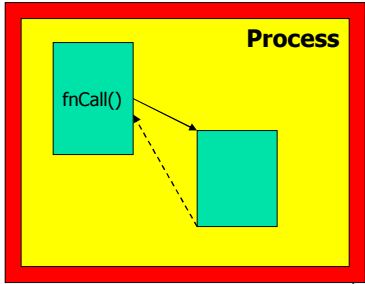
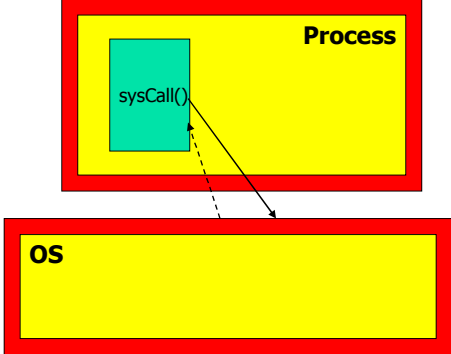


Brief OS Review

Tarek Abdelzaher

Copyright ©: Nahrstedt, Angrave, Abdelzaher 1

System Calls

| Function Call | System Call |
|---|--|
|  |  |
| <p>Caller and callee are in the same Process</p> <ul style="list-style-type: none">- Same user- Same "domain of trust" | <ul style="list-style-type: none">- OS is trusted; user is not.- OS has super-privileges; user does not- Must take measures to prevent abuse 2 |

Examples of System Calls

- Example:
 - `getuid()` //get the user ID
 - `fork()` //create a child process
 - `exec()` //executing a program
- Don't mix system calls with standard library calls
 - Differences?
 - Is `printf()` a system call?
 - Is `rand()` a system call?

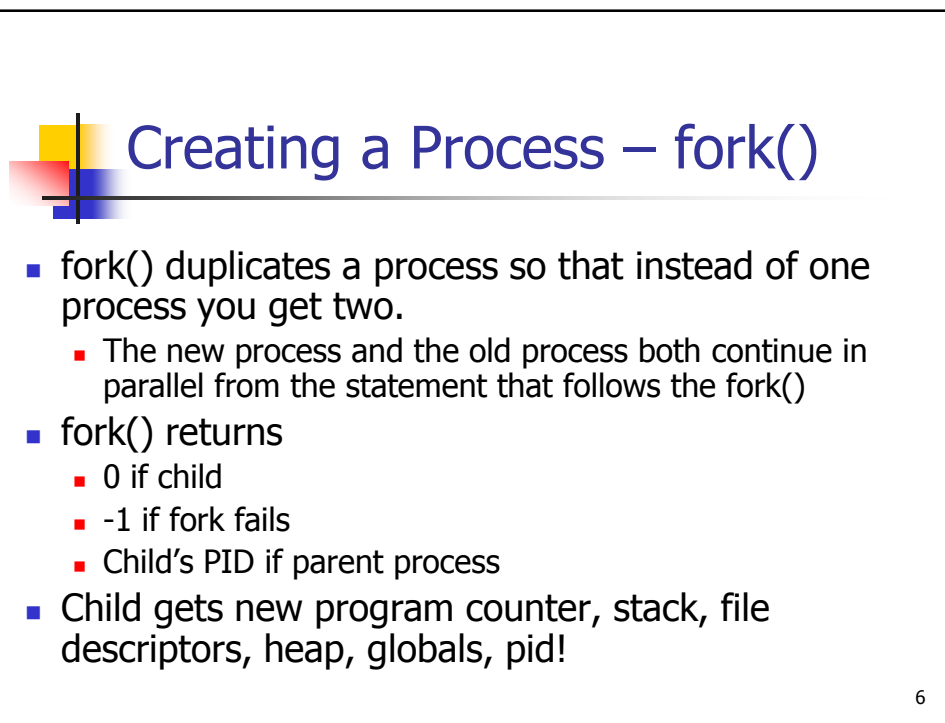
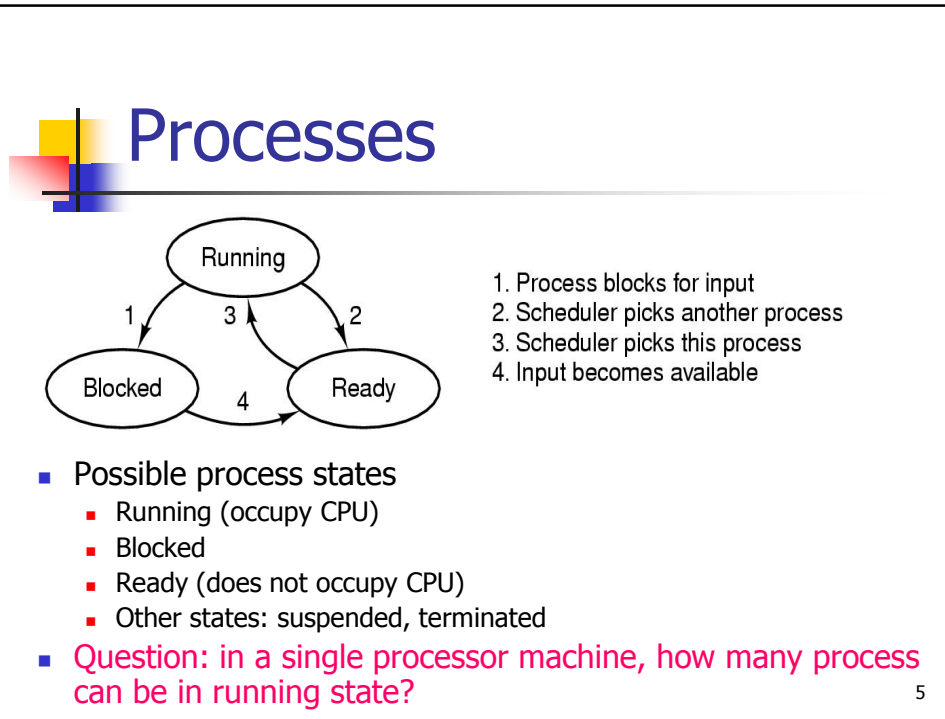
3

