

CS 471 Operating Systems

Yue Cheng

George Mason University
Fall 2019

Announcement

- OS/161 Project 0 released on Blackboard
- Please complete the Google Form for OS/161 team composition

What is a Process?

What is a Process?

- **Programs** are code (static entity)
- **Processes** are running programs

- Java analogy
 - class -> “program”
 - object -> “process”

What is in a Process?


Process



What things change as a program runs?

What is in a Process?


Process



What things change as a program runs?

What is in a Process?


Process



What things change as a program runs?

What is in a Process?

Process



What things change as a program runs?


Peeking Inside

- Processes share code, but each has its own “context”
- CPU
 - Instruction pointer (Program Counter)
 - Stack pointer
- Memory
 - Set of memory addresses (“address space”)
 - `cat /proc/<PID>/maps`
- Disk
 - Set of file descriptors
 - `cat /proc/<PID>/fdinfo/*`

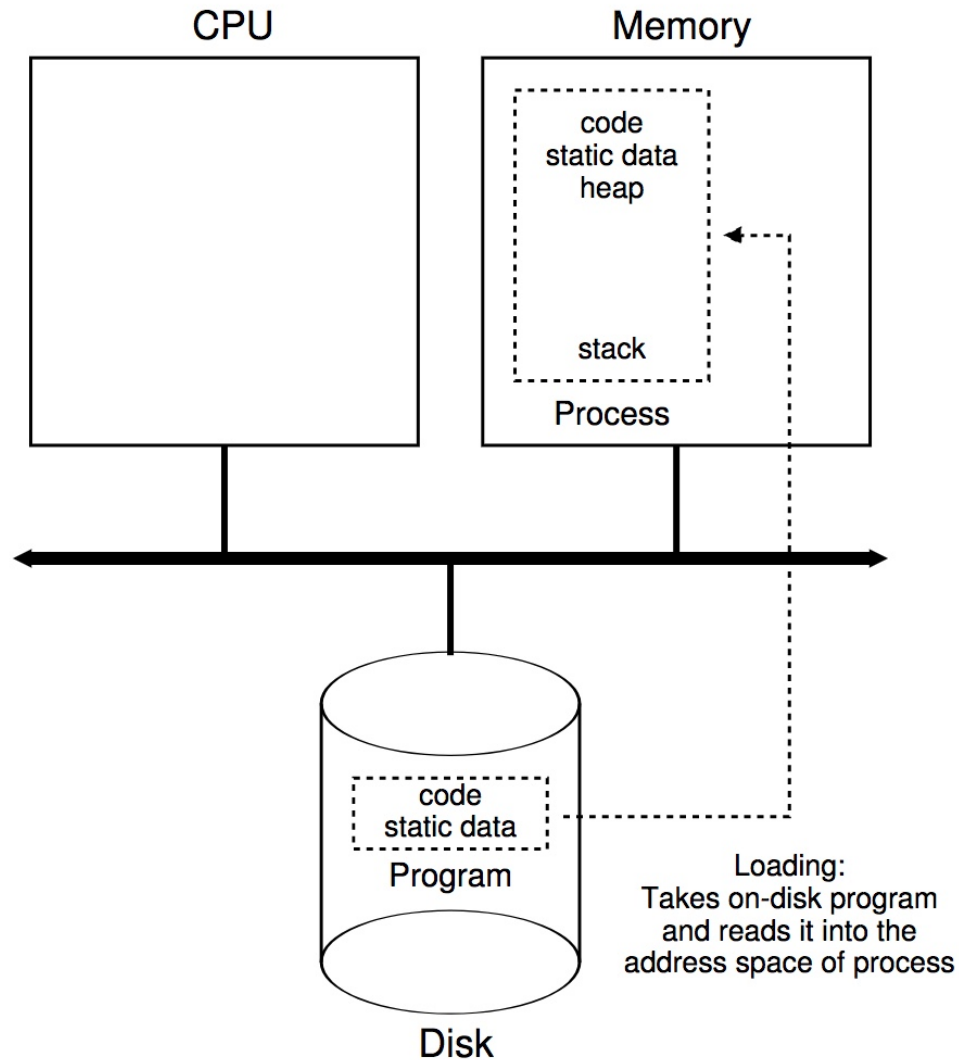
Process Creation

- Principle events that cause process creation
 - System initialization
 - Execution of a process creation system call by a running process
 - User request to create a process

