


Starting Reading #3

More OS – processes

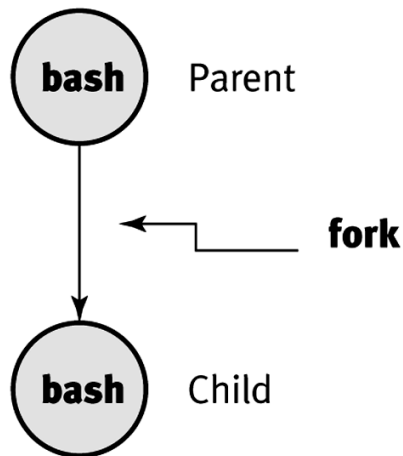
Processes

- A process is an executable, machine language program that the OS (Linux) has been asked to run
 - Copied to memory, and assigned a process ID (PID)
 - Scheduled for execution by the CPU
- Processes create other processes via system calls
 - A program (e.g., in C or C++) creates a new process and terminates itself with a call to `exec`
 - A program creates a child process by calling `fork`
 - e.g.: `$> ./myscript`
 - First line is: `#!/bin/bash`
 - `bash` runs (interprets script)




Steps to execute a program (sort)

Step 1: Shell uses **fork** to create a child



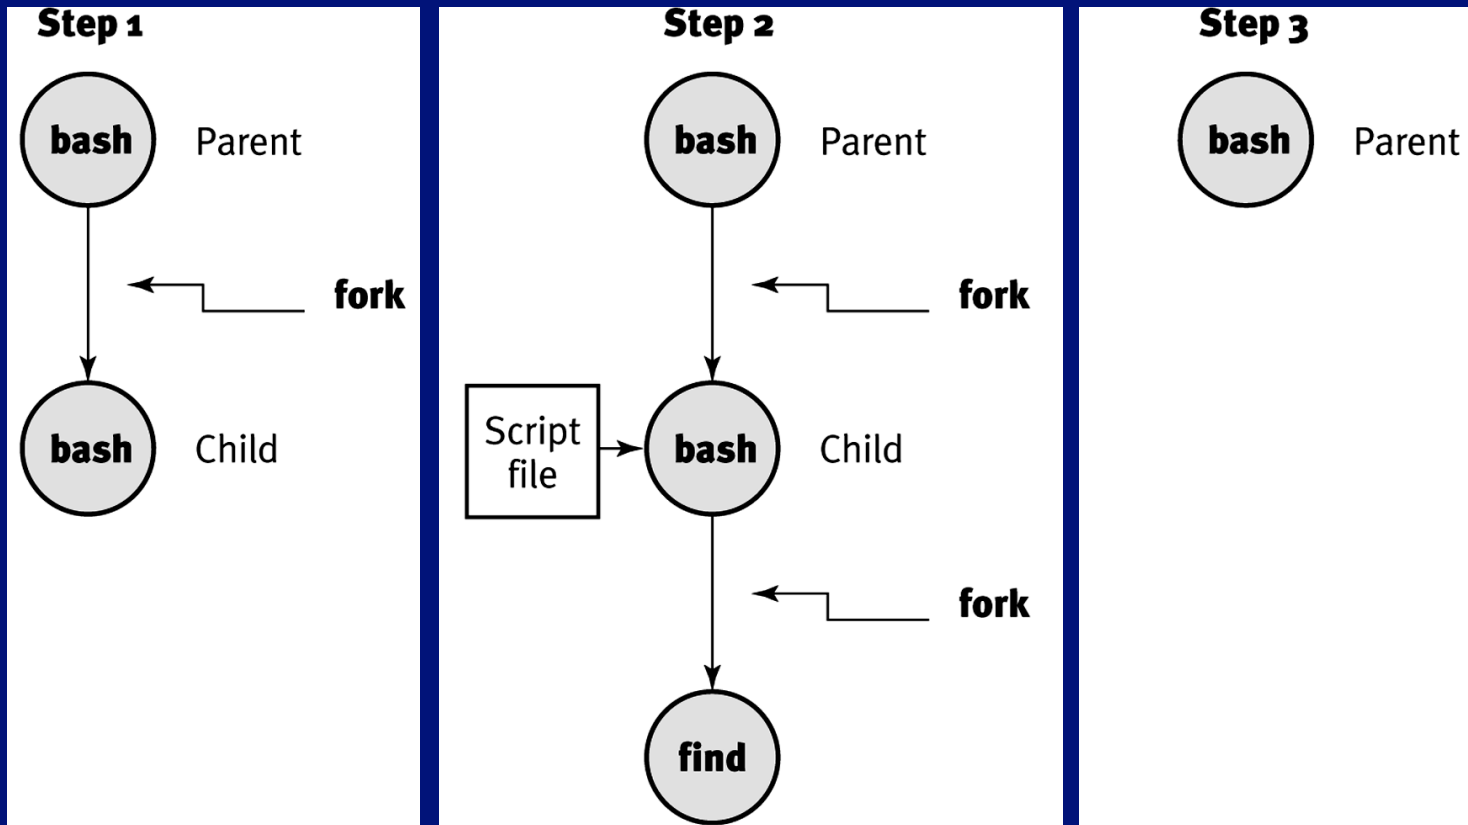
Step 2: Child uses **exec** to overwrite itself with the executable file corresponding to the **sort** command.



