



CS 3214 lecture #4 more on processes!

{UTF} - unicode transformation format

Ex 0

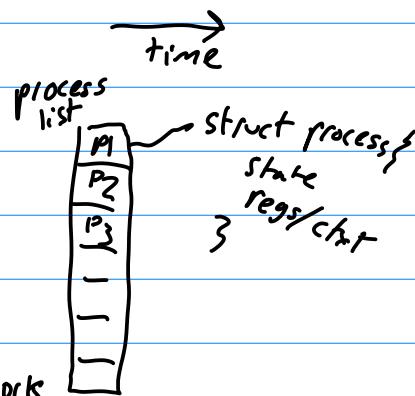
My office hours Zoom T/Th 1-2



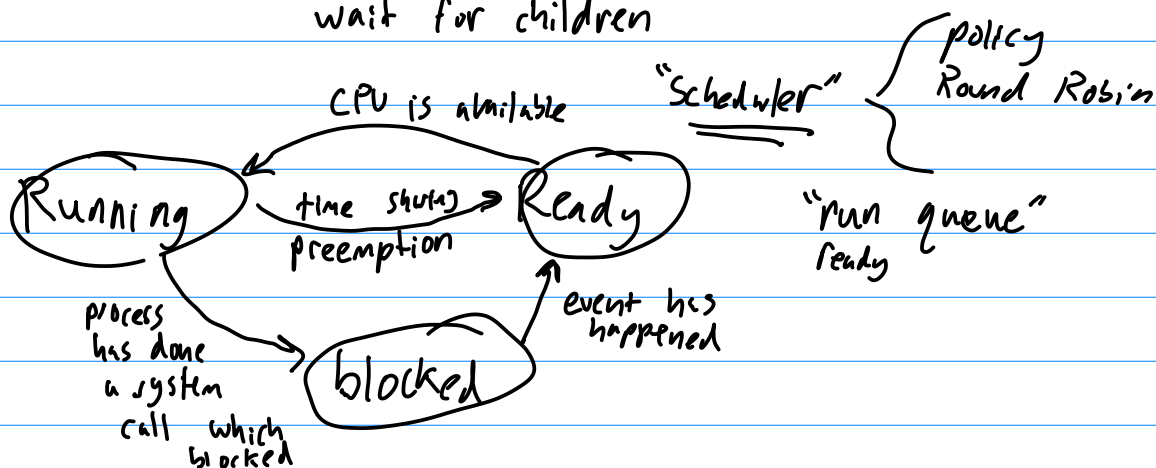
How does the kernel manage processes?

process states

time sharing {
 Running - on CPU [fetch/decode/execute]
 'Ready' ← could be on CPU
 Blocked ← user input / I/O devices, net, lock
 data from a pipe
 cat foo | head | wc



sleep
 wait for children




what if:

- n CPU + n ready processes \leftarrow all will run on CPU
- n CPU + 0 ready processes \leftarrow "idle" process "hlt"
- n CPU + $\underline{\underline{k < n}}$ ready
- n CPU + $2n$ ready \leftarrow get to run half as often
my job takes twice as long
- n CPU $m \gg n$ ready \leftarrow laggy

typical laptop: 150-500 blocked
0-2 running

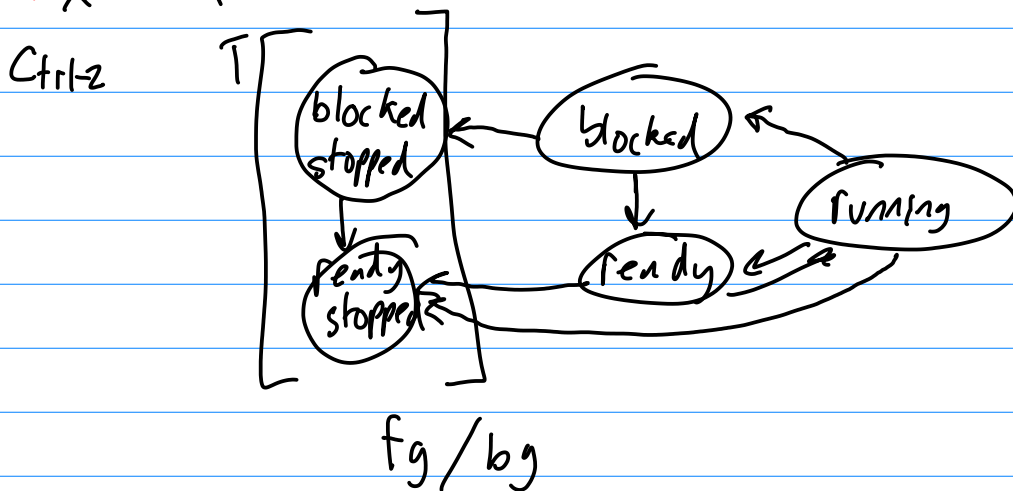
Real States in Linux

- $\left\{ \begin{array}{l} S - \text{sleep "interruptible"} \\ D - \text{uninterruptible sleep} \end{array} \right.$
- $\left\{ \begin{array}{l} R - \text{running/ready "runnable"} \\ I - \text{"idle kernel thread"} \end{array} \right.$
- T - job control stop
- t - debugger



~~X~~
~~X - dead~~

Z - defunct "zombie" process (terminated but parent has not waited for it)



→ poll ← wastes CPU cycles
↙ ↘
block

measure CPU load

- load average

- time

Generally prefer blocking *
↳ use the facilities your kernel gives you

Transitions are unpredictable

Next time: process management!
creation/termination/etc.