

**MECHANISM FOR MANAGING RESOURCE LOCKING
IN A MULTI-THREADED ENVIRONMENT**

Inventor(s): Mario D. Nemirovsky, Jeffrey T. Huynh

Related Applications

[0001] This application is a continuation-in-part of U.S. Application Serial No. 10/254,377, filed September 24, 2002, which claims the benefit of U.S. Provisional Application Serial No. 60/325,638, filed September 28, 2001, U.S. Provisional Application Serial No. 60/341,689, filed December 17, 2001, and U.S. Provisional Application Serial No. 60/388,278, filed June 13, 2002. The contents of all of these applications are incorporated in their entirety herein by this reference.

Background

[0002] In a multi-threaded environment, resources may be shared by multiple threads. For example, a section of memory that stores the value of a global variable may be accessed and updated by many different threads. Whenever a resource such as a memory section is shared, it is important to ensure that updates to that resource are performed atomically. If they are not, then data consistency could be compromised. For this reason, multi-threaded environments usually implement some type of access management methodology to ensure that updates to shared resources are carried out atomically.

[0003] One way to ensure atomicity is to implement locking. With locking, whenever a thread wishes to access a shared resource, it makes a request to obtain a lock on the resource. If the resource is not currently locked by another thread, a lock is granted to the thread. Thereafter, the thread can read and update the contents of the resource with full assurance that no other thread will update the contents of the resource

while the thread has the lock. When the thread is through updating the resource, it releases the lock and another thread is allowed to obtain a lock on the resource. Under the locking approach, a thread cannot access and update a resource until it has a lock on the resource. Thus, if a thread requests a lock on a resource and that resource is already locked by another thread, then the thread has to make another lock request. The thread continues making lock requests until it finally obtains a lock on the resource. In some instances, when a large number of threads are trying to access the same resource at the same time, many threads could end up making many lock requests before they obtain a lock on the resource. While the threads are making these repeated lock requests, they are unnecessarily consuming power and processing resources. Thus, in an environment in which there are many concurrently executing threads, locking can lead to waste and inefficiency.

[0004] Another approach that has been used to ensure atomicity is the load link/store conditional approach. Under this approach, exclusive access to a resource is checked not at the time a resource is accessed but rather on the back end when the contents of the resource are to be updated. Under the load link/store conditional approach, when a thread wishes to access a resource, it obtains a reservation. With a reservation, the thread is allowed to access and read the contents of the resource. If the thread updates the contents and then wishes to write the updated contents back into the resource, then the thread has to check to see if its reservation is still valid. If no other thread has updated the contents of the resource after the reservation was obtained, then the reservation is still valid, in which case, the thread is allowed to update the resource. However, if any other thread updated the resource after the reservation was obtained, then the reservation is no longer

valid. In this case, the thread has to start the process all over again, namely, it has to obtain another reservation, access the resource to read the current contents, update the contents, and then try to write the updated contents to the resource again. In a scenario where many threads are trying to update the same resource at the same time, many threads could end up repeating this process many times. While the threads are repeating this process, they are unnecessarily consuming power and processing resources. Thus, like the locking approach, the load link/store condition approach can lead to significant waste and inefficiency in an environment in which there are many concurrently executing threads.

Summary

[0005] In accordance with one embodiment of the present invention, there is provided an improved mechanism for managing the locking of shared resources. This mechanism (referred to hereinafter as the lock manager) enables atomicity to be ensured in an environment in which there is a large number of concurrently executing threads (i.e. a massively multi-threaded environment) without suffering the waste and inefficiency of the prior approaches.

[0006] In one embodiment, the lock manager manages resource locking as follows. Initially, the lock manager receives, from a particular stream executing a particular thread, a request to obtain a lock on a resource. As used herein, the term stream refers broadly to any set of components that cooperate to execute a thread. In one embodiment, the lock manager resides on a processor that comprises many streams. In response to this lock request, the lock manager determines whether the resource is currently locked by another

stream. If the resource is currently locked by another stream, the lock manager does not simply return an indication that a lock is currently unavailable, as was done in the prior locking approach. Instead, the lock manager adds the particular stream to a wait list of streams waiting to obtain a lock on the resource. By doing so, the lock manager in effect gives the particular stream a reservation to obtain a lock on the resource in the future. Thus, the particular stream need not submit another lock request. In one embodiment, in addition to adding the particular stream to the wait list, the lock manager also causes the particular stream to halt execution of the particular thread. That way, the particular stream does not consume any processing resources while it is waiting for a lock on the resource.

[0007] The lock manager allows each stream on the wait list, in turn, to obtain a lock on the resource, access and optionally update the contents of the resource, and release the lock on the resource. At some point, based upon the wait list, the lock manager determines that it is the particular stream's turn to obtain a lock on the resource. At that point, the lock manager grants the particular stream a lock on the resource. In one embodiment, in addition to granting the lock, the lock manager also causes the particular stream to resume execution of the particular thread. Thereafter, the particular stream can access the resource, perform whatever processing and updates it wishes, and release the lock when it is done. In this manner, the lock manager enables the particular stream to reserve and to obtain a lock on the resource. By implementing locking in this way, streams do not need to repeatedly request a lock on a resource. Instead, they submit only one lock request, and when it is their turn to obtain a lock, they are granted that lock. By

eliminating the need to repeatedly request locks, the lock manager makes it possible to implement resource locking efficiently in a massively multi-threaded environment.

Brief Description of the Drawings

- [0008] Fig. 1 is an architectural overview for a packet processing engine in an embodiment of the present invention.
- [0009] Fig. 2 is a memory map for the packet processing engine in an embodiment of the present invention.
- [0010] Fig. 3 illustrates detail of the address space for the packet processing engine in an embodiment of the present invention.
- [0011] Figs. 4a through Fig. 4d comprise a list of configuration registers for a packet processing engine according to an embodiment of the present invention.
- [0012] Fig. 5 illustrates hashing function hardware for the packet processing engine in an embodiment of the present invention.
- [0013] Fig. 6 is a table that lists performance events for the packet processing engine in an embodiment of the present invention.
- [0014] Fig. 7 lists egress channel determination for the packet processing engine in an embodiment of the present invention.
- [0015] Fig. 8 lists egress port determination for the packet processing engine in an embodiment of the present invention.
- [0016] Fig. 9 indicates allowed degree of interleaving for the packet processing engine in an embodiment of the present invention.

[0017] Fig. 10 is an illustration of Global block architecture in an embodiment of the invention.

[0018] Fig. 11 is an expanded view showing internal components of the Global block.

[0019] Fig. 12 is an illustration of a Routing block in an embodiment of the invention.

[0020] Fig. 13 is a table indicating migration protocol between tribes.

[0021] Fig. 14 is a block diagram of the Network Unit for an embodiment of the invention.

[0022] Fig. 15 is a diagram of a Port Interface block in the Network Unit in an embodiment.

[0023] Fig. 16 is a diagram of a Packet Loader Block in the Network Unit in an embodiment of the invention.

[0024] Fig. 17 is a diagram of a Packet Buffer Control Block in an embodiment.

[0025] Fig. 18 is a diagram of a Packet Buffer Memory Block in an embodiment.

[0026] Fig. 19 is a table illustrating the interface between a Tribe and the Interconnect Block.

[0027] Fig. 20 is a table illustrating the interface between the Network Block and the Interconnect Block.

[0028] Fig. 21 is a table illustrating the interface between the Global Block and the Interconnect Block.

[0029] Fig. 22 is a diagram indicating migration protocol timing in the Interconnect Block.

- [0030] Fig. 23 is a table illustrating the interface between a Tribe and a Memory Interface block in an embodiment of the invention.
- [0031] Fig. 24 is a table illustrating the interface between the Global Block and a Memory Interface block in an embodiment of the invention.
- [0032] Fig. 25 is a table illustrating the interface between a Memory Controller and a Memory Interface block in an embodiment of the invention.
- [0033] Fig. 26 shows tribe to memory interface timing in an embodiment of the invention.
- [0034] Fig. 27 shows tribe memory interface to controller timing.
- [0035] Fig. 28 shows tribe memory interface to Global timing.
- [0036] Fig. 29 shows input module stall signals in a memory block.
- [0037] Fig. 30 is a table illustrating the interface between a Tribe and a Memory Block in an embodiment of the invention.
- [0038] Fig. 31 is a table illustrating the interface between a Tribe and the Network Block in an embodiment of the invention.
- [0039] Fig. 32 is a table illustrating the interface between a Tribe and the Interconnect block in an embodiment of the invention.
- [0040] Fig. 33 is a block diagram of an embodiment of the invention.
- [0041] Fig. 34 is a Tribe microarchitecture block diagram.
- [0042] Fig. 35 is shows a fetch pipeline in a tribe in an embodiment of the invention.
- [0043] Fig. 36 is a diagram of a Stream pipeline in tribe architecture.
- [0044] Fig. 37 is a stream pipeline, indicating operand write.
- [0045] Fig. 38 is a stream pipeline, indicating branch execution.

- [0046] Fig. 39 illustrates an execute pipeline.
- [0047] Fig. 40 illustrates interconnect modules.
- [0048] Fig. 41 illustrates a matching matrix for the arbitration problem.
- [0049] Fig. 42 illustrates arbiter stages.
- [0050] Fig. 43 illustrates deadlock resolution.
- [0051] Fig. 44 is an illustration of a crossbar module.
- [0052] Fig. 45 illustrates the tribe to memory interface modules.
- [0053] Fig. 46 illustrates the input module data path.
- [0054] Fig. 47 illustrates a write buffer module.
- [0055] Fig. 48 illustrates the return module data path.
- [0056] Fig. 49 is an illustration of a request buffer and issue module.
- [0057] Fig. 50 is a high level block diagram of a set of components that participate in a lock reservation methodology, in accordance with one embodiment of the present invention.
- [0058] Fig. 51 is a block diagram that shows one of the memory controller blocks of Fig. 50 in greater detail, in accordance with one embodiment of the present invention.
- [0059] Figs. 52-56 show the contents of a lock management storage and a wait list storage structure as they are updated in the process of putting streams on a wait list and taking streams off of a wait list, in accordance with one embodiment of the present invention.
- [0060] Fig. 57 shows an augmented lock management storage which further comprises a value portion, value valid portion, and a dirty bit portion, in accordance with one embodiment of the present invention.

[0061] Fig. 58 shows an augmented lock management storage which further comprises a reservation still valid portion, in accordance with one embodiment of the present invention.

Detailed Description of Embodiment(s)

Overview of Porthos Multi-Threaded Packet Processing Engine

[0062] In a preferred embodiment of the present invention a multithreaded packet processing engine that the inventors term the Porthos chip is provided for stateful packet processing at bit rates up to 20Gbps in both directions. Fig. 1 is an architectural overview for a packet processing engine 101 in an embodiment of the present invention.

[0063] A two bi-directional network port 102 is provided with maximum input and output rates of 10 Gbps each. Packet Buffer 103 is a first-in-first-out (FIFO) buffer that stores individual packets from data streams until it is determined whether the packets should be dropped, forwarded (with modifications if necessary), or transferred off chip for later processing. Packets may be transmitted from data stored externally and may also be created by software and transmitted.

[0064] In preferred embodiments, processing on packets that are resident in the chip occurs in stages, with each stage associated with an independent block of memory. In the example of Fig. 1 here are eight stages 104, labeled (0-7), each associated with a particular memory block 105, also labeled (0-7).

[0065] Each stage 104, called by the inventors a *tribe*, can execute up to 32 software threads simultaneously. A software thread will typically, in preferred embodiments of the

invention, execute on a single packet in one tribe at a time, and may jump from one tribe to another.

[0066] A two HyperTransport interface 106 is used to communicate with host processors, co-processors or other Porthos chips.

[0067] In preferred embodiments of the invention each tribe executes instructions to accomplish the necessary workload for each packet. The instruction set architecture (ISA) implemented by Porthos is similar to the well-known 64-bit MIPS-IV ISA with a few omissions and a few additions. The main differences between the Porthos ISA and MIPS-IV are summarized as follows:

1. Memory Addressing and Register Size

[0068] The Porthos ISA contains 64-bit registers, and utilizes 32-bit addresses with no TLB. There is no 32-bit mode, thus all instructions that operate on registers operate on all 64-bits. The functionality is the same as the well-known MIPS R4000 in 64-bit mode. All memory is treated as big-endian and there is no mode bit that controls endianness. Since there is no TLB, there is no address translation, and there are no protection modes implemented. This means that all code has access to all regions of memory. This would be equivalent to a MIPS processor where all code was running in kernel mode and the TLB mapped the entire physical address space. The physical address space of Porthos is 32-bits, so the upper 32 bits of a generated 64-bit virtual address are ignored and no translation takes place. There are no TLB-related CP0 registers and no TLB instructions.

2. Omitted Instructions

[0069] In preferred embodiments there is no floating-point unit in the Porthos chip, and therefore no floating-point instructions. However the floating-point registers are implemented. Four instructions that load, store, and move data between the regular registers and the floating-point registers (CP1 registers) are implemented (LDC1, SDC1, DMFC1, DMTC1). No branches on CP1 conditions are implemented. Coprocessor 2 registers are also implemented along with their associated load, store and move instructions (LDC2, SDC2, DMFC2, DMTC2). The unaligned load and store instructions are not implemented. The CACHE instruction is not implemented.

3. Synchronization support has been enhanced

[0070] The SC, SCD, LL and LLD instructions are implemented. Additionally, there is an ADDM instruction that atomically adds a value to a memory location and returns the result. In addition there is a GATE instruction that stalls a stream to preserve packet dependencies. This is described in more detail in a following section on flow gating.

4. Timers and Interrupts are changed

[0071] External events and timer interrupts are treated such that new threads are launched. These global events are not thread-specific and are thus not delivered to an active thread. Thus, a thread has no way to enable or disable these events itself, they are configured globally. This is explained in detail in a section below on timers and interrupts.

5. New Set of CP0 Registers

CP7	Sequence Number
CP21	Tribe/Stream Number
CP22	FlowID
CP23	GateVector

6. Thread control instructions

DONE	Terminates a thread
FORK	Forks a new thread
NEXT	Thread migration

7. Special purpose ALU instructions

[0072] Support for string search, including multiple parallel byte comparison, has been provided for in new instructions. In addition there are bit field extract and insert instructions. Finally, an optimized ones-complement add is provided for TCP checksum acceleration.

8. Memory Map

[0073] Porthos has eight ports 107 (Fig. 1) to external memory devices. Each of these ports represents a distinct region of the physical address space. All tribes can access all memories, although there is a performance penalty for accessing memory that is not local to the tribe in which the instructions are executed. A diagram of the memory map is shown in Fig. 2.

[0074] The region of configuration space is used to access internal registers including packet buffer configuration, DMA configuration and HyperTransport port configuration space. More details of the breakdown of this space are provided later in this document.

9. Tribe Migration

[0075] A process in embodiments of the present invention by which a thread executing on a stream in one tribe is transferred to a stream in another tribe is called *migration*. When migration happens, a variable amount of context follows the thread. The CPU registers that are not transferred are lost and initialized to zero in the new tribe. Migration may occur out of order, but it is guaranteed to preserve thread priority as defined by a SeqNum register. Note, however, that a lower priority thread may migrate ahead of a higher priority thread if it has a different destination tribe.

[0076] A thread migrating to the tribe that it is currently in is treated as a NOP. A thread may change its priority by writing to the SeqNum register.

[0077] The thread migration instruction: NEXT specifies a register that contains the destination address and an immediate that contains the amount of thread context to preserve. All registers that are not preserved are zeroed in the destination context. If a thread migrates to the tribe it is already in, the registers not preserved are cleared.

Flow Gating

[0078] Flow gating is a unique mechanism in embodiments of the present invention wherein packet seniority is enforced by hardware through the use of a *gate* instruction

that is inserted into the packet processing workload. When a gate instruction is encountered, the instruction execution for that packet is stalled until all older packets of the same flow have made progress past the same point. Software manually advances a packet through a gate by updating a GateVector register. Multiple gates may be specified for a given packet workload and serialization occurs for each gate individually.

[0079] Packets are given a sequence number by the packet buffer controller when they are received and this sequence number is maintained during the processing of the packet.

[0080] A configurable hardware pre-classifier is used to combine specified bytes from the packet and generate a FlowID number from the packet itself. The FlowID is initialized by hardware based on the hardware hash function, but may be modified by software. The configurable hash function is also be used to select which tribe a packet is sent to. Afterward, tribe to tribe migration is under software control.

[0081] A new instruction is utilized in a preferred embodiment of the invention that operates in conjunction with three internal registers. In addition to the FlowID register and the PacketSequence register discussed above, each thread contains a GateVector register. Software may set and clear this register arbitrarily, but it is initialized to 0 when a new thread is created for a new packet. A new instruction, named GATE, is implemented. The GATE instruction causes execution to stall until there is no thread with the same FlowID, a PacketSequence number that is lower, and with a GateVector in which any of the same bits are zero. This logic serializes all packets within the same flow at that point such that seniority is enforced.

[0082] Software is responsible for setting a bit in the GateVector register when it leaves the critical section. This will allow other packets to enter the critical section. The GateVector register represents progress through the workload of a packet. Software is responsible for setting bits in this register manually if a certain packet skips a certain gate, to prevent younger packets from unnecessarily stalling. If the GateVector is set to all 1s, this will disable flow gating for that packet, since no younger packets will wait for that packet. Note that forward progress is guaranteed since the oldest packet in the processing system will never be stalled and when it completes, another packet will be the oldest packet.

[0083] In a preferred embodiment a seniority scheduling policy is implemented such that older packets are always given priority for execution resources within a processing element. One characteristic of this strictly implemented seniority scheduling policy is that if two packets are executing the exact same sequence of instructions, a younger packet will never be able to overtake an older packet. In certain cases, the characteristic of no overtaking may simplify handling of packet dependencies in software. This is because a no-overtaking processing element enforces a pipelined implementation of packet workloads, so the oldest packet is always guaranteed to be ahead of all younger packets. However, a seniority based instruction scheduler and seniority based cache replacement can only behave with no overtaking if packets are executing the exact same sequence of instructions. If conditional branches cause packets to take different paths, a flow gate would be necessary. Flow gating in conjunction with no-overtaking processing elements allow a clean programming model to be presented that is efficient to implement in hardware.

Event Handling

[0084] Events can be categorized into three groups: triggers from external events, timer interrupts, and thread-to-thread communication. In the first two groups, the events are not specific to any specific physical thread. In the third group, software can signal between two specific physical threads.

Packet Buffer Overview

In this section the following nomenclature is used:

Port - Physically independent full-duplex interface

Channel - Tag associated to each of the packets that arrive or leave through a port.

Interleaving degree - The maximum number of different packets or frames that are in the process of being received or transmitted out.

[0085] The packet buffer (103 Fig. 1) is an on-chip 256K byte memory that holds packets while they are being processed. The packet buffer is a flexible FIFO that keeps all packet data in the order it was received. Thus, unless a packet is discarded, or consumed by the chip (by transfer into a local memory), the packet will be transmitted in the order that it was received. The packet buffer architecture allows for an efficient combination of pass-through and re-assembly scenarios. In a pass-through scenario, packets are not substantially changed; they are only marked, or modified only slightly before being transmitted. The payload of the packets remains substantially the same. Pass-through scenarios occur in TCP-splicing, monitoring and traffic management applications. In a re-assembly scenario, packets must be consumed by the chip and buffered into memory where they are re-assembled. After re-assembly, processing occurs

on the reliable data stream and then re-transmission may occur. Re-assembly scenarios occur in firewalls and load balancing. Many applications call for a combination of pass-through and re-assembly scenarios

[0086] The Packet Buffer module in preferred embodiments interacts with software in the following ways:

- Providing the initial values of some GPRs and CP0 registers at the time a thread is scheduled to start executing its workload.
- Satisfying the requests to the packet buffer memory and the shared memory
- Satisfying the requests to the configuration registers, for instance
 - Hash function configuration
 - Packet table read requests
 - Packet status changes (packet to be dropped, packet to be transmitted out)
 - Performance counters reads
- Allocating space in the packet buffer for software to construct packets.

Frames of packets arrive to the Packet Buffer through a configurable number of ingress ports and leave the Packet Buffer through the same number of egress ports. The maximum ingress/egress interleave degree depends on the number and type of ports, but it does not exceed 4.

[0087] The ingress/egress ports can be configured in one of the following six configurations (all of them full duplex):

- 1 channelized port
- 2 channelized ports
- 4 channelized ports
- 1 non-channelized port

- 2 non-channelized ports
- 4 non-channelized ports

[0088] The channelized port is intended to map into an SPI4.2 interface, and the non-channelized port is intended to map into a GMII interface. Moreover, for the 1-port and 2-port channelized cases, software can configure the egress interleaving degree as follows:

- 1 channelized port: egress interleave degree of 1, 2, 3 or 4.
- 2 channelized ports: egress interleave degree of 1 or 2 per port.

[0089] Software is responsible to complete the processing of the oldest packets that the Packet Buffer module keeps track of in a timely manner, namely before:

1. The subsequent newest packets fill up the packet buffer so that no more packets can be fit into the buffer. At 300MHz core frequency, peak rate of ingress data of 10Gbps and a packet buffer size of 256KB, this will occur in approximately 200 microseconds; and
2. There are 512 total packets in the system, from the oldest to the newest, no matter whether packets in between the oldest and the newest have been dropped (or DMA out to external memory) by software. Otherwise the Packet Buffer module will drop the incoming frames.

[0090] If software does not complete the packets before any of the previous two events occurs, the Packet Buffer module will start dropping the incoming packets until both conditions are no longer met. Note that in this mode of dropping packet data, no flow control will occur on the ingress path, i.e. the packet will be accepted at wire speed but the packets will never be assigned to any tribe, nor its data will be stored in the packet buffer. More details on packet drops is provided below.

Packet Buffer Address Space

[0091] Two regions of the Porthos chip 32-bit physical address space are controlled directly by the Packet Buffer module. These are shown in Fig. 3:

- the packet buffer memory: 256KB of memory where the packets are stored as they arrive. Software is responsible to take them out of this memory if needed (for example, in applications that need re-assembly of the frames)
- the configuration register space: 16KB (not all used) that contains the following sections:
 - the configuration registers themselves: are used to configure some functionality of the Packet Buffer module.
 - the packet table: contains status information for each of the packets being kept track of.
 - the get room space: used for software to request consecutive chunks of space within the packet buffer.

