JEOPARDY!

Midterm 2 Review

By your lovely TAs

JEOPARDY BOARD

FINAL JEOPARDY

Midterm 1	Memory Management	Misc	File System	Concurrency
$100	$100	$100	$100	$100
$200	$200	$200	$200	$200
$300	$300	$300	$300	$300
$400	$400	$400	$400	$400
$500	$500	$500	$500	$500

Midterm 1 - $100 Question

An advantage of threads over processes

Click to see answer

Midterm 1 - $100 Answer

Shared memory, no context switching

Click to return to Jeopardy Board

Midterm 1 - $200 Question

Makes an Array faster than a linked list

Click to see answer 

Midterm 1 - $200 Answer

Caching

Click to return to Jeopardy Board

Midterm 1 - $300 Question

Best algorithm for dealing with page faults

Click to see answer

Midterm 1 - $300 Answer

LRU (Least Recently Used)

Click to return to Jeopardy Board

Midterm 1 - $400 Question

A 32 bit system with byte addressability and page size with 2048 and 2 GiB of physical memory. What is the page offset bits, physical page number, and physical address

Click to see answer

Midterm 1 - $400 Answer

Virtual page: 21 bits
Page offset: 11 bits
Physical Page Number: 20 bits
Physical Address: 31 bits

Click to return to Jeopardy Board

Midterm 1 - $500 Question

What is the difference between TLB and L1 Cache?

Click to see answer

Midterm 1 - $500 Answer

TLB is for speeding up address translation (don't need to use page table)
Cache is for speeding up main memory access (no need to access RAM)

Click to return to Jeopardy Board

Memory Management - $100 Question

What are some example I/O devices?

Click to see answer

Memory Management - $100 Answer

Keyboard, mouse, disk, etc.

Click to return to Jeopardy Board

Memory Management - $200 Question

If there are 2048 addresses in a page, how many bits are required to specify one such address?

Click to see answer 

Memory Management - $200 Answer

11

Click to return to Jeopardy Board 

Memory Management - $300 Question

How does the multi level page table take advantage of both spatial locality?

Click to see answer 

Memory Management - $300 Answer

If we access a page, most likely that we will access pages and memory addresses that are adjacent to each other

Click to return to Jeopardy Board

Memory Management - $400 Question

Rank the following in terms of access speed: (1) RAM (2) Disk (3) CPU registers (4) L1 Cache (5) L3 Cache

Click to see answer 

Memory Management - $400 Answer

Speed-wise: CPU Registers > L1 Cache > L3 Cache > RAM > Disk

Click to return to Jeopardy Board

Memory Management - $500 Question

You would be better off not using buffered reads/writes if you want better…
(2 main things)

Click to see answer

Memory Management - $500 Answer

1) Reliability – buffer isn't actually written until flushed
2) Performance – buffer takes time to copy

Click to return to Jeopardy Board

Misc - $100 Question

What Scheduling Protocol is the PennOS Scheduler most like?

Click to see answer

Misc - $100 Answer

Round Robin

Click to return to Jeopardy Board

Misc - $200 Question

Why is atomicity important in concurrency?

Hint: What if TSL was not atomic?

Click to see answer 

Misc - $200 Answer

Atomicity prevents other operations from interrupting the operation.

Click to return to Jeopardy Board

Misc - $300 Question

What is Happening at Each Point?

