

Process Groups & Terminal Control

Computer Operating Systems, Spring 2024

Instructor: Travis McGaha

Head TAs: Nate Hoaglund & Seungmin Han

TAs:

Adam Gorka	Haoyun Qin	Kyrie Dowling	Ryoma Harris
Andy Jiang	Jeff Yang	Oliver Hendrych	Shyam Mehta
Charis Gao	Jerry Wang	Maxi Liu	Tom Holland
Daniel Da	Jinghao Zhang	Rohan Verma	Tina Kokoshvili
Emily Shen	Julius Snipes	Ryan Boyle	Zhiyan Lu



Poll Everywhere

pollev.com/tqm

- ❖ How was parser and shredder?
Are you excited for more C and penn-shell? 😊

Administrivia

- ❖ Peer Evaluation: out now, due Saturday 2/10 @ 11:59 pm
 - Please do it, it shouldn't take long
 - Mostly completion, don't just say "this is fine" for everything

- ❖ Penn-shell is out!
 - Milestone is due a week from tomorrow (2/14 @ 11:59pm)
 - Full thing is due a week and half later (2/23 @ 11:59 pm)
 - Demo in second half of this class
 - Done in partners
 - Should have everything you need to complete the assignment in this class

Administrivia

- ❖ Recitation
 - On Monday!
 - A few tips about dealing with pipe() and the rest will be office hours to help finish milestone 1

- ❖ Partners have been randomly assigned!
 - If you need to contact your partner, let us know and we can email both of you

- ❖ Check-in Quiz due before next lecture
 - Will release soon

Penn-Shell Compatibility

❖ From the `signal(2)` man page

Portability

The only portable use of `signal()` is to set a signal's disposition to `SIG_DFL` or `SIG_IGN`. The semantics when using `signal()` to establish a signal handler vary across systems (and POSIX.1 explicitly permits this variation); do not use it for this purpose.

❖ If you want to have better help from TA's put this at the top of your file before you `#include` anything

- This **should** get signals to behave as we expect, so TAs can better help
- If you got it working another way, that is OK. Auto-grader **should** still accept it

```
#ifndef _POSIX_C_SOURCE
#define _POSIX_C_SOURCE 200809L
#endif

#ifndef _DEFAULT_SOURCE
#define _DEFAULT_SOURCE 1
#endif
```



pollev.com/tqm

- ❖ Any questions, comments or concerns from last lecture?

Lecture Outline

- ❖ **Process Groups**
 - `setpgid()`
- ❖ Terminal Control
 - `tcsetpgrp()`
- ❖ SIGSTOP
- ❖ Project 1: Synch vs Asynch wait
 - SIGCHLD

Process Groups

- ❖ Processes are associated together into Process Groups.
 - A process always is in a process group
- ❖ Allows for convenient process & signal management:
 - If ctrl + C (SIGINT) is sent to a process via the keyboard, it is also sent to all processes within its group
- ❖ When we create a process with `fork()`, the child belongs to the same process group as the parent
- ❖ Shell has the notion of a **job**: “commands” started interactively. All processes in a job are in the same group
- ❖ Relevant for **penn-shell**

Process Group ID

- ❖ The process group ID is equal to a process ID
 - The process ID of the first process to exist in the group
 - If a process group “leader” terminates, can its process ID be reused by another process? Even if the old group is still going?
 - Answer: no, that process ID will be reserved until the group is done

- ❖

```
int setpgid(pid_t pid, pid_t pgid);
```

- ❖ Sets page group id of the specified process to the new value
 - Only works if pgid specifies an existing process group
 - Or if pgid == pid, creates a new process group of that id

Process Group ID


- ❖ `pid_t getpgid(pid_t pid);`
- ❖ Gets the process group id of the specified process
- ❖ If `0` is passed in, get the group ID of the calling process
- ❖ `-1` returned on error

