


TCSS 422: OPERATING SYSTEMS

INTRODUCTION TO  
OPERATING SYSTEMS,  
PROCESSES



Wes J. Lloyd

School of Engineering and Technology

University of Washington - Tacoma


April 2, 2020

TCSS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

Tacoma

1

CHAPTER 2:  
INTRODUCTION TO  
OPERATING SYSTEMS



April 2, 2020

TCSS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

Tacoma

L2.2

2

MATERIAL / PACE

- Please classify your perspective on material covered in today's class (52 respondents):
- 1-mostly review, 5-equal new/review, 10-mostly new
- **Average – 5.83 (-)**
- Please rate the pace of today's class:
- 1-slow, 5-just right, 10-fast
- **Average – 5.17 (-)**

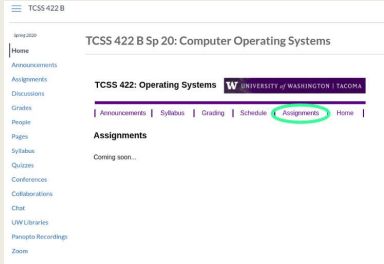
April 2, 2020

TCSS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.3

3

FEEDBACK FROM 3/31

- **What are the actual programming Assignments?**
- In Canvas...
  - 

April 2, 2020

TCSS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.4

4

FEEDBACK - 2

- **I am not sure if I could just use VSCode to develop the program since I prefer It over VM?**
- How to install VSCode on Ubuntu 18.04:  
<https://linuxize.com/post/how-to-install-visual-studio-code-on-ubuntu-18-04/>

April 2, 2020

TCSS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.5

5

FEEDBACK - 3

- **How to Invoke concurrency through the use of PIDs?**
- In Linux, concurrency (*multiple things happening at the same time*) is implementing using either **PROCESSES** or **THREADS**
- When we create a new **PROCESS** or **THREAD** Linux assigns a Process ID (PID) as a unique identifier
- Linux then creates data records that capture lots of state information regarding **PROCESSES** and **THREADS** that are indexed by the PID
- This data is exposed using "virtual files" that are generated on-the-fly by Linux which can be found under a directory on the filesystem, (one for each PID) here → `/proc/{pid}/`
  - `cd /proc/1`
  - `ls`

April 2, 2020

TCSS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.6

6

WORKING WITH PROCESSES IN LINUX

- **CTRL-C** - CANCEL/EXIT a process
- **CTRL-Z** - SUSPEND/PAUSE process, return to command prompt
- **bg** - SEND paused process to background and RESUME
  - Disconnects the standard input (keyboard)
  - Standard output still written to console
- **fg** - BRINGS top most process from jobs list to foreground
  - Reconnects the standard input (keyboard)
- **jobs** - shows list of suspended/backgrounded jobs
- **jobs -p** - shows PID of suspended/backgrounded jobs
- **top** - "task manager" like User Interface that shows PIDs
- **htop** - another "task manager" alternative
- **ps** - command to inspect processes in Linux

April 2, 2020

TCCS422: Operating Systems (Spring 2020)  
School of Engineering and Technology, University of Washington - Tacoma

L2.7

7

COURSE BACKGROUND SURVEY

- Please complete the Course Background Survey:  
  
Computer science, demographics, employment, goals, covid-19, etc.
- <https://forms.gle/MucS87eDQsS4B3328>

April 2, 2020

TCCS422: Operating Systems (Spring 2020)  
School of Engineering and Technology, University of Washington - Tacoma

L2.8

8

"QUIZ" 0 – C PROGRAMMING BACKGROUND SURVEY

- Available via Canvas System
- Under:  
Assignments → Tutorials/Quizzes/In-class Activities
- Please disregard grade assigned by Canvas
- All submissions will receive 10 pts after assignment closes

April 2, 2020

TCCS422: Operating Systems (Spring 2020)  
School of Engineering and Technology, University of Washington - Tacoma

L2.9

9

VIRTUAL MACHINE SURVEY

- Please complete the Virtual Machine Survey to request a "School of Engineering and Technology" remotely hosted Ubuntu VM
- Requires log-in to UW Google for verification:
- <https://forms.gle/R8N4HTjx6qKf1VJ88>

April 2, 2020

TCCS422: Operating Systems (Spring 2020)  
School of Engineering and Technology, University of Washington - Tacoma

L2.10

10

Have you previously used Oracle Virtual Box to create a Virtual Machine?

