


TCCS 422: OPERATING SYSTEMS

Processes, Process API, Limited Direct Execution



Wes J. Lloyd
School of Engineering and Technology,
University of Washington - Tacoma

October 3, 2018 TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma Tacoma

FEEDBACK FROM 10/1

- How do processes start threads?
 - Done in code or by compiler?
- VM Survey - results submitted
- Assignment 0 questions
- C Tutorial - to be posted



October 3, 2018 TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma L3.2

OBJECTIVES

- Chapter 4 - Processes
- Chapter 5 - Process API
- Chapter 6 - Limited Direct Execution
- Chapter 7 - Introduction to Scheduling
- Chapter 8 - Multi-level Feedback Queue

October 3, 2018 TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma L3.3

CHAPTER 4: PROCESSES



October 3, 2018 TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma L3.4

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

October 3, 2018 TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma L3.5

PROCESS

A process is a running program.

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

October 3, 2018 TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma L3.6

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)

October 3, 2018 TCCS422: Operating Systems [Fall 2018]
 School of Engineering and Technology, University of Washington - Tacoma L3.7

PROCESS API: CREATE

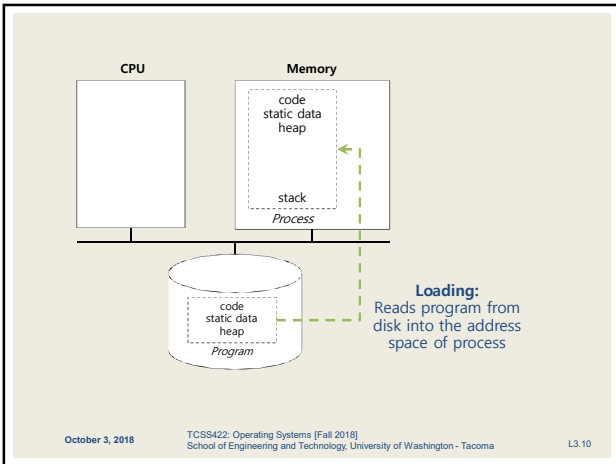
1. Load program code (and static data) into memory
 - Program executable code (binary): loaded from disk
 - Static data: also loaded/created in address space
 - **Eager loading:** Load entire program before running
 - **Lazy loading:** Only load what is immediately needed
 - Modern OSes: Supports paging & swapping
2. Run-time stack creation
 - Stack: local variables, function params, return address(es)

October 3, 2018 TCCS422: Operating Systems [Fall 2018]
 School of Engineering and Technology, University of Washington - Tacoma L3.8

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

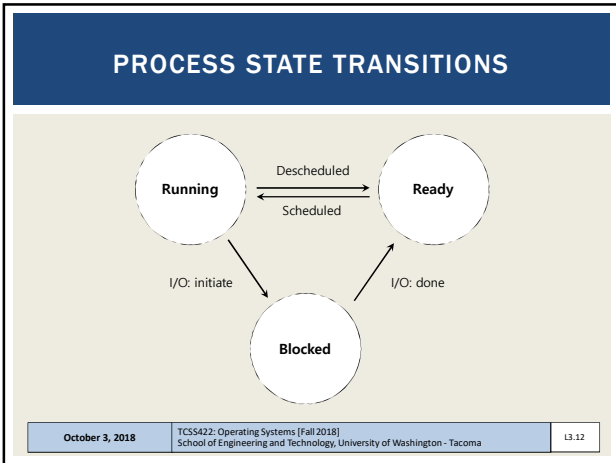
October 3, 2018 TCCS422: Operating Systems [Fall 2018]
 School of Engineering and Technology, University of Washington - Tacoma L3.9



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

October 3, 2018 TCCS422: Operating Systems [Fall 2018]
 School of Engineering and Technology, University of Washington - Tacoma L3.11



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


October 3, 2018 TCCS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma L3.20

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**

October 3, 2018 TCCS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma L3.20


**CHAPTER 5:
 C PROCESS API**



October 3, 2018 TCCS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma L3.21

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



October 3, 2018 TCCS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma L3.22

