

Project Shell 2

Due: November 8, 2017 at 11:59pm

1 Introduction	2
2 Assignment	2
2.1 Stencil	2
2.2 Jobs vs. Processes	2
2.3 Foreground vs. Background	3
2.4 Specification	3
3 Signals	5
3.1 Ignoring Signals	5
3.2 Signal Handling in the Child Process	5
3.3 Process Groups	6
4 Multiple Jobs	7
4.1 Process State	7
5 Reaping	7
5.1 Reaping a Child Process	8
6 fg and bg	9
7 Support Code	10
8 Library Functions	11
9 Testing Your Shell	11
9.1 Format of the Tracefiles	12
9.2 Test Executables	12
9.3 Demo	12
9.4 Error Checking	12
9.5 Shell Clean Up	12
10 GDB Tips for Shell 2	13
10.1 handle signal	13
10.2 catch signal	13
11 Handing In	13

1 Introduction

Scream harvesting is a dangerous job, so Monsters, Inc. wants to maintain good business practices and keep their scarers safe. To do this, they need to update their management software in order to monitor the status of all the scaring jobs going on, and make sure that no monsters get lost out in the human world. Your task is to create a new and improved shell, a “conch shell,” which allows Monsters Inc. to keep track of all the jobs the scarers need to complete.

Now, your shell can run only one program at a time since it runs programs in the foreground, occupying the entire shell. This isn't especially useful: most of the time, you want to be able to run multiple processes concurrently. Additionally, if that process were to hang or become unresponsive, it would be impossible to end it without also killing the shell.

2 Assignment

In this assignment, you will be expanding your C Shell with a basic job control system: your shell will be able to handle multiple processes by running some in the background, as well as managing processes with signal forwarding.

2.1 Stencil

There is some stencil code for this assignment. You can install it in your home directory by running

```
cs0330_install shell_2
```

However, most of the support code for this project is your C Shell. Make sure to copy all of your C Shell source to your Shell 2 directory. You can use the Makefile from C Shell to run Shell 2, but you will want to compile with the `-D_GNU_SOURCE` and `-std=gnu99` flags to avoid potential headaches, and you will need to add the new files in the Shell 2 stencil to your Makefile.

2.2 Jobs vs. Processes

Suppose you execute your shell within your shell. Then the inner shell itself spawns its own child processes in response to the commands it receives.

In this situation, the shell is not running just a single foreground process. The command it has executed now comprises multiple processes: one for the inner shell, and additional processes spawned by the inner shell.

Still, from the point of view of the outer shell, the inner shell and its child processes should be considered a single unit. This unit is called a *job*. A *job* is a process or group of processes created by your shell from a single command.

2.3 Foreground vs. Background

The notion of *foreground* and *background* jobs is crucial to a job control system.

A *foreground* job is the group of processes with which the user interacts. When the current version of your shell executes a command, that process receives all commands; the shell itself must wait for that process to terminate before it can resume receiving user input.

A *background* job is the opposite. Rather than take over the shell's interface, background processes run behind it. The shell still receives commands while the background processes execute without interrupting the shell.

Because background jobs do not take over the shell's interface, it is possible for a shell to execute many jobs at once.

2.4 Specification

In addition to maintaining its current behavior, your shell should now exhibit the following additional behaviors:

- If **CTRL-C**, **CTRL-Z**, or **CTRL-** is typed into your shell, the resulting signal (**SIGINT**, **SIGTSTP**, or **SIGQUIT**, respectively) should be sent to the currently-running foreground job (if one exists) instead of to your shell. If no foreground job is running, then nothing should happen.
- If a command ends with the character **&**, that command should run in the background. Otherwise, that command should be run in the foreground. Note: the '**&**' character must be the last thing on the command line.
- When starting a job in the background, a message indicating its job and process ID should be printed to standard output. In Shell 2, this is formatted as:

```
[<job id>] (<process id>)
```

- Jobs started by your shell that have terminated should be *reaped* (see section 5).
- Whenever a job (foreground OR background) is terminated by a signal, your shell should print a brief message indicating the cause of termination, which can be determined by

the status set by `waitpid()`. The message should be formatted as follows (`jid` is the job id and `pid` is the process id):

```
[<jid>] (<pid>) terminated by signal <signal number>
```