Accesses to the packet buffer address space

[0092] Software can perform any byte read/write, half-word (2-byte) read/write, word (4-byte) read/write or double word (8-byte) read/write to the packet buffer. Single quad-word (16-byte) and octo-word (32-byte) read requests are also allowed, but not the single quad-word and octo-word writes. To write 2 or 4 consecutive (in space) double words, software has to perform, respectively, 2 or 4 double-word writes. The Packet Buffer will not guarantee that these consecutive writes will occur back to back; however, no other access from the same tribe will sneak in between the writes (but accesses from other tribes can).

[0093] Even though the size of the packet buffer memory is 256KB, it actually occupies 512KB in the logical address space of the streams. This has been done in order to help minimizing the memory fragmentation that occurs incoming packets are stored into the packet buffer. This mapping is performed by hardware; packets are always stored consecutively into the 512KB of space from the point of view of software.

[0094] Software should only use the packet buffer to read the packets that have been stored by the Packet Buffer module, and to modify these packets. The requests from the 8 tribes are treated fairly; all the tribes have the same priority in accessing the packet buffer.

Accesses to the configuration register physical address space

[0095] The configuration registers are logically organized as double words. Only double word reads and writes are allowed to the configuration register space. Therefore, if software wants to modify a specific byte within a particular configuration register, it needs to read that register first, do the appropriate shifting and masking, and write the whole double word back.

Writes to the reserved portion of the configuration register space will be disregarded. Reads within this portion will return a value of 0.

[0096] Some bits of the configuration registers are reserved for future use. Writes to these bits will be disregarded. Reads of these bits will return a value of 0.

[0097] Unless otherwise noted, the configuration registers can be both read and written. Writes to the packet table and to the read-only configuration registers will be disregarded.

[0098] Software should change the contents of the configuration registers when the Packet Buffer is in quiescent mode, as explained below, and there are no packets in the system, otherwise results will be undefined. Software can monitor the contents of the 'packet_table_packets' configuration register to figure out whether the Packet Buffer is still keeping packets or not.

Configuration register list

[0099] All the configuration registers have an after-reset value of 0x0 unless otherwise specified. Figs. 4a-4d comprise a table listing all of the configuration registers. The following sections provide more details on some of the configuration registers.

Hashing function

[0100] Fig. 5 illustrates the hash function hardware structured into two levels, each containing one (first level) or four (second level) hashing engines. The result of the hashing engine of the first level is two-fold:

- a 16-bit value, named the flow identifier (or flowId for short). This value will be provided to the tribe as part of the initial migration of the packet. Software may use this value, for example, as an initial classification of the packet into a flow.
- a 2-bit value, that is used by the hardware to select the result of one of the 4 hashing functions that compose the second level of the hashing hardware.

[0101] Each of the four hashing functions in the second level generates a 3-bit value that corresponds to a tribe number. One of these four results is selected by the first level, and becomes the number of the tribe that the packet is going to initially migrate into.

[0102] All four hashing engines in the second level are identical, and the single engine in the first level is almost also the same as the ones in the second level. Each hashing engine can be configured independently. The following is the configuration features that are common to all the hashing engines:

- select vector [0..i..63] configuration register: each bit of this vector determines whether byte i of the packet will be selected to compute the result of the hashing engine (1) or not (0).
- position vector [0..i..63] configuration register: the 16-bit result of the hashing engine is computed using two 8-bit XOR functional units, one for the upper 8-bits and one for the lower 8-bits. In the case that byte i was selected by the select vector, bit i in the position vector determines whether the byte will be used to compute the lower 8 bits of the 16-bit flowId result (0) or the upper 8 bits (1). If the byte was not selected in the select vector, the corresponding bit in the position vector is a don't care.

[0103] For the first level hashing engine, there exists a skip configuration register that specifies how many LSB bits of the 16-bit result will be skipped to determine the chosen second level hashing engine. If the skip value is, for instance, two, then the second level hashing engine will be chosen using bits [2..3] of the 16-bit result. Note that the skip configuration register is only used to select the second level hashing function and it does not modify the 16-bit result that becomes the flowId value.

[0104] For each of the second level hashing engines there also exists a skip configuration register performing the same manipulation of the result as in the first level.

After this shifting of the result, another manipulation is performed using two other configuration registers; the purpose of this manipulation is to generate a tribe number out of a set of possible tribe numbers. This total number of tribes in this set is a power of 2 (i.e. 1, 2, 4 or 8), and the set can start at any tribe number. Example of sets are [0,1,2,3], [1,2,3,4], [2,3], [7], [0,1,2,3,4,5,6,7], [4,5,6,7], [6,7,0,1], [7,0,1,2], etc. This manipulation is controlled by two additional configuration registers, one per each of the second-level hashing engines:

- first: 3-bit value that specifies which is the first tribe of the set (0: tribe 0, .. 7: tribe 7)
- total: 2-bit vector that specifies how many consecutive tribes the set has (0:1 tribe, 1:2 tribes, 2:4 tribes, 3:8 tribes)

[0105] The maximum depth that the hashing hardware will look into the packet is 64 bytes from the start of the packet. If the packet is smaller than 64 bytes and more bytes are selected by the select vectors, results will be undefined.

[0106] Software should be careful in configuring the hashing function hardware since only non-variant bytes across all the packets of the same flow should be selected to perform the hashing computation; otherwise, different flow identifiers for the packets of the same flow might be generated.

Quiescent mode

[0107] The Packet Buffer module is considered to be in *quiescent mode* whenever it is not receiving (and accepting) any packet and software has written a 0 in the 'continue' configuration register. Note that the Packet Buffer can be in quiescent mode and still have valid packets in the packet table and packet buffer. Also note that all the

transmission-related operations will be performed normally; however any incoming packet will be dropped since the 'continue' configuration register is zero.

[0108] When the contents of the 'continue' configuration register toggles from 0 to 1, the Packet Buffer module will perform the following operation:

- any new incoming packet that starts arriving after the setting of the 'continue' configuration register takes place physically will be accepted (it may be eventually dropped for other reasons as explained below).

When the toggling is from 1 to 0, the following operation takes place:

- any packet that was currently being received when the clearing of the 'continue' configuration register occurs will be fully received.
- any new incoming packet that starts arriving after the setting of the configuration register takes place will be fully dropped.

[0109] Software should put the Packet Buffer module in quiescent mode whenever it wants to modify configurable features (note that the Packet Buffer comes out of reset in quiescent mode). The following are the steps software should follow:

1. Write a 0 into the 'continue' configuration register.
2. Monitor the 'status' register until bit 1 is set. When this occurs, the quiescent state has been entered.
3. Configure the desired feature
4. Write a 1 into the 'continue' configuration register to allow new incoming packets to be accepted.

[0110] If the above steps are followed, there will be no packets being received when the 0 to 1 transition happens on the 'continue' configuration register. This is not true if software does not wait for quiescent mode before setting the 'continue' configuration register; in this case, the Packet Buffer may keep receiving the packet it was receiving when the 1 to 0 transition took place.

Performance counters

[0111] There are a total of 128 performance events in the Packet Buffer module (63 defined, 65 reserved) that can be monitored by software. Out of these events, a total of 8 can be monitored at the same time. a 48-bit counter is assigned to one particular event and increments the value of the counter by the proper quantity each time the event occurs. Events are tracked by hardware every cycle.

[0112] Software can configure which event to monitor in each of the 8 counters. The contents of the counters are made visible to software through a separate set of configuration registers.

Fig. 6 is a table showing the performance events that can be monitored.

Internal state probes

[0113] Software can probe the internal state of the Packet Buffer module using the 'internal_state_number' configuration register. When software reads this configuration register, the contents of some internal state are provided. The internal state that is provided to software is yet TBD. It is organized in 64-bit words, named internal state words. Software writes an internal state word into the 'internal_state_number'

configuration register previously to reading the same configuration register to get the contents of the state. This feature is intended only for low level debugging.

Egress channel determination

[0114] When software writes into the 'done' or 'egress_path_determined' configuration register it provides, among other information, the egress channel associated to the transmission. This channel ranges from 0 to 255, and software actually provides a 9-bit quantity, named the encoded egress channel, that will be used to compute the actual egress channel. Fig. 7 is a table that specifies how the actual egress channel is computed from the encoded egress channel.

[0115] The egress channel information is only needed in the case of channelized ports. Otherwise, this field is treated as a don't care.

Egress port determination

[0116] When software writes into the 'done' or 'egress_path_determined', it provides, along with other information, the egress port associated to the transmission. This port number ranges from 0 to 3 (depending on how the Packet Buffer module has been configured), and software actually provides a 5-bit quantity, named the encoded egress port, that will be used to compute the actual egress port. Fig. 8 is a table that shows how the actual egress port is computed from the encoded egress port.

Completing and dropping packets

[0117] Software eventually has to decide what to do with a packet that sits in the packet buffer, and it has two options:

- Complete the packet: the packet will be transmitted out whenever the packet becomes the oldest packet in the packet buffer.
- Drop the packet: the packet will be eventually removed from the packet buffer.

In both cases, the memory that the packet occupies in the packet buffer and the entry in the packet table will be made available to other incoming packets as soon as the packet is fully transmitted out or dropped. Also, in both cases, the Packet Buffer module does not guarantee that the packet will be either transmitted or dropped right away. Moreover, there is also no upper limit on the time the packet might sit in the packet buffer before it gets transmitted out or dropped. An example of a large period of time between software requests a packet to be transmitted and the actual start of the transmission occurs in an egress-interleave of 1 case when software completes a packet that is not the oldest one, and the oldest packet is not completed nor dropped for a long time.

[0118] Software completes and drops packets by writing into the 'done' and 'drop' configuration registers, respectively. The information provided in both cases is the sequence number of the packet. For the completing of packets, the following information is also provided:

- Header growth offset: an 10-bit value that specifies how many bytes the start of the packet has either grown (positive value) or shrunk (negative value) with respect the original packet. The value is encoded in 2's complement. If software does not move the start of the packet, this value should be 0.
- Encoded egress channel.
- Encoded egress port.

[0119] The head of a packet is allowed to grow up to 511 bytes and shrink up to the minimum of the original packet size and 512).

[0120] Software should either complete or drop the packet. Results will be undefined if multiple completions/drops occur for the same packet. Moreover, there is no guarantee that the packet data stored in the packet buffer will be coherent after software has completed or drop the packet.

Egress path determination

[0121] The egress path information (egress port and, in case of channelized port, the egress channel) is mandatory and needs to be provided when software notifies that the processing of a particular packet has been completed. However, software can at any time communicate the egress path information of a packet, even if the processing of the packet still has not been completed. Software does this by writing into the 'egress_path_determination' configuration register the sequence number of the packet, the egress port and, if needed, the egress channel.

[0122] Of course, the packet will not be transmitted out until software writes into the 'done' command, but the early knowledge of the egress path allows the optimization of the scheduling of packets to be transmitted out in the following cases:

- 1-port channelized with egress interleave of 2, 3 or 4.
- 2-port with egress interleave of 1 or 2
- 4-port with egress interleave of 1

Note that even if software notified the egress path information through the 'egress_path_determination' configuration register, it needs to provide it again when notifying the completion of the processing through the 'done' configuration register.

GetRoom Command

[0123] Software can transmit a packet that it has generated through a GetRoom mechanism. This mechanism works as follows:

- Software requests some space to be set aside in the packet buffer. This is done through a regular read to the GetRoom space of the configuration space. The address of the load is computed by adding the requested size in bytes to the base of the GetRoom configuration space.
- ThePacket Buffer module will reply to the load:
 - Unsuccessfully: it will return a '1' in the MSB bit and '0' in the rest of the bits
 - Successfully: it will return in the 32 LSB bits the physical address of the start of the space that has been allocated, and in bits [47..32] the corresponding sequence number associated to that space.
- Software, upon the successful GetRoom command, will construct the packet into the requested space.
- When the packet is fully constructed, software will complete it through the packet complete mechanism explained before.

[0124] Note that for software-created packets, it is expected the delta to be always 0 when the packet is completed since the header growth offset is not taken into account when the size is allocated.

Configuring the number and type of ports

[0125] Software can configure the number of ports, whether they are channelized or not, and the degree of interleaving. All the ports will have the same channelization and interleaving degree properties (ie it can no happen that one port is channelized and another port not).

[0126] A port is full duplex, thus there is the same number of ingress and egress ports. Fig. 9 shows six different configurations in terms of number of ports and type of port (channelized/non-channelized). For each configuration, it is shown the interleaving degree allowed per port (which can also be configured if more than an option exists). The channelization and interleaving degree properties applies to both the ingress and egress paths of the port.

[0127] Software determines the number of ports by writing into the 'total_ports' configuration register, and the type of ports by writing into the 'port_type' configuration register.

[0128] For the 1-port and 2-port channelized cases, software can configure the degree of egress interleaving. The ingress interleaving degree can not be configured since it is determined by the sender of the packets, but in any case it can not exceed 4 in the 1-port channelized case and 2 in the 2-port channelized case; for the other cases, it has to be 1 (the Packet Buffer module will silently drop any packet that violates the maximum ingress interleaving degree restriction).

[0129] The egress interleaving degree is configured through the 'islot_enabled' and 'islot_channel_0..3' configuration registers. An "islot" stands for interleaving slot, and is

associated to one packet being transmitted out, across all the ports. Thus, the maximum number of islots at any time is 4 (for the 1-port channelized case, all 4 islots are used when the egress port is configured to support an interleaving degree of 4; for the 1-port case, up to 4 packets can be simultaneously being transmitted out -one per port-). Note that the number of enabled islots should coincide with the number of ports times the egress interleaving degree.

[0130] Notification about how many ports there are is made through the 'total_ports' configuration register. It will also be notified about the type of the ports (all have to be of the same type) through the 'port_type' configuration register.

[0131] For channelized (ie SPI4.2) ports, software will configure a range of channel numbers that will be transmitted in each of the 4 outbound "interleaving slots" ("islot" for short). This configuration is performed through the 'islot_channels_0', ..., 'islot_channels_3'. For example, if there is one single SPI4.2 port and

- islot_channels_0: 0-63
- islot_channels_1: 64-127
- islot_channels_2: 128-191
- islot_channels_3: 192-255

then the output packet data may have, for example, channels 0, 54, 128 and 192 interleaved (or channels 0, 65, 190, 200, etc.) but it will never have channels 0 and 1 interleaved.

[0132] For the 2-port SPI4.2 scenario, islot0 and islot1 are assigned to port 0, and islot2 and islot3 are assigned to port 1. Thus, the maximum interleaving degree per port is 2. With the same channel range example above, port 0 will never see channels 128-255,

and it will never see channels 70 and 80 interleaved. The following configuration is a valid one that covers all the channels in each egress port:

- islot0_channels: 0-127
- islot1_channels: 128-255
- islot2_channels: 0-127
- islot3_channels: 128-255

[0133] Note that if software fails to cover a particular channel with an islot assigned to the channel and packets with that particular channel have to be transmitted to that port, results will be undefined.

[0134] Software can also disable the islots to force no interleaving on the SPI4.2 ports. This is done through the 'islot_enable' configuration register. For example, in the 1-port SPI4.2 case, if 'islot_enable' is 0x4 (islot2 enabled and the rest disabled), then an interleaving of just 1 will happen on the egress port, and for the range of channels specified in the 'islot2_channels' configuration register.

[0135] For the 4-port GMII case, the channel-range associated to each of the islots is meaningless since a GMII port is not channelized. An interleaving degree of 1 will always occur at each egress port.

Software can complete all packets in any order, even those that will change its ingress port or channel. There is no need for software to do anything special when completing a packet with a change of its ingress port or channel, other than notifying the new egress path information through the 'done' configuration register.

Initial migration

[0136] When packets have been fully received by the Packet Buffer module and they have been fully stored into the packet buffer memory, the first migration of those packets into one of the tribes will be initiated. The migration process consists of a request of a migration to a tribe, waiting for the tribe to accept the migration, and providing some control information of the packet to the tribe. This information will be stored by the tribe in some software visible registers of one of the available streams.

[0137] The Packet Buffer module assigns to each packet a flow identification number and a tribe number using the configurable hashing hardware, as described above. The packets will be migrated to the corresponding initial tribes in exactly the same order of arrival. This order of arrival is across all ingress ports and, if applicable, all channels. If a packet has to be migrated to a tribe that has all its streams active, the tribe will not accept the migration. The Packet Buffer module will keep requesting the migration until the tribe accepts it.

[0138] After the migration has taken place, the following registers are initialized in one of the streams of the tribe:

- PC: initialized with the value in the corresponding 32-bit program_counter configuration register. Note that all the streams within a tribe that are activated due to an initial migration will start executing code at the same initial program counter.
- CP0.22: the flow identification number, a 16-bit value obtained by the hashing hardware.

- CP0.7: the sequence number, a 16-bit value that contains the order of arrival of the packet. If a packet A fully arrived right after a packet B, the sequence number of A will be the sequence number of B plus 1 (assuming no other packet from other port completed nor a GetRoom command successfully happened in between). The sequence number wraps around at 0xFFFF.
- GPR.30: the ingress port (bits 9-10) and channel of the packet (bits 0-7).
- GPR.31: the 32-bit logical address where the packet resides. This address points to the first byte of the packet. If the NET module left space at the front of the packet (specified by the header_growth_offset configuration register), this address still points to the first byte that arrived of the packet, not to the first byte of the added space.

Hardware-initiated packet drops

[0139] There are two types of packet drops:

- Software-initiated drops: software explicitly requests a particular packet to be dropped.
- Hardware-initiated drops: a packet is dropped because there is no space to store the packet or its control information.

Furthermore, the cause of a hardware-initiated drop could be one of the following:

- The packet buffer is full. If the occupancy of the packet buffer when a new packet starts arriving is such that it cannot be guaranteed that a maximum-size packet could be fit in, the hardware will drop that incoming packet.
- The packet table is full. If the table that is used to store the packet descriptors (control information) of the packets is full when a new packet starts arriving, the hardware will drop that incoming packet. The packet table is considered to be full when there are less than 4 entries available in the packet table upon a packet arrival.

- The 'continue' configuration register is 0. The Packet Buffer module comes out of reset with a 0 in the 'continue' configuration register. Until software writes a 1 in there, any incoming packet will be dropped.
- Interleaving degree violation. If an ingress port violates the maximum degree of packet interleaving that the NET supports.
- The size of the packet being received exceeds the maximum allowed packet size. The maximum packet size that can be accepted is 65536 bytes. Software can override this maximum size to a lower value, from 1KB to 64KB, always in increments of 1K (see configuration register 'max_packet_size')¹. If an incoming packet is detected that it may be over the maximum packet size allowed when the next valid data of the packet arrives, the packet is forced to finish right away and the rest of the data that eventually will come from that packet will be dropped by the hardware. Therefore, a packet that exceeds the maximum allowed packet size will be seen by software as a valid packet of a size that varies between the maximum allowed size and the maximum allowed size minus 7 (since up to 8 valid bytes of packet data can arrive every cycle).
- The ingress port notifies that the packet currently being received has an error. This notification can occur at any time during the reception of the packet.

[0140] Note that entire packets are dropped. When a packet is dropped by hardware, there is no interrupt generated. Software can check at any given time the total number of packets that have been dropped due to each of the hardware-initiated causes by monitoring specific performance events.

Porthos Instruction Set

[0141] The Porthos instruction set in a preferred embodiment of the present invention is as follows:

ALU

Arithmetic

ADD, ADDU, SUB, SUBU, ADDI, ADDIU, SLT, SLTU, SLTI, SLTIU
DADD, DADDU, DSUB, DSUBU, DADDI, DADDIU,

Logical

AND, OR, XOR, NOR, ANDI, ORI, XORI, NORI

Shift

SLL, SRL, SRA, SLLV, SRLV, SRAV,
DSLL, DSRL, DSRA, DSLLV, DSRLV, DSRAV, DSLL32, DSRL32

Multiply/Divide

MULT, MULTU, DIV, DIVU,
DMULT, DMULTU, DDIV, DDIVU

Memory

Load

LB, LH, LHU, LW, LWU, LD

Store

SB, SH, SW, SD

Synchronization

LL, LLD, SC, SCD

SYNC

ADDM

Control

Branch

BEQ, BNE, BLEZ, BGTZ, BLTZ, BGEZ, BLTZAL, BGEZAL

Jump

J, JR, JAL, JALR

Trap

TGE, TGEU, TLT, TLTU, TEQ, TNE,

TGEI, TGEIU, TLTI, TLTIU, TEQI, TNEI

Miscellaneous

SYSCALL, BREAK, ERET, *NEXT*, *DONE*, *GATE*, *FORK*

Miscellaneous

MFHI, MTHI, MFLO, MTLO,

MTC0, MFC0

CP0 Registers

The CP0 registers in a preferred embodiment of the invention are as follows:

Config

TribeNum, StreamNum (CP0 Register 21)

Status

EPC

Cause

FlowID (CP0 Register 22)

GateVector (CP0 Register 23)

SeqNum (CP0 Register 7)

Microarchitecture Description of the Global Block of the Porthos Chip

Overview of the Global Block

[0142] Referring again the Fig. 1, a Global Unit 108 provides a number of functions, which include hosting functions and global memory functions. Further, interconnections of global unit 108 with other portions of the Porthos chip are not all indicated in Fig. 1, to

keep the figure relatively clear and simple. Global block 108, however, is bus-connected to Network Unit and Packet Buffer 103, and also to each one of the memory units 105.

[0143] The global (or “GBL”) block 106 of the Porthos chip is responsible for the following functions:

- Implements a memory controller for external EPROM
- Interfaces with two HyperTransport IP blocks
- Provides input and output paths for the general purpose I/Os
- Satisfies external JTAG commands
- Generates interrupts as a result of HT, GPIO or JTAG activity
- Interfaces with the network block to satisfy HyperTransport requests to packet buffer memory
- Provides a path for memory interconnection among the different local memories of the tribes

Nomenclature for Global Block processes:

- Request: an access from a source to a destination to obtain a particular address (read request) or to modify a particular address (write request)
- Response: a petition from a source to a destination to provide the requested data (in case of a read request) or to acknowledge that the request has been fulfilled (in case of a write request)
- Transaction: composed of the request initiated by the source A to destination B and the corresponding response initiated by the source B to the destination A. Note that a transaction is always composed of a request and a response; if the request is for a write, the response will provide just an acknowledge that the write has been fulfilled.

Fig. 10 is a top-level module diagram of the GBL block 108.

