

Project 2 – Basic Shell Part 2: Job Control

START THIS PROJECT EARLY
READ THIS ENTIRE DOCUMENT
(TWICE)

The Big Idea

In this assignment you will be building upon the basic shell you developed in Project 1. Specifically, we will be focusing on Process Groups, Signals, and Signal Handlers in order to implement a standard job control system similar to the ones found in most shells. In this context, a *job* is the execution of a single command line entry. For example:

```
$ ls -l
```

is a single job consisting of one process. Likewise,

```
$ ls -l | grep .c
```

is a single job consisting of two pipelined processes.

Upon successful completion of this project, your shell will have the following additional functionality:

- You will be able to check the status of jobs with the new built-in jobs command
- You will be able to send signals to specific processes and job numbers from your command line using the new built-in kill command (*Tip: Port over your code from Homework 4*)
- You will be able to suspend the foreground job by hitting Ctrl+z

```
$ sleep 100
^Z[0] + suspended sleep 100
$ jobs
[0] + stopped sleep 100
$
```

- You will be able to continue a stopped job in the background using the new bg command:

```
$ jobs
[0] + stopped sleep 100
[1] + stopped sleep 500
$ bg %1
[1] + continued sleep 500
$ jobs
[0] + stopped sleep 100
[1] + running sleep 500
$
```

- You will be able to start a job in the background by ending a command with the ampersand (&)

character. Doing this causes the shell to display the job number and the involved process IDs separated by spaces:

```
$ frame_grabber cam0 | encode -o awesome_meme.mp4 &
[0] 3626 3627
$ jobs
[0] + running frame_grabber cam0 | encode -o awesome_meme.mp4 &
$
```

- You will be able to move a background or stopped job to the foreground using the new built-in fg command:

```
$ jobs
[0] + running frame_grabber cam0 | encode -o awesome_meme.mp4 &
$ fg %0
encoding frame 42239282 [OK]
encoding frame 42239283 [OK]
encoding frame 42239284 [OK]
encoding frame 42239285 [OK]
encoding frame 42239286 [OK]^Z
[0] + suspended frame_grabber cam0 | encode -o awesome_mem.mp4 & $
```

- You will also be able to kill all processes associated with a job using the new built-in kill command:

```
$ jobs
[0] + running frame_grabber cam0 | encode -o awesome_meme.mp4 &
$ kill %0
[0] + done frame_grabber cam0 | encode -o awesome_meme.mp4 &
$
```

2

Using Process Groups

Process Groups are, as the name would imply, a group of processes. Jobs are based on process groups. Each process group has a *Process Group ID (PGID)*, which is generally chosen to be the same as the PID of the first process placed in the group. This process is sometimes referred to as the *process group leader*, but it has no special significance. Process groups are convenient because they provide a simple means to send a group of related processes the same signal — generally for control purposes (i.e. SIGTSTP, SIGCONT, etc). More importantly, however, process groups are used to determine which processes can read and write to stdin and stdout as well as which processes receive control signals from the terminal. For example, when you send ^C (SIGINT) or ^Z (SIGTSTP) using the keyboard, the terminal catches the keystroke and sends the corresponding signal to every process in the foreground process group. For this reason, interactive shells (bash, zsh, etc), put the processes comprising a job into their own process group. Let's look at an illustration:

```

                                Session 400
                                process group leaders
                                bash      find      sort
session leader
PID = 400          PID = 658 PID = 660
```

	Process	= 660 Controlling	
	group 400	SID = 400	foreground
		PPID = 400	
	background	PGID = 658	process group
	process groups	SID = 400	PPID = 400 PGID = 660 SID = 400
		wc	
		PID = 659	uniq
		PPID = 400	PID = 661 PPID = 400 PGID = 660
controlling process	Controlling	PGID = 658	SID = 400
PPID = 399	terminal	SID = 400	
PGID = 400		Process	Process
SID = 400	Foreground PGID	group 658	group 660

Figure 1: The controlling terminal has a session (set of process groups). The foreground process group receives interactive control signals from the terminal sent by the keyboard. Only the foreground process group may read from stdin and write to stdout. If a background process reads from stdin or writes to stdout, all the processes in its process group will receive the SIGTTIN or SIGTTOU signal, respectively.

Figure 1 shows a situation that can be easily reproduced with the following commands:

```
$ find . -name "*.c" | wc -l &
[1] 658 659
$ sort numbers.txt | uniq
```

It is important to note that the shell (bash in this case) places jobs in their own process groups immediately after fork(ing). This ensures that the shell is never a member of a child process group. This is important because, otherwise, interrupting the foreground sort number.txt | uniq job in our example via ^C would send SIGINT to bash (PID 400) along with PIDs 660 and 661, which would kill not only the foreground process group but our shell as well! By putting jobs into their own process groups, the shell protects itself from receiving such control signals intended only for the current foreground process group.

