

Topic Notes: Processes and Threads

What is a process? Our text defines it as “an abstraction of a running program.”

Definitions from other textbooks:

- A program in execution (good one)
- An asynchronous activity
- the “locus of control” of a program in execution
- that which is manifest by the existence of a process control block in the OS
- that entity which is assigned to processors
- the “dispatchable” unit
- the “animated spirit” of a procedure

A process is sequential.


Parts of a process:

- program code (text section)
- program counter and other registers
- stack (local variables and function call information)
- data (global variables)

A typical multiprogrammed system has many processes at any time. Try `ps -aux` or `ps -ef` to see the processes on your favorite Unix system. Only one of these can be on the processor at a time.

So if a process is in the system but not executing on the CPU, where is it?

States of a process:



What information do we need to let a process transition among these states?

Think about it in terms of what a person needs to do to get back to what he or she was doing before being interrupted.

- you're sitting in the lab working hard on your OS project and someone interrupts you
- your attention shifts and you go off and do something else for a while
- then when you need to come back to the work, you need to remember *what* you were doing, and *where you were* in the process of doing it

We do this all the time, and many of us are really pretty good at it. A processor can't just pick up where it left off, unless we carefully remember everything it was doing when it was so rudely interrupted.

Process Control Block (PCB) – information needed to save and restore a process.

- Process state (running, waiting, ready)
- Process identifier (PID)
- Program counter
- Other CPU registers
- CPU scheduling information
- Memory-management information
- Accounting information
- I/O status information

In many Unix systems, the PCB is divided between two structures:

- The *proc* structure, which holds the process information needed at all times, even when the process is not active (swapped out).
- The *user* structure, which holds other information about the process.

Historically, it was important to keep as much in the user structure and as little in the proc structure as possible, because of memory constraints. As memory has gotten larger, this has become less important.

In FreeBSD, for example, most of the information is now in the proc structure. We can see it on your favorite FreeBSD system in `/usr/src/sys/sys/proc.h`. See `struct proc` to see the structure at line 495.

The user structure is in `/usr/src/sys/sys/user.h`, and it is very small. It contains per-thread information. Part of this is a `struct pcb`, which is an architecture-dependent structure defined in `/usr/src/sys/i386/include/pcb.h` (line 44). Here you find the actual x86 registers that need to be saved when a process is removed from the CPU. For comparison, check out the Alpha version in `/usr/src/sys/alpha/include`.

The transition is called a *context switch*. This is pure overhead, so it needs to be fast. Some systems do it faster than others. Hardware support helps - more on that later.

This requires at least two levels of privilege, or modes of operation:

- User mode: execution of user processes. Low privilege. Can only access its own memory, etc.
- Monitor mode (kernel mode, system mode): operating system execution. Can do things like switch among user processes, shut down the system, etc.

Hardware support is needed. Privileged instructions can be executed only while in monitor mode. We'll talk more about this as we go forward.

Process Creation/Deletion

In Unix, every process (except for the first) is created from an existing process. See `ps` again for examples. The processes form a tree, with the root being the `init` process (PID=1).

When a new process is created, what information does it inherit from or share with its parent?

- Does it get any resources that were allocated to the parent?
- Does the parent wait for the child to complete, or do they execute concurrently?
- Is the child a duplicate of the parent, or is it something completely different?
- If it's a duplicate, how much context do they share?
- Can the parent terminate before the child?

Creation of a new process is a highly privileged operation. It requires the allocation of a new PCB, insertion of this PCB into system data structures, among other operations that we would not trust to regular "user" programs.

The typical solution to this is that such operations are provided through *system calls*. These system calls are functions that are part of the operating system and are permitted to perform some tasks that a normal process is not able to do on its own. In effect, the program temporarily gains a higher privilege level while executing the system call.

In Unix, the `fork()` system call duplicates a process. The child is a copy of the parent - in execution at the same point, the statement after the return from `fork()`.