The GBL is composed of the following modules:

- Local memory queues 1001 (LMQ): there is one LMQ per each local memory block. The LMQ contains the logic to receive transactions from the attached local memory to another local memory, and the logic to send transactions from a local memory to the attached local memory.
- Routing block 1002 (RTN): routes requests from the different sources to the different destinations.
- EPROM controller 1003 (EPC): contains the logic to interface with the external EPROM and the RTN
- HyperTransport controller 1004 (HTC): there is one HTC per HyperTransport IP block.
- General purpose I/O controller 1005 (IOC): contains the logic to receive activity from the GPIO input pins and to drive the GPIO output pins.
- JTAG controller 1006 (JTC): contains the logic to convert JTAG commands to the corresponding requests to the different local memories.
- Interrupt handler 1007 (INT): generates interrupts to the tribes as a result of HT, JTAG or GPIO activity
- Network controller 1008 (NTC): logic that interfaces to the network block to satisfy HT commands that affect the packet buffer memory without software intervention.

Local Memory Queues block 1001 (LMQ)

[0144] Block 1001 contains the logic to receive transactions from the attached local memory to another local memory, and the logic to send transactions from a local memory to the attached local memory. Fig. 11 is an expanded view showing internal components of block 1001.

Description of LMQ 1001:

[0145] LMQ block 1001 receives requests and responses from the local memory block. A request/response contains the following fields (the number in parentheses is the width in bits):

- valid (1): asserted when the local memory block sends a request or response. If de-asserted, the rest of the fields are "don't care".
- data (64): the data associated to a write request or a read response; otherwise (read request or write response) is "don't care".
- stream (5): in case of a request, this field contains the number of the stream within the tribe that performs the request to the local memory. In case of a response, this field contains the same value received on the corresponding request.
- regdest (5): in case of a read request, this field contains the register number where the requested data will be stored. In case of a response, this field contains the same value received on the corresponding request.
- type (3): specifies the type of the request (signed read , unsigned read, write) or response (signed read, unsigned read, write).
- address (32): in case of a request, this field contains the physical address associated to the read or write. In case of a response, contains the same value received on the corresponding request.

[0146] LMQ block 1001 looks at the type field to figure out into which of the input queues the access from the local memory will be inserted into.

[0147] The LMQ block will independently notify to the local memory when it can not accept more requests or responses. The LMQ block guarantees that it can accept one more request/response when it asserts the request/response full signal.

[0148] The LMQ block sends requests and responses to the local memory block. A request/response contains the same fields as the request/response received from the local memory block. However the address bus is shorter (23 bits) since the address space of each of the local memories is 8MB.

[0149] The requests are sent to the local memory in the same order are received from the RTN block. Similarly for the responses. When there is an available request and an available response to be sent to the local memory, the LMQ will give priority to the response. Thus, newer responses can be sent before than older requests.

Routing block 1002 (RTN)

[0150] This block contains the paths and logic to route requests from the different sources to the different destinations. Fig. 12 shows this block (interacting only to the LMQ blocks).

Description

[0151] The RTN block 1002 contains two independent routing structures, one that takes care of routing requests from a LMQ block to a different one, and another one that routes responses from a LMQ block to a different one. The two routing blocks are independent and do not communicate. The RTN can thus route a request and a response originating from the same LMQ in the same cycle.

[0152] The result of routing of a request/response from a LMQ to the same LMQ is undefined.

Microarchitecture Description of the Network Block 103 of the Porthos Chip

Overview

[0153] The network (or “NET”) block 103 (Fig. 1) of the Porthos chip is responsible for the following functions:

- Receiving the packets from 1, 2 or 4 ports and storing them into the packet buffer memory.
- Notifying one of the tribes that a new packet has arrived, and providing information about the packet to the tribe.
- Satisfying the read and write requests to the packet buffer memory performed by the different tribes and the global block.
- Keeping track of the status of a packet.
- Monitoring the oldest packet to each of the egress ports and sending it out to the corresponding port if it has already been processed
- Providing a DMA mechanism to the tribes to transfer data out of the packet buffer memory and into the external global memory.

[0154] The NET block will always consume the data that the ingress ports provide at wire speed (up to 10Gbps of aggregated ingress bandwidth), and will provide the processed packets to the corresponding egress ports at the same wire speed. The NET block will not perform flow control on the ingress path.

[0155] The data will be dropped by the NET block if the packet buffer memory can not fit in any more packets, or the total number of packets that the network block keeps track of reaches its maximum of 512, or there is a violation by the SPI4 port on the maximum number of interleaving packets that it sends, or there is a violation on the maximum packet size allowed, or software requests incoming packets to be dropped.

[0156] Newly arrived packets will be presented to the tribes at a rate no lower than a packet every 5 clock cycles. This provides the requirement of assigning a 40-byte packet to one of the tribes (assuming that there are available streams in the target tribe) at wire speed. The core clock frequency of the NET block is 300MHz.

[0157] Frames of packets arrive to the NET through a configurable number of ingress ports and leave the NET through the same number of egress ports. The maximum ingress/egress interleave degree depends on the number and type of ports, but it does not exceed 4.

[0158] The ingress/egress ports can be configured in one of the following six configurations (all of them full duplex):

- 1 channelized port
- 2 channelized ports
- 4 channelized ports
- 1 non-channelized port
- 2 non-channelized ports
- 4 non-channelized ports

[0159] The channelized port is intended to map into an SPI4.2 interface, and the non-channelized port is intended to map into a GMII interface. Moreover, for the 1-port and 2-port channelized cases, software can configure the egress interleaving degree as follows:

- 1 channelized port: egress interleave degree of 1, 2, 3 or 4.
- 2 channelized ports: egress interleave degree of 1 or 2 per port.

[0160] The requirement of the DMA engine is to provide enough bandwidth to DMA the packets out of the packet buffer to the external memory (through the global block) at wire speed.

Block diagram

[0161] Fig. 14 shows the block diagram of the NET block 103. The NET block is divided into 5 sub-blocks, namely:

- PortInterface (PIF): responsible for receiving the packets on the different ingress ports (1,2 or 4) and deciding to which of the 4 ingress interleaving slots the data of the packet belongs to, and responsible for interfacing with the egress ports also on the egress path.
- PacketLoader (PLD): responsible for:
 - Applying a hash function to the packet being received for the purpose of flow identification and for deciding to which tribe the packet will be eventually assigned to

- Deciding where to store the packet into the packet buffer, and performing all the necessary writes
 - Allocating an entry in the packet table with the control information of the newly arrived packet
 - Providing the information of newly arrived packets, in the order of arrival across all ingress ports, to the different tribes for processing
 - Maintaining the status of each of the packets in the packet table, in particular, whether the packets have been completely processed by the tribes or not yet.
 - Monitoring the oldest packet in the packet table to decide what to do with it (skip it if the packet is not valid –ie software has explicitly requested to the NET block to drop the packet-; transmit it out to the corresponding egress port if the packet has been completed; or nothing if the packet is still active), and do this for each of the egress interleaving slots.
- PacketBufferController (PBC): its function is to provide some buffering for the requests of each of the sources of accesses to the packet buffer memory, and perform the scheduling of these requests to the different banks of the packet buffer memory. The different sources are: the network ingress path, the network egress path, the DMA engine (TBD), the global block and the 8 tribes. The scheduler implements a fixed priority scheme in the order listed before (ingress path having the highest priority). The 8 tribes are treated fairly among them.
 - PacketBufferMemory (PBM): it contains the packet buffer memory, divided into 8 interleaved banks. Performs the different accesses that the PBC has scheduled to each of the banks, and routes the result to the proper source. This block also performs the configuration register reads and writes, thus interacting with the different sub-blocks to access the corresponding configuration register.

[0162] The following sections provide detailed information about each of the blocks in the Network block. The main datapath busses are shown in bold and they are 64-bit

wide. Moreover, all busses are unidirectional. Unless otherwise noted, all the signals are point-to-point and asserted high.

All outputs of the different sub-blocks (PIF, PLD, PBM and PBC) are registered.

PortInterface block 1401 (PIF)

Detailed description

[0163] The PIF block has two top-level functions: ingress function and egress function. Fig. 15 shows its block diagram.

Ingress function

[0164] The ingress function interfaces with the SPI4.2/GMII ingress port, with the PacketLoader (PLD) and with the PacketBufferMemory (PBM).

SPI4.2/GMII port

[0165] From a SPI4.2/GMII port, it receives the following information:

- Valid (1): if asserted, validates the rest of the inputs. It specifies that SPI4 is sending valid data in the current cycle.
- Data (64): contains the 64 bits of packet data provided by the SPI4. This 64-bit vector is logically divided into 8 bytes.
- End_of_packet (1): if asserted, it specifies that valid data is the last data of the packet.
- Last_byte (3): pointer to the last valid MSB byte in 'data'. If all 8 bytes are valid, 'last_byte' is 7; if only 1 byte is valid, 'last_byte' is 0. If 1 or more bytes are valid, they are right aligned (first valid byte is byte0, then byte1, etc.). It can not occur that,

for example, byte 0 and 2 are valid, but not byte 1. In other words, if the data is not the end of the packet, then 'last_byte' should be 3; if the data is the end of the packet, then 'last_byte' can take any value.

- Channel (8): the channel associated to the packet data received. The SPI4 protocol allows up to 256 channels. This field is a don't care in case of a GMII port.

[0166] Every cycle, a port may send data (of a single packet only). But packets can arrive interleaved (in cycle x, packet data from a packet can arrive, and in cycle x+1 data from a different –or same- packet may arrive). The ingress function will know to which of the packets being received the data belongs to by looking at the channel number. Note that packets can not arrive interleaved in a GMII port.

[0167] In a SPI4.2 port, up to 256 packets (matching the number of channels) can be interleaved. However, the Porthos chip will only handle up to 4. Therefore, any packet interleaving violation will be detected and the corresponding packet data will be dropped by the ingress function.

[0168] The ingress function monitors the packets and the total packet data dropped due to the interleaving violation.

[0169] The number of total ports is configurable by software. There can be 1, 2 or 4 ingress ports. In case of a single SPI4.2 port, the maximum interleaving degree is 4. In case of 2 SPI4.2 ports, the maximum interleaving degree is 2 per port. In the case of 4 ports, no interleaving is allowed in each port.

[0170] For SPI4.2 ports, when valid data of a packet arrives, the ingress function performs an associative search into a 4-entry table (the channel_slot table). Each entry of these table (called slot), corresponds to one of the packets that is being received. Each entry contains two fields: active (1) and channel (8). If the associative search finds that the channel of the incoming packet matches the channel value stored in one of the active entries of the table, then the packet data corresponds to a packet that is being received. If the associative search does not find the channel in the table, then two situations may occur:

- There is at least one non-active entry in the portion of the table associated to the ingress port: in this case, the valid data received is the start of a new packet. The entry is marked as active and the incoming channel is stored into that entry.
- All the entries in the portion of the table associated to the ingress port are active. This implies a protocol violation and the packet data will be dropped. The hardware sets a Xth bit in a 256-bit array (where X is the incoming channel number) called violating_channels.

[0171] For the 1-SPI4 port, all the 4 entries of the table are available for the port; for the 2-SPI4 port, the first 2 entries are allocated for port 0, and the second two entries for port 1, thus forcing a maximum ingress interleaving degree of 2 per port.

[0172] The incoming channel associated to every valid data is looked up in the violating_channels array to figure out whether the packet data needs to be dropped (ie whether the valid data corresponds to a packet that, when it first arrived, violated the interleave restriction). If the corresponding bit in the violating_channels is 0, then the

channel is looked up in the channel_slot table, otherwise the packet data is dropped. If the packet data happens to be the last data of the packet, the corresponding bit in the violating_channels array is cleared.

[0173] There is no flow control between the SPI4 ingress port and the ingress function.

PLD interface:

[0174] If the packet data is not dropped, it is inserted into a 2-entry FIFO. Each entry of this FIFO contains the information that came from the SPI4 ingress port: data (64), end_of_packet (1), last_byte (3), channel (8), and information generated by the ingress function: slot (2), start_of_packet (1).

[0175] Only valid packet data of packets that comply with interleave restriction will be stored into the FIFO. If the FIFO is not empty, the contents of the head entry of the FIFO are provided to the PLD and the head entry is removed.

[0176] A logic exists that will monitor the head of each of the 4 fifos and will send valid data to the PLD in a round-robin fashion. This logic is capable of sending up to 8 bytes of valid data to the PLD per cycle. At a core frequency of 300MHz, it implies that the network block can absorb packet data at a peak close to 20Gbps.

[0177] There is no flow control between the ingress function and the PLD block. This implies that the aggregated bandwidth of across all ingress ports should be less than 19.2Gbps (for 300MHz core frequency operation).

There are no configuration registers affecting the ingress function.

Performance events 0-11 are monitored by the ingress function.

Egress function

[0178] The egress function interfaces with the egress ports, the PacketLoader (PLD) and the PacketBufferMemory (PBM).

PBM interface

[0179] The egress function receives packet data from the PBM of packets that reside in the packet buffer. There is an independent interface for each of the egress interleaving slots, as follows:

- Valid (1): if asserted, validates the rest of the inputs. It specifies that valid data is sent in the current cycle or not.
- Data (64): contains the 64 bits of packet data provided. This 64-bit vector is logically divided into 8 bytes.
- End_of_packet (1): if asserted, it specifies that valid data is the last data of the packet.
- Last_byte (3): pointer to the last valid MSB byte in 'data'. If all 8 bytes are valid, 'last_byte' is 7; if only 1 byte is valid, 'last_byte' is 0. If 1 or more bytes are valid, they are right aligned (first valid byte is byte0, then byte1, etc.). It can not occur that, for example, byte 0 and 2 are valid, but not byte 1. In other words, if the data is not

the end of the packet, then 'last_byte' should be 3; if the data is the end of the packet, then 'last_byte' can take any value.

- Port (2): the outbound port.
- Channel (8): the outbound channel associated to the packet data. Meaningless if the egress port is not channelized.

[0180] A total of up 4 FIFOs, one associated to each egress interleaving slots, store the incoming information. Each FIFO has 8 entries.

[0181] Whenever the number of occupied entries in the PBM FIFO is 5 or more, a signal is provided to the PBC block as a mechanism of flow control. xx There could be at most 5 chunks of packet data already read and in the process of arriving to the egress function.

Egress port interface:

[0182] A logic exists that will look at the head of each of the 4 FIFOs and, in a round-robin fashion, will send the valid data to the corresponding egress port. Note that if 4 egress ports exist, then there is a 1-to-1 correspondence between a fifo and a port. If 2 channelized ports exist, then the round-robin logic is applied between fifo 0 and fifo 1 for port 0 and fifo2 and fifo3 for port1. In the case of 2 non-channelized ports, either islot 0 or islot 1 is disabled (implying that either fifo 0 or fifo 1 is empty), and similarly for islot2 and islot3 (for fifo 2 and fifo 3). In the case of 1 channelized port, the round robin prioritization is applied among all the fifos; for the 1 non-channelized port case, all except one fifo should be empty.

[0183] The round robin logic works in parallel for each of the egress ports.

[0184] The valid contents of the head of the FIFO that the prioritization logic chooses are sent to the corresponding egress port. This information is structured in the same fields as in the ingress port interface.

[0185] There is an extra 1-bit signal from the egress port to egress function, 'advance' that is used for flow control between the port and the egress function in case the egress port can not accept data. If this is the case, the port de-asserts 'advance'. Whenever 'advance' is asserted, the egress function is allowed to send valid data to the port. If de-asserted, the egress function will not send any valid data, even though there might be valid data ready to be sent. If the egress port de-asserts 'advance' in cycle x, it still may receive valid packet data in cycle x+1 since the 'advance' signal is assumed to be registered at the port side.

[0186] The egress function could send valid packet data at a peak rate of 8 bytes per cycle, which translates approximately to 19.2 Gbps (@ 300MHz core frequency). Thus, a mechanism is needed for a port to provision for flow control.

No configuration registers exist in this subblock.

Performance events 12-23 are monitored by the egress function.

PacketLoader block 1402 (PIF)

Detailed description

[0187] The PIF block performs four top-level functions: packet insertion, packet migration, packet transmission and packet table access. Fig. 16 shows its block diagram.

Packet insertion function

[0188] This function interfaces with the PortInterface (PIF) and PacketBufferController (PBC). The function is pipelined (throughput 1) into 3 stages (0, 1 and 2).

Stage 1:

[0189] Packet data is received from the PLD along with the slot number that the PLD computed. If the packet data is not the start of a new packet (this information is provided by the PLD), the slot number is used to look up a table (named slot_state) that contains, for each entry or slot, whether the packet being received has to be dropped. There are three reasons why the incoming packet has to be dropped, and all of them happened when the first data of the packet arrived at Stage A of the PLD:

- The 'continue' configuration register was 0.
- The total number of entries in the packet table (that holds the packet descriptors) was more than 508.
- The packet buffer memory (that holds the data of the packets) was not able to guarantee the storage of a packet of the maximum allowed size.

[0190] If the packet data is the start of the packet, some logic decides whether to drop the packet or not. If any of the above three conditions holds, the packet data is dropped and the slot entry in the slot_table is marked so that future packet data that arrives for that packet is also dropped.

[0191] This guarantees that the whole packet will be dropped, no matter whether the above conditions hold or not when any of the rest of the data of the packet arrives.

[0192] For the purpose of determining at stage 1 whether the packet table is full or not, the threshold number of entries is 512 (the maximum number of entries) minus the maximum packets that can be received in an interleaved way, which is 4. Therefore, if the number of entries when the first data of the packet arrives is more than 508, the packet will be dropped.

[0193] To determine whether the packet buffer will be able to hold the packet or not, some state is looked up that contains information regarding how full the packet buffer is. Based on this information, the decision to drop the packet due to packet buffer space is performed. To understand how this state is computed, first let us describe how the packet buffer is logically organized by the hardware to store the packets.

[0194] The 256KB of the packet buffer are divided into four chunks (henceforth named sectors) of 64KB. Sector 0 starts at byte 0 and ends at byte 0xFFFF, and sector 3 starts at byte 0x30000 and ends at byte 0x3FFFF. The number of sectors matches the number of maximum packets that at any given time can be in the process of being received.

[0195] As will be seen later on, when the packet first arrives, it is assigned one of the sectors, and the packet will be stored somewhere in that sector. That sector becomes

active until the last data of the packet is received. No other packet will be assigned to that sector if it is active.

[0196] Thus, when a new packet arrives and all the sectors are active, then the packet will not be able to be stored. Another reason why the packet might not be accepted is if the total available space in each of the non-active sectors is smaller than the maximum allowed packet size. This maximum allowed packet size is determined by the 'max_packet_size' configuration register, and it ranges from 1KB to 64KB, in increments of 1KB. When the start of a new packet is received, no information regarding the size of the packet is provided up front (the NET block is protocol agnostic, and no buffering of the full packet occurs to determine its size). Therefore, it has to be assumed that the size of the packet is the maximum size allowed in order to figure out whether there is enough space in the sector or not to store the packet.

[0197] In stage 1, the information of whether each sector is active or not, and whether each sector can accept a maximum size packet or not is available. This information is then used to figure out whether the first data of the packet (and eventually the rest of the data) has to be dropped.

[0198] In stage 1, the logic maintains, for all the packets being received, the total number of bytes that have been received so far. This value is compared with the allowed maximum packet size and, if the packet size can exceed the maximum allowed size when the next valid data of the packet arrives, the packet is forced to finish right away (its

end_of_packet bit is changed to 1 when sent to stage 2) and the rest of the data that eventually will come from that packet will be dropped. Therefore, a packet that exceeds the maximum allowed packet size will be seen by software as a valid packet of a size that varies between the maximum allowed size and the maximum allowed size minus 7 (since up to 8 valid bytes of packet data can arrive every cycle). No additional information is provided to software (no interrupt or error status).

[0199] Some information from PLD is propagated into stage 2: valid, start_of_packet, data, port, channel, slot, error, and the following results from stage 1: revised end_of_packet, current_packet_size. If the packet data is dropped in stage 0, no valid information is propagated into stage 2.

Stage 2:

[0200] In this stage, the state information for each of the four sectors is updated, and the hashing function is applied to the packet data.

[0201] When the first data of a packet arrives at stage 2, a non-active sector (guaranteed by stage 1 to exist) is assigned to the packet. The sector that is less occupied is chosen. This is done to minimize the memory fragmentation that occurs at the packet buffer. This implies that some logic will maintain, for each of the sectors, the total number of 8-byte chunks that the sector holds of packets that are kept in the network block (ie packets that have been received but not yet migrated, packets that are being

processed by the tribes, and packets that have been processed but still not been transmitted or dropped).

[0202] Each of the four sectors is managed as a FIFO. Therefore, a tail and head pointer are maintained for each of them. The incoming packet data will be stored at the position within the sector pointed by the tail pointer.

The head and tail pointers point to double words (8 bytes) since the incoming data is in chunks of 8 bytes.

[0203] The tail pointer for the first data of the packet will become (after converted to byte address and mapped into the global physical space of Porthos) the physical address where the packet starts, and it will be provided to one of the tribes when the packet is first migrated (this will be covered on the migration function).

[0204] The tail pointer of each a sector is incremented every time a new valid packet data has arrived (of a packet assigned to that sector). Note that the tail pointer may wrap around and start at the beginning of the sector. This implies that the packet might physically be stored in a non-consecutively manner (but with at most one discontinuity, at the end of the sector). However, as it will be seen as part of the stage 3 logic, a re-mapping of the address is performed before providing the starting address of the packet to software.

[0205] Whenever valid data of a packet is received, the occupancy for the corresponding sector is incremented by the number of bytes received. Whenever a packet

is removed from the packet buffer (as will be seen when the transmission function is explained) the occupancy is decremented by the amount of bytes that the packet was occupying in the packet buffer.

[0206] In stage 2 the hashing function is applied to the incoming packet data. The hashing function and its configuration is explained above. The hashing function applies to the first 64 bytes of the packet. Therefore, when a chunk of data (containing up to 8 valid bytes) arrives at stage 2, the corresponding configuration bits of the hashing function need to be used in order to compute the partial hashing result.

[0207] The first-level hashing function and all the second-level hashing functions are applied in parallel on the packet data received.

[0208] Both partial hashing results and the configuration bits to apply to the next chunk of valid bytes are kept for each of the four ingress interleaving slots.

In this state, if there is a pending GetRoom command, it is served. The GetRoom command is generated by software by writing into the 'get_room' configuration register, with the offset of the address being the amount of space that software requests. The NET will search for a chunk of consecutive space of that size (rounded to the nearest 8-byte boundary) in the packet buffer. The result of the command will be unsuccessful if:

- there are no available entries in the packet table
- there is no space available in the packet buffer to satisfy the request

[0209] A pending GetRoom command will be served only if there is no valid data in State 2 from ingress and there is no valid data in Stage 1 that corresponds to a start of packet.

