

TCSS 422: OPERATING SYSTEMS

Process API,  
Limited Direct Execution



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


October 3, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.1

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, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.3

OBJECTIVES

- Process API – Ch. 5
- Limited Direct Execution – Ch. 6

October 3, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.2

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, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.4

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, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.5

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

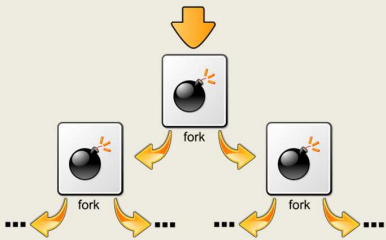


October 3, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.7

:() { | : & } ; :



October 3, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.6

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, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.8

Slides by Wes J. Lloyd

L3.3

Slides by Wes J. Lloyd

L3.4

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, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.9

exec()

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

October 3, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.11

FORK EXAMPLE

- Linux example

October 3, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.10

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 “black box”
- We don't want to know what is inside... 😊

```
graph LR
    Input --> BlackBox[Black Box]
    BlackBox --> Output
    subgraph Unknown
        BlackBox
    end
```

October 3, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.12

## 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, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.13

## 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, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.15

## 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, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.14

## 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, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.16

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 {
    int wc = wait(NULL); // parent goes down this path (main)
}
return 0;
}
```

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

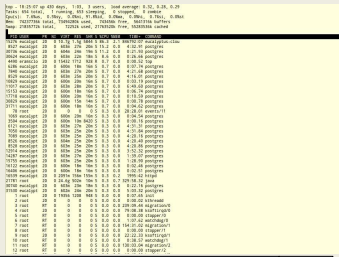
October 3, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.17

VIRTUALIZING THE CPU

- How does the CPU support running so many jobs simultaneously?
- Time Sharing**
- Tradeoffs:**
  - Performance
    - Excessive overhead
  - Control
    - Fairness
    - Security
- Both HW and OS support is used

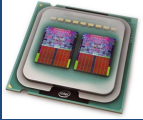


October 3, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.19

LIMITED DIRECT EXECUTION



October 3, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.18

DIRECT EXECUTION

- What if programs could directly control the CPU / system?

OS	Program
1. Create entry for process list 2. Allocate memory for	
<div>Computer BOOT Sequence: OS with Direct Execution</div>	
6. Execute call main()	8. Execute return from main()
9. Free memory of process 10. Remove from process list	

October 3, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.20

Slides by Wes J. Lloyd

L3.9

Slides by Wes J. Lloyd

L3.10

## 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, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.21

## CONTEXT SWITCHING OVERHEAD



October 3, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.23

## CONTROL TRADEOFF

### ■ Too much control:

- No security
- No time sharing

### ■ Too little control:

- Too much OS overhead
- Poor performance for compute & I/O
- Complex APIs (system calls), difficult to use

October 3, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.22

## LIMITED DIRECT EXECUTION

- OS implements LDE to support time/resource sharing
- Enabled by **protected (safe) control transfer**
- CPU supported context switch
- Provides data isolation

October 3, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.24

### 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, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.25

### SYSTEM CALLS

- Enable restricted “OS” operations
- Kernel exposes key functions through an API:
  - Device I/O
  - Task swapping: context switch
  - Memory management/allocation: malloc()
  - Creating/destroying processes

October 3, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.27

### 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, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.26

### TRAPS: SYSTEM CALLS, EXCEPTIONS, INTERRUPTS

- Trap: any transfer to kernel mode
- Three kinds of traps
  - Sys 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

