



UNIVERSITY OF OREGON

CIS 415:

Operating Systems

Deadlock

Spring 2012
Prof. Kevin Butler

- **Last class:**
 - ▶ Synchronization
- **Today:**
 - ▶ Deadlocks

- A set of processes is deadlocked if each process in the set is waiting for an event that only another process in the set can cause.
- An event could be:
 - ▶ Waiting for a critical section
 - ▶ Waiting for a condition to change
 - ▶ Waiting for a physical resource

Conditions for Deadlock



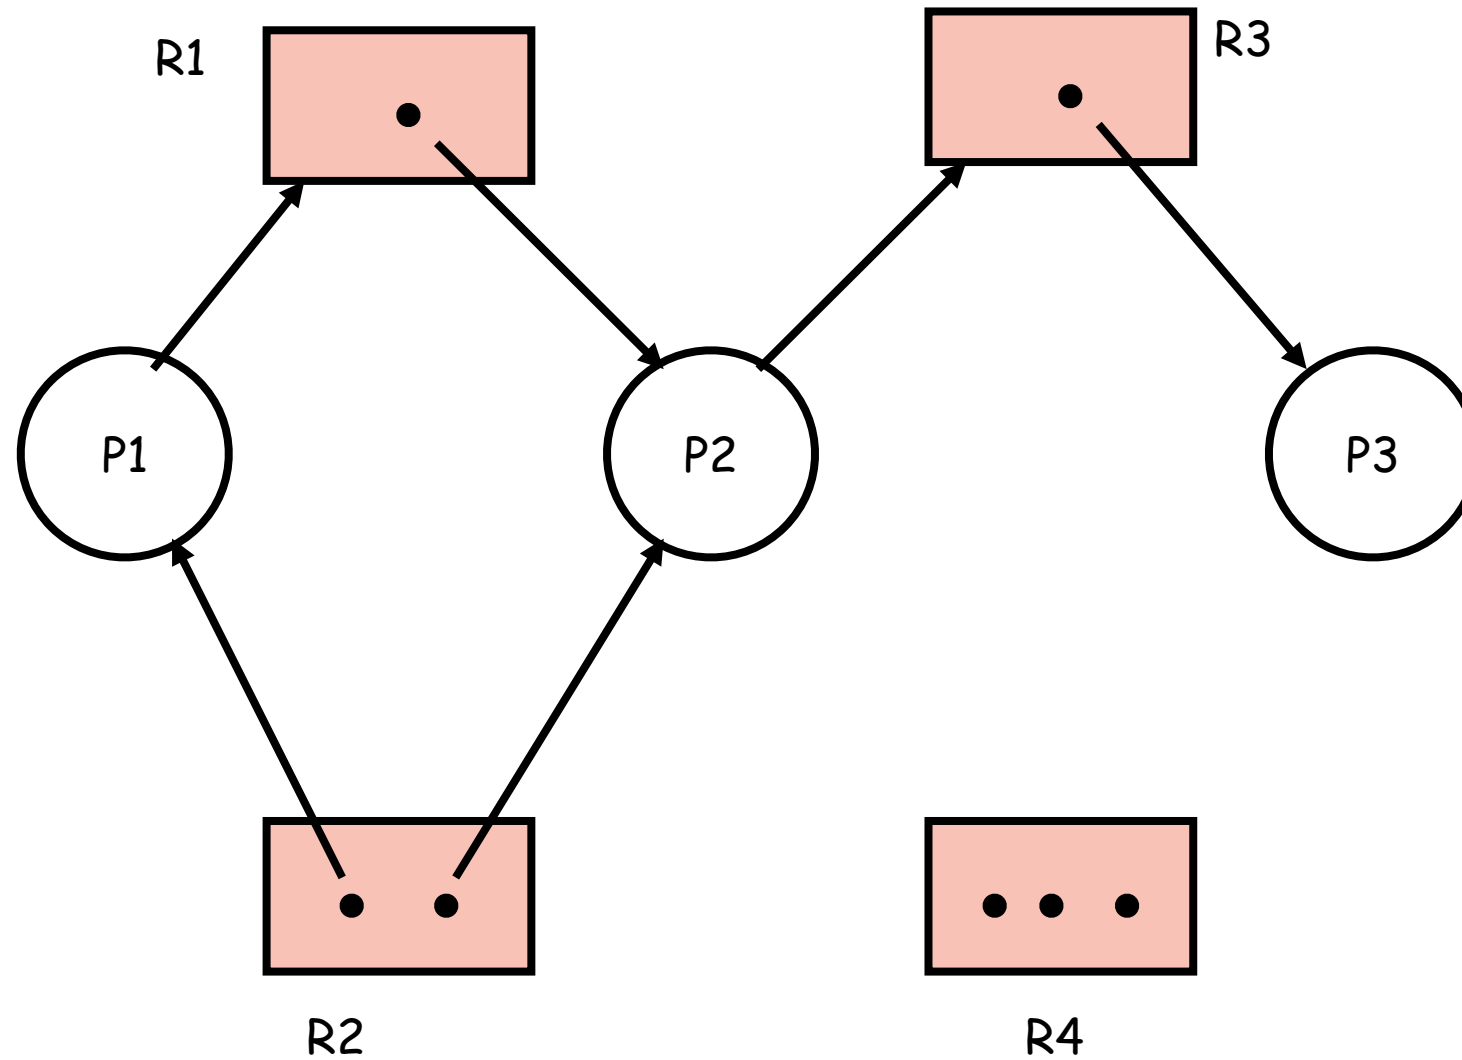
- **Mutual exclusion:** The requesting process is delayed until the resource held by another is released.
- **Hold and wait:** A process must be holding at least 1 resource and must be waiting for 1 or more resources held by others.
- **No preemption:** Resources cannot be preempted from one and given to another.
- **Circular wait:** A set (P_0, P_1, \dots, P_n) of waiting processes must exist such that P_0 is waiting for a resource held by P_1 , P_1 is waiting for ... by P_2 , ... P_n is waiting for ... held by P_0 .

