

*Concurrent Architectures:
UNIX Part #1*

Alistair Rendell

UNIX

- Origins – The Multics Project
 - MIT and GE
 - Ken Thomson and Dennis Ritchie
 - The C language
 - PDP 11
- Features
 - Hierarchical file system
 - Compatible I/O (file, device & inter-process)
 - Command language choice
 - Portability
- Decendents
 - Version 7, BSD 4.2, SystemV, Solaris, XINU, Minix, Linux

The Single UNIX[®] Specification History & Timeline

1969	The Beginning	The history of UNIX starts back in 1969, when Ken Thompson, Dennis Ritchie and others started working on the "little-used PDP-7 in a corner" at Bell Labs and what was to become UNIX.
1971	First Edition	It had an assembler for a PDP-11/20, file system, fork(), roff and ed. It was used for text processing of patent documents.
1972	First UNIX Installations	The first installations had 3 users, no memory protection, and a 500 KB disk.
1973	Fourth Edition	It was rewritten in C. This made it portable and changed the history of OS's.
1975	Sixth Edition	UNIX leaves home. Also widely known as Version 6, this is the first to be widely available outside of Bell Labs. The first BSD version (1.x) was derived from V6.
1979	Seventh Edition	It was an "improvement over all preceding and following Unices" [Bourne]. It had C, UUCP and the Bourne shell. It was ported to the VAX and the kernel was more than 40 Kilobytes (K).
1980	Xenix	Microsoft introduces Xenix. 32V and 4BSD introduced.
1982	System III	AT&T's UNIX System Group (USG) release System III, the first public release outside Bell Laboratories. SunOS 1.0 ships. HP-UX introduced. Ultrix-11 introduced.
1983	System V	Computer Research Group, UNIX System Group (USG) and a third group merge to become UNIX System Development Lab. AT&T announces UNIX System V, the first supported release. Installed base 45,000.
1984	4.2BSD	University of California at Berkeley releases 4.2BSD, includes TCP/IP, new signals and much more. X/Open formed.
1984	SVR2	System V Release 2 introduced. At this time there are 100,000 UNIX installations around the world.
1986	4.3BSD	4.3BSD released, including internet name server. SVID introduced. NFS shipped. AIX announced. Installed base 250,000.
1987	SVR3	System V Release 3 including STREAMS, TLI, RFS. At this time there are 750,000 UNIX installations around the world. IRIX introduced.
1988		POSIX.1 published. Open Software Foundation (OSF) and UNIX International (UI) formed. Ultrix 4.2 ships.
1989		AT&T UNIX Software Operation formed in preparation for spinoff of UNIX development group. Motif 1.0 ships.
1989	SVR4	UNIX System V Release 4 ships, unifying System V, BSD and Xenix. Installed base 1.2 million.
1990	XPG3	X/Open launches XPG3 Brand. OSF/1 debuts. Plan 9 from Bell Labs ships.

"The Single UNIX Specification brings all the benefits of a single standard operating system, namely application and information portability, scalability, flexibility and freedom of choice for customers"

Allen Brown, President and CEO, The Open Group

The Story of the License Plate...

In 1983 Digital Equipment Corporation (DEC) was ramping up their engineering group to create and ship their first UNIX system product. One of the stalwarts of the group was Armando P. Stetter. Armando was a UNIX system devotee. He lived and he breathed the UNIX system. When he got his new car, it was natural that he got vanity license plates that said "UNIX" on them. And it only made it better that the state motto of New Hampshire was "Live Free or Die". Armando often got requests from people along the lines of "When will we be able to get our UNIX system license directly from DEC?" And Armando kept saying "Real Soon Now" (RSN). Armando was going to a conference and he was dreading having to say "RSN" many more times, so he had a bright idea. Armando went prepared to give out "UNIX licenses". On stage, when the question came up, "When will we be able to get our UNIX system license directly from DEC?", Armando yelled "Right Now!" and produced incriminating of his license plate, holding it up for all to see. It was an almost perfect likeness of his license plate, with the trademark "UNIX" in



the middle of it, but instead of having "Live Free or Die" across the bottom of the plate (as in the real case), it had it across the top. Across the bottom was the trademark acknowledgement.

DEC made UNIX license plates up in small numbers and handed them out at events. They usually ran out. The demand for the license plates never did abate. People saw them on an office wall, or heard about them somewhere, and wanted one of their own.

Armando left the state for the sunny climes of California, and had taken his car and license plate with him. Or so many people thought.

In 1989 Jon "maddog" Hall was purchasing a new car, a Jeep Wrangler. And of course the license plate had to be relevant. So Jon, a long time DEC employee and UNIX system guru, submitted his application with many variations and the clerk said "I think we can give you your first choice..." and gave him the temporary paper plates (to be used on the car until the metal plates were manufactured) with "UNIX" on them. And so it has been ever since. Jon's Jeep has been the holder of the UNIX license plate.

The Open Group thanks Jon "maddog" Hall for sharing the story of the UNIX license plate.

The UNIX Brand

The UNIX Brand is used to identify products that have been certified as conforming to the Single UNIX Specification, initially UNIX 93, followed subsequently by UNIX 95, UNIX 98 and now UNIX 03.

