** can include attributes for each interrupt, as shown in Attribute **314J**. CPU Mask **316-1J** can be used to facilitate load balancing by allowing for certain CPUs and/or threads to be masked out of the interrupt handling. In one embodiment, CPU Mask may be 32-bits wide to allow for any combination of 8 cores, each having 4 threads, to be masked. As an example, Core-2 **210c** and Core-7 **210h** of FIG. **2A** may be intended to be high availability processors, so CPU Mask **316-1J** of FIG. **3J** may have its corresponding bits set to “1” for each interrupt in IRT **308J** so as to disallow any interrupt processing on Core-2 or Core-7.

[0098] Further, for both CPUs/cores as well as threads, a round-robin scheme (e.g., by way of a pointer) can be employed among those cores and/or threads that are not masked for a particular interrupt. In this fashion, maximum programmable flexibility is allowed for interrupt load balancing. Accordingly, operation 300J of FIG. 3J allows for two levels of interrupt scheduling: (i) the scheduling of 306J, as discussed above; and (ii) the load balancing approach including CPU/core and thread masking.

[0099] As another aspect of embodiments of the invention, thread-to-thread interrupting is allowed whereby one thread can interrupt another thread. Such thread-to-thread interrupting may be used for synchronization of different threads, as is common for telecommunications applications. Also, such thread-to-thread interrupting may not go through any scheduling according to embodiments of the invention.

C. DATA SWITCH AND L2 CACHE

[0100] Returning now to FIG. 2A, the exemplary processor may further include a number of components that promote high performance, including: an 8-way set associative on-chip level-2 (L2) cache (2 MB); a cache coherent Hyper Transport interface (768 Gbps); hardware accelerated Quality-of-Service (QOS) and classification; security hardware acceleration—AES, DES/3DES, SHA-1, MD5, and RSA; packet ordering support; string processing support; TOE hardware (TCP Offload Engine); and numerous IO signals. In one aspect of an embodiment of the invention, data switch interconnect 216 may be coupled to each of the processor cores 210a-h by its respective data cache 212a-h. Also, the messaging network 222 may be coupled to each of the processor cores 210a-h by its respective instruction cache 214a-h. Further, in one aspect of an embodiment of the invention, the advanced telecommunications processor can also include an L2 cache 208 coupled to the data switch interconnect and configured to store information accessible to the processor cores 210a-h. In the exemplary embodiment, the L2 cache includes the same number of sections (sometimes referred to as banks) as the number of processor cores. This example is described with reference to FIG. 4A, but it is also possible to use more or fewer L2 cache sections.

[0101] As previously discussed, embodiments of the invention may include the maintenance of cache coherency using MOSI (Modified, Own, Shared, Invalid) protocol. The addition of the “Own” state enhances the “MSI” protocol by allowing the sharing of dirty cache lines across process cores. In particular, an example embodiment of the invention may present a fully coherent view of the memory to software that may be running on up to 32 hardware contexts of 8 processor cores as well as the I/O devices. The MOSI protocol may be used throughout the L1 and L2 cache (e.g., 212a-h and 208, respectively, of FIG. 2A) hierarchy. Further, all external references (e.g., those initiated by an I/O device) may snoop the L1 and L2 caches to ensure coherency and consistency of data. In one embodiment, as will be discussed in more detail below, a ring-based approach may be used to implement cache coherency in a multiprocessing system. In general, only one “node” may be the owner for a piece of data in order to maintain coherency.

[0102] According to one aspect of embodiments of the invention, an L2 cache (e.g., cache 208 of FIG. 2A) may be a 2 MB, 8-way set-associative unified (i.e., instruction and data) cache with a 32B line size. Further, up to 8 simultaneous references can be accepted by the L2 cache per cycle. The L2

arrays may run at about half the rate of the core clock, but the arrays can be pipelined to allow a request to be accepted by all banks every core clock with a latency of about 2 core clocks through the arrays. Also, the L2 cache design can be “non-inclusive” of the L1 caches so that the overall memory capacity can be effectively increased.

