

## Assignment Two

## 1 Code Review

As was the case for the first OS/161 assignment, you should begin with a careful review of the existing OS/161 code with which you will be working. The rest of this section of the assignment identifies some important files for you to consider. There are also a number of questions that you should be able to answer after you have read and understood the code. You should read these files and prepare short answers to the questions. You are expected to prepare a brief document, in PDF format, containing your answers. In your document, please number each of your answers and ensure that your answer numbers correspond to the question numbers. Your PDF document should be placed in a file called `codeanswers2.pdf`.

**In kern/userprog**

This directory contains the files that are responsible for loading and running user-level programs. Currently, the only files in the directory are `loadelf.c`, `runprogram.c`, and `uio.c`, although you may want to add more of your own during this assignment. Understanding these files is the key to getting started with the implementation of multiprogramming. Note that to answer some of the questions, you will have to look in files outside this directory.

**loadelf.c:** This file contains the functions responsible for loading an ELF executable from the filesystem into virtual memory space. (ELF is the name of the executable format produced by `cs350-gcc`.) As provided to you, OS/161 has a very limited virtual memory implementation. Virtual addresses are translated to physical addresses using a very simple translation scheme and the total size of the virtual address spaces of all running processes cannot be any larger than the size of physical memory.

**runprogram.c:** This file contains only one function, `runprogram()`, which is responsible for running an application program from the kernel menu. It is a good base for writing the `execv()` system call, but only a base. You will need to determine what is required for `execv()` that `runprogram()` does not concern itself with. You will also be expected to modify the `runprogram` command as part of this assignment, as described in Section 2.3.

**uio.c:** This file contains functions for moving data between kernel and user space. Knowing when and how to cross this boundary is critical to properly implementing user-level programs, so this is a good file to read very carefully. You should also examine the code in `kern/lib/copyinout.c` and figure out when these functions should be used in implementing your system calls.

**Question 1.** What are the ELF magic numbers?

**Question 2.** What is the difference between `UIO_USERSPACE` and `UIO_USERSPACE`? When should one use `UIO_SYSSPACE` instead?

**Question 3.** Why can the struct `uio` that is used to read in a segment be allocated on the stack in `load_segment()` (i.e., where does the memory read actually go)?

**Question 4.** In `runprogram()`, why is it important to call `vfs_close()` before going to user mode?

**Question 5.** Which function forces the processor to switch into user mode? Is this function machine dependent?

**Question 6.** In which file are `copyin` and `copyout` defined? `memmove`? Why can't `copyin` and `copyout` be implemented as simply as `memmove`? What is the purpose of `copyin` and `copyout`?

**Question 7.** What (briefly) is the purpose of `userptr.t`?

### In kern/arch/mips/mips

Exceptions are the key to operating systems; they are the mechanism that enables the OS to regain control of execution and therefore do its job. When the OS boots, it installs an “exception handler” (carefully crafted assembly code) at a specific address in memory. When the processor raises an exception, it invokes this handler, which sets up a “trap frame” and calls into the operating system. Since “exception” is such an overloaded term in computer science, operating system lingo for an exception is a “trap” – when the OS traps execution. System calls are implemented as one type of exception, and `syscall.c` handles traps that happen to be system calls.

**trap.c:** `mips_trap()` is the key function for returning control to the operating system. This is the C function that gets called by the assembly exception handler. `md_usermode()` is the key function for returning control to user programs. `kill_curthread()` is the function for handling broken user programs; when the processor is in usermode and hits something it can't handle (say, a bad instruction), it raises an exception. There's no way to recover from this, so the OS needs to kill off the process. Part of this assignment will be to write a useful version of this function.

**syscall.c:** `mips_syscall()` is the function that delegates the actual work of a system call to the kernel function that implements it. You will need to write those functions, and ensure that they are invoked from `mips_syscall()`. Notice that `reboot()` is the only case currently handled. You will also find a function, `md_forkentry()`, which is intended to be the function that run by the new thread (implementing the new process) that your kernel will need to create in response to a `fork()` system call.