The return value indicates if you are child or parent.

0 is child, > 0 means parent, -1 means failure (limit reached, permission denied)

Example C program:

```
pid=fork();
if (pid) {
    parent stuff;
}
else {
    child stuff;
}
```

A more complete program that uses `fork()` along with three other system calls (`wait()`, `getpid()`, and `getppid()`) is in `forking.c`:

See Example:

`/cluster/examples/forking`

Some comments about this program:

First, we run it to observe what happens. Note that there is only one copy of the printout before the `fork()`, two of the one after. `fork()` is a very unusual function - you call it once, but it returns twice!

How do we know how these work? See the man pages! The `fork` page tells us what we need to include, and how to use it.

How about `wait`? If we issue the command `man wait`, we get the man page for `builtin`, as there is a built-in `wait` command in the shell. To get the page we want, we need to specify a manual “section.” Section 2 is the system calls section. The command `man 2 wait` will get us the page we want.

Are there really two processes? Let’s look at the output of `ps` as the program runs.

Again, these system calls let you, as a normal user, do things that only the system can really do. Your “user mode” process can get access to “kernel mode” functionality through these calls.

How many processes can we create on various systems? Where does this limit come from? Can we create enough processes to take down a system?

See Example:

`/cluster/examples/forkbomb`

If we look at the `ps` output on a Unix system, we will see a lot of processes owned by `root`. Many of these are essential parts of the system and are intended to continue running as long as the system is up. These processes are called *daemons* and are the motivation for the BSD logo.

The FreeBSD implementation of `fork()` is in `/usr/src/sys/kern/kern_fork.c`.

Things to note here (for an example – don’t worry about the details):

- the action is happening in `fork1()`
- Line 271: actually allocate the new proc structure with `uma_zalloc`, a special `malloc` basically (see `zalloc(9)`).
- Fill in some of the structure.
- Line 289: check if a new process is allowed based on global system limits.
- `nprocs` is the kernel variable that stores the number of active processes in the system.
- Line 322: find an available pid. Note the `gotos`.
- Line 446: initialize proc table entry
- Line 622: attach new process to parent – or to init process if a certain flag is set to say that the parent should not wait for its child.
- Line 644: deal with some memory for the new child.
- Line 679: set on run queue, so the child can start executing.

We will see that in many cases, you will want to use the `vfork(2)` system call instead of `fork()`. This one doesn't copy the address space of the process – it assumes you are going to replace the newly created process' program with a new one.

In the Windows world, there is a `CreateProcess()` Win32 call (Win32 is like a POSIX for Windows) that creates a new process and loads the correct program into that new process.

Later, we will consider more system calls, including ones that let you do things more interesting than making a copy of yourself.

We also will consider later how all this gets started – how does that first process get started that forks everything else?

Processes may need to communicate with each other in some more interesting way than making copies of themselves. We will see a number of ways this can be done, but in your lab this week, you will make use of a small chunk of POSIX shared memory.

The textbook has examples of how to create a shared memory segment that can be accessed by your parent and child process.

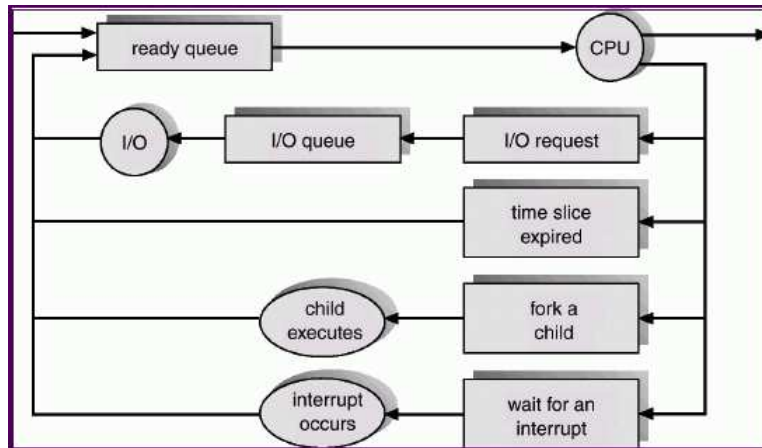
Another possibility is to have processes pass messages to each other over the network or through the file system.