CTRL +C, same group


pgid = 100

User Processes

./example
pid = 100



/bin/sleep
pid = 101




OS

CTRL +C, same group


pgid = 100

User Processes

./example
pid = 100



/bin/sleep
pid = 101



OS

CTRL + C





CTRL +C, same group

pgid = 100

User Processes


./example
pid = 100

/bin/sleep
pid = 101



OS

CTRL + C

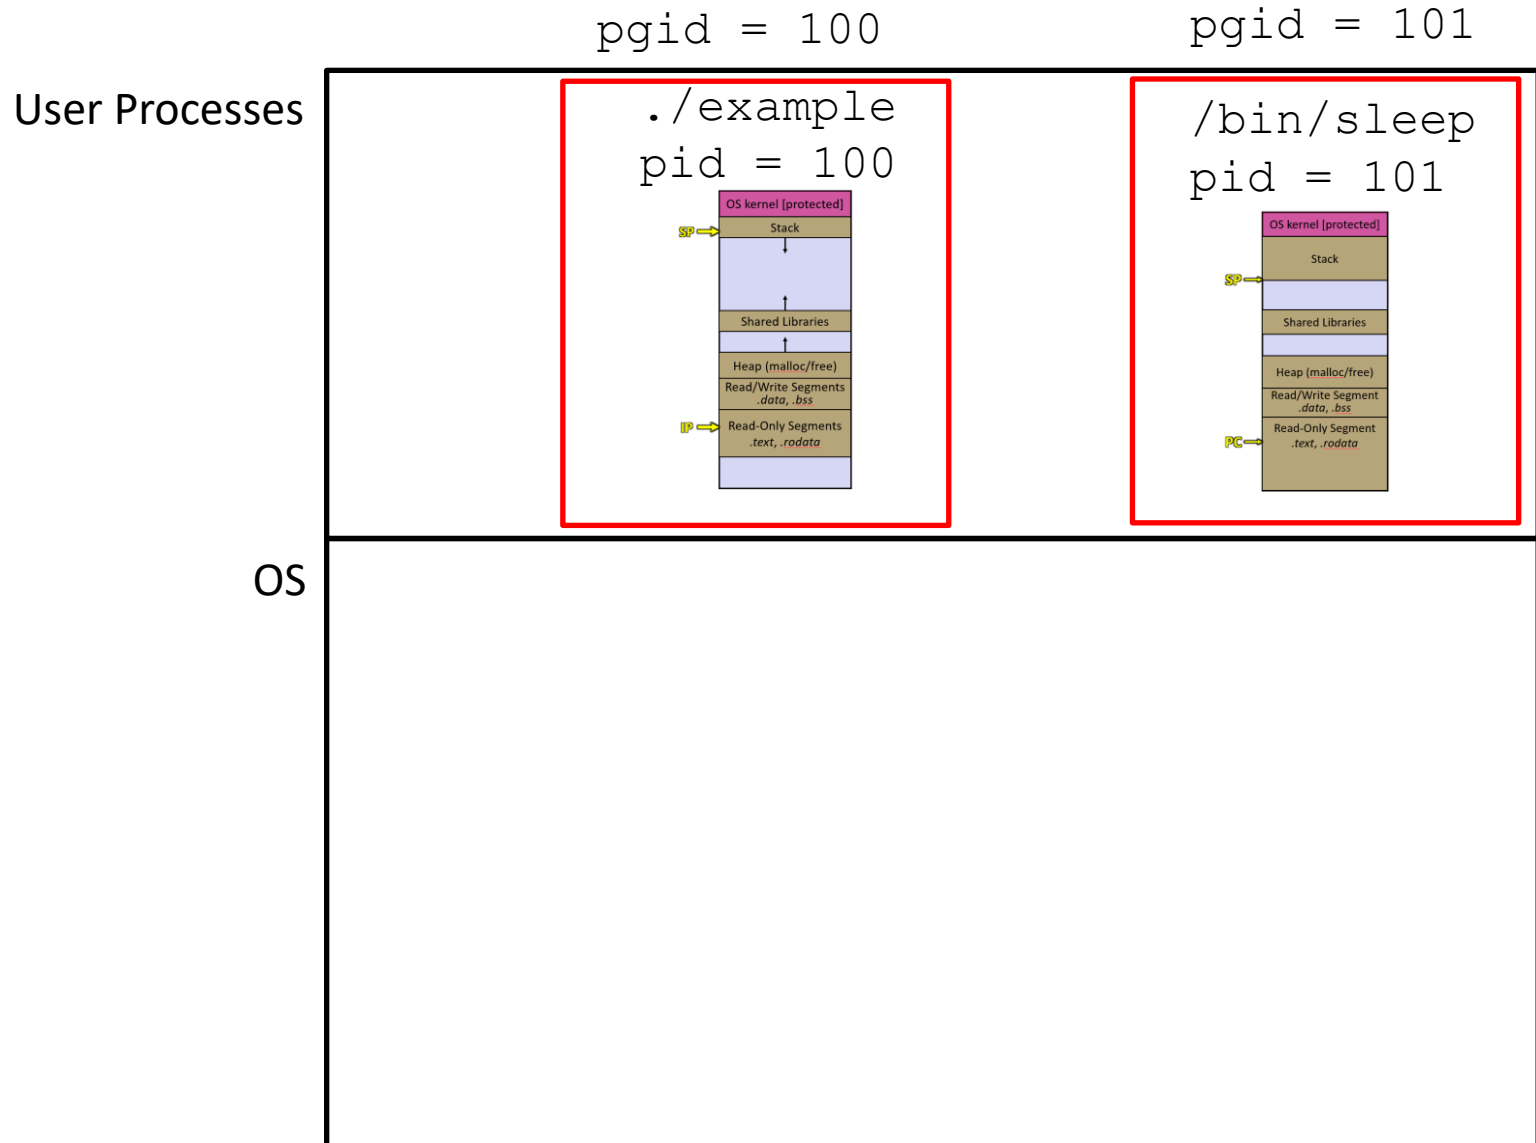


SIGINT is sent to every Process in the process group



GAP SLIDE: MOVING ON TO DIFFERENT EXAMPLE

CTRL +C, different group



CTRL +C, different group



User Processes

pgid = 100

pgid = 101

./example
pid = 100

/bin/sleep
pid = 101



OS

CTRL + C