[0103] As to ECC protection for an L2 cache implementation, both cache data and cache tag arrays can be protected by SECDED (Single Error Correction Double Error Detection) error protecting codes. Accordingly, all single bit errors are corrected without software intervention. Also, when uncorrectable errors are detected, they can be passed to the software as code-error exceptions whenever the cache line is modified. In one embodiment, as will be discussed in more detail below, each L2 cache may act like any other “agent” on a ring of components.

[0104] According to another aspect of embodiments of the invention, “bridges” on a data movement ring may be used for optimal redirection of memory and I/O traffic. Super Memory I/O Bridge 206 and Memory Bridge 218 of FIG. 2A may be separate physical structures, but they may be conceptually the same. The bridges can be the main gatekeepers for main memory and I/O accesses, for example. Further, in one embodiment, the I/O can be memory-mapped.

[0105] Referring now to FIG. 4A, a data switch interconnect (DSI) ring arrangement according to an embodiment of the invention is shown and indicated by the general reference character 400A. Such a ring arrangement can be an implementation of DSI 216 along with Super Memory I/O Bridge 206 and Memory Bridge 218 of FIG. 2A. In FIG. 4A, Bridge 206 can allow an interface between memory & I/O and the rest of the ring. Ring elements 402a-j each correspond to one of the cores 210a-h and the memory bridges of FIG. 2A. Accordingly, element 402a interfaces to L2 cache L2a and Core-0 210a, and element 402b interfaces to L2b and Core 210b, and so on through 402h interfacing to L2h and Core 210h. Bridge 206 includes an element 402i on the ring and bridge 218 includes an element 402j on the ring.

[0106] As shown in FIG. 4A, four rings can make up the ring structure in an example embodiment: Request Ring (RQ), Data Ring (DT), Snoop Ring (SNP), and Response Ring (RSP). The communication on the rings is packet based communication. An exemplary RQ ring packet includes destination ID, transaction ID, address, request type (e.g., RD, RD_EX, WR, UPG), valid bit, cacheable indication, and a byte enable, for example. An exemplary DT ring packet includes destination ID, transaction ID, data, status (e.g., error indication), and a valid bit, for example. An exemplary SNP ring packet includes destination ID, valid bit, CPU snoop response (e.g., clean, shared, or dirty indication), L2 snoop response, bridge snoop response, retry (for each of CPU, bridge, and L2), AERR (e.g., illegal request, request parity), and a transaction ID, for example. An exemplary RSP ring packet includes all the fields of SNP, but may represent a “final” status, as opposed to the “in-progress” status of the RSP ring.

[0107] Referring now to FIG. 4B, a DSI ring component according to an embodiment of the invention is shown and indicated by the general reference character 400B. Ring component 402b-0 may correspond to one of the four rings RQ, DT, SNP, or RSP, in one embodiment. Similarly, ring components 402b-1, 402b-2, and 402b-3 may each correspond to

one of the four rings. As an example, a “node” can be formed by the summation of ring components **402b-0**, **402b-1**, **402b-2**, and **402b-3**.

[0108] Incoming data or “Ring In” can be received in flip-flop **404B**. An output of flip-flop **404B** can connect to flip-flops **406B** and **408B** as well as multiplexer **416B**. Outputs of flip-flops **406B** and **408B** can be used for local data use. Flip-flop **410B** can receive an input from the associated L2 cache while flip-flop **412B** can receive an input from the associated CPU. Outputs from flip-flops **410B** and **412B** can connect to multiplexer **414B**. An output of multiplexer **414B** can connect to multiplexer **416B** and an output of multiplexer **416B** can connect to outgoing data or “Ring Out.” Also, ring component **402b-0** can receive a valid bit signal.

