DATA CONSISTENCY MODELS, LOCKS, AND LOCK-FREE SYNCHRONIZATION

Ivan Jibaja CS 395T – Topics in Multicore Programming

OUTLINE

- × Recall: PRAM Model
- × New Model
- × Architecture Details
- * Consistency Models (Relaxed, Release)
- × Locks
 - + Disadvantages
- × Lock-free Synchronization
- × Lock-free Data Structures

RECALL: PRAM MODEL RECALL: PRAM MODEL × PRAM Model: × Parallel Random Access + No lock-step execution of global instructions Machine (PRAM) No global shared-memory for reads/writes by all of the + Natural extension of RAM Shared memory processors in each cycle model × New Model: × Processors operate Each processor executes its own program at its own P1 P2 Pp speed synchronously (in lockstep) Only one processor can access memory at a time × Each processor has private How does a processor know what other processors are doing? We need <u>synchronization</u> memory Let's look at our new model in more detail Are PRAM Models any good?



MULTITHREADED ON MULTIPROCESSOR

P

MEMORY



- Each processors executes one thread
 Operations in each thread are executed in order
- One processor at a time can access global memory to perform load/store/atomic operations (no caching of global data)

With these assumptions, you can show that running a multi-threaded program on a multiprocessor does not change possible output from the uniprocessor case

EXAMPLE (I)

Code: Initially A = Flag = 0

P1 A = 23;

Flag = 1;

- Semantics:
 - + P1 writes data into A and sets Flag to tell P2 that data value can be read from A.

P2

... = A;

while (Flag != 1) {;}

+ P2 waits till Flag is set and then reads data from A.

EXAMPLE (II) Code: (similar to Dekker's algorithm) Initially Flag1 = Flag2 = 0 P1 P2 Flag1 = 1; Flag2 = 1; If (Flag2 == 0) If (Flag1 == 0) critical section critical section • What is the problem with our model? We are making unrealistic architectural assumptions. Note whether the problems require turbatere support in the form of atomic instructions. We take the take in the turbatered

ARCHITECTURE DETAILS

- × We have some architectural constrains with 2 of our assumptions:
- Processors do not cache global data:
 - For execution efficiency, processors are allowed to cache global data:
 - Leads to cache coherence problems, which can be solved using coherent caches
- Instructions within each thread are executed in order
 - For execution efficiency, processors are allowed to execute instructions out of order subject to data/control dependances:
 - . Changes the semantics of the program
 - To prevent this requires attention to memory consistency models

CURRENT PROCESSORS

		Loads Reordered After Loads?	Loads Reordered After Stores?	Stores Reordered After Stores?	Stores Reordered After Loads?	Atomic Instructions Reordered With Loads?	Atomic Instructions Reordered With Stores?	Dependent Loads Reordered?	Incoherent Instruction Cache/Pipeline?		
	Alpha	Y	Y	Y	Y	Y	Y	Y	Y		
	AMD64	Y			Y						
	IA64	Y	Y	Y	Y	Y	Υ		Y		
	(PA-RISC)	Y	Y	Y	Y						
	PA-RISC CPUs										
	POWER	Y	Y	Y	Y	Y	Υ		Y		
	SPARC RMO	Y	Y	Y	Y	Y	Y		Y		
	(SPARC PSO)			Y	Y		Υ		Y		
	SPARC TSO				Y				Y		
	x86	Y	Y		Y				Y		
	(x86 OOStore)	Y	Y	Y	Y				Y		
	zSeries				Y				Y		

RECALL: UNIPROCESSOR EXECUTION

- Adding new constrains to our uniprocessor model:
 - + Processors reorder operations to improve performance
 - + Constraint on reordering: must respect dependences data dependences must be respected: in particular,
 - loads/stores to a given memory address must be executed in program order
 - control dependences must be respected
 - Reorderings can be performed either by the compiler or the processor

