

CS503A

Lecture #3

2014-04-08

Administrivia

Our Story So Far (with a bit of new stuff to keep you awake)

- operating systems operate at lower levels of abstraction
 - closer to the metal
 - performance is critical

- kernel operates in "supervisor" mode

- Context switching
- application programs make system calls for services
 - timer interrupts processing so kernel has chance to do housekeeping tasks & schedule other processes (time slice)
 - context switch to kernel for hardware interrupt
source hardware interrupts

- init (process #1): system startup
 - secondary function: reap zombie orphans

- kernel concerns

- process management
- file systems
- device drivers
- networking
- memory management
- accounting
- protection/security
- error detection

OSSF (cont.)

- process: (essentially) running program

- executable file

- argv

- pre-opened files

no intrinsic meaning to system. means
whatever the receiving program wants it to
mean - but: conventions

exit status

0 - Success/true
non-zero - fail/false

- convention:

0

Standard input

cin

1

Standard output

cout

2

Standard error

cerr

most programs expect
these files to be pre-opened
languages bound this to their
formatted/bulked I/O library
e.g. cin/cout/cerr

- C/C++ provide standard libraries for formatting & buffering.

- actual I/O operations performed by

read/write system calls

OSSF (cont.)

- system call/trap

- typically: arguments stored in registers
(ints & pointers)

result in register: int

-ve: fail

0: ok, if no other meaningful result

≥ 0 : result (e.g. fd, pid)

- global variable `errno` holds integer error code

- `perror`: library function

} more info

OSSF (cont.)

Unix/Linux/Posix: separate process creation from program execution

pid_t pid = fork()

→ clones process

< 0: fail

= 0: child process

> 0: parent process (return result is pid of child)

- child process inherits same open files as parent
- parent collects exit status of child: wait/waitpid

⇒ process free (new)

- init ^{inherits} options

- fork is useful for distributing work
 - assignment 1

- But what if you want to run a different program?
- stay tuned

* process free

OSSF (cont.)

inter-process communication

- pipe: original IPC mechanism

- producer - consumer

- producer writes data to memory buffer in kernel

- consumer reads data from buffer

Both processes operate concurrently

- when buffer is full, writing is blocked

- when buffer is empty, reading is blocked*

* unless non-blocking I/O flag is set

- pipe is unidirectional

- for bidirectional communication, use 2 pairs of endpoints

- but watch out for deadlocks (nm)

- pipe instance is file descriptor (small integer)

- use same read/write system calls

- file descriptor is object (in C++ sense)

- but implemented in C

- uses table at function pointer, just like

C++ vtbl

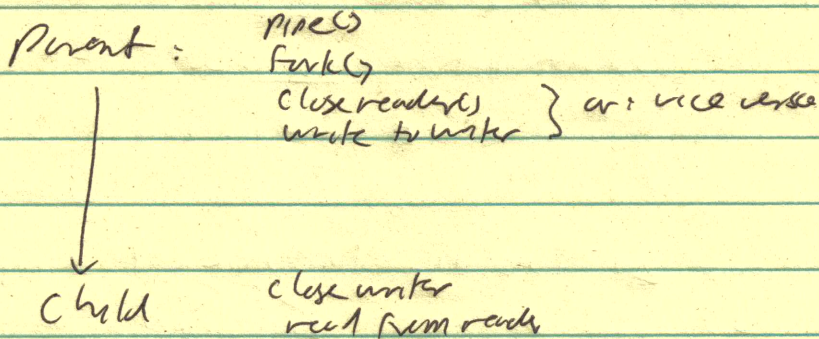
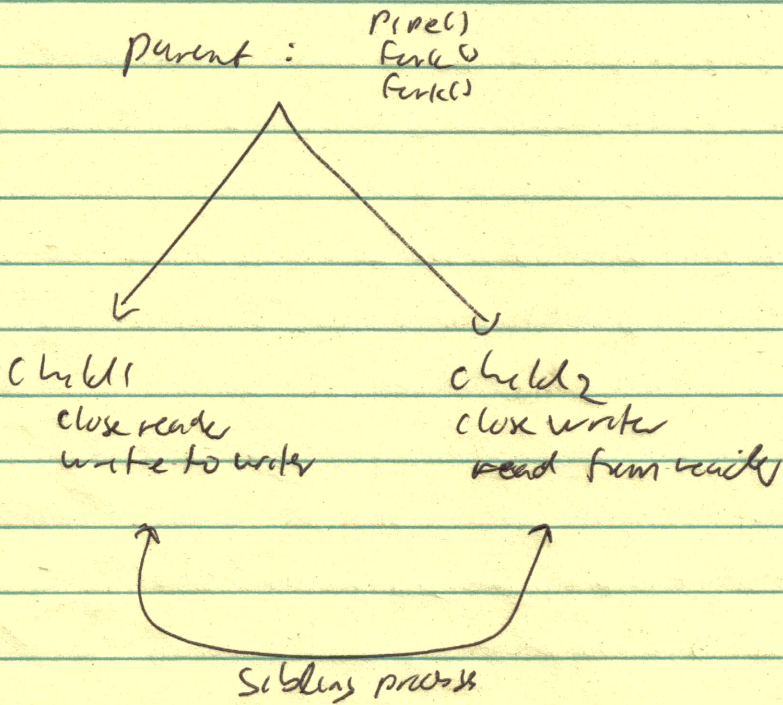
- from a programming POV, pipe behaves just like any ordinary file