1991		UNIX System Laboratories (USL) becomes a company - majority-owned by AT&T. Linus Torvalds commences Linux development. Solaris 1.0 debuts.
1992	SVR4.2	USL releases UNIX System V Release 4.2 (Destiny). October - XPG4 Brand launched by X/Open. December 22nd - Novell announces intent to acquire USL. Solaris 2.0 and HP-UX 9.0 ship.
1993	4.4BSD	4.4BSD the final release from Berkeley. June 16 - Novell acquires USL.
Late 1993	SVR4.2MP	Novell decides to get out of the UNIX business. Rather than sell the business as a single entity, Novell transfers the rights to the UNIX trademark and the specification to X/Open Company. COSE Initiative delivers "Spec 1170" to X/Open for fast-track. In December Novell ships SVR4.2MP, the final USL OEM release of System V.
1994	Single UNIX Specification	BSD 4.4-Lite eliminated all code claimed to infringe on USL/Novell. As the owner of the UNIX trademark, X/Open introduces the Single UNIX Specification (formerly Spec 1170) which separates the UNIX trademark from any actual code stream itself, thus allowing multiple implementations.
1995	UNIX 95	X/Open introduces the UNIX 95 branding program for implementations of the Single UNIX Specification. Novell sells UnixWare business to SCO. Digital UNIX introduced. UnixWare 2.0 ships. OpenServer 5.0 debuts.
1996		The Open Group forms as a merger of the Open Software Foundation (OSF) and X/Open. UnixWare 2.1, HP-UX 10.20 and IRIX 6.2 ship.
1997	Single UNIX Specification, Version 2	The Open Group introduces Version 2 of the Single UNIX Specification, including support for real-time, threads and 64-bit and larger processors. The specification is made freely available on the web. IRIX 6.4, AIX 4.3 and HP-UX 11 ship.
1998	UNIX 98	The Open Group introduces the UNIX 98 family of brands, including Base, Workstation and Server. First UNIX 98 registered products shipped by Sun, IBM and NCR. The Open Source movement starts to take off with announcements from Netscape and IBM. UnixWare 7 and IRIX 6.5 ship.
1999	UNIX at 30	The UNIX system reaches thirty. Solaris 7 ships. Linux 2.2 kernel released. The Open Group and the IEEE commence joint development of a revision to POSIX and the Single UNIX Specification. First LinuxWorld conferences. Dot com fever on the stock markets. Tru64 UNIX ships.
2001	Single UNIX Specification, Version 3	Version 3 of the Single UNIX Specification unites IEEE POSIX, The Open Group and the industry efforts. Linux 2.4 kernel released. The value of procurements of open systems referencing the UNIX brand exceeds \$55 billion. AIX 5L ships.
2003	ISO/IEC 9945	The core volumes of Version 3 of the Single UNIX Specification are approved as an international standard. "Westwood" test suites shipped for UNIX 03 brand. Solaris 9.0 E ships. Linux 2.6 kernel released.

UNIX and the UNIX logo are registered trademarks of The Open Group in the United States and other countries. Linux is a registered trademark of Linus Torvalds. All other trademarks are the property of their respective owners. Copyright © 2003 The Open Group. All rights reserved.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Unix Process Reproduction

- Unix process creation is by cloning
- Standard Unix C library contains function `fork`
- Typical usage:

```
pid = fork();
```
- `fork` spawns a duplicate
 - same memory, regs, PC
- `fork` returns
 - 0 to child
 - and to the parent
 - PID (of child) if it works
 - -1 if it fails
 - never 0!

Unix Processes: Overwriting

- Maybe another program is more suitable?
 - Passing responsibility
 - Many computations have quite different phases
 - A child process may need to do a very different job
- Use exec family

```
exec(filename, args);
```
- Does following
 - Memory is replaced by the image given in the executable file
 - Registers initialized in the usual way for a program
 - Parameters are passed as if the new program was called from the shell
 - Only returns in case of a failure

Exec Example

```
if (Fork()==0) {
    execl("/usr/bin/grep",
          arg2("grep", keyword));
    abort();
} else {
    <rest of PARENT program>
}
```

- Questions

- what is the child?
- what is concurrent?
- why does grep appear twice?
- why might it abort?

Process Requiem!

- Co-ordination can be a problem
 - A child was spawned to do something in parallel (concurrently)
 - At some point the parent must be sure it is complete
- The `wait` primitive solves it
 - waits for termination of a (any) child
 - returns the PID of the child that has died
 - argument gives exit code for the child

Example of Wait

```
if (Fork() == 0) {  
    <Child's task>  
    exit(0);  
} else {  
    <Parent's task>  
    wait (Codes);  
}
```

- Concurrency?

Unix Processes Dozing

- Why wait for time to pass?
 - It may want to do something periodically
 - It can make sure other processes get a chance
- The `sleep` primitive solves it

```
int sleep(int secs);
```

- The supplied argument is the duration (in seconds)
- May return early if woken up (by a signal)
- Results indicate sleep-time remaining
- Use of `sleep(0)`

UNIX Shells

- Job control language
 - Pre-Unix JCLs were proprietary and idiosyncratic
- Not part of Operating System
 - But must reflect its facilities
 - Can have a shell per user
- There are many common ones
 - sh – Bourne shell
 - csh – C shell
 - ksh – Korn shell
 - tcsh – Extended C-shell
 - bash – Bourne-again shell

UNIX Shells (2)

- Features of the Shell
 - A virtual machine
 - Has a language
 - can invoke programs
 - types, variables, control
 - Can write programs
- Access to Unix's hierarchical file system supported by `cd`, `pwd`, `ls`, etc
- Uniform I/O is supported by re-direction of `stdin` and `stdout`

```
ls -als > DIR
```

```
ls -als > /dev/null
```

```
ls -als | grep def
```

UNIX Shells (3)

- Processes supported by shells

```
ps  
emacs assignment1.adb &  
ps -aux | grep u9876543
```

- Can write executable programs in shell code
 - Usual termed a script
- Interpreted, not compiled
- First line of file indicates the appropriate interpreter

C-Shell Scripts (Example)