- Whenever a job is stopped by a signal, you should print a similar message formatted as follows:

```
[<jid>] (<pid>) suspended by signal <signal number>
```

- Whenever a **background** job (not a foreground job) exits normally, print yet another similar message as follows:

```
[<jid>] (<pid>) terminated with exit status <status>
```

- Whenever a job is resumed via a **SIGCONT** signal, print another similar message with the following formatting:

```
[<jid>] (<pid>) resumed
```

- The following additional built-in commands should be supported, offering the same functionality and syntax as **bash** and most other shells:
 - **jobs** lists all the current jobs, listing each job's job ID, state (running or suspended), and command used to execute it. Foreground jobs should be added if they are ever suspended. Background jobs should always be added to the jobs list (as long as the background job does not immediately fail because it is not a valid command). We have implemented this printing for you in the support code. You must call `jobs()` to print all the jobs in your job list. See Section 7 for more information on the support code. For simplicity's sake, assign each job its ID and add it to the job list when it is created. The first job should have ID 1, the second ID 2, and so on.
 - **bg %<job>** resumes `<job>` (if it is suspended) and runs it in the background, where `<job>` is a job ID.
 - **fg %<job>** resumes `<job>` (if it is suspended) and runs it in the foreground. `<job>` is a job ID, just as in the **bg** command.

The % above simply tells your shell that you are referencing a job ID. To maximize compatibility with the provided testing program (described in Section 9), your shell should print no other output unless it is compiled with **PROMPT** defined (as in C Shell).

The following sections provide information that you will find invaluable in producing the above behaviors. We recommend completing this project in a linear manner, fully finishing each subsection of a given section before moving on to the next section.

3 Signals

In C Shell, when a job is launched, the shell can do nothing but wait until that job finishes. But what if the shell needs to perform some other task, or that job will never finish? There needs to be some way to interrupt execution of that job.

Signals are delivered asynchronously by the operating system to do precisely that. When you type characters on the keyboard, those characters are delivered to the current terminal. The terminal processes those characters and then performs some action upon its *foreground process group*. Most characters are simply passed along to be read by that process, but some, such as **CTRL-C**, generate a signal which is sent to every process in the terminal's foreground process group instead.

Therefore, in order to properly forward a signal to a job running in the foreground, your shell must set the process group ID of the terminal to be that of the foreground job. If there is no current foreground job, then the terminal's process group ID should instead be set to that of the shell.

3.1 Ignoring Signals

If you type **CTRL-C** when your shell has no foreground job, the controlling terminal will send **SIGINT** to your shell only. To avoid having the shell exit (or suspend execution), you must override its default behavior when it receives those signals so that it ignores the signal instead. You should **not** install a signal handler that does nothing. Instead, you should set the response to the signal to be **SIG_IGN**. Instead of using the `sigaction()` system call from the lab, you can more concisely invoke the `signal()` system call.

We have not yet covered the **SIGTTOU** signal, which must also be ignored here in addition to **SIGINT**, **SIGTSTP**, and **SIGQUIT**. **SIGTTOU** is sent to background processes if they attempt to write to **STDOUT** to prevent them from doing so. **SIGTTOU** should be ignored as well because when we run a foreground process from within our shell, then from the perspective of bash, our shell becomes a background process. This means that when the shell tries to take control back after the foreground process is terminated using `tcsetpgrp()`, it will be unable to do so since it receives **SIGTTOU**, which is not the behaviour we want.

3.2 Signal Handling in the Child Process

After forking into the child process and before calling `execv()`, you should set up signal handlers to set the behaviour of all the signals ignored in the parent process back to their default behaviour using **SIG_DFL**. This is because we want the child process to be able to accept and act on signals sent to it.

3.3 Process Groups

- `int setpgid(pid_t pid, pid_t pgid)`
- `int tcsetpgrp(int fd, pid_t pgrp)`

When you `fork()` off a new child process, that process begins execution as a copy of its parent, also inheriting its parent's *process group ID*. For signals to be properly sent to the new child but not also your shell, the process group ID of that child must be changed.¹

The `setpgid()` system call fortunately exists to make this change. All that remains, then, is to ensure that each job has a unique process group ID. This is simple to do — individual process IDs are guaranteed to be unique, so all you need to do is set the process group ID of a job to its process ID. The child process must do this immediately after it returns from `fork()` before `execv()` is called.