Yes

no

April 2, 2020

TCCS422: Operating Systems (Spring 2020)  
Start the pre School of Engineering and Technology, University of Washington - Tacoma

Total Res: L2.11

11

Have you previously used Oracle Virtual Box to create a Virtual Machine?

Yes

no

April 2, 2020

TCCS422: Operating Systems (Spring 2020)  
Start the pre School of Engineering and Technology, University of Washington - Tacoma

L2.12

12

### Have you previously used Oracle Virtual Box to create a Virtual Machine?

yes no

April 2, 2020 TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma L2.13

13

### OBJECTIVES – 4/2

- **Chapter 2: Operating Systems – Three Easy Pieces**
  - Introduction to operating systems
  - Management of resources
  - Concepts of virtualization/abstraction
  - **THREE EASY PIECES:**
    - Virtualizing the CPU
    - Virtualizing Memory
    - Virtualizing I/O
  - Operating system design goals
- **Chapter 4: Processes**
  - Process states, context switches
  - Kernel data structures for processes and threads

April 2, 2020 TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma L2.14

14

### RESOURCE MANAGEMENT

- The OS is a resource manager
- Manages CPU, disk, network I/O
- Enables many programs to
  - **Share the CPU**
  - **Share the underlying physical memory (RAM)**
  - **Share physical devices**
    - Disks
    - Network Devices
    - ...

April 2, 2020 TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma L2.15

15

### VIRTUALIZATION

- Operating systems present **physical resources** as **virtual representations** to the programs sharing them
  - Physical resources: CPU, disk, memory, ...
  - The virtual form is **“abstract”**
  - The OS presents an illusion that each user program runs in isolation on its own hardware
  - This virtual form is general, powerful, and easy-to-use

April 2, 2020 TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma L2.16

16

### ABSTRACTIONS

- What form of abstraction does the OS provide?
  - CPU
    - Processes and threads
  - Memory
    - Address space
    - → large array of bytes
    - All programs see the same “size” of RAM
  - Disk
    - Files, file systems

April 2, 2020 TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma L2.17

17

### ABSTRACTION CHALLENGES

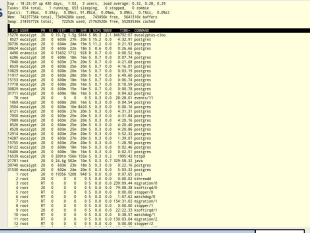
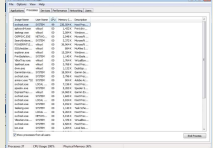
- What level of abstraction?
  - How much of the underlying hardware should be exposed?
    - What if **too much**?
    - What if **too little**?
- What are the correct abstractions?
  - Security concerns

April 2, 2020 TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma L2.18

18

VIRTUALIZING THE CPU

- Each running program gets its own "virtual" representation of the CPU
- Many programs seem to run at once
- Linux: "top" command shows process list
- Windows: task manager



April 2, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.19

19

VIRTUALIZING THE CPU - 2

- Simple Looping C Program

```
1 #include <stdio.h>
2 #include <stdlib.h>
3 #include <sys/time.h>
4 #include <assert.h>
5 #include "common.h"
6
7 int
8 main(int argc, char *argv[])
9 {
10     if (argc != 2) {
11         fprintf(stderr, "usage: cpu <string>\n");
12         exit(1);
13     }
14     char *str = argv[1];
15     while (1) {
16         Spin(1); // Repeatedly checks the time and
17                 // returns once it has run for a second
18         printf("%s\n", str);
19     }
20     return 0;
}
```

April 2, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.20

20

VIRTUALIZING THE CPU - 3

```
prompt> gcc -o cpu cpu.c -Wall
prompt> ./cpu "A"
A
A
A
A
^C
prompt>
```

- Runs forever, must Ctrl-C to halt...

April 2, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.21

21

VIRTUALIZATION THE CPU - 4

```
prompt> ./cpu A & ; ./cpu B & ; ./cpu C & ; ./cpu D &
[1] 7353
[2] 7354
[3] 7355
[4] 7356
A
B
D
C
A
B
D
C
A
C
B
D
...
```

Even though we have only one processor, all four instances of our program seem to be running at the same time!

