



UNIVERSITY OF OREGON

CIS 415:
Operating Systems
Final Review

Spring 2014
Prof. Kevin Butler

- Friday, June 13 at 8 AM
- 2 hours
- Structure: similar to previous exam in terms of question layout, may be similar length or possibly a bit longer
- Questions:
 - ▶ some technical details
 - ▶ some conceptual questions
- Material: slides, text, class discussion, homeworks

What we've covered



- Everything from the first part of the course is fair game
 - ▶ but probably de-emphasized a bit - bit twiddling probably not required but have a strong idea of the concepts
- Chapter 6 -- Synchronization
- Chapter 7 -- Deadlocks
- Chapter 8 -- Main Memory (Physical)
- Chapter 9 -- Virtual Memory
- Chapter 10 -- File System Interface & Chapter 14 -- Protection
- Chapter 11 -- File System Implementation
- Chapter 12 -- Storage
- Chapter 21 - Linux

- Problems
- Synchronization Requirements
- Disabling Interrupts
- Busy-wait/Spinlock solutions
 - ▶ Related to properties
- Hardware-enabled Solutions
- OS-supported

Requirements for Solution



- 1. Mutual Exclusion - If process P_i is executing in its critical section, then no other processes can be executing in their critical sections
- 2. Progress - If no process is executing in its critical section and there exist some processes that wish to enter their critical section, then the selection of the processes that will enter the critical section next cannot be postponed indefinitely
- 3. Bounded Waiting - A bound must exist on the number of times that other processes are allowed to enter their critical sections after a process has made a request to enter its critical section and before that request is granted
 - ▶ Assume that each process executes at a nonzero speed
 - ▶ No assumption concerning relative speed of the N processes

- Hardware Enabled Solutions
- OS-supported Solutions
 - ▶ Mutex
 - ▶ Semaphores
 - ▶ Condition Variables
- Apply these to code
- Classic Synchronization Problems

- You are given a data-type Semaphore_t.
- On a variable of this type, you are allowed
 - ▶ P(Semaphore_t) -- wait
 - ▶ V(Semaphore_t) – signal
- Intuitive Functionality:
 - ▶ Logically one could visualize the semaphore as having a counter initially set to 0.
 - ▶ When you do a P(), you decrement the count, and need to block if the count becomes negative.
 - ▶ When you do a V(), you increment the count and you wake up 1 process from its blocked queue if not null.

- Necessary Conditions
- Safe States
- Resource Allocation Graph
- Deadlock Prevention
 - ▶ Safe States
- Deadlock Detection
 - ▶ Detection Algorithm
 - ▶ Recovery

- **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 .

Deadlock Prevention 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.

