

Lecture 4: threads & Scheduling

- More on PCB
- Threads
- Show & tell
- Scheduling

What goes in the PCB?

- everything the OS needs to know to manage the process

- registers & processor state (flags)

- process state, if waiting, what for?

- scheduling info (e.g. priority, history)

- permissions (e.g. userid, groupid)

- state of resources (e.g. open files, network connections)
file descriptors

- identifier (process ID), parent

- configuration (example: signal handlers:

location of code (in proc)
to handle signals from OS

(like interrupts)

- o Processes have their own address spaces
 - completely isolated (except via. inter-process-comm syscalls)
 - need context switches to communicate expensive (cache invalidation)

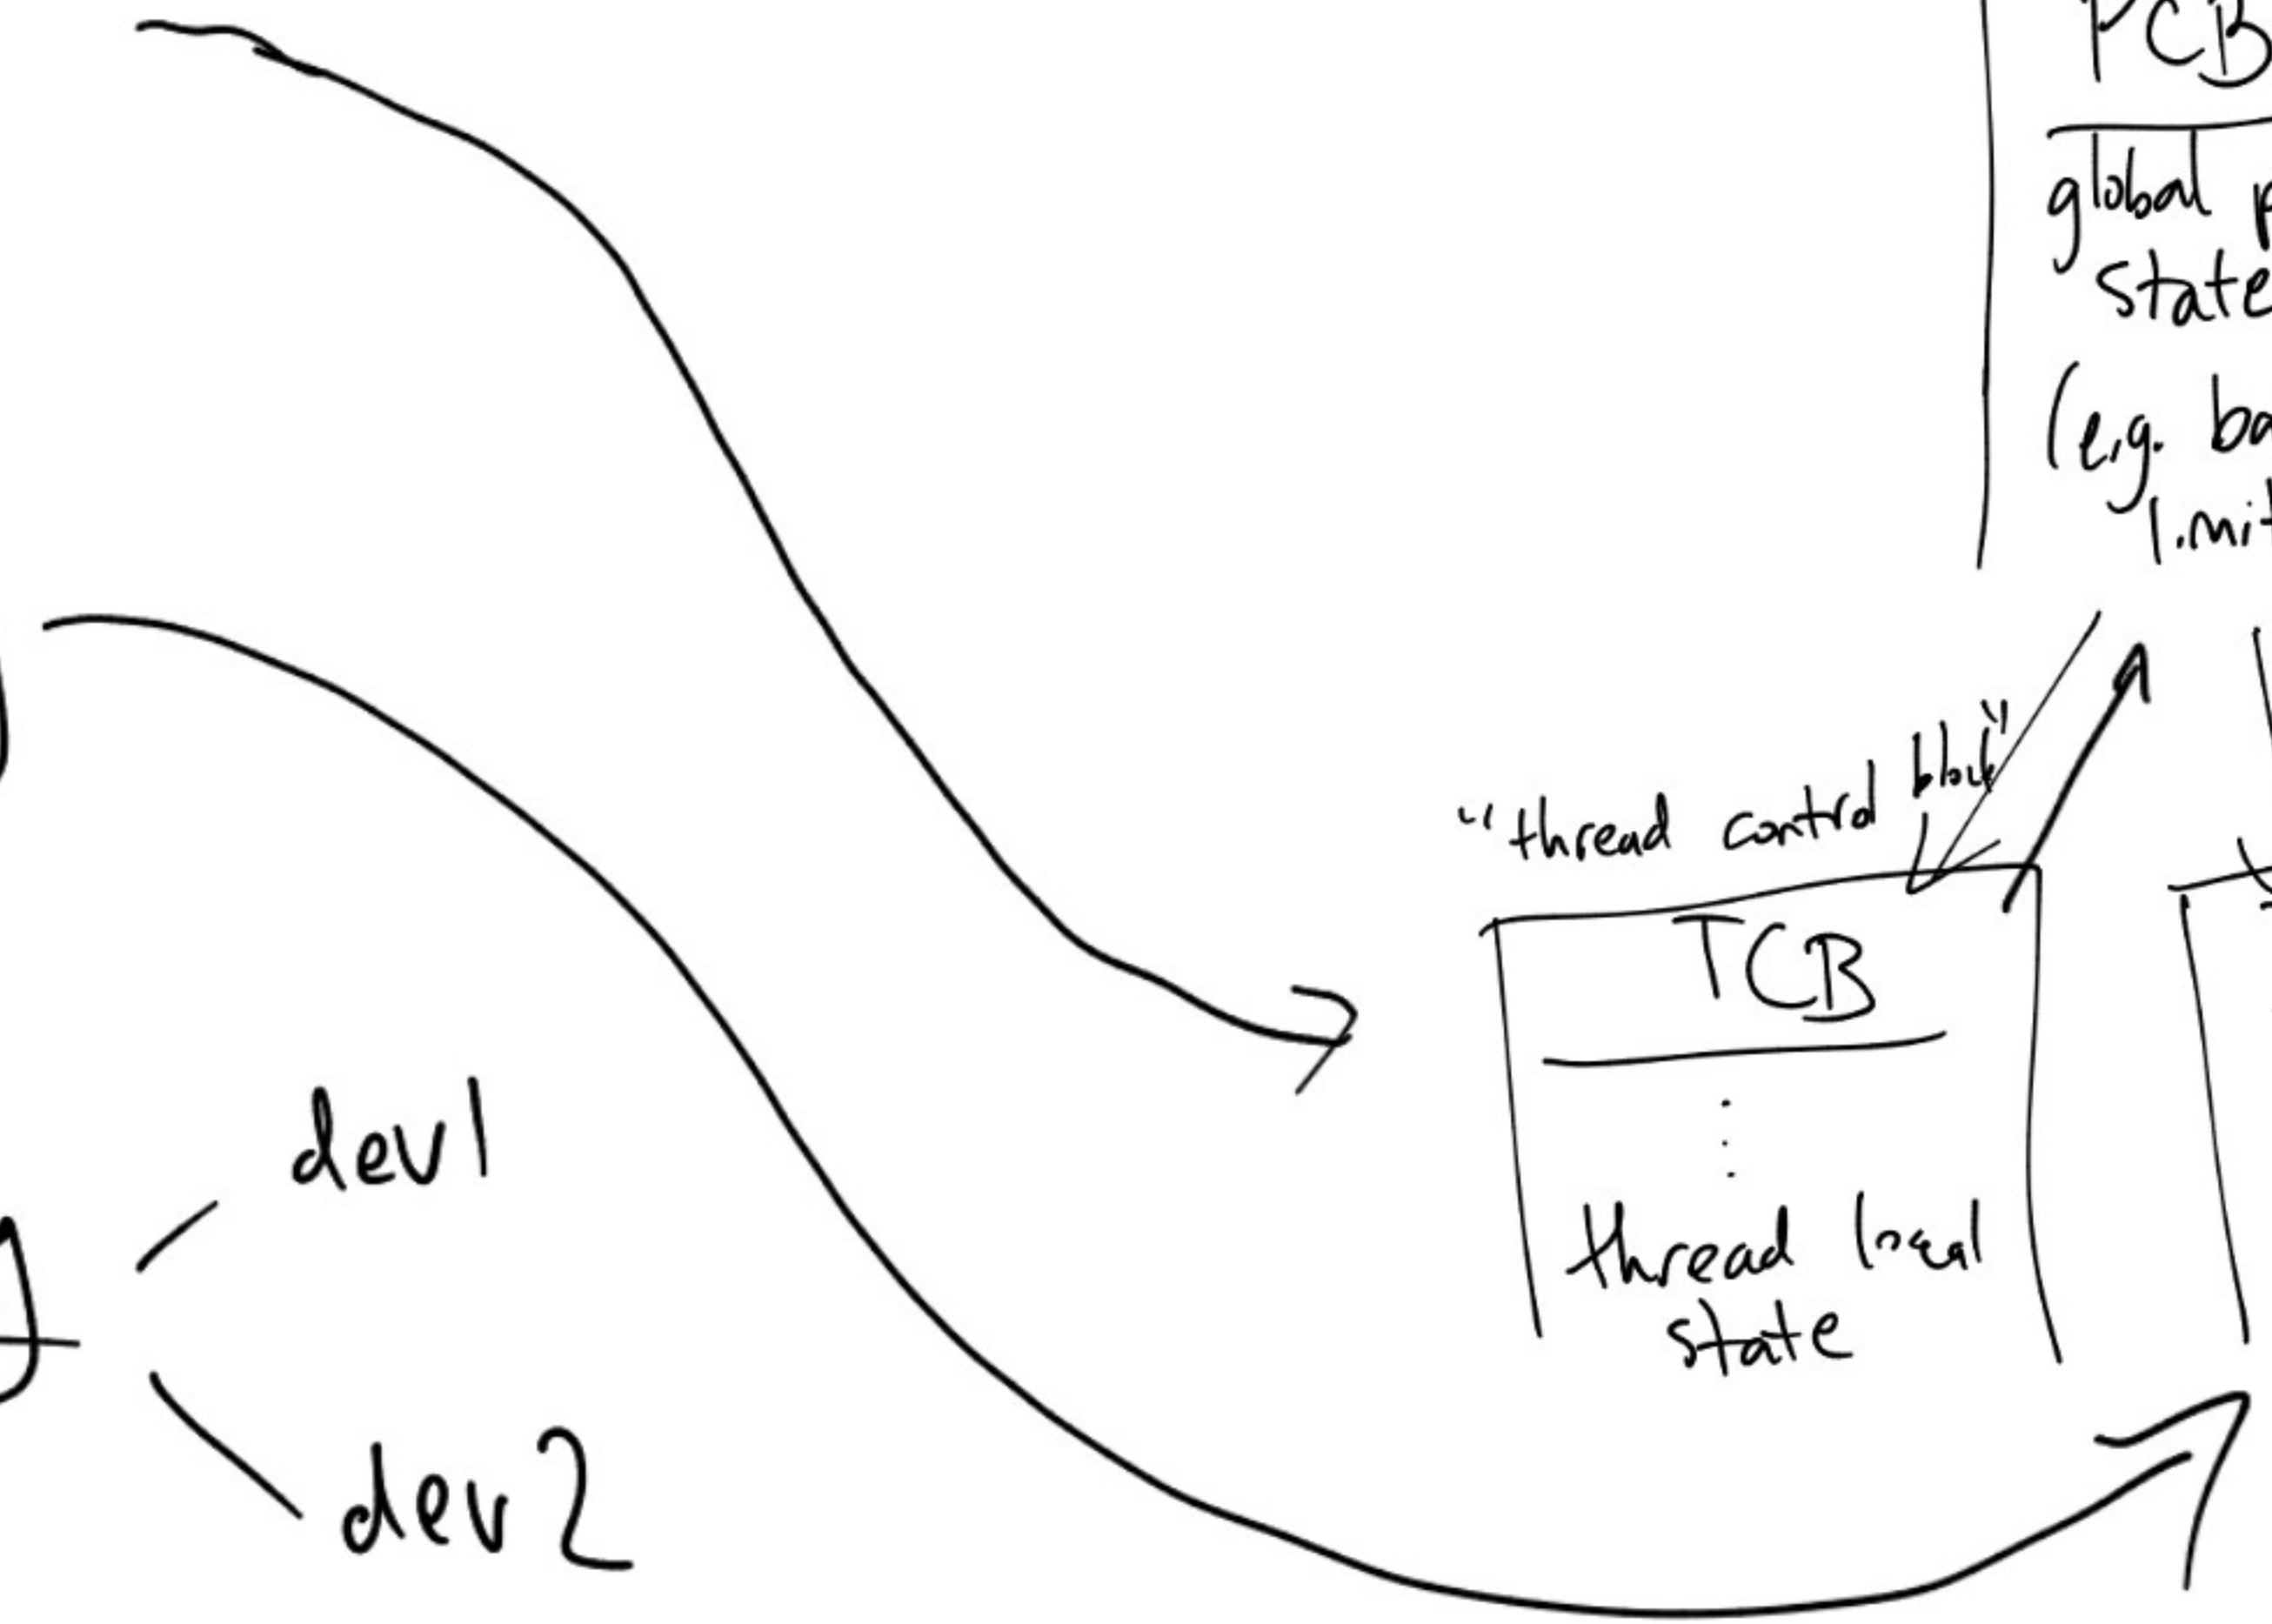
