

CS 537

Section 7

Programming Assignment 4

Michael Swift

© 2004-2007 Ed Lazowska, Hank Levy, Andrea and
Remzi Arpaci-Dusseau, Michael Swift

1

Quiz 5 answers

- Making threads be created one at a time
 - Call `pthread-join(&threads[i]);`
 - Call `sem_wait()` in `main()`, `sem_signal()` in `print_hello`
- Fixing locks:
 - No need to spin;
 - Initialize mutex to `FALSE`
 - Release mutex before returning
- More locking
 - Causes deadlock

© 2004-2007 Ed Lazowska, Hank Levy, Andrea and
Remzi Arpaci-Dusseau, Michael Swift

2

P4: File System Measurement

- Due date: Tuesday 12/15
- Goal: measure FS characteristics
- This is a group project
 - Email me by Wednesday, 12/2, if you want to change groups

© 2004-2007 Ed Lazowska, Hank Levy, Andrea and
Remzi Arpaci-Dusseau, Michael Swift

3

Measuring FS characteristics


- Block size
 - What is the smallest unit of data the file system will transfer to and from disk?
- Prefetch size
 - How much data does the FS read in advance when you read sequentially?
 - When you read sequential data, the FS may fetch more data than you want
 - Usually it gets your data, then keeps fetching more in the background
 - How much it fetches depends on:
 - How far into the file you are
 - Whether the block number is a multiple of 2,4,8,16,32, etc

© 2004-2007 Ed Lazowska, Hank Levy, Andrea and
Remzi Arpaci-Dusseau, Michael Swift

4

Multi-Level Indexed Files

- Variation of Indexed Allocation
 - Dynamically allocate hierarchy of pointers to blocks as needed
 - Meta-data: Small number of pointers allocated statically
 - Additional pointers to blocks of pointers
 - Examples: UNIX FFS-based file systems



12/1/2009

© 2004-2007 Ed Lazowska, Hank Levy, Andrea and
Remzi Arpaci-Dusseau, Michael Swift

5

Measuring FS characteristics

- Number of direct block pointers
 - File system stores indexes of some blocks directly with file, others are stored with indirect pointers
 - How many blocks can be directly referenced from the file?

© 2004-2007 Ed Lazowska, Hank Levy, Andrea and
Remzi Arpaci-Dusseau, Michael Swift

6

How to do this

- Disk access takes longer than memory access
 - ~1 ms for disk
 - ~1-10 microseconds for memory
- Figure out when the file system has to access the disk
 - Write code to force this to happen
 - Use time measurements to figure out when it does happen

© 2004-2007 Ed Lazowska, Hank Levy, Andrea and
Remzi Arpaci-Dusseau, Michael Swift

7

Timers

- You need precise timers to measure the difference between.
 - Use rdtsc() instruction on x86.

```
extern __inline__ unsigned long long int
rdtsc() {
    unsigned long long int x;
    __asm__ volatile (".byte 0x0f, 0x31" : "=A"
        (x));
    return x;
}
```

- This returns # of processor cycles since boot
 - Divide by MHz from /proc/cpu_info to convert

© 2004-2007 Ed Lazowska, Hank Levy, Andrea and
Remzi Arpaci-Dusseau, Michael Swift

8

Timers

- Other timers
 - Gettimeofday() returns data with microsecond precision but may not be accurate (microsecond value may not change)
 - Clock_gettime() can also be used
 - Use CLOCK_MONOTONIC _HR

© 2004-2007 Ed Lazowska, Hank Levy, Andrea and
Remzi Arpaci-Dusseau, Michael Swift

9

Test platforms

- Run tests on instructional linux machines
 - Measure files in /tmp
 - This uses ext3 file system
- You can develop/test code on any machine, but results must be from these machines
 - Don't try this on a Mac

© 2004-2007 Ed Lazowska, Hank Levy, Andrea and
Remzi Arpaci-Dusseau, Michael Swift

10

Hints

- The file system reads in extra block when you read sequentially
 - To bypass this, read backwards
- You need to control what is in the cache
 - Use O_DIRECT flag to open()
 - Use fadvice perl script
 - You may need some experimentation to get this right

© 2004-2007 Ed Lazowska, Hank Levy, Andrea and
Remzi Arpaci-Dusseau, Michael Swift

11

Working in groups

- Meet early to divide work
 - Schedule meetings and don't blow them off; that will hurt you grade
- Use source code control
 - Subversion, CVS, git, mercurial, etc.
- Make sure code compiles before you check it in
 - Don't make others do your dirty work
- Commit to when you will have things ready for others to use/test
 - And tell them in advance if you will not make the deadline
- Others?

© 2004-2007 Ed Lazowska, Hank Levy, Andrea and
Remzi Arpaci-Dusseau, Michael Swift

12