Deadlock Detection 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

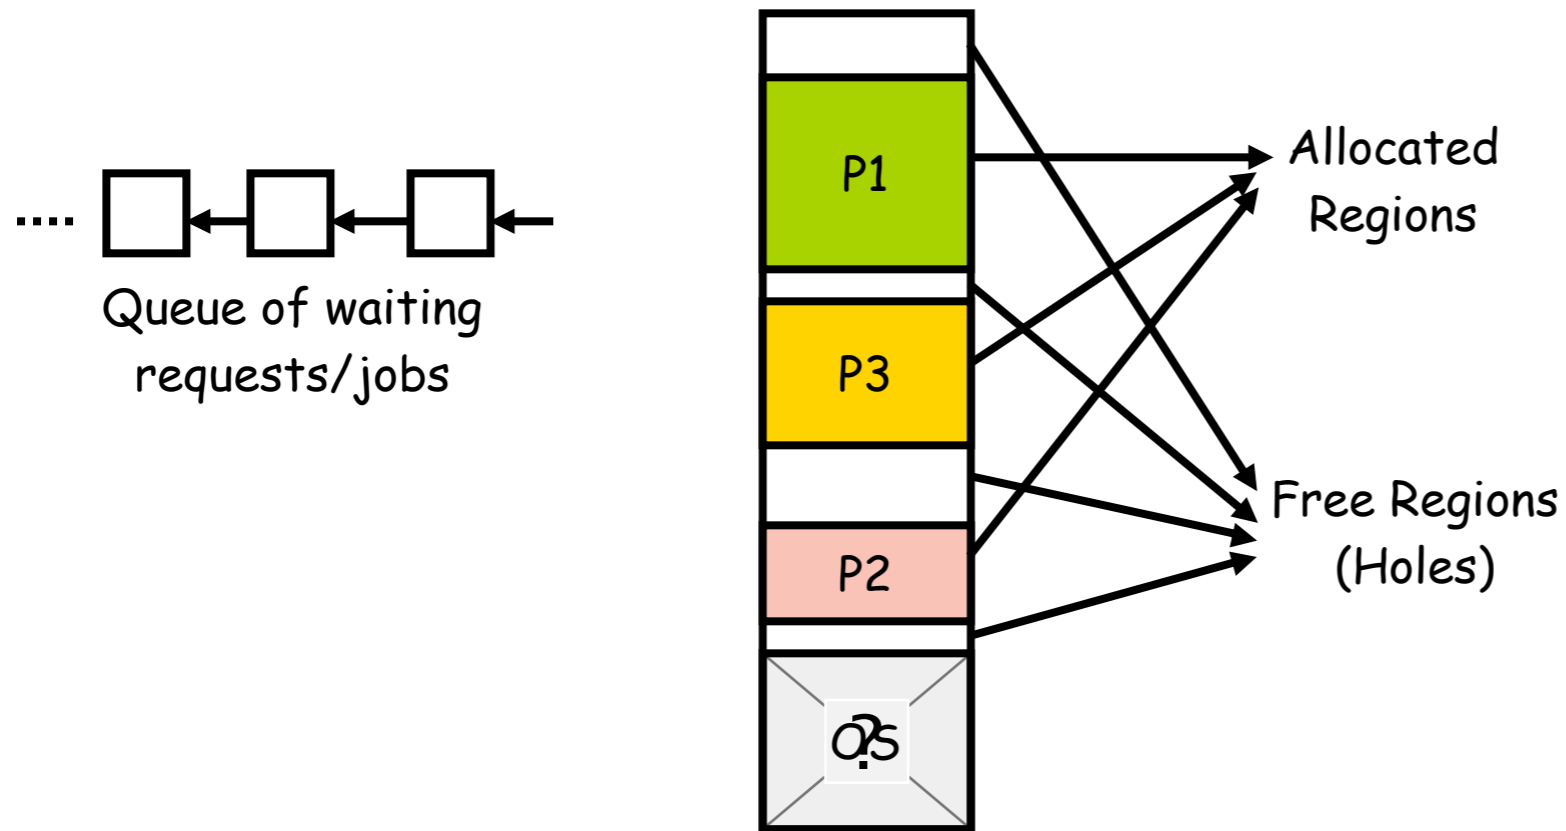
Free		
A	B	C
0	0	0

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.

- Swapping
- Allocation
 - ▶ Contiguous, Non-contiguous (paging)
 - ▶ Algorithms
- Fragmentation
 - ▶ Internal, External
- Page-tables, TLBs
 - ▶ virtual-physical translation
 - ▶ Page table structure, entries