[0109] Generally, higher priority data received on Ring In will be selected by multiplexer **416B** if the data is valid (e.g., Valid Bit=“1”). If not, the data can be selected from either the L2 or the CPU via multiplexer **414B**. Further, in this example, if data received on Ring In is intended for the local node, flip-flops **406B** and/or **408B** can pass the data onto the local core instead of allowing the data to pass all the way around the ring before receiving it again.

[0110] Referring now to FIG. 4C, a flow diagram of an example data retrieval in the DSI according to an embodiment of the invention is shown and indicated by the general reference character **400C**. The flow can begin in Start **452** and a request can be placed on the request ring (RQ) (**454**). Each CPU and L2 in the ring structure can check for the requested data (**456**). Also, the request can be received in each memory bridge attached to the ring (**458**). If any CPU or L2 has the requested data (**460**), the data can be put on the data ring (DT) by the node having the data (**462**). If no CPU or L2 has found the requested data (**460**), the data can be retrieved by one of the memory bridges (**464**). An acknowledgement can be placed on the snoop ring (SNP) and/or the response ring (RSP) by either the node that found the data or the memory bridge (**466**) and the flow can complete in End (**468**). In one embodiment, the acknowledgement by the memory bridge to the SNP and/or RSP ring may be implied.

[0111] In an alternative embodiment, the memory bridge would not have to wait for an indication that the data has not been found in any of the L2 caches in order to initiate the memory request. Rather, the memory request (e.g., to DRAM), may be speculatively issued. In this approach, if the data is found prior to the response from the DRAM, the later response can be discarded. The speculative DRAM accesses can help to mitigate the effects of the relatively long memory latencies.

D. MESSAGE PASSING NETWORK

[0112] Also in FIG. 2A, in one aspect of an embodiment of the invention, the advanced telecommunications processor can include Interface Switch Interconnect (ISI) **224** coupled to the messaging network **222** and a group of communication ports **240a-f**, and configured to pass information among the messaging network **222** and the communication ports **240a-f**.

[0113] Referring now to FIG. 5A, a fast messaging ring component or station according to an embodiment of the invention is shown and indicated by the general reference character **500A**. An associated ring structure may accommodate point-to-point messages as an extension of the MIPS architecture, for example. The “Ring In” signal can connect to both Insertion Queue **502A** and Receive Queue (RCVQ) **506A**. The insertion queue can also connect to multiplexer

504A, the output of which can be “Ring Out.” The insertion queue always gets priority so that the ring does not get backed-up. Associated registers for the CPU core are shown in dashed boxes **520A** and **522A**. Within box **520A**, buffers RCV Buffer **510A-0** through RCV Buffer **510A-N** can interface with RCVQ **506A**. A second input to multiplexer **504A** can connect to Transmit Queue (XMTQ) **508A**. Also within box **520A**, buffers XMT Buffer **512A-0** through XMT Buffer **512A-N** can interface with XMTQ **508A**. Status **514A** registers can also be found in box **520A**. Within dashed box **522A**, memory-mapped Configuration Registers **516A** and Credit Based Flow Control **518A** can be found.

[0114] Referring now to FIG. 5B, a message data structure for the system of FIG. 5A is shown and indicated by the general reference character **500B**. Identification fields may include Thread **502B**, Source **504B**, and Destination **508B**. Also, there can be a message size indicator Size **508B**. The identification fields and the message size indicator can form Sideboard **514B**. The message or data to be sent itself (e.g., MSG **512B**) can include several portions, such as **510B-0**, **510B-1**, **510B-2**, and **510B-3**. According to embodiments, the messages may be atomic so that the full message cannot be interrupted.