A separate routine is necessary to make sure that the terminal sends signals to the right processes, however. `tcsetpgrp()` conveniently does this, changing the process group that the terminal sends typed characters and signals to, in other words, the controlling process group.

- Keep in mind that there can be at most one foreground job that your shell may be running. If there is a foreground job then signals such as **CTRL+C** must be sent to and affect the running foreground process. If there is no foreground job running in your shell then these signals should be sent to the shell process itself.
- Since `tcsetpgrp()` is used to ensure that the correct process group receives signals, you should use this call whenever a foreground process begins and make sure to set terminal control back when the foreground process terminates. Additionally, be wary of terminal control when continuing stopped background or foreground jobs.
- The file descriptor passed into `tcsetpgrp()` should be that of the controlling terminal of the current foreground job, i.e. just the **STDIN** file descriptor. The second parameter is the process group ID of the process that terminal control should be set to. Therefore, it must be the process group ID of the foreground job your shell is running (if any), or the process group ID of your shell itself.
- We recommend calling `tcsetpgrp()` in your child process, after your call to `setpgid()` and before your signal handling and input/output redirection code.

¹ This will also matter later on when your shell begins to support multiple jobs; changing the process groups of those jobs is necessary so that signals will be sent to only the foreground job.

4 Multiple Jobs

Once your shell can manage signals and its foreground job, it is time to augment it so that it can run jobs in the background as well as the foreground. If your shell reads a ‘&’ character at the end of a line of input, it should execute that command in the background. You can assume that ‘&’ will only appear as the last non-whitespace character on a line.

Because a process running in the background immediately returns control of the program interface to the shell, you should be able to run a large number of background jobs (only limited by system resources). You should not wait for a background job to complete before launching another job.

Keep in mind that background jobs are run in separate process groups. You will need to track each job using a list. We have provided several functions that you can use to do this. For each job you are responsible for tracking its job ID, process ID, command, and state. See Section 7: Support.

4.1 Process State

At any time, a process can be in one of several states: it may be running, stopped, or terminated². Stopped processes have been suspended, but are still active and can be resumed by the user at any time. Terminated processes, referred to as *zombies*, have stopped execution and can no longer be executed, but persist on the system, continuing to consume resources until they are explicitly *reaped*.

A process’ initial state is running, and its state can be altered using signals. For example, **SIGINT** terminates a process, and **SIGTSTP** stops a process. **SIGCONT** can be sent to resume execution of a stopped process.

5 Reaping

Once your shell is capable of running background jobs, it is time to properly track them. Waiting for a background job to finish defeats the purpose of the background job, however. What’s needed is a means for knowing when any child process is terminated, then getting them out of the zombie state to clean up the process table. At present, your shell uses the `wait()` system call to wait for its foreground process to finish before displaying another prompt and running another command. For this assignment, you will want to use `waitpid()`, which suspends execution of the current process until a *specific* process changes state, based on its process ID.

² There are other states, but your shell need not track them.

Once the process has finished, `waitpid()` also *reaps* the process, freeing any resources that persist on the system.

5.1 Reaping a Child Process

- `int waitpid(pid_t pid, int *status, int options)`

The `waitpid()` system call allows your shell to do exactly that: examine a child process which has changed state, and instruct the operating system to clean up any resources associated with that child if it terminated.

The options parameter may be set to any combination of the following flags, or to 0:

- **WNOHANG** instructs the function to return immediately without hanging if no child process has terminated, returning 0 and leaves the contents of the integer pointed to by `status` unchanged. If any child process indicated by the `pid` argument *has* changed state, `waitpid` returns the process id of that child and updates `status` as usual.
- **WUNTRACED** instructs the function to return if a child process has been stopped, even if no child process has terminated; and
- **WCONTINUED** instructs the function to return if a child has been restarted from a stopped state.

`waitpid()` stores information about what happened to a child process in the location pointed to by `status` (if `status` is not `NULL`). You'll want to access this information so your shell can print an informative message about what happened. You can use the following macros, all of which have a single integer `status` as an argument, to access this information:

- **WIFEXITED(`status`)**, which returns true if the process terminated normally and 0 otherwise;
- **WEXITSTATUS(`status`)**, which returns the exit status of a normally-terminated process;
- **WIFSIGNALED(`status`)**, which returns true if the process was terminated by a signal and 0 otherwise;
- **WTERMSIG(`status`)**, which returns the signal which terminated the process if it terminated by a signal;
- **WIFSTOPPED(`status`)**, which returns true if the process was stopped by a signal;

- **WSTOPSIG(status)**, which returns the signal which stopped the process if it was stopped by a signal;
- **WIFCONTINUED(status)**, which returns true if the process was resumed by **SIGCONT** and 0 otherwise.