Resource Allocation Graph

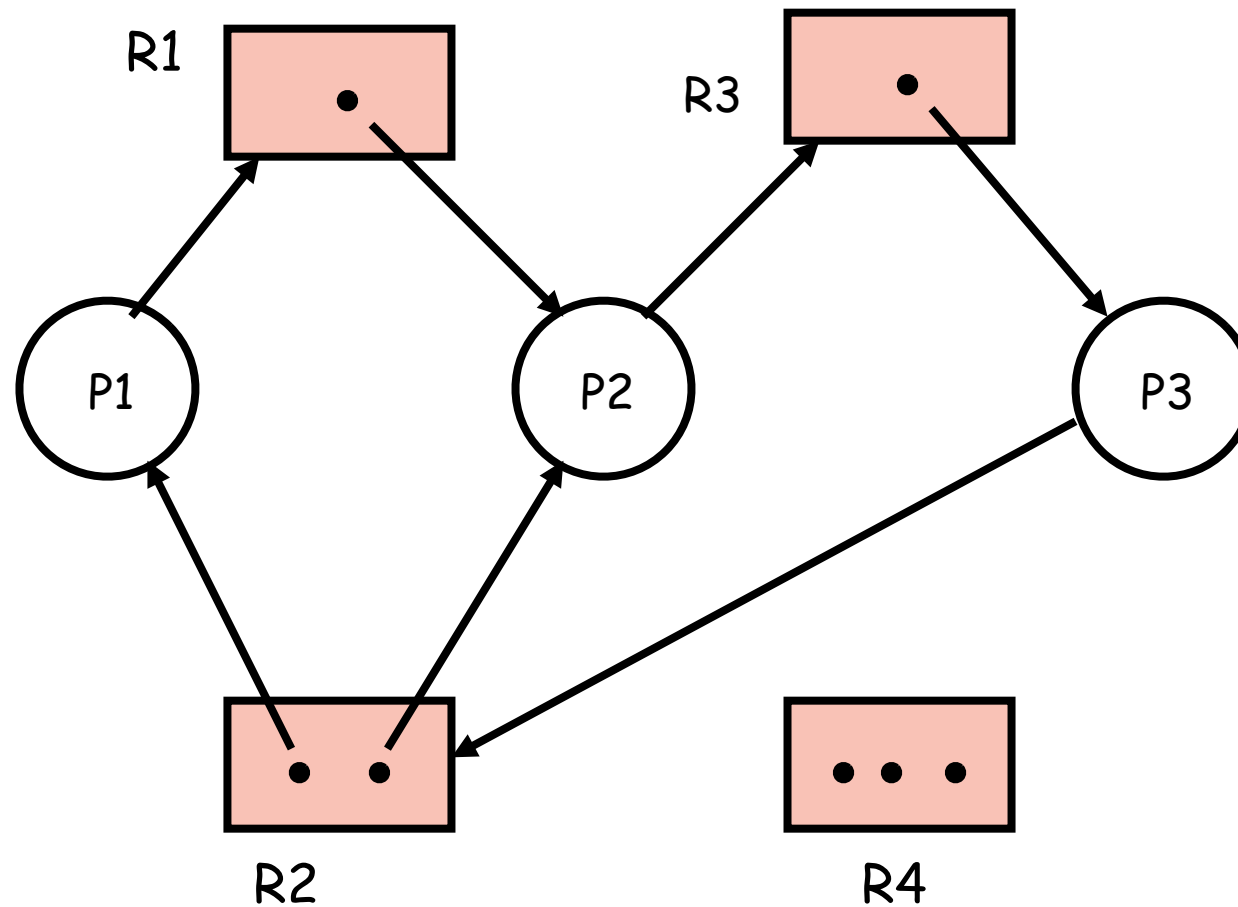


- Vertices (V) = **Processes** (P_i) and **Resources** (R_j)
- Edges (E) = **Assignments** ($R_j \rightarrow P_i$, R_j is allocated to P_i) and **Request** ($P_i \rightarrow R_j$, P_i is waiting for R_j).
- For each Resource R_j , there could be multiple instances.
- A requesting process can be granted any one of those instances if available.