[0210] The following information is provided to stage 3: valid, data, port, channel, slot, end of packet, start of packet, size of the packet, the dword address, error, get room result, and the current result of the first level of hashing function.

Stage 3:

[0211] In stage 3, the valid packet data is sent to the PBC in order to be written into the packet buffer, and, in case the valid data corresponds to the end of a packet, a new entry in the packet table is written with the descriptor of the packet.

[0212] If the packet data is valid, the 64-bit data is sent to the PBC using the double word address (that points to a double word within the packet buffer). All the 8 bytes will be written (even if less than 8 bytes are actually valid). The PBC is guaranteed to accept this request for write into the packet buffer. There is no flow control between the PBC and stage 3.

[0213] If the valid data happens to be the last data of a packet, a new entry in the packet table is initialized with the packet descriptor. Stage 1 guaranteed that there would be at least one entry in the packet table.

The packet table entries are managed like a FIFO, and the entry number corresponds to the 9 LSB bits of the sequence number, a 16-bit value that is assigned by stage 3 to each

packet. Thus, it is not possible that two packets exist with a sequence number having the 9 LSB bits the same.

[0214] The packet descriptor is composed of the following information:

- Dword address (16): the “expanded” dword address within the packet buffer where the first 8 bytes of the packet reside. The expanded dword address consists on performing the following manipulation of the original dword address computed in stage 2:
 - Bit[15] becomes bit[14]
 - Bit[14] becomes bit[13]
 - Bit[13] becomes 0

[0215] This expanded dword address is compressed back following the inverse procedure when the packet is transmitted out (as will be explained in the transmission function).

- Tribe (3): the tribe number to which the packet will be first migrated into. This value is derived from the second level hashing result generated in stage A and after applying in stage 3 some of the configuration bits of the hashing function.
- FlowId (16): the result of the first level of the hashing function, computed in stage 2.
- Sequence number (16): the value that is assigned by stage 3 to each packet at the end of the packet, ie when the packet has fully been received. After a sequence number has been provided, the register that contains the current sequence number is incremented. The sequence number wraps around at 0xFFFF.
- Inbound port (2): the port number associated to the incoming packet.
- Inbound channel (8): the channel number associated to the incoming packet.

- Outbound port (5): this field will be eventually written with the software-provided value when the 'done' or 'egress_path_determined' configuration registers are written. At stage 3, this field is initialized to 0.
- Outbound channel (9): this field will be eventually written with the software-provided value when the 'done' or 'egress_determined' configuration registers are written. At stage 3, this field is initialized to 0.
- Status (2): it is initialized with 1 (Active). This status will eventually change to either 0 (Invalid) if software requests the packet to be dropped, or to 2 (Done) if software requests the packet to be transmitted out.
- Size (19): the size in bytes of the packet. The maximum allowed size is the size of a sector, ie 65536 bytes (but software can override the maximum allowed size to a lower value with the 'max_packet_size' configuration register).
- Header growth delta (8): initialized with 0. Eventually this field will contain the amount of bytes that the head of the packet has grown or shrunk, and it will be provided by software when the packet is requested to be transmitted out.
- Scheduled (1): specifies whether the egress path information is known for this packet. At stage 3, this bit is initialized to 0 (ie not scheduled).
- Launch (1): bit that indicates whether the packet will be presented to one of the tribes for processing. At stage 3, this bit is initialized to 1 (ie the packet will be provided to one of the tribes for processing).
- Error (1): bit that indicates that the packet arrived with an error notification from the ingress port.

[0216] The following are the valid combination of the 'launch' and 'error' bits in the packet descriptor:

- launch = 1, error = 0. The normal case in which an error-free packet arrives and the NET block will eventually migrate into a tribe.
- launch = 0, error = 0. A packet descriptor originated through a GetRoom command (explained later on). The packet associated to the descriptor will not be migrated.

- launch = 0, error = 1. A packet arrived with an error notification. The packet is allowed to occupy space in the packet buffer and packet table for simplicity reasons (since the error can come in the middle of the packet, it is easier to let the packet reside in the already allocated packet buffer than recovering that space; besides, errors are rare, so the wasted space should have a minimal impact). The packet descriptor is marked with an Invalid status, and therefore the space that it occupies in the packet buffer will be eventually reclaimed when the packet descriptor becomes the oldest one controlled by one of the egress interleaving slots.
- launch = 1, error = 1. Will never occur.

[0217] The descriptor will be read at least twice: once by the migration function to get some of the information and provide it to the initial tribe, and by the transmit logic, to figure out whether the packet needs to be transmitted out or dropped.

[0218] And the descriptor will be (partially) written at least once: when software decides what to do with the packet (transmit it out or drop it). The path and channel information (for the egress path) might be written twice if software decides to notify this information to the NET block before it notifies the completion of the packet.

Configuration register interface

[0219] The following configuration registers are read and/or written by the packet insertion function:

- 'max_packet_size': to cap the maximum packet size in order to minimize the memory fragmentation in the packet buffer.
- 'continue': if 0, the new incoming packets will be dropped.

- ‘packet_table_packets’: the total number of packets that the packet table keeps track of.
- ‘status’: specifies whether the network block is in reset mode and whether it is in quiescent mode.
- Hashing engine configuration registers
 - First level (l1_selection, l1_position, l1_skip)
 - Second level (l2_selection[0..3], l2_position[0..3], l2_skip[0..3], l2_first[0..3], l2_total[0..3])

Performance events numbers 32-36 are monitored by the packet insertion function.

Packet migration function

[0220] The purpose of this function is to monitor the oldest packet in the packet table that still has not been migrated into one of the tribes and perform the migration. The migration protocol is illustrated in the table of Fig. 13.

[0221] This function keeps a counter with the number of packets that have been inserted into the packet table but still have not been migrated. If this number is greater than 0, the state machine that implements this function will request to read the oldest packet (pointer by the ‘oldest to process’ pointer). When the requested information is provided by the packet table access function (explained later on) the packet migration function requests the interconnect block to migrate a packet into a particular tribe (the tribe number was stored into the packet table by the packet insert function). When the interconnect accepts the migration, the packet migration function will send, in 3

consecutive cycles, information of the packet that one of the streams of the selected tribe will store in some general purpose and some CP0 registers.

[0222] The following information is provided in each of the 3 cycles in which data is transferred from the packet migration function to the interconnect block (all the information is available from the information stored in the packet table by the packet insertion function):

- First cycle:
 - PC (32): address where the stream of the tribe will start executing instructions
 - FlowId (16): the result of the first level of hashing
- Second cycle:
 - Sequence number (16)
- Third cycle:
 - Address (32): physical address where the first packet of the packet resides.
 - Ingress port (2): the ingress port of the packet.
 - Ingress channel (8): the ingress channel of the packet.

[0223] Note that the same amount of information could be sent in only two cycles, but the single write port of the register file of the stream along with the mapping of this information into the different GPR and CP0 registers, requires a total of 3 cycles.

[0224] The migration interface with the interconnect block is pipelined in such a way that, as long as the interconnect always accepts the migration, every cycle the packet migration function will provide data. This implies that a migration takes a total of 3 cycles.

[0225] To maintain the 3-cycle throughput, there is a state machine that always tries to read the oldest packet to be migrated and put it into a 4-entry FIFO. Another state machine will consume the entries in this FIFO and perform the 3-cycle data transfer and complying with the Interconnect protocol. The FIFO is needed to squash the latency in accessing the packet table. As it will be seen later on when describing the packet table access function, requests performed by the packet migration function to the packet table might not be served right away. Figure 13 shows a timing diagram of the interface between the packet migration function and the Interconnect module. The 'last' signal is asserted by the packet migration function when sending the information in the second data cycle. If the Interconnect does not grant the request, the packet migration function will keep requesting the migration until granted.

[0226] The migration protocol suffers from the following performance drawback: if the migration request is to a tribe x that can not accept the migration, the migration function will keep requesting for this migration, even if the following migration is available for request to a different tribe that would accept it. With a different, more complex interface, migrations could occur in a different order other than the order of arrival of packets into the packet table, improving the overall performance of the processor.

[0227] The following configuration registers are read and/or written by the packet migration function:

- 'program_counter[0..7]': the initial PC from where the stream that will be associated to the packet will start fetching instructions.

Packet transmission function

[0228] The purpose of this function is to monitor the oldest packet that the packet table keeps track and decide what to do based on its status (drop it or transmit it). This is performed for each of the four egress interleaving slots.

[0229] There is an independent state machine associated to each of the four egress interleaving slots. Each state machine has a pointer to the oldest packet it keeps track of. When appropriate, each state machine requests to a logic to read the entry pointed by its pointer. The logic will schedule the different requests in a round-robin fashion and whenever the packet table access function allows it.

[0230] Whenever software requests to transmit or drop the oldest packet in the packet table a bit (name oldest_touched) is set. Whenever the state machine reads the entry pointed by its pointer, it resets the bit (logic exists to prevent both the set and reset at the same time).

[0231] The state machine will read the entry pointed by its pointer whenever the total number of packet in the table is greater than 1 and

1. 'oldest_touched' is 1, or
2. the previous packet read was dropped or transmitted out ('oldest_processed' = 1).

[0232] This algorithm prevents the state machine to continuously reading the entry of the packet table with the information of the oldest packet, thus saving power.

[0233] The result of the reading of the packet table is presented to all of the state machines, so each state machine needs to figure out whether the provided result is the requested one (by comparing the entry number of the request and result). Moreover, in the case that the entry was indeed requested by the state machine, it might occur that the packet descriptor is not controlled by it since each state machine controls a specific egress port, and for channelized ports, a specific range of channels. In the case that the requested entry is not controlled by the state machine, it is skipped and the next entry is requested (the pointer is incremented and wrapped around at 512 if needed).

[0234] The port that each state machine controls is fixed given the contents of the 'total_ports' configuration register, as follows:

- total_ports = 1. All state machines control port 0
 - total_ports = 2. State machine 0 and 1 control port 0, and state machines 2 and 3 control port 1.
 - total_ports = 4. There is a 1-to-1 correspondence between state machine and port.
- Any other value of 'total_ports' will render undefined results.

[0235] The range of channels that each state machine controls is provided by the 'islot0_channels', ..., 'islot3_channels' configuration registers.

[0236] The status field of the packet descriptor indicates what to do with the packet: drop (status is invalid), transmit (status is done), scheduled (the egress path information is known) or nothing (status is active).

If the packet descriptor is controlled by the state machine, then:

- if the status field is invalid, the state machine will update the pointer (it will be incremented by 1), and it will decrement the occupancy figure of the sector in which the packet resides by the size of the packet, including the offset for header growth, if any. It will also set the 'oldest_processed' bit and decrement the total number of packets.
- if the status field is completed, the state machine will start requesting the PBC to read the packet memory, and it will perform as many reads as necessary to completely read out the packet. These requests are requested to a logic that receives these requests from all the state machines, and will schedule them to the PBC in a round robin fashion. If this logic can not schedule the request of a particular state machine or if the PBC can not accept the requests, it will let the state machine know, and the state machine will need to hold the generation of the requests until the logic can schedule the requests. The request to the PBC includes the following information:
 - the address of the double word to be read out from the packet buffer
 - the channel number and port number
 - whether the request is for the last data of the packet or not
 - which bytes are valid
- if the packet is not completed, the state machine will take no action and will wait until software resolves the corresponding packet by either writing into the 'done' or 'drop' configuration register.

If the packet descriptor is not controlled by the state machine, then

- if the status field is invalid or completed, the state machine skips the packet, and the next entry is requested.
- if the status field not completed and the 'scheduled' bit is 1, the state machine also skips the packet and reads the next entry.
- if the status field not completed and the 'scheduled' bit is 0, the state machine will take no action and will wait until software resolves the corresponding packet by either writing into the 'done' or 'drop' configuration register, or until software notifies the egress path information by writing into the 'egress_path_determination' configuration register.

[0237] Any request to the packet buffer will go to the PBC sub-block, and eventually the requested data will arrive to the PIF sub-block. Part of the request to the PBC contains the state machine number, or egress interleaving slot number, so that the PIF sub-block can enqueue the data into the proper FIFO.

[0238] When a read request is performed (up to 8 bytes worth of valid data), the occupancy of the corresponding sector is decremented by the number of valid bytes. When all the necessary read requests have been done, the 'oldest_processed' bit is set and the total number of packets is decremented.

The 'oldest_processed' bit is reset when a new packet table entry is read.

[0239] The following configuration registers are read and/or written by the packet migration function:

- 'default_egress_channel': this is the egress channel in case the encoded egress channel in the packet descriptor is 0x1.

- 'to_transmit_ptr': the pointer to the oldest packet descriptor in the packet table
- 'head_growth_space': the amount of space reserved for each packet so that its head can grow. This information is needed by the packet transmission function to correctly update the occupancy figure when a packet is dropped or transmitted out.

[0240] There are no performance events associated to packet transmission function.

Packet table access function

[0241] The purpose of this function is to schedule the different accesses to the packet table. The access can come from the packet insertion function, the packet migration function, the packet transmission function, and from software (through the PBM interface).

[0242] This function owns the packet table itself. It has 512 entries; therefore, the maximum number of packets that can be kept in the network block is 512. See the packet insertion function for the fields in each of the entries. The table is single ported (every cycle only a read or a write can occur). Since there are several sources that can request accesses simultaneously, a scheduler will arbitrate and select one request at a time.

[0243] The scheduler has a fixed-priority scheme implemented, providing the highest priority to the packet insertions from the packet insertion tribe. Second highest priority are the requests from software through the PBM interface, followed by the requests from the packet migration function and finally the requests from the packet transmit function.

The access to the packet table takes one cycle, and the result is routed back to the source of the request.

[0244] The requests from software to the packet table can be divided into two types:

- Direct accesses. The packet table is part of the address space; software can perform reads and writes to it.
- Indirect accesses. Whenever software writes into the 'drop' or 'done' configuration registers, the hardware generates a write access to appropriate packet table entry with the necessary information to update the status of the packet.

[0245] All the reads/writes performed by software to the configuration registers of the PLD block are handled by the packet table access function. The only configuration registers not listed above are:

- 'done': software writes in this register to notify that the processing of the packet is completed. The sequence number, egress channel and head growth delta are provided.
- 'drop': software writes in this register to notify that a packet has to be dropped. The sequence number is provided.

[0246] Performance events numbers 37-43 are monitored by the packet table access function..

PacketBufferController block 1403 (PBC)

[0247] The PBC block performs two top-level functions: requests enqueueing and requests scheduling. The requests enqueueing function buffers the requests to the packet

buffer, and the requests scheduling performs the scheduling of the oldest request of each source into the 8 banks of the packet memory. Figure 17 shows its block diagram.

Requests enqueueing function

[0248] The purpose of this function is to receive the requests from all the different sources and put them in the respective FIFOs. There are a total of 10 sources (8 tribes, packet stores from the ingress path, packet reads from the egress path) [and DMA and tribe-like requests from the GLB block - TBD]. Only one request per cycle is allowed from each of the sources.

[0249] With the exception of the requests from the ingress path (named 'network in') all the requests from the other sources are enqueued into corresponding FIFOs. The request from the ingress path is stored in a register because the scheduling function (described later) will always provide priority to these requests and, therefore, they are guaranteed to be served right away.

[0250] All the FIFO's have 2 entries each, and whenever they get 1 or 2 entries with valid requests, a signal is sent to the corresponding source for flow control purposes.

[0251] For the requests coming from the tribes [and the GLB DMA and tribe-like requests - TBD] block, the requests enqueueing function performs a transformation of the address as follows:

- If the address falls into the configuration register space containing the configuration registers and the packet table), the upper 18 bits of the address are zero'ed out (only the 14 LSB bits are kept, which correspond the configuration register number). The upper 1024 configuration registers correspond to the 512 entries in the packet table (2 consecutive configuration registers compose one entry).
- If the address falls into the packet buffer space, the address is modified as follows:
 - Bit 16 becomes Bit 17
 - Bit 17 becomes Bit 18
 - Bit 19 is reset.

[0252] This is done to convert the 512KB logical space of the packet buffer that software sees to the physical 256KB space. Also, a bit is generated into the FIFO that specifies whether the access is to the packet buffer or the configuration register space.

This function does not affect or is affected by any configuration register.

Performance events numbers 64-86 and 58 are monitored by the packet insertion function.

Requests scheduling function

[0253] This function looks at the oldest request in the FIFOs and schedules them into the 8 banks of the packet memory. The goal is to schedule as many requests (up to 8, one per bank). It will also schedule up to one request per cycle that access the configuration register space.

[0254] The packet buffer memory is organized in 8 interleaved banks. Each bank is then 64KB in size and its width is 64 bits. The scheduler will compute into which bank

the different candidate requests (the oldest requests in each of the FIFOs and the network in register) will access. Then, it will schedule one request to each bank, if possible.

[0255] The scheduler has a fixed-priority scheme implemented as follows (in order of priority):

- Ingress requests
- Egress requests
- Global requests - TBD
- Tribe requests. The tribe requests are treated fairly among themselves. Even banks will pick the access of the tribe with the lowest index, whereas odd banks will pick the access of the tribe with the highest index. Since the accesses of a tribe are expected to be usually sequential, consecutive accesses will visit consecutive banks, thus providing a balanced priority to each tribe.

[0256] Whenever a tribe or GBL request accesses the configuration register space, no other configuration space access will be scheduled from any of the tribes of GBL until the previous access has been performed.

This function does not get affected nor affects any configuration register.

Performance events numbers 32-39, 48-55, 57 and 59 are monitored by the packet insertion function

PacketBufferMemory block 1404 (PBM)

[0257] The purpose of this block is to perform the request to the packet buffer memory or the configuration register space. When the result of the access is ready, it will

route the result (if needed) to the corresponding source of the request. The different functions in this block are the configuration register function and the result routing function. Fig. 18 shows its block diagram.

[0258] The packet buffer is part of this block. The packet buffer is 256KB in size and it is physically organized in 8 interleaved banks, each bank having one 64-bit port. Therefore, the peak bandwidth of the memory is 64 bytes per cycle, or 2.4G bytes/sec.

Configuration register function

[0259] The PBC scheduled up to 1 request to the configuration register space. This function serves this request. If the configuration register number falls into the configuration registers that this function controls ('perf_counter_event[0..7]' and 'perf_counter_value[0..7]'), this function executes the request; otherwise, the request is broadcast to both the PIF and PLD blocks. One of them will execute the request, whoever controls the corresponding configuration register.

[0260] This function keeps track of the events that the PIF, PLD and PBC blocks report, and keeps a counter for up to 8 of those events (software can configure which ones to keep track).

Result routing function

[0261] The result routing function has the goal of receiving the result of both the packet memory access and the configuration register space access and rout it to the source of the request.

[0262] To do that, this function stored some control information of the request, which is later on used to decide where the result should go. The possible destinations of the result of the request are the same sources of requests to the PBC with the exception of the egress path (network out requests) that do not need confirmation of the writes.

[0263] The results come from the packet buffer memory and the configuration register function.

[0264] No performance events nor configuration registers are associated to this function.

Interconnect Block of the Porthos Chip

Overview

[0265] The migration interconnect block of the Porthos chip (see Fig. 1, element 109) arbitrates stream migration requests and directs migration traffic between the network block and the 8 tribes. It also resolves deadlocks and handles events such as interrupts and reset.

Interfaces

[0266] Interface names follow the convention SD_name, where S is source block code and D is destination block code. The block codes are:

T: Tribe

I: Migration interconnect

G: Controller

[0267] Fig. 19 is a table providing Interface to tribe # (ranging from 0 to 7), giving name and description.

[0268] Fig. 20 is a table providing Interface to Network block, with name and description.

[0269] Fig. 21 is a table providing interface to global block, with name and description.

Tribe full codes are:

TRIBE_FULL	3
TRIBE_NEARLY_FULL	2
TRIBE_HALF_FULL	1
TRIBE_EMPTY	0

Migration Protocol Timing

[0270] A requester sends out requests to the interconnect, which replies with grant signals. If a request is granted, the requester sends 64-bit chunks of data, and finalizes the

transaction with a finish signal. The first set of data must be sent one cycle after "grant." The signal "last" is sent one cycle before the last chunk of data, and a new request can be made in the same cycle. This allows the new data transfer to happen right after the last data has been transferred.

[0271] Arbitration is ongoing whenever the destination tribe is free. Arbitration for a destination tribe is halted when "grant" is asserted and can restart one cycle after "last" is asserted for network-tribe / interrupt-tribe migration, or the same cycle as "last" for tribe-tribe migration.

[0272] There is a race condition between the "last" signal and the "full" signal. The "last" signal can be sent as soon as one cycle after "grant" while the earliest "full" arrives 4 cycles after "grant" from tribe. To avoid this race condition and prevent overflow, the "almost full to full" is used for 3 cycles after a grant for a destination tribe.

[0273] The Network-Tribe / Tribe-tribe migration protocol timing is shown in Fig. 22.

Interconnect modules

[0274] Fig. 40 illustrates interconnect modules. The interconnect block consists of 3 modules. An Event module collects event information and activate a new stream to

process the event. An Arbiter module performs arbitration between sources and destinations. A Crossbar module directs data from sources to destinations.

Arbiter

Arbitration problem

[0275] There are 11 sources of requests, the 8 tribes, the network block, the event handling module and transient buffers. Each source tribe can make up to 7 requests, one for each destination tribe. The network block, event handling module, and transient buffers each can make one request to one of the 8 tribes.

[0276] If there's a request from transient buffers to a tribe, that request has the highest priority and no arbitration is necessary for that tribe. If transient buffers are not making request, then arbitration is necessary.

[0277] Fig. 41 illustrates a matching matrix for the arbitration problem. Each point is a possible match, with 1 representing a request, and X meaning illegal match (a tribe talking to itself). If a source is busy, the entire row is unavailable for consideration in prioritization and appear as zeroes in the matching matrix. Likewise, an entire column is zeroed out if the corresponding destination is busy.

[0278] The arbiter needs to match the requester to the destination in such a way as to maximize utilization of the interconnect, while also preventing starvation.

[0279] A round-robin prioritizing scheme is used in an embodiment. There are two stages. The first stage selects one non-busy source for a given non-busy destination. The second stage resolves cases where the same source was selected for multiple destinations.

[0280] At the end of the first stage, a crossbar mux selects can be calculated by encoding the destination columns. At the end of the second stage, the "grant" signals can be calculated by OR-ing the entire destination column.