Step 3: **sort** starts execution while 'bash' waits for the command to finish. When **sort** finishes, the child process terminates and 'bash' starts execution again, waiting for the user to give it another command to execute.



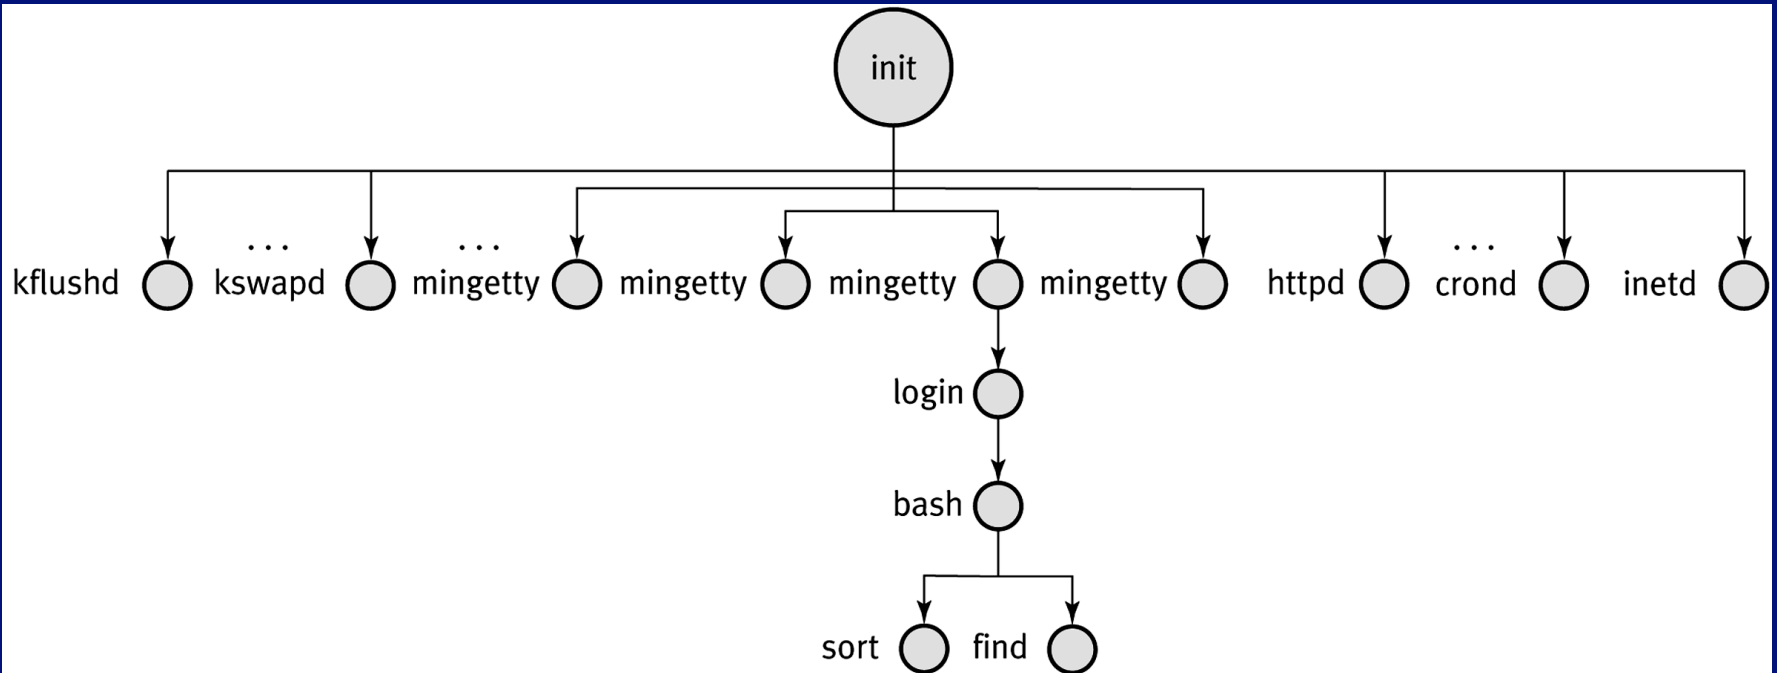
Steps to execute a shell script



Process hierarchy

- init – is PID 1, but all other processes have parents (so PPID)
 - The process hierarchy's depth is limited only by available virtual memory
- A process may control the execution of any of its descendants