3

Putting processes into a process group is easily accomplished. For example, consider the following:

```
1 for (t=0; t<tasks_in_job; t++) {
2 job_pids[t] = fork ();
3 setpgid (job_pids[t], job_pids[0]);
4
5 if (job_pids[t] == 0) {
6 /* child logic */
7 } else {
8 /* parent logic */
9 }
10 }
```

Obviously, there is nuance to this. First, keep in mind that passing a zero to either argument of setpgid() has special meaning—the following are all equivalent:

```
setpgid (0, 0);
setpgid (getpid(), 0);
setpgid (getpid(), getpid());
```

So, what we are doing in our `fork()` example above is to have both the parent and the child attempt to put the child into a new process group. This is necessary because we don't know in what order the scheduler will choose to execute these two processes after the `fork()`. Note, that there is no harm in placing a process into a process group it is already a member of—simply nothing happens. In this way, regardless of which process (parent or child) gets scheduled first, the child will get kicked out of the parent's process group as quickly as possible.

Important SIGCHLD Details

Keep in mind that when a child process changes its execution state, the parent process will receive a SIGCHLD signal. This state change may be caused by signals sent to the child such as SIGTSTP, SIGCONT, SIGTERM, SIGKILL, etc. When the child receives one of these signals (and changes state), the kernel subsequently sends the parent a SIGCHLD signal so that it may act accordingly. The reason the kernel sent the SIGCHLD signal can be determined by the parent process by investigating the status returned by `waitpid`. This requires the use of the WIF* macros detailed in the `waitpid` man page. For example:

```
int status;
...
chld_pid = waitpid (-1, &status, WNOHANG | WUNTRACED | WCONTINUED)
...
if (WIFSTOPPED (status)) {
    /* child received SIGTSTP */
} else if (WIFCONTINUED (status)) {
    /* child received SIGCONT */
} else ...
```

This is an important opportunity for the parent shell to change which process group is in the foreground—there are other important opportunities as well, such as job creation. Setting the foreground process group is accomplished by calling `tcsetpgrp()` (see `man tcsetpgrp`). Keep in mind that when a foreground job completes, the shell's process group does not automatically get set to the foreground! It is the shell's responsibility to ensure that happens.

Other Important Signals: SIGTTIN and SIGTTOU

We have already discussed the fact that only the foreground process group can read from `stdin` and write to `stdout`. So... what happens when a process not in the foreground process group tries to do one of these two activities? Well, the offending background process will receive the SIGTTIN signal if it tries to read from `stdin` or the SIGTTOU signal if it tries to write to `stdout`. The default handler for these signals stops the offending process (it does not terminate it!). This will obviously be a big issue for your main shell process, which will potentially attempt to do both of these things in the background (e.g. the prompt!) while a foreground job is running. Your main shell process should never be in a stopped state – if that happens, it's game over: the manager of all the process groups (i.e. jobs) is out of commission.

Hint: It may be smart for the shell to check if it is the foreground process using `tcgetpgrp()` when in this scenario. Keep in mind, a process can always intelligently and judiciously `pause()` until the time is right to take back the foreground.

Process Groups are Not Jobs (but they help!)

In implementing your job system, you will find that process groups alone do not provide everything you need to define a job. Job numbers, for example, will have to be tracked by you as well as the name of the job (i.e. the command entered at the prompt). How will you know when a job is done? You will need to keep track of the child PIDs comprising a job so that you can check for job completion every time a child terminates. If all child PIDs comprising a job have terminated, then the job is complete. For simplicity, you only need to be able to support a maximum of 100 simultaneously managed jobs. I suggest keeping an array of the following struct to manage jobs:

```
typedef enum {
    STOPPED,
    TERM,
    BG,
    FG,
} JobStatus;

typedef struct {
    char* name;
    pid_t* pids;
    unsigned int npids;
    pid_t pgid;
    JobStatus status;
} Job;
```

I also suggest writing a small family of functions (i.e. an API) for doing common job activities (e.g. adding, removing, killing, etc). How you actually end up doing this, however, is up to you. If you want to support an arbitrary number of simultaneously managed jobs, feel free to implement a linked list of Job structures (but I think you already have enough to deal with). Just make sure what you design here works.

The New Job Management Built-In Commands

You will need to write four (4) built-in commands: fg, bg, kill, and jobs. Unlike the which command, I recommend running these job control commands inside the shell process so that they have the ability to easily mutate elements of the Job array. To be clear, these job management commands do not need to support redirection or piping. Except for the jobs command, all of these commands can accept a job number as an argument. As you saw in the examples on Page 1, when specified as a command argument job numbers are decorated with a leading percent (%) character to indicate that the number that follows is a job number.