We will spend a good chunk of time later this month on cooperating processes.

Once we have all of these processes, we need to manage them. Recall the process states.

Basically, each process in the system has to be in one of several queues. These are *process scheduling queues*

The system maintains a *process table* and the PCBs are stored in these process scheduling queues.



The selection of a process from the ready queue to run on the CPU is the topic for our CPU scheduling lectures.

A fundamental job of the OS is to manage these queues.

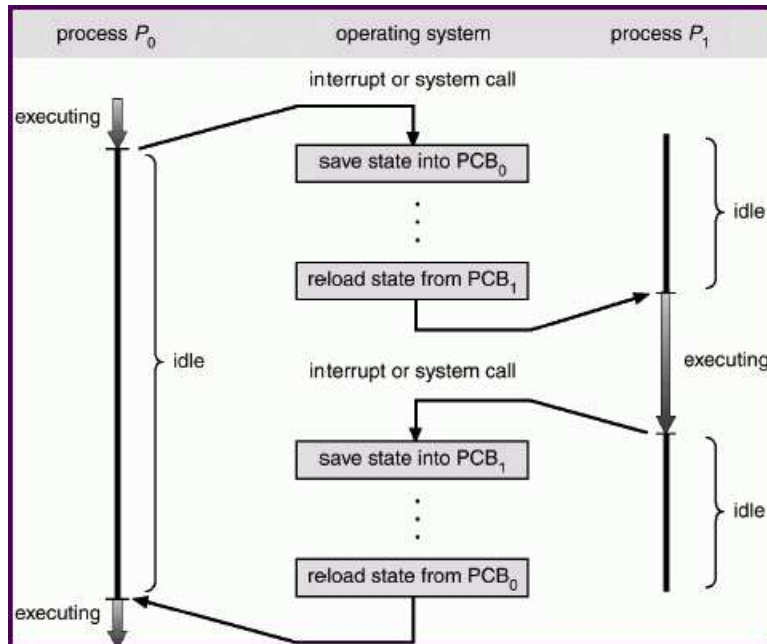
FreeBSD defines its run/ready queues in `/usr/src/sys/kern/kern_switch.c`

There are actually 64 regular queues (this number is defined in `/usr/src/sys/sys/runq.h`). We'll think more about what these are about when we talk about CPU scheduling.

As processes go through the system, processes will be assigned to the CPU (dispatched) or removed from the CPU dozens or even hundreds of times per second. Each of these swaps is a *context switch*, saving a CPU state in a PCB, then restoring another PCB.

Remember, the context switch is pure overhead. The system is not doing any useful computation when it's working on a context switch. It had better be fast.

Diagram of context switch:



FreeBSD does this in what is necessarily an architecture-dependent routine.

In fact, it's an assembly source file, not C! `/usr/src/sys/i386/i386/switch.s`

Notes about the context switch:

- Line 95: an assembly code entry point (essentially function) `cpu_switch`.
- Line 108: save hardware registers into PCB.
- Line 144: save floating point stuff – note “big C function”
- Line ???: choose a new proc: come back next week for more
- Lots of other stuff....
- Line 253: restore PCB of new proc
- ends up returning to the new proc, picking up where it left off

Registers are just one part of a process' context. What about memory? What about stuff in a cache?

We'll talk about main memory around spring break. Each process in the ready/run queues has some main memory allocated to it.