An example

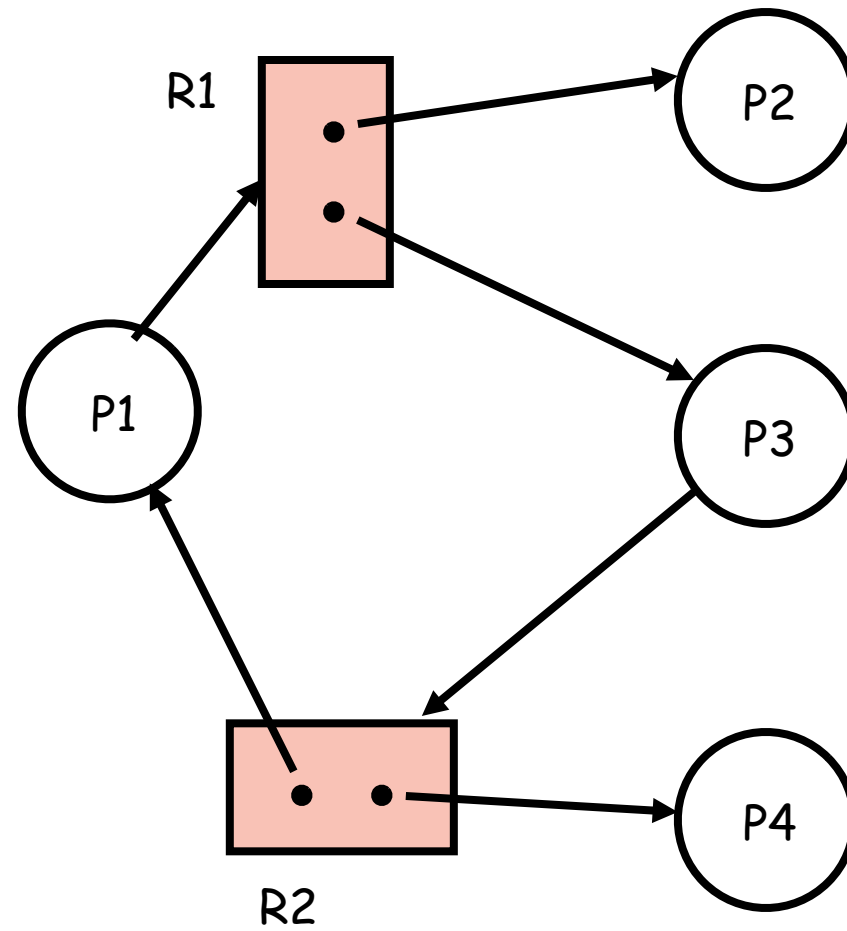


A deadlock



**If there is a deadlock, there will be a cycle
(Necessary Condition).**

Cycle is NOT sufficient



- Ignore the problem altogether (ostrich algorithm) since it may occur very infrequently, cost of detection/prevention may not be worth it.
- Detect and recover after its occurrence.
- Avoidance by careful resource allocation
- Prevention by structurally negating one of the four necessary conditions