5

Built-in Command: fg

If no arguments are supplied, the following should be printed to the screen and no job states are

```
modified: Usage: fg %<job number>
```

When a properly formatted job number (that exists) is supplied, the corresponding job is moved to the foreground (and continued if necessary). If the supplied job number is not properly formatted or corresponds to a job that does not exist, the following is printed to stdout:

```
pssh: invalid job number: [job number]
```

where [job number] is the supplied (malformed or invalid) argument.

Built-in Command: bg

If no arguments are supplied, the following should be printed to the screen and no job states are

modified: Usage: bg %<job number>

When a properly formatted job number (that exists) is supplied, the corresponding job is continued but not moved to the foreground. If the supplied job number is not properly formatted or corresponds to a job that does not exist, the following is printed to stdout:

pssh: invalid job number: [job number]

where [job number] is the supplied (malformed or invalid) argument.

Built-in Command: kill

If no arguments are supplied, the following should be printed to the screen and no job states are

modified: Usage: kill [-s <signal>] <pid> | %<job> ...

This means that a mixed list of PIDs and job numbers can be supplied to the kill command, and the desired signal will be sent to each one specified. (The bar simply means the user can supply a pid OR a job number; the ellipsis just means the user can supply as many of these as they like). By default SIGTERM is sent. If the optional -s argument is provided, then the specified signal number is sent instead.

If an invalid job number is supplied, the following is printed to stdout:

pssh: invalid job number: [job number]

where [job number] is the supplied job number.

If an invalid PID is supplied, the following is printed to stdout:

pssh: invalid pid: [pid number]

where [pid number] is the supplied (malformed or invalid) argument.

6

Built-in Command: jobs

This command takes no arguments. All active jobs are simply printed to stdout in the following format:

[job number] + state cmdline

where job number is the job number, state is either stopped or running, and cmdline is the command line used to start the job. Here is an example output showing all possible states:

```
[0] + stopped foo | bar | baz  
[2] + running pi_computation
```

[3] + stopped top &

notice that it is absolutely possible for jobs to terminate in an order different than they were started in (i.e. job 1 is already done)! The next job that is launched should therefore be job number 1 since it is the lowest available job number.

Providing Job Status Updates

In addition to being able to run the built-in jobs command to see if background jobs are stopped or running, the user should receive feedback from your shell when:

1. A foreground job is suspended via ^Z (SIGTSTP). For example:

```
$ ./my_program
Running...
^Z[1] + suspended ./my_program
$ jobs
[0] + running pi_computation
[1] + stopped ./my_program
$
```

2. A stopped background job is continued (either via bg or by sending it SIGCONT directly using kill). For example:

```
$ jobs
[0] + running pi_computation
[1] + stopped ./my_program
$ bg %1
[1] + continued ./my_program
$ jobs
[0] + running pi_computation
[1] + running ./my_program
```

3. A background job completes or is terminated. For example:

```
$ jobs
[0] + running pi_computation
[1] + stopped ./my_program
$ kill -s 9 %0
[0] + done pi_computation
```

This “done” status message should not be displayed when a foreground process completes/terminates— only background processes.

7

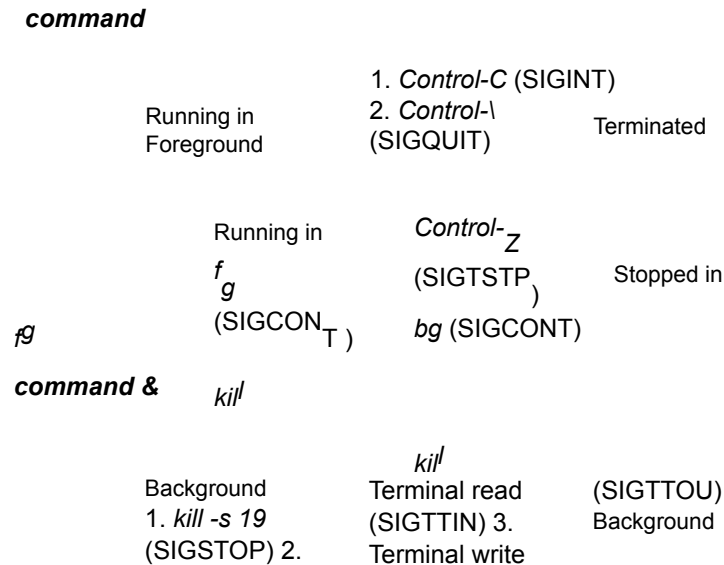
Hint: These messages to stdout are reports generated by the shell when one of its child process groups changes state. Consequently, these messages will be triggered within your shell’s SIGCHLD signal handler. Be careful though—some jobs will consist of multiple processes! If a job with 5 processes changes state, you don’t want the shell to report 5 status changes when the job is continued, for example—you only want it to notify the user once that the entire job has continued.

