



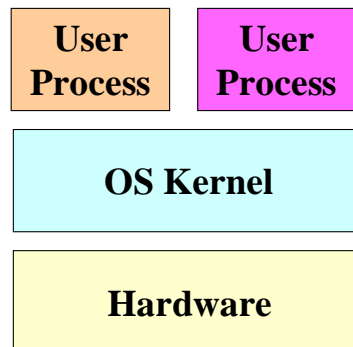
Processes

CS 217



Operating System


- Supports virtual machines
 - Promises each process the illusion of having whole machine to itself
- Provides services:
 - Protection
 - Scheduling
 - Memory management
 - File systems
 - Synchronization
 - etc.



What is a Process?




- A process is a running program with its own ...
 - Processor state
 - » PC, PSR, registers
 - Address space (memory)
 - » Text, bss, data, heap, stack



Operating System



- Resource allocation
 - Sharing
 - Protection
 - Fairness
 - Higher-level abstractions
- Common strategies
 - Chop up resources into small pieces and allocate small pieces at fine-grain level
 - Introduce level of indirection and provide mapping from virtual resources to physical ones
 - Use past history to predict future behavior



Example: Process Scheduling




- We have a single physical CPU and a whole lot of processes/jobs to run
 - Which process do we run next?
 - For how long do we run it?

Process	Time
1	10
2	29
3	3
4	7
5	12

CPU-bound processes

- Solution 1:
 - Run each of them to completion in first-come first-served order



Average wait time of processes: $(0 + 10 + 39 + 42 + 49) / 5 = 28$

Example: Process Scheduling




- We have a single physical CPU and a whole lot of processes/jobs to run
 - Which process do we run next?
 - For how long do we run it?

Process	Time
1	10
2	29
3	3
4	7
5	12

CPU-bound processes

- Another solution:
 - Run them to completion in shortest-first order




Average wait time of processes: $(0 + 3 + 10 + 20 + 32) / 5 = 13$

CPU-I/O Burst Cycle




- Typical execution of process:



CPU-I/O Burst Cycle



- Schedule CPU burst for process B while process A is waiting for I/O
 - Better utilize CPU




Time Slicing



- Divide up time into quantums
 - Schedule quantums, not complete jobs
 - Schedule another process if perform I/O
 - Preempt process at end of quantum
- Motivations
 - CPU-I/O Burst Cycle
 - Interactive response