 - Can suspend or resume it
 - Can alter its relative priority
 - Can even terminate it completely
- By default, terminating a process will terminate all of its descendants too
 - So terminating the root process will terminate the session

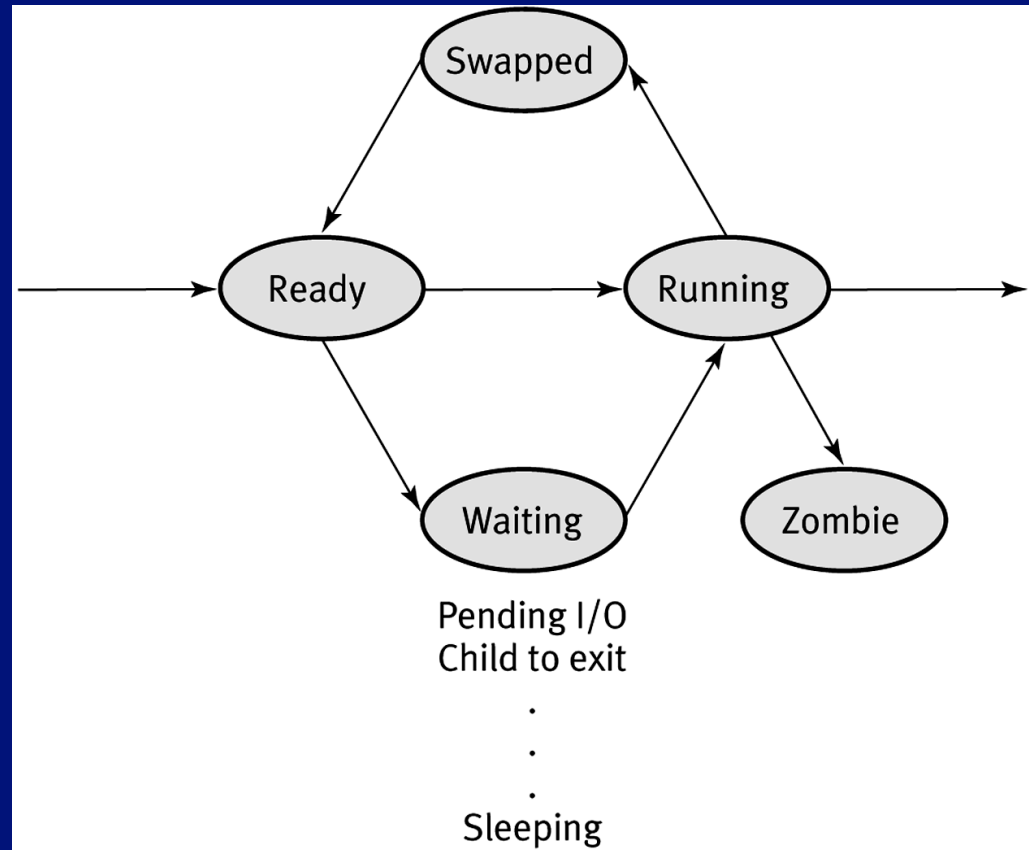
Example Linux process hierarchy



(From Linux-specific version of Sarwar et al. text)

Linux process states

- Just one process can be "running" at any one time
- OS has other processes in various states
- A process may be cycled through many states before it terminates