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;
}
    
```

October 3, 2018 TCCS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma L3.23

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**

October 3, 2018 TCCS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma L3.24

:(){ :| & }::

fork

fork

fork

fork

October 3, 2018 TCCS422: Operating Systems [Fall 2018] L3.25
 School of Engineering and Technology, University of Washington - Tacoma

wait()

- wait(), waitpid()
- Called by parent process
- Waits for a child process to finish executing
- Not a sleep() function
- Provides some ordering to multi-process execution

October 3, 2018 TCCS422: Operating Systems [Fall 2018] L3.26
 School of Engineering and Technology, University of Washington - Tacoma

FORK WITH WAIT

```

#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <sys/wait.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)
        int wc = wait(NULL);
        printf("hello, I am parent of %d (wc:%d) (pid:%d)\n",
            rc, wc, (int) getpid());
    }
    return 0;
}
    
```

October 3, 2018 TCCS422: Operating Systems [Fall 2018] L3.27
 School of Engineering and Technology, University of Washington - Tacoma

FORK WITH WAIT - 2

- Deterministic ordering of execution

```

prompt> ./p2
hello world (pid:29266)
hello, I am child (pid:29267)
hello, I am parent of 29267 (wc:29267) (pid:29266)
prompt>
    
```

October 3, 2018 TCCS422: Operating Systems [Fall 2018] L3.28
 School of Engineering and Technology, University of Washington - Tacoma

FORK EXAMPLE

- Linux example

October 3, 2018 TCCS422: Operating Systems [Fall 2018] L3.29
 School of Engineering and Technology, University of Washington - Tacoma

exec()

- Supports running an external program
- 6 types: execl(), execlp(), execl(), execvp(), execvpe()
- execl(), execlp(), execl(): const char *arg**
 Command arguments provided as **LIST of pointers** to strings provided as arguments... (arg0, arg1, .. argn) (terminated by a null pointer)
- execvp(), execvp(), execvpe()**
 Command arguments provided as an **ARRAY of pointers** to strings as arguments
- Strings are null-terminated
- First argument is name of file being executed

October 3, 2018 TCCS422: Operating Systems [Fall 2018] L3.30
 School of Engineering and Technology, University of Washington - Tacoma

EXEC() - 2

- Common use case:
 - Write a new program which wraps a legacy one
 - Provide a new interface to an old system: Web services
 - Legacy program thought of as a "black box"
- May not want to know what is inside the black box... ☹️
- FORTRAN ???

October 3, 2018
TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma
L3.31

EXEC EXAMPLE

```

#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <string.h>
#include <sys/wait.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());
        char *myargs[3];
        myargs[0] = strdup("wc"); // program: "wc" (word count)
        myargs[1] = strdup("p3.c"); // argument: file to count
        myargs[2] = NULL; // marks end of array
        ...
    }
}
    
```

October 3, 2018
TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma
L3.32

EXEC EXAMPLE - 2

```

...
execvp(myargs[0], myargs); // runs word count
printf("this shouldn't print out");
} else { // parent goes down this path (main)
    int wc = wait(NULL);
    printf("hello, I am parent of %d (wc:%d) (pid:%d)\n",
           rc, wc, (int) getpid());
}
return 0;
}
    
```

```

prompt> ./p3
hello world (pid:29383)
hello, I am child (pid:29384)
29 107 1030 p3.c
hello, I am parent of 29384 (wc:29384) (pid:29383)
prompt>
    
```

October 3, 2018
TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma
L3.33

EXEC WITH FILE REDIRECTION (OUTPUT)

```

#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <string.h>
#include <fcntl.h>
#include <sys/wait.h>

int
main(int argc, char *argv[]){
    int rc = fork();
    if (rc < 0) { // fork failed; exit
        fprintf(stderr, "fork failed\n");
        exit(1);
    } else if (rc == 0) { // child: redirect standard output to a file
        close(STDOUT_FILENO);
        open("./p4.output", O_CREAT|O_WRONLY|O_TRUNC, S_IRWXU);
        ...
    }
}
    