Memory Allocation



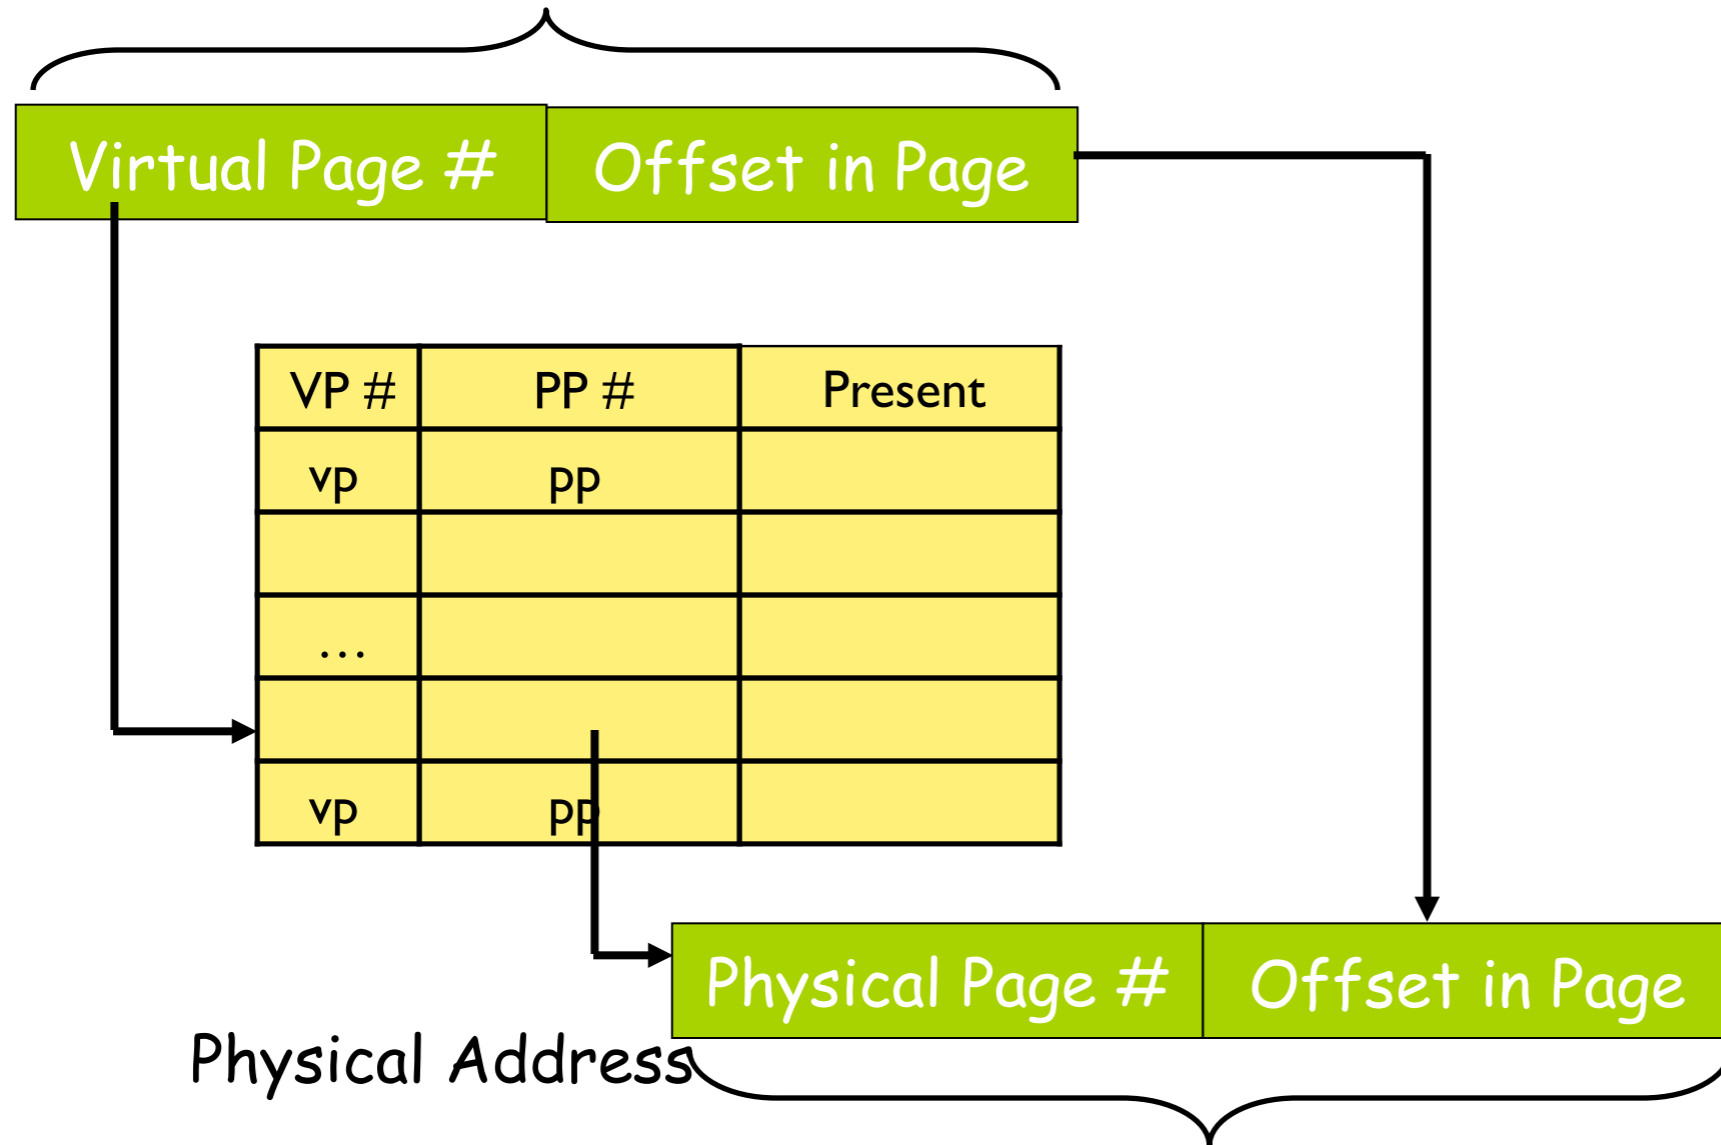
Question: How do we perform this allocation?