- System call to set non-blocking I/O flag

Pipes (cont.)

- `pipe` system call returns pair of open file descriptors for producer (writer) & consumer (reader)
 - each fd is one-way (endpoints)
- single process can't really use pipe descriptors
 - read empty pipe would block, then you're stuck (similarly, write until full, then blocked)
 - besides, why would you? *why would you do such a thing*
- file descriptor is index into kernel data structure
 - table is per process (each process has fd 0, 1, 2, ...)
 - ⇒ so how do you get a pipe across two processes?
- files stay open across `fork`
 - two processes that share a pipe must have a common ancestor that created the pipe
- typically, one process closes the reader & keeps the writer while the other process does the opposite

Pipes (cont.)



Pipes (cont.)

- Shell syntax: do-this | do-that
 - pipe symbol (vertical bar)
 - also used for or operator
- Shell only supports pipelines with **FD 0/1**
standard input / standard output
(can redirect stdout to stdin - shell magic)
- pipe system call allows more complex
unix patterns (not stdin/stdout)
⇒ subject of assignment 7.1

"Everything Is a File"

- basic file ops: read byte(s), write byte(s)
 - => that's a good abstraction for a lot of stuff
 - ordinary file
 - tty (keyboard/screen) - dumb terminal
 - pipe (original IPC mechanism)
 - network connections
 - speakers, microphone, camera
 - blinking lights
 - pseudo devices (dev/null dev/zero dev/random)
 - other devices
- stdin is not keyboard } default for
std is not screen } interactive programs
- shell syntax for redirecting input/output
 - design principle: policy is mechanism
 - > >> < |

Process Management (cont.)

- fork(2) allows a program to divide work across multiple processes, each running the same program
 - pipe is just one IPC mechanism
 - we'll look at others later in the quarter

* subject of assignment 1

- but what if you want to run a different program?

⇒ execve (plus several standard library variants that let you marshal arguments differently — all eventually call execve)

- exec* replaces current running program in process with new program, new argv

* same process, new program

- if call to exec is successful, it does not return (you can't go home again)

