

Homework 1
Operating Systems, V22.0202
Fall 2007, Professor Yap

Due: Wed Sep 19, in class

INSTRUCTIONS:

- Please read questions carefully. When in doubt, please ask.
 - The written homework is to be submitted in hardcopy during class, but the programming part sent to us in a single file (as detailed below).
-

1. (5 Points each)

- (i) Illustrate with concrete examples the following four views of an OS:
- (a) Manager of Processes
 - (b) Manager of Resources
 - (c) An extension of the computer hardware
 - (d) An abstract computer

SOLUTION: (a) As a manager of processes, the OS ensures that each process is fairly scheduled on the CPU, is protected from other processes, and has access to system resources. E.g., while one process is waiting for I/O, the OS can switch to another process. E.g., when two processes tries to gain access to the same file, the OS must be able to ensure that the actions of the two processes are not interleaved (one solution is to make a copy of the file for each process).

(b) As a manager of resources, the OS provides protected access to various system resources via system calls. For instance, the file system provides users access to storage space.

(c) As an extender of computer hardware, the OS can make provide convenience services via a useful API. E.g., read/write on disks at the lowest level of detail is very difficult to use. The OS provides higher level views of read/write, and to the user, this appears as an extension of the hardware.

(d) An OS becomes the abstract view of a computer because we no longer think in terms of the underlying hardware (whether it is an Intel or AMD chipset) but in terms of the system services in the OS API (whether it is a Windows or Unix environment). This abstract level greatly simplifies our computer interaction.

- (ii) Exercise 1.11 (p.36). (DMA Question) Direct memory access is used for high-speed I/O devices in order to avoid increasing the CPU's execution load.
- a. How does the CPU interface with the device to coordinate the transfer?
 - b. How does the CPU know when the memory operations are complete?
 - c. The CPU is allowed to execute other programs while the DMA controller is transferring data. Does this process interfere with the execution of the user programs? If so, describe what forms of interference are caused.

SOLUTION: a. All devices have special hardware controllers. Normally, the OS has device drivers (which are kernel programs) that communicate with the controllers. The device drivers have registers, counters and buffers to store arguments and results. Normally, these drivers sits in a tight loop to see the I/O through. But these would tie up the CPU during the I/O. With DMA, The CPU first loads them, and then the device controller takes over.
b. The device controller sends an interrupt to the CPU.
c. You might say there is no interference with user programs, provided you discount interrupts. The DMA controller sends an interrupt when it is done, and this can cause a user process to be suspended.

(iii) Exercise 1.13 (p.36). (Cache question) Give two reasons why caches are useful. What problems do they solve? What problems do they cause? If a cache can be made as large as the device for which it is caching (e.g., a cache as large as a disk), why not make it that large and eliminate the device?

SOLUTION: Reason 1: caches are useful when two or more components need to exchange data, and the components perform transfers at different speeds. E.g., the CPU registers and the disk are two such components. Caches solve this problem by providing a buffer of intermediate speed to store part of the data in the slow component. If the fast component finds the data it needs in the cache, it need not wait for the slower component. The data in the cache must be kept consistent with the data
Reason 2: cache is used in prefetching of instructions in the CPU execution cycle – we introduce a cache to hold the next few instructions, so that we do not have to wait for slow memory access. In this case, the cache is more like a pipeline.
Problems solved by caches: improve speed for communication between components with different speeds. Problems created by caches: since cache hold temporary information for a slower memory component, we need to maintain consistency between the copy in the cache with its actual value in the memory component. This requires extra checking when writing to memory, etc.
Why not just replace the slow component by a large cache? First problem is economic – cache memory is usually much more expensive than the slower memory that it is a surrogate for. Second is difference in technology – usually cache memory is dynamic (it does not retain its information when power is turned off, and thus it cannot be a complete substitute for permanent memory such as a disk).

(iv) Exercise 2.2 (p.73). Five services provided by an OS question that are designed to make it more convenient for users to use the computer system. In what cases would it be impossible for user-level programs to provide these services? Explain. HINT: Think of the system calls an OS provides.

SOLUTION: 1. Program execution. ("exec") The OS loads the executable file into memory and starts execution. [In multiprogramming, we need to schedule the processes, and only an impartial OS can be trusted to do this correctly.]
2. I/O operations. ("read/write") These operations require very low level communication and control. The OS shields the users from having to know these details.
3. File Manipulation ("creation", "deletion", "allocation", "renaming", etc). We need protection and enforcement of file access privileges – these cannot be left to user programs.
4. Communication. ("send", "receive", etc) This requires low level protocols and proper synchronization. So the integrity of the communication system requires kernel protection.
5. Error Handling. Data transfers may be imperfect, so error needs to be detected and corrected. Files need to be checked after a reboot, etc. [All these require system-wide knowledge, and such knowledge is often beyond any user's competence]

2. (10 Points)

An alert reviewer notices a consistent spelling error in the manuscript on Operating Systems textbook, about to go to press. The book has about 700 pages, each with 50 lines of 80 characters each. How long would it take to process (i.e., make one pass through) the manuscript, assuming that the manuscript is in the following levels of the memory hierarchy:

	TYPE	CAPACITY	ACCESS TIME
(a)	Registers	< 1KB	1 nsec
(b)	Cache	1 MB	2 nsec
(c)	Main Memory	1 GB	10 nsec
(d)	Magnetic Disk	50 GB	10 msec
(e)	Magnetic Tape	100 GB	100 sec

For internal storage methods, assume the access time is per character, for disk devices, assume the access time is per block of 1024 characters, and for tape assume the access time is to the time to get to the start of the data, with subsequent access time similar to disk access.

