Processes

CSE 4001 Operating Systems Concepts

E. Ribeiro

January 26, 2022

Outline

A process is a program in execution

What is a process?

A *process* is an abstraction of *a program in execution*.

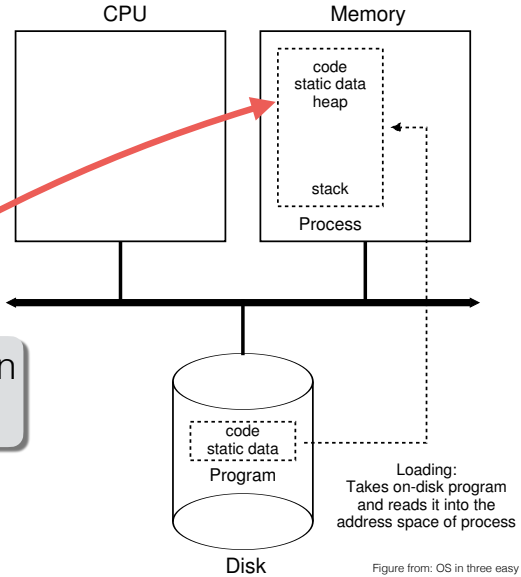

Figure from: OS in three easy pieces

Main components of the address space

What is a process?

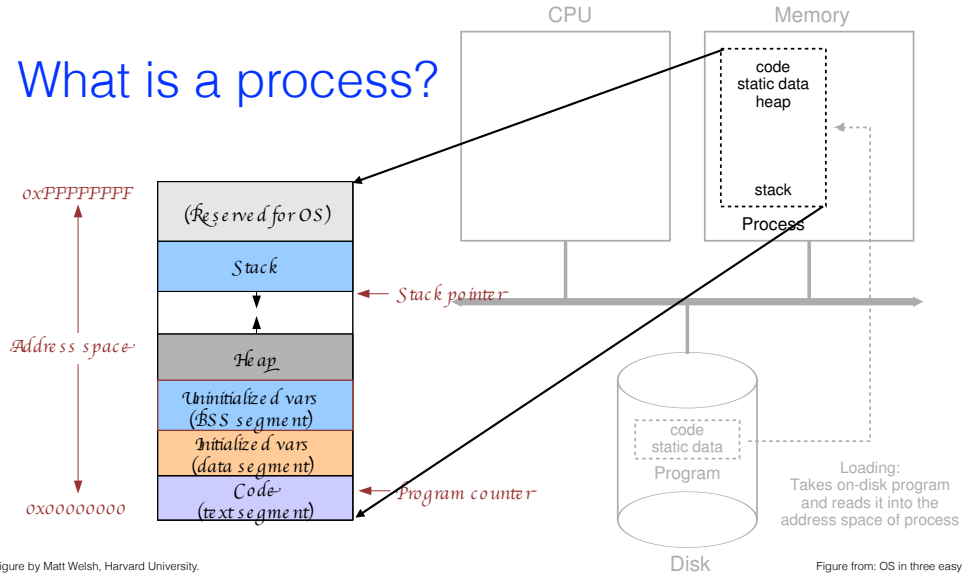

Figure by Matt Welsh, Harvard University.

Figure from: OS in three easy pieces

The code part of the address space

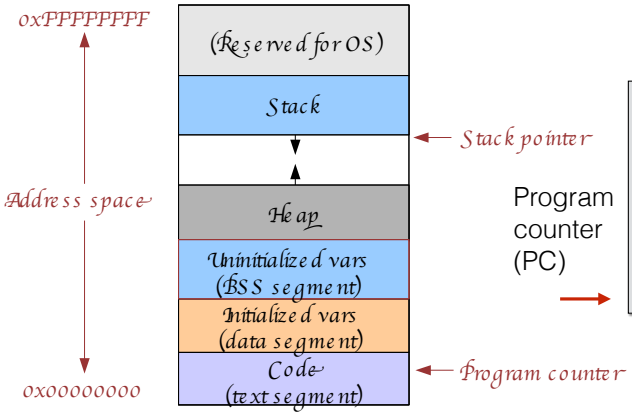

Program counter (PC)

```
/* Hello World program */

#include<stdio.h>

main()
{
    printf("Hello World");

}
```

Figure by Matt Welsh, Harvard University.

The OS view of a process

- Process state (ready, running, blocked, ...)
- The **address space** (how many possible addresses)
- The **code** of the running program
- The **data** of the running program
- An execution **stack** encapsulating the state of procedure calls
- The **program counter (PC)** indicating the address of the next instruction.
- A set of general-purpose **registers** with current values
- A set of operating system **resources**
  - open files, network connections, signals, etc.
- CPU scheduling info: process **priority**
- Each process is identified by its **process ID (PID)**

All these information is stored in a construct called **Process Control Block (PCB)**

The Process Control Block (PCB)

The OS maintains a PCB for each process. It is a data structure with many fields.