[0115] The credit-based flow control can provide a mechanism for managing message sending, for example. In one embodiment, the total number of credits assigned to all transmitters for a target/receiver cannot exceed the sum of the number of entries in its receive queue (e.g., RCVQ **506A** of FIG. 5A). For example, 256 may be the total number of credits in one embodiment because the size of the RCVQ of each target/receiver may be 256 entries. Generally, software may control the assignment of credits. At boot-up time, for example, each sender/xmitter or participating agent may be assigned some default number of credits. Software may then be free to allocate credits on a per-transmitter basis. For example, each sender/xmitter can have a programmable number of credits set by software for each of the other targets/receivers in the system. However, not all agents in the system may be required to participate as targets/receivers in the distribution of the transmit credits.

[0116] In one embodiment, Core-0 credits can be programmed for each one of Core-1, Core-2, . . . Core-7, RGMII_0, RGMII_1, XGMII/SPI-4.2_0, XGMII/SPI-4.2_1, POD0, POD1, . . . POD4, etc. The Table 1 below shows an example distribution of credits for Core-0 as a receiver:

TABLE 1

Transmit Agents	Allocated Credits (Total of 256)
Core-0	0
Core-1	32
Core-2	32
Core-3	32
Core-4	0
Core-5	32
Core-6	32
Core-7	32
POD0	32
RGMII_0	32
All Others	0

[0117] In this example, when Core-1 sends a message of size 2 (e.g., 2 64-bit data elements) to Core-0, the Core-1 credit in Core-0 can be decremented by 2 (e.g., from 32 to 30). When Core-0 receives a message, the message can go into the

RCVQ of Core-0. Once the message is removed from the RCVQ of Core-0, that message storage space may essentially be freed-up or made available. Core-0 can then send a signal to the sender (e.g., a free credit signal to Core-1) to indicate the amount of space (e.g., 2) additionally available. If Core-1 continues to send messages to Core-0 without corresponding free credit signals from Core-0, eventually the number of credits for Core-1 can go to zero and Core-1 may not be able to send any more messages to Core-0. Only when Core-0 responds with free credit signals could Core-1 send additional messages to Core-0, for example.

[0118] Referring now to FIG. 5C, a conceptual view of how various agents may be attached to the fast messaging network (FMN) according to an embodiment of the invention is shown and indicated by the general reference character 500C. The eight cores (Core-0 502C-0 through Core-7 502C-7) along with associated data caches (D-cache 504C-0 through 504C-7) and instruction caches (I-cache 506C-0 through 506C-7) can interface to the FMN. Further, Network I/O Interface Groups can also interface to the FMN. Associated with Port A, DMA 508C-A, Parser/Classifier 512C-A, and XGMII/SPI-4.2 Port A 514C-A can interface to the FMN through Packet Distribution Engine (PDE) 510C-A. Similarly, for Port B, DMA 508C-B, Parser/Classifier 512C-B, and XGMII/SPI-4.2 Port B 514C-B can interface to the FMN through PDE 510C-B. Also, DMA 516C, Parser/Classifier 520C, RGMII Port A 522C-A, RGMII Port B 522C-B, RGMII Port C 522C-C, RGMII Port D 522C-D can interface to the FMN through PDE 518C. Also, Security Acceleration Engine 524C including DMA 526C and DMA Engine 528C can interface to the FMN.

[0119] As an aspect of embodiments of the invention, all agents (e.g., cores/threads or networking interfaces, such as shown in FIG. 5C) on the FMN can send a message to any other agent on the FMN. This structure can allow for fast packet movement among the agents, but software can alter the use of the messaging system for any other appropriate purpose by so defining the syntax and semantics of the message container. In any event, each agent on the FMN includes a transmit queue (e.g., 508A) and a receive queue (e.g., 506A), as discussed above with reference to FIG. 5A. Accordingly, messages intended for a particular agent can be dropped into the associated receive queue. All messages originating from a particular agent can be entered into the associated transmit queue and subsequently pushed on the FMN for delivery to the intended recipient.