**Question 8.** What is the numerical value of the exception code for a MIPS system call?

**Question 9** Why does `mips_trap()` set `cur脾` to `SPL_HIGH` “manually”, instead of using `splhigh()`?

**Question 10** How many bytes is an instruction in MIPS? Answer this by reading `mips_syscall()` carefully, not by looking somewhere else, and explain how you arrived at the answer.

**Question 11** Why do you “probably want to change” the implementation of `kill_curthread()`?

**Question 12** What would be required to implement a system call that took more than 4 arguments?

### In kern/include

These are the include files that the kernel needs. The `kern` subdirectory contains include files that are visible not only to the operating system itself, but also to user-level programs. (Think about why it's named “kern” and where the files end up when installed.)

**Question 13.** What is the standard interface to a file system (i.e., what functions must you implement to implement a new file system)?

**Question 14.** How large are OS/161 pids? How many processes do you think OS/161 could run concurrently as you have it now (ASST2)? A sentence or two of justification is fine.

**Question 15.** A `vnode` is a kernel abstraction that represents a file. What operations can you do on a `vnode`? If two different processes open the same file, do we need to create two `vnodes`?

**Question 16.** What is the system call number for a reboot? Is this value available to userspace programs? Why or why not.

### In kern/lib

These are library routines that are used throughout the kernel, e.g., for managing queues and for kernel memory allocation.

**Question 17.** Operating systems usually have two separate implementations of `malloc`, one for use by the kernel, and one for use by application programs. In OS/161, the kernel's `malloc` (`kmalloc`) is in `kern/lib/kheap.c`. The application `malloc` is missing, but it belongs in `lib/libc/malloc.c`. Explain why there needs to be two separate implementations and describe the implications (for applications) of the missing implementation.

**In `kern/fs`**

The file system implementation has two subdirectories. `kern/fs/vfs` contains the file-system independent layer (`vfs` stands for “Virtual File System”). It establishes a framework into which you can add new file systems easily. You will want to look at `vfs.h` and `vnode.h` before looking at this directory. `kern/fs/sfs` contains the implementation of the simple file system that OS/161 uses by default.

**Question 18.** What does a device pathname in OS/161 look like?

**Question 19.** What does a raw device name in OS/161 look like?

**Question 20.** What lock protects the `vnode` reference count?

**In `lib/crt0`**

This is the user program startup code. There's only one file in here, `mips-crt0.S`, which contains the MIPS assembly code that receives control first when a user-level program is started. It calls the user program's `main()`. This is the code that your `execv()` implementation will be interfacing to, so be sure to check which values it expects to appear in each register, and so forth.

**In `lib/libc`**

This is the user-level C library. There's obviously a lot of code here. We don't expect you to read it all. For present purposes you can focus on the code that implements the user-level side of system calls.

`errno.c`: This is where the global variable `errno` is defined.

`syscalls-mips.S`: This file contains the machine-dependent code necessary for implementing the user-level side of MIPS system calls.

`syscalls.S`: This file is created from `syscalls-mips.S` at compile time and is the actual file assembled into the C library. The actual names of the system calls are placed in this file using a script called `callno-parse.sh` that reads them from the kernel's header files. This avoids having to make a second list of the system calls. In a real system, typically each system call stub is placed in its own source file, to allow selectively linking them in. OS/161 puts them all together to simplify the makefiles.

## 2 Implementation Requirements

For this assignment, you are expected to implement exception handling and several OS/161 system calls, and to enhance the kernel's `runprogram` command.

All code changes for this assignment should be enclosed in `#if OPT_A2` statements, in much the same way that you used `#if OPT_A1` for Assignment 1:

```
#if OPT_A2
    // code you created or modified for ASST2 goes here
#else
    // old (pre-A2) version of the code goes here,
    // and is ignored by the compiler when you compile ASST2
    // the ‘else’ part is optional and can be left
    // out if you are just inserting new code for ASST2
#endif /* OPT_A2 */
```

For this to work, you must add `#include "opt-A2.h"` at the top of any file for which you make changes for this assignment.

By default, any code changes that you made for Assignment 1 (which should already be wrapped with `#if OPT_A1`) will *also* be included in your build when you compile for Assignment 2.

