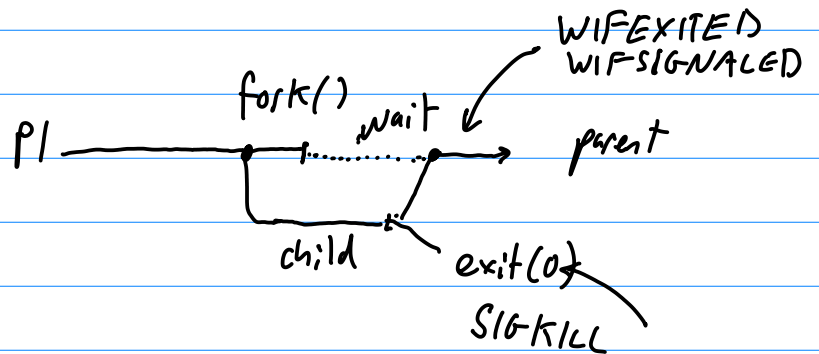


# CS 3214 lecture #6 "fork/exec fds & pipes"


Ex 1, project 1

Recap: fork



1. speed up usage of fork  
parallelization usage

2. "exec" usage of fork



fork

new process, old program  
clone of parent  
call once return 2x

exec

same process, new program  
everything\* is re-initialize  
heap, stack, regs

→ fds are kept

call once never return

"file descriptors"

how to interact/use kernel object

integers ← low number

file descriptor  
↙ handle  
↘ sequence of bytes

"everything is a file"

- 0 - stdin
- 1 - stdout
- 2 - stderr
- 3

cat foo > file

### unified\* interface


```
int fd = open("/bin/ls")
read(fd, ...
write(fd, ...
close(fd, ...
```

lseek(fd



**dup2** ← "rename / rewrite"

types of things the kernel controls

- file on disk
- devices (abstracted)
- socket
- terminals (pty)
- IPC inter process communication pipes




stdin  
stdout  
stderr



fd table

fd table


exec



redirection of fd survived exec!

copy...

pty



cat foo | grep | wc

Pipes - interprocess communication

kernel managed way  
for processes to communicate  
safely

blocks  
What if reader is too fast?

What if writer is too fast?  
blocks