[0281] Each source and each destination has a round-robin pointer. This points to the source or destination with the highest priority. The round-robin prioritization logic begin searching for the first available source or destination beginning at the pointer and moving in one direction.

[0282] Fig. 42 illustrates arbiter stages. The arbitration scheme described above is "greedy," meaning it attempts to pick the requests that can proceed, skipping over sources and destinations that are busy. In other words, when a connection is set up between a source and a destination, the source and destination are locked out from later arbitration. With this scheme, there are cases when the arbiter starves certain context. It could happen that two repeated requests, with overlapping transaction times, can prevent other requests from being processed. To prevent this, the arbitration operates in two modes. The first mode is "greedy" mode as described above. For each request that cannot proceed, there is a counter that keeps track of the number of times that request has been

skipped. When the counter reaches a threshold, the arbitration will not skip over this request, but rather wait at the request until the source and destination become available. If multiple requests reach this priority for the same source or destination, then one-by-one will be allowed to proceed in a strict round-robin fashion. The threshold can be set via the Greedy Threshold configuration register.

Utilization

[0283] Utilization of the interconnect depends on the nature of migration. If only one source is requesting all destinations (say tribe0 wants tribe1-7) or if all sources are requesting one destination, then the maximum utilization is 12.5% (1 out of 8 possible simultaneous connections). If the flow of migration is unidirectional, (say network to tribe0, tribe0 to tribe1, etc.), then the maximum utilization is 100%.

Deadlock Resolution

[0284] Fig. 43 illustrates deadlock resolution. Deadlock occurs when the tribes in migration loops are all full, i.e. tribe 1 requests migration to tribe 2 and vice versa and both tribes are full. The loops can have up to 8 tribes.

[0285] To break a deadlock, Porthos uses two transient buffers in the interconnect, with each buffer capable of storing an entire migration (66 bits times maximum migration cycles). The migration request with both source and destination full (with destination wanting to migrate out) can be sent to a transient buffer. The transient stream becomes

highest priority and initiate a migration to the destination, while at the same time the destination redirect a migration to the second transient buffer. Both of these transfers need to be atomic, meaning no other transfer is allowed to the destination tribe and the tribe is not allowed to spawn new stream within itself. This process is indicated to the target tribe by the signal IT_transient_swap_valid_#. The migrations into and out of a transient buffers use the same protocol as tribe-tribe migrations.

[0286] This method begins by detecting only possibility of deadlock and not the actual deadlock condition. It allows forward progress while looking for the actual deadlock, although there maybe cases where no deadlock is found. It also substantially reduces the hardware complexity with minimal impact on performance.

[0287] A migration that uses the transient buffers will incur an average of 2 migration delays (a migration delay is the number of cycles needed to complete a migrate). The delays don't impact performance significantly since the migration is already waiting for the destination to free up.

[0288] Using transient buffers will suffice in all deadlock situations involving migration:

- Simple deadlock loops involving 2 to 8 tribes
- Multiple deadlock loops with 1 or more shared tribes
- Multiple deadlock loops with no shared tribe
- Multiple deadlock loops that are connected

[0289] In the case of multiple loops, the transient buffers will break one loop at a time. The loop is broken when the transient buffers are emptied.

[0290] Hardware deadlock resolution cannot solve the deadlock situation that involve software dependency. For example, a tribe in one deadlock loop waits for some result from a tribe in the another deadlock loop that has no tribe in common with the first loop. Transient buffers will service the first deadlock loop and can never break that loop.

Event module

[0291] Upon hardware reset, an event module spawns a new stream in tribe 0 to process reset event. This reset event comes from global block. The reset vector is PC = 0xBFC00000.

[0292] Event module spawns a new stream via the interconnect logic based on external and timer interrupts. The default interrupt vector is 0x80000180.

[0293] Each interrupt is maskable by writing to Interrupt Mask configuration registers in configuration space. There are two methods an interrupt can be directed. In the first method, the interrupt is directed to any tribe that is not empty. This is accomplished by the event module making requests to all 8 destination tribes. When there is a grant to one tribe, the event module stops making requests to the other tribes and start a migration for the interrupt handling stream. In the second method, the interrupt is being directed to a

particular tribe. The tribe number for the second method as well as which method are specified using Interrupt Method configuration registers for each interrupt.

[0294] The event module has a 32-bit timer which increments every cycle. When this timer matches the Count configuration register, it activates a new stream via the migration interconnect.

[0295] The interrupt vectors default to 0x80000180 and are changeable via Interrupt Vector configuration registers.

[0296] External interrupt occurs when the external interrupt pin is asserted. If no thread is available to accept the interrupt, the interrupt is pending until a thread becomes available.

[0297] In order to reserve some threads for event-based activations, migrations from network to a tribe can be limited. These limits are set via Network Migration Limit configuration registers (there is one per tribe). When the number of threads in a tribe reaches its corresponding limit, new migrations from network to that tribe are halted until the number of threads drops below the limit.

Crossbar module

[0298] Fig. 44 is an illustration of the crossbar module. This is a 10 inputs, 8 outputs crossbar. Each input is comprised of a "valid" bit, 64-bit data, and a "last" bit. Each

output is comprised of the same. For each output, there's a corresponding 4-bit select input which selects one of 10 inputs for that particular output. Also, for each output, there's a 1-bit input which indicates whether the output port is being selected or busy. This "busy" bit is ANDed with the selected "valid" and "last" so that those signals are valid only when the port is busy. The output is registered before being sent to the destination tribe.

Performance Counters

[0299] With performance counters the performance of the interconnect can be determined. An event is selected by writing to the Interconnect Event configuration registers (one per tribe) in configuration space. Global holds the selection via the selection bus, and the tribe memory interface returns to global the selected event count every cycle via the event bus. The events are:

- Total number of requests and total number of grants in a period of time
- Number of requests and number of grants for each destination in a period of time
- Average time from request to grant overall
- Average time from request to grant for each destination
- Average time from request to grant for each source
- Average migration time overall
- Average migration time per destination
- Average migration time per source

Configuration Registers

[0300] The configuration registers for interconnect and their memory locations are:

Interrupt Masks	0x70000800
Interrupt Pending	0x70000804
Timer	0x70000808
Count	0x70000810
Timer Interrupt Vector	0x70000818
External Interrupt Vector	0x70000820
Greedy Threshold	0x70000828
Network Migration Limit	0x70000830

Memory Interface Block Porthos Chip

Overview

[0301] This section describes the microarchitecture of the memory interface block, which connects the memory controller to the tribe and the global block. Fig. 45 illustrates the tribe to memory interface modules.

Interfaces

[0302] Interface names follow the convention SD_signal_name, where S is source block code and D is destination block code. The block codes are:

T: Tribe

M: Tribe memory interface

L: Tribe memory controller

G: Controller

Fig. 23 is a table illustrating Interface to Tribe.

Fig. 24 is a table illustrating interface to Global.

Fig. 25 is a table illustrating interface to Tribe Memory Controller.

Request types (not command type) are:

MEM_UREAD	0
MEM_SREAD	1
MEM_WRITE	2
MEM_UREAD_RET	4
MEM_SREAD_RET	5
MEM_WRITE_RET	6
MEM_ERROR_RET	7

Memory size codes are:

MEM_SIZE_8	0
MEM_SIZE_16	1
MEM_SIZE_32	2
MEM_SIZE_64	3
MEM_SIZE_128	4
MEM_SIZE_256	5

Interface timings

Tribe to tribe memory interface timing:

[0303] Tribe sends all memory requests to tribe memory interface. The request can be either access to tribe's own memory or to other memory space. If a request accesses tribe's own memory, the request is saved in the tribe memory interface's request queue. Else, it is sent to global block's request queue. Each of these queues have a corresponding full signal, which tells tribe block to stop sending request to the corresponding memory space.

[0304] A request is valid if the valid bit is set and must be accepted by the tribe memory interface block. Due to the one cycle delay of the full signal, the full signal must be asserted when there's one entry left in the queue.

Figure 26 illustrates tribe to tribe memory interface timing.

Tribe memory interface to controller timing:

[0305] This interface is different from other interfaces in that if memory controller queue full is asserted, the memory request is held until the full signal is de-asserted.

Fig. 27 illustrates tribe memory interface to controller timing

Tribe memory interface to global timing:

[0306] Tribe memory interface can send request or return over the transaction bus (MG_transaction*). Global can send request or return over the GM_transaction set of signals.

Fig. 28 illustrates tribe memory interface to global timing

Tribe memory interface block modules

Input module:

[0307] This module accepts requests from tribe and global. If the request from tribe has address that falls within the range of tribe memory space, the request is valid and can be sent to request queue module. If the request has address that falls outside that range, the request is directed to the global block. The tribe number is tagged to the request that goes to global block.

[0308] Global block only send valid request to a tribe if the request address falls within the range of that tribe's memory space.

[0309] The input module selects one valid request to send to request buffer, which has only one input port. The selection is as follows:

- Pick saved tribe request if there are 2 saved requests and an incoming tribe request, or if there's only saved tribe request
- Else pick saved global request if there's only saved global request
- Else pick incoming tribe request if it exists
- Else pick incoming global request if it exists

The module sends flow control signals to tribe and global:

- Stall tribe requests if there's incoming global request
- Stall global requests if there's incoming tribe request

[0310] Save global input if input is not selected during mux selection and saved input slot is free. Else keep old saved input. Similarly for tribe input.

Fig. 29 illustrates input module stall signals

Fig. 46 illustrates the input module data path.

Request buffer and issue module:

[0311] Fig. 49 is an illustration of a request Buffer and Issue Module. There are 16 entries in the queue. When there is a new request, the address and size of the request is compared to all the addresses and sizes in the request buffer. Any dependency is saved in the dependency matrix, with each bit representing a dependency between two entries. When an entry is issued to memory controller, the corresponding dependency bits are cleared in the other entries.

[0312] The different dependencies are:

- Write-after-write: the second write overwrites data written by the first write, so the second write is not allowed to be processed before the first write.
- Write-after-read: The read should not be affected by the write. Thus, the write is not be allowed to be processed before the read.

- Read-after-write: If there are not enough bytes in the writebuffer to forward to the read, then there's no forwarding and the read is not allowed to be processed before the write.

[0313] For read-after-read, if there's no write to the same address between the reads, the reads can be reordered.

Each entry has:

- 1-bit valid
- 27-bit byte address
- 3-bit size code of read data requested
- 5-bit stream number
- 3-bit tribe number
- 5-bit register destination (read)
- 64-bit data (write)
- 16-bit dependency vector

[0314] This module also reorders and issues the requests to the memory controller. The reordering is necessary so that the memory bus is better utilized. The reordering algorithm is as follows:

- If an entry is dependent on another entry, it is not considered for issue until the dependency is cleared.
- Find all entries with address in a bank different from any issued request in the previous n cycles, where n is the striping distance (i.e. 8 for RLD RAM at 300MHz).
- Separate the eligible entries into reads and writes

- Try to issue up to x number of the same type (read or write) before switching to another type. If the other type is not available, continue issuing the same type.
- Save the bank number of the issued request in history table.

[0315] A count register keeps track of the number of valid entries. If the number reaches a watermark level, both the MT_int_request_queue_full and MG_req_transaction_full signals are asserted.

Write buffer:

Fig. 47 Illustrates a write buffer module.

[0316] The write buffer stores 16 latest writes. If subsequent read is to one of these addresses, then the data stored in the write buffer can be forwarded to the read.

Each entry has:

- 8-bit valid bits (one for each byte of data)
- 32-bit address
- 64-bit data
- 4-bit LRU

[0317] When there is a new read, the address of the read is compared to all the addresses in the write buffer. If there is a match, the data from the buffer can be forwarded.

[0318] When there is a new write, the address of the write is compared to all the addresses in the write buffer. If there is a match, the write data replaces the write buffer data. If there's no match, the write data replaces one of the writebuffer entry. The replacement entry is picked based on LRU bits, described below.

[0319] To prevent frequent turning over of writebuffer entries, only writes from local tribe are allowed to replace an entry. Writes from other tribes are only used to overwrite an entry with the same address.

[0320] LRU field indicates how recent the entry has been accessed. The higher the number, the less recently used. A new entry has LRU value of zero. Everytime there is an access to the same entry (forward from the entry or overwrite of entry), the value is reset to zero while the other entries' LRU are increased by one. When a LRU value reaches maximum, it is unchanged until the entry is itself being accessed.

[0321] The replacement entry is picked from the entries with the higher LRUs before being picked from entries with lower LRUs.

Return module:

[0322] There are 3 possible sources for returns to tribe: tribe memory, global, and forwarding. Returns from tribe memory bound for tribe are put into an 8 entry queue. Memory tag information arrives first from the request queue. If it's a write return, it can

be returned to tribe immediately. If it's a read, it must wait for the read data returning from the tribe memory.

[0323] If global is contending for the return to tribe bus, memory block asserts MG_rsp_transaction_full signal to temporarily stop the response from global so tribe memory returns and/or forwarded returns bound for tribe can proceed.

[0324] There are 2 possible sources for returns to global: tribe memory and forwarding. These must contend with tribe requests for the transaction bus. Returns from tribe memory bound for global are put into another set of 8 entry queue. This queue is the similar to the queue designated for returns to tribe.

[0325] If tribe is contending for the return to global bus, memory block asserts MT_ext_request_queue_full signal to stop the external requests from tribe so tribe memory returns and/or forwarded returns bound for global block can proceed.

[0326] All memory accesses are returned to the original tribe that made the requests. Writes are returned to acknowledge completion of writes. Reads are returned with the read data. Returned information include the information send with the request originally. These are stream, regdest, type, size, offset, and data. Offset is lower 3 bits of the original address. Regdest, offset, size and data are relevant only for reads.

[0327] Stream, regdeset, size and offset are unchanged in all returns. Type is changed to the corresponding return type. If there is ECC uncorrectable error or non-existing memory error, the type MEM_ERROR_RET is returned with read return.

[0328] Read data results are 64-bit aligned, so the tribe needs to perform shifting and sign-extension if needed to get the final results.

[0329] Fig. 48 illustrates the return module data path.

Tribe memory configuration registers

[0330] The memory controllers are configured by writing to configuration registers during initialization. These registers are mapped to configuration space beginning at address 0x70000000. Global must detect the write condition and broadcast to all the tribe memory blocks. It needs to assert the GM_initialize_controller while placing the register address and data to be written on the memory transaction bus. Please see Denali specification for descriptions of controller registers.

Assumptions about memory controller

[0331] Memory controller IP is expected to have the following functionalities:

- ECC is enabled, so read-modify-write is included for unaligned accesses
- 8-entry ingress queue (data and command)
- 1-entry egress queue

Can process up to 256-bit memory requests.

Doesn't include reordering or forwarding features.

Performance counters

[0332] This block generates event counts for performance counters. The event is selected by writing to Tribe MI Event configuration registers (one per tribe) in configuration space. Global holds the selection via the selection bus, and the tribe memory interface returns to global the selected event count every cycle via the event bus.

The events counted are:

- length of request queue
- length of return queue
- write/read issued
- forwarded from write buffer
- global request stall
- global response stall
- tribe request stall

Tribe Block Microarchitecture Porthos Chip

Overview

[0333] A Tribe block 104 (See Fig. 1) contains a multithreaded pipeline that implements the processing of instructions. It fits into the overall Porthos chip microarchitecture as shown in Fig. 33. The Tribe microarchitecture is shown in Fig. 34,

which illustrates the modules that implement the Tribe and the major data path connections between those modules.

[0334] A Tribe contains an instruction cache and register files for 32 threads. The tribe block interfaces with the Network block (for handing packet buffer reads and writes), the Interconnect block (for handling thread migration) and the Memory block (for handling local memory reads and writes).

Interfaces

Fig. 30 shows interface to the Memory block.

Fig. 31 shows interface to the Network block.

Fig. 32 shows interface to the Interconnect block.

Tribe Detailed Description

[0335] The Tribe block consists of three decoupled pipelines. The fetch logic and instruction cache form a fetch unit that will fetch from two threads per cycle according to thread priority among the threads that have a fetch available to be performed. The Stream block within the Tribe contains its own state machine that sequences reads and writes to the register file and executes certain instructions. Finally the scheduler, global ALU and memory modules form an execute unit that schedules operations based on global priority among the set of threads that have an instruction available for scheduling. At most one instruction per cycle is scheduled from any given thread. Globally, up to three instructions can be scheduled in a single cycle, but some instructions can be fully executed within the Stream block, not requiring global scheduling. Thus, the maximum

rate of instruction execution is actually determined by the fetch unit, which can fetch up to eight instructions each cycle. A sustained execution rate of five to six instructions per cycle is expected.

Instruction Fetch

[0336] The instruction fetch mechanism fetches four instructions from two threads for a total fetch bandwidth of eight instructions. The fetch unit includes decoders so that four decoded instructions are delivered to two different stream modules in each cycle. There is a 16K byte instruction cache shared by all threads that is organized as 1024 lines of 16 bytes each, separated into four ways of 256 lines. The fetch mechanism is pipelined, with the tags accessed in the same cycle as the data. The fetch pipeline is illustrated in Fig. 35. In an alternative embodiment, the tag read for all ways is pipelined with the data read for only the way that contains valid data. This increases the overall fetch pipeline by one cycle, but it would significantly reduce the amount of power and the wiring required to support the instruction cache.

Stream Modules

[0337] The Stream modules (one per stream for a total of 32 within the Tribe block) are responsible for sequencing reads and writes to the register files, executing branch instructions, and handling certain other arithmetic and logic operations. A Stream module receives two decoded instructions at a time from the Fetch mechanism and saves them for later processing. One instruction is processed at a time, with some instructions

taking multiple cycles to process. Since there is only a single port to the register file, all reads and writes must be sequenced by the Stream block. The basic pipeline of the Stream module is shown in Fig. 36. Note that in cases where only a single operand needs to be read from the register file, the instruction would be available for global scheduling with only a single RF read stage. Each register contains a ready bit that is used to determine if the most recent version of a register is in the register file, or it will be written by an outstanding memory load or ALU instruction.

[0338] Writes returning from the Network block and the Memory block must also be sequenced to the register file. The register write sequencing pipeline of the Stream block is shown in Fig. 37. When a memory instruction, or an instruction for the global ALU is encountered, the operation matrix, or OM, register is updated to reflect a request to the global scheduling and execute mechanism.

[0339] Branch instructions are executed within the Stream module as illustrated in Fig. 38. Branch operands can come from the register file, or can come from outstanding memory or ALU instructions. The branch operand registers are updated in the same cycle in which the write to the register file is scheduled. This allows the execution of the branch to take place in the following cycle. Since branches are delayed, the instruction after the branch instruction must be processed before the target of the branch can be fetched. The earliest that a branch delay slot instruction can be processed is the same cycle that a branch is executed. Thus, a fetch request can be made at the end of this cycle

at the earliest. The processing of the delay slot instruction would occur later than this if it was not yet available from the Fetch pipeline.

Scheduling and Execute

[0340] The scheduling and execute modules schedule up to three instructions per cycle from three separate streams and handle register writes back to the stream modules. The execute pipeline is shown in Fig. 39. Streams are selected based on what instruction is available for execution (only one instruction per stream is considered a candidate), and on the overall stream priority. Once selected, a stream will not be able to be selected in the following cycle since there is a minimum two cycle feedback to the Stream block for preparing another instruction for execution.

Thread Migration

[0341] The thread migration module is responsible for migrating threads into the Tribe block and out of the Tribe block. A thread can only be participating in migration if it is not actively executing instructions. During migration, a single register read or write per cycle is processed by the Stream module and sent to the Interconnect block. A migration may contain any number of registers. When an inactive stream is migrated in, all registers that are not explicitly initialized are set invalid. An invalid register will always return 0 if read. A single valid bit per register allows the register file to behave as if all registers are initialized to zero when a thread is initialized.

[0342] In an alternative embodiment, thread migration is automatic and under hardware control. Hardware in each of the tribes monitors the frequency of accesses to a remote local memory vs. accesses to its own local memory. If a certain threshold is reached, or based on a predictive algorithm, the thread is automatically migrated by the hardware to another tribe for which a higher percentage of local accesses will occur. In this case migration is transparent to software.

Thread Priority and Flow Gating

[0343] Thread priority (used for fetch scheduling and execute scheduling) are maintained by the "LStream" module. This module also maintains a gateability vector used to implement FlowGating. The LStream module is responsible for determining for each thread whether or not it should stall upon the execution of a "GATE" instruction, or should stall. This single bit per thread is exported to each Stream block. Any time a change is made to any CP0 register that can potentially affect gateability, the LStream module will export all 0's on its gateability vector (indicating no thread can proceed past a GATE), until a new gateability vector is computed.

Changes that affect gateability are rare. They are as follows:

1. A new thread is created, it will be migrated in with its own sequence number, gate vector and flow ID register;
2. An existing thread is deactivated, either due to a DONE instruction or a NEXT instruction (migration out to another tribe);
3. A thread explicitly updates one of its gateability CP0 registers (sequence number, gate vector, flow ID) using the MTC0 instruction.

Debugging and Performance Monitoring

[0344] The Tribe block contains debugging hardware to assist software debugging and performance counters to assist in architecture modeling.

[0345] All of the above description and teaching is specific to a single implementation of the present invention, and it should be clear to the skilled artisan that there are many alterations and amendments that might be made to the example provided, without departing from the spirit and scope of the invention. For example, the aggressively multi-threaded architecture may be accomplished with more or fewer tribes. Many unique and novel features stand alone without the limitation of a tribe architecture at all. Interconnection and communication among the many parts of the Porthos chip may be accomplished in a variety of ways within the spirit and scope of the invention.

[0346] In addition to the above, in some embodiments of the Porthos chip a portion of the packet buffer memory can be configured as "shared" memory for all the tribes. This portion will not be used by the logic that decides where the incoming packet will be stored into. Therefore, this portion of shared memory is available for the tribes for any storage purpose. In addition the ports to the packet buffer can be used for both types of accesses (to packet data and to the shared portion).

In some embodiments software can configure the size of the shared portion of the packet buffer. One implementation of this configuration mechanism allows software to set aside either half, one fourth or none of the packet buffer as shared memory. The

shared memory can be used to store data that is global to all the processing cores, but it can also be divided into the different cores so that each core has its own space, thus no mutually exclusive operation is needed to allocate memory space.

