


TCSS 422: OPERATING SYSTEMS

Three Easy Pieces
Processes, Process API



Wes J. Lloyd
Institute of Technology
University of Washington - Tacoma

March 28, 2018

TCSS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

OBJECTIVES

- Assignment 0 – Introduction to Linux
- Feedback from 3/26
- Chapter 2: Operating Systems – Three Easy Pieces
 - ✓ Easy piece #1: CPU Virtualization
 - ✓ Easy piece #2: Memory Virtualization
 - Easy piece #3: I/O Virtualization
 - Operating system design goals
- Processes – Ch. 4
- C Linux Process API – Ch. 5
- Limited Direct Execution – Ch. 6
 - Virtualizing the CPU

March 28, 2018

TCSS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L2.2

VIRTUAL MACHINE SURVEY

- Please complete the Virtual Machine Survey is wanting an Institute of Technology hosted Ubuntu 16.04 VM
- <https://goo.gl/forms/w9VWqkX756yXBUBt1>

March 28, 2018

TCSS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L2.3

SELECTED FEEDBACK FROM 3/26

- What would be the shared memory for threads?
- Are processes running really fast one at a time, so that it looks like multiple processes are running at the same time?
- How do you prevent race conditions?
- How do we determine whether to use a thread or a process?

March 28, 2018

TCSS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L2.4

FEEDBACK - 2


- Will the in class quizzes be known in advance?
- Do you also record the lectures done in class?

March 28, 2018

TCSS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L2.5

CH. 2: INTRODUCTION TO OPERATING SYSTEMS



March 28, 2018

TCSS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L2.6

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

QUESTION: 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

March 28, 2018
TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma
L2.8

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

March 28, 2018
TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma
L2.9

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;

```

- `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"

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma
L2.10

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: see **ZFS**
 - Carefully order writes on disk

March 28, 2018
TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma
L2.11

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

March 28, 2018
TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma
L2.12

SUMMARY: OPERATING SYSTEM DESIGN GOALS - 2

- **RELIABILITY**
 - OS must not crash, 24/7 Up-time
 - Poor user programs must not bring down the system:

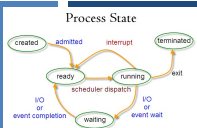

Blue Screen



- Other Issues:
 - Energy-efficiency
 - Security (of data)
 - Cloud: Virtual Machines

March 28, 2018	TCS422: Operating Systems [Spring 2018] Institute of Technology, University of Washington - Tacoma	L2.13
----------------	---	-------

CHAPTER 4: PROCESSES

March 28, 2018	TCS422: Operating Systems [Spring 2018] Institute of Technology, University of Washington - Tacoma	L2.14
----------------	---	-------

CPU VIRTUALIZING

- How should the CPU be shared?
- Time Sharing:
Run one process, pause it, run another
- How do we SWAP processes in and out of the CPU efficiently?
 - Goal is to minimize **overhead** of the swap

March 28, 2018	TCS422: Operating Systems [Spring 2018] Institute of Technology, University of Washington - Tacoma	L2.15
----------------	---	-------

PROCESS

A process is a running program.

- Process comprises of:
 - Memory
 - Instructions ("the code")
 - Data (heap)
 - Registers
 - PC: Program counter
 - Stack pointer

March 28, 2018	TCS422: Operating Systems [Spring 2018] Institute of Technology, University of Washington - Tacoma	L2.16
----------------	---	-------

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)

March 28, 2018	TCS422: Operating Systems [Spring 2018] Institute of Technology, University of Washington - Tacoma	L2.17
----------------	---	-------

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)

March 28, 2018	TCS422: Operating Systems [Spring 2018] Institute of Technology, University of Washington - Tacoma	L2.18
----------------	---	-------

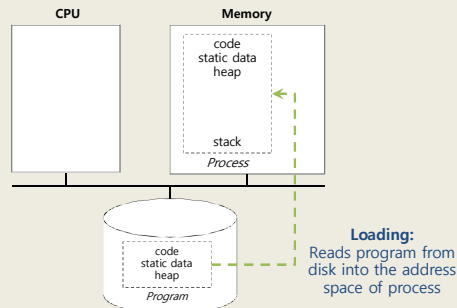
PROCESS API: CREATE

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

March 28, 2018

TCS5422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L2.19



March 28, 2018

TCS5422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L2.20

PROCESS STATES

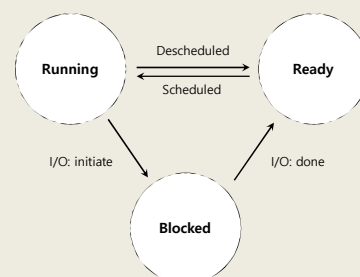
- **RUNNING**
 - Currently executing instructions
- **READY**
 - Process is ready to run, but has been preempted
 - CPU is presently allocated for other tasks
- **BLOCKED**
 - Process is **not** ready to run. It is waiting for another event to complete:
 - Process has already been initialized and run for awhile
 - Is now waiting on I/O from disk(s) or other devices

March 28, 2018

TCS5422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L2.21

PROCESS STATE TRANSITIONS



March 28, 2018

TCS5422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L2.22

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

March 28, 2018

TCS5422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L2.23

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 };
```