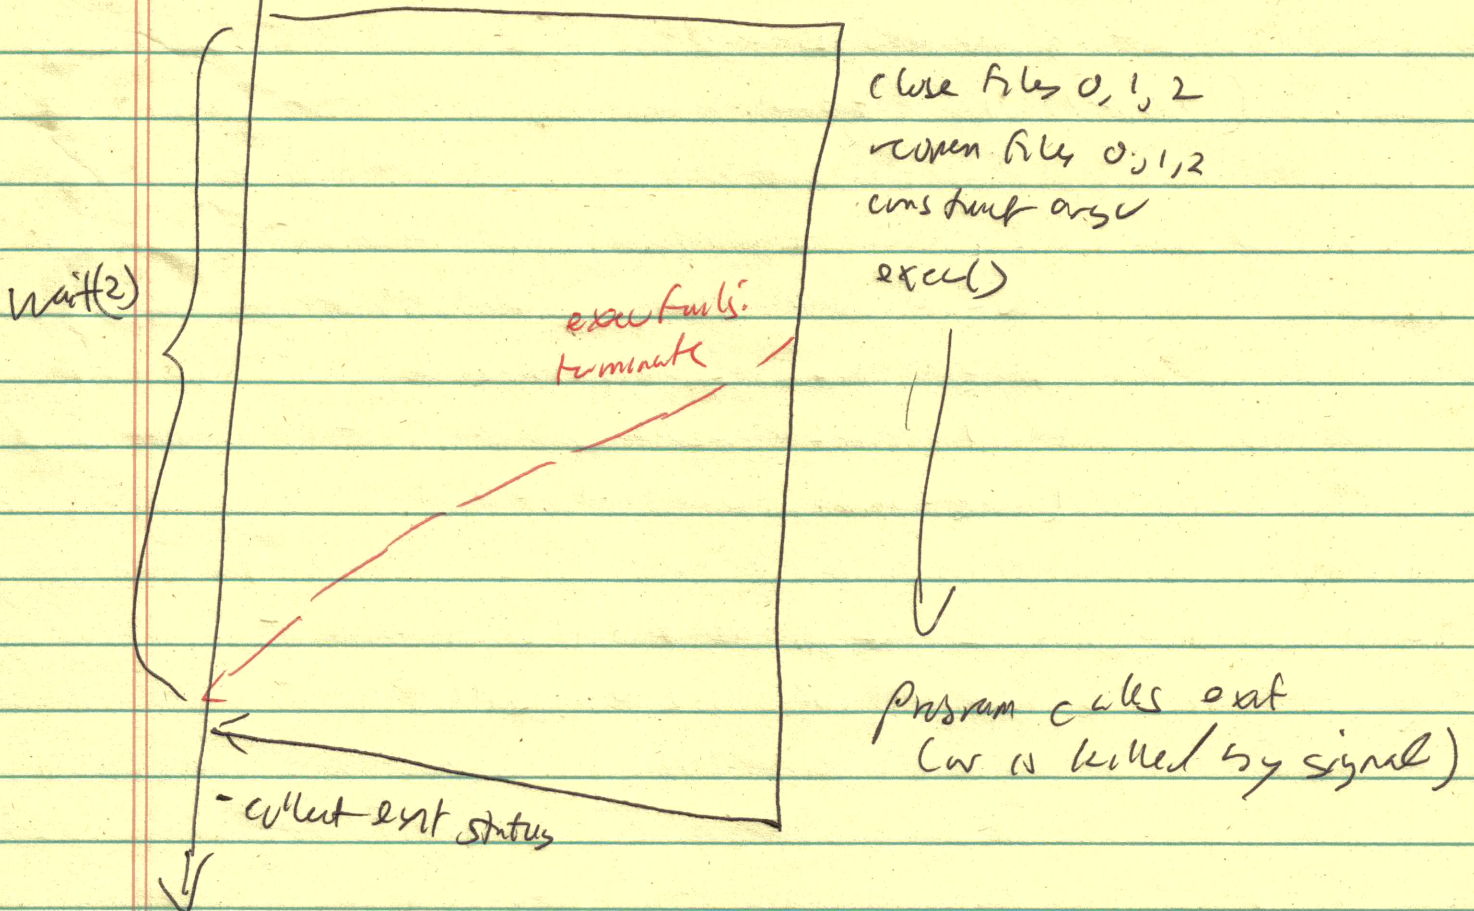
- open files (file descriptors) stay open across exec*
(unless you set the file's close-on-exec flag)

Shell

- Why separate calls for `fork()` & `exec()`?
⇒ allows shell & other processes to do some house keeping after `fork()`, before `exec()`