SIGINT is sent to every Process in the process group

Child is in a separate group

GAP SLIDE: MOVING ON TO DIFFERENT EXAMPLE

Process Groups: utility

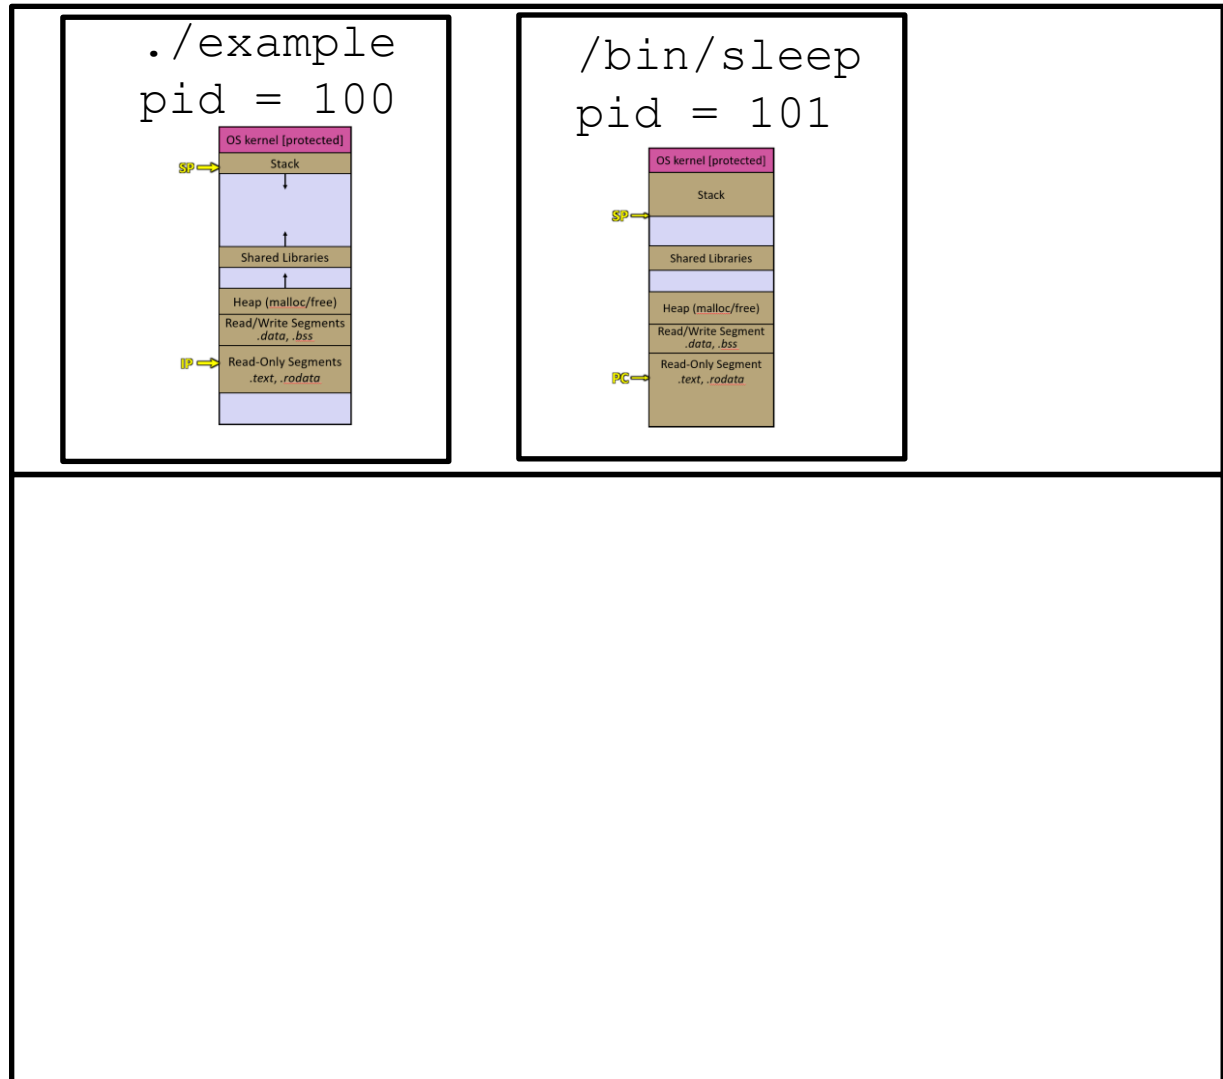
- ❖ Can pass in `-PGID` (negative PGID) to `kill()` and `waitpid()`
- ❖ Doing so for `kill()` will send the signal to all processes in the group
- ❖ Doing so for `waitpid()` will wait for any process in the group
- ❖ You may find this useful for proj1: `penn-shell`

Diagram: signals between process groups

User Processes

pgid = 100

pgid = 101



OS

Let's say we have the parent process on the left

a child process in its own group on the right.