March 28, 2018

TCS5422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L2.24

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[NOFILE]; // Open files
    struct inode *cwd;    // Current directory
    struct context context; // Switch here to run process
    struct trapframe *tf; // Trap frame for the
                        // current interrupt
};
```

March 28, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L2.25

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
 - 1227 – 1587
- **struct thread_info**, provides “context”
 - thread_info.h is at:
/usr/src/linux-headers-[kernel version]/arch/x86/include/asm/

March 28, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L2.26

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 asynchronous flags */
    __u32 cpu;                          /* current CPU */
    int preempt_count;                 /* 0 => preemptable,
                                        <0 => BUG */
    mm_segment_t addr_limit;           /* user/kernel address limit */
    struct restart_block restart_block; /* restart block */
    void *user;                        /* user space pointer */
#ifdef CONFIG_X86_32
    unsigned long previous_esp;        /* ESP of the previous stack in
                                        case of nested (IRQ) stacks */
#endif
    __u8 supervisor_stack[0];          /* supervisor stack */
#ifdef CONFIG_X86_32
    int uaccess_err;                  /* uaccess error */
#endif
};
```

March 28, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L2.27

LINUX STRUCTURES - 2

- List of Linux data structures:
<http://www.tldp.org/LDP/tlk/ds/ds.html>
- Description of process data structures:
<http://www.makelinux.net/books/lkd2/ch03lev1sec1>
2nd edition is online (dated from 2005):
Linux Kernel Development, 2nd edition
Robert Love
Sams Publishing

March 28, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L2.28

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

State	Count
RUNNING	1
READY	2
BLOCKED	42
All of the above	4
None of the above	5

March 28, 2018

TCCS422: Operating Systems [Spring 2018]
Start the presentation at the Institute of Technology, University of Washington - Tacoma
Total Runs: L2.29

L2.29

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

March 28, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L2.30

CHAPTER 5:
C PROCESS API




March 28, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L2.31

fork()

- Creates a new process - think of “a fork in the road”
- “Parent” process is the original
- Creates “child” process of the program from the **current execution point**
- Book says “pretty odd”
- Creates a **duplicate** program instance (these are **processes!**)
- Copy** of
 - Address space (memory)
 - Register
 - Program Counter (PC)
- Fork returns
 - child PID to parent
 - 0 to child



March 28, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L2.32

FORK EXAMPLE

p1.c

```
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>

int main(int argc, char *argv[]){
    printf("hello world (pid:%d)\n", (int) getpid());
    int rc = fork();
    if (rc < 0) { // fork failed; exit
        fprintf(stderr, "fork failed\n");
        exit(1);
    } else if (rc == 0) { // child (new process)
        printf("hello, I am child (pid:%d)\n", (int) getpid());
    } else { // parent goes down this path (main)
        printf("hello, I am parent of %d (pid:%d)\n",
            rc, (int) getpid());
    }
    return 0;
}
```

March 28, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L2.33

FORK EXAMPLE - 2

Non deterministic ordering of execution

```
prompt> ./p1
hello world (pid:29146)
hello, I am parent of 29147 (pid:29146)
hello, I am child (pid:29147)
prompt>
```

or

```
prompt> ./p1
hello world (pid:29146)
hello, I am child (pid:29147)
hello, I am parent of 29147 (pid:29146)
prompt>
```

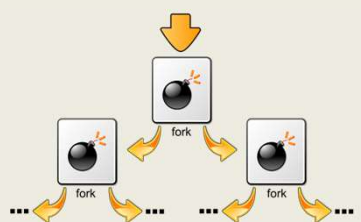
CPU scheduler determines which to run first

March 28, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L2.34

:(){|: &};:



March 28, 2018

TCCS422: Operating Systems [Spring 2018]
Institute of Technology, University of Washington - Tacoma

L2.35

QUESTIONS

