

CSCE 351

Operating System Kernels

The UNIX/MINIX Operating System Processes & IPC

Steve Goddard
goddard@cse.unl.edu

<http://www.cse.unl.edu/~goddard/Courses/CSCE351>

1

UNIX History

- ◆ The original UNIX
 - » An experimental operating system developed by Ken Thompson & Dennis Ritchie in the late '60s
- ◆ Main variants (“Standards”)
 - » System V (1983)
 - ◆ developed by AT&T
 - » 4.4 BSD (1993)
 - ◆ Open Software Foundation
- ◆ Commercial products
 - » Ultrix, DEC UNIX — DEC
 - » SunOS, Solaris — Sun
 - » HP/UX — Hewlett Packard
 - » POSIX
 - ◆ IEEE/ISO
 - » FreeBSD
 - » Linux
 - » AIX — IBM
 - » Xenix — Microsoft
 - » ...

2

MINIX History

- ◆ The original MINIX was developed by Andrew Tanenbaum after AT&T closed the source code (1978)
 - » An educational operating system
 - » Version 1.0 was released circa 1987
 - » Version 2.0 was released circa 1997
- ◆ Version 2.0
 - » POSIX compliant
 - » Fully documented
 - » Requires only 30MB
 - » Micro-kernel design
- ◆ “Looks like UNIX from the outside”

3


Processes

- ◆ A process is created by the `fork()` system call
 - » creates a new address space that is a duplicate of the callers

```
main (argc, argv)
int childpid;
{
  switch (childpid = fork()) {
    case 0:          /* child */
      child_func();
      exit(0);
    default:        /* parent */
      parent_func();
      while(wait((int *) 0) != childpid);
      exit(0);
    case -1:        /* oops */
      error("fork:%s\n", sys_errlist[errno]);
  }
}
```

Parent address space

childpid = 1




Child address space

childpid = 0


4

Processes

- ◆ Alternatively, processes can be “created” by an `execve()`
 - » replaces the memory image of the caller with a new program




- ◆ This is how the shell executes commands
 - » a `fork()` followed by an `exec()`



5

Example: How users logs in




- ◆ There exists a “master process” in UNIX: `init`
- ◆ `init` forks a process for each terminal port
- ◆ each `init` copy execs `getty` which prints the login prompt and then reads the login and password
- ◆ `getty` then execs `login` which verifies the login
- ◆ `login` then execs `csH` which forks new processes for each command

6


(Simple) Interprocess Communication

- ◆ Like message passing except more general
- ◆ Pipes — a shared, in-memory file
 - » a queue of 4K bytes
 - » buffered, asynchronous message passing
 - ❖ blocks reader when queue is empty
 - ❖ blocks writer when queue is full



7

(Simple) Interprocess Communication

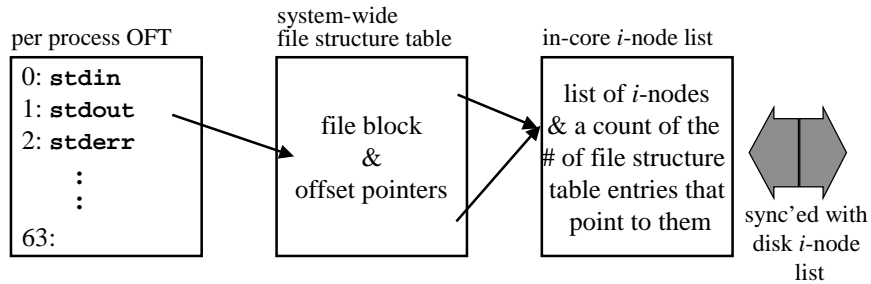


```
main() {
  int pipe1[2], pipe2[2];
  if (pipe(pipe1) == -1 || pipe(pipe2) == -1) error(...
  switch (childpid = fork()) {
    case 0:
      close(pipe1[1]);          /* write descriptor for pipe1 */
      close(pipe2[0]);          /* read descriptor for pipe2 */
      client(pipe1[0],pipe2[1]); /* client program */
    default :
      close(pipe1[0]);          /* read descriptor for pipe1 */
      close(pipe2[1]);          /* write descriptor for pipe2 */
      server(pipe2[0],pipe1[1]); /* server program */
      while (wait((int *) 0) != childpid); /* wait for child */
  }
}
```

8

Open File Table

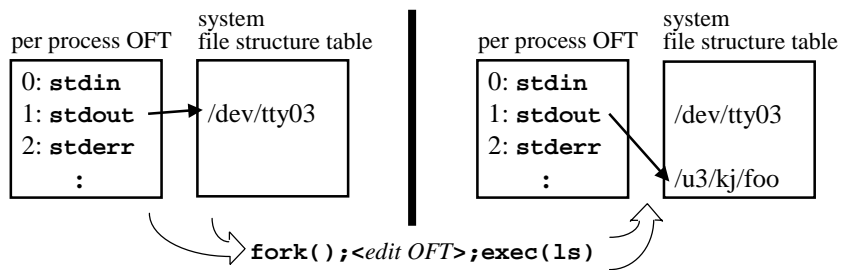
- ◆ UNIX maintains an *open file table* for each process which lists each file in use by the process
 - » the OFT is copied when processes are forked



The UNIX File System

Open file table examples

- ◆ I/O redirection — `ls > foo`
 - » just change a pointer in the OFT



(Primary) MINIX Kernel Interprocess Communication

- ◆ Message Passing is used by Kernel tasks

- » Blocking send and receive primitives

- ❖ Syntax:

- `send(dest, &message)`

- `receive(dest, &message)`

- `sendrec(src_dest, &message)`

- Note: page 97 shows this as `send_rec(src_dest, &message)`

- but page 560, file `include/minix/syslib.h` shows it as `sendrec ()`