PERMITTED MEMORY-OP REORDERINGS

- * Stores to different memory locations can be performed out of program order store b1, flag store v1, data \leftrightarrow store b1. flag
 - store v1. data
- × Loads from different memory locations can be performed out of program order load data,r2

load flag, r1 load data, r2 \leftrightarrow load flag, r1

Load and store to different memory locations can be performed out of program order



PROBLEM IN MULTIPROCESSOR CONTEXT

× Our first model (Canonical ordering)

- + Operations from given processor are executed in program order
- + Memory operations from different processors appear to be interleaved in some order at the memory

• Our revisited model:

- + If a processor is allowed to reorder independent operations in its <u>own</u> instruction stream, will the execution always produce the same results as the canonical model?
- + Answer: ?

EXAMPLE (I) REVISITED

Code: Initially A = Flag = 0

P1 A = 23;

Flag = 1;

- Semantics:
 - + P1 writes data into A and sets Flag to tell P2 that data value can be read from A.

P2

... = A;

while (Flag != 1) {;}

+ P2 waits till Flag is set and then reads data from A.

EXECUTION SEQUENCE FOR (I) Code Initially A = Flag = 0 while (Flag != 1) {;} ... = A; A = 23; Flag = 1; Possible execution sequence on each processor: P1 P2 Write A 23 Write Flag 1 Read Flag //get 0 Read Flag Read A //get 1 //what do you get? Problem: If the two writes on processor P1 can be reordered, it is possible for processor P2 to read 0 from variable A. Can happen on most modern processors

EXAMPLE (II) REVISITED Code: (similar to Dekker's algorithm) Initially Flag1 = Flag2 = 0 P2 P1 Flag1 = 1; Flag2 = 1;If (Flag2 == 0)If (Flag1 == 0) critical section critical section Possible execution sequence on each processor: P1 P2 Write Flag1, 1 Write Flag2, 1 Read Flag2 //get 0 Read Flag1 //what could you get?

SUMMARIZING THE PROBLEM

- Uniprocessors can reorder instructions subject only to control and data dependence constraints
- These constraints are not sufficient in shared-memory context
 + simple parallel programs may produce counter-intuitive results
- Question: what constraints must we put on uniprocessor instruction reordering so that:
 - + shared-memory programming is intuitive?
 - + but we do not lose uniprocessor performance?
- Many answers to this question:
- + answer is called memory consistency model supported by the processor

CONSISTENCY MODELS

- Consistency models are <u>not</u> about memory operations <u>from different processors</u>.
- Consistency models are <u>not</u> about <u>dependent</u> <u>memory operations in a single processor's</u> instruction stream (these are respected even by processors that reorder instructions).

CONSISTENCY MODELS

 Consistency models are all about ordering constraints on <u>independent memory operations</u> <u>in a single processor's</u> instruction that should be respected to obtain intuitively reasonable results.

SIMPLEST MEMORY CONSISTENCY MODEL

- × Sequential consistency (SC) [Lamport]
 - + Our canonical model: processor is not allowed to reorder reads and writes to global memory



SEQUENTIAL CONSISTENCY

- * SC constrains all memory operations:
 - \checkmark Write \rightarrow Read
 - Write \rightarrow Write
 - × Read \rightarrow Read, Write
 - Simple model for reasoning about parallel programs
 - You can verify that the examples considered earlier work correctly under sequential consistency.
 - However, this simplicity comes at the cost of uniprocessor performance.

Question: how do we reconcile sequential consistency model with the demands of performance?

RELAXED MODEL: WEAK CONSISTENCY

- * Programmer specifies regions within which global memory operations can be reordered
- Processor has fence instruction:
- + all data operations before fence in program order must complete before fence is executed
 + all data operations after fence in program order must wait for fence to complete
- + all data operations after fence in program order must wait for fence to comp
 + fences are performed in program order
- Implementation of fence:
- + processor has counter that is incremented when data op is issued, and decremented when data op is completed
- Example: PowerPC has SYNC instruction
- Language constructs:
- + OpenMP: flush
 - + All synchronization operations like lock and unlock act like a fenc