[0120] In another aspect of embodiments of the invention, all threads of the core (e.g., Core-0 502C-0 through Core-7 502C-7 or FIG. 5C) can share the queue resources. In order to ensure fairness in sending out messages, a “round-robin” scheme can be implemented for accepting messages into the transmit queue. This can guarantee that all threads have the ability to send out messages even when one of them is issuing messages at a faster rate. Accordingly, it is possible that a given transmit queue may be full at the time a message is issued. In such a case, all threads can be allowed to queue up one message each inside the core until the transmit queue has room to accept more messages. As shown in FIG. 5C, the networking interfaces use the PDE to distribute incoming packets to the designated threads. Further, outgoing packets for the networking interfaces can be routed through packet ordering software.

[0121] Referring now to FIG. 5D, network traffic in a conventional processing system is shown and indicated by the

general reference character 500D. The Packet Input can be received by Packet Distribution 502D and sent for Packet Processing (504D-0 through 504D-3). Packet Sorting/Ordering 506D can receive the outputs from Packet Processing and can provide Packet Output. While such packet-level parallel-processing architectures are inherently suited for networking applications, but an effective architecture must provide efficient support for incoming packet distribution and outgoing packet sorting/ordering to maximize the advantages of parallel packet processing. As shown in FIG. 5D, every packet must go through a single distribution (e.g., 502D) and a single sorting/ordering (e.g., 506D). Both of these operations have a serializing effect on the packet stream so that the overall performance of the system is determined by the slower of these two functions.

[0122] Referring now to FIG. 5E, a packet flow according to an embodiment of the invention is shown and indicated by the general reference character 500E. This approach provides an extensive (i.e., scalable) high-performance architecture enabling flow of packets through the system. Networking Input 502E can include and of RGMII, XGMII, and/or SPI-4.2 interface configured ports. After the packets are received, they can be distributed via Packet Distribution Engine (PDE) 504E using the Fast Messaging Network (FMN) to one of the threads for Packet Processing 506E: Thread 0, 1, 2, and so on through Thread 31, for example. The selected thread can perform one or more functions as programmed by the packet header or the payload and then the packet on to Packet Ordering Software 508E. As an alternative embodiment, a Packet Ordering Device (POD), as shown in box 236 of FIG. 2A, for example, may be used in place of 508E of FIG. 5E. In either implementation, this function sets up the packet ordering and then passes it on to the outgoing network (e.g., Networking Output 510E) via the FMN. Similar to the networking input, the outgoing port can be any one of the configured RGMII, XGMII, or SPI-4.2 interfaces, for example.

E. INTERFACE SWITCH

[0123] In one aspect of embodiments of the invention, the FMN can interface to each CPU/core, as shown in FIG. 2A. Such FMN-to-core interfacing may include push/pop instructions, waiting for a message instruction, and interrupting on a message arrival. In the conventional MIPS architecture, a co-processor or “COP2” space is allocated. However, according to embodiments of the invention, the space designated for COP2 is instead reserved for messaging use via the FMN. In one embodiment, software executable instructions may include message send (MsgSnd), message load (MsgLd), message-to-COP2 (MTC2), message-from-COP2 (MFC2), and message wait (MsgWait). The MsgSnd and MsgLd instructions can include target information as well as message size indications. The MTC2 and MFC2 instructions can include data transfers from/to local configuration registers, such as Status 514A and registers 522A of FIG. 5A. The MsgWait instruction can include the operation of essentially entering a “sleep” state until a message is available (e.g., interrupting on message arrival).

[0124] As another aspect of embodiments of the invention, fast messaging (FMN) ring components can be organized into “buckets.” For, example, RCVQ 506A and XMTQ 508A of FIG. 5A may each be partitioned across multiple buckets in similar fashion to the thread concept, as discussed above.

[0125] In one aspect of embodiments of the invention, a Packet Distribution Engine (PDE) can include each of the

XGMII/SPI-4.2 interfaces and four RGMII interfaces to enable efficient and load-balanced distribution of incoming packets to the processing threads. Hardware accelerated packet distribution is important for high throughput networking applications. Without the PDE, packet distribution may be handled by software, for example. However, for 64B packets, only about 20 ns is available for execution of this function on an XGMII type interface. Further, queue pointer management would have to be handled due to the single-producer multiple-consumer situation. Such a software-only solution is simply not able to keep up with the required packet delivery rate without impacting the performance of the overall system.