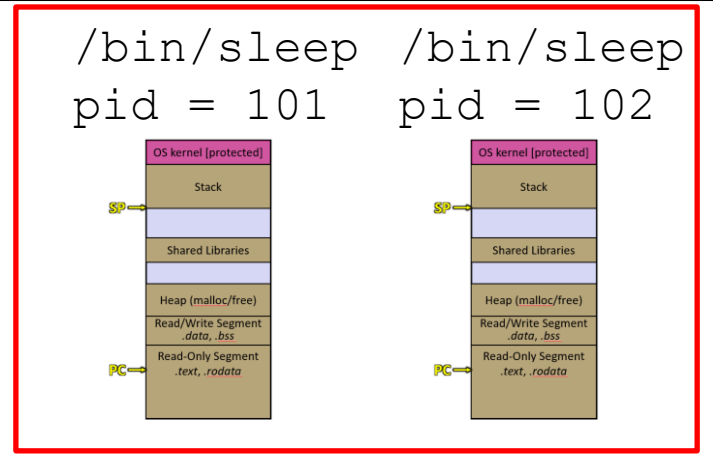
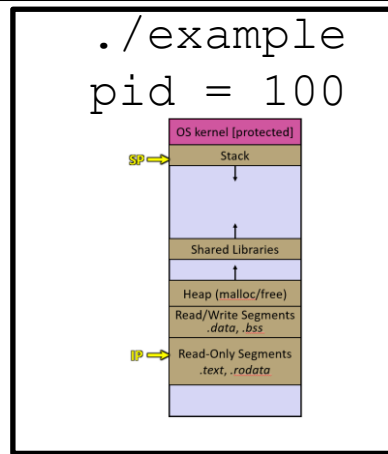
What if the parent forks a second child and adds it to the other child's group? (or the child calls fork directly, but don't do this in penn-shell)

Diagram: signals between process groups

User Processes

pgid = 100

pgid = 101



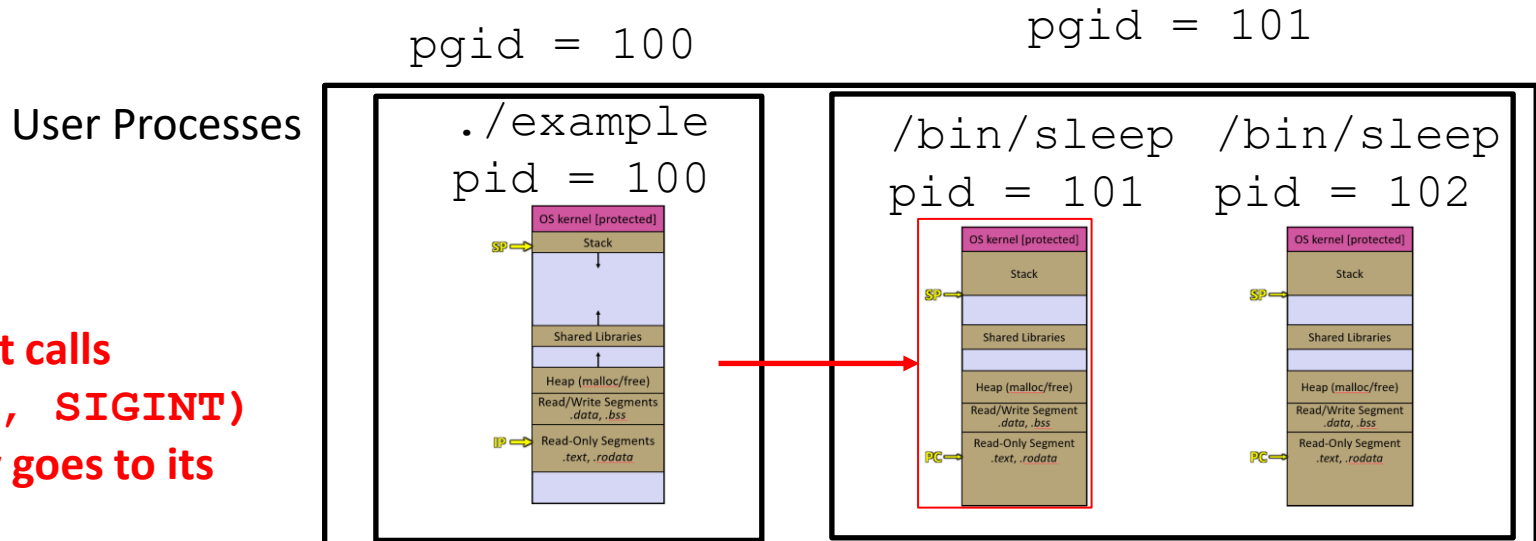
OS

Let's say we have the parent process on the left

a child process in its own group on the right.