April 2, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.22

22

VIRTUALIZING MEMORY

- Computer memory is treated as a large array of bytes
- Programs store all data in this large array
  - Read memory (load)
    - Specify an address to read data from
  - Write memory (store)
    - Specify data to write to an address

April 2, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.23

23

VIRTUALIZING MEMORY - 2

- Program to read/write memory (mem.c):

```
1 #include <unistd.h>
2 #include <stdio.h>
3 #include <stdlib.h>
4 #include "common.h"
5
6 int
7 main(int argc, char *argv[])
8 {
9     int *p = malloc(sizeof(int)); // a1: allocate some
10                                     memory
11     assert(p != NULL);
12     printf("(%) address of p: %08x\n",
13            getpid(), (unsigned) p); // a2: print out the
14                                     address of the memory
15     *p = 0; // a3: put zero into the first slot of the memory
16     while (1) {
17         Spin(1);
18         *p = *p + 1;
19         printf("(%) p: %d\n", getpid(), *p); // a4
20     }
21     return 0;
}
```

April 2, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.24

24

VIRTUALIZING MEMORY - 3

- Output of mem.c

```
prompt> ./mem
(2134) memory address of p: 00200000
(2134) p: 1
(2134) p: 2
(2134) p: 3
(2134) p: 4
(2134) p: 5
^C
```

- int value stored at virtual memory address: 00200000
- program increments int value

April 2, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.25

25

VIRTUALIZING MEMORY - 4

- Multiple instances of mem.c

```
prompt> ./mem & ./mem &
[1] 24113
[2] 24114
(24113) memory address of p: 00200000
(24114) memory address of p: 00200000
(24113) p: 1
(24114) p: 1
(24114) p: 2
(24113) p: 2
(24113) p: 3
(24114) p: 3
...
```

- THE BOOK IS WRONG – Linux has changed !!
- What could be wrong about having malloc() return the same virtual memory address for every program instance?

April 2, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.26

26

VIRTUALIZING MEMORY - 5

- Multiple instances of mem.c

```
prompt> ./mem & ./mem &
[1] 24113
[2] 24114
(24113) memory address of p: 00200000
(24114) memory address of p: 00200000
(24113) p: 1
(24114) p: 1
(24114) p: 2
(24113) p: 2
(24113) p: 3
(24114) p: 3
...
```

- ORIGINALLY: (int\*)p receives the same memory location 00200000
- Why does modifying (int\*)p in program #1 (PID=24113), not interfere with (int\*)p in program #2 (PID=24114) ?
  - The OS has “virtualized” memory. Each program has it’s own virtual address space

April 2, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.27

27

INSPECTING THE VIRTUAL MEMORY MAP OF A PROCESS

- cat /proc/\$\$/maps
- \$\$ is the current process, can replace with an PID

```
llloyd@llloyd:/$ cat /proc/30402/maps
00400000-00401000 r-xp 00000000 fc:00 6611568      /home/llloyd/Dropbox/courses/tccs422/examples/mem
00000000-00001000 r-p 00000000 fc:00 6611568      /home/llloyd/Dropbox/courses/tccs422/examples/mem
00001000-00002000 rw-p 00001000 fc:00 6611568      /home/llloyd/Dropbox/courses/tccs422/examples/mem
02483000-02484000 rw-p 00000000 00:00 0          [heap]
7f2b456c000-7f2b456d000 r-xp 00000000 fc:00 15738979 /lib/x86_64-linux-gnu/libc-2.23.so
7f2b456d000-7f2b456e000 r-p 0001c000 fc:00 15738979 /lib/x86_64-linux-gnu/libc-2.23.so
7f2b456e000-7f2b456f000 r-p 0001c000 fc:00 15738979 /lib/x86_64-linux-gnu/libc-2.23.so
7f2b456f000-7f2b4570000 r-p 00000000 00:00 0          /lib/x86_64-linux-gnu/libpthread-2.23.so
7f2b4570000-7f2b4571000 r-p 00018000 fc:00 15738966 /lib/x86_64-linux-gnu/libpthread-2.23.so
7f2b4571000-7f2b4572000 r-p 00018000 fc:00 15738966 /lib/x86_64-linux-gnu/libpthread-2.23.so
7f2b4572000-7f2b4573000 r-p 00018000 fc:00 15738966 /lib/x86_64-linux-gnu/libpthread-2.23.so
7f2b4573000-7f2b4574000 r-p 00000000 00:00 0          /lib/x86_64-linux-gnu/lib-2.23.so
7f2b4574000-7f2b4575000 r-p 00026000 fc:00 15738965 /lib/x86_64-linux-gnu/lib-2.23.so
7f2b4575000-7f2b4576000 r-p 00000000 00:00 0          [stack]
7f2b4576000-7f2b4577000 r-p 00000000 00:00 0          [var]
7f2b4577000-7f2b4578000 r-xp 00000000 00:00 0          [vdso]
ffffffffff0000-ffffffffff0100 r-xp 00000000 00:00 0
```