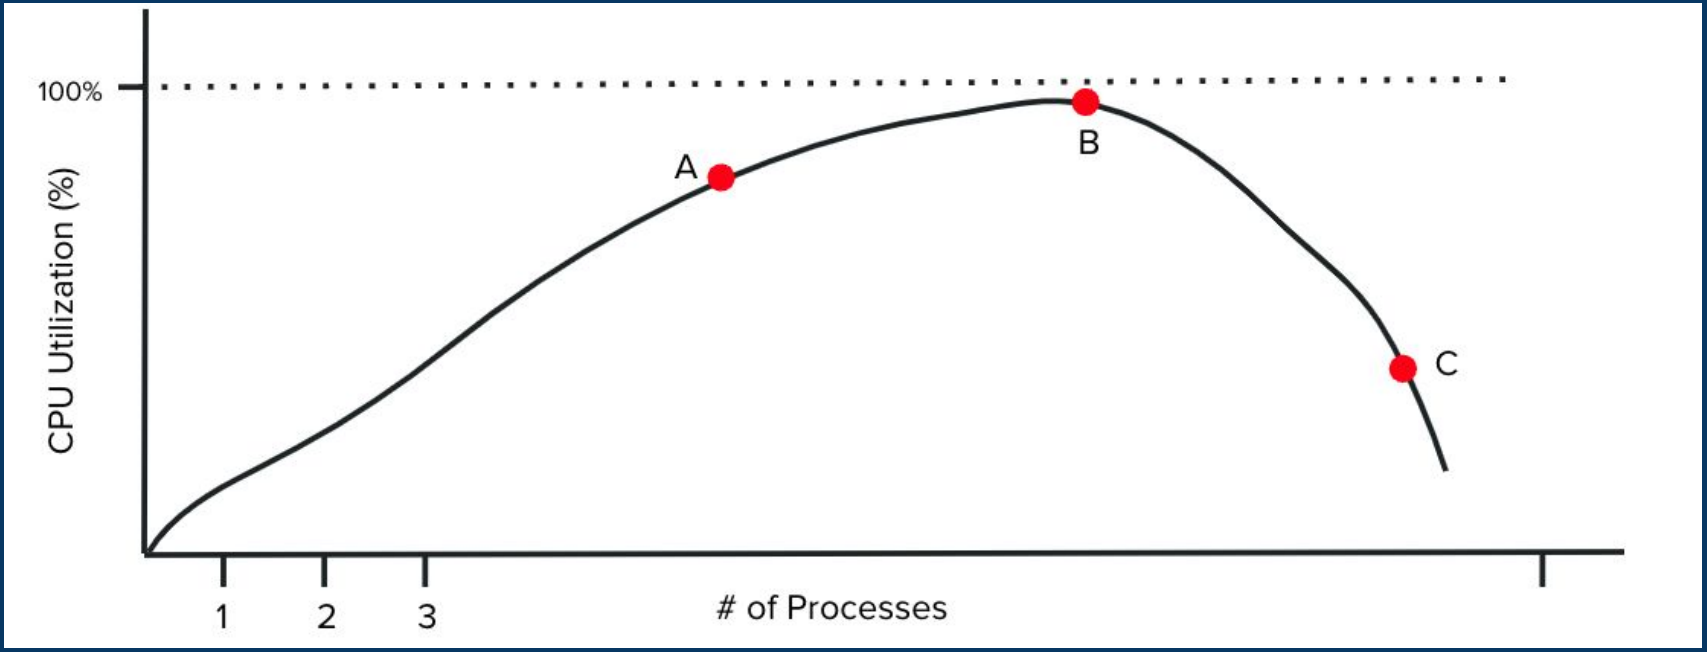

Click to see answer

Misc - $300 Answer

A - Context Switching is a problem

B - We are starting to access the disk

C - Every context switch needs access to disk memory

Click to return to Jeopardy Board

Misc - $400 Question

True or False
"The scheduling of both kernel level threads and user level threads can be customized by the user application"

Click to see answer 

Misc - $400 Answer

False, Kernel level threads need to use OS specific scheduling. User level threads can be customized by application

Click to return to Jeopardy Board

Misc - $500 Question

How many hi?

```
int main() {
  for (i is [1,3]) {
    pid = fork();
   if (pid == 0) print("hi\n");
  }
}
```

Click to see answer

Misc - $500 Answer

11. Due to stdout buffering

Follow up. How can we change this so we get 7 hi?

Click to return to Jeopardy Board

File System - $100 Question

What does RAID stand for?

Click to see answer 

File System - $100 Answer

Redundant Array of Independent Disks

Click to return to Jeopardy Board

File System - $200 Question

Why are there only N-1 blocks that we can store actual file data in a FAT when the FAT table has N entries?

Click to see answer 

File System - $200 Answer

The first entry of the FAT Table (FAT[0]) holds metadata about the FAT and does not correspond with a data block.

Click to return to Jeopardy Board

File System - $300 Question

A FAT Table file system generally suffers from what type of Fragmentation?

Click to see answer 

File System - $300 Answer

Internal Fragmentation

Click to return to Jeopardy Board

File System - $400 Question

Why are Inodes generally better than FAT in terms of Memory Usage?

Click to see answer

File System - $400 Answer

You only need to store in memory the inodes of the open files instead of the whole FAT

Click to return to Jeopardy Board

File System - $500 Question

A FAT spans 8 blocks, the block size is 8 kilobytes, and each FAT entry is 2 bytes. How many entries are in the FAT Table?