What if the parent forks a second child and adds it to the other child's group? (or the child calls fork directly, but don't do this in penn-shell)

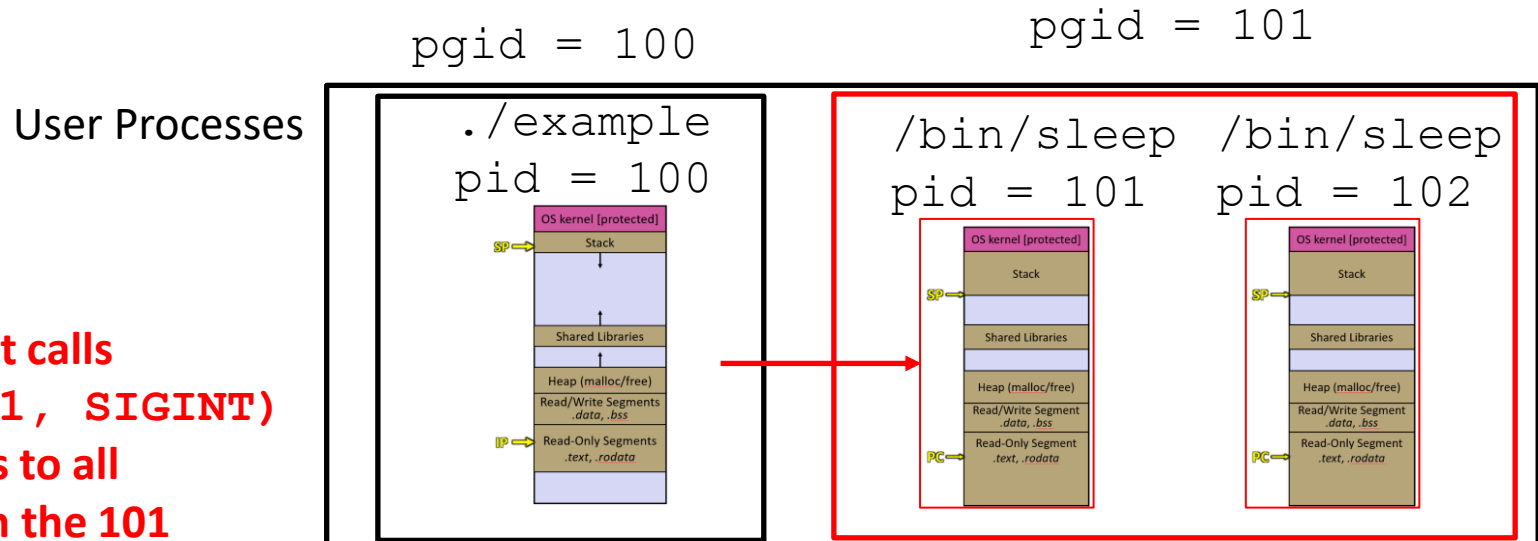
Diagram: signals between process groups



If the parent calls
`kill(101, SIGINT)`
Then it only goes to its
child

OS

Diagram: signals between process groups



If the parent calls
`kill(-101, SIGINT)`
Then it goes to all
processes in the 101
group

OS

Demo: pgrp_signals.c

- ❖ See code demo: `pgrp_signals.c`
 - Handler registered for SIGINT in both child and parent
 - Parent puts child in its own group
 - CTRL + C is input -> parent signal handler is invoked -> parent relays the signal to the child
 - What happens if we don't call kill in parent handler?
 - What happens if we then don't put child in its own group?

Lecture Outline

- ❖ Process Groups
 - `setpgid()`
- ❖ **Terminal Control**
 - `tcsetpgrp()`
- ❖ SIGSTOP
- ❖ Project 1: Synch vs Asynch wait
 - SIGCHLD

What if the child tried to use the terminal?

❖ Demo!

- Modify the `pgrp_signals.c` so that the child does “cat” (read from stdin, echo it to stdout until EOF)
- it does not work?