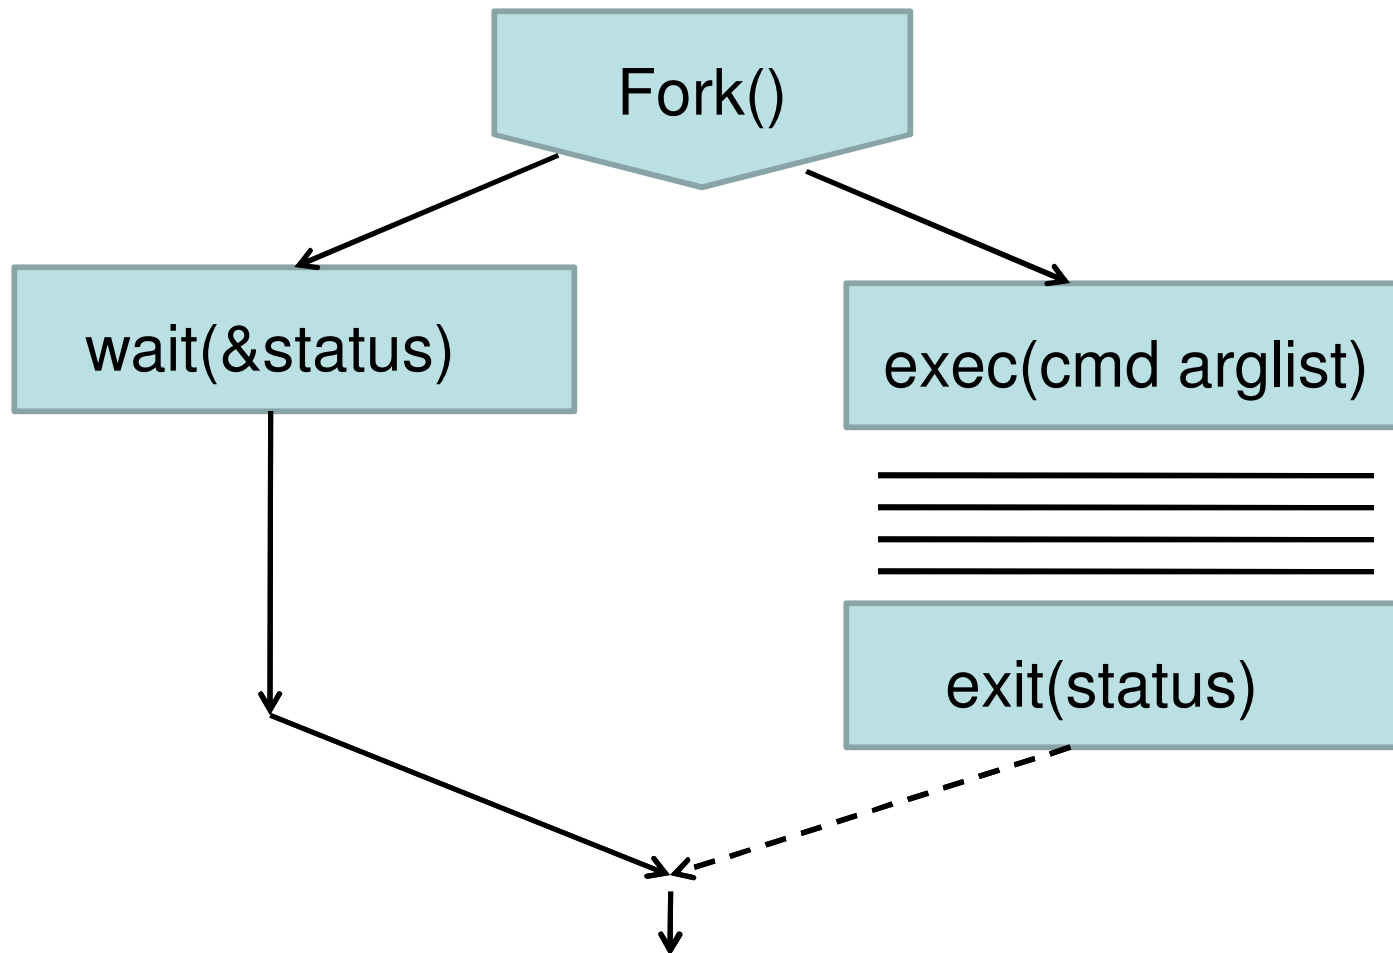
- Example of shell script

```
#!/bin/csh
set filist = 'cat $1'
foreach nam ( $filist )
    if (-e $nam) then
        cp $nam $2/$nam.bak
    else
        echo $nam "does not exist"
    endif
end
```

- What does it do?
- What does \$1 and \$2 refer to?