April 2, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.28

28

VIRTUAL MEMORY

- Key take-aways:
- Each process (program) has its own **virtual address space**
- The OS maps virtual **address spaces** onto **physical memory**
- A memory reference from one process can not affect the address space of others.
  - Isolation
- Physical memory, a **shared resource**, is managed by the OS

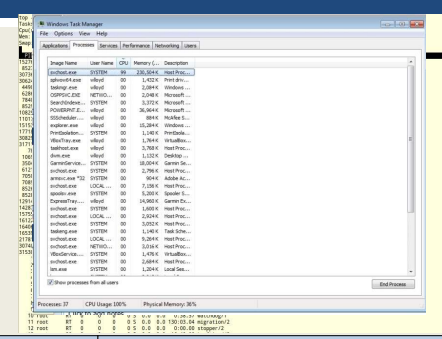
April 2, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.29

29

CONCURRENCY



April 2, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.30

30

CONCURRENCY

- Linux: 654 tasks
- Windows: 37 processes
- The OS appears to run many programs at once, juggling them
- Modern multi-threaded programs feature concurrent threads and processes
- What is a key difference between a process and a thread?

April 2, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.31

31

CONCURRENCY - 2

```
1  #include <stdio.h>
2  #include <stdlib.h>
3  #include "common.h"
4
5  volatile int counter = 0;
6  int loops;
7
8  void *worker(void *arg) {
9
10
11
12
13
14 }
15 ...
```

Not the same as Java volatile:  
Provides a compiler hint that an object may change value unexpectedly (in this case by a separate thread) so aggressive optimization must be avoided.

thread.c listing continues ...

Good article on Java volatile keyword:  
(hint – not enough to ensure correctness w/ concurrent threads in JAVA)  
<http://tutorials.jenkov.com/java-concurrency/volatile.html>

April 2, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.32

32

CONCURRENCY - 3

```
16  int
17  main(int argc, char *argv[])
18  {
19      if (argc != 2) {
20          fprintf(stderr, "usage: threads <value>\n");
21          exit(1);
22      }
23      loops = atoi(argv[1]);
24      pthread_t p1, p2;
25      printf("Initial value : %d\n", counter);
26
27      pthread_create(&p1, NULL, worker, NULL);
28      pthread_create(&p2, NULL, worker, NULL);
29      pthread_join(p1, NULL);
30      pthread_join(p2, NULL);
31      printf("Final value : %d\n", counter);
32      return 0;
33  }
```

- Program creates two threads
- Check documentation: “man pthread\_create”
- worker() method counts from 0 to argv[1] (loop)

April 2, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.33

33

Linux “man” page example

PTHREAD\_CREATE(3)

Linux Programmer's Manual

PTHREAD\_CREATE(3)

NAME

pthread\_create - create a new thread

SYNOPSIS

#include <pthread.h>
int pthread\_create(pthread\_t \*thread, const pthread\_attr\_t \*attr, void \*(\*start\_routine)(void \*), void \*arg);
Compile and link with -pthread.

DESCRIPTION

The pthread\_create() function starts a new thread in the calling process. The new thread starts execution by invoking start\_routine(). arg is passed as the sole argument of start\_routine().
The new thread terminates in one of the following ways:
\* It calls pthread\_exit(), specifying an exit status value that is available to another thread in the same process that calls pthread\_join().
\* It returns from start\_routine(). This is equivalent to calling pthread\_exit() with the value supplied in the return statement.
\* It is canceled (see pthread\_cancel()).
\* Any of the threads in the process calls exit(), or the main thread performs a return from main(). This causes the termination of all threads in the process.
The attr argument points to a pthread\_attr\_t structure whose contents are used at thread creation time to determine attributes for the new thread; this structure is initialized using pthread\_attr\_init() and related functions. If attr is NULL, then the thread is created with default attributes.