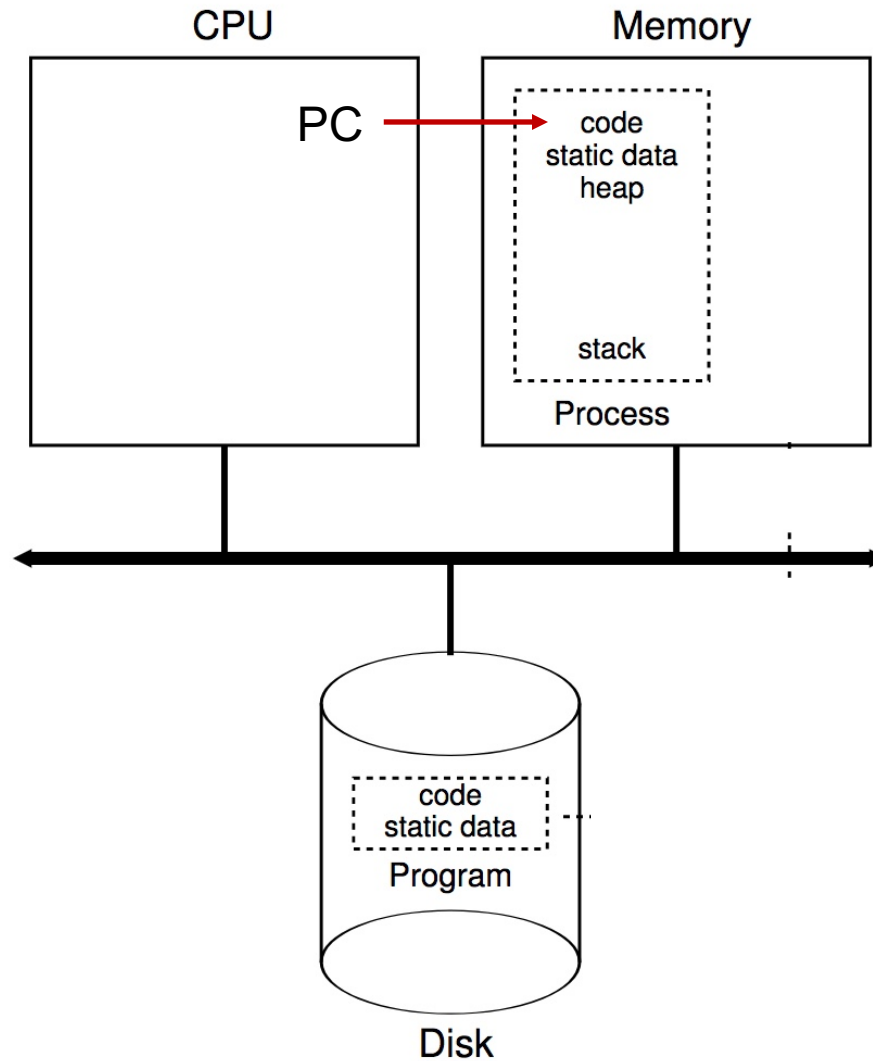
Process Creation



Process Creation




Process Creation



Process Creation (cont.)

- Parent process creates children processes, which, in turn create other processes, forming a tree (**hierarchy**) of processes
- **Questions:**
 - Will the parent and child execute **concurrently**?
 - How will the **address space** of the child be related to that of the parent?
 - Will the parent and child **share some resources**?

An Example Process Tree



How to View Process Tree in Linux?

- `% ps auxf`
 - 'f' is the option to show the process tree
- `% pstree`

Process Creation in Linux

- Each process has a **process identifier (pid)**
- The parent executes `fork()` system call to spawn a child
- The child process has a **separate copy** of the parent's address space
- Both the parent and the child continue execution at the instruction following the `fork()` system call
- The return value for the `fork()` system call is
 - **zero** value for the new (**child**) process
 - **non-zero pid** for the **parent** process
- Typically, a process can execute a system call like `exec1()` to load a binary file into memory

This is really the pid of the child process

Simply the return value of `fork()` in the context of the new child proc

man page of fork ()

<http://man7.org/linux/man-pages/man2/fork.2.html>

RETURN VALUE [top](#)

On success, the PID of the child process is returned in the parent, and 0 is returned in the child. On failure, -1 is returned in the parent, no child process is created, and [errno](#) is set appropriately.