RELAXED MODEL: RELEASE CONSISTENCY

- × Further relaxation of weak consistency
- Synchronization accesses are divided into
- + Acquires: operations like lock
- + Release: operations like unlock
- Semantics of acquire:
- + Acquire must complete before all following memory accesses Semantics of release:
- + All memory operations (read/write) before release are complete However,
- + Acquire does not wait for accesses preceding it
- Accesses after release in program order do not have to wait for release
 Operations which follow release and which need to wait must be protected by an acquire

COMMENTS

- × There are a lot of other consistency models out there:
 - + Causal consistency
 - + Processor consistency
 - + Delta consistency....
- It is important to remember that these are concerned with reordering of independent memory operations within a processor.
- Easy to come up with shared-memory programs that behave differently for each consistency model.

MEMORY CONSISTENCY

- What instructions is compiler or hardware allowed to reorder?
 * Preserve the illusion that there is a single logical memory location
- Nothing really to do with memory operations from different processors/threads
- Sequential consistency:
- perform global memory operations in program order Relaxed consistency models: all of them rely on some notion of a fence operation that
- single logical memory location corresponding to each program variable even though there may be lots of physical memory locations where the variable is stored

MEMORY COHERENCE

demarcates regions within which reordering is permissible

CONSISTENCY VS. COHERENCE

OTHER TOOLS FOR SYNCHRONIZATION

LOCKS

3

× In a Uniprocessor, locking was achieved by disabling interrupts for the smallest possible number of instructions that will access shared data ("the critical section"):

while (true) { /* disable interrupts */ /* critical section */ /* enable interrupts */ /* remainder */

This approach does not work with multiprocessors. Why?

LOCK IMPLEMENTATION IN HARDWARE × Test and Set, here on, TS: + TS on a boolean variable flag #atomic // The two lines below will be executed one after the other without interruption If(flag == false) bool lock = false; // shared lock variable // Process i flag = true; Init **i;** while(true) { #end atomic while (lock==false){ // entry protocol TS(lock)}; lock = false; // exit protocol //Remainder of code;}

DISADVANTAGES OF LOCKS

- × Notice the while loop in the algorithm
- If process 0 waits a lot of time to enter the critical section, it continually checks the flag and turn to see it can or not, while <u>not doing any useful work</u>
- This is termed <u>busy waiting</u>, and locking mechanisms have a major disadvantage in that regard.
- Locks that employ continuous checking mechanism for a flag are called <u>Spin-Locks</u>.
- Spin locks are good when the you know that the wait is not long enough.

EXAMPLE: A LOCK BASED STACK

- Stack: A list or an array based data structure that enforces last-in-first-out ordering of elements
- × Operations
 - + Void Push(T data) : pushes the variable data on to the stack
 - + T Pop() : removes the last item that was pushed on to a stack. Throws a stackEmptyException if the stack is empty
 - + Int Size() : returns the size of the stack
- All operations are synchronized using one common lock object.

CODE : JAVA

Class Stack<T> {

ArrayList<T> _container = new ArrayList<T>(); RentrantLock _lock = new ReentrantLock();

public void push(T data){ _lock.lock(); _container.add(data); _lock.unlock();}

public int size(){
 int retVal; lock.Lock(); retVal = _container.size();
 _lock.unlock();
 return retVal;

return retval,

public T pop(){ _lock.lock();

if(_container.empty()) { _lock.unlock();

__introver new Exception("Stack Empty");}
T retVal _container.get(_container.size() - 1);
_lock.unlock(); return retVal;

PROBLEMS WITH LOCKS

- Stack is simple enough. There is only one lock. The overhead isn't that much. But there are data structures that could have multiple locks
- Problems with locking
 - + Deadlock
 - + Priority inversion
 - + Convoying
 - + Preemption tolerance
 - + Overall performance

PROBLEMS WITH LOCKING 2

× Priority inversion:

- + Assume two threads:
 - ×T2 with very low priority
 - ×T1 with very high priority
- Both need to access a shared resource R but T2 holds the lock to R
 - xT2 takes longer to complete the operation leaving the higher priority thread waiting, hence by extension T1 has achieved a lower priority

PROBLEMS WITH LOCKING 3

 Deadlock: Processes can't proceed because each of them is waiting for the other release a needed resource.

