





Conditions for a clock	۵۰ ۵۷ ۵۵
C1: If a h are events in process D and	ດີ ດີ ດ ເຈັ້ນ ເຈັ້ນ ເຈັ້ນ ເຈັ້ນ
• C1. If a, b are events in process P_i and	
a comes before <i>b</i> , then $C_i \langle a \rangle < C_i \langle b \rangle$	Patra ga
• C2: If event a is the sending of a	
message by process P_i and event <i>b</i> is	$P_{2} = \frac{1}{2} = \frac{1}{2$
the receipt of that message by process \mathbf{D}	
P_{j} , then $C_{I}\langle a \rangle < C_{I}\langle b \rangle$	P2 4
Space-Time Diagram	
 Add ticks for every count in each process 	
 Draw "tick lines" between equally numbered ticks 	Remarks
• C1 \Rightarrow tick line between two events	lines without changing logical clocks ⇒
• C2 \Rightarrow every msg must cross tick line	logical time is different than physical
, ,	time

Constructing Clocks

Implementation rule

- IR1: Each process P_i increments C_i between any two successive events
- IR2:
 - (a) If event *a* is the sending a message *m* by process P_i, then the message *m* contains a timestamp $T_m = C_h (a)$
 - (b) Upon receiving a message m, process P_j sets C_j greater than or equal to its present value and greater than T_m

Remarks:

- Gives an easy algorithm for constructing a clock
- Note that C⟨a⟩ < C⟨b⟩ does not imply a → b. Still only a partial order (can only compare certain elements)



Total order

- Order events according to logical clocks
- Break ties using process number
- Allows any two events to be compared
- Total ordering is not unique (depends on choice of clocks)

HYCON-EECI, Mar 08

R. M. Murray, Caltech CDS

Example: Resource allocation Algorithm P4 P5 P2 P31. P_1 sends message $T_m:P_i$ request to every other process and puts message on its queue 2. P_i queues all requests and sends R timestamped acknowledgement to sender 3. Process P_i uses resource when **Problem description** Fixed processes P_i sharing resource R 1.T_m:P_i request is ordered before any other request in queue (according to Once a process grabs a resource, it total order) must release it before it is use again 2.P, has been received ack from · Requests granted in order they were everyone with timestamp > T_m requested · Every request is eventually granted 3. P_i removes T_m:P_i request message from queue and sends T_m:P_i release Solve in *distributed* way; processes message to everyone agree on who goes next When P_i receives a $T_m:P_i$ release · Problem is non-trivial, even with central message, it removes message from its scheduling (see Lamport paper) queue

5





Message Reliability ("Extended Virtual Synchrony")

Reliability

- Unreliable Message may be dropped or lost and will not be recovered.
- *Reliable* Message will be reliably delivered to all recipients who are in group to which message was sent.
- Safe The message will ONLY be delivered to a recipient if everyone currently in the group definitely has the message

Remarks

- Key issue is keeping track of reliability in groups. Reliable messages should be received by everyone (eventually).
- Requires agreement algorithm across computers (who has what)
- HW: find an example where reliable messages are not safe.











Solution using Spread

- Assume totally ordered, reliable messages ("agreed" message type)
- All processes and resource in single spread group

Algorithm

- 1. P, sends multcast message to group requesting resource
- 2. P_i queues all requests and sends ack
- 3. Process P_i uses resource when
 - P, request is at top of queue
 - · Ack has been received from everyone
- P_i sends release message when done
- 1. P_i dequeues release when message received
- 2. Note: spread provides single order

HYCON-EECI, Mar 08



13











Intoud	Usuge	
 When to use threads Main usage is when the program has to wait on a process or resource Eliminate threads if they aren't needed (eg, tight interlocking with no waits) 	 Conditional variables Allows a thread to sleep until a certain condition is met Used in conjunction with a mutex 	
 Avoiding deadlocks Never put a mutex around a call that might itself block (I/O call, mutex, etc) If you have to use nested mutex's, make sure they are in the same order whenever they are invoked Performance improvements Try to keep critical sections as small as possible (avoids excessive waiting) Combine accesses to same variables in nearby sections Use buffers to minimize lock times 	<pre>1 thread1() { 2 mutex_lock(xmtx); 3 while (!condition) 4 cond_wait(&cond, xmtx); 5 do_something(); 6 mutex_unlock(xmtx); 7 }</pre>	