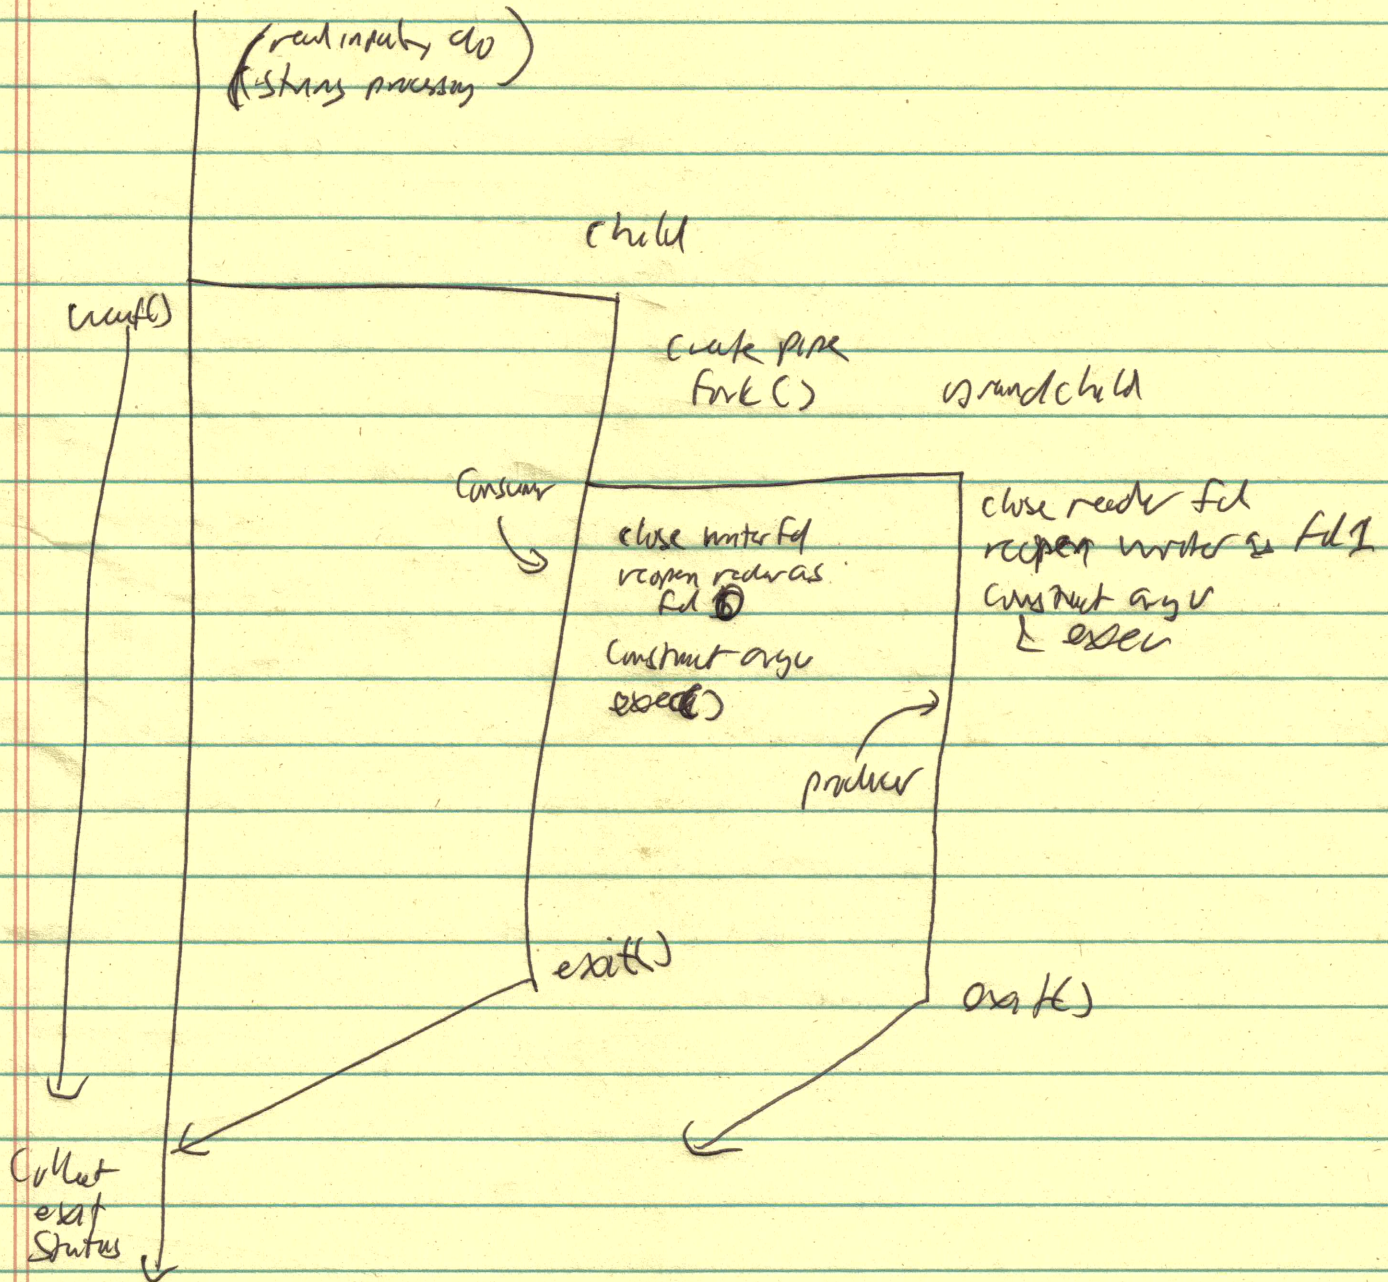
- we now have enough information to construct a rudimentary shell