April 2, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.34

34

CONCURRENCY - 4

- Command line parameter argv[1] provides loop length
- Defines number of times the shared counter is incremented
- Loops: 1000

```
prompt> gcc -o thread thread.c -Wall -pthread
prompt> ./thread 1000
Initial value : 0
Final value : 2000
```

- Loops 100000

```
prompt> ./thread 100000
Initial value : 0
Final value : 149012 // huh??
prompt> ./thread 100000
Initial value : 0
Final value : 137298 // what the??
```

April 2, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.35

35

CONCURRENCY - 5

- When loop value is large why do we not achieve 200000 ?
- C code is translated to (3) assembly code operations
  - Load counter variable into register
  - Increment it
  - Store the register value back in memory
- These instructions happen concurrently and VERY FAST
- (P1 || P2) write incremented register values back to memory, While (P1 || P2) read same memory
- Memory access here is unsynchronized (non-atomic)
- Some of the increments are lost

April 2, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.36

36

W

To perform parallel work, a single process may:

Launch multiple threads to execute code in parallel while sharing global data in memory

Launch multiple processes to execute code in parallel without sharing global data in memory

Both A and B

None of the above

Start the presentation to see live content. Still no live content? Install the app or get help at [PellEv.com/app](#)

Total Revenue

37

PARALLEL PROGRAMMING

- To perform parallel work, a single process may:
- A. Launch multiple threads to execute code in parallel while sharing global data in memory
- B. Launch multiple processes to execute code in parallel without sharing global data in memory
- C. Both A and B
- D. None of the above

April 2, 2020

TCSS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.38

38

PERSISTENCE

- DRAM: Dynamic Random Access Memory: DIMMs/SIMMs
  - Stores data while power is present
  - When power is lost, data is lost (*volatile*)
- Operating System helps "persist" data more **permanently**
  - I/O device(s): hard disk drive (HDD), solid state drive (SSD)
  - File system(s): "catalog" data for storage and retrieval

April 2, 2020

TCSS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.39

39

PERSISTENCE - 2

```
1  #include <stdio.h>
2  #include <unistd.h>
3  #include <assert.h>
4  #include <fcntl.h>
5  #include <sys/types.h>
6
7  int
8  main(int argc, char *argv[])
9  {
10     int fd = open("/tmp/file", O_WRONLY | O_CREAT
11                  | O_TRUNC, S_IRWXU);
12     assert(fd > -1);
13     int rc = write(fd, "hello world\n", 13);
14     assert(rc == 13);
15     close(fd);
16     return 0;
17 }
```

- open(), write(), close(): OS system calls for device I/O
- Note: man page for open(), write() require page number: "man 2 open", "man 2 write", "man close"

April 2, 2020

TCSS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.40

40

PERSISTENCE - 3

- To write to disk, OS must:
  - Determine where on disk data should reside
  - Perform sys calls to perform I/O:
    - Read/write to file system (*inode record*)
    - Read/write data to file
- Provide fault tolerance for system crashes
  - Journaling: Record disk operations in a journal for replay
  - Copy-on-write - replicating shared data - see *ZFS*
  - Carefully order writes on disk

April 2, 2020

TCSS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.41

41

SUMMARY:  
OPERATING SYSTEM DESIGN GOALS

- ABSTRACTING THE HARDWARE**
  - Makes programming code easier to write
  - Automate sharing resources – save programmer burden
- PROVIDE HIGH PERFORMANCE**
  - Minimize overhead from OS abstraction (Virtualization of CPU, RAM, I/O)
  - Share resources fairly
  - Attempt to tradeoff performance vs. fairness → consider priority
- PROVIDE ISOLATION**
  - User programs can't interfere with each other's virtual machines, the underlying OS, or the sharing of resources

April 2, 2020

TCSS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.42

42

SUMMARY:  
OPERATING SYSTEM DESIGN GOALS - 2

■ **RELIABILITY**

■ OS must not crash, 24/7 Up-time

■ Poor user programs must not bring down the system:


Blue Screen

■ Other Issues:

■ Energy-efficiency

■ Security (of data)

■ Cloud: Virtual Machines



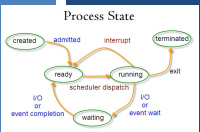
April 2, 2020


TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.43

43

CHAPTER 4:  
PROCESSES





January 9, 2019

TCCS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L2.44

44

VIRTUALIZING THE CPU

■ How should the CPU be shared?