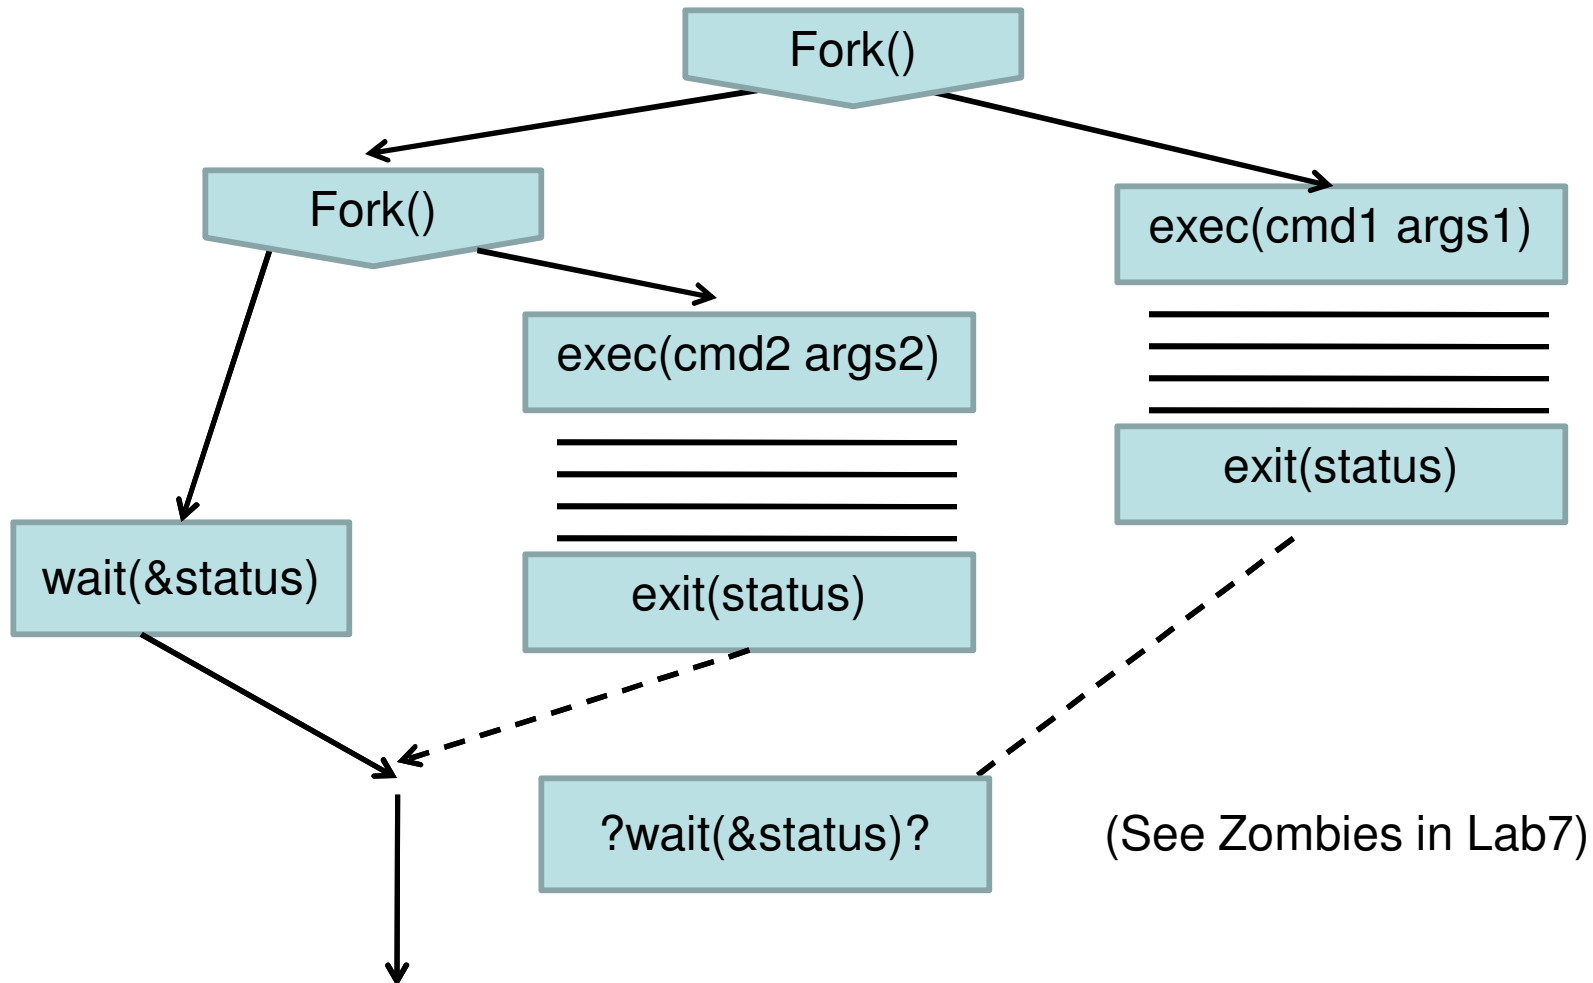
Command-Line Interpretation

- `cmd arglist`



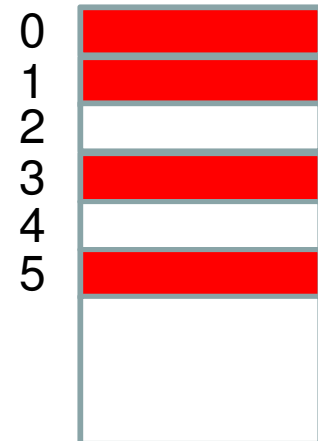
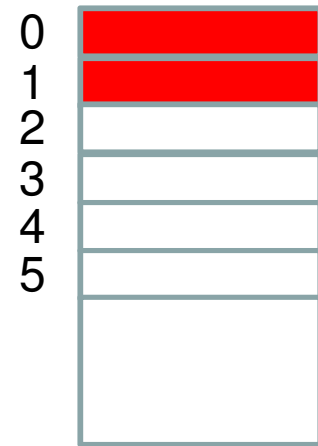
Backgrounding a Command

- `cmd1 args1 & cmd2 args2`



UNIX File Descriptors

- What are they?
 - Non-negative integers denoting an open I/O channel
 - They index the File Descriptor Table (FDT) associated with each current process
 - The indexes (0,1,2,...etc.) are file descriptors
 - The corresponding table entries mean something to the operating system!