I/O Library Calls versus System Calls

- Each system call has analogous procedure calls from the standard I/O library:

| System Call | Standard I/O call |
|---------------------------|------------------------------|
| ■ <code>open</code> | <code>fopen</code> |
| ■ <code>close</code> | <code>fclose</code> |
| ■ <code>read/write</code> | <code>getchar/putchar</code> |
| ■ | <code>getc/putc</code> |
| ■ | <code>fgetc/fputc</code> |
| ■ | <code>fread/fwrite</code> |
| ■ | <code>gets/puts</code> |
| ■ | <code>fgets/fputs</code> |
| ■ | <code>scanf/printf</code> |
| ■ | <code>fscanf/fprintf</code> |
| ■ <code>lseek</code> | <code>fseek</code> |

4





exec() Function

- Exec function allows child process to execute code that is different from that of parent
- Exec family of functions provides a facility for overlaying the process image of the calling process with a new image.
- Exec functions return -1 and sets errno if unsuccessful

7

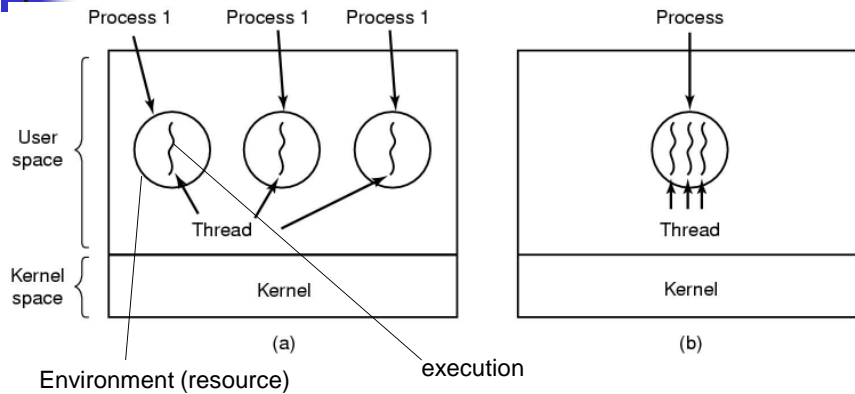


Threads

| POSIX function | description |
|----------------|---------------------------------------|
| pthread_create | create a thread |
| pthread_detach | set thread to release resources |
| pthread_equal | test two thread IDs for equality |
| pthread_exit | exit a thread without exiting process |
| pthread_kill | send a signal to a thread |
| pthread_join | wait for a thread |
| pthread_self | find out own thread ID |

8

Threads: Lightweight Processes



- (a) Three processes each with one thread
- (b) One process with three threads


9

Trade-offs?