In some embodiments the division of the shared space into the different processing cores and/or threads may provide storage for the stack of each thread. For those threads in which life corresponds to the life of the packet, the header growth offset mechanism may be used to provide storage space for the stack. For those threads that operate on more than a packet, or that need the stack after completing and sending out the processed packet, a persistent space is needed for the stack; for these threads, space in the external memory (long latency) or in the shared portion of the packet buffer (short latency) is required.

[0347] Further to the above, in some embodiments the header growth offset mechanism is intended for software to have some empty space at the beginning of the packet in case the header of the packet has to grow in a few bytes. Note that software may also use this mechanism to guarantee that there is space at the end of a packet A by using the header growth offset space that will be set aside for a future incoming packet B that will be stored after packet A. Even if packet B has still not arrived, software can use the space at the end of packet A since it is guaranteed that either that space has still not been assigned to any packet, or will be assigned to packet B without modifying its content when this occurs. The header growth offset can also be shared among the incoming packet B and the packet stored right above A, as long as the upper space of the growth offset used as tail growth offset of packet A does not overlap with the lower space of the growth offset used as head growth offset of packet B.

There are similarly many other alterations that may be made within the spirit and scope of the invention.

Resource Locking Conceptual Overview

[0348] As disclosed previously, each tribe 104 (Fig. 1) in this embodiment of the packet processing engine 101 comprises 32 streams. Each stream is capable of executing instructions for a single thread. As used herein, the term stream refers broadly to any set of components that cooperate to execute a thread. Figs. 34, 35, and 36 show an example of the components that can make up a stream, in accordance with one embodiment of the present invention. With 32 streams, each tribe 104 is able to support the concurrent execution of 32 threads. With 8 tribes 104, and 32 streams in each tribe, the packet processing engine 101 is able to support 256 concurrently executing threads.

[0349] During execution, these threads may contend for shared resources. Used in this context, the term resource refers broadly to anything that contains information that can be accessed and updated by multiple threads. Examples of resources may be memory/storage locations that contain the values of shared variables, registers that contain information used by multiple threads, etc. Because these resources may be accessed and updated by multiple threads, it is important to ensure that updates to these resources are done atomically. If they are not, then data consistency may be compromised. In one embodiment, atomicity is ensured by way of a lock reservation methodology. This methodology is carried out by one or more lock managers with cooperation from the various streams, in accordance with one embodiment of the present invention.

[0350] In one embodiment, the lock reservation methodology is implemented as follows. Initially, a lock manager receives, from a particular stream executing a particular thread, a request to obtain a lock on a resource. In response to this lock request, the lock manager determines whether the resource is currently locked by another stream. If the resource is currently locked by another stream, then the lock manager adds the particular stream to a wait list of streams waiting to obtain a lock on the resource. By doing so, the lock manager in effect gives the particular stream a reservation to obtain a lock on the resource in the future. Thus, the particular stream does not need to submit another lock request. In one embodiment, in addition to adding the particular stream to the wait list, the lock manager also causes the particular stream to halt execution of the particular thread. That way, the particular stream does not consume any processing resources while it is waiting for a lock on the resource. This is advantageous in the processing engine architecture described above because the various streams share numerous hardware processing resources. If the particular stream does not consume any processing resources while it is waiting for a lock, then those processing resources can be used by other streams that are doing useful work. Thus, causing the particular thread to halt execution of the particular streams enables processing resources to be used more efficiently.

[0351] After it is put on the wait list, the particular stream waits for a lock on the resource. In the meantime, the lock manager allows each stream on the wait list, in turn, to obtain a lock on the resource, access and optionally update the contents of the resource, and release the lock on the resource. At some point, based upon the wait list, the lock manager determines that it is the particular stream's turn to obtain a lock on the resource. At that point, the lock manager grants the particular stream a lock on the resource. In one

embodiment, in addition to granting the lock, the lock manager also causes the particular stream to resume execution of the particular thread. Thereafter, the particular stream can access the resource, perform whatever processing and updates it wishes, and release the lock when it is done. In this manner, a lock reservation methodology is implemented. By implementing locking in this way, streams do not need to repeatedly request a lock on a resource. Instead, they submit only one lock request, and when it is their turn to obtain a lock, they are granted that lock. By eliminating the need to repeatedly request locks, the lock reservation methodology makes it possible to implement resource locking efficiently in a massively multi-threaded environment.

Sample Implementation

[0352] To facilitate a complete understanding of the lock reservation methodology, a sample implementation will now be described. In the following description, it will be assumed, for illustrative purposes, that the resource to be locked is one or more memory/storage locations and that the lock manager coordinates access to a memory.

Overview

[0353] Fig. 50 shows a high level block diagram of the components that are involved to some extent in the lock reservation methodology. These components include the 8 tribes 104. As shown, each tribe 104 comprises 32 streams. Each stream has a stream number or stream ID. The stream ID allows each stream in the processing engine 101 to be uniquely identified. In one embodiment, the stream ID is a number between 0 and 255 (because there are 256 streams in the processing engine). In the example shown, streams

0 through 31 are in tribe 0 104(0), streams 32 through 63 are in tribe 1 104(1), and so forth, with streams 224 through 255 being in tribe 7 104(7).

[0354] Each tribe 104 has an associated local external memory 5004. The streams in a tribe 104 can access the local external memory 5004 associated with that tribe through a lock management block 5002. For example, the streams in tribe 0 104(0) can access local external memory 5004(0) through lock management block 5002(0). Likewise, the streams in tribe 7 104(7) can access local external memory 5004(7) through lock management block 5002(7). The streams in a tribe 104 can also access the local external memories 5004 associated with other tribes; however, to do so, the stream would have to go through the global unit 108. Thus, for example, if the streams in tribe 0 104(0) wanted to access local external memory 5004(7), they would have to go through the global unit 108 and the lock management block 5002(7).

[0355] Each lock management block 5002 controls accesses to its associated local external memory 5004. For example, lock management block 5002(0) controls accesses to local external memory 5004(0), lock management block 5002(1) controls accesses to local external memory 5004(1), and so forth. In one embodiment, it is the lock management blocks 5002 that implement, in large part, the lock reservation methodology.

Lock Management Block

[0356] Fig. 51 shows one of the lock management blocks 5002 in greater detail. For the sake of example, the lock management block shown in Fig. 51 is lock management block 5002(0) that is local to tribe 0 104(0) and memory 5004(0). In one embodiment, all of the lock management blocks 5002 have basically the same structure as that shown in

Fig. 51. As shown, lock management block 5002(0) comprises a lock manager 5102, a plurality of lock management storages 5106, and a wait list storage structure 5104. In one embodiment, it is the lock manager 5102 that implements, in large part, the lock reservation methodology. For purposes of the present invention, the functionality of the lock manager 5104 may be implemented in any desired way. For example, the lock manager 5104 may be implemented using hardware logic components. Its functionality may also be derived by having a processor execute software or firmware instructions. In addition, the functionality of the lock manager 5104 may be realized with a combination of hardware logic components and software-executing processing components. These and other implementations of the lock manager 5104 are within the scope of the present invention.

[0357] As shown in Fig. 51, the lock manager 5102 is coupled to a scheduler 5130. In turn, the scheduler 5130 is coupled to receive requests from various streams. These requests may come from local streams (since the lock management block 5002(0) is local to tribe 0 104(0), the local streams are stream 0 through stream 31). The requests may also come from non-local streams. In such a case, the requests would come through the global unit 108. The scheduler 5130 receives all such requests and selectively passes them to the lock manager 5102 for processing. Depending upon the type of request, the lock manager 5102 may implement the lock reservation methodology on the request.

[0358] In the course of implementing the lock reservation methodology, the lock manager 5102 stores and updates information in the lock management storages 5106 and the wait list storage structure 5104. In one embodiment, the lock management storages 5106 store information pertaining to the locks (if any) that have been imposed on certain

memory locations, and the wait list storage structure 5104 maintains the wait lists (if any) of streams that are waiting on those locks.

[0359] As shown in Fig. 51, each lock management storage 5106 comprises a plurality of portions. A first portion 5110 stores a lock bit. This bit indicates whether the lock management storage 5106 is currently being used to lock a memory location. A second portion 5112 stores an address. This is the address of the memory location that is being locked. A third portion 5114 stores the stream ID of the stream that currently holds the lock on the memory location, and a fourth portion 5116 stores the stream ID of the stream that is currently last on the wait list (if any) to obtain a lock on the memory location. The use of the lock management storages 5106 and the significance of the information stored therein will be elaborated upon in a later section.

[0360] For purposes of the present invention, a lock management block 5002 may comprise any n number of lock management storages 5106, where n is an integer. In one embodiment, each lock management block 5002 comprises 8 lock management storages 5106. With 8 storages 5106, it is possible for the lock manager 5102 to manage locks on 8 different memory locations concurrently.

[0361] Each of the memory locations locked by each of the lock management storages 5106 may have associated therewith a wait list of streams waiting to obtain a lock on that memory location. These wait lists are maintained in the wait list storage structure 5104. In one embodiment, the wait list storage structure 5104 comprises an entry corresponding to each stream in the processing engine 101. As noted previously, this embodiment of the processing engine 101 has 256 streams; thus, the wait list storage structure 5104 has 256 entries. Entry 0 corresponds to stream 0, entry 1 corresponds to

stream 1, and so on. Each entry can store the ID of a stream. Since there are 256 streams, 8 bits are needed to uniquely identify a stream; thus, in one embodiment, each entry is 8 bits in size.

[0362] When a stream ID is stored in an entry of the wait list storage structure 5104, it indicates that that stream follows the stream that corresponds to that entry on a wait list. For example, if the stream ID of stream 4 (S4 for short) is stored in entry 8, it means that stream 4 follows stream 8 on a wait list. Likewise, if the stream ID of stream 120 (S120) is stored in entry 4, then it means that stream 120 follows stream 4 in a wait list. Thus, if a wait list consists of S8, S4, and S120, in that order, then S4 would be stored in entry 8 (the entry that corresponds to stream 8), and S120 would be stored in entry 4 (the entry that corresponds to stream 4). With information stored in this way, it is possible to determine which stream follows a particular stream on a wait list by just consulting the entry corresponding to the particular stream.

[0363] In one embodiment, the wait list storage structure 5104 can be used to maintain all of the wait lists for all of the memory locations locked by the lock management storages 5106. Thus, if lock management storage 5106(0) is locking an address X and lock management storage 5106(n) is locking an address Y, the wait lists (if any) for both of these locks can be maintained within the same wait list storage structure 5104. For example, suppose that the wait list for address X is S8, S4, and S120, and the wait list for address Y is S50, S90, and S0. In this case, S4 would be stored in entry 8, S120 would be stored in entry 4, S90 would be stored in entry 50, and S0 would be stored in entry 90. As this example shows, only one wait list storage structure 5104 is needed to

maintain the wait lists for all of the memory locations locked by all of the lock management storages 5106 in a lock management block 5002.

Stream Cooperation

[0364] As noted previously, it is the lock manager 5102 that implements, in large part, the lock reservation methodology. While this is true, the streams do also participate in the process. In one embodiment, the streams participate in the lock reservation methodology by executing two special instructions: the LL and the SC instructions.

[0365] To elaborate, each stream executes the instructions for a thread. The thread may include the LL and the SC instructions. The LL instruction is executed to request and obtain a lock on a memory location. The SC instruction is executed to update the contents of the memory location and to release the lock. If the software of the thread is written properly, the SC instruction will come after an LL instruction (a stream cannot logically release a lock until it has requested and obtained a lock). Any number of other instructions may come between the LL and the SC instructions. These other instructions would represent the processing that is performed using the information that is read from the memory location that is locked.

[0366] When a stream encounters and executes an LL instruction, it submits a request to the lock manager 5102 to access and read information from a particular address. As part of this request, the stream provides a set of information. This information may include the stream's ID, the memory address to be accessed, and the type of access desired. In this case, the type of access is LL, which the lock manager 5102 recognizes as a locking request. Thus, in response to this request, the lock manager 5102 will

implement the lock reservation methodology. As part of executing the LL instruction, the stream will wait for a response from the lock manager 5102 before executing another instruction for the thread. If the stream receives a wait signal from the lock manager 5102 (thereby indicating that the desired address is currently locked by another stream), then the stream halts execution of the thread. In effect, the stream goes to "sleep" and waits to be awakened. If the stream receives a proceed signal from the lock manager 5102 (thereby indicating that the stream has been granted a lock on the desired address), then the stream continues or resumes (if execution was halted) execution of the thread. In this manner, the stream cooperates with the lock manager 5102 to implement the lock requesting part of the lock reservation methodology.

[0367] When a stream encounters and executes an SC instruction, it submits a request to the lock manager 5102 to write data into a particular address and to release the lock on that address. As part of this request, the stream provides a set of information. This information may include the stream's ID, the memory address to be written into and released, the data that is to be written into the address or a reference to a register in which that data is stored, and the type of access desired. In this case, the type of access is SC, which the lock manager 5102 recognizes as a write access and a lock release. Thus, in response to this request, the lock manager 5102 will write the information into the address, release the lock held by the stream on the address, and allow the next stream on a wait list (if any) to obtain a lock on the address. In this manner, the stream cooperates with the lock manager 5102 to implement the lock releasing part of the lock reservation methodology.

Sample Operation

[0368] To facilitate a complete understanding of the present invention, a sample operation illustrating the lock reservation methodology will now be described with reference to Figs. 50-56.

[0369] Suppose that stream 3 in tribe 0 104(0) (Fig. 50) encounters an LL instruction in the thread that it is executing. Upon executing this instruction, stream 3 determines that it needs to access an address X. For the sake of example, it will be assumed that address X is an address within memory 5004(0); thus, access to this memory location is managed by lock management block 5002(0). As a result of executing the LL instruction, stream 3 submits an access request (through scheduler 5130 in Fig. 51) to lock manager 5102. This request includes stream 3's ID (denoted as S3 for ease of expression), address X, and an indication that this is an LL type of request. After submitting the request, stream 3 waits for a response before executing any other instruction in its thread.

[0370] At some point, the access request is received by lock manager 5102 from the scheduler 5130. From the information in the request, lock manager 5102 knows that this is an LL type of request (which is a locking request); thus, it knows to implement the lock reservation methodology. To do so, lock manager 5102 extracts the address X from the request, and compares it with the addresses stored in portion 5112 of the lock management storages 5106. Lock manager 5102 also looks at the lock bit stored in portion 5110 of each lock management storage 5106. Unless a lock management storage 5106 has both a match for the address X and its lock bit set, the address X is not currently locked by any other stream. In the current example, it will be assumed that address X is not currently locked. Thus, lock manager 5102 can grant a lock on address X to stream 3.

[0371] To do so, lock manager 5102 selects one of the lock management storages 5106 that is not currently being used (for the sake of example, it will be assumed that storage 5106(0) is selected). Lock manager 5102 then sets the lock bit in portion 5110 to 1, and writes the address X into portion 5112. In addition, lock manager 5102 stores the ID (S3) of stream 3 into portion 5114 to indicate that stream 3 currently has the lock on address X. Furthermore, lock manager 5102 stores the ID of stream 3 into portion 5116 to indicate that stream 3 is the last stream on the wait list of streams waiting to obtain a lock on address X. Since there are currently no streams waiting to obtain a lock on address X, stream 3 is technically the last stream on the wait list. Fig. 52 shows the contents of lock management storage 5106(0) after the lock has been granted. Since there is currently no wait list for the lock on address X, no information is currently stored in the wait list storage structure 5104 for this lock.

[0372] After granting the lock, lock manager 5102 accesses address X in memory 5004(0), reads the information therefrom, and provides the information to stream 3. In addition, lock manager 5102 sends a proceed signal to stream 3 to cause stream 3 to proceed with the execution of its thread. Thereafter, stream 3 can execute other instructions in its thread that use the information from address X.

[0373] Suppose now that while stream 3 has a lock on address X, stream 8 of tribe 0 104(0) encounters an LL instruction in the thread that it is executing. Suppose further that execution of this instruction causes stream 8 to determine that it also needs to access address X. As a result of executing this LL instruction, stream 8 submits an access request (through scheduler 5130 in Fig. 51) to lock manager 5102. This request includes stream 8's ID (S8), address X, and an indication that this is an LL type of request. After

submitting the request, stream 8 waits for a response before executing any other instruction in its thread.

[0374] At some point, the access request is received by lock manager 5102 from the scheduler 5130. From the information in the request, lock manager 5102 knows that this is an LL type of request; thus, it knows to implement the lock reservation methodology. To do so, lock manager 5102 extracts the address X from the request, and compares it with the addresses stored in portion 5112 of the lock management storages 5106. Lock manager 5102 also looks at the lock bit stored in portion 5110 of each lock management storage 5106. This time, lock manager 5102 finds that lock management storage 5106(0) has both a match for the address X and its lock bit set; thus, address X is currently locked by another stream. In such a circumstance, lock manager 5102 will add stream 8 to a wait list of streams waiting to obtain a lock on address X.

[0375] To do so, lock manager 5102 reads portion 5116 of lock management storage 5106(0) to determine which stream is currently the last stream in the wait list of streams waiting for a lock on address X. In the present example, the last stream is currently stream 3. Upon learning this, lock manager 5102 accesses entry 3 (the entry that corresponds to stream 3) in the wait list storage structure 5104, and stores in that entry the stream ID (S8) for stream 8. This is shown in Fig. 53. By doing so, lock manager 5102 adds stream 8 (and hence, the thread that stream 8 is executing) to the wait list. In addition, lock manager 5102 updates portion 5116 of lock management storage 5106(0) with the ID (S8) for stream 8 to indicate that stream 8 is now the last stream on the wait list. After that is done, lock manager 5102 sends a wait signal to stream 8. This signal

causes stream 8 to halt execution of its thread, and to wait for the lock on address X to become available.

[0376] Suppose now that while stream 3 still has a lock on address X, stream 60 of tribe 1 104(1) encounters an LL instruction in the thread that it is executing. Suppose further that execution of this instruction causes stream 60 to determine that it also needs to access address X. As a result of executing this LL instruction, stream 60 submits an access request (through global unit 108 and scheduler 5130) to lock manager 5102. This request includes stream 60's ID (S60), address X, and an indication that this is an LL type of request. After submitting the request, stream 60 waits for a response before executing any other instruction in its thread.

[0377] At some point, the access request is received by lock manager 5102 from the scheduler 5130. From the information in the request, lock manager 5102 knows that this is an LL type of request; thus, it knows to implement the lock reservation methodology. To do so, lock manager 5102 extracts the address X from the request, and compares it with the addresses stored in portion 5112 of the lock management storages 5106. Lock manager 5102 also looks at the lock bit stored in portion 5110 of each lock management storage 5106. Again, lock manager 5102 finds that lock management storage 5106(0) has both a match for the address X and its lock bit set; thus, address X is currently locked by another stream. In such a circumstance, lock manager 5102 will add stream 60 to the wait list of streams waiting to obtain a lock on address X.

[0378] To do so, lock manager 5102 reads portion 5116 of lock management storage 5106(0) to determine which stream is currently the last stream in the wait list of streams waiting for a lock on address X. In the present example, the last stream is currently

stream 8. Upon learning this, lock manager 5102 accesses entry 8 (the entry that corresponds to stream 8) in the wait list storage structure 5104, and stores in that entry the stream ID (S60) for stream 60. This is shown in Fig. 54. By doing so, lock manager 5102 adds stream 60 (and hence, the thread that stream 60 is executing) to the wait list. In addition, lock manager 5102 updates portion 5116 of lock management storage 5106(0) with the ID (S60) for stream 60 to indicate that stream 60 is now the last stream on the wait list. After that is done, lock manager 5102 sends a wait signal to stream 60. This signal causes stream 60 to halt execution of its thread, and to wait for the lock on address X to become available.

[0379] Suppose now that stream 3 encounters an SC instruction in the thread that it is executing. Upon executing this instruction, stream 3 determines that it needs to write to and release the lock on address X. As a result of executing the SC instruction, stream 3 submits a write request (through scheduler 5130) to lock manager 5102. This request includes stream 3's ID (S3), address X, the actual write data to be written into address X or a reference to a register that currently contains the write data, and an indication that this is an SC type of request. In one embodiment, after submitting the write request, stream 3 can proceed to execute other instructions in its thread. It does not need to wait for a response from lock manager 5102. Of course, if it is desirable for stream 3 to wait for a response, it may do so.

[0380] At some point, the write request is received by lock manager 5102 from the scheduler 5130. From the information in the request, lock manager 5102 knows that this is an SC type of request; thus, it knows to implement the lock reservation methodology. To do so, lock manager 5102 first verifies that the write request is from a stream that

currently has a lock on the address X. To do so, lock manager 5102 extracts the address X and the stream ID (S3) from the request, and compares the address X with the addresses stored in portion 5112 of the lock management storages 5106, and compares the stream ID with the stream ID's stored in portion 5114 of the lock management storages 5106. In addition, lock manager 5102 also looks at the lock bit stored in portion 5110 of each lock management storage 5106. In the present example, lock manager 5102 finds, from the information stored in lock management storage 5106(0), that there currently is a lock on address X, and that the lock is currently held by stream 3. Thus, the write request is from a proper stream. In response to this write request, lock manager 5102 accesses the address X in memory 5004(0), and stores the write data therein. The memory location at address X is thus updated with updated content.

[0381] Lock manager 5102 recognizes that the write request (because it is of type SC) is also a request by stream 3 to release its lock on address X. Thus, in response to this write request, lock manager 5102 checks to see which stream (if any) is the next stream to be granted the lock on address X. To do so, lock manager 5102 compares the stream ID stored in portion 5114 of lock management storage 5106(0) with the stream ID stored in portion 5116. If these are the same (thereby indicating that the stream with the current lock is the last stream on the wait list), then there are no streams waiting for a lock on the address X. In the current example, however, these ID's are not the same. Thus, lock manager 5102 knows that at least one stream is waiting for a lock on address X.

[0382] To determine which stream is the next stream to be granted a lock on address X, lock manager 5102 reads portion 5114 of lock management storage 5106(0) to ascertain which stream currently has the lock on address X. In the current example, that

stream is stream 3. Thereafter, lock manager 5102 accesses entry 3 (the entry that corresponds to stream 3) in the wait list storage structure 5104 and obtains therefrom a thread ID. In the current example, the ID (S8) of stream 8 is stored in entry 3; thus, lock manager 5102 determines that stream 8 is the next stream to be granted a lock on address X. To grant the lock to stream 8, lock manager 5102 writes stream 8's ID (S8) into portion 5114 of lock management storage 5106(0) to indicate that stream 8 now has the lock on address X. This is shown in Fig. 55. After that is done, lock manager 5102 accesses the memory location at address X of memory 5004(0), and reads the current contents therefrom. Lock manager 5102 then provides the current contents of address X to stream 8, and sends a proceed signal to stream 8. This signal causes stream 8 to resume execution of its thread. In this manner, lock manager 5102 grants the lock to stream 8 and wakes it up.