SOLUTION: The file is $700 \times 50 \times 80 = 2.8\text{MB}$ bytes large.

1. There is no way it could fit into the registers of a machine.
2. Cache. It could fit into the cache of a big machine. The time to read this file at 2 ns./byte would take

$$(2 \times 10^{-9} \times (2.8 \times 10^6)) = 5.6 \times 10^{-3} = 5.6\text{ms}.$$
3. Main Memory. This is 10 ns/byte and so is 5 times slower, at 28 ms.
4. Disk. Each block is 1024 Bytes, so the file takes 2800 blocks. Time to access each block is 10ms, so the total time is $10 \times 10^{-3} \times 2800 = 28\text{sec}.$
5. Tape. We need 100 sec to access the tape, and thereafter, it has the same speed as the disk. Thus the time is 128sec.

3. (70 Points)

This question assumes that you have downloaded Cygwin on your computer. Also, you have downloaded gcc, make and tar programs. Please look at our class links for this information.

Enclosed is a simple program called `fork.c` and `test.c` that calls `fork()` and `exec*()`. They are put into a tar file called "hw1.tar". You are to write a variation of these programs.

- 0. Read our quick introductions to Cygwin, C programming and Make program. These are links under RESOURCES on class page.
- 1. Untar the file homework1.tar (move the file to a directory and type "tar -xvf homework1.tar"). Among the files, you will see "fork.c" and "Makefile". Please compile fork.c by typing "make". If gcc compiler creates an executable file called "a.out", you can execute the program by typing "a.out". [Note: a.out is the standard compiler output on UNIX]
- 2. Please create a variation of the fork.c program which will create an arbitrary number of processes until your system crashes (there might be some other possibilities, depending on your system). For safety, we suggest stopping all other programs before you run the test.

REQUIREMENTS:

- A. Be sure that all your child processes are kept around for the duration of the execution, not just created to die immediately.
- B. Only the parent process can fork.
- C. Your program must use the following system calls: `exec*`, `getpid`.
- D. All processes must print useful messages to the screen so that we have a record of how these processes are created, etc. It is up to you to design this as creatively as possible. We will grade you on this.

- 3. There is a simple make program also provided here. You must use this to help you programming and testing.
- 4. Read the following submission rules carefully, as up to 20 percent of this question may be deducted for non-compliance. These rules apply for all programming assignments (with the appropriate modifications).
 - Make sure that your program is well-documented (long and short comments)
 - Include a README file. This file should contain
 - (1) your name,
 - (2) NYU ID,
 - (3) email,
 - (4) phone Contact,
 - (5) the operating system environment for your work,
 - (6) acknowledgement of sources (otherwise it would be plagiarism),
 - (7) any collaboration with others (if none, please say so),
 - (8) any other information you like us to know,
 - (9) and include the following statement: This submission represents my own own.
 - Your C program file name MUST be called h1.c. (We want standard program names to simplify our grading process.)
 - Include an OUTPUT file showing the output of running your h1.c program. Include some description of what happens.
 - You must put h1.c, README, OUTPUT, Makefile, and any other necessary files into a tar file called h1.tar. (Type "tar cvf h1.tar h1.c Makefile ..." to create this tar file).
 - Send the tar file to the grader (but cc to me). After we untar your file, we expect to type "make" to compile it, and to type "make test" to run the compiled program.

SOLUTION:

We grade your conformance with the instructions above.

In particular, we want you to keep the processes around after creation. To do this, you would need to call "exec" to sleep (as in the sample program).

It would be nice if you also make each process announce its presence to the world (echo or printf).

```

// file: test.c
// Comp.Sys.Org.II, Spring 2005, Yap.
// The following code is a simple modification of the "fork.c" code from
//
// Dr Ian G Graham, Department of Information Technology, GUGC
// Copyright 2000-2003
//
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <string.h>
#include <errno.h>
#include <sys/wait.h>

extern int errno;    /* system error number */

void syserr(char *); /* error report, abort routine */

int main(int argc, char *argv[])
{
    pid_t pid;
    int rc;

    pid = getpid();
    printf("Process ID before fork: %d\n", pid);

    switch (fork()) {
        case -1:
            syserr("fork");
        case 0:

/* execution in child process */
            pid = getpid();
            printf("Process ID in child after fork: %d\n", pid);
            execlp("echo", "0", "Child to world: Hello!", NULL);
            syserr("execl");
        }

/* continued execution in parent process */

        pid = getpid();
        printf("Process ID in parent after fork: %d\n", pid);

        //getchar(); // hang until I type something (alternatively,
        // let parent wait for return of child process.
        exit(0);
    }

void syserr(char * msg)
{
    fprintf(stderr, "%s: %s", strerror(errno), msg);
    //abort(errno); // variant
    abort();
}

```