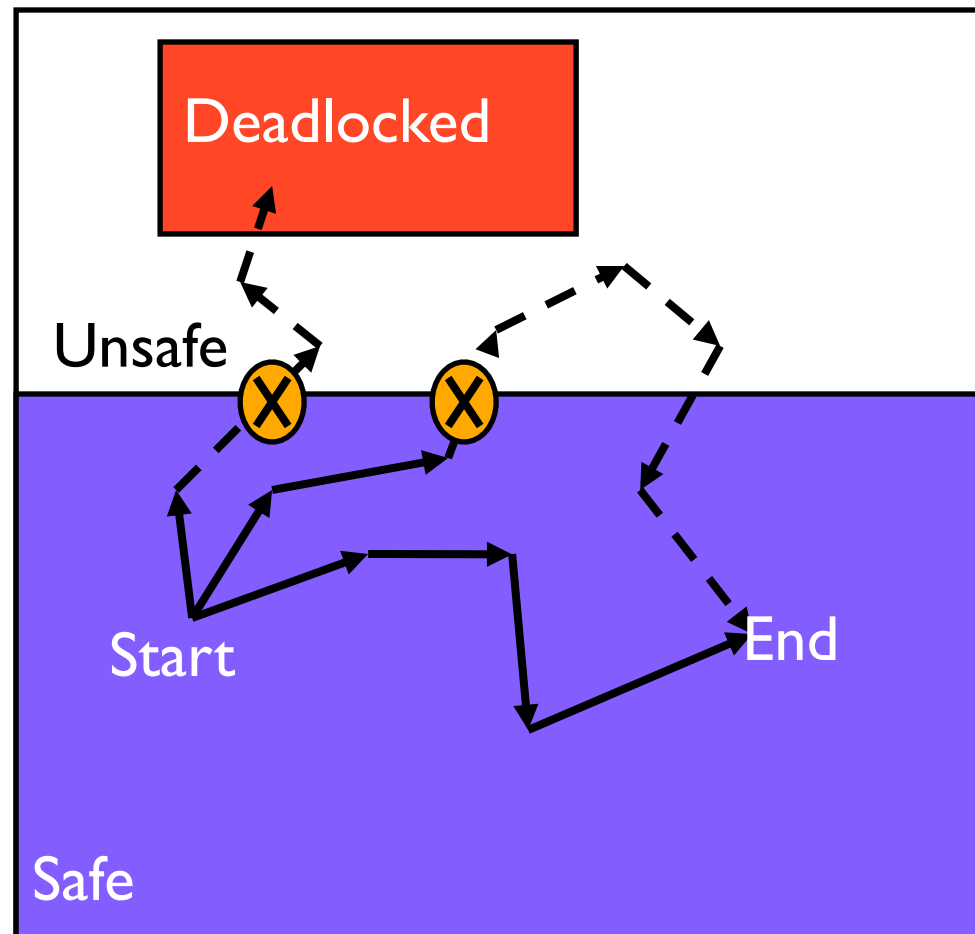
- Note that all 4 necessary conditions need to hold for deadlock to occur.
- We can try to disallow one of them from happening:
 - ▶ Mutual exclusion: This is usually not possible to avoid with many resources.
 - ▶ No preemption: This is again not easy to address with many resources. Possible for some resources (e.g. CPU)
 - ▶ Hold and Wait:
 - Allow at most 1 resource to be held/requested at any time
 - Make sure all requests are made at the same time.
 - ...

▶ Circular Wait

- Number the resources, and make sure requests are always made in increasing/decreasing order.
- Or make sure you are never holding a lower numbered resource when requesting a higher numbered resource.

- Avoid actions that may lead to a deadlock.
- Visualize the system as a state machine moving from 1 state to another as each instruction is executed.
- A state can be: *safe*, *unsafe* or *deadlocked*.
- Safe state is one where
 - it is not a deadlocked state
 - there is some sequence by which all requests can be satisfied.
- To avoid deadlocks, we try to make only those transitions that will take you from one safe state to another.
- This may be a little conservative, but it avoids deadlocks

State Transitions



1 resource with 12 units of that resource available.

Current State: Free = (12 - (5 + 2 + 2)) = 3

	Max. Needs	Currently Allocated	Still Needs
P0	10	5	5
P1	4	2	2
P2	9	2	7

Free = 3

After reducing P1,
Free = 5

After reducing P0,
Free = 10
Then reduce P2.

This state is **safe** because, there is a sequence (P1 followed by P0 followed by P2) by which max needs of each process can be satisfied. This is called the **reduction sequence**.

Unsafe State



What if P2 requests 1 more and is allocated 1 more?

New State:

	Max. Needs	Currently Allocated	Still Needs
P0	10	5	5
P1	4	2	2
P2	9	3	6

Free = 2

This is unsafe.

Only P1 can be reduced. If P0 and P2 then come and ask for their full needs, the system can become deadlocked. Hence, by granting P2's request for 1 more, we have moved from a safe to **unsafe state.**

Deadlock avoidance algorithm will NOT allow such a transition, and will not grant P2's request immediately.

- Deadlock avoidance essentially allows requests to be satisfied only when the allocation of that request would lead to a safe state.
- Else do not grant that request immediately.

Banker's algorithm



- When a request is made, check to see if after the request is satisfied, there is (at least one!) sequence of moves that can satisfy all possible requests. ie. the new state is safe.
- If so, satisfy the request, else make the request wait.

Checking for Safe State



N processes and M resources

Data Structures:

MaxNeeds[N][M];

Allocated[N][M];

StillNeeds[N][M];

Free[M];

Temp[M];

Done[N];

```
while () {
    Temp[j]=Free[j] for all j
    Find an i such that
        a) Done[i] = False
        b) StillNeeds[i,j] <= Temp[j]
    if so {
        Temp[j] += Allocated[i,j] for all j
        Done[i] = TRUE /* release Allocated[i] */
    }
    else if Done[i] = TRUE for all i then state is safe
    else state is unsafe
}
```

$M \cdot N^2$ steps to detect if a state is safe!