## 2.1 System Calls

You are expected to implement the following OS/161 system calls:

- `open`, `read`, `write`, `lseek`, `close`
- `getpid`
- `fork`, `execv`, `waitpid`, `_exit`

There is a man (manual) page for each OS/161 system call. These manual machines describe the expected behaviour of the system calls and specify the values expected to be returned by the system calls, including the error numbers that they may return. **You should consider these manual pages to be part of the specification of this assignment, since they describe the way that that system calls that you are implementing are expected to behave.** The system call man pages are located in the OS/161 source tree under `os161-1.11/man/syscall`. They are also available on-line through the course web page.

It is expected that your system calls will correctly and gracefully handle all system call error conditions, and that they will properly return the error codes as described on the man pages. This is because application programs, including those used to test your kernel for this assignment, depend on the behaviour of the system calls as specified in the man pages. Under no circumstances should an incorrect system call parameter cause your kernel to crash.

Integer codes for system calls are listed in `kern/include/kern/callno.h`. The file `include/unistd.h` contains the user-level function prototypes for OS/161 system calls. These describe how a system call is made from within a user-level application. The file `kern/include/syscall.h` contains the kernel's prototypes for its internal system call handling functions. You will want to create a kernel handler function for each system call that your kernel handles, and place a prototype for that function in `kern/include/syscall.h`.

For example, as provided to you, the OS/161 kernel supports only one system call: `reboot`. The user-level function prototype for `reboot`, which is found in `include/unistd.h`, looks like this:

```
int reboot(int code);
```

The kernel's internal handler function for the `reboot` system call is called `sys_reboot`. Its prototype, found in `kern/include/syscall.h`, looks like this:

```
int sys_reboot(int code);
```

## 2.2 Exception Handling

There are a number of types of exceptions that may be generated by the MIPS processor and that must be handled by the kernel. Some of these exceptions, such as those related to virtual memory, are already handled by the OS/161 kernel as it is provided to you. However, other types of "fatal" exceptions that may occur during the execution of a user-level program are not handled well. Instead of terminating the process that caused the exception, the kernel (as provided to you) simply panics and shuts down.

For this assignment, you are expected to modify the kernel's exception handling so that fatal exceptions (e.g., illegal instructions) that occur in user-level code are handled by terminating the process that caused the exception. Under no circumstances should such an exception cause your kernel to crash or shutdown.

## 2.3 Modifying `runprogram`

The kernel's `runprogram` command, which you will find implemented in `kern/userprog/runprogram.c`, can be used to launch an application program from the kernel menu. One important thing that `runprogram` cannot do is to pass arguments to the application program that is being launched. You are required to modify

`runprogram` so that it is capable of passing arguments from the kernel command line to the application that `runprogram` launches. For example, you should be able to run the `add` program from the `testbin` directory and get the proper result:

```
% sys161 kernel "p testbin/add 2 4;q"
Answer: 6
```

`runprogram` should set up arguments for the new program the same way that `execv` is supposed to set them up. Specifically, the new program's `main()` function should receive two parameters: `argv` and `argc`. `argv` is an array of strings (i.e., an array of pointers to character arrays), and `argc` indicates the number of arguments. The on-line man page for the `execv` system call describes `argv` and `argc` in more detail. Read it carefully.

## 3 Discussion

This section of the assignment provides some additional discussion of aspects of OS/161 that are relevant to the assignment.

### 3.1 open, read, write, lseek, close

Applications do their input and output using so-called *file descriptors* that are provided by the kernel. A file descriptor may refer to an open file or to a device such as the console. The `open` system call creates a new file descriptor referring to a specified file. Subsequent `read` or `write` system calls using that file descriptor cause data to be read from or written to that file.

Whenever a new process is created, the kernel should arrange that three file descriptors are created automatically and are open at the time that process's application starts running. These are standard input, standard output, and standard error (descriptors 0, 1, and 2, respectively). The kernel should arrange that these file descriptors start out attached to the console device ("con:").

These system calls are largely about manipulation of file descriptors, or process-specific filesystem state. This part of this assignment is about designing and implementing a system to track this state. You need to think about the state you need to maintain, how to organize it, and when and how it has to change.

### 3.2 getpid

A pid, or process ID, is a unique number that identifies a process. pid allocation and reclamation are the important concepts that you must implement. It is not OK for your kernel to crash because over the lifetime of its execution you've used up all of the pids. Design your pid system, implement all the tasks associated with pid maintenance, and only then implement `getpid()`.