- Programs are provided with a virtual address space (say 1 MB).
- Role of the OS to fetch data from either physical memory or disk.
 - ▶ Done by a mechanism called (demand) paging.
- Divide the virtual address space into units called “virtual pages” each of which is of a fixed size (usually 4K or 8K).
 - ▶ For example, 1M virtual address space has 256 4K pages.
- Divide the physical address space into “physical pages” or “frames”.
 - ▶ For example, we could have only 32 4K-sized pages.

Page Tables



Virtual Address




- Page Fault Handling
 - ▶ Performance Estimations
- Memory Initialization
- Page Replacement
 - ▶ Algorithms
 - ▶ Belady's Anomaly
- Uses of Virtual Memory
 - ▶ COW, Shared Pages, Memory-mapped Files
- Thrashing



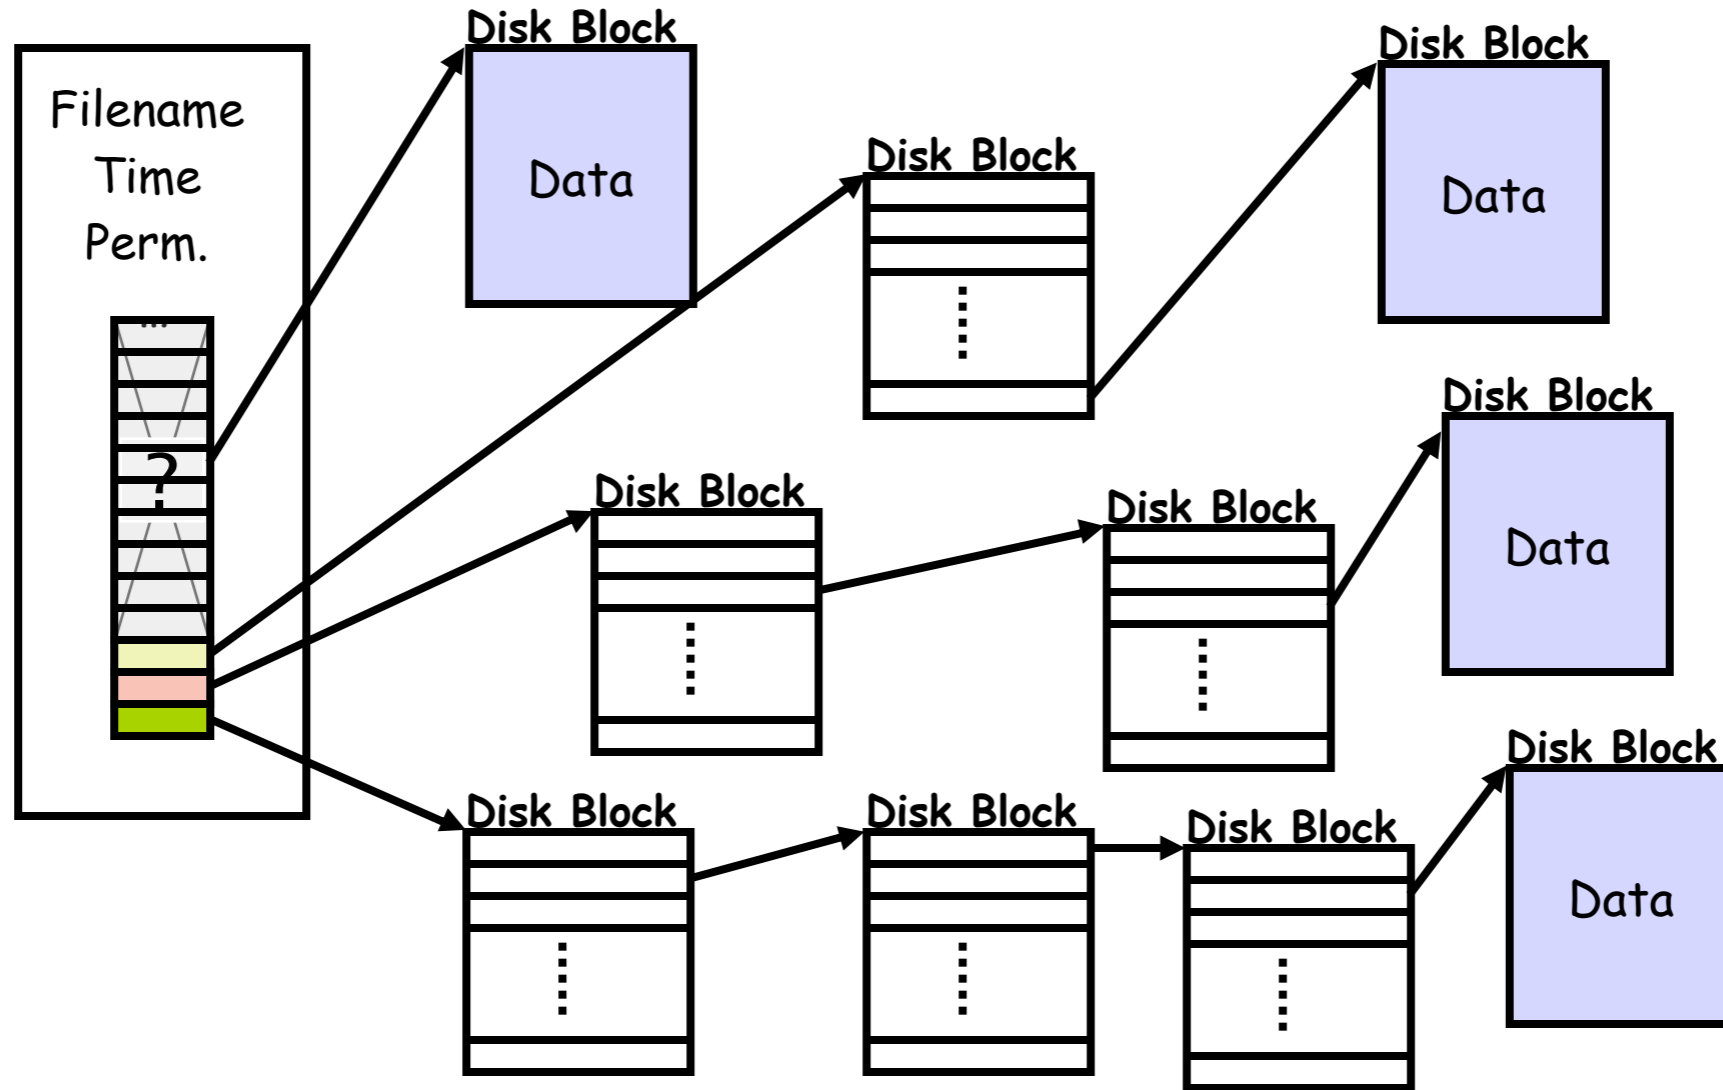
- If there is a reference to a page, first reference to that page will trap to operating system:
 - ▶ page fault
- Operating system looks at another table to decide:
 - ▶ Invalid reference -- abort
 - ▶ Just not in memory
- Get empty frame
- Swap page into frame
- Reset tables
- Set validation bit = v
- Restart the instruction that caused the page fault