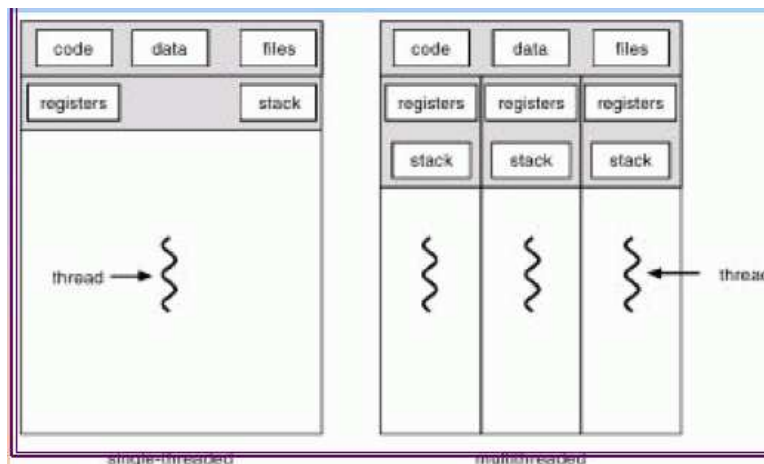
But cache is another story. Remember how a typical cache is set up. Closest to the CPU and registers in the memory hierarchy.

As the CPU requests memory, lines of values from memory are brought into the cache. If all goes well, these lines will be reused.

But when a context switch occurs, the things from the process that was on the CPU that are now in cache seem unlikely to survive there until that process gets another chance on the CPU. We may force a lot of cache misses, adding more overhead to the context switch cost.

Threads/Lightweight Processes

Those of you who have taken or TA'ed for the “objectdraw” version of the intro class here know about Java threads. You may have called them `ActiveObjects`. These gave you the ability to have your programs doing more than one thing at a time. Such a program is called *multithreaded*.



Threads are a stripped-down form of a process:

- a *task* or *process* can contain multiple *threads*
- Threads share process context
- Threads are asynchronous
- Threads have less context than processes
- Threads can be created/destroyed at a lower cost than processes
- Threads cooperate to do a process in parallel with (relatively) fine granularity of parallelism
- Threads are well-suited to symmetric multiprocessors (SMPs).

Threads can be useful even in single-user multitasking systems:

- Foreground and background work: one thread to deal with user input, another to update display, etc such as in a spreadsheet or word processor.
- Asynchronous processing: auto-save features, background computation.

- Speed execution: compute with one batch of data while reading another. **Your** process can be in I/O and on the CPU at the same time!
- Organizing programs: cleaner and more elegant program design.

A “normal” single-thread process has its program text, global address space, a stack, and a PC.

A multithreaded process has its program text, a single shared global address space, but a stack and a PC for each thread.

See SG&G Figure 4.1.

Each thread also needs to be able to access the registers when it’s on the CPU, so while the context switch among threads doesn’t have to worry about everything, it will still need to save registers.

Sample thread application: An NFS (network file system) server.

A single thread of execution would be unreasonable for all but the most lightly-loaded servers. Realistically, a new process or thread needs to be created for each request.

With “heavyweight” processes, each request from the network would require a heavyweight process creation. A series of requests would result in frequent, short-duration processes being created and destroyed. Significant overhead would result.

With multithreading, a single NFS server task (a process) can exist, and a new thread would be created within the task for each request (“**pop-up**” threads). There is less overhead, plus if there are multiple CPUs available, threads can be assigned to each for faster performance.

User Threads vs. Kernel Threads

Threads may be implemented within a process (in “user space,” also often referred to as “userland”) or by the kernel (in “kernel space”).

With user threads, the kernel sees only the main process, and the process itself is responsible for the management (creation, activation, switching, destruction) of the threads.

With kernel threads, the kernel takes care of all that. The kernel schedules not just the process, but the thread.

There are significant advantages and disadvantages to each. Some systems provide just one, some systems, such as Solaris, provide both.

