

## **Administrivia**

- **Submit paper writeups in plain text**
  - No attachments. No .doc, .pdf, .rtf, HTML, etc.
- **Presentation assignments going out tomorrow**
- **Email: `mfreed+cos518@cs.princeton.edu`**
- **Question: Cut some older papers for more recent?**

# Schedule for Sept 20, 2007

- **Today (MJF)**

- Review basic system programming
  - Thanks to David Mazières for slides and examples
- Discuss receive livelock paper

- **Tuesday (MJF)**

- Review programming with threads and events
- Flash: Async I/O (events) + process pools
- Tame: Simplifying event-based programming

- **Thursday**

- RPCs + Remote Objects (One of you)
- Review SUN-RPC (MJF)

# System calls

- **Problem: How to access resources other than CPU**
  - Disk, network, terminal, other processes
  - CPU prohibits instructions that would access devices
  - Only privileged OS “kernel” can access devices
- **Applications request I/O operations from kernel**
- **Kernel supplies well-defined *system call* interface**
  - Applications set up syscall arguments and *trap* to kernel
  - Kernel performs operation and returns result
- **Higher-level functions built on syscall interface**
  - `printf`, `scanf`, `gets`, etc. all user-level code

# I/O through the file system

- **Applications “open” files/devices by name**
  - I/O happens through open files
- `int open(char *path, int flags, ...);`
  - flags: `O_RDONLY`, `O_WRONLY`, `O_RDWR`
  - `O_CREAT`: create the file if non-existent
  - `O_EXCL`: (w. `O_CREAT`) create if file exists already
  - `O_TRUNC`: Truncate the file
  - `O_APPEND`: Start writing from end of file
  - mode: final argument with `O_CREAT`
- **Returns file descriptor—used for all I/O to file**

## Error returns

- **What if open fails? Returns -1 (invalid fd)**
- **Most system calls return -1 on failure**
  - Specific kind of error in global int errno
- **#include <sys/errno.h> for possible values**
  - 2 = ENOENT “No such file or directory”
  - 13 = EACCES “Permission Denied”
- **perror function prints human-readable message**
  - perror ("initfile");  
→ “initfile: No such file or directory”

# Operations on file descriptors

- `int read (int fd, void *buf, int nbytes);`
  - Returns number of bytes read
  - Returns 0 bytes at end of file, or -1 on error
- `int write (int fd, void *buf, int nbytes);`
  - Returns number of bytes written, -1 on error
- `off_t lseek (int fd, off_t pos, int whence);`
  - whence: 0 – start, 1 – current, 2 – end
  - Returns previous file offset, or -1 on error
- `int close (int fd);`
- `int fsync (int fd);`
  - Guarantee that file contents is stably on disk

# File descriptor numbers

- **File descriptors are inherited by processes**
  - When one process spawns another, same fds by default
- **Descriptors 0, 1, and 2 have special meaning**
  - 0 – “standard input” (stdin in ANSI C)
  - 1 – “standard output” (stdout, printf in ANSI C)
  - 2 – “standard error” (stderr, perror in ANSI C)
  - Normally all three attached to terminal
- **Example:** type .c

# The rename system call

- `int rename(const char *from, const char *to);`
  - Changes name to to reference file from
  - Removes file name from
  - Returns 0 on success, -1 on error
- **Guarantees that to will exist despite any crashes**
  - to may still be old file
  - from and to may both be new file
  - but to will always be old or new file
- **fsync/rename idiom used extensively**
  - E.g., emacs: Writes file `.#file#`
  - Calls `fsync` on file descriptor
  - `rename (".#file#", "file");`

# Creating processes

- `int fork (void);`
  - Create new process that is exact copy of current one
  - Returns *process ID* of new proc. in “parent”
  - Returns 0 in “child”
- `int waitpid (int pid, int *stat, int opt);`
  - `pid` – process to wait for, or -1 for any
  - `stat` – will contain exit value, or signal
  - `opt` – usually 0 or `WNOHANG`
  - Returns process ID or -1 on error

# Deleting processes

- `void exit (int status);`
  - Current process ceases to exist
  - `status` shows up in `waitpid` (shifted)
  - By convention, `status` of 0 is success, non-zero error
- `int kill (int pid, int sig);`
  - Sends signal `sig` to process `pid`
  - `SIGTERM` most common value, kills process by default (but application can catch it for “cleanup”)
  - `SIGKILL` stronger, kills process always

# Running programs

- `int execve (char *prog, char **argv, char **envp);`
  - prog – full pathname of program to run
  - argv – argument vector that gets passed to main
  - envp – environment variables, e.g., PATH, HOME
- **Generally called through a wrapper functions**
- `int execlp (char *prog, char **argv);`
  - Search PATH for prog
  - Use current environment
- `int execlp (char *prog, char *arg, ...);`
  - List arguments one at a time, finish with NULL
- **Example:** `minish.c`

# Manipulating file descriptors

- `int dup2 (int oldfd, int newfd);`
  - Closes `newfd`, if it was a valid descriptor
  - Makes `newfd` an exact copy of `oldfd`
  - Two file descriptors will share same offset (1seek on one will affect both)
- `int fcntl (int fd, F_SETFD, int val)`
  - Sets *close on exec* flag if `val = 1`, clears if `val = 0`
  - Makes file descriptor non-inheritable by spawned programs
- **Example:** `redirsh.c`

# Pipes

- `int pipe (int fds [2] );`
  - Returns two file descriptors in `fds [0]` and `fds [1]`
  - Writes to `fds [1]` will be read on `fds [0]`
  - When last copy of `fds [1]` closed, `fds [0]` will return EOF
  - Returns 0 on success, -1 on error
- **Operations on pipes**
  - `read/write/close` – as with files
  - When `fds [1]` closed, `read (fds [0] )` returns 0 bytes
  - When `fds [0]` closed, `write (fds [1] )`:
    - Kills process with SIGPIPE, or if blocked
    - Fails with EPIPE
- **Example:** `pipesh.c`