An example



5 processes, 3 resource types A (10 instances), B (5 instances), C (7 instances)

MaxNeeds				Allocated				StillNeeds				Free		
	A	B	C		A	B	C		A	B	C	A	B	C
P0	7	5	3	P0	0	1	0	P0	7	4	3	3	3	2
P1	3	2	2	P1	2	0	0	P1	1	2	2			
P2	9	0	2	P2	3	0	2	P2	6	0	0			
P3	2	2	2	P3	2	1	1	P3	0	1	1			
P4	4	3	3	P4	0	0	2	P4	4	3	1			

This state is safe, because there is a reduction sequence $\langle P1, P3, P4, P2, P0 \rangle$ that can satisfy all the requests.

Exercise: Formally go through each of the steps that update these matrices for the reduction sequence.

An example



If P1 requests 1 more instance of A and 2 more instances of C
can we safely allocate these? – Note these are all allocated together!
and we denote this set of requests as (1,0,2)

If allocated the resulting state would be:

MaxNeeds				Allocated				StillNeeds				Free		
	A	B	C		A	B	C		A	B	C	A	B	C
P0	7	5	3	P0	0	1	0	P0	7	4	3	2	3	0
P1	3	2	2	P1	3	0	2	P1	0	2	0			
P2	9	0	2	P2	3	0	2	P2	6	0	0			
P3	2	2	2	P3	2	1	1	P3	0	1	1			
P4	4	3	3	P4	0	0	2	P4	4	3	1			

**This is still safe since there is a reduction sequence
<P1,P3,P4,P0,P2> to satisfy all the requests. (work this out!)
Hence the requested allocations can be made.**

An example



After this allocation, P0 then makes a request for (0,2,0).
If granted the resulting state would be:

MaxNeeds				Allocated				StillNeeds				Free		
	A	B	C		A	B	C		A	B	C	A	B	C
P0	7	5	3	P0	0	3	0	P0	7	2	3	2	1	0
P1	3	2	2	P1	3	0	2	P1	0	2	0			
P2	9	0	2	P2	3	0	2	P2	6	0	0			
P3	2	2	2	P3	2	1	1	P3	0	1	1			
P4	4	3	3	P4	0	0	2	P4	4	3	1			

This is an UNSAFE state.

So this request should NOT be granted.

- Ignore the problem altogether (ostrich algorithm) since it may occur very infrequently, cost of detection/prevention may not be worth it.
- • Detect and recover after its occurrence.
- Avoidance by careful resource allocation
- Prevention by structurally negating one of the four necessary conditions

- If there is only 1 instance of each resource, then a cycle in the resource-allocation graph is a “sufficient” condition for a deadlock, i.e. you can run a cycle-detection algorithm to detect a deadlock.
- With multiple instances of each resource, ???

Detection Algorithm

N processes, M resources

Data structures:

```
Free[M];
Allocated[N][M];
Request[N][M];
Temp[M];
Done[N];
```

1. `Temp[i] = Free[i]` for all `i`
`Done[i] = FALSE` unless there is
no resources allocated to it.
2. Find an index `i` such that both
 - (a) `Done[i] == FALSE`
 - (b) `Request[i] <= Temp` (vector comp.)If no such `i`, go to step 4.
3. `Temp = Temp + Allocated[i]` (vector add)
`Done[i] = TRUE; /* release Allocated[i] */`
Go to step 2.
4. If `Done[i]=FALSE` for some `i`, then
there is a deadlock.

Basic idea is that there
is at least 1 execution
that will unblock all
processes.

M*N² algorithm!

Example



5 processes, 3 resource types A (7 instances), B (2 instances), C (6 instances)

Allocated			
	A	B	C
P0	0	1	0
P1	2	0	0
P2	3	0	3
P3	2	1	1
P4	0	0	2

Request			
	A	B	C
P0	0	0	0
P1	2	0	2
P2	0	0	0
P3	1	0	0
P4	0	0	2

Free		
A	B	C
0	0	0

This state is NOT deadlocked.

By applying algorithm, the sequence $\langle P0, P2, P3, P1, P4 \rangle$ will result in Done[i] being TRUE for all processes.

- Once deadlock is detected what should we do?
 - ▶ Preempt resources (whenever possible)
 - ▶ Kill the processes (and forcibly remove resources)
 - ▶ Checkpoint processes periodically, and roll them back to last checkpoint (relinquishing any resources they may have acquired since then).