[0126] According to an embodiment of the invention, the PDE can utilize the Fast Messaging Network (FMN) to quickly distribute packets to the threads designated by software as processing threads. In one embodiment, the PDE can implement a weighted round-robin scheme for distributing packets among the intended recipients. In one implementation, a packet is not actually moved, but rather gets written to memory as the networking interface receives it. The PDE can insert a "Packet Descriptor" in the message and then send it to one of the recipients, as designated by software. This can also mean that not all threads must participate in receiving packets from any given interface.

[0127] Referring now to FIG. 6A, a PDE distributing packets evenly over four threads according to an embodiment of the invention is shown and indicated by the general reference character 600A. In this example, software may choose threads 4 through 7 for possible reception of packets. The PDE can then select one of these threads in sequence to distribute each packet, for example. In FIG. 6A, Networking Input can be received by Packet Distribution Engine (PDE) 602A, which can select one of Thread 4, 5, 6, or 7 for packet distribution. In this particular example, Thread 4 can receive packet 1 at time t_1 and packet 5 at time t_5 , Thread 5 can receive packet 2 at time t_2 and packet 6 at time t_6 , Thread 6 can receive packet 3 at time t_3 and packet 7 at time t_7 , and Thread 7 can receive packet 4 at time t_4 and packet 8 at time t_8 .

[0128] Referring now to FIG. 6B, a PDE distributing packets using a round-robin scheme according to an embodiment of the invention is shown and indicated by the general reference character 600B. As describe above with reference to the FMN, software can program the number of credits allowed for all receivers from every transmitter. Since the PDE is essentially a transmitter, it can also use the credit information to distribute the packets in a "round-robin" fashion. In FIG. 6B, PDE 602B can receive Networking Input and provide packets to the designated threads (e.g., Thread 0 through Thread 3), as shown. In this example, Thread 2 (e.g., a receiver) may be processing packets more slowly than the other threads. PDE 602B can detect the slow pace of credit availability from this receiver and adjust by guiding packets to the more efficiently processing threads. In particular, Thread 2 has the least number of credits available within the PDE at cycle t_{11} . Although the next logical receiver of packet 11 at cycle t_{11} may have been Thread 2, the PDE can identify a processing delay in that thread and accordingly select Thread 3 as the optimal target for distribution of packet 11. In this particular example, Thread 2 can continue to exhibit processing delays relative to the other threads, so the PDE can avoid distribution to this thread. Also, in the event that none of the receivers has room to accept a new packet, the PDE can extend the packet queue to memory.

[0129] Because most networking applications are not very tolerant of the random arrival order of packets, it is desirable to deliver packets in order. In addition, it can be difficult to combine features of parallel processing and packet ordering in a system. One approach is to leave the ordering task to software, but it then becomes difficult to maintain line rate. Another option is to send all packets in a single flow to the same processing thread so that the ordering is essentially automatic. However, this approach would require flow identification (i.e., classification) prior to packet distribution and this reduces system performance. Another drawback is the throughput of the largest flow is determined by the performance of the single thread. This prevents single large flows from sustaining their throughput as they traverse the system.

[0130] According to an embodiment of the invention, an advanced hardware-accelerated structure called a Packet Ordering Device (POD) can be used. An objective of the POD is to provide an unrestricted use of parallel processing threads by re-ordering the packets before they are sent to the networking output interface. Referring now to FIG. 6C, a POD placement during packet lifecycle according to an embodiment of the invention is shown and indicated by the general reference character 600C. This figure essentially illustrates a logical placement of the POD during the life cycle of the packets through the processor. In this particular example, PDE 602C can send packets to the threads, as shown. Thread 0 can receive packet 1 at time t_1 , packet 5 at time t_5 , and so on through cycle t_{n-3} . Thread 1 can receive packet 2 at time t_2 , packet 6 at time t_6 , and so on through cycle t_{n-2} . Thread 2 can receive packet 3 at time t_3 , packet 7 at time t_7 , and so on through time t_{n-1} . Finally, Thread 3 can receive packet 4 at time t_4 , packet 8 at time t_8 , and so on through time t_n .