File Descriptors

- Three channels are standard

<code>stdin</code>	file descriptor 0
<code>stdout</code>	file descriptor 1
<code>stderr</code>	file descriptor 2

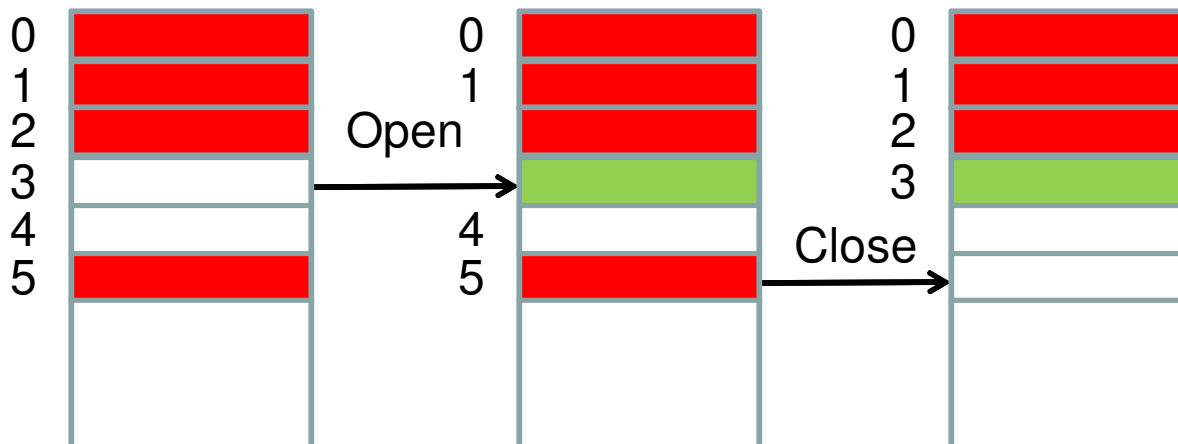
- **Initially**

<code>stdin</code>	keyboard
<code>stdout</code>	current xterm window
<code>stderr</code>	current xterm window

- FDT is preserved across `exec`
- FDT copied during `fork`

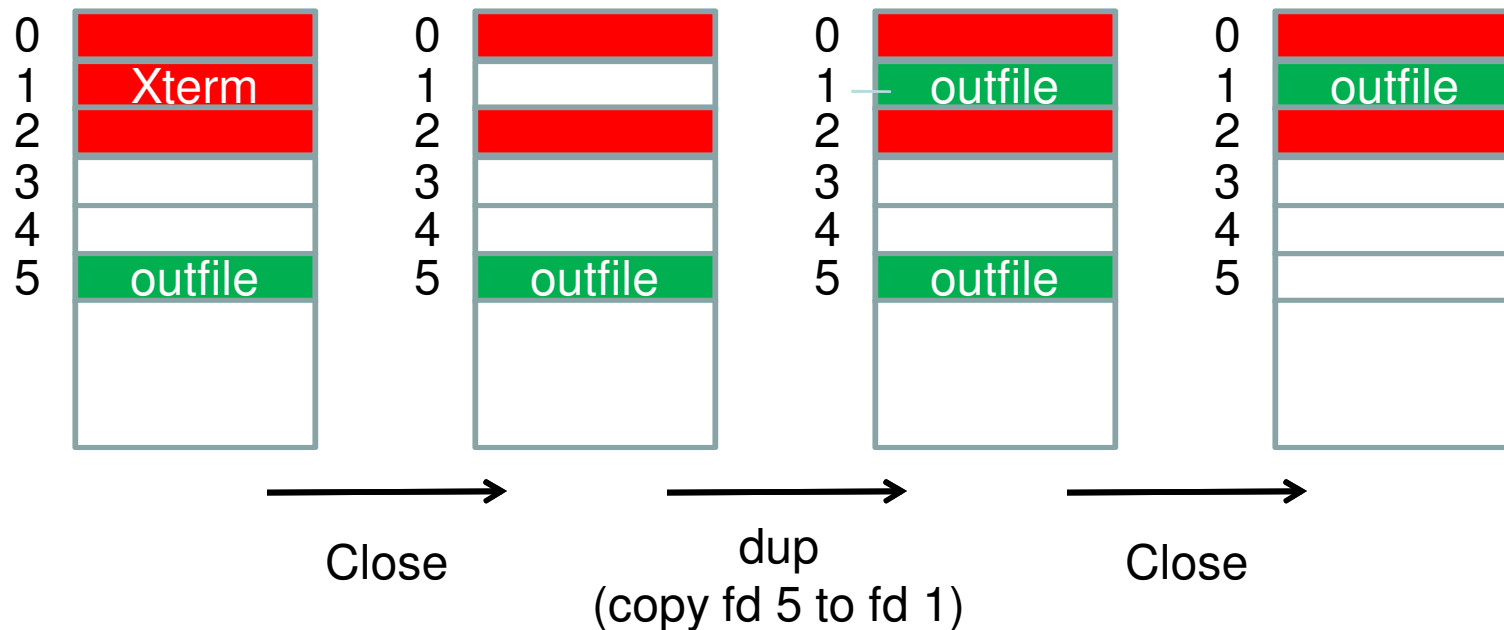
Opening and Closing Files

- Opening a file
 - first free entry in the FDT is chosen
- Closing a file
 - Entry in FDT is marked free
 - Opposite to open!



Redirecting I/O

- Redirecting for example stdout
 - User cannot change FDT directly
 - There is a primitive `dup` for this purpose



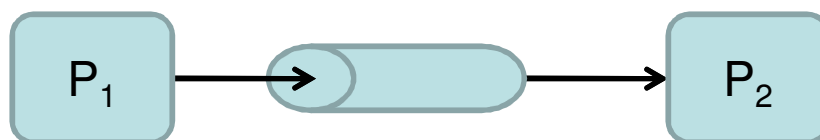
UNIX Pipes

- Consider

```
ps -aux | grep 666
```

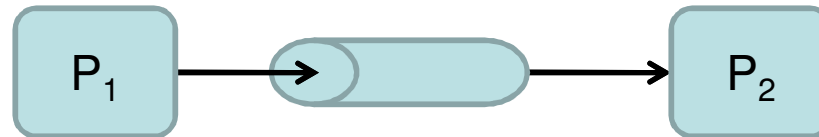
- `ps` writes to `stdout`
- `grep` reads from `stdin`
- I/O has been re-directed by the shell
- Shell doesn't handle the characters

- Mental model of a pipe



- It's a buffer for characters
- Handled by I/O sub-system of Unix
- May only be one-way (Solaris has two-way)

UNIX Pipes Dependencies



- P₂ will wait for production
- P₁ will not wait for consumption
- There may also be buffering in P₁
- The importance of timely flushing!
 - what's this mean?
- So what sort of communication primitive is a pipe?

