

Project #1 – Exceptions and Simple System Calls¹

Introduction to Operating Systems
Assigned: September 5, 2003

CSE421
Due: September 26, 2003 11:59:59 PM

The first project is designed to further your understanding of the relationship between the operating system and user programs. In this assignment, you will implement simple system call traps. In Nachos, an exception handler handles all system calls. You are to handle user program run time exceptions as well as system calls for IO processing. We give you some of the code you need; your job is to complete the system and enhance it.

Phase 1: Understand the Code

The first step is to read and understand the part of the system we have written for you. Our code can run a single user-level ‘C’ program at a time. As a test case, we’ve provided you with a trivial user program, ‘halt’; all halt does is to turn around and ask the operating system to shut the machine down. Run the program ‘**nachos -rs 1023 -x ../test/halt**’. As before, trace what happens as the user program gets loaded, runs, and invokes a system call.

The files for this assignment are:

progtest.cc	test routines for running user programs.
syscall.h	the system call interface: kernel procedures that user programs can invoke.
exception.cc	the handler for system calls and other user-level exceptions, such as page faults. In the code we supply, only the ‘halt’ system call is supported.
bitmap.*	routines for manipulating bitmaps (this might be useful for keeping track of physical page frames)
fileys.h	
openfile.h	(found in the fileys directory) a stub defining the Nachos file system routines. For this assignment, we have implemented the Nachos file system by directly making the corresponding calls to the UNIX file system; this is so that you need to debug only one thing at a time.
translate.*	translation table routines. In the code we supply, we assume that every virtual address is the same as its physical address -- this restricts us to running one user program at a time. You will generalize this to allow multiple user programs to be run concurrently in a later lab.
machine.*	emulates the part of the machine that executes user programs: main memory, processor registers, etc.
mipssim.cc	emulates the integer instruction set of a MIPS R2/3000 processor.
console.*	emulates a terminal device using UNIX files. A terminal is (i) byte oriented, (ii) incoming bytes can be read and written at the same time, and (iii) bytes arrive asynchronously (as a result of user keystrokes), without being explicitly requested.
synchconsole.*	routine to synchronize lines of I/O in Nachos. Use the synchconsole class to ensure that your lines of text from your programs are not intermixed.
../test/*	C programs that will be cross-compiled to MIPS and run in Nachos

¹ This project was originally suggested by the designers of Nachos. Original version of this document was created by Dr. Chris Egert. Modified by Bina Ramamurthy.

Phase 2: Design Considerations

In order to fully realize how an operating system works, it is important to understand the distinction between kernel (system space) and user space. If we remember from class, each process in a system has its own local information, including program counters, registers, stack pointers, and file system handles. Although the user program has access to many of the local pieces of information, the operating system controls the access. The operating system is responsible for ensuring that any user program request to the kernel does not cause the operating system to crash. The transfer of control from the user level program to the system call occurs through the use of a “system call” or “software interrupt/trap”. Before invoking the transfer from the user to the kernel, any information that needs to be transferred from the user program to the system call must be loaded into the registers of the CPU. For pass by value items, this process merely involves placing the value into the register. For pass by reference items, the value placed into the register is known as a “user space pointer”. Since the user space pointer has no meaning to the kernel, we will have to translate the contents of the user space into the kernel such that we can manipulate the information. When returning information from a system call to the user space, information must be placed in the CPU registers to indicate either the success of the system call or the appropriate return value.

In this assignment we are giving you a simulated CPU that models a real CPU. In fact, the simulated CPU is the same as the real CPU (a MIPS chip), but we cannot just run user programs as regular UNIX processes, because we want complete control over how many instructions are executed at a time, how the address spaces work, and how interrupts and exceptions (including system calls) are handled.

Our simulator can run normal programs compiled from C -- see the Makefile in the ‘test’ subdirectory for an example. The compiled programs must be linked with some special flags, and then converted into Nachos format, using the program “coff2noff”. The only caveat is that floating-point operations are not supported.

NOTE : For this project you should ***NOT*** alter the code within the machine directory, only the code within the *userprog* directory.

Phase 3: [80%] Exceptions and IO System Calls

Implement exception handling and handle the basic system calls for file IO. (All system calls are listed in *syscall.h*) We have provided you an assembly-language routine, ‘*syscall*’, to provide a way of invoking a system call from a C routine (UNIX has something similar -- try ‘*man syscall*’). You will need to do the following steps:

- a) Alter *exception.cc* to handle all of system exceptions as listed in *machine/machine.h* Most of the exceptions listed in this file are comprised of run time errors, from which the user program will be unable to recover. The only special cases are *no exception*, which will return control to the operating system and *syscall exception*, which will handle our user system calls. For all other exceptions, the operating system should print an error message and Halt the simulation.
- b) Create a control structure that can handle the various Nachos system calls. Test your control structure by re-implementing the **void Halt()** system call. Make sure that this call operates in the same manner as we discussed during the Nachos walkthrough; it should cause the Nachos simulation to terminate immediately. Test the call’s accuracy with the test user program.