[0131] Packet Ordering Device (POD) 604C can be considered a packet sorter in receiving the packets from the different threads and then sending to Networking Output. All packets received by a given networking interface can be assigned a sequence number. This sequence number can then be forwarded to the working thread along with the rest of the packet information by the PDE. Once a thread has completed processing the packet, it can forward the packet descriptor along with the original sequence number to the POD. The POD can release these packets to the outbound interface in an order strictly determined by the original sequence numbers assigned by the receiving interface, for example.

[0132] In most applications, the POD will receive packets in a random order because the packets are typically processed by threads in a random order. The POD can establish a queue based on the sequence number assigned by the receiving interface and continue sorting packets as received. The POD can issue packets to a given outbound interface in the order assigned by the receiving interface. Referring now to FIG. 6D, a POD outbound distribution according to an embodiment of the invention is shown and indicated by the general reference character 600D. As can be seen in Packet Ordering Device (POD) 602D, packets 2 and 4 can be initially sent to the POD by executing threads. After several cycles, a thread can complete work on packet 3 and place it in the POD. The packets may not yet be ordered because packet 1 is not yet in place. Finally, packet 1 is completed in cycle t_7 and placed in the POD accordingly. Packets can now be ordered and the POD can begin issuing packets in the order: 1, 2, 3, 4. If packet 5 is received next, it is issued in the output following packet 4. As the remaining packets are received, each can be stored in the queue (e.g., a 512-deep structure) until the next

higher number packet is received. At such time, the packet can be added to the outbound flow (e.g., Networking Output).

[0133] It is possible that the oldest packet may never arrive in the POD, thus creating a transient head-of-line blocking situation. If not handled properly, this error condition would cause the system to deadlock. However, according to an aspect of the embodiment, the POD is equipped with a time-out mechanism designed to drop a non-arriving packet at the head of the list once a time-out counter has expired. It is also possible that packets are input to the POD at a rate which fills the queue capacity (e.g., 512 positions) before the time-out counter has expired. According to an aspect of the embodiment, when the POD reaches queue capacity, the packet at the head of the list can be dropped and a new packet can be accepted. This action may also remove any head-of-line blocking situation as well. Also, software may be aware that a certain sequence number will not be entered into the POD due to a bad packet, a control packet, or some other suitable reason. In such a case, software control may insert a “dummy” descriptor in the POD to eliminate the transient head-of-line blocking condition before allowing the POD to automatically react.

[0134] According to embodiments of the invention, five programmable PODs may be available (e.g., on chip) and can be viewed as generic “sorting” structures. In one example configuration, software control (i.e., via a user) can assign four of the PODs to the four networking interfaces while retaining one POD for generic sorting purposes. Further, the PODs can simply be bypassed if so desired for applications where software-only control suffices.

F. MEMORY INTERFACE AND ACCESS

[0135] In one aspect of embodiments of the invention, the advanced telecommunications processor can further include memory bridge **218** coupled to the data switch interconnect and at least one communication port (e.g., box **220**), and configured to communicate with the data switch interconnect and the communication port.

[0136] In one aspect of the invention, the advanced telecommunications processor can further include super memory bridge **206** coupled to the data switch interconnect (DSI), the interface switch interconnect and at least one communication port (e.g., box **202**, box **204**), and configured to communicate with the data switch interconnect, the interface switch interconnect and the communication port.