Running a Shell as a Subprocess

- Can have a shell to do a job
 - Unix allows easy creation of subprocesses
 - A subprocess can run a shell script

- An easy-to-code method

```
    csh (cmd)
```

- `cmd` is shell code
- A process is created to run `cmd`
- Result is an input stream on which output of `cmd` is delivered

Creating UNIX Pipes

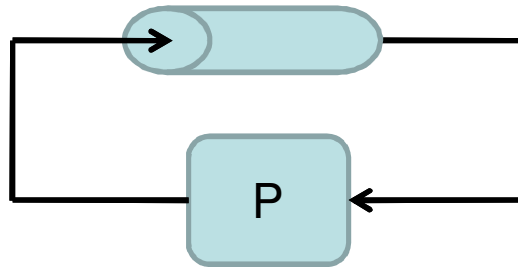
- Opening a pipe

```
void pipe (int pipefd[2])
```

- `pipefd[0]` set to file descriptor of input end of pipe

- `pipefd[1]` set to file descriptor of output end of pipe

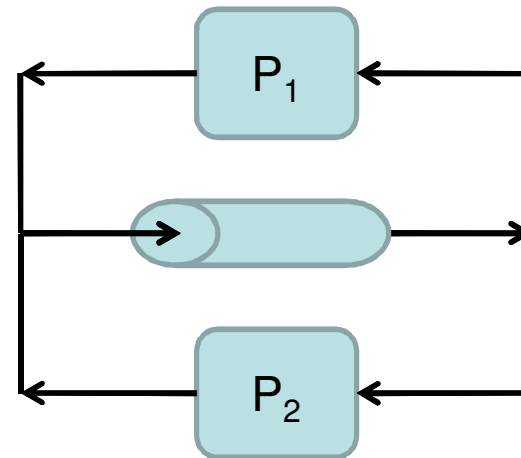
- The result



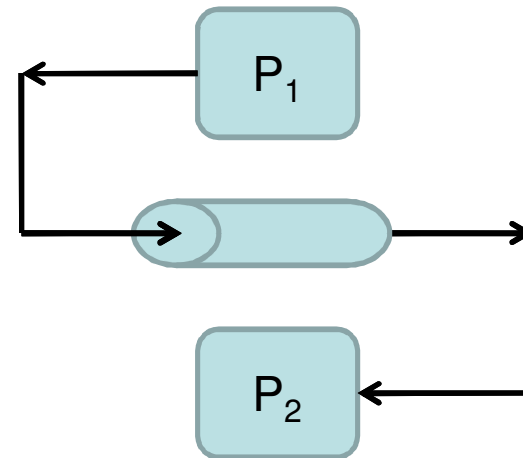
does not connect two processes!