■ Time Sharing:  
Run one process, pause it, run another

■ The act of swapping process A out of the CPU to run process B is called a:

■ **CONTEXT SWITCH**

■ How do we SWAP processes in and out of the CPU efficiently?

■ Goal is to minimize **overhead** of the swap

■ **OVERHEAD** is time spent performing OS management activities that don't help accomplish real work

January 9, 2019

TCCS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L2.45

45

PROCESS

A process is a running program.

■ Process comprises of:

■ **Memory**

■ Instructions ("the code")

■ Data (heap)

■ **Registers**

■ PC: Program counter

■ Stack pointer

January 9, 2019

TCCS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L2.46

46

PROCESS API

■ Modern OSes provide a Process API for process support

■ Create

■ Create a new process

■ Destroy

■ Terminate a process (ctrl-c)

■ Wait

■ Wait for a process to complete/stop

■ Miscellaneous Control

■ Suspend process (ctrl-z)

■ Resume process (fg, bg)

■ Status

■ Obtain process statistics: (top)

January 9, 2019

TCCS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L2.47

47

PROCESS API: CREATE

1. Load program code (and static data) into memory

■ Program executable code (binary): loaded from disk

■ Static data: also loaded/created in address space

■ **Eager loading:** Load entire program before running

■ **Lazy loading:** Only load what is immediately needed

■ Modern OSes: Supports paging & swapping

2. Run-time stack creation

■ Stack: local variables, function params, return address(es)

January 9, 2019

TCCS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L2.48

48



PROCESS API: CREATE

3. Create program's heap memory

▪ For dynamically allocated data

4. Other initialization

▪ I/O Setup

▪ Each process has three open file descriptors:  
Standard Input, Standard Output, Standard Error

5. Start program running at the entry point: `main()`

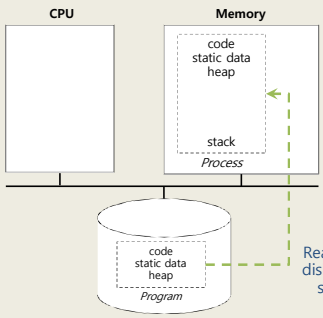
▪ OS transfers CPU control to the new process

January 9, 2019

TCCS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L2.49

49



January 9, 2019

TCCS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L2.50

50

PROCESS STATES

▪ **RUNNING**

▪ Currently executing instructions

▪ **READY**

▪ Process is ready to run, but has been preempted

▪ CPU is presently allocated for other tasks

▪ **BLOCKED**

▪ Process is **not** ready to run. It is waiting for another event to complete:

▪ Process has already been initialized and run for awhile

▪ Is now waiting on I/O from disk(s) or other devices

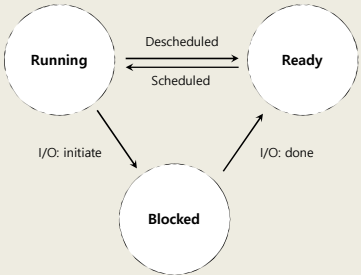
January 9, 2019

TCCS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L2.51

51

PROCESS STATE TRANSITIONS



January 9, 2019

TCCS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L2.52

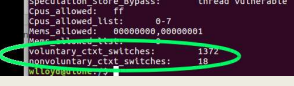
52

OBSERVING PROCESS META-DATA

▪ Can inspect the number of **CONTEXT SWITCHES** made by a process

▪ Let's run `mem.c` (from chapter 2)

▪ `cat /proc/{process-id}/status`



▪ `proc "status"` is a virtual file generated by Linux

▪ Provides a report with process related meta-data

▪ What appears to happen to the number of context switches the longer a process runs? (`mem.c`)

April 2, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.53

53

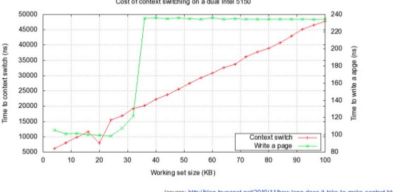
CONTEXT SWITCH

▪ **How long does a context switch take?**

▪ 10,000 to 50,000 ns (.01 to .05 ms)

▪ 2,000 context switches is near 100ms

**Without CPU affinity**



April 2, 2020

TCCS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L2.54

54

**W** When a process is in this state, it is advantageous for the Operating System to perform a CONTEXT SWITCH to perform other work