[0383] Suppose now that stream 8, at some point, encounters an SC instruction in the thread that it is executing. Upon executing this instruction, stream 8 determines that it needs to write to and release the lock on address X. As a result of executing the SC instruction, stream 8 submits a write request (through scheduler 5130) to lock manager 5102. This request includes stream 8's ID (S8), address X, the actual write data to be written into address X or a reference to a register that currently contains the write data, and an indication that this is an SC type of request. In one embodiment, after submitting the write request, stream 8 proceeds to execute other instructions in its thread.

[0384] At some point, the write request is received by lock manager 5102 from the scheduler 5130. From the information in the request, lock manager 5102 knows that this is an SC type of request; thus, it knows to implement the lock reservation methodology.

To do so, lock manager 5102 first verifies that the write request is from a stream that currently has a lock on the address X. In the present example, lock manager 5102 finds, from the information stored in lock management storage 5106(0), that there currently is a lock on address X, and that the lock is currently held by stream 8. Thus, the write request is from a proper stream. In response to this write request, lock manager 5102 accesses the address X in memory 5004(0), and stores the write data therein. The memory location at address X is thus again updated with updated content.

[0385] Lock manager 5102 recognizes that the write request (because it is of type SC) is also a request by stream 8 to release its lock on address X. Thus, in response to this write request, lock manager 5102 checks to see which stream (if any) is the next stream to be granted the lock on address X. To do so, lock manager 5102 compares the stream ID stored in portion 5114 of lock management storage 5106(0) with the stream ID stored in portion 5116. If these are the same, then there are no streams waiting for a lock on the address X. In the current example, however, these ID's are not the same. Thus, lock manager 5102 knows that at least one stream is waiting for a lock on address X.

[0386] To determine which stream is the next stream to be granted a lock on address X, lock manager 5102 reads portion 5114 of lock management storage 5106(0) to ascertain which stream currently has the lock on address X. In the current example, that stream is stream 8. Thereafter, lock manager 5102 accesses entry 8 (the entry that corresponds to stream 8) in the wait list storage structure 5104 and obtains therefrom a thread ID. In the current example, the ID (S60) of stream 60 is stored in entry 8; thus, lock manager 5102 determines that stream 60 is the next stream to be granted a lock on address X. To grant the lock to stream 60, lock manager 5102 writes stream 60's ID

(S60) into portion 5114 of lock management storage 5106(0) to indicate that stream 60 now has the lock on address X. This is shown in Fig. 56. After that is done, lock manager 5102 accesses the memory location at address X of memory 5004(0), and reads the current contents therefrom. Lock manager 5102 then provides the current contents of address X to stream 60, and sends a proceed signal to stream 60. This signal causes stream 60 to resume execution of its thread. In this manner, lock manager 5102 grants the lock to stream 60 and wakes it up.

[0387] Suppose now that stream 60, at some point, encounters an SC instruction in the thread that it is executing. Upon executing this instruction, stream 60 determines that it needs to write to and release the lock on address X. As a result of executing the SC instruction, stream 60 submits a write request (through the global unit 108 and scheduler 5130) to lock manager 5102. This request includes stream 60's ID (S60), address X, the actual write data to be written into address X or a reference to a register that currently contains the write data, and an indication that this is an SC type of request. In one embodiment, after submitting the write request, stream 60 proceeds to execute other instructions in its thread.

[0388] At some point, the write request is received by lock manager 5102 from the scheduler 5130. From the information in the request, lock manager 5102 knows that this is an SC type of request; thus, it knows to implement the lock reservation methodology. To do so, lock manager 5102 first verifies that the write request is from a stream that currently has a lock on the address X. In the present example, lock manager 5102 finds, from the information stored in lock management storage 5106(0), that there currently is a lock on address X, and that the lock is currently held by stream 60. Thus, the write

request is from a proper stream. In response to this write request, lock manager 5102 accesses the address X in memory 5004(0), and stores the write data therein. The memory location at address X is thus again updated with updated content.

[0389] Lock manager 5102 recognizes that the write request (because it is of type SC) is also a request by stream 60 to release its lock on address X. Thus, in response to this write request, lock manager 5102 checks to see which stream (if any) is the next stream to be granted the lock on address X. To do so, lock manager 5102 compares the stream ID stored in portion 5114 of lock management storage 5106(0) with the stream ID stored in portion 5116. If these are the same, then there are no streams waiting for a lock on the address X. In the current example, the same ID is stored in both portion 5114 and portion 5116. Thus, there are no streams on the wait list. In this case, lock manager 5102 resets the lock bit in portion 5110 of lock management storage 5106(0) to release the lock on address X. Lock management storage 5106(0) may thereafter be reused to lock the same or a different address.

Management of Concurrent Locks

[0390] In the above sample operation, the locking of only one address (address X) is discussed. It should be noted, though, that lock manager 5102 can manage the concurrent locking of multiple addresses. For example, while lock management storage 5106(0) is used to store the locking information for address X, another lock management storage, such as storage 5106(n), may be used to store the locking information for an address Y. To manage the lock on address Y, lock management storage 5106(n) and wait list storage structure 5104 may be manipulated and updated in the same manner as that described

above for address X. Overall, each of the lock management storages 5106 may be used to manage a lock on a particular address, and the wait list storage structure 5104 may be used to maintain the wait lists (if any) for all of these locks.

Memory Access Optimization

[0391] In the sample operation described above, when lock manager 5102 processes a write request (of type SC) from a stream, lock manager 5102 accesses the address X in memory 5004(0) two times. The first time is to write the updated contents received from the stream into address X. The second time is to read the updated contents from address X to provide those contents to the next stream on the wait list. Because these two accesses are performed one after the other, the information written into and read from the address X are the same. Thus, the read operation is somewhat redundant.

[0392] To eliminate the need for the read operation, several additional portions may be added to each lock management storage 5106(0). This is shown in Fig. 57, wherein the lock management storage 5106(0) has been augmented to further comprise a value portion 5702, a value valid portion 5704, and a dirty bit portion 5706. The value portion 5702 may be used to store the most current contents of the address locked by the lock management storage 5106. With the current contents available in portion 5702, it is no longer necessary to access the memory to read the current contents of the address therefrom. The value valid portion 5704 stores a single bit. If this bit is set (i.e. is a "1"), it means that the contents stored in portion 5702 are valid, and hence, can be used. If this bit is not set (i.e. is a "0"), then it means that the contents in portion 5702 are not valid and cannot be used. The dirty bit portion 5706 stores a bit that indicates whether the

contents stored in portion 5702 are currently consistent with the contents stored in the address locked by the lock management storage 5106(0). Sometimes, it may be desirable to store the most current contents in portion 5702 without also storing those contents in memory (this avoids a memory access). The dirty bit portion 5706 provides an indication as to whether this has been done. Specifically, if a "1" is stored in portion 5706, then it means that the contents of portion 5702 have not been written into memory (thereby meaning that the contents of 5702 are "dirty" or inconsistent with the contents of the memory). On the other hand, if a "0" is stored in portion 5706, then it means that the contents of portion 5702 and the contents of memory are consistent.

[0393] To illustrate how these portions 5702, 5704, 5706 may be used, reference will be made to an example. Suppose that stream 3 executes an LL instruction, and as a result thereof, submits a request to lock manager 5102 to lock address X of memory 5004(0). Suppose further that address X is not currently locked by another stream. Thus, stream 3 is granted the lock on address X. In granting the lock, lock manager 5102 selects one of the lock management storages (assume storage 5106(0), for the sake of example), and updates it as follows: (1) writes a "1" into portion 5110; (2) stores X into portion 5112; (3) stores the ID (S3) of stream 3 into portions 5114 and 5116; (4) writes a "0" into portion 5704 to indicate that the data in value portion 5702 is currently not valid; and (5) writes a "0" into portion 5706.

[0394] After updating lock management storage 5106(0) in the above manner, lock manager 5102 accesses address X in memory 5004(0), and reads the current contents therefrom. Lock manager 5102 then stores the current contents of address X into portion 5702, and sets the bit in portion 5704 to "1" to indicate that the data in value portion 5702

is now valid. After that is done, lock manager 5106 provides the current contents to stream 3, and sends a proceed signal to stream 3 to cause stream 3 to proceed with the execution of its thread.

[0395] Suppose now that while stream 3 has a lock on address X, stream 8 submits a request to lock address X. In the manner described previously, lock manager 5102 adds stream 8 to a wait list of streams waiting for a lock on address X. As part of this process, lock manager 5102 writes the ID (S8) of stream 8 into portion 5116 to indicate that stream 8 is now the last stream on the wait list. Thereafter, lock manager 5102 sends a wait signal to stream 8 to cause stream 8 to wait for a lock on address X.

[0396] Suppose now that lock manager 5102 receives a write request (of type SC) from stream 3. In response to this write request, lock manager 5102 stores a set of updated contents (the updated contents are provided by stream 3 as part of the write request) into value portion 5702. In addition, lock manager 5102 sets the dirty bit in portion 5706 to "1" to indicate that the contents of value portion 5702 are now no longer consistent with the contents in address X. Thereafter, lock manager 5102 releases the lock held by stream 3 on address X, and grants the lock on address X to stream 8. In granting the lock to stream 8, lock manager 5102 writes the ID (S8) of stream 8 into portion 5114 to indicate that stream 8 now has the lock on address X. After the lock is granted to stream 8, lock manager 5102 provides to stream 8 the current contents of address X. To do so, lock manager 5102 does not access address X (in fact, if it did, it would obtain stale data). Rather, lock manager 5102 reads the value portion 5702 of lock management storage 5106(0), and provides the contents contained therein to stream 8 (lock manager 5102 knows the contents of value portion 5702 are valid because the bit in

portion 5704 has been set). In this manner, the memory read is eliminated. In addition to providing the contents of portion 5702 to stream 8, lock manager 5102 also sends a proceed signal to stream 8 to wake it up.

[0397] Suppose now that lock manager 5102 receives a write request (of type SC) from stream 8. In response to this write request, lock manager 5102 stores a set of updated contents (the updated contents are provided by stream 8 as part of the write request) into value portion 5702. In addition, lock manager 5102 sets the dirty bit in portion 5706 to "1" (if it is not already set) to indicate that the contents of value portion 5702 are now no longer consistent with the contents in address X. Thereafter, lock manager 5102 releases the lock held by stream 8 on address X, and grants the lock on address X to the next stream on the wait list. In the current example, there is no stream waiting for a lock on address X. Thus, lock manager 5102 releases the lock on address X by resetting the lock bit in portion 5110 to "0". Also, lock manager 5102 sees that the dirty bit in 5706 is set, which means that address X has not been updated with the most current contents. Thus, lock manager 5102 reads the current contents from value portion 5702, accesses address X in memory 5004(0), and stores the current contents therein. The address X in memory 5004(0) is thus updated. The lock management storage 5106(0) may thereafter be used to lock another address (or address X again).

[0398] In the above example, when lock manager 5102 releases the lock on address X (by resetting the lock bit in portion 5110), it updates address X with the current contents in value portion 5702 (assuming the value valid bit in portion 5704 and the dirty bit in portion 5706 are set). As an alternative, lock manager 5102 may forego this update. In such a case, the address X would continue to have stale data.

[0399] If address X is not updated after the lock on address X is released, then two possible events might occur. First, lock manager 5102 may receive a request from a stream to lock address X again. In such a case, since lock management storage 5106(0) already has address X stored in portion 5112, lock manager 5102 would use lock management storage 5106(0) again to lock address X. To grant the lock to the requesting stream, lock manager 5102 sets the bit in lock portion 5110 to "1", and writes the ID of the stream into portions 5114 and 5116. In addition, lock manager 5102 checks the bit in value valid portion 5704. If this bit is set (and it should be), thereby indicating that the data in value portion 5702 is still valid, then lock manager 5102 obtains the data from value portion 5702 and provides that data to the requesting stream. In this manner, lock manager 5102 avoids accessing address X in memory 5004(0).

[0400] A second possibility is that lock manager 5102 receives a request from a stream to lock another address (assume address Y for the sake of example), and lock manager 5102 decides to use lock management storage 5106(0) to manage the lock on that address. In such a case, before lock manager 5102 writes any information into lock management storage 5106(0), it checks the bits in the value valid portion 5704 and the dirty bit portion 5706. If both of these bits are "1", then lock manager 5102 obtains the data from value portion 5702, and writes that data into address X (the address that is currently stored in address portion 5112). By doing so, lock manager 5102 ensures that address X is kept up to date. Thereafter, lock manager 5102 can use the lock management storage 5106(0) to manage the lock on address Y in the same manner as that described previously.

[0401] In addition to being used for the lock release/grant process, portion 5702 may also be used for other reads and writes. Basically, whenever a read or write request is received by lock manager 5102 to read from or write to an address that is currently in address portion 5112 of a lock management storage 5106(0) (regardless of whether the bit in lock portion 5110 is set), lock manager 5102 does not need to access the memory 5004. Rather, in the case of a read, it can obtain the most current contents of the address from portion 5702 of the lock management storage 5106(0) instead of going to memory (assuming the bit in value valid portion 5704 is set). In the case of a write, it can write the write data into value portion 5702, set the bit in value valid portion 5704, and set the bit in dirty bit portion 5706 instead of going to memory. Thus, portion 5702 may be used like a cache for the address. This and other uses are within the scope of the present invention.

Accommodating Non-Locking Storage

[0402] In the embodiment described thus far, it has been assumed that once a stream has a lock on an address, no other stream is allowed to write into that address. While this assumption is generally true, there may be implementations in which it is desirable to allow a stream to write into an address even when that address is locked by another stream. This may be done, for example, for override purposes. To accommodate such implementations, several changes may be made to the embodiment described above.

[0403] One change is to incorporate another portion into each lock management storage 5106. As shown in Fig. 58, a lock management storage 5106(0) may be further augmented to include a lock broken portion 5802. In one embodiment, this portion 5802

stores a single bit. When this bit is not set (i.e. is a "0"), it indicates that the address specified in portion 5112 has not been written into since the stream with the current lock acquired the lock on the address. Thus, when this bit is not set, the stream with the current lock still has exclusivity on the address (thus, the lock has not been broken). On the other hand, if the bit in portion 5802 is set (i.e. is a "1"), then it means that another stream has written into the address specified in portion 5112 (thus, the lock has been broken). This in turn means that the stream with the current lock on the address no longer has exclusivity on the address and hence, is working with stale data (the data in the address has been updated since the stream obtained the data). In such a case, the stream should not be allowed to write into the address since doing so would result in data consistency being compromised. Depending upon the state of the bit in portion 5802, the lock manager 5102 will behave differently. Use of portion 5802 by the lock manager 5102 will be discussed further in a later section.

[0404] Another change that can be made is to the manner in which each stream executes the SC instruction. As described thus far, when a stream executes an SC instruction, it submits a write request to the lock manager 5102. The stream then assumes that the write operation succeeded and proceeds to execute other instructions in its thread without waiting for a response. This assumption can be made if exclusivity to the address is guaranteed. However, in an implementation where a stream is allowed to write into an address even when another stream has a lock on that address, exclusivity is not guaranteed; thus, this assumption cannot be made. For such an implementation, each stream executes an SC instruction slightly differently. Specifically, after submitting a write request, a stream waits for a response from the lock manager 5102. If the response

indicates a success, then the stream proceeds to execute other instructions in its thread.

On the other hand, if the response indicates a failure, then the stream may need to take another course of action (for example, request another lock on the address, re-update the contents of the address, and try again to write the updated contents into the address).

[0405] To illustrate how lock manager 5102 may operate in the presence of portion 5802, reference will now be made to an example. Suppose that: (a) stream 3 currently has a lock on address X in memory 5004(0); (b) stream 8 is the next and last stream on the wait list for address X; (c) the current contents of address X is V; and (d) no other thread has written into address X since stream 3 obtained the lock on address X. In such a case, the bit in portion 5110 (Fig. 58) would be set (i.e. a "1"), portion 5112 would contain address X, portion 5114 would contain stream 3's ID (S3), portion 5116 would contain stream 8's ID (S8), portion 5702 would contain V, portion 5704 would contain a "1", portion 5706 would contain a "0" (it is assumed that address X has been updated with the most current contents), and the bit in portion 5802 would be a "0".

[0406] Suppose now that lock manager 5102 receives a write request from stream 191 to write an updated value V' into address X. Suppose further that this write request is of a type that allows an address to be written into even when that address is locked by another stream. In response to this request, lock manager 5102 checks to see if address X is currently locked by another stream. In the current example, address X is locked by stream 3. In such a case, lock manager 5102 updates the value in portion 5702 to V', and sets the bit in portion 5704 to "1". In addition, lock manager 5102 sets the bit in portion 5802 to "1" to indicate that the contents of address X have been updated. Thus, stream 3 no longer has exclusivity on address X, meaning that the lock held by stream 3 on address

X has been broken. Thereafter, lock manager 5102 may optionally access address X in memory 5004(0), and store the updated value V' therein. If it does so, it resets the bit in portion 5706 to "0". If it does not, it sets the bit in portion 5706 to "1".

[0407] Suppose now that stream 3 executes an SC instruction. As a result of executing this instruction, stream 3 submits a write request (of type SC) to lock manager 5102 to write a set of updated content into address X. After submitting this write request, stream 3 waits for a response from lock manager 5102 before executing any other instructions. In response to this write request, lock manager 5102 determines that stream 3 currently has a lock on address X. However, based on the information in portion 5802, lock manager 5102 sees that stream 3's lock on address X has been broken. Thus, lock manager 5102 does not write the updated content provided by stream 3 into address X. Lock manager 5102 also does not write the updated content provided by stream 3 into portion 5702. Instead, lock manager 5102 sends a failure message to stream 3 to indicate that the attempted write failed. In response to this failure message, stream 3 may take any desired course of action, including for example, requesting another lock on address X, updating the contents of address X again, and trying again to write the updated content into address X.

[0408] Even though lock manager 5102 does not write the updated content from stream 3 into address X, lock manager 5102 still recognizes the write request (of type SC) as a request from stream 3 to release its lock on address X. Thus, in response to the write request, lock manager 5102 releases stream 3's lock on address X. In addition, lock manager 5102 determines which stream (if any) is the next stream to obtain a lock on address X, and grants a lock to that stream. In the current example, stream 8 is the next

stream on the wait list; thus, lock manager 5102 grants the lock on address X to stream 8. In granting the lock, lock manager 5102 updates portion 5114 with the ID (S8) of stream 8. In addition, lock manager 5102 sets the bit in portion 5802 back to "0" to indicate that the current stream (stream 8) has exclusivity on address X (i.e. the lock held by the current stream has not been broken). Thereafter, lock manager 5102 reads the current value V' of address X from portion 5702, provides that value to stream 8 (it knows the value is valid because the bit in portion 5704 has been set to "1"), and sends a proceed signal to stream 8. In this manner, lock manager 5102 grants stream 8 the lock on address X, and wakes it up.

[0409] In the manner described, the lock reservation methodology may be implemented even in an implementation in which an address may be written into even when another stream has a lock on the address.

[0410] At this point, it should be noted that although the invention has been described with reference to a specific embodiment, it should not be construed to be so limited. Various modifications may be made by those of ordinary skill in the art with the benefit of this disclosure without departing from the spirit of the invention. Thus, the invention should not be limited by the specific embodiments used to illustrate it but only by the scope of the issued claims and the equivalents thereof.

What is claimed is:

1 1. A method implemented by a lock manager in a multi-threaded
2 environment, comprising:
3 receiving, from a particular stream executing a particular thread, a request to
4 obtain a lock on a resource;
5 determining whether the resource is currently locked by another stream; and
6 in response to a determination that the resource is currently locked by another
7 stream, adding the particular stream to a wait list of streams waiting to obtain a lock on
8 the resource.

1 2. The method of claim 1, further comprising:
2 determining, at a later time based upon the wait list, that it is the particular
3 stream's turn to obtain a lock on the resource; and
4 granting the particular stream a lock on the resource.

1 3. The method of claim 2, wherein the resource comprises one or more
2 storage locations.

1 4. The method of claim 2, further comprising:
2 in response to a determination that the resource is currently locked by another
3 stream, causing the particular stream to halt execution of the particular thread.

1 5. The method of claim 4, wherein causing the particular stream to halt
2 execution of the particular thread comprises:
3 instructing the particular stream to wait for a lock on the resource before
4 continuing with execution of the particular thread.

1 6. The method of claim 4, further comprising:
2 after determining that it is the particular stream's turn to obtain a lock on the
3 resource, causing the particular stream to resume execution of the particular thread.

1 7. The method of claim 6, wherein the resource comprises one or more
2 storage locations, and wherein causing the particular stream to resume execution of the
3 particular thread comprises:
4 providing a set of current contents of the one or more storage locations to the
5 particular stream; and
6 instructing the particular stream to proceed with execution of the particular thread.

1 8. The method of claim 2, wherein adding the particular stream to the wait
2 list of streams comprises:
3 accessing a lock management storage that contains information pertaining to
4 locking of the resource;
5 ascertaining from information in the lock management storage that a certain
6 stream is a last stream on the wait list;
7 accessing, in a wait list storage structure, an entry that corresponds to the certain
8 stream; and
9 storing into the entry a set of information identifying the particular stream.

1 9. The method of claim 8, wherein adding the particular stream to the wait
2 list of streams further comprises:
3 updating information in the lock management storage to indicate that the
4 particular stream is now the last stream on the wait list.

1 10. The method of claim 9, wherein determining that it is the particular
2 stream's turn to obtain a lock on the resource comprises:
3 determining that the certain stream no longer needs a lock on the resource;
4 accessing the entry in the wait list storage structure that corresponds to the certain
5 stream; and
6 obtaining from the entry the set of information identifying the particular stream.

1 11. The method of claim 10, wherein granting the particular stream a lock on
2 the resource comprises:

3 updating information in the lock management storage to indicate that the
4 particular stream now has a lock on the resource.

1 12. The method of claim 11, further comprising:
2 in response to a determination that the resource is currently locked by another
3 stream, causing the particular stream to halt execution of the particular thread.

1 13. The method of claim 12, further comprising:
2 after determining that it is the particular stream's turn to obtain a lock on the
3 resource, causing the particular stream to resume execution of the particular thread.