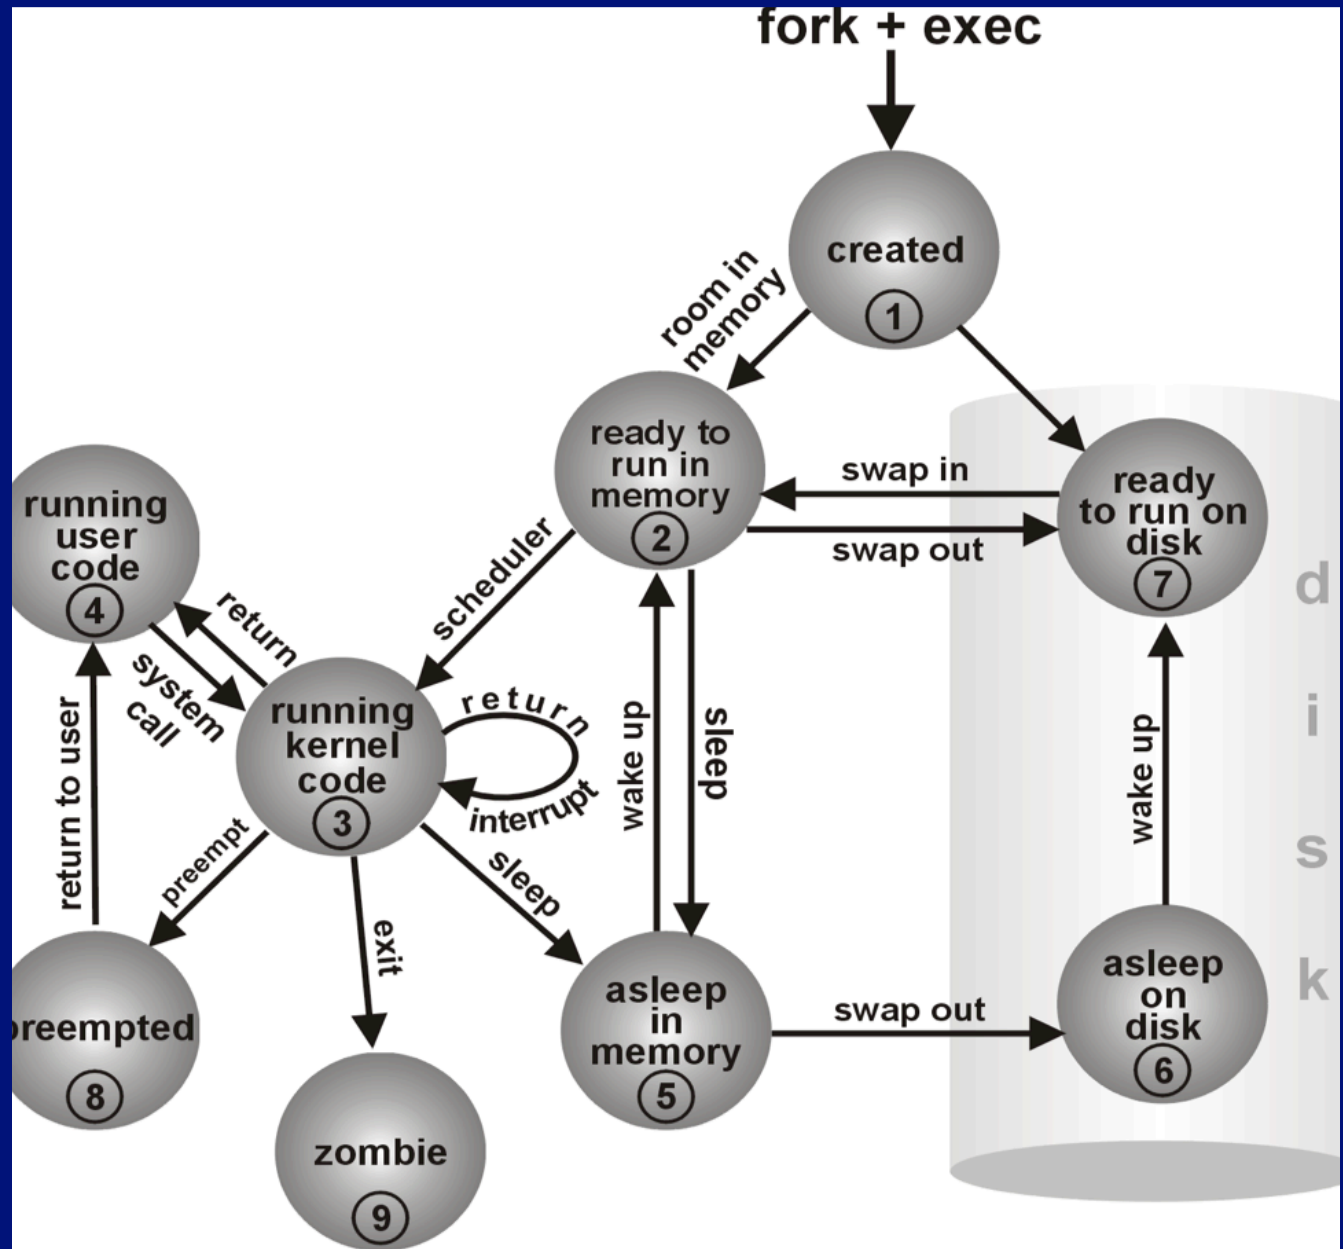


Meanings of Linux process states

State	Description
Ready	The process is ready to run but doesn't have the CPU. Based on the scheduling algorithm, the scheduler decided to give the CPU to another process. Several processes can be in this state, but on a machine with a single CPU, only one can be executing (using the CPU).
Running	The process is actually running (using the CPU).
Waiting	The process is waiting for an event. Possible events are an I/O (e.g., disk/terminal read or write) is completed, a child process exits (parent waits for one or more of its children to exit), or the sleep period expires for the process.
Swapped	The process is ready to run, but it has been temporarily put on the disk (on the swap space); perhaps it needs more memory and there isn't enough available at this time.
Zombie	When the parent of a process terminates before it executes the exit call, it becomes a zombie process. The process finishes and finds that the parent isn't waiting. The zombie processes are finished for all practical purposes and don't reside in the memory, but they still have some kernel resources allocated to them and can't be taken out of the system. All zombies (and their live children) are eventually adopted by the granddaddy, the init process, which removes them from the system. In general, any dying process is said to be in the zombie state.

More states

From: *Bulletproof Unix* by Tim Gottleber, 2003



Foreground and background

- When a command is executed from the prompt and runs to completion at which time the prompt returns, it is said to run in the **foreground**
- When a command is executed from the prompt followed by the token '&' on the command line, the prompt immediately returns while the command is said to run in the **background**
- Programs that interact with a user should be run in the foreground
- Programs that execute slowly and without intervention belong in the background – so other work can get done!
 - e.g., daemons (background processes for system administration)

User control of process state

- Terminate a foreground process with ctrl-C
- Send running foreground process to background by ctrl-Z

```
-bash-4.2$ find / *.txt > /dev/null 2> /dev/null
```

← entered ctrl-Z here

```
[1]+  Stopped  find / *.txt > /dev/null 2> /dev/null
```

```