- c) All system calls beyond **Halt()** will require that Nachos increment the program counter before the system call returns. If this is not properly done, Nachos will execute the system call forever. Since the MIPS emulator handles look ahead program counters, as well as normal ones, you will have to emulate the program counter increment code as found in the machine directory. You will have to copy the code into your *syscall* exception handler and insert it at the proper place. For now, you will have to use the **Halt()** system call at the end of each of your user programs.
- d) Implement the **int CreateFile(char *name)** system call. The createfile system call will use the Nachos Filesystem Object Instance to create a zero length file. Remember, the filename exists in user space. This means the buffer that the user space pointer points to must be *translated* from user memory space to system memory space. The createfile system call returns 0 for successful completion, -1 for an error.
- e) Implement the **OpenFileID Open(char *name, int type)** and **int Close(OpenFileID id)** system calls. The user program can open two types of “files” files that can be read only, and others that can be read and write. Each process will allocate a fixed size file descriptor table. For now, set this size to be 10 file descriptors. The first two file descriptors, 0 and 1, will be reserved for console input and console output respectively. The open file system call will be responsible for translating the user space buffers when necessary and allocating the appropriate kernel constructs. For the case of actual files, you will use the filesystem objects provided to you in the filesystem directory. (NOTE : We are using the FILESYSTEM_STUB code). The calls will use the Nachos Filesystem Object Instance to open and close files. The **Open** system call returns the file descriptor id (OpenFileID == an integer number), or -1 if the call fails. Open can fail for several reasons, such as trying to open a file that does not exist or if there is not enough room in the file descriptor table. If the type parameter is set to any other value, the system call should fail. The close system call will take a file descriptor as the parameter. The system call will return -1 on failure and 0 on success.
- f) Implement the **int Read(char *buffer, int charcount, OpenFileID id)** and **int Write(char *buffer, int charcount, OpenFileID id)** system calls. These system calls respectively read and write to a file descriptor ID. Remember, you must translate the character buffers appropriately and you must differentiate between console IO (OpenFileID 0 and 1) and File (read only) or File (read and write) or (any other valid OpenFileID). The read and write interaction will work as follows:

For console read and write, you will use the SynchConsole class, instantiated through the gSynchConsole global variable. You will use the default SynchConsole behaviors for read and write, however you will be responsible for returning the correct types of values to the user. Read and write to Console will return the *actual* number of characters read or written, not the requested number of characters. In the case of read or write failure to console, the return value should be -1. If an end of file is reached for a read operation from the console, the return value should be -2. End of file from the console is returned when the user types in Control-A. Read and write for console will use ASCII data for input and output. (Remember, ASCII data is NULL (\0) terminated)

For file read and write, you will use the supplied classes in file system. You will use the default filesystem behaviors, however, you will return the same type of return values as for synchconsole. Both read and write will return the *actual* number of characters read and written.

Both system calls will return -1 for failure and read will return -2 if the end of file is reached. Read and write for files will use binary data for input and output.

- g) Implement the **int Seek(int pos, OpenFileID id)** system call. Seek will move the file cursor to a specified location. The parameter *pos* will be the absolute character position within a file. If *pos* is a -1 , the position will be moved to the end of file. The system call will return the actual file position upon success, -1 if the call fails. Seeks on console IO will fail. Seeks on the mailbox will move the mailbox pointer to the proper position.
- h) Implement a **createfile** user program to test the createfile system call. You are not going to pass command line arguments to the call, so you will have to either use a fixed filename, or prompt the user for one when you have console IO working. You will have to test for both types of files (read only and read and write).
- i) Implement an **echo** user program. For each line of input from the console, the line gets echoed back to the console.
- j) Implement a **cat** user program. Ask for two filenames, concatenate the files and return it in the second file.
- k) Implement a **reverse** user program. Ask for a source and destination filename; reverse the source file and write into the destination file.
- l) Implement a **interleave** user program. This program asks for a source file, and a destination file, takes input from the console, and writes source file character-interleaved with console input to the destination.
- m) Implement any other tests you feel are necessary to ensure the correctness of your solution. BIG HINT: Implement tests to cover some of the things we changed that are not tested by parts h and l through .

NOTE: A large portion of your grade will depend not only on the correctness of your implementation but how accurately your code conforms to the system call specifications presented in this document. As we are building a robust operating system, the user should not be able to perform any system call that will crash Nachos. The worse case scenario is that the user program might generate a Nachos runtime exception, which you will be handling within part a.

Phase 4: [10%] Adding a new syscall.

In some of the test programs we felt it would have been convenient to have a “remove file” operation especially when temporary files were needed. We decided to add a syscall to remove a nachos file. We do not have to open a file to remove a file. We will use the `fileSystem->Remove` function that uses the filename rather than open file id. Implement a new system call

```
/* Remove the file indicated
 * Parameter is the filename
 *
 */
int RemoveFile(char *filename);
```

RemoveFile will have code 12. You will have to update the **Start.s** file in the **test** directory and the **syscall.h** file in the **userprog** directory. Recompile and test it with a sample program **fileLifecycle**.

Phase 5: [10%] Documentation

This includes internal documentation (comments) and a **BRIEF, BUT COMPLETE** external document (read as: paper) describing what you did to the code and why you made your choices. **DO NOT PARAPHRASE THIS LAB DESCRIPTION AND DO NOT TURN IN A PRINTOUT OF YOUR CODE AS THE EXTERNAL DOCUMENTATION.**

Deliverables and grading

When you complete your project, remove all executables and object files. If you want me to read a message with your code, create a README.NOW file and place it in the nachos *code* directory. Tar and compress the code, and submit the file using the online submission system. It is important that you follow the design guidelines presented for the system calls. I will be running my own shells and test programs against your code to determine how accurately you designed your lab, and how robust your answers are. Grading for the implementation portion will depend on how robust and how accurate your solution is. Remember, the user should not be able to do anything to corrupt the system, and system calls should trap as many error conditions as possible.