```

October 3, 2018
TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma
L3.34

FILE MODE BITS

- S_IRWXU
read, write, execute/search by owner
- S_IRUSR
read permission, owner
- S_IWUSR
write permission, owner
- S_IXUSR
execute/search permission, owner
- S_IRWXG
read, write, execute/search by group
- S_IRGRP
read permission, group
- S_IWGRP
write permission, group
- S_IXGRP
execute/search permission, group
- S_IRWXO
read, write, execute/search by others
- S_IROTH
read permission, others
- S_IWOTH
write permission, others

October 3, 2018
TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma
L3.35

EXEC W/ FILE REDIRECTION (OUTPUT) - 2

```

// now exec "wc"...
char *myargs[3];
myargs[0] = strdup("wc"); // program: "wc" (word count)
myargs[1] = strdup("p4.c"); // argument: file to count
myargs[2] = NULL; // marks end of array
execvp(myargs[0], myargs); // runs word count
} else { // parent goes down this path (main)
    int wc = wait(NULL);
}
return 0;
}
    
```

```

prompt> ./p4
prompt> cat p4.output
32 109 846 p4.c
prompt>
    
```

October 3, 2018
TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma
L3.36

DIRECT EXECUTION - 2

- **With direct execution:**
 - How does the OS stop a program from running, and switch to another to support **time sharing**?
 - How do programs share disks and perform I/O if they are given direct control? Do they know about each other?
 - With direct execution, how can dynamic memory structures such as linked lists grow over time?

October 3, 2018
TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma
L3.43

CONTROL TRADEOFF

- **Too little control:**
 - No security
 - No time sharing
- **Too much control:**
 - Too much OS overhead
 - Poor performance for compute & I/O
 - Complex APIs (system calls), difficult to use

October 3, 2018
TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma
L3.44

CONTEXT SWITCHING OVERHEAD

The diagram shows three horizontal bars representing execution time. The top bar, 'Multitasking', consists of interleaved colored segments (yellow, green, blue, orange). The middle bar, 'Multitasking with context switching', has the same colored segments but with thin black bars between them, representing the overhead of switching. The bottom bar, 'Sequential', has the colored segments in a single order without interleaving. A bracket above the middle bar is labeled 'Total cost of context switching'.

October 3, 2018
TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma
L3.45

LIMITED DIRECT EXECUTION

- OS implements LDE to support time/resource sharing
- Limited direct execution means “only limited” processes can execute **DIRECTLY** on the CPU in **trusted** mode
- **TRUSTED** means the process is trusted, and it can do anything... (e.g. it is a system / kernel level process)
- Enabled by **protected (safe) control transfer**
- CPU supported context switch
- Provides data isolation

October 3, 2018
TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma
L3.46

CPU MODES

- Utilize CPU Privilege Rings (Intel x86)
 - rings 0 (kernel), 1 (VM kernel), 2 (unused), 3 (user)

access ←———— no access

- **User mode:**
Application is running, but w/o direct I/O access
- **Kernel mode:**
OS kernel is running performing restricted operations

October 3, 2018
TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma
L3.47

CPU MODES

- **User mode: ring 3 - untrusted**
 - Some instructions and registers are disabled by the CPU
 - Exception registers
 - HALT instruction
 - MMU instructions
 - OS memory access
 - I/O device access
- **Kernel mode: ring 0 – trusted**
 - All instructions and registers enabled

October 3, 2018
TCCS422: Operating Systems [Fall 2018]
School of Engineering and Technology, University of Washington - Tacoma
L3.48

SYSTEM CALLS

- Implement restricted “OS” operations
- Kernel exposes key functions through an API:
 - Device I/O (e.g. file I/O)
 - Task swapping: context switching between processes
 - Memory management/allocation: malloc()
 - Creating/destroying processes

October 3, 2018 TCCS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma L3.49

TRAPS: SYSTEM CALLS, EXCEPTIONS, INTERRUPTS

```

Mainline Code
loop() {
  instruction 1
  instruction 2
  instruction 3
  instruction 4
  instruction 5
}

Interrupt Service Routine
ISR() {
  instruction 1
  instruction 2
  instruction 3
}
            