RUNNING READY BLOCKED All of the above None of the above

January 9, 2019 TCS5422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma L2.55

55

**QUESTION: WHEN TO CONTEXT SWITCH**

- When a process is in this state, it is advantageous for the Operating System to perform a CONTEXT SWITCH to perform other work:
  - (a) RUNNING
  - (b) READY
  - (c) BLOCKED
  - (d) All of the above
  - (e) None of the above

January 9, 2019 TCS5422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma L2.56

56

**PROCESS DATA STRUCTURES**

- OS provides data structures to track process information
  - Process list
    - Process Data
    - State of process: Ready, Blocked, Running
  - Register context
- PCB (Process Control Block)
  - A C-structure that contains information about each process

January 9, 2019 TCS5422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma L2.57

57

**XV6 KERNEL DATA STRUCTURES**

- xv6: pedagogical implementation of Linux
- Simplified structures

```
// the registers xv6 will save and restore
// to stop and subsequently restart a process
struct context {
    int eip; // Index pointer register
    int esp; // Stack pointer register
    int ebx; // Called the base register
    int ecx; // Called the counter register
    int edx; // Called the data register
    int esi; // Source index register
    int edi; // Destination index register
    int ebp; // Stack base pointer register
};

// the different states a process can be in
enum proc_state { UNUSED, EMBRYO, SLEEPING,
    RUNNABLE, RUNNING, ZOMBIE };
```

January 9, 2019 TCS5422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma L2.58

58

**XV6 KERNEL DATA STRUCTURES - 2**

```
// the information xv6 tracks about each process
// including its register context and state
struct proc {
    char *mem; // Start of process memory
    uint sz; // Size of process memory
    char *kstack; // Bottom of kernel stack
    // for this process
    enum proc_state state; // Process state
    int pid; // Process ID
    struct proc *parent; // Parent process
    void *chan; // If non-zero, sleeping on chan
    int killed; // If non-zero, have been killed
    struct file *ofile[NFILE]; // Open files
    struct inode *cwd; // Current directory
    struct context context; // Switch here to run process
    struct trapframe *tf; // Trap frame for the current interrupt
};
```

January 9, 2019 TCS5422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma L2.59

59

**LINUX: STRUCTURES**

- struct task\_struct**, equivalent to struct proc
  - Provides process description
  - Large: 10,000+ bytes
  - /usr/src/linux-headers-[kernel version]/include/linux/sched.h
    - ~ LOC 1391 – 1852 (4.4.0-170)
    - earlier was LOC 1227 – 1587
- struct thread\_info**, provides “context”
  - thread\_info.h is at:
    - /usr/src/linux-headers-[kernel version]/arch/x86/include/asm/

January 9, 2019 TCS5422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma L2.60

60

LINUX: THREAD\_INFO

```
struct thread_info {
    struct task_struct *task; /* main task structure */
    struct exec_domain *exec_domain; /* execution domain */
    __u32 flags; /* low level flags */
    __u32 status; /* thread synchronous flags */
    __u32 cpu; /* current CPU */
    int preempt_count; /* 0 => preemptable,
                        <0 => BUG */
    mm_segment_t addr_limit;
    struct restart_block restart_block;
    void __user *sysenter_return;
#ifdef CONFIG_X86_32
    unsigned long previous_esp; /* ESP of the previous stack in
                                case of nested (IRQ) stacks */
#endif
    __u8 supervisor_stack[0];
#ifdef __u8
    int uaccess_err;
};
```

January 9, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L2.61

61

LINUX STRUCTURES - 2

- List of Linux data structures:  
<http://www.tldp.org/LDP/tlk/ds/ds.html>
- Description of process data structures:  
<https://learning.oreilly.com/library/view/linux-kernel-development/9780768696974/cover.html>  
3rd edition is online (dated from 2010):  
See chapter 3 on Process Management  
  
Safari online – accessible using UW ID SSO login  
Linux Kernel Development, 3<sup>rd</sup> edition  
Robert Love  
Addison-Wesley


January 9, 2019

TCSS422: Operating Systems [Winter 2019]  
School of Engineering and Technology, University of Washington - Tacoma

L2.62

62

QUESTIONS



March 31, 2020

TCSS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L1.64

63

WE WILL RETURN AT  
2:40PM



March 31, 2020

TCSS422: Operating Systems [Spring 2020]  
School of Engineering and Technology, University of Washington - Tacoma

L1.64

64