PO	SIX Threads (Pthreads)	
Thread creationpthread_createpthread_exitpthread_join	call a function as a new thread of execution terminate the current thread wait for a specific thread to exit	
Mutexes		
 pthread_mutex_init 	initialize a mutex	
 pthread_mutex_lock 	lock a mutex (blocks until mutex is available)	
 pthread_mutex_unlock 	unlock a mutex (and unblock first blocked threads)	
 pthread_mutex_destroy 	free up resources associated with a mutex	
Conditional variables		
 pthread_cond_init 	initialize a condition	
 pthread_cond_wait 	wait until condition is satisfied (paired with a mutex)	
 pthread_cond_signal 	signal that a condition is now satisfied	
 pthread_cond_destroy 	free up resources associate with a mutex	
Read/write locks		
 Variation on mutexes that 	allow multiple unblocking reads	
		01



Thread scheduling policies	Homework	
 FIFO - threads are called in first in, first out order within each priority level Thread continues to run until a higher priority thread is runnable Threads at same priority must block in order for other threads to run 	 Write a simple multi-threaded program using pthreads that reads numbers from an input stream (terminal), averages all numbers that have been read, and prints out the average once a second [use three threads] 	
 Round-robin - each thread is called in sequence, within priority level Thread runs for fixed period of time 	 Project Ideas Expand sparrow to allow multi-threaded servo and channel (I/O) execution 	
before it is pre-empted	 Analyze how to best use mutex's for minimizing control latency when 	
 Other - implementation specific Operating system defines how 	reading inputs via the network and writing outputs via (slow) serial ports	
threads are scheduled This is the default (and undefined!) 	 Convert a controller from continuous to discrete for a multi-threaded control system using FIFO or round-robin scheduling 	

HYCON-EECI, Mar 08



Example: State Client	
 Asynchronously reads actuator state from adrive via thread Passes data back to calling module (eg, trajFollower) 	
<pre>void CStateClient::getActuatorStateThread() { int actuatorstatesocket = m_skynet.listen(SNactuatorstate, ALLMODULES); while(m_bRunThreads) { if(m_skynet.get_msg(actuatorstatesocket, &m_rcvdActuatorstate, sizeof(m_rcvdActuatorstate), 0, &pActuatorstateMutex) != sizeof(m_rcvdActuatorstate)) skynet_error(); DGCSetConditionTrue(condNewActuatorState);</pre>	 Thread to read msgs Infinite loop Read msg (blocks until available) Unblock anv-
<pre>} } void CStateClient::UpdateActuatorState() { DGClockMutex(&m_actuatorstateMutex); memcpy(&m_actuatorState, &m_rcvdActuatorstate, sizeof()); DGCunlockMutex(&m_actuatorstateMutex); }</pre>	 One waiting Copy state into buffer Use mutex to insure completeness
<pre>void CStateClient::WaitForNewActuatorState() { DGCWaitForConditionTrue(condNewActuatorState); UpdateActuatorState(); condNewActuatorState.bCond = false; }</pre>	 Block until new state msg arrives
HYCON-EECI, Mar 08 R. M. Murray, Caltech CDS	26

Verifying Multi-Threaded Programs

SPIN (Holzmann)

- Model system using PROMELA (Process Meta Language)
 - Asynchronous processes
 - Buffered and unbuffered message channels
 - Synchronizing statements
 - Structured data -
- Simulation: Perform random or iterative simulations of the modeled system's execution
- Verification: Generate a C program that performs a fast exhaustive verification of the system state space
- Check for deadlocks, livelocks, unspecified receptions, and unexecutable code, correctness of system invariants, non-progress execution cycles
- Also support the verification of linear time temporal constraints

HYCON-EECI, Mar 08

TLA/TLC (Lamport et al)

- Temporal Logic of Actions (TLA): Leslie Lamport, 1980's
- Behavior (a sequence of states) is described by an initial predicate and an action

Spec = Init $\land \Box$ Action

- Specify a system by specifying a set of possible behaviors
- Theorem: A temporal formula satisfied by every behavior

Theorem \equiv Spec \Rightarrow \Box Properties

TLA+

- Can be used to write a precise, formal description of almost any sort of discrete system
- Especially well suited to describing asynchronous systems
- Tools: Syntactic Analyzer, TLC model checker

Richard M. Murray, Caltech CDS

Summary: Embedded Systems Programming Feeder 1 (wait) (wait) Feeder 2 read meas ctrl write ctrl computation sionMapper Safety Mon. (wait) (wait) **Advantages Open Issues for Control Theory** · Increased modularity How do we best implement controllers in this setting? Simplified programming* How do we verify that programs Cautions satisfy the specifications and Asynchronous execution design intent Race conditions How do we implement multi-rate Deadlocking controllers using threaded process Debugging and distributed computing? HYCON-FECI Mar 08

27