```
graph LR
    subgraph Mainline_Code [Mainline Code]
        direction TB
        I1[Instruction 1]
        I2[Instruction 2]
        I3[Instruction 3]
        I4[Instruction 4]
        I5[Instruction 5]
    end
    subgraph ISR [Interrupt Service Routine]
        direction TB
        IS1[Instruction 1]
        IS2[Instruction 2]
        IS3[Instruction 3]
    end
    Mainline_Code -- "Interrupt" --> ISR
    ISR -- "Return" --> Mainline_Code
```

October 3, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.28

EXCEPTION TYPES					
Exception type	Synchronous vs. asynchronous	User request vs. coerced	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
Tracing 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
Floating-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, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.29

MULTITASKING

How/when should the OS regain control of the CPU to switch between processes?

Cooperative multitasking (mostly pre 32 bit)

- < Y
- Op
  - When performing I/O
  - Illegal operations

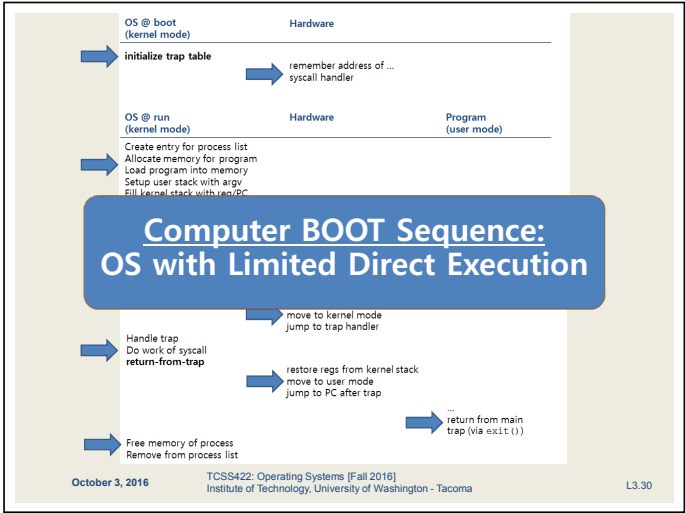
A process gets stuck in an infinite loop.  
→ Reboot the machine

What problems could you see with this approach?

October 3, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.31



MULTITASKING - 2

Preemptive multitasking (32 & 64 bit OSes)

>= Mac OSX, Windows 95+

Timer

- Raise
- Inter
  - 1. Current program is halted
  - 2. Program states are saved
  - 3. OS Interrupt handler is run (kernel mode)

A timer interrupt gives OS the ability to run again on a CPU.

What is a good interval for the timer interrupt?

October 3, 2016

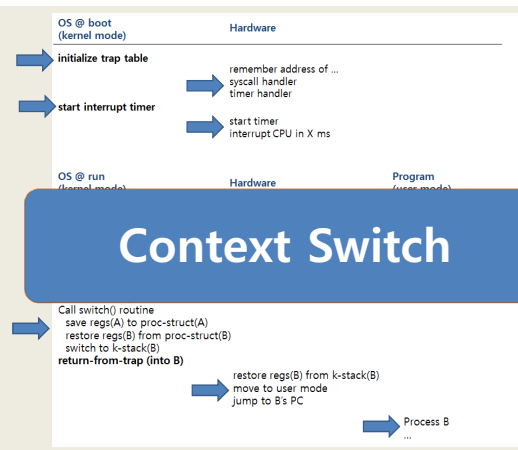
TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.32



# CONTEXT SWITCH

- Preemptive multitasking initiates “trap” into the OS code to determine:
  - Whether to continue running the **current process**, or switch to a **different one**.
  - If the decision is made to switch, the OS performs a context switch swapping out the current process for a new one.



# CONTEXT SWITCH - 2

1. Save register values of the current process to its kernel stack
  - General purpose registers
  - PC: program counter (instruction pointer)
  - kernel stack pointer
2. Restore soon-to-be-executing process from its kernel stack
3. Switch to the kernel stack for the soon-to-be-executing process

# INTERRUPTED INTERRUPTS

- What happens if during an interrupt (trap to kernel mode), another interrupt occurs?
- Linux
  - < 2.6 kernel: non-preemptive kernel
  - >= 2.6 kernel: preemptive kernel

PREEMPTIVE KERNEL


- Use “locks” as markers of regions of non-preemptibility (non-maskable interrupt)
- Preemption counter (`preempt_count`)
  - begins at zero
  - increments for each lock acquired (not safe to preempt)
  - decrements when locks are released
- Interrupt can be interrupted when `preempt_count=0`
  - It is safe to preempt (maskable interrupt)
  - the interrupt is more important

October 3, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L3.37

QUESTIONS



October 3, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L1.38