Sessions

- ❖ A **Session** is a collection of process groups
 - A session can be attached to a controlling terminal
 - Or not attached to any terminal (daemon's)

- ❖ You can think of a session as mostly associated with a “login” or instance of a terminal application. Each login/terminal is a session

- ❖ Within a session (that has a controlling terminal) there are
 - Background processes
 - Foreground processes

Foreground Process Groups

- ❖ Foreground process groups (i.e., Foreground Jobs) can read from STDIN and the processes in that group receive the signals from the keyboard (e.g., CTRL + C)
- ❖ A foreground group can make another group the foreground with the function:
- ❖

```
int tcsetpgrp(int fd, pid_t pgrp);
```

 - **fd** is a file descriptor associated with the terminal (stdin)
 - Sets the process group specified by **pgrp** to be the foreground process group
 - **-1** returned on error, **0** when successful

Background Process

- ❖ If a background process tries to read from `stdin`, it gets sent the signal `SIGTTIN`
- ❖ If a background process tries to take control of the terminal with `tcsetgpgrp`, then the group gets sent `SIGTTOU`, which will stop the process group
- ❖ Writing to `stdout` from the background is ok, but can be configured so that background processes get `SIGTTOU`

Demo: tc.c

- ❖ See code demo: `tc.c`
 - Fixed our process group code so that it can run `cat` 😊
 - Parent can print to `stdout` even if has given away the terminal

- How can we make the parent take back the terminal control?

Poll Everywhere

pollev.com/tqm

- ❖ What is the intention of this code? Does it do what it intends to do? How can we fix it?

```

13 int main() {
14     while (true) {
15         fprintf(stderr, "give command: ");
16         char c;
17         ssize_t bytes = read(STDIN_FILENO, &c, 1);
18         if (bytes == -1) {
19             perror("read\n");
20             exit(EXIT_FAILURE);
21         } else if (bytes == 0) {
22             break;
23         }
24
25         if (c == 'c') {
26             pid_t pid = fork();
27
28             if (pid == 0) {
29                 // child
30                 // reads from the terminal and
31                 // prints what it reads until EOF
32                 char* args[] = {"cat", NULL};
33                 execvp(args[0], args);
34                 exit(EXIT_FAILURE);
35             }
36             // parent
37

```

```

36         // parent
37
38         // put the child in its own process group
39         if (setpgid(pid, pid) == -1) {
40             perror("setpgid\n");
41             exit(EXIT_FAILURE);
42         }
43
44         // give terminal to the child
45         if (tcsetpgrp(STDIN_FILENO, pid) == -1) {
46             perror("tcsetpgrp\n");
47             exit(EXIT_FAILURE);
48         }
49         printf("starting to wait\n");
50
51         int wstatus;
52         waitpid(pid, &wstatus, 0);
53     } else if (c == 's') {
54         printf("sleeping...\n");
55         sleep(5);
56         printf("awake\n");
57     } else if (c == 'p') {
58         printf("HOWDY\n");
59     }
60 }

```

Demo: tc_loop.c

- ❖ See code demo: `tc_loop.c`
 - The code from the poll
 - Let's try to fix it...

- How can we make the parent take back the terminal control?