-bash-4.2$ ← can execute more commands while find works
```

– If enter fg 1 now, job 1 will execute in foreground again

- Use ps to find PIDs of running processes

```
-bash-4.2$ ps
```

PID	TTY	TIME	CMD
20637	pts/4	00:00:00	bash
21581	pts/4	00:00:02	find
21632	pts/4	00:00:00	ps

- Terminate a background process with kill command

```
bash-4.2$ kill -9 21581 ← -9 is the "sure kill" signal number
```

```
-[1]+  Killed      find / *.txt > /dev/null 2> /dev/null
```

Fields of `ps -l` output (cont. next slide)

Field	Meaning
F	Flags: Flags associated with the process. It indicates things like whether the process is a user or kernel process, and why the process stopped or went to sleep.
UID	User ID: Process owner's user ID
PID	Process ID: Process ID of the process
PPID	Parent PID: PID of the parent process
PRI	Priority: Priority number of a process that dictates when the process is scheduled.
NI	Nice value: The nice value of a process; another parameter used in the computation of a process's priority number.
VSZ	Virtual size: The number in this field is the size of the memory image of a process (code+data+stack) in blocks.

Fields of `ps -l` output (cont.)

RSS	Resident set size: The amount of physical memory in kilobytes; it does not include space taken by the page table and kernel task structure for the process.
WCHAN	Wait channel: Null for running processes, or processes that are ready to run and are waiting for the CPU to be given to them. For a waiting or sleeping process, this field lists the event the process is waiting for—the kernel function where the process resides.
STAT	Process state: See next slide.
TTY	Terminal: The terminal name a process is attached to
TIME	Time: The time (in minutes and seconds) a process has currently been running for, or previously ran for before sleeping or stopping.
COMMAND	Command: Lists the command line that was used to start this process. The <code>-f</code> option is needed to see the full command line; otherwise only the last component of the pathname is displayed.

Process state abbreviations

Process State	Meaning
D	Uninterruptible sleep (usually doing I/O or waiting for it)
N	Low-priority process (a process that has been niced)
R	Runnable process: waiting to be scheduled to use CPU
S	Sleeping
T	Traced or stopped
Z	A zombie (defunct) process
W	A process that is completely swapped on the disk (no resident pages)