Defined in:
`/include/linux/sched.h`

Figure by Matt Welsh, Harvard University.

Life cycle of a process

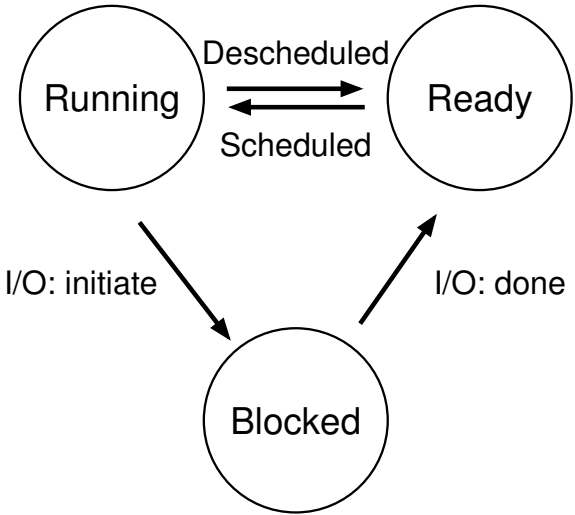

Figure from: OS in three easy pieces

Two processes running, no I/O

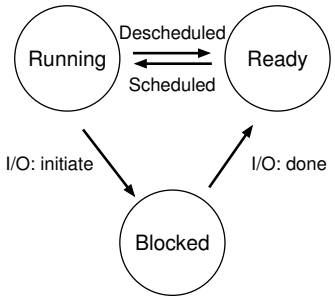

Time	Process_0	Process_1	Notes
1	Running	Ready	
2	Running	Ready	
3	Running	Ready	
4	Running	Ready	Process_0 now done
5	–	Running	
6	–	Running	
7	–	Running	
8	–	Running	Process_1 now done

Figure from: OS in three easy pieces

Two processes running, with I/O

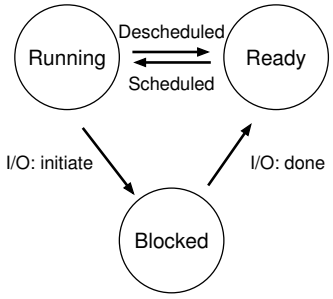

Time	Process_0	Process_1	Notes
1	Running	Ready	
2	Running	Ready	
3	Running	Ready	Process_0 initiates I/O
4	Blocked	Running	Process_0 is blocked,
5	Blocked	Running	so Process_1 runs
6	Blocked	Running	
7	Ready	Running	I/O done
8	Ready	Running	Process_1 now done
9	Running	–	
10	Running	–	Process_0 now done

Figure from: OS in three easy pieces

Process states (Unix)

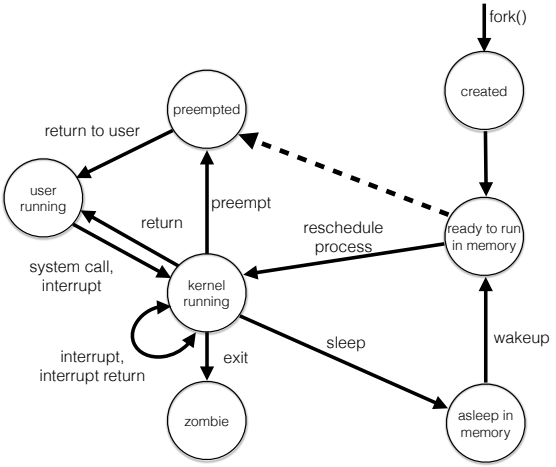

Figure adapted from Stallings' book

Created: Process is newly created but it is not ready to run yet.

Preempted: Process is returning from kernel to user mode, but the kernel preempts it and does a process switch to schedule another process.

Zombie: Process is no longer exists, but it leaves a record for its parent process to collect.