### 3.3 fork, execv, waitpid, \_exit

These system calls are probably the most difficult part of the assignment. `fork` and `execv` enable multiprocessing and make OS/161 much more useful. `_exit` and `waitpid` are closely related to each other, since `_exit` allows the terminating process to specify an exit status code, and `waitpid` allows another process to obtain that code. Read the `waitpid` man page carefully.

### 3.4 Testing and runprogram

As provided to you, the kernel implements the `runprogram` operation, which can be invoked from the kernel menu. `runprogram` causes the kernel to launch a specified user-level program. This is sufficient to allow you to run one user-level application program at a time, which is handy for testing that your system calls work.

OS/161 includes a number of application programs that you can use. The `bin` and `sbin` directories contain a number of standard utility programs, such as a command shell. In addition, the `testbin` directory contains a variety of relatively simple programs that are intended primarily to test the OS/161 kernel. These are the programs that you are most likely to find useful for the assignment. Any of these programs can

be launched directly from the kernel using the `runprogram` command. An example of this was given in Section 2.3.

The source for all of the OS/161 application programs is under `cs350-os161/os161-1.11` in the `bin`, `sbin`, and `testbin` subdirectories. Once they are compiled, the application binaries are installed into `cs350-os161/root`, into the `bin`, `sbin` and `testbin` subdirectories there.

### 3.5 Configuring and Building

Before you do any coding for Assignment 2, you will need to reconfigure your kernel for this assignment. Follow the same procedure that you used to configure for Assignment 1, but use the Assignment 2 configuration file instead:

```
% cd cs350-os161/os161-1.11/kern/conf
% ./config ASST2
% cd ../compile/ASST2
% make depend
% make
% make install
```

This will configure, build and install your Assignment 2 kernel. Note that you build your kernel in `kern/compile/ASST2`, not `kern/compile/ASST1`.

To build the OS/161 user-level applications, you need to run `make` in the top-level directory of the OS/161 source tree:

```
% cd cs350-os161/os161-1.11
% make
```

You should not make any changes to the existing user-level applications. Generally, you should not have to rebuild those applications every time you build a new kernel. However, there are certain header files, e.g. in `kern/include/kern` that are used by the kernel *and* by the user-level application programs. If you make changes to these files, you must rebuild the user-level code. You do so like this:

```
% cd cs350-os161/os161-1.11
% make clean
% make
```

It is always OK to rebuild the user-level applications. If you are getting any weird, unexpected behaviour from those applications, it is a good idea to rebuild them just to be on the safe side. More importantly, you should **build both your kernel and user-level applications before you submit them to us**. We will be testing your kernel using our copies of the user-level applications - if we cannot build those applications, then we cannot test your code!

## 4 Strategy

Because of the the nature of the system calls you will be implementing, and because of the way that your work will be tested, there are some dependencies among the different parts of the assignment. Because of this, you should work on the different parts of this assignment in a specific order, which we describe in the rest of this section.

### 4.1 console write and `_exit`

We strongly recommend that you start with `_exit` and with `write` to standard output. You should get both of these calls working *before* you continue with the rest of the assignment. One important reason for this is that this is the bare minimum that you will need to be able to do anything useful in Assignment 3. A second reason is that the testing of other parts of *this* assignment depends on `write` (to the console) and `_exit` working properly. If we are unable to test these other parts of the assignment, then you will not get credit for the work that you put into them! So: get these two system calls working first.

- `_exit`: To get started, don't worry about storing the exit code and making it available to other processes. Focus instead on ensuring that `_exit` cleanly terminates the process the calls it. Later you can ensure that `_exit` meets its full specification.
- `write`: Be sure that the `write` system call is able to at least correctly write data to standard output, which should be connected to the console device. This will permit user-level programs to call `printf` to print output to `stdout`. This, in turn, will make it possible to demonstrate that user-level programs in this and future assignments are executing as expected (or not). For starters, don't worry about being able to write to other file descriptors, only worry about writing to `stdout`. Later, you can ensure that `write` meets its full specification.