Note: You will not be able to catch and print *every* change in status for a given job. Specifically, if a job resumes and quickly finishes before reaping occurs, reaping will simply see that the job terminated and will print accordingly. This is expected behavior.

Shells commonly post changes in the status of jobs to the terminal as they are about to try and print another prompt. This is where you should check which jobs have changed state in your shell through reaping.

In theory, reaping can also be achieved by installing **waitpid()** in the **SIGCHLD** signal handler. The **SIGCHLD** signal is sent to the parent of a child process when it exits, is interrupted, or resumes after being interrupted. We strongly advise **not** using this method to reap process as this can create concurrency issues. You should be using just the jobs list and **waitpid()** to reap processes.

6 fg and bg

Now that you've added all of the other functionality, it is time to add the last piece of the puzzle: restarting a stopped job in the foreground or background using the commands **fg** and **bg** respectively, or moving a background job into the foreground with **fg**.

For both of these commands you should send a **SIGCONT** signal to the stopped job.

Some programs that might help you test management of foreground and background processes include:

- `/bin/sleep`
- `/usr/bin/yes`
- `/usr/bin/find`

Note that if you try to run any program that reads from standard input (such as `/bin/cat`) in the background, it will receive the signal **SIGTTIN** from the terminal and be suspended. You should be able to run (and restart) such processes in the foreground with no problems, however.

7 Support Code

You are provided with support code which will help you manage the job list. The `jobs.h` file contains declarations of the functions that you can use to this; those functions are defined in the `jobs.c` file. You are not responsible for understanding how the functions in `jobs.c` are implemented, although you are encouraged to read through it.

Here are the functions you are provided with:

- `init_job_list()`: initializes a job list (a `job_list_t`) and returns a pointer to it.
- `cleanup_job_list(job_list_t *job_list)`: cleans up the job list, de-allocating any system resources associated with it. Make sure you always call this before your program exits!
- `add_job(job_list_t *job_list, int jid, pid_t pid, process_state_t state, char *command)`: adds a job to the job list. Each job in the job list has a job id, process id, state, and command (which is the command used to start that job (e.g. `/bin/sleep`)).
- `remove_job_jid(job_list_t *job_list, int jid)` and `remove_job_pid(job_list_t *job_list, pid_t pid)`: remove a job from the job list, indicated by the job's job id or process id respectively.
- `update_job_jid(job_list_t *job_list, int jid, process_state_t state)` and `update_job_pid(job_list_t *job_list, pid_t pid, process_state_t state)`: update the process state of the job indicated by the given job id or process id respectively.
- `get_job_pid(job_list_t *job_list, int jid)`: gets the process id of a job indicated by a job id.
- `get_job_jid(job_list_t *job_list, pid_t pid)`: gets the job id of a job indicated by a process id.
- `get_next_pid(job_list_t *job_list)`: returns the process id of the next job in the job list, or `-1` if the end of the list has been reached. Future calls to this function will then resume from the start of the list.
- `jobs(job_list_t *job_list)`: prints out the contents of the job list. Useful for implementing the `jobs` builtin command.

You are also provided with two macros, `_STATE_RUNNING` and `_STATE_STOPPED` which correspond to process states. Use these macros whenever you need to initialize or update a process state.

To use these functions in your program, you must do two things. First, you must include `jobs.h` in any file in which you reference any of these functions. Second, you must make sure to compile your program with `jobs.c`. You can do so by updating the Makefile from C Shell to include `jobs.c` in the list of files to compile (update “`sh.c`” to use “`sh.c jobs.c`” wherever applicable).

8 Library Functions

The same restriction on non-syscall functions from C Shell applies, but you are also allowed to use the support code. Remember that `atoi()` and `sprintf()` are allowed functions; these functions should be sufficient to do any string manipulation or formatting that you might need.