One big thing is that with user threads, it is often the case that when a user thread blocks (such as for I/O service), the entire process blocks, even if some other threads could continue. With kernel threads, one thread can block while others continue.

pthreads, POSIX threads

We saw how to use `fork()` to create threads in Unix. There is a POSIX standard interface to create threads.

A google search for “pthread tutorial” yields many results.

Instead of creating a copy of the process like `fork()`, create a new thread, and that thread will call a function *within the current task*.

The basic pthread functions are:

Simplest form:

```
int pthread_create(pthread_t *thread, const pthread_attr_t *attr,
                  void * (*start_routine)(void *),
                  void *arg);
```

```
int pthread_join(pthread_t thread, void **status);
```

```
void pthread_exit(void *value_ptr);
```

- `pthread_create(3THR)` – Well, this creates a new thread. It takes 4 arguments. The first is a pointer to a variable of type `pthread_t`. Upon return, this contains a thread identifier that is used later in `pthread_join()`. The second is a pointer to a `pthread_attr_t` that specifies thread creation attributes. In the `pthreadhello` program, we pass in `NULL`, which specifies the system default attributes. The third argument is a pointer to a function that will be called when the thread is started. This function must take a single parameter of type `void *` and return `void *`. The fourth parameter is the pointer that will be passed as the argument to the thread function.
- `pthread_exit(3THR)` – This causes the calling thread to exit. This is called implicitly if the thread function returns. Its argument is a return status value, which can be retrieved by `pthread_join()`.
- `pthread_join(3THR)` – This causes the calling thread to block until the thread with the identifier passed as the first argument to `pthread_join()` has exited. The second argument is a pointer to a location where the return status passed to `pthread_exit()` can be stored. In the `pthreadhello` program, we pass in `NULL`, and hence ignore the value.

Prototypes for pthread functions are in `pthread.h` and programs need to link with `libpthread.a` (use `-lpthread` at link time). When using the Sun compiler, the `-mt` flag should also be specified to indicate multithreaded code.

To compile under FreeBSD (lab machines), need to pass a special flag `-pthread` to the compiler and linker to tell it to use the reentrant C library `libc_r.a`.

It is also a good idea to do some extra initialization, to make sure the system will allow your threads to make use of all available processors. It may, by default, allow only one thread in your program to be executing at any given time. If your program will create up to n concurrent threads, you should make the call:

```
pthread_setconcurrency(n+1);
```

somewhere before your first thread creation. The “+1” is needed to account for the original thread plus the n you plan to create.

You may also want to specify actual attributes as the second argument to `pthread_create()`. To do this, declare a variable for the attributes:

```
pthread_attr_t attr;
```

and initialize it with:

```
pthread_attr_init(&attr);
```

and set parameters on the attributes with calls such as:

```
pthread_attr_setscope(&attr, PTHREAD_SCOPE_PROCESS);
```

I recommend the above setting for threads in Solaris.

Then, you can pass in `&attr` as the second parameter to `pthread_create()`.

Any global variables in your program are accessible to all threads. Local variables are directly accessible only to the thread in which they were created, though the memory can be shared by passing a pointer as part of the last argument to `pthread_create()`.

A pthread “hello, world” program is in `pthreadhello.c`.

See Example:

```
/cluster/examples/pthreadhello
```

We said that threads share more context than processes. To see this, try out `what_shared.c`.

See Example:

```
/cluster/examples/what_shared
```

A slightly more interesting example is `proctree_threads`.

See Example:

```
/cluster/examples/proctree_threads
```

This example builds a “tree” of threads to a depth given on the command line. It includes calls to `pthread_self()`. This function returns the thread identifier of the calling thread.

This is worth a look, since it’s a pthreads version of the program you will be writing for Lab 1, but you’ll be using processes created with `fork` instead of threads created with `pthread_create`.

Other Popular Thread Implementations

We will not say much about other thread implementations, but many of the concepts and challenges are similar:

- Windows has a Win32 thread interface which is similar in capability to pthreads.
- Java threads are a fundamental part of the language and JVM.