

# CSE421/521 Introduction to Operating Systems Spring 2008

## Project #1: Concurrency using processes and threads

Bina Ramamurthy

January 22, 2008

### 1 Objective

To familiarize the students with:

- Programming in C/C++.
- Unix system/library calls, especially *fork*, *pipe*, *exec*, *wait*, *kill* and *exit*.
- Concurrent execution of user and system processes.
- Using POSIX *threads* for concurrency.

### 2 Problem Statement

This project is to be developed in several small steps to help you understand the concepts better.

1. (30 points) The goal of this part is to be able to spawn processes to realize basic concurrent processing and to be able to pass data between the processes spawned. The problem is loosely based on Conway's problem and communicating sequential processes. You will read in a text and print it out in lines of 25 characters after squeezing out redundant blanks, including new-line characters and tabs. The way you will do this is to write a program that takes a text file as a command-line argument. You will accomplish this task using three alternative strategies:
  - (a) Basic C program: Write a program that reads in the text from a file, performs the required operations on the text and prints it out.
  - (b) Spawn process and communicate: Spawn a single child process that will perform the task of reading characters from a file a character-at-a-time and returns the result to the parent using pipe. Parent prints the result.
  - (c) Multiple processes: Spawn three child processes from the parent and let each perform part of the work as follows:
    - i. The first child will open an input file name specified in the command line (`argv[1]`). It will then read the input character-at-a-time. Each character read will be written to a pipe (connected to the second child) except that tabs and new-line characters, when encountered, will be written as blanks.
    - ii. The second child will read the pipe from the first child a character-at-a-time, and pass these (via a second pipe) to the third child. Characters will be read and written as is with one exception: only the first blank of a sequence of blanks will be written.
    - iii. The third child will read the characters from the pipe (from the second child) and print the characters, 25 per line, to the screen.

Use *fork* to spawn the child processes. Use a *pipe* to pass that value between children and the parent.

- (5 + 5 points) The goal of this part is to illustrate the `exec` command and to study sequential and concurrent processes. Write a program that uses the UNIX fork system call to create two processes. Execute from the children appropriate UNIX commands that will clearly provide observable differences in concurrent and sequential mode.

Write your program in two versions:

- Concurrent Mode.* The parent process will fork two child processes almost concurrently (one immediately after another).
  - Sequential Mode.* The parent process will force a sequential execution of the two child processes, to prevent interleaving of outputs. In other words the output should be the concatenation from the two child processes.
- (10 points) Write a shell-like program that illustrates how UNIX spawns processes. This simple program will provide its own prompt to the user, read the command from the input and execute the command. It is sufficient to handle just “argument-less” commands, such as `ls` and `date`.
  - (10 points) Make the mini-shell (from the previous part) a little more powerful by allowing arguments to the commands. For example, it should be able to execute commands such as `more filename` and `ls -l /tmp` etc.
  - (10 points) Add to the mini-shell ability to execute command lines with commands connected by pipes. Example: `ls -l | wc`
  - (30 points) Quicksort is a very efficient algorithm for sorting large data sets. Implement quicksort using a single threaded program and then with multiple threads (one of each divided component) and compare the times of the two implementations. Use Pthreads for realizing the threads.

### 3 Implementation Details

- In general, the execution of any of the programs above will be carried out by specifying the executable program name followed by the command line arguments.
- See the man pages for more details about specific system or library calls and commands: UNIX `fork(2)`, `pipe(2)`, `execve(2)`, `execl(3)`, `execlp(3)`, `cat(1)`, `wait(2)` etc.
- When using system or library calls you have to make sure that your program will exit gracefully when the requested call cannot be carried out.
- One of the dangers of learning about forking processes is leaving unwanted processes active and wasting system time. Make sure each process terminates cleanly when processing is completed. Parent process should wait until the child processes complete, print a message and then quit.
- Your program should be robust. If any of the calls fail, it should print an error message and exit with appropriate error code. Always check for failure when invoking a system or library call.

### 4 Material to be Submitted

- Submit online the **source code each of the programs**. Use meaningful names for the file so that the contents of the file is obvious. Submit a README file that lists the files you have submitted along with a one phrase explanation.
- Only internal documentation is needed. Please state clearly the purpose of each program at the start of the program. Add comments to explain your program. (-5 points, if insufficient.)
- Test runs: It is very important that you show that your program works for all possible inputs. Submit online a **single typescript** file clearly showing the working of all the programs for correct input as well as graceful exit on error input.
- Due date: 2/19/2008 (Saturday), submit on-line before midnight.