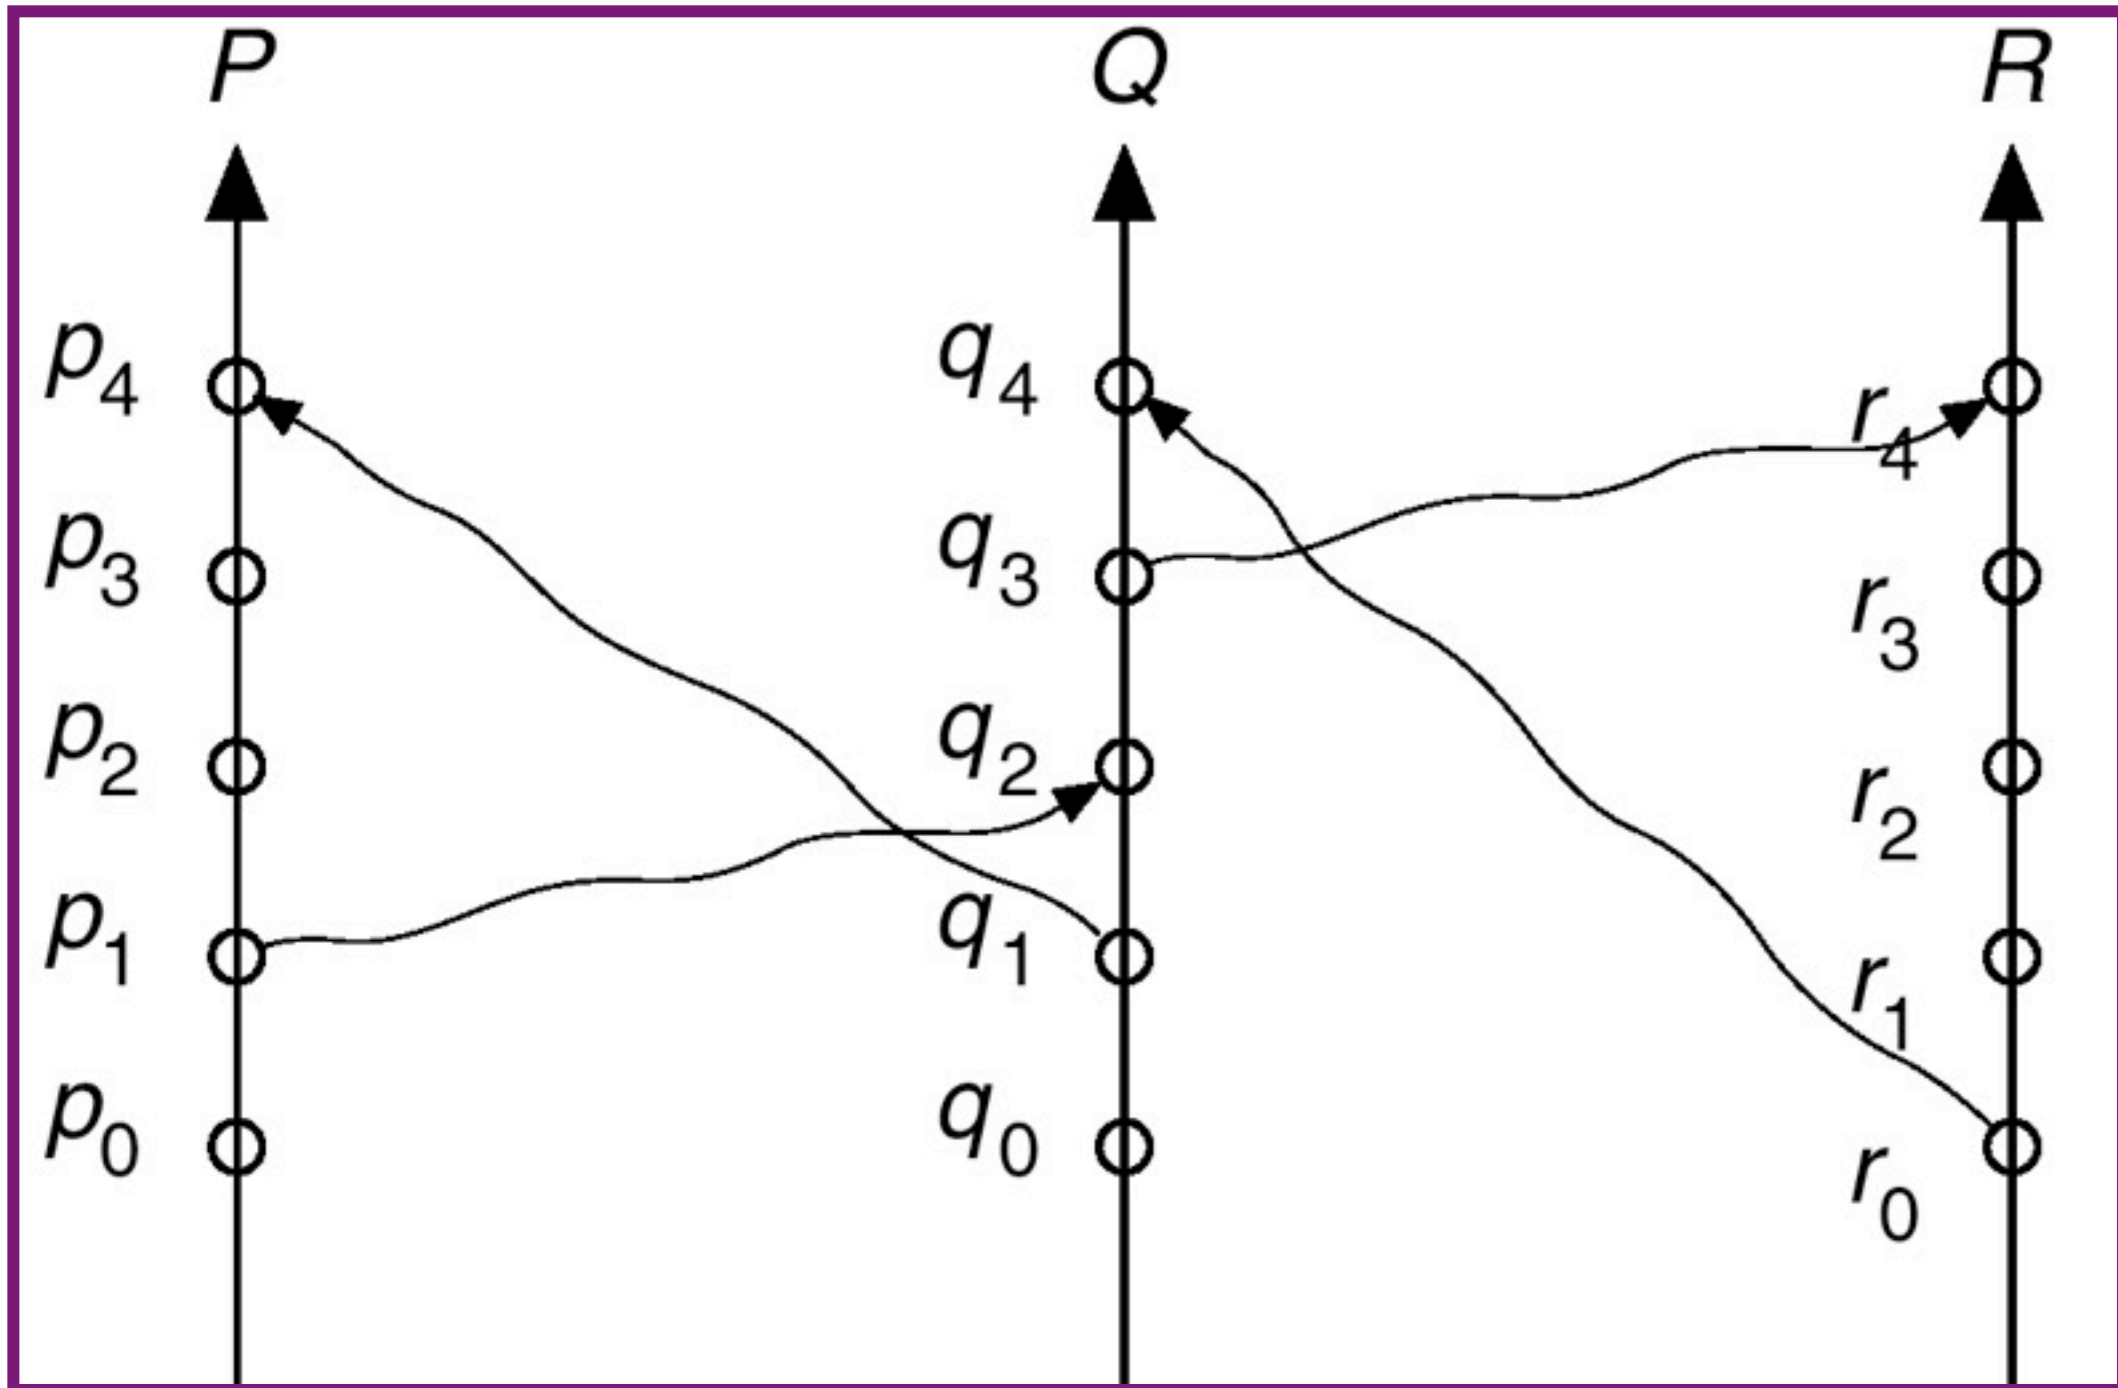
- To date: we've thought about how to order events in order to provide synchronization and prevent deadlock
- What have we been relying on in order to get ordering?
 - ▶ A consistent clock amongst processes
- What happens when processes aren't sharing the same clock?
 - ▶ Distributed system