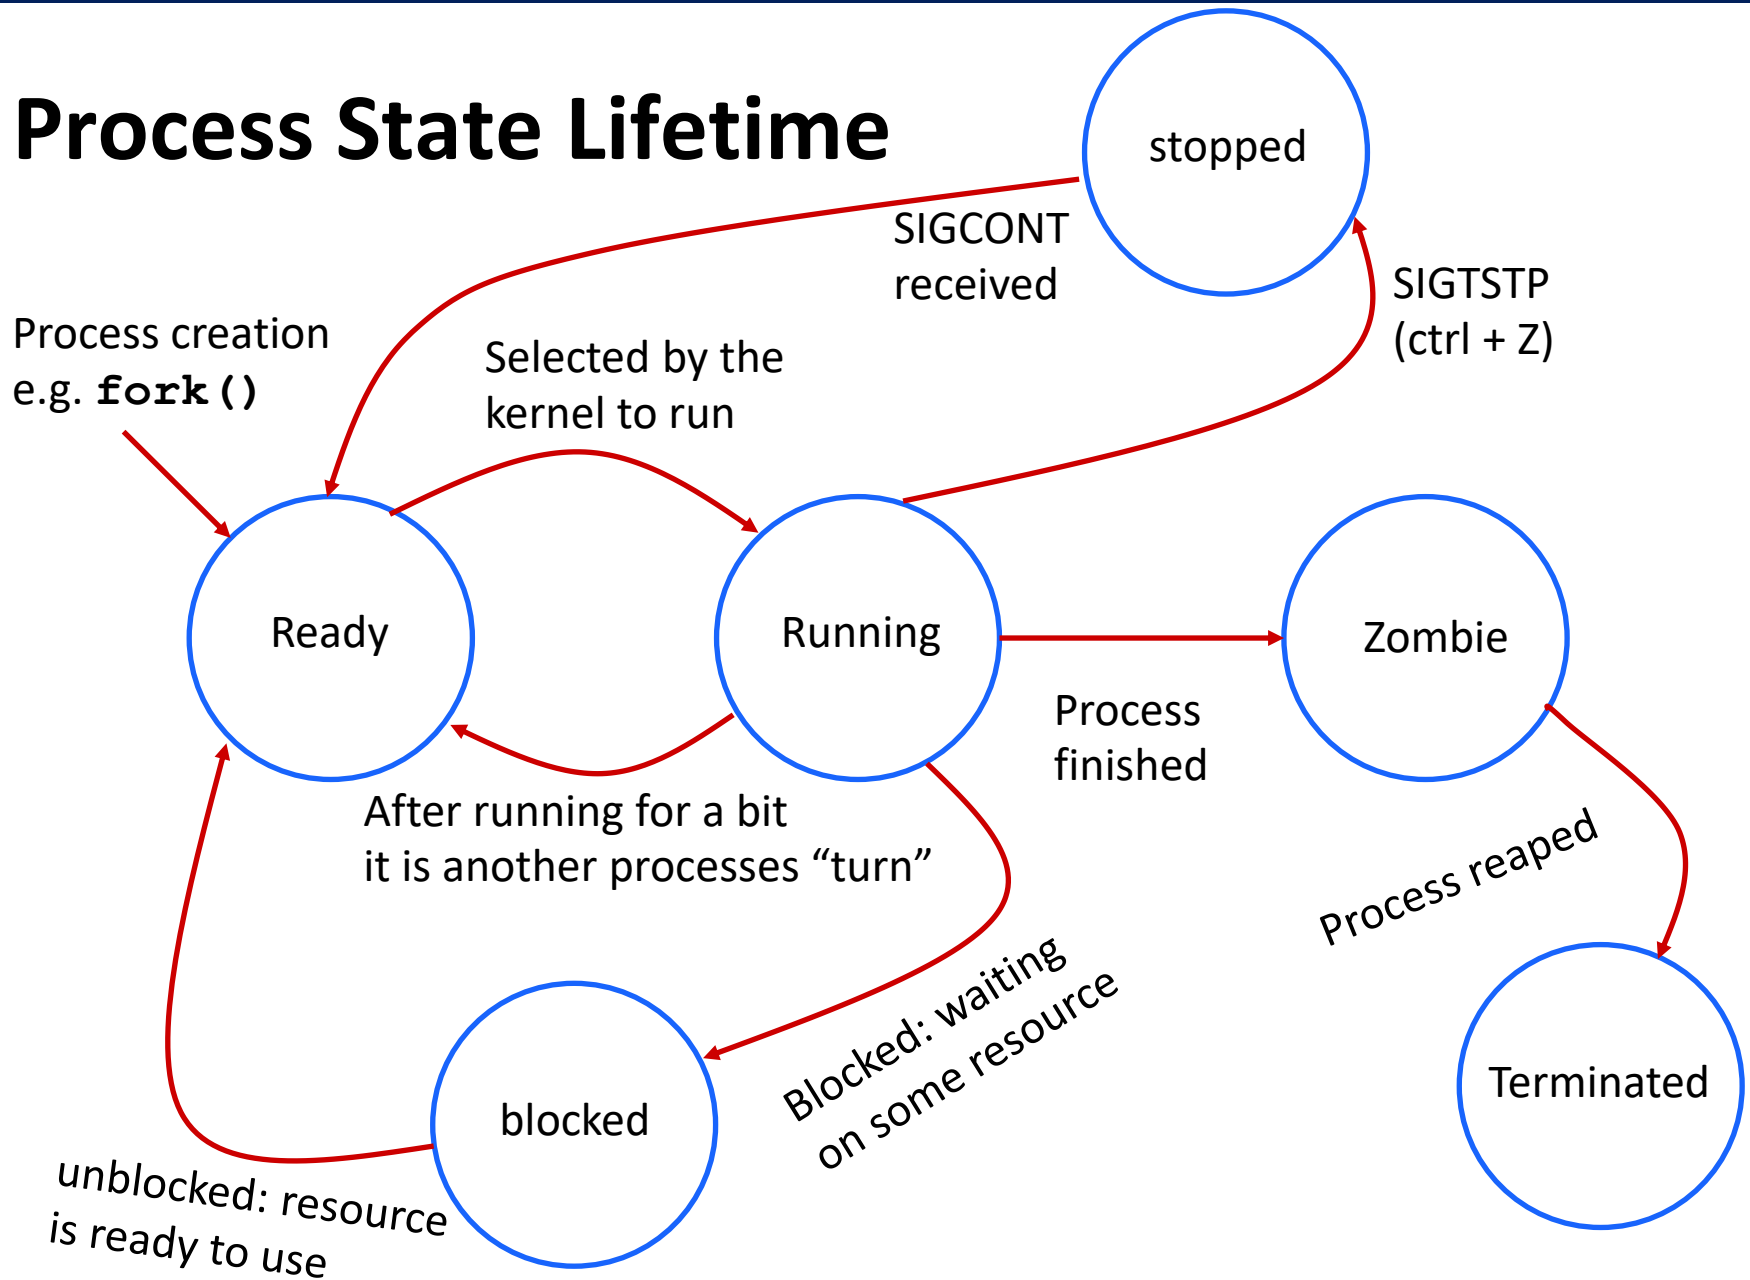
Lecture Outline

- ❖ Process Groups
 - `setpgid()`
- ❖ Terminal Control
 - `tcsetpgrp()`
- ❖ **SIGSTOP**
- ❖ Project 1: Synch vs Asynch wait
 - `SIGCHLD`