- read line of input
- string processing to get filename of executable file
- string processing to get args
- string processing to get stdin/out/err
- `fork()`



Shell (cont.)


- producer-consumer pipeline (simplified)



Assignment #1: Convex Hull

- Convex Polygon

convex



non-convex



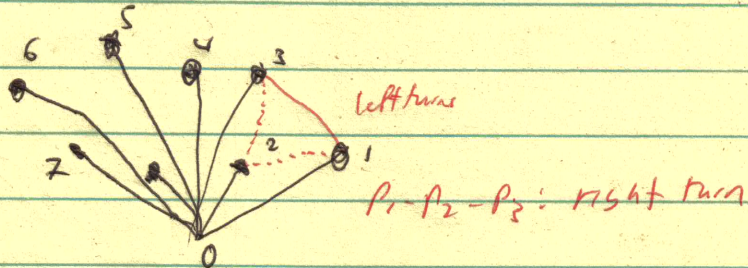
- Convex hull

- important (seems) problem in computational geometry
- much-studied, many algorithms

- Graham scan algorithm


- Sort points by angle w/ lowest point
- tricks to avoid trig ($O(N)$, but expensive), divide-by-zero
- graham scan: build stack of points, removing points that make right turns (only works in the plane - 2D)
- find min: $O(N)$
- sort: $O(N \log N)$
- graham scan: $O(N)$

Assignment #1: Convex Hull (Cont.)




- Merge 2 convex

- given 2 Convex Hulls, find C.H of union



- Points in P are ordered around p_0
 Points in Q are ordered around q_0

- assume p_0 is lower



- upper & lower chains
 - lower chain cannot be in merged hull
 - upper chain is sorted around p_0

Assignment 1: Convex Hull (cont.)

- divide & conquer algorithm
 - divide points into 2 sets
 - compute hulls of both sets
 - find upper chain of upper hull ($O(n)$)
 - merge with lower hull ($O(n)$)
 - Graham Scan of merged sets ($O(n)$)

- recursion (divide in half): $O(\log N)$
→ total running time $O(N \log N)$

- but the assignment is to fork()
compute both halves in separate processes
 - concurrent & independent operations
 - multicore architecture: 2 cores running at the same time
- child process must communicate results back to parent (pipe)

- ~~demo~~ ^{Skeleton} program does everything except multiprocessing part
 - don't need to worry about the algorithm, just the system programming part