ERRORS [top](#)

EAGAIN A system-imposed limit on the number of threads was encountered. There are a number of limits that may trigger this error:

- * the **RLIMIT_NPROC** soft resource limit (set via [setrlimit\(2\)](#)), which limits the number of processes and threads for a real user ID, was reached;
- * the kernel's system-wide limit on the number of processes and threads, [/proc/sys/kernel/threads-max](#), was reached (see [proc\(5\)](#));
- * the maximum number of PIDs, [/proc/sys/kernel/pid_max](#), was reached (see [proc\(5\)](#)); or
- * the PID limit ([pids.max](#)) imposed by the cgroup "process number" (PIDs) controller was reached.

Example Program with `fork()`

```
void main () {
    int pid;

    pid = fork();
    if (pid < 0) { /* error_msg */}
    else if (pid == 0) { /* child process */
        execl("/bin/lis", "lis", NULL); /* execute lis */
    } else { /* parent process */
        /* parent will wait for the child to complete */
        wait(NULL);
        exit(0);
    }
    return;
}
```

A Very Simple Shell using `fork()`

```
while (1) {
    type_prompt();
    read_command(cmd);
    pid = fork();
    if (pid < 0) { /* error_msg */
    else if (pid == 0) { /* child process */
        execute_command(cmd);
    } else { /* parent process */
        wait(NULL);
    }
}
```

More example: fork 1

```
forkexample.c *
1  #include <sys/types.h>
2  #include <stdio.h>
3  #include <stdlib.h>
4  #include <unistd.h>
5
6  int number = 7;
7
8  int main(void) {
9      pid_t pid;
10     printf("\nRunning the fork example\n");
11     printf("The initial value of number is %d\n", number);
12
13     pid = fork();
14     printf("PID is %d\n", pid);
15
16     if (pid == 0) {
17         number *= number;
18         printf("\tIn the child, the number is %d -- PID is %d\n", number, pid);
19         return 0;
20     } else if (pid > 0) {
21         wait(NULL);
22         printf("In the parent, the number is %d\n", number);
23     }
24
25     return 0;
26 }
27
```

What happens to the value of number?

Results

```
./forkexample1
```

Running the fork example

The initial value of number is 7

PID is 2137

PID is 0

In the child, the number is 49 -- PID is 0

In the parent, the number is 7

Further more example: fork 2

```
forkexample2.c *
1 #include <sys/types.h>
2 #include <stdio.h>
3 #include <stdlib.h>
4 #include <unistd.h>
5
6 int number = 7;
7
8 int main(void) {
9     pid_t pid;
10    printf("\nRunning the fork example\n");
11    printf("The initial value of number is %d\n", number);
12
13    pid = fork();
14    printf("PID is %d\n", pid);
15
16    if (pid == 0) {
17        number *= number;
18        fork();
19        printf("\tIn the child, the number is %d -- PID is %d\n", number, pid);
20        return 0;
21    } else if (pid > 0) {
22        wait(NULL);
23        printf("In the parent, the number is %d\n", number);
24    }
25
26    return 0;
27 }
28
```

What happens to the value of number?

Results

```
./forkexample2
```

Running the fork example

The initial value of number is 7

PID is 2164

PID is 0

In the child, the number is 49 -- PID is 0

In the child, the number is 49 -- PID is 0

In the parent, the number is 7

exec1 vs. fork

```
execlexample.c *
1  #include <sys/types.h>
2  #include <stdio.h>
3  #include <stdlib.h>
4  #include <unistd.h>
5
6  int number = 7;
7
8  int main(void) {
9      pid_t pid;
10     printf("\nRunning the exec1 example\n");
11     pid = fork();
12     printf("PID is %d\n", pid);
13
14     if (pid == 0) {
15         printf("\tIn the exec1 child, PID is %d\n", pid);
16         execl("./forkexample2", "forkexample2", NULL);
17         return 0;
18     } else if (pid > 0) {
19         wait(NULL);
20         printf("In the parent, done waiting\n");
21     }
22
23     return 0;
24 }
```

Results

./execlexample

Running the execl example

PID is 2179

PID is 0

In the execl child, PID is 0

Running the fork example

The initial value of number is 7

PID is 2180

PID is 0

In the child, the number is 49 -- PID is 0

In the child, the number is 49 -- PID is 0

In the parent, the number is 7

In the parent, done waiting



forkexample2