- File System Concepts
 - ▶ Files, Directories, File Systems
 - ▶ Operations and Usage
 - ▶ Remote File Systems
- File System Implementation
 - ▶ What's on the disk? How's it formatted?
 - ▶ What's in memory? How's it represented?
- File System Usage
 - ▶ Get a file
 - ▶ Caching
 - ▶ Free Space
 - ▶ Recovery

File System Mounting



i-node



Access Control: Mode Bits

- Three classes of users: public, group, owner
- Three types of access permissions:
 - ▶ read, write, execute

- Example:

	Owner	Group	Public
<code>rwX=</code>	111	101	101
Octal	7	5	5

What if no exec access and only owner can read/write?

- An ***access control*** system determines what ***rights*** a particular ***entity*** has for a set of ***objects***
- It answers the question
 - ▶ E.g., do ***you*** have the right to ***read /etc/passwd***
 - ▶ Does ***Alice*** have the right to ***view*** the ***CIS website?***
 - ▶ Do ***students*** have the right to ***share project data?***
 - ▶ Does ***Prof. Butler*** have the right to ***change*** your ***grades?***
- An ***Access Control Policy*** answers these questions

- Disk scheduling algorithms
- Access and transfer time (and their components)
- Real schedulers and their differences
- Buses
- DMA

- What is the future of operating system design and research?
 - ▶ Mobile phones (Android is a big research platform right now)
 - ▶ Clouds (Amazon AWS and other services)
 - ▶ Both cases: distributed operation is becoming ever more critical
 - IPC within system and to other components that comprise logical systems
 - ubiquitous computing and prevalent network connectivity
- **SECURITY**

- If you are interested in what you've learned in this class and want to consider learning more about systems concepts, you may also think about the following courses:
 - ▶ CIS 432: Networking
 - ▶ CIS 433: Computer and Network Security
- Also: seminars and reading groups
- Think about your future and what you want
 - ▶ Take advantage of resources you have at your disposal while you're a student

Thanks to...



- Colleagues here at UO, the Pennsylvania State University, the University of Pennsylvania, North Carolina State University, and Columbia University
 - ▶ provided basis for many course materials and course ideas
 - ▶ special thanks: Trent Jaeger and Adam Aviv
- Amir Farzad and Zelong Li: GTF and grader
- Braden Hollembaek: course oracles
- You! Hope you learned something from it all

- Have a great summer! Good luck next Friday!