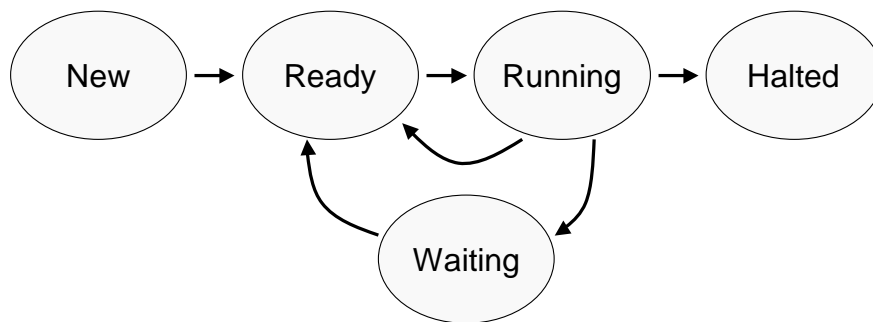
Process	Time
1	10
2	29
3	3
4	7
5	12



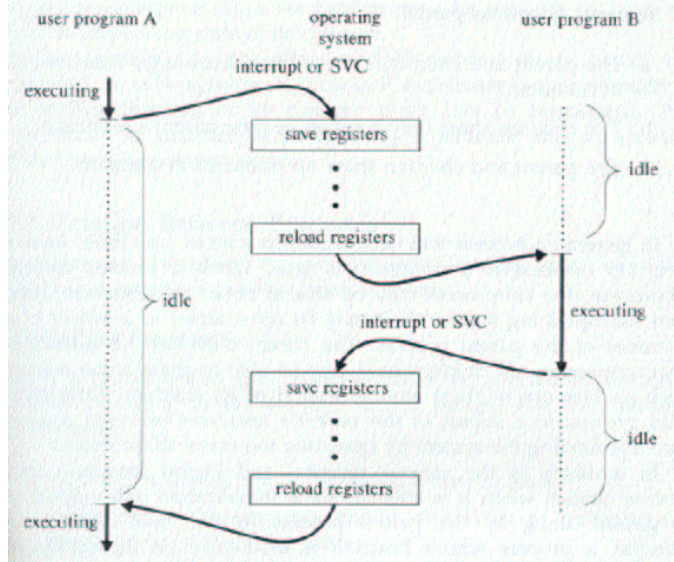
Life Cycle of a Process



- Running: instructions are being executed
- Waiting: waiting for some event (e.g., i/o finish)
- Ready: ready to be assigned to a processor



Context Switch



Process Control Block



- For each process, the kernel keeps track of ...
 - Process state (new, ready, waiting, halted)
 - CPU registers (PC, PSR, global, local, ...)
 - CPU scheduling information (priority, queues, ...)
 - Memory management information (page tables, ...)
 - Accounting information (time limits, group ID, ...)
 - I/O status information (open files, I/O requests, ...)

Fork



- Create a new process (system call)
 - child process inherits state from parent process
 - parent and child have separate copies of that state
 - parent and child share access to any open files

```
pid = fork();
if (pid != 0) {
    /* in parent */
    ...
}
/* in child */
...
```

Fork



- Inherited:
 - user and group IDs
 - environment
 - close-on-exec flag
 - signal handling settings
 - supplementary group IDs
 - set-user-ID mode bit
 - set-group-ID mode bit
 - profiling on/off/mode status
 - debugger tracing status
 - nice value
 - stdin
 - scheduler class
 - all shared memory segments
 - all mapped files
 - file pointers
 - non-degrading priority
 - process group ID
 - session ID
 - current working directory
 - root directory
 - file mode creation mask
 - resource limits
 - controlling terminal
 - all machine register states
 - control register(s)
- Separate in child
 - process ID
 - address space (memory)
 - file descriptors
 - active process group ID.
 - parent process ID
 - process locks, file locks, page locks, text locks and data locks
 - pending signals
 - timer signal reset times
 - share mask

Exec



- Overlay current process image with a specified image file (system call)
 - affects process memory and registers
 - has no affect on file table
- Example:

```
execlp("ls", "ls", "-l", NULL);  
fprintf(stderr, "exec failed\n");  
exit(1);
```

Exec (cont)



- Many variations of **exec**

```
int execlp(const char *file,  
           const char *arg, ...)  
int execl(const char *path,  
           const char *arg, ...)  
int execv(const char *path,  
           char * const argv[])  
int execl(const char *path,  
           const char *arg, ...,  
           char * const envp[])
```
- Also **execve** and **execvp**

Fork/Exec



- Commonly used together by the shell

```
... parse command line ...
pid = fork()
if (pid == -1)
    fprintf(stderr, "fork failed\n");
else if (pid == 0) {
    /* in child */
    execvp(file, argv);
    fprintf(stderr, "exec failed\n");
}
else {
    /* in parent */
    pid = wait(&status);
}
... return to top of loop ...
```

Wait



- Parent waits for a child (system call)
 - blocks until status of a child changes
 - returns pid of the child process
 - returns -1 if no children exist (already exited)

```
pid_t wait(int *status);
```

System



- Convenient way to invoke fork/exec/wait
 - Forks new process
 - Execs command
 - Waits until it is complete

```
int system(const char *cmd);
```

- Example:

```
int main()
{
    system("echo Hello world");
}
```

Summary



- Operating systems manage resources
 - Divide up resources (e.g., quantum time slices)
 - Allocate them (e.g., process scheduling)
 - A processes is a running program with its own ...
 - Processor state
 - Address space (memory)
 - Create and manage processes with ...
 - `fork`
 - `exec`
 - `wait`
 - `system`
- } Used in shell