1 14. The method of claim 9, further comprising:
2 receiving, from a second particular stream executing a second particular thread, a
3 second request to obtain a lock on a second, different resource;
4 determining whether the second resource is currently locked by another stream;
5 in response to a determination that the second resource is currently locked by
6 another stream, adding the second particular stream to a second wait list of streams
7 waiting to obtain a lock on the second resource;
8 determining, at a later time based upon the second wait list, that it is the second
9 particular stream's turn to obtain a lock on the second resource; and
10 granting the second particular stream a lock on the second resource.

1 15. The method of claim 14, wherein adding the second particular stream to
2 the second wait list of streams comprises:
3 accessing a second lock management storage that contains information pertaining
4 to locking of the second resource;
5 ascertaining from information in the second lock management storage that a
6 second certain stream is a last stream on the second wait list;



7 accessing, in the wait list storage structure, an entry that corresponds to the second
8 certain stream; and
9 storing into the entry that corresponds to the second certain stream a set of
10 information identifying the second particular stream.

1 16. The method of claim 15, wherein adding the second particular stream to
2 the second wait list of streams further comprises:
3 updating information in the second lock management storage to indicate that the
4 second particular stream is now the last stream on the second wait list.

1 17. The method of claim 16, wherein determining that it is the second
2 particular stream's turn to obtain a lock on the second resource comprises:
3 determining that the second certain stream no longer needs a lock on the second
4 resource;
5 accessing the entry in the wait list storage structure that corresponds to the second
6 certain stream; and
7 obtaining from the entry that corresponds to the second certain stream the set of
8 information identifying the second particular stream.

1 18. A lock manager, comprising:
2 means for receiving, from a particular stream executing a particular thread, a
3 request to obtain a lock on a resource;
4 means for determining whether the resource is currently locked by another stream;
5 and
6 means for adding, in response to a determination that the resource is currently
7 locked by another stream, the particular stream to a wait list of streams waiting to obtain a
8 lock on the resource.

1 19. The lock manager of claim 18, further comprising:



2 means for determining, at a later time based upon the wait list, that it is the
3 particular stream's turn to obtain a lock on the resource; and
4 means for granting the particular stream a lock on the resource.

1 20. The lock manager of claim 19, wherein the resource comprises one or
2 more storage locations.

1 21. The lock manager of claim 19, further comprising:
2 means for causing, in response to a determination that the resource is currently
3 locked by another stream, the particular stream to halt execution of the particular thread.

1 22. The lock manager of claim 21, wherein the means for causing the
2 particular stream to halt execution of the particular thread comprises:
3 means for instructing the particular stream to wait for a lock on the resource
4 before continuing with execution of the particular thread.

1 23. The lock manager of claim 21, further comprising:
2 means for causing, after determining that it is the particular stream's turn to obtain
3 a lock on the resource, the particular stream to resume execution of the particular thread.

1 24. The lock manager of claim 23, wherein the resource comprises one or
2 more storage locations, and wherein the means for causing the particular stream to
3 resume execution of the particular thread comprises:
4 means for providing a set of current contents of the one or more storage locations
5 to the particular stream; and
6 means for instructing the particular stream to proceed with execution of the
7 particular thread.

1 25. The lock manager of claim 19, wherein the means for adding the particular
2 stream to the wait list of streams comprises:



means for accessing a lock management storage that contains information pertaining to locking of the resource;

means for ascertaining from information in the lock management storage that a certain stream is a last stream on the wait list;

means for accessing, in a wait list storage structure, an entry that corresponds to the certain stream; and

means for storing into the entry a set of information identifying the particular stream.

26. The lock manager of claim 25, wherein the means for adding the particular stream to the wait list of streams further comprises:

means for updating information in the lock management storage to indicate that the particular stream is now the last stream on the wait list.

27. The lock manager of claim 26, wherein the means for determining that it is the particular stream's turn to obtain a lock on the resource comprises:

means for determining that the certain stream no longer needs a lock on the resource;

means for accessing the entry in the wait list storage structure that corresponds to the certain stream; and

means for obtaining from the entry the set of information identifying the particular stream.

28. The lock manager of claim 27, wherein the means for granting the particular stream a lock on the resource comprises:

means for updating information in the lock management storage to indicate that the particular stream now has a lock on the resource.

29. The lock manager of claim 28, further comprising:

means for causing, in response to a determination that the resource is currently locked by another stream, the particular stream to halt execution of the particular thread.

30. The lock manager of claim 29, further comprising:

means for causing, after determining that it is the particular stream's turn to obtain a lock on the resource, the particular stream to resume execution of the particular thread.

31. The lock manager of claim 26, further comprising:

means for receiving, from a second particular stream executing a second particular thread, a second request to obtain a lock on a second, different resource;

means for determining whether the second resource is currently locked by another stream;

means for adding, in response to a determination that the second resource is currently locked by another stream, the second particular stream to a second wait list of streams waiting to obtain a lock on the second resource;

means for determining, at a later time based upon the second wait list, that it is the second particular stream's turn to obtain a lock on the second resource; and

means for granting the second particular stream a lock on the second resource.

32. The lock manager of claim 31, wherein the means for adding the second particular stream to the second wait list of streams comprises:

means for accessing a second lock management storage that contains information pertaining to locking of the second resource;

means for ascertaining from information in the second lock management storage that a second certain stream is a last stream on the second wait list;

means for accessing, in the wait list storage structure, an entry that corresponds to the second certain stream; and

means for storing into the entry that corresponds to the second certain stream a set of information identifying the second particular stream.



1 33. The lock manager of claim 32, wherein the means for adding the second
2 particular stream to the second wait list of streams further comprises:

3 means for updating information in the second lock management storage to
4 indicate that the second particular stream is now the last stream on the second wait list.

1 34. The lock manager of claim 33, wherein the means for determining that it is
2 the second particular stream's turn to obtain a lock on the second resource comprises:

3 means for determining that the second certain stream no longer needs a lock on
4 the second resource;

5 means for accessing the entry in the wait list storage structure that corresponds to
6 the second certain stream; and

7 means for obtaining from the entry that corresponds to the second certain stream
8 the set of information identifying the second particular stream.

1 35. A machine implemented method, comprising:
2 receiving, from a first stream executing a first thread, a request to obtain a lock on
3 a resource;

4 granting the first stream a lock on the resource;

5 receiving, from a second stream executing a second thread, a request to obtain a
6 lock on the resource;

7 determining that the resource is currently locked; and

8 in response to a determination that the resource is currently locked, adding the
9 second stream to a wait list of streams waiting to obtain a lock on the resource, wherein
10 the wait list indicates that the second stream follows the first stream in obtaining a lock
11 on the resource.

1 36. The method of claim 35, further comprising:

2 receiving, from the first stream, an indication that the first stream is releasing the
3 lock on the resource;



4 determining, based upon the wait list, that it is the second stream's turn to obtain a
5 lock on the resource; and
6 granting the second stream a lock on the resource.

1 37. The method of claim 36, wherein granting the first stream a lock on the
2 resource comprises:
3 storing, into a lock management storage, a set of information pertaining to locking
4 of the resource, the set of information indicating that the first stream currently has a lock
5 on the resource and that the first stream is currently the last stream on the wait list of
6 streams to obtain a lock on the resource.

1 38. The method of claim 37, wherein adding the second stream to the wait list
2 of streams comprises:
3 ascertaining, from the set of information in the lock management storage, that the
4 first stream is currently the last stream on the wait list of streams;
5 accessing, in a wait list storage structure, an entry that corresponds to the first
6 stream;
7 storing into that entry a set of information identifying the second stream; and
8 updating the set of information in the lock management storage to indicate that the
9 second stream is now the last stream on the wait list of streams.

1 39. The method of claim 38, further comprising:
2 receiving, from a third stream executing a third thread, a request to obtain a lock
3 on the resource;
4 determining that the resource is currently locked; and
5 in response to a determination that the resource is currently locked, adding the
6 third stream to the wait list of streams by:
7 ascertaining, from the set of information in the lock management storage,
8 that the second stream is currently the last stream on the wait list of
9 streams;

10 accessing, in the wait list storage structure, an entry that corresponds to the
11 second stream;
12 storing into that entry a set of information identifying the third stream; and
13 updating the set of information in the lock management storage to indicate
14 that the third stream is now the last stream on the wait list of
15 streams.

1 40. The method of claim 38, wherein determining that it is the second stream's
2 turn to obtain a lock on the resource comprises:
3 accessing, in the wait list storage structure, the entry that corresponds to the first
4 stream; and
5 obtaining from that entry the information identifying the second stream.

1 41. The method of claim 40, wherein granting the second stream a lock on the
2 resource comprises:
3 updating the set of information in the lock management storage to indicate that the
4 second stream currently has a lock on the resource.

1 42. The method of claim 36, wherein the resource comprises one or more
2 storage locations.

1 43. The method of claim 36, further comprising:
2 in response to a determination that the resource is currently locked, causing the
3 second stream to halt execution of the second thread.

1 44. The method of claim 43, further comprising:
2 after determining that it is the second thread's turn to obtain a lock on the
3 resource, causing the second stream to resume execution of the second thread.



1 45. The method of claim 35, wherein the resource comprises one or more
2 storage locations, and wherein the method further comprises:

3 receiving, from the first stream, a request to write a set of updated contents into
4 the resource and to release the lock on the resource;

5 in response to this request:

6 storing the set of updated contents into a lock management storage that

7 stores information pertaining to locking of the resource;

8 determining, based upon the wait list, that it is the second stream's turn to

9 obtain a lock on the resource;

10 granting the second stream a lock on the resource;

11 obtaining, from the lock management storage and not from the resource,

12 the set of updated contents; and

13 providing the set of updated contents to the second stream.

1 46. The method of claim 35, wherein the resource comprises one or more
2 storage locations, and wherein the method further comprises:

3 receiving, from the first stream, a request to write a set of updated contents into
4 the resource and to release the lock on the resource;

5 in response to this request:

6 determining whether the resource has been updated by another stream after

7 the first stream was granted the lock on the resource;

8 in response to a determination that the resource has been updated by

9 another stream after the first stream was granted the lock on the

10 resource, sending a failure indication to the first stream to indicate

11 that the set of updated contents received from the first stream was

12 not stored into the resource;

13 determining, based upon the wait list, that it is the second stream's turn to

14 obtain a lock on the resource; and

15 granting the second stream a lock on the resource.

1 47. The method of claim 35, wherein the resource comprises one or more
2 storage locations, and wherein the method further comprises:
3 receiving, from a third stream, a request to write a set of updated contents into the
4 resource;
5 storing the set of updated contents into the resource;
6 accessing a lock management storage that contains information pertaining to
7 locking of the resource; and
8 updating the information in the lock management storage to indicate that the
9 resource has been updated after the first stream was granted the lock on the resource.

1 48. The method of claim 47, further comprising:
2 receiving, from the first stream, a request to write a second set of updated contents
3 into the resource and to release the lock on the resource;
4 in response to this request:
5 accessing the lock management storage;
6 ascertaining, from the information in the lock management storage, that
7 the resource has been updated by another stream after the first
8 stream was granted the lock on the resource;
9 in response to this determination, sending a failure indication to the first
10 stream to indicate that the second set of updated contents received
11 from the first stream was not stored into the resource;
12 determining, based upon the wait list, that it is the second stream's turn to
13 obtain a lock on the resource;
14 granting the second stream a lock on the resource; and
15 updating the information in the lock management storage to indicate that
16 the resource has not been updated after the second stream was
17 granted the lock on the resource.

1 49. A lock manager, comprising:

2 means for receiving, from a first stream executing a first thread, a request to
3 obtain a lock on a resource;
4 means for granting the first stream a lock on the resource;
5 means for receiving, from a second stream executing a second thread, a request to
6 obtain a lock on the resource;
7 means for determining that the resource is currently locked; and
8 means for adding, in response to a determination that the resource is currently
9 locked, the second stream to a wait list of streams waiting to obtain a lock on the
10 resource, wherein the wait list indicates that the second stream follows the first stream in
11 obtaining a lock on the resource.

1 50. The lock manager of claim 49, further comprising:
2 means for receiving, from the first stream, an indication that the first stream is
3 releasing the lock on the resource;
4 means for determining, based upon the wait list, that it is the second stream's turn
5 to obtain a lock on the resource; and
6 means for granting the second stream a lock on the resource.

1 51. The lock manager of claim 50, wherein the means for granting the first
2 stream a lock on the resource comprises:
3 means for storing, into a lock management storage, a set of information pertaining
4 to locking of the resource, the set of information indicating that the first stream currently
5 has a lock on the resource and that the first stream is currently the last stream on the wait
6 list of streams to obtain a lock on the resource.

1 52. The lock manager of claim 51, wherein the means for adding the second
2 stream to the wait list of streams comprises:
3 means for ascertaining, from the set of information in the lock management
4 storage, that the first stream is currently the last stream on the wait list of streams;

means for accessing, in a wait list storage structure, an entry that corresponds to the first stream;

means for storing into that entry a set of information identifying the second stream; and

means for updating the set of information in the lock management storage to indicate that the second stream is now the last stream on the wait list of streams.

53. The lock manager of claim 52, further comprising:

means for receiving, from a third stream executing a third thread, a request to obtain a lock on the resource;

means for determining that the resource is currently locked; and

means for adding, in response to a determination that the resource is currently locked, the third stream to the wait list of streams, wherein the means for adding the third stream to the wait list of streams comprises:

means for ascertaining, from the set of information in the lock management storage, that the second stream is currently the last stream on the wait list of streams;

means for accessing, in the wait list storage structure, an entry that corresponds to the second stream;

means for storing into that entry a set of information identifying the third stream; and

means for updating the set of information in the lock management storage to indicate that the third stream is now the last stream on the wait list of streams.

54. The lock manager of claim 52, wherein the means for determining that it is the second stream's turn to obtain a lock on the resource comprises:

means for accessing, in the wait list storage structure, the entry that corresponds to the first stream; and

means for obtaining from that entry the information identifying the second stream.

1 55. The lock manager of claim 54, wherein the means for granting the second
2 stream a lock on the resource comprises:

3 means for updating the set of information in the lock management storage to
4 indicate that the second stream currently has a lock on the resource.

1 56. The lock manager of claim 50, wherein the resource comprises one or
2 more storage locations.

1 57. The lock manager of claim 50, further comprising:
2 means for causing, in response to a determination that the resource is currently
3 locked, the second stream to halt execution of the second thread.

1 58. The lock manager of claim 57, further comprising:
2 means for causing, after determining that it is the second thread's turn to obtain a
3 lock on the resource, the second stream to resume execution of the second thread.

1 59. The lock manager of claim 49, wherein the resource comprises one or
2 more storage locations, and wherein the lock manager further comprises:

3 means for receiving, from the first stream, a request to write a set of updated
4 contents into the resource and to release the lock on the resource;

5 means for responding to this request, comprising:

6 means for storing the set of updated contents into a lock management
7 storage that stores information pertaining to locking of the
8 resource;

9 means for determining, based upon the wait list, that it is the second
10 stream's turn to obtain a lock on the resource;

11 means for granting the second stream a lock on the resource;

12 means for obtaining, from the lock management storage and not from the
13 resource, the set of updated contents; and



14 means for providing the set of updated contents to the second stream.

1 60. The lock manager of claim 49, wherein the resource comprises one or
2 more storage locations, and wherein the lock manager further comprises:

3 means for receiving, from the first stream, a request to write a set of updated
4 contents into the resource and to release the lock on the resource;

5 means for responding to this request, comprising:

6 means for determining whether the resource has been updated by another
7 stream after the first stream was granted the lock on the resource;

8 means for sending, in response to a determination that the resource has
9 been updated by another stream after the first stream was granted
10 the lock on the resource, a failure indication to the first stream to
11 indicate that the set of updated contents received from the first
12 stream was not stored into the resource;

13 means for determining, based upon the wait list, that it is the second
14 stream's turn to obtain a lock on the resource; and

15 means for granting the second stream a lock on the resource.

1 61. The lock manager of claim 49, wherein the resource comprises one or
2 more storage locations, and wherein the lock manager further comprises:

3 means for receiving, from a third stream, a request to write a set of updated
4 contents into the resource;

5 means for storing the set of updated contents into the resource;

6 means for accessing a lock management storage that contains information
7 pertaining to locking of the resource; and

8 means for updating the information in the lock management storage to indicate
9 that the resource has been updated after the first stream was granted the lock on the
10 resource.

1 62. The lock manager of claim 61, further comprising:

2 means for receiving, from the first stream, a request to write a second set of
3 updated contents into the resource and to release the lock on the resource;
4 means for responding to this request, comprising:
5 means for accessing the lock management storage;
6 means for ascertaining, from the information in the lock management
7 storage, that the resource has been updated by another stream after
8 the first stream was granted the lock on the resource;
9 means for sending, in response to this determination, a failure indication to
10 the first stream to indicate that the second set of updated contents
11 received from the first stream was not stored into the resource;
12 means for determining, based upon the wait list, that it is the second
13 stream's turn to obtain a lock on the resource;
14 means for granting the second stream a lock on the resource; and
15 means for updating the information in the lock management storage to
16 indicate that the resource has not been updated after the second
17 stream was granted the lock on the resource.

1 63. A multi-threaded processing engine, comprising:
2 a first stream capable of executing a first thread;
3 a second stream capable of executing a second thread; and
4 a lock manager, wherein the lock manager comprises:
5 means for receiving, from the first stream, a request to obtain a lock on a resource;
6 means for granting the first stream a lock on the resource;
7 means for receiving, from the second stream, a request to obtain a lock on the
8 resource;
9 means for determining that the resource is currently locked; and
10 means for adding, in response to a determination that the resource is currently
11 locked, the second stream to a wait list of streams waiting to obtain a lock on the
12 resource, wherein the wait list indicates that the second stream follows the first stream in
13 obtaining a lock on the resource.



1 64. The processing engine of claim 63, wherein the lock manager further
2 comprises:

3 means for receiving, from the first stream, an indication that the first stream is
4 releasing the lock on the resource;

5 means for determining, based upon the wait list, that it is the second stream's turn
6 to obtain a lock on the resource; and

7 means for granting the second stream a lock on the resource.

1 65. The processing engine of claim 64, wherein the processing engine further
2 comprises a lock management storage, and wherein the means for granting the first
3 stream a lock on the resource comprises:

4 means for storing, into the lock management storage, a set of information
5 pertaining to locking of the resource, the set of information indicating that the first stream
6 currently has a lock on the resource and that the first stream is currently the last stream on
7 the wait list of streams to obtain a lock on the resource.

1 66. The processing engine of claim 65, wherein the processing engine further
2 comprises a wait list storage structure, and wherein the means for adding the second
3 stream to the wait list of streams comprises:

4 means for ascertaining, from the set of information in the lock management
5 storage, that the first stream is currently the last stream on the wait list of streams;

6 means for accessing, in the wait list storage structure, an entry that corresponds to
7 the first stream;

8 means for storing into that entry a set of information identifying the second
9 stream; and

10 means for updating the set of information in the lock management storage to
11 indicate that the second stream is now the last stream on the wait list of streams.

1 67. The processing engine of claim 66, wherein the processing engine further
2 comprises a third stream capable of executing a third thread, and wherein the lock
3 manager further comprises:

4 means for receiving, from the third stream, a request to obtain a lock on the
5 resource;

6 means for determining that the resource is currently locked; and

7 means for adding, in response to a determination that the resource is currently
8 locked, the third stream to the wait list of streams, wherein the means for adding the third
9 stream to the wait list of streams comprises:

10 means for ascertaining, from the set of information in the lock

11 management storage, that the second stream is currently the last
12 stream on the wait list of streams;

13 means for accessing, in the wait list storage structure, an entry that
14 corresponds to the second stream;

15 means for storing into that entry a set of information identifying the third
16 stream; and

17 means for updating the set of information in the lock management storage
18 to indicate that the third stream is now the last stream on the wait
19 list of streams.

1 68. The processing engine of claim 66, wherein the means for determining that
2 it is the second stream's turn to obtain a lock on the resource comprises:

3 means for accessing, in the wait list storage structure, the entry that corresponds to
4 the first stream; and

5 means for obtaining from that entry the information identifying the second stream.

1 69. The processing engine of claim 68, wherein the means for granting the
2 second stream a lock on the resource comprises:

3 means for updating the set of information in the lock management storage to
4 indicate that the second stream currently has a lock on the resource.

1 70. The processing engine of claim 64, wherein the resource comprises one or
2 more storage locations.

1 71. The processing engine of claim 64, wherein the lock manager further
2 comprises:
3 means for causing, in response to a determination that the resource is currently
4 locked, the second stream to halt execution of the second thread.

1 72. The processing engine of claim 71, wherein the lock manager further
2 comprises:
3 means for causing, after determining that it is the second thread's turn to obtain a
4 lock on the resource, the second stream to resume execution of the second thread.

1 73. A method implemented by a lock manager in a multi-threaded
2 environment, comprising:
3 receiving, from a particular thread, a request to obtain a lock on a resource;
4 determining whether the resource is currently locked by another thread;
5 in response to a determination that the resource is currently locked by another
6 thread, adding the particular thread to a wait list of threads waiting to obtain a lock on the
7 resource;
8 determining, at a later time based upon the wait list, that it is the particular
9 thread's turn to obtain a lock on the resource; and
10 granting the particular thread a lock on the resource.

1 74. The method of claim 73, wherein the resource comprises one or more
2 storage locations.

1 75. The method of claim 73, further comprising:

2 in response to a determination that the resource is currently locked by another
3 thread, causing execution of the particular thread to be halted.

1 76. The method of claim 75, further comprising:
2 after determining that it is the particular thread's turn to obtain a lock on the
3 resource, causing execution of the particular thread to be resumed.

1 77. A lock manager, comprising:
2 means for receiving, from a particular thread, a request to obtain a lock on a
3 resource;
4 means for determining whether the resource is currently locked by another thread;
5 means for adding, in response to a determination that the resource is currently
6 locked by another thread, the particular thread to a wait list of threads waiting to obtain a
7 lock on the resource;
8 means for determining, at a later time based upon the wait list, that it is the
9 particular thread's turn to obtain a lock on the resource; and
10 means for granting the particular thread a lock on the resource.

1 78. The lock manager of claim 77, wherein the resource comprises one or
2 more storage locations.

1 79. The lock manager of claim 77, further comprising:
2 means for causing, in response to a determination that the resource is currently
3 locked by another thread, execution of the particular thread to be halted.

1 80. The lock manager of claim 79, further comprising:
2 means for causing, after determining that it is the particular thread's turn to obtain
3 a lock on the resource, execution of the particular thread to be resumed.