× Scenario:

- + There are two locks A and B
- + Process 1 needs A and B in that order to safely execute
- + Process 2 needs B and A in that order to safely execute
 + Process 1 acquires A and Process two acquires B
- + Now Process 1 is waiting for Process 2 to release B and Process 2 is waiting for process 1 to release A

PROBLEMS WITH LOCKING 4

- Convoying, all the processes need a lock A to proceed however, a lower priority process acquires A it first. Then all the other processes slow down to the speed of the lower priority process.
- × Think of a freeway:
 - + You are driving an Aston Martin but you are stuck behind a beat up old pick truck that is moving very slow and there is no way to overtake him.

PROBLEMS WITH LOCKING 5

- × Overall performance
 - + Arguable
 - + Efficient lock-based algorithms exist
 - + Constant struggle between simplicity and efficiency
 - + Example. thread-safe linked list with lots of nodes
 - × Lock the whole list for every operation?
 - × Reader/writer locks?
 - × Allow locking individual elements of the list?

LOCK-FREE SYNCHRONIZATION ("NON-BLOCKING")

LOCK-FREE DATA STRUCTURES

A data structure wherein there are no explicit locks used for achieving synchronization between multiple threads, and the progress of one thread doesn't block/impede the progress of another.

 Doesn't imply starvation freedom (Meaning one thread could potentially wait forever). But nobody starves in practice

- Advantages:
 - + You don't run into all the problems that you would with using locks
- × Disadvantages: To be discussed later

LOCK-FREE SYNCHRONIZATION

- Think in terms of Algorithms + Data Structure = Program
- Thread safe access to shared data without the use of locks, mutexes etc.
- Possible but not practical/feasible in the absence of hardware support
- × So what do we need?
 - + Design algorithm to avoid critical sections
 - + A compare and set primitive (CAS) from the hardware

BUILDING LOCK-FREE DATA STRUCTURES

- * What do we need to build these lockfree data structures?
- * Limit the scope of changes to a single atomic variable
 - + Stack : head
 - + Queue: head or tail depending on enque or deque

A SIMPLE LOCK-FREE EXAMPLE

× A lock-free Stack

- * Adopted from Geoff Langdale at CMU
 - Intended to illustrate the design of lock-free data structures and problems with lock-free synchronization
 - + There is a primitive operation we need:
 - * Compare and Set (CAS)
 - Available on most modern machines
 - X86 assembly: xchg
 PowerPC assembly: LL(load linked), SC (Store Conditional)

LOCK-FREE STACK WITH INTS IN C

A stack based on a singly linked list. Not particularly good design!

```
struct NodeEle {
    int data;
    Node *next;
};
```

```
typedef NodeEle Node;
```

Node* head; // The head of the list

Now that we have the nodes let us proceed to body of the stack

LOCK-FREE STACK PUSH

void push(int t) {
 Node* node = malloc(sizeof(Node));
 node->data = t;
 do {
 node->next = head;
 } while (!cas(&head, node, node->next));
}
Let us see how this works!



























SOLUTIONS:

x Double word compare and set.

- + One 32 bit word for the address
- One 32 bit word for the update count which is incremented every time a node is updated
- + Compare and Set iff both of the above match
- + Java provides AtomicStampedReference
- Use the lower address bits of the pointer (if the memory
- is 4/8 byte Aligned) to keep a counter to update
 But the probability of a false positive is still greater than doubleword compareandset because instead of 2^32 choices for the counter you have 2^2 or 2^3 choices for the counter

DISADV. OF LOCK-FREE DATA STRUCTURES

- Current hardware limits the amount of bits available in CAS operation to 32/64 bits.
- Imagine the implementation of data structures like BST's pose a problem
 - + When you need to balance a tree you need update several nodes all at once.
- Way to get around it
 - + Transactional memory based systems