Stopped Jobs

- ❖ Processes can be in a state slightly different than being blocked. *// This is relevant for `penn-shell`*
 - When a process gets the signal `SIGSTOP`, the process will not run on the CPU until it is resumed by the `SIGCONT` signal
 - Other signals can still stop a program by default, like `SIGTSTP` or `SIGTTOU`
- ❖ Demo:
 - In terminal: `ping google.com`
 - Hit `CTRL + Z` to stop
 - Command: `"jobs"` to see that it is still there, just stopped
 - Can type either `"%<job_num>"` or `"fg"` to resume it

Process State Lifetime



Lecture Outline

- ❖ Process Groups
 - `setpgid()`
- ❖ Terminal Control
 - `tcsetpgrp()`
- ❖ SIGSTOP
- ❖ **Project 1: Synch vs Asynch wait**
 - **SIGCHLD**

Background in the shell

- ❖ To start a background job in the shell (and in penn-shell) run the command with a `&` at the end.
 - `sleep 10 &`
- ❖ While a command is running in the background, we can run other commands in the shell
- ❖ Can use the `jobs` command to see the status of the jobs we have started

Penn-shell

- ❖ Part of what you do in HW1 (after the milestone) is to make a shell that manages process groups in the foreground and background
- ❖ This means your code will have to handle multiple process groups at once, keeping track of the state of all of them.
- ❖ Need to maintain a linked list of the current jobs to handle job control

"Normal" approach Pseudo Code

❖ Discuss: what does this do?

❖ Is there a flaw in this?
Not in correctness but maybe

- Responsiveness
- Resource utilization
- etc.

```
int main(int argc, char* argv[]) {
    while(...) {
        printf(PROMPT);

        getline(&user_input);

        pid = fork_exec(user_input);

        waitpid(pid, &wstatus, 0);

        for (pid_t p : background) {
            // check status of background
            waitpid(p, &wstatus, WNOHANG);
            // if there is an update,
            // need to update the lists...
        }
        // re-prompt user
    }
}
```

Analysis: "Normal"

- ❖ The “normal”: check background processes before re-prompting the user
 - may not be responsive to background processes finishing
 - Consider we have many background processes then the user runs `sleep 1000000` in the foreground...
 - those background processes will not be reaped until foreground finishes

"Polling" approach Pseudo Code

- ❖ Discuss: what does this do?
- ❖ How does this compare to the previous attempt?

```
int main(int argc, char* argv[]) {
    while (...) {
        printf(PROMPT);
        getline(&user_input);
        pid = fork_exec(user_input);

        while (waitpid(pid, &wstatus, WNOHANG) == 0) {
            for (pid_t p : background) {
                // check status of background
                waitpid(p, &wstatus, WNOHANG);
                // if there is an update,
                // need to update the lists...
            }
        }
        // re-prompt user
    }
}
```


Analysis: Polling

- ❖ Polling is a term used to describe when we check to see if something is ready, but do not block if it is not ready
- ❖ This approach is more responsive than the previous one...
- ❖ but it busy waits... consuming CPU cycles...

Aside: SIGCHLD

- ❖ This approach registers **SIGCHLD** as a handler, **SIGCHLD** is a signal that is sent when a child process stops or is terminated
 - Is ignored by default

"async" approach Pseudo Code

- ❖ Discuss: what does this do?
- ❖ How does this compare to the previous attempt?

```
void handler(int signo) {
    for (pid_t p : background) {
        // check status of background
        waitpid(p, &wstatus, WNOHANG);
        // if there is an update,
        // need to update the lists...
    }
}

int main(int argc, char* argv[]) {
    signal(SIGCHLD, handler);
    while (...) {
        printf(PROMPT);
        getline(&user_input);
        pid = fork_exec(user_input);
        waitpid(pid, &wstatus, 0);
        // re-prompt user
    }
}
```

Analysis: Async

- ❖ This approach registers **SIGCHLD** as a handler, **SIGCHLD** is a signal that is sent when a child process stops or is terminated
 - Is ignored by default
- ❖ This allows us to respond quickly to the background children terminating
- ❖ No busy waiting! Main process instead is mostly blocked waiting on the foreground job
- ❖ Must use signal handlers and handle critical sections ;_;
- ❖ **Handling this ASYNC is your extra credit**
pass the normal autograder first PLEASE

Reminder: sigsuspend

- ❖ Another way to approach handling async is to use `sigsuspend()`
 - May be a little harder to reason about
 - Don't have to do much in the signal handler if this is the case!