- Without sharing a clock, it's not possible to get a total ordering over events
 - ▶ Instead, we get partial ordering
- Within a sequential process, all events are executed in a totally ordered fashion
- Message can only be received after it's sent
- *Happened-before* relation → reflects partial ordering

- Properties of the happened-before relation
 - ▶ If A, B are events in the same process and A executes before B then $A \rightarrow B$
 - ▶ If A is a send message event from a process and B is a receive message event from another process then $A \rightarrow B$
 - ▶ If $A \rightarrow B$ and $B \rightarrow C$ then $A \rightarrow C$ (*what property is this?*)
- If events A and B are not related by \rightarrow then they can execute *concurrently* (no effect of A on B)

Concurrent vs \rightarrow



- Associate a timestamp with each system event
 - ▶ Require that for every pair of events A and B , if $A \rightarrow B$, then the timestamp of A is less than the timestamp of B
- Associate *logical (Lamport) clock* LC_i with process P_i
 - ▶ Implement as a simple counter incremented between any two successive events executed within a process.
- Process advances logical clock when receiving message with timestamp $>$ current value of LC
- If $TS_A == TS_B$, events are concurrent (use process ID to break ties and create a total ordering)

- **Deadlocks**
 - ▶ **Necessary and sufficient conditions**
 - Resource allocation graph
 - ▶ **Strategies**
 - Ignore
 - Prevention
 - ▶ Safe States
 - Avoidance
 - Detection and recovery
 - Distributed ordering

- **Exam Structure**
 - ▶ **(12) Short Answer (1-3 sentences)**
 - 3-4 pts each
 - ▶ **(4) Long Answer (2 paragraphs max)**
 - 7 pts each
 - ▶ **(3) Constructions (several related, small questions)**
 - 10-12 pts each

- **Exam Structure**
 - ▶ **(12) Short Answer (1-3 sentences)**
 - How and what questions
 - ▶ How does X work?
 - ▶ What is Y?
 - ▶ **(4) Long Answer (2 paragraphs max)**
 - How and why questions
 - ▶ How and why does X work that way?
 - ▶ The 'why' may be implicit, but do not assume that I know that you know how these work
 - ▶ **(3) Constructions (several related, small questions)**
 - Specific questions about OS mechanisms/concepts

- Scope
 - ▶ About 1/3 to 1/2 is related to HWs/quizzes/project
 - ▶ Test covers all topics that we discussed in class
 - But, not all the answers are specified in the slides
 - ▶ And related sections in book
 - We followed pretty closely in Chs. 3-5
 - ▶ Hopefully, your notes are good (or you have a good memory of what we discussed)

- **Scope**
 - ▶ Chapter 1-5, some chapter 6, possibly a bit on chapter 7 (also a little on chapter 16 and parts of chapter 21)
 - ▶ More emphasis on chapters 3-5 (others were review or pretty new)

- Chapter I
 - ▶ Hardware concepts
 - CPU
 - ▶ Internals
 - Memory
 - ▶ Memory hierarchy
 - I/O devices
 - ▶ Interaction
 - ▶ Communication mechanisms
 - Interrupts
 - Bus

- Chapter 2
 - ▶ OS structure
 - ▶ What is the OS?
 - Functions
 - ▶ OS API
 - System call processing
 - Process and file system calls
 - (we'll come back to mmap later)
 - ▶ Process structure (address space)
 - ▶ OS structures
 - Monolithic and microkernel

- Chapter 3 -- Processes
 - ▶ Process Structure (Address space)
 - Process creation (fork/exec)
 - Process loading (executables and libraries)
 - ▶ Process representation in kernel (structure)
 - Context switch
 - Hierarchy
 - ▶ Process states
 - ▶ Interprocess communication
 - Shared memory
 - Message passing
 - In detail -- studying actual systems will help understand the concepts
 - ▶ Remote procedure calls

- Chapter 4 -- Threads
 - ▶ Purpose of threads
 - ▶ Threading models
 - ▶ Thread context switch
 - ▶ Thread system issues
 - ▶ Threading system basics
 - Clone system call
 - Linux threads
 - Pthreads invocation, termination

- Chapter 5 -- Scheduling
 - ▶ Concepts
 - Bursts, preemption, basic criteria
 - ▶ I/O bound and CPU bound processes
 - ▶ Algorithms
 - FCFS, SJF, RR, priority
 - Exponential average
 - Multiqueue scheduling (combinations of above)
 - ▶ Study Linux/Solaris to better understand scheduling

- Chapter 6 -- Synchronization
 - ▶ Concepts
 - Concurrency, critical section problem & requirements
 - ▶ S/W and H/W approaches to synchronization
 - ▶ Semaphores and Monitors
 - ▶ Synchronization Problems
 - Bounded-buffer
 - Readers-Writers
 - Dining Philosophers