

Symmetric Multiprocessors Part 1 (Chapter 5)

Copyright 2003 Mark D. Hill University of Wisconsin-Madison

Slides are derived from work by Sarita Adve (Illinois), Babak Falsafi (CMU), Alvy Lebeck (Duke), Steve Reinhardt (Michigan), and J. P. Singh (Princeton). Thanks!









Why Shared Memory? • Pluses - For applications looks like multitasking uniprocessor - For OS only evolutionary extensions required - Easy to do communication without OS - Software can worry about correctness first then performance Minuses - Proper synchronization is complex - Communication is implicit so harder to optimize - Hardware designers must implement • Result - Symmetric Multiprocessors (SMPs) are the most success parallel machines ever - And the first with multi-billion-dollar markets

UPC Parallel Computer Architecture

(C) 2003 Mark D. Hill from Adve, Falsafi, Lebeck, Reinhardt, & Singh



































			-		
roc Action	P1 State	P2 state	P3 state	Bus Act	Data from
. P1 read u	S			BusRd	Memory
. P3 read u	S		S	BusRd	Memory
. P3 write u	1		М	BusRdX	Memory or not
. P1 read u	S		S	BusRd	P3's cache
. P2 read u	S	S	S	BusRd	Memory
Why Mo What if	odified to not in ar Write pro	o Sharec ny cache duces 2 bi	l? l? us transact	tions!	











FROM/TO	NP	Ι	Е	S	М
NP			BusRd	BusRd	BusRdX
			6+64	6+64	6+64
I			BusRd	BusRd	BusRdX
			6+64	6+64	6+64
E					
S			NA		BusUpgr
М	BusWB	BusWB	NA	BusWB	
	6 + 64	6+64		6 + 64	







Invalidate vs. Update
Pattern 1:
for i = 1 to k
P1(write, x); // one write before reads
P2PN-1(read, x);
end for i
Pattern 2:
for i = 1 to k
for j = 1 to m
P1(write, x); // many writes before reads
end for j
P2(read, x);
end for i
(C) 2003 Mark D. Hill from Adve, Falsafi, Lebeck, Reinhardt, & Singh UPC Parallel Computer Architecture 25



Qualitative Sharing Patterns • [Weber & Gupta, ASPLOS3] · Read-Only • Migratory Objects Maniputalated by one processor at a time Often protected by a lock - Usually a write causes only a single invalidation Synchronization Objects Often more processors imply more invalidations Mostly Read - More processors imply more invalidations, but writes are rare • Frequently Read/Written - More processors imply more invalidations (C) 2003 Mark D. Hill from Adve, Falsafi, Lebeck, Reinhardt, & Singh UPC Parallel Computer Architecture 27



- Intuition says loads should return latest value
 what is latest?
- Coherence concerns only one memory location
 Consistency concerns apparent ordering for all
- Iocations
 A Memory System is Coherent if
 can serialize all operations to that location such that,
 - can serialize all operations to that location such that,
 operations performed by any processor appear in program order
 » program order = order defined by program text or assembly
 code
 - value returned a read is value written by last store to that location

28

(C) 2003 Mark D. Hill from Adve, Falsafi, Lebeck, Reinhardt, & Singh UPC Parallel Computer Architecture

Why Coherence != Consistency				
/* initial A = I	B = flag = 0 */			
<u>P1</u>	<u>P2</u>			
A = 1;	while (flag == 0);			
B = 1;	print A;			
flag = 1;	print B;			
Intuition says print	ed A = B = 1			
Coherence doesn't	say anything, why?			
Consider coalescir	ng write buffer			
	0			
(C) 2003 Mark D. Hill from Adve,	UPC Parallel Computer Architecture	29		

Sequential Consistency	\sum
• Lamport 1979	
"A multiprocessor is sequentially consistent if the result of any execution is the same as if the operations of all the processors were executed in some sequential order, and the operations of each individual processor appear in this sequence in the order specified by its program"	
C) 2003 Mark D. Hill from Adve, Falsafi, Lebeck, Reinhardt, & Singh UPC Parallel Computer Architecture	30























Alpha test & set Bit in 32-bit word				
t_and_set	t:		# arg reg a0 holds bit# to test	
1:	ldl_l	t0, (a1)	#get word to test	
	bis	zero, 1, t1	# logical OR set t1=1	
	sll	t1, a0, t2	# create bitmask	
	and	t0, t2, t3	# isolate bit	
	bne	t3, 2f	# branch if already set	
	bis	t0, t2, v0	# or in bit to be set	
	stl_q	v0, (a1)	# conditional store	
	beq	v0, 3f	# branch if store failed	
	mb		# memory barrier	
	ret	zero, (ra)	# ret 1 (v0) if set success	
2:	bis	zero,zero,v0	# ret 0 if bit already set	
	ret	zero, (ra)		
3:	br	zero, 1b	# retry interlocked update	
(C) 2003 Mark D. Hill from Adve, Falsafi, Lebeck, Reinhardt, & Singh UPC Parallel Computer Architecture				40