- Kernel thread packages
 - Each thread can make blocking I/O calls
 - Can run concurrently on multiple processors
- Threads in User-level
 - Fast context switch
 - Customized scheduling

10

Hybrid Implementations (Solaris)



- Multiplexing user-level threads onto kernel-level threads

11

Synchronization

- Processes and threads can be preempted at arbitrary times, which may generate problems.
- Example: What is the execution outcome of the following two threads (initially $x=0$)?

Thread 1:

Read X
Add 1
Write X

Thread 2:

Read X
Add 1
Write X

12



Critical Region (Critical Section)

```

Process {
    while (true) {
        ENTER CRITICAL SECTION
        Access shared variables;
        LEAVE CRITICAL SECTION
        Do other work
    }
}

```

13



Semaphores

```

wait (sem_t *sp)
    if (sp->value >0) sp->value--;
    else {
        Add process to sp->list;
        <block>
    }

signal (sem_t *sp)
    if (sp->list != NULL)
        remove next process from sp->list;
    else sp->value++;

```

14



Mutex

- Simplest and most efficient thread synchronization mechanism
- A special variable that can be either in
 - **locked state**: a distinguished thread that *holds* or *owns* the *mutex*; or
 - **unlocked state**: no thread holds the *mutex*
- When several threads compete for a *mutex*, the losers block at that call
 - The *mutex* also has a queue of threads that are waiting to hold the *mutex*.
- POSIX does not require that this queue be accessed FIFO.

15



POSIX Mutex-related Functions

- int **pthread_mutex_init**(pthread_mutex_t *restrict mutex, const pthread_mutexattr_t *restrict attr);
 - Also see PTHREAD_MUTEX_INITIALIZER
- int **pthread_mutex_destroy**(pthread_mutex_t *mutex);
- int **pthread_mutex_lock**(pthread_mutex_t *mutex);
- int **pthread_mutex_trylock**(pthread_mutex_t *mutex);
- int **pthread_mutex_unlock**(pthread_mutex_t *mutex);

16

Scheduling


- Basic scheduling algorithms
 - FIFO (FCFS)
 - Shortest job first
 - Round Robin
 - Priority Scheduling

17

Preemptive vs. Non-preemptive scheduling

■ Non-preemptive scheduling:

- The running process keeps the CPU until it **voluntarily** gives up the CPU
 - process exits
 - switches to blocked state
 - 1 and 4 only (no 3)



■ Preemptive scheduling:

- The running process can be interrupted and must release the CPU (can be **forced** to give up CPU)

18



Signals

- Signal is *generated* when the event that causes it occurs.
- Signal is *delivered* when a process receives it.
- The *lifetime* of a signal is the interval between its generation and delivery.
- Signal that is generated but not delivered is *pending*.
- Process *catches* signal if it executes a *signal handler* when the signal is delivered.
- Alternatively, a process can *ignore* a signal when it is delivered, that is to take no action.
- Process can temporarily prevent signal from being delivered by *blocking* it.
- *Signal Mask* contains the set of signals currently blocked.