Click to see answer

File System - $500 Answer

8 * 2^13 / 2 = 2^15 entries

Click to return to Jeopardy Board

Concurrency - $100 Question

This occurs when 2+ threads access the same memory location concurrently and at least 1 thread is writing

Click to see answer 

Concurrency - $100 Answer

Data race

Click to return to Jeopardy Board

Concurrency - $200 Question

This snippet best demonstrates what synchronization tool?

```
consumer() {
  lock(&mutex);
  while (arr.size == 0) {
    cond_wait(&cond, &mutex);
  }
  arr.consume()
  unlock(&mutex);
}
```

```
producer() {
  lock(&mutex);
  arr.produce();
  cond_signal(&cond);
  unlock(&mutex);
}
```

Click to see answer

Concurrency - $200 Answer

Condition variable

Click to return to Jeopardy Board

Concurrency - $300 Question

If we detect a ____ in the graph, there's a deadlock

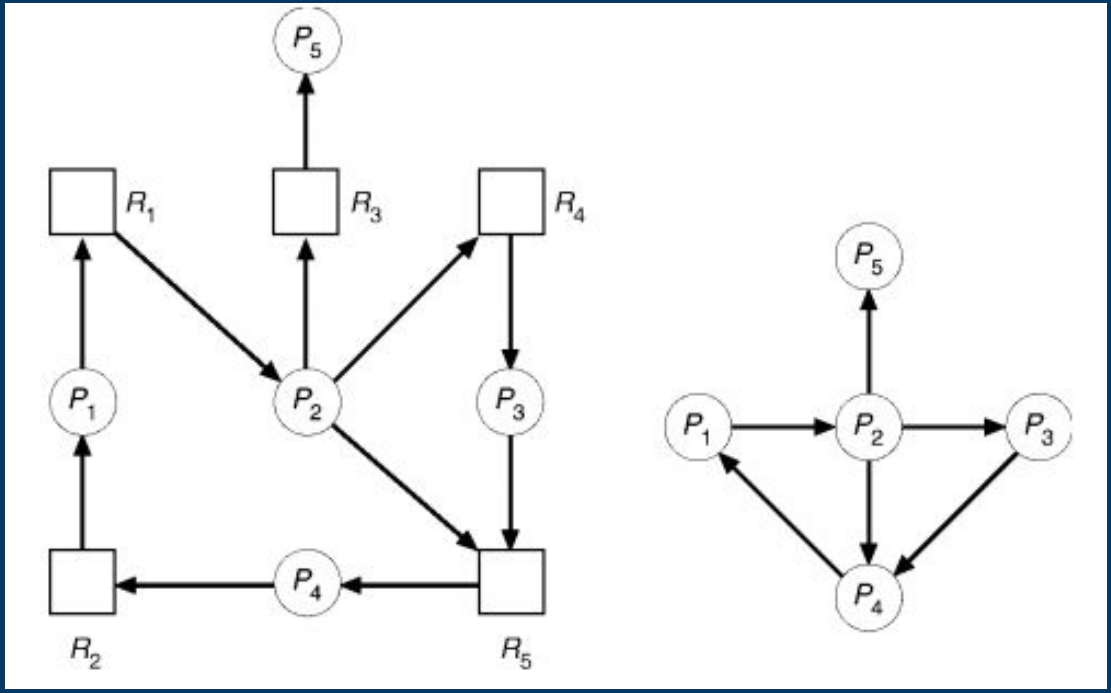

Click to see answer

Concurrency - $300 Answer

Cycle

Click to return to Jeopardy Board

Concurrency - $400 Question

If 40% of a program can be parallelized, what's the theoretical max speedup possible (vs sequential)?

$$\frac{T_1}{T_p} = \frac{1}{1-S+\frac{S}{P}}$$

Click to see answer

Concurrency - $400 Answer

~1.667x (or 5/3)

T1/TP = 1 / (1 - 0.4 + 0)

= 1 / ⅗ = 5/3

$$\frac{T_1}{T_p} = \frac{1}{1-S+\frac{S}{P}}$$

Click to return to Jeopardy Board

Concurrency - $500 Question

In high-traffic web servers, a pool of threads is ***reused*** for new client connections. This reduces overhead of what two operations?

Click to see answer 

Concurrency - $500 Answer

Thread creation + destruction

Click to return to Jeopardy Board

FINAL

JEOPARDY!

Topic: Concurrency

Click to see question 

Final Jeopardy Question

Consider one solution to the Dining Philosopher's problem (where each philosopher grabs the left chopstick followed by the right chopstick before eating). Explain how process priorities (to remove the "resource preemption" pre-condition for deadlocks) can be used to eliminate deadlocks if applied to this problem. Do not worry about the level of concurrency for this problem.

Click to see answer 

Final Jeopardy Answer

We can modify the dining philosopher's problem by ordering philosophers from 0 to n-1 by priority (strict ordering), using the philosopher identifier for ordering purposes. When a philosopher with higher priority tries to get a chopstick owned by a lower priority philosopher, the lower priority philosopher will release its chopstick for the higher priority process. Hence, there is resource preemption as resource holders may be made to relinquish their resources to other processes.

Click to return to Jeopardy Board