ECE 353: Systems Software
Lecture 3

Process Information

1.0.0

Jon Eyolfson

January 12, 2023

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License

A Process is an Instance of a Running Program

A *program* (or application) is just a static definition, including:

- Instructions
- Data
- Memory allocations
- Symbols it uses

A *process* executes a program, and the kernel manages processes

Process is like a Combination of all the Virtual Resources

If we consider a “virtual CPU”, the OS needs to track all registers

It also contains all other resources it can access (memory and I/O)

Every execution runs the same code (part of the program)

- An execution is running some specific code (part of the process)

A Process is More Flexible

A process contains both the program and information specific to its execution

It allows multiple executions of the same program

It even allows a process to run multiple copies of itself

A Process Control Block (PCB) Contains All Execution Information

Specifically, in Linux, this is the `task_struct` you can browse on GitHub

It contains:

- Process state
- CPU registers
- Scheduling information
- Memory management information
- I/O status information
- Any other type of accounting information

Uniprogramming is for Old Batch Processing Operating Systems

Uniprogramming: only one process running at a time

- Two processes are not parallel and not concurrent, no matter what

Multiprogramming: allow multiple processes

- Two processes can run in parallel or concurrently

Modern operating systems try to run everything in parallel and concurrently

Process State Diagram (You Could Rename Waiting to Ready)

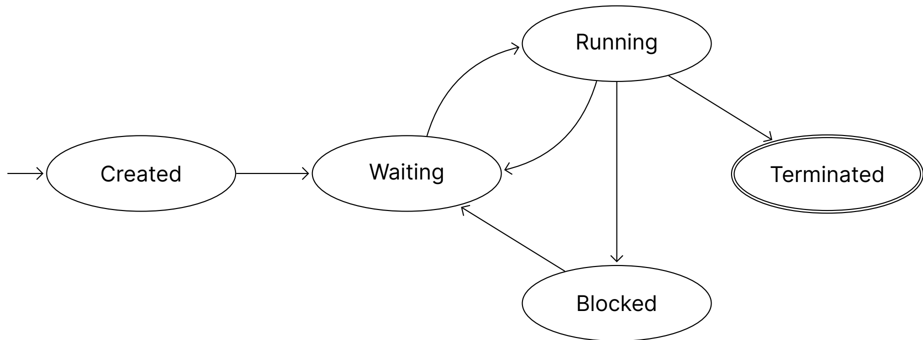

On Linux, You Can Read Process State Using the “proc” Filesystem

There’s a standard `/proc` directory that represents the kernel’s state

- These aren’t real files, they just look like it!

Every directory that’s a number (process ID) in `/proc` represents a process

There’s a file called `status` that contains the state (used for Lab 1)

The Scheduler Decides When To Switch

To create a process, the operating system has to at least load it into memory

When it’s waiting, the scheduler (coming later) decides when it’s running

We’re going to first focus on the mechanics of switching processes

The Core Scheduling Loop Changes Running Processes

1. Pause the currently running process
2. Save its state, so you can restore it later
3. Get the next process to run from the scheduler
4. Load the next process’ state and let that run

We Can Let Processes Themselves, or the Operating System Pause

Cooperative multitasking

The processes use a system call to tell the operating system to pause it

True multitasking

The operating system retains control and pauses processes

For true multitasking the operating system can:

- Give processes set time slices
- Wake up periodically using interrupts to do scheduling

Swapping Processes is called Context Switching

We've said that at minimum we'd have to save all of the current registers

- We have to save all of the values, using the same CPU as we're trying to save

There's hardware support for saving state, however you may not want to save everything

Context switching is pure overhead, we want it to be as fast as possible

Usually there's a combination of hardware and software to save as little as possible

System Calls are Rare in C

Mostly you'll be using functions from the C standard library instead

Most system calls have corresponding function calls in C, but may:

- Set `errno`
- Buffer reads and writes (reduce the number of system calls)
- Simplify interfaces (function combines two system calls)
- Add new features

C exit Has Additional Features

System call exit (or exit_group): the program stops at that point

C exit: there's a feature to register functions to call on program exit (atexit)

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

void fini(void) {
  puts("Do fini");
}

int main(int argc, char **argv) {
  atexit(fini);
  puts("Do main");
  return 0;
}
```

`execve` Replaces the Process with Another Program, and Resets

`execve` has the following API:

- `pathname`: Full path of the program to load
- `argv`: Array of strings (array of characters), terminated by a null pointer
  - Represents arguments to the process
- `envp`: Same as `argv`
  - Represents the environment of the process
- Returns an error on failure, does not return if successful

execve-example.c Turns the Process into ls

```c
int main(int argc, char *argv[]) {
  printf("I'm going to become another process\n");
  char *exec_argv[] = {"ls", NULL};
  char *exec_envp[] = {NULL};
  int exec_return = execve("/usr/bin/ls", exec_argv, exec_envp);
  if (exec_return == -1) {
    exec_return = errno;
    perror("execve failed");
    return exec_return;
  }
  printf("If execve worked, this will never print\n");
  return 0;
}
```

The Operating System Creates and Runs Processes

The operating system has to:

- Loads a program, and create a process with context
- Maintain process control blocks, including state
- Switch between running processes using a context switch
- Replace programs running in processes (Unix)