19



Examples of POSIX Required Signals

| Signal | Description | default action |
|----------------|--|-----------------------------|
| SIGABRT | process abort | implementation dependent |
| SIGALRM | alarm clock | abnormal termination |
| SIGBUS | access undefined part of memory object | implementation dependent |
| SIGCHLD | child terminated, stopped or continued | ignore |
| SIGILL | invalid hardware instruction | implementation dependent |
| SIGINT | interactive attention signal (usually ctrl-C) | abnormal termination |
| SIGKILL | terminated (cannot be caught or ignored) | abnormal termination |

20

Examples of POSIX Required Signals

| Signal | Description | default action |
|----------------|--|-----------------------------------|
| SIGSEGV | Invalid memory reference | implementation dependent |
| SIGSTOP | Execution stopped | stop |
| SIGTERM | termination | Abnormal termination |
| SIGTSTP | Terminal stop | stop |
| SIGTTIN | Background process attempting read | stop |
| SIGTTOU | Background process attempting write | stop |
| SIGURG | High bandwidth data available on socket | ignore |
| SIGUSR1 | User-defined signal 1 | abnormal termination 21 |

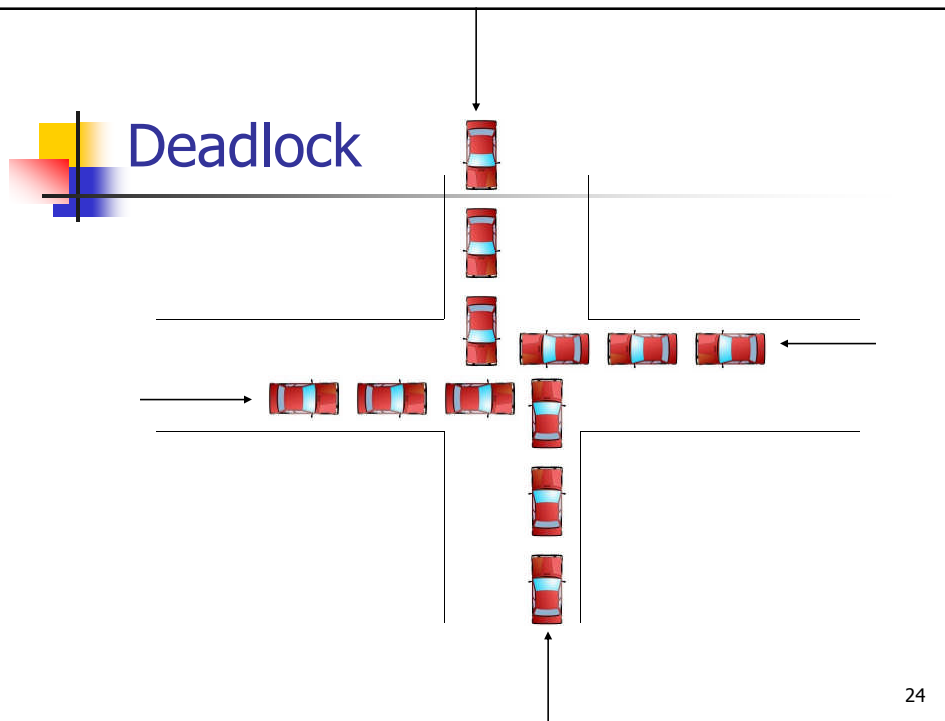
Command Line Generates Signals

- You can send a signal to a process from the command line using **kill**
- `kill -l` will list the signals the system understands
- `kill [-signal] pid` will send a signal to a process.
 - The optional argument may be a name or a number (default is SIGTERM).
- To unconditionally kill a process, use:
 - `kill -9 pid` which is `kill -SIGKILL pid`.

Signal Masks

- Process can temporarily prevent signal from being delivered by *blocking* it.
- *Signal Mask* contains a set of signals currently blocked.
- **Important!** Blocking a signal is different from ignoring signal. Why?
- When a process blocks a signal, the OS does not deliver signal until the process unblocks the signal
 - A *blocked* signal is not delivered to a process until it is unblocked.
- When a process ignores signal, signal is delivered and the process handles it by throwing it away.

23



24


Deadlock

- Mutual exclusion
- Hold and wait condition
- No preemption condition
- Circular wait condition

- Original scenario & our proposed ritual had all four of these properties.


25

Resource Allocation Graph



- resource R assigned to process A
- process B is requesting/waiting for resource S
- process C and D are in deadlock over resources T and U


26



Strategies

- Strategies for dealing with Deadlocks
 - detection and recovery
 - dynamic avoidance (at run-time)
 - prevention (by off-line design)
 - negating one of the four necessary conditions

27



Other Issues

- Memory management
- Files
- I/O

28