Softball: For the purposes of this project, feel free to put the necessary printf() statements directly in your signal handler (although, this is bad practice and should never be done in production code because printf() is non-reentrant due to its output buffering!). If you are feeling ambitious, you could use write()

directly instead of `printf()`, but this will require you to do string formatting manually prior to calling `write()`, which will further increase your code complexity.

Summary of Job States: The Big Picture

Now that we have discussed the details, let's step back and consider the possible states a job can be in and how it may transition between these states by examining the following diagram:



Keep in mind that signal events such as *Control-Z*, *Control-C*, *Control-*, SIGTTIN, SIGTTOU, etc will directly change a job's state and the shell must be reactive in handling these events through its SIGCHLD handler.

Meanwhile, other events are caused directly by the shell as a result of commands typed by the user (e.g. *fg*, *bg*, and *kill*). In these cases, the shell itself is responsible for sending the necessary signals to the corresponding job's process group. Once a job receives a signal sent by the shell, the child processes comprising the job may change state, which may result in required reactive behavior of the shell, again, through its SIGCHLD handler.

The key to this project is properly managing signals received by the shell and sending the correct signals to child processes groups (i.e. jobs) when necessary to ensure proper state.

For example, if a foreground job receives SIGTSTP from the controlling terminal, the child processes in the process group for that job will change state, causing the shell to receive SIGCHLD. In reacting to the SIGCHLD signal, the shell must update its job management data structure, retake the foreground, and then print job status feedback to stdout informing the user that the job was suspended.

Grading Rubric & Convenient Feature Checklist

When assessing the following table prior to submission (notice the fancy check boxes for your benefit) please keep in mind that although partial credit may be possible in extremely select circumstances, it is not very likely. Do not halfway implement a feature in a way that is completely non-functional, give up, and expect partial credit. Your features must actually work to receive credit.

Feature Points

- A process group is created for each new job. The PID of the first process in the job is used as the PGID. 10
- A new job intended to run in the foreground is made the foreground process group. 10
- A new job intended to run in the background is correctly setup and launched in the background. Job number and associated PIDs are printed to stdout by the shell as specified in the problem description. 10
- The shell process does not become stopped by SIGTTIN or SIGTTOU due to attempted background reads or writes to stdin/stdout. 10
- The shell process properly waits on all child processes comprising jobs (i.e. the shell does not generate a legion of zombie processes). 10
- A new background job is given the lowest available job number. 5
- Ctrl+Z (SIGTSTP) properly suspends the foreground job. The suspended message specified in the project description is printed to stdout by the shell informing the user accordingly. The shell itself is not suspended and is moved to the foreground. 5
- Ctrl+C (SIGINT) properly terminates the foreground job. The done message specified in the project description is NOT printed to stdout by the shell. The shell itself is not terminated and is moved to the foreground. 5
- Completed/terminated background jobs result in the shell printing the done message specified in the project description to stdout. The job is appropriately removed from the job management system. 5
- A job continued via either bg or SIGCONT properly resumes running in the background and the continued message specified in the project description is printed to stdout by the shell. 5
- The fg built-in command functions as specified in the project description. 5 □ The bg built-in command functions as specified in the project description. 5 □ The jobs built-in command functions as specified in the project description. 5 □ The kill built-in command can send signals to specified PIDs. 5
- The kill built-in command can send signals to an entire process group associated with a specified job number. 5

- The kill built-in command can send a user specified signal number as specified by the optional -s parameter.

5

9

A Few Final Tips

- Always assume the user has their tty set to properly respond to background stdout writes with SIGTTOU. Test your shell under this behavior by ensuring your terminal is properly set by running the following command prior to invoking your pssh shell:

```
$ stty tostop
```

Your project will be graded under an environment with this behavior.

- When calling tcsetpgrp() as a background process, you may find that the terminal is sending your shell SIGTTOU. This is because background processes are not allowed to modify the terminal. This includes changing the foreground process group! However, your shell must take back the foreground when the current foreground job is either suspended or terminated. So, how does the shell take back the foreground?

The answer is to temporarily ignore the SIGTTOU signal very briefly while performing the tcsetpgrp() call (as we discussed in Lecture 14).

Do NOT ignore SIGTTOU all the time, though—this will yield undesirable behavior.

- The kill() function provided by signal.h has the ability to send a signal to all processes within a specified process group. This is very useful! I suggest you read the man page:

```
$ man 2 kill
```

- The following man pages also make for good reading and could prove to be extremely useful references:

```
$ man 2 waitpid
```

```
$ man 2 signal
```

```
$ man 2 setpgid
```

```
$ man 3 tcgetpgrp
```

```
$ man 2 pause
```

- Use the provided job info.c program to help you check that you are properly putting job processes into their own process group. You can also use job info.c to check that you are setting the foreground process correctly. Details and examples of how to do this are in the header comments of job info.c.