## 4.2 Argument Passing

The second thing that you should do is to modify `runprogram` so that it is capable of passing arguments (see Section 2.3). Many of the test programs that we will use to test other parts of this assignment expect to be passed command line arguments. If argument passing does not work then we cannot test those parts of your work, and you will not receive credit for it. Make sure that argument passing is working properly *before* you start working on the file related system calls (e.g., `open`, `close`), and *before* you start working on `fork`, `execv`, and `waitpid`.

## 4.3 Other System Calls, Except `execv`

Once argument passing, console write, and `_exit` are finished, you should begin working on the remaining system calls, except for `execv`. These calls come in two groups. One group consists of the file-related system calls, i.e., `open`, `close`, `read`, `write`, and `lseek`. For the most part, these will be tested as a group. For example, we simply cannot test reads and writes from a file, or seeking in a file, if the file cannot be opened and closed. The second group includes `fork`, `waitpid` and `_exit`. Again, there are strong interdependencies among the calls in this group, e.g., we can't test `waitpid` if `fork` and `_exit` do not also work.

You can tackle these two groups of calls in either order. However, we strongly recommend that you get one group working before moving on to the other. Having two half-finished groups of system calls will be much worse, mark-wise, than having one group finished and working and other unstarted.

## 4.4 Exception Handling and `execv`

Nothing depends on these, so do them last. Note that some of the difficulty of `execv` comes from argument passing. If you have already implemented argument passing for `runprogram`, you'll have a much easier time with `execv`.

# 5 Design Document

You are expected to prepare a short design document describing your implementation of these system calls. Your document should describe the major issues that you encountered in designing and implementing the mechanisms that support the required system calls in your kernel, and how you resolved those issues in your design. In particular, your design document must address at least the following issues:

**file descriptors:** How are these implemented? What kernel data structures were created to manage file descriptors, and how are they used?

**process identifiers (pids):** How are processes uniquely identified? How do you generate unique pids for new processes? How and under what circumstances do you reclaim and reuse pids?

**process exit codes:** How and where are these recorded? For how long are they maintained?

**waiting for processes:** How did you implement `waitpid`? Which synchronization mechanism do you use to cause a process to block (if necessary) in `waitpid`? What restrictions, if any, have you imposed on which processes a process is permitted to wait for.

**process arguments:** How did you implement `argc` and `argv` for `execv` and `runprogram`? Where are the arguments placed when they are passed to the new process?

**synchronization:** Have you created kernel data structures that are shared by multiple threads? If so, how is access to these structures synchronized? Which synchronization primitives did you use?

Be sure to justify your design decisions.

It is OK if portions of your design are not implemented, and you may receive some credit for a high-quality design even if it is not implemented. However, if your implementation does not match your design, **your design document must clearly indicate which parts of the design are not implemented, or are not implemented as described in the design document. Submission of a design document that does not match the submitted implementation (and which does not flag any such differences) will be considered an act of academic dishonesty.**

Please prepare your design document as a PDF file called `design2.pdf`. Your design document should be *at most* three pages long using 1 inch margins and a minimum font size of 11 points.

## 6 What to Submit

You should submit your kernel source code, your code question answers, and your design document using the `cs350_submit` command, as you did for Assignment 1. It is important that you use the `cs350_submit` command - do not use the regular `submit` command directly.

Assuming that you followed the standard OS/161 setup instructions, your OS/161 source code will be located in `$HOME/cs350-os161/os161-1.11`. To submit your work, you should

1. place `codeanswers2.pdf` in the directory `$HOME/cs350-os161/`
2. place `design2.pdf` in the directory `$HOME/cs350-os161/`
3. run `/u/cs350/bin/cs350_submit 2` in the directory `$HOME/cs350-os161/`

This will package up your OS/161 kernel code and submit it, along with the two PDF files, to the course account.

**Important:** The `cs350_submit` script packages and submits everything under the `os161-1.11/kern` directory, except for the subtree `os161-1.11/kern/compile`. You are permitted to make changes to the OS/161 source code outside the `kern` subdirectory. For example, you might create a new test program under `testbin`. However, such changes will not be submitted when you run `cs350_submit`. Only your kernel code, under `os161-1.11/kern`, will be submitted.