9 Testing Your Shell

To help test your shell, we have provided you with a tester. To use the tester, run

```
cs0330_shell_2_test -s <33noprompt>
```

Optional arguments to the tester include:

<code>-m</code>	Run the tests in valgrind to check for memory leaks
<code>-p</code>	Run multiple tests concurrently
<code>-t <trace num></code>	Specify a certain trace to run (default: check all)
<code>-v</code>	Trace progress by printing trace descriptions, input, and output
<code>-w</code>	Show <i>more</i> information (such as the entire valgrind output if <code>-m</code> was run)
<code>-h</code>	Print help message

The tests can be found in `/course/cs0330/pub/shell_2`.

We highly recommend using the `-p` flag to speed up the testing. However, this can occasionally cause a trace to fail unexpectedly -- run that trace with `-t <tracenum>` to check.

Also, remember you can (and should) use `cs0330_shell_1_test` to make sure all functionality from Shell 1 remains intact, but you will have to use the `-p` flag for it to work on Shell 2.

9.1 Format of the Tracefiles

Trace files contain a sequence of commands for your shell to run, interspersed with special commands that the driver interprets to mean special actions are to be taken.

TSTP	Send a SIGTSTP signal to the shell
INT	Send a SIGINT signal to the shell
QUIT	Send a SIGQUIT signal to the shell
KILL	Send a SIGKILL signal to the shell
CLOSE	Close writer (sends EOF to shell)
WAIT	Wait for shell to terminate
SLEEP<n>	Sleep for n seconds

9.2 Test Executables

The testing scripts utilize small executables that are located (with their source code) in the `/course/cs0330/pub/shell_2/programs` directory. To test just one of the traces, run the tester with `'-t <tracenum>'` using the number, not the whole filename. It will be hard to understand what the tests are doing without understanding these programs, so we recommend examining the README contained in that directory as well as the source code for some of the programs.

9.3 Demo

When in doubt, follow the implementation provided in the demos. The demos can be invoked using `cs0330_shell_2_demo` and `cs0330_noprompt_shell_2_demo`. The demos are produced from identical source code, but the first is compiled with the macro `PROMPT` defined, and consequently produces additional output.

9.4 Error Checking

Make sure to error check **any** system calls that can throw errors. Check the man pages if you're not sure.

9.5 Shell Clean Up

Make sure you run `cs0330_cleanup_shell` after every time you work on Shell 2 on department machines. This is because it is possible that your shell could leave running processes after your shell has terminated (known as zombie processes). This slows down the department machines.

10 GDB Tips for Shell 2

As with all other projects, GDB will be very useful in this project. Specifically, since you are now working with signals, GDB has several commands that can help debug signals. Note that `set follow-fork-mode`, explained in the “C Shell” handout, will still be useful.

10.1 handle signal

GDB also provides several commands related to signals that will be helpful. Firstly, there is `handle <signal> [keywords...]`. This command allows you to manage what happens when you send a signal while in GDB. For example, when gdb receives a `SIGINT` signal, it prints a warning to the terminal, passes it to the program, and stops the program. If you wanted to print information that the signal was received without stopping the program, you would type `handle SIGINT nostop` in GDB. For more useful `handle` commands, visit the documentation [here](#).

10.2 catch signal

If, however, all you want is to simple break on a signal (for example, to step through your signal handler), you would execute the `catch signal [signal... — all]` command. For example, to break on a `SIGINT`, you would run `catch signal SIGINT`, or `catch signal all` if you wanted to catch all signals. Again, for more information, refer [here](#).

11 Handing In

To hand in the second part of your shell, run:

```
cs0330_handin shell_2
```

from your project working directory. You should hand in all of your code, the Makefile used to compile your program, and a README documenting the structure of your program, any bugs you have in your code, any extra features you added, and how to compile it.

Note: If you had **functionality** bugs on Shell 1 that are fixed in your Shell 2 code, document this in your README for up to 50% points back on Shell 1 deductions. Be as specific as possible (i.e., name tests that you were failing before and are now passing and refer to your grade report). Not fixing style deductions from Shell 1 can result in further style deductions for Shell 2.

Important note: If you have handed in but plan to hand in again after the TAs start grading (11/11 at noon), in addition to running the regular handin script, you must run `cs0330_grade_me_late shell_2` to inform us not to start grading you yet. You must run the script by **Saturday 11/11 at noon**.