[0137] In another aspect of embodiments of the invention, memory ordering can be implemented on a ring-based data movement network, as discussed above with reference to FIGS. **4A**, **4B**, and **4C**.

G. CONCLUSION

[0138] Advantages of the invention include the ability to provide high bandwidth communications between computer systems and memory in an efficient and cost-effective manner.

[0139] Having disclosed exemplary embodiments and the best mode, modifications and variations may be made to the disclosed embodiments while remaining within the subject and spirit of the invention as defined by the following claims.

1-30. (canceled)

31. A processor, comprising:

a packet processing system to execute a plurality of threads to process packets;

a packet distribution engine to: receive a plurality of packets in a first order; assign packet descriptors to respective packets of said plurality of packets, the packet descriptors comprising respective sequence numbers; provide the respective packets to memory; and distribute the packet descriptors to threads of the packet processing system without moving the respective packets from memory; and

a packet ordering device to receive the packet descriptors from the packet processing system in a second order different from the first order and to establish an order for distributing the respective packets based, at least in part, upon the respective sequence numbers of the packet descriptors.

32. The processor of claim **31**, further comprising an interface coupled to the packet distribution engine and selected from the group consisting of an XGMII interface, an SPI-4.2 interface and an RGMII interface.

33. The processor of claim **31**, wherein the packet distribution engine is configured to receive the plurality of packets in the first order from a networking input.

34. The processor of claim **33**, further comprising at least two interfaces coupled to the packet distribution engine and selected from the group consisting of an RGMII interface, an XGMII interface, and an SPI-4.2 interface.

35. The processor of claim **31**, wherein the packet distribution engine is separate from said packet ordering device.

36. The processor of claim **31**, wherein the packet processing system comprises a plurality of processor cores, a respective processor core for executing multiple threads.

37. The processor of claim **31**, wherein the packet processing system comprises a plurality of processor cores, each processor core for supporting a plurality of operating systems.

38. The processor of claim **31**, wherein the order for distributing the respective packets is the same as the first order.

39. The processor of claim **31**, wherein the packet ordering device is configured to drop a packet descriptor after a time-out counter expires such that the order established by the packet ordering device for distributing the respective packets does not include the packet associated with the dropped packet descriptor, and the next packet descriptor in the first order replaces the dropped packet descriptor.

40. The processor of claim **31**, wherein the packet ordering device is configured to distribute the respective packets to a networking output in the order established by the packet ordering device.

41. A method of controlling a flow of packets, the method comprising:

receiving a plurality of packets in a first order;

storing respective packets of said plurality of packets in memory;

assigning, by a packet distribution engine, packet descriptors including respective sequence numbers to the respective packets of said plurality of packets;

distributing, by said packet distribution engine, said packet descriptors to threads of a packet processing system, without moving the respective packets from the memory;

receiving, at a packet ordering device, the packet descriptors from the packet processing system in a second order different from the first order; and

establishing, by said packet ordering device, an output order for distributing the respective packets based, at least in part, upon the respective sequence numbers of the packet descriptors.

42. The method of claim **41**, wherein said output order is the same as said first order.

43. The method of claim **41**, wherein said output order is the same as said first order except that a non-arriving packet descriptor is dropped after a time-out counter has expired.

44. The method of claim **41**, further comprising:
dropping a packet descriptor after a time-out counter expires such that the output order does not include the packet associated with the dropped packet descriptor, and the next packet descriptor in the first order replaces the dropped packet descriptor.

45. The method of claim **41**, further comprising distributing the respective packets in the output order via an interface selected from the group consisting of an XGMII interface, an SPI-4.2 interface and an RGMII interface.

46. The method of claim **41**, wherein the packet processing system comprises a plurality of processor cores, each processor core executing multiple threads.

47. The method of claim **41**, wherein the packet processing system comprises a plurality of processor cores, and each processor core supports a plurality of operating systems.

48. The method of claim **41**, wherein the packet ordering device distributes the respective packets to a networking output in the output order.

* * * * *