Connecting Two Processes by a Pipe

- Fork the process, but do it after opening a pipe

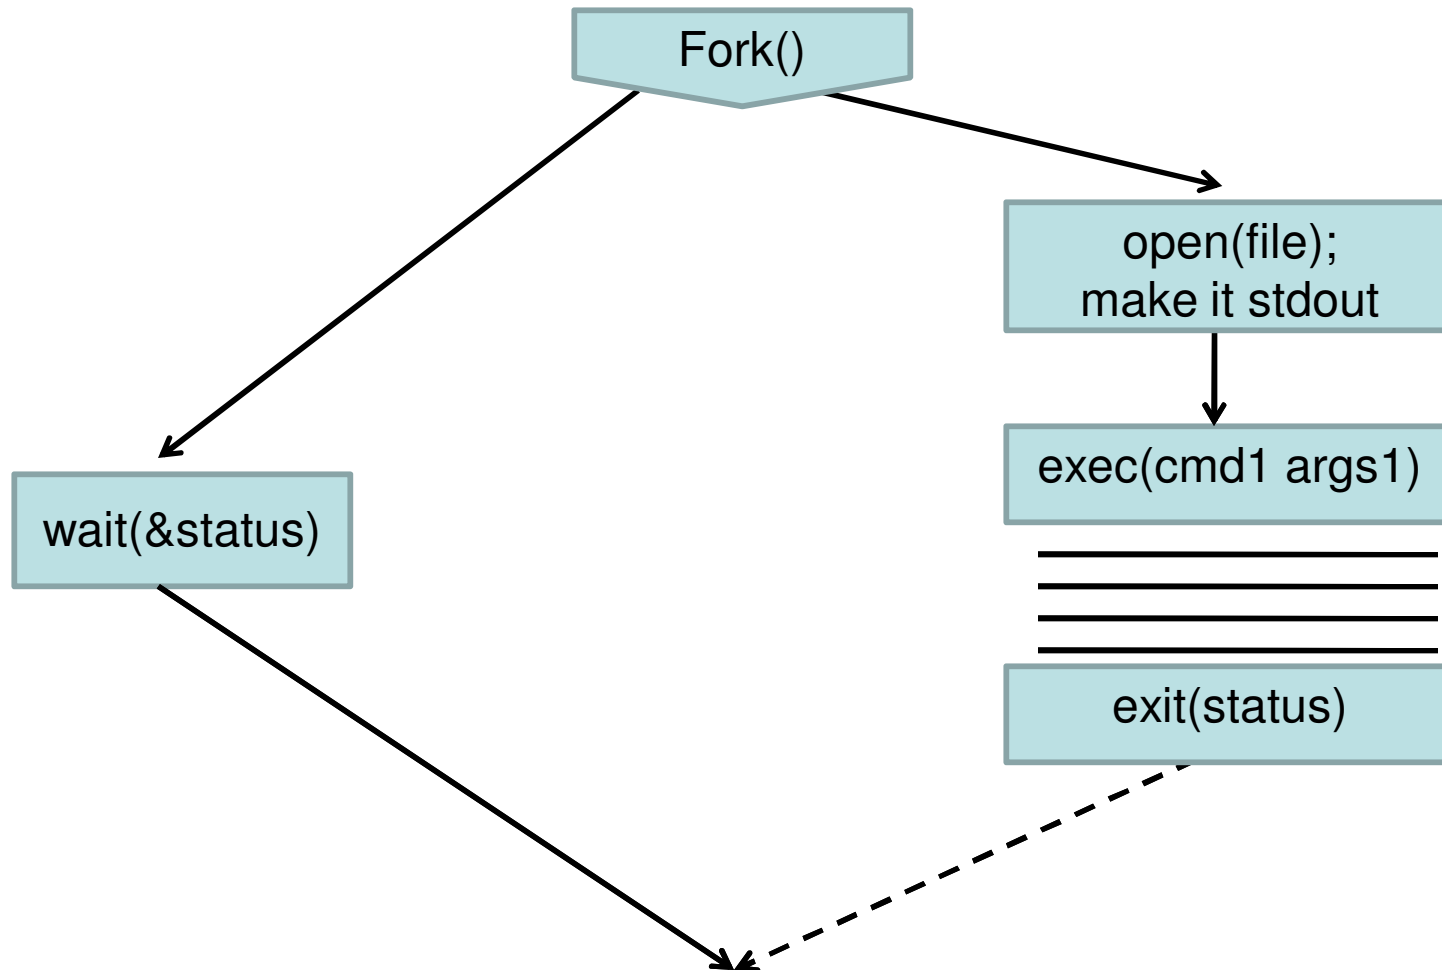


- Close unwanted descriptors



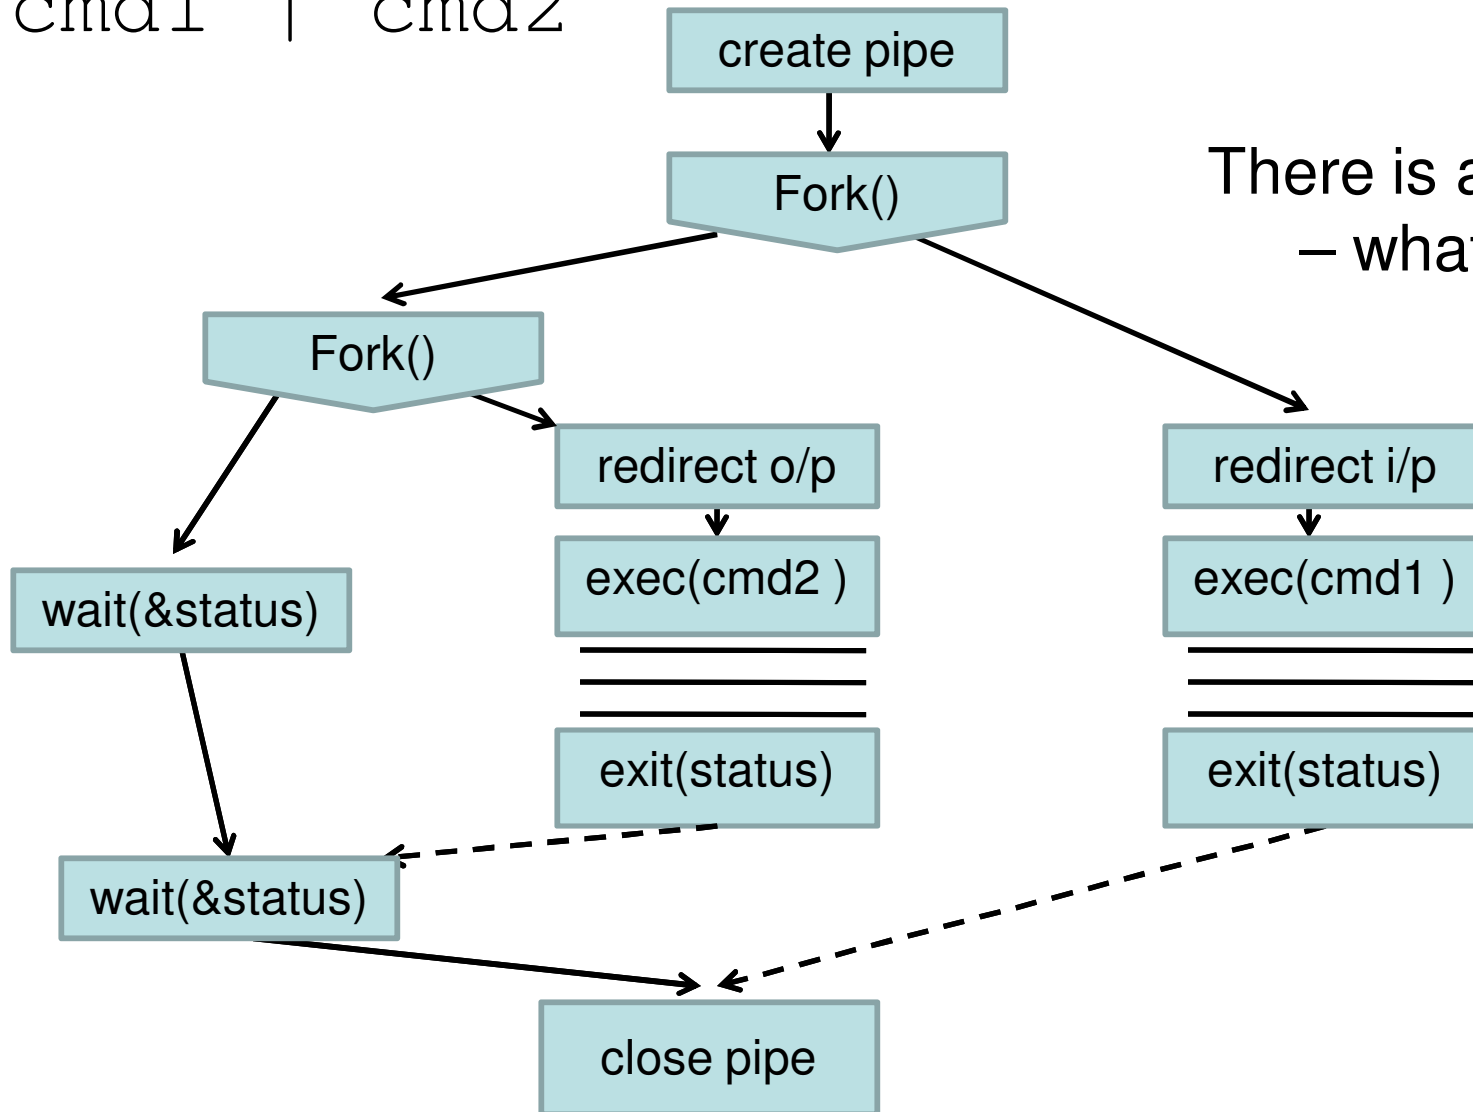
Command line I/O Redirection

- `cmd1 args1 > file`



Command Line Pipes

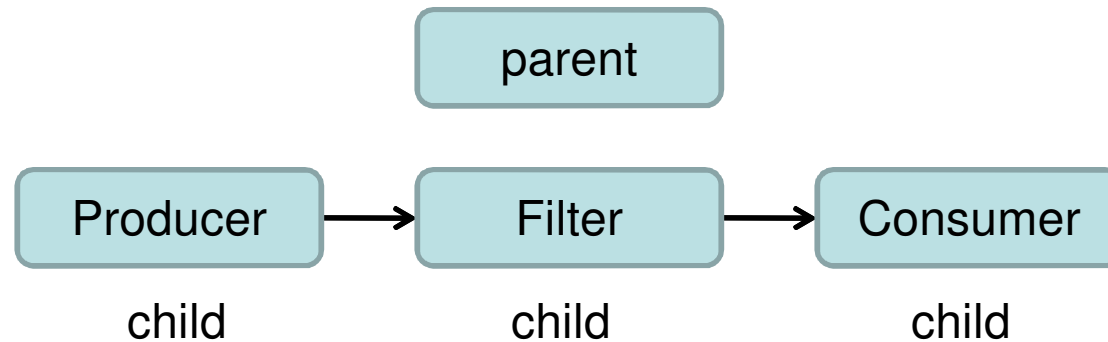
- `cmd1 | cmd2`



There is a mistake
– what is it?

Subprocess Pipelines

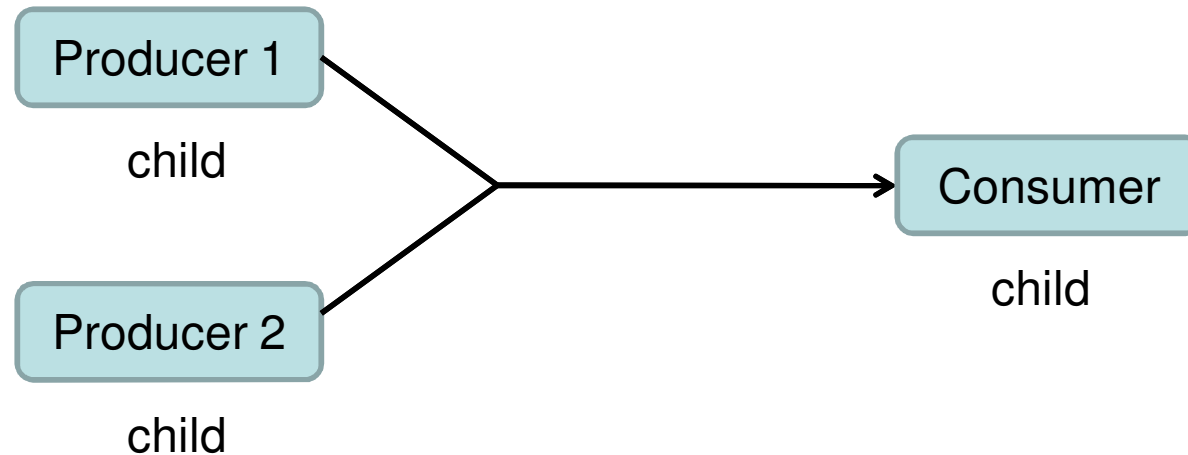
- Desired state



- make pipes (before forking)
- fork two times
- close pipes in parent
- each child closes descriptors that it won't use
- each child does an exec if appropriate

Sharing Pipes

- Pipes can be shared



- Integrity of messages
 - Pipe is just a stream of chars
 - No interleaving BUT care required with buffering
 - Messages less than 2k

Recall Slide from CDS-L3

```
int data_pipe [2], c, rc;

if (pipe (data_pipe)==-1){
    perror ("no pipe");exit(1);
}

if (fork () == 0) {
    close (data_pipe [1]);
    while ((rc = read
        (data_pipe [0],&c,1))>0){
        putchar (c);
    }
    if (rc == -1) {
        perror ("pipe broken");
        close (data_pipe [0]);
        exit (1);
    }
    close (data_pipe [0]);
    exit (0);
} else {

    close (data_pipe [0]);
    while ((c = getchar ()) > 0) {
        if (write
            (data_pipe[1],&c,1)==-1){
            perror ("pipe broken");
            close (data_pipe [1]);
            exit (1);
        }
    }
    close (data_pipe [1]);
    pid = wait ();
}
```