Process states (Unix) without hard drive

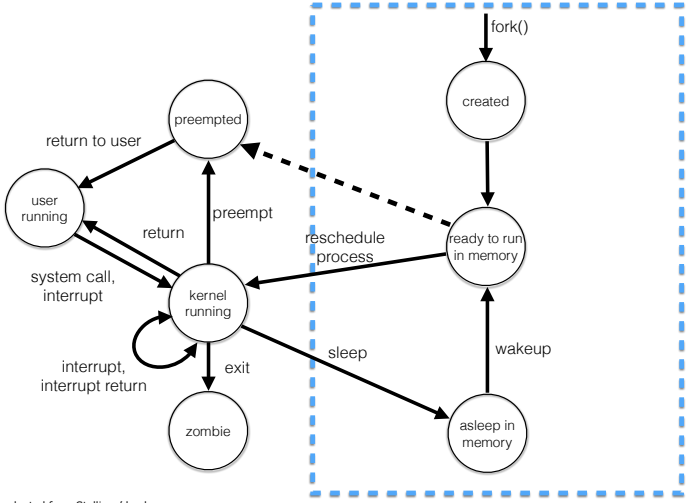

Figure adapted from Stallings' book

Process states (Unix) with hard drive

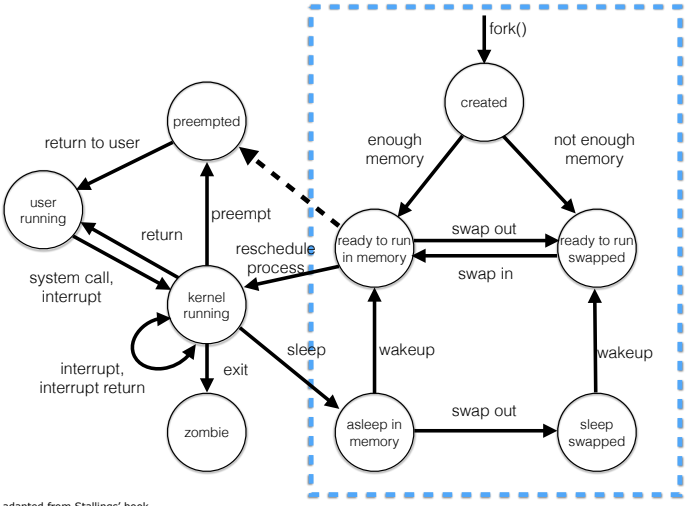

Figure adapted from Stallings' book

Ready queue and various I/O queue: process waiting

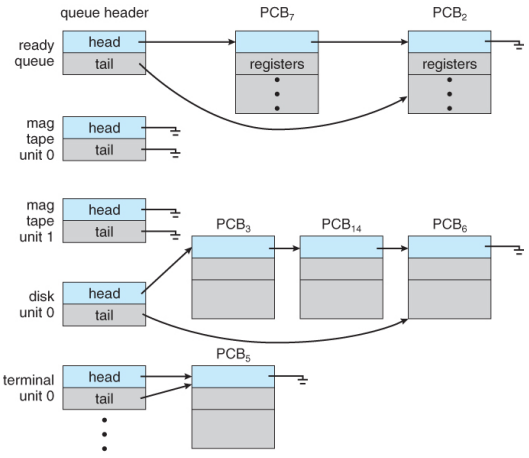

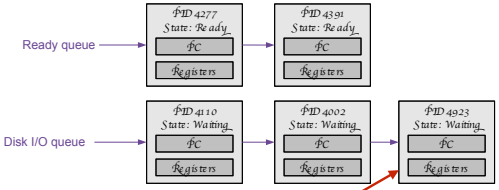

- OS maintains a set of queues
- Each PCB is queued on a state queue based on the process' current state.
- As processes change states, PCBs are unlinked from one queue and linked into another.

Adapted from Silberschatz, Galvin, and Gagne, 2009.

Ready queue and various I/O queue: process moved to ready

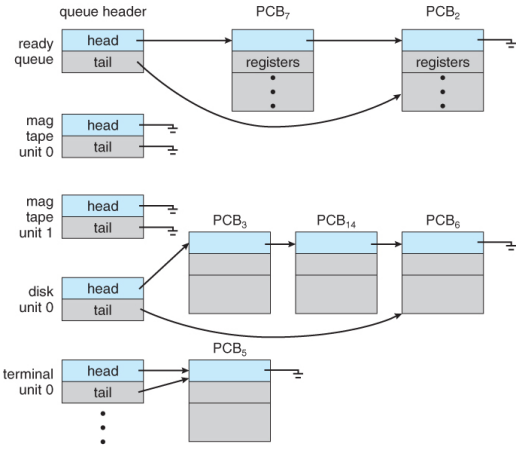

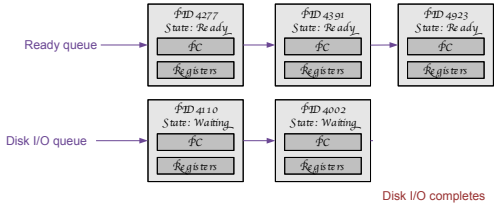

- OS maintains a set of queues
- Each PCB is queued on a state queue based on the process' current state.
- As processes change states, PCBs are unlinked from one queue and linked into another.

Adapted from Silberschatz, Galvin, and Gagne, 2009.