```

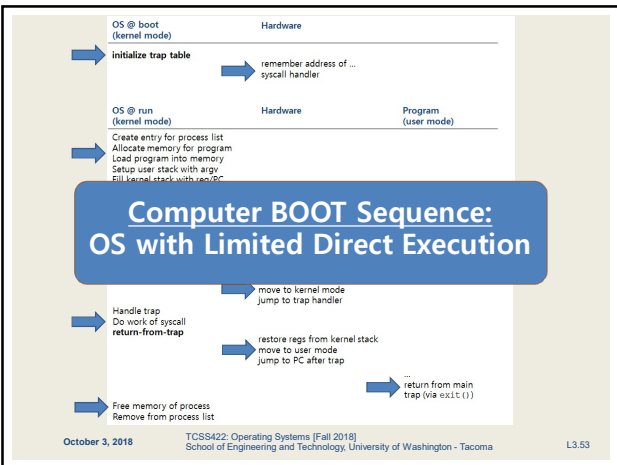
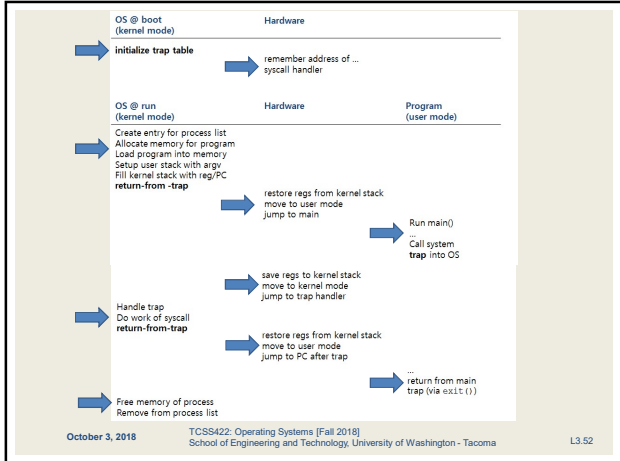
- Trap: any transfer to kernel mode
- Three kinds of traps
 - System call:** (planned) user → kernel
 - SYSCALL for I/O, etc.
 - Exception:** (error) user → kernel
 - Div by zero, page fault, page protection error
 - Interrupt:** (event) user → kernel
 - Non-maskable vs. maskable
 - Keyboard event, network packet arrival, timer ticks
 - Memory parity error (ECC), hard drive failure

October 3, 2018 TCCS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma L3.50

EXCEPTION TYPES

Exception type	Synchronous vs. asynchronous	User request vs. system	User maskable vs. nonmaskable	Within vs. between instructions	Resume vs. terminate
I/O device request	Asynchronous	Coerced	Nonmaskable	Between	Resume
Invoke operating system	Synchronous	User request	Nonmaskable	Between	Resume
Trapping instruction execution	Synchronous	User request	User maskable	Between	Resume
Breakpoint	Synchronous	User request	User maskable	Between	Resume
Integer arithmetic overflow	Synchronous	Coerced	User maskable	Within	Resume
Trapping point arithmetic overflow or underflow	Synchronous	Coerced	User maskable	Within	Resume
Page fault	Synchronous	Coerced	Nonmaskable	Within	Resume
Misaligned memory accesses	Synchronous	Coerced	User maskable	Within	Resume
Memory protection violation	Synchronous	Coerced	Nonmaskable	Within	Resume
Using undefined instruction	Synchronous	Coerced	Nonmaskable	Within	Terminate
Hardware malfunction	Asynchronous	Coerced	Nonmaskable	Within	Terminate
Power failure	Asynchronous	Coerced	Nonmaskable	Within	Terminate

October 3, 2018 TCCS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma L3.51



MULTITASKING

- How/when should the OS regain control of the CPU to switch between processes?
- Cooperative multitasking (mostly pre 32-bit)
 - < Windows 95, Mac OSX
 - Opportunistic: running programs must give up control
 - User programs must call a special **yield** system call
 - When performing I/O
 - Illegal operations
- (POLLEV)
 What problems could you see with this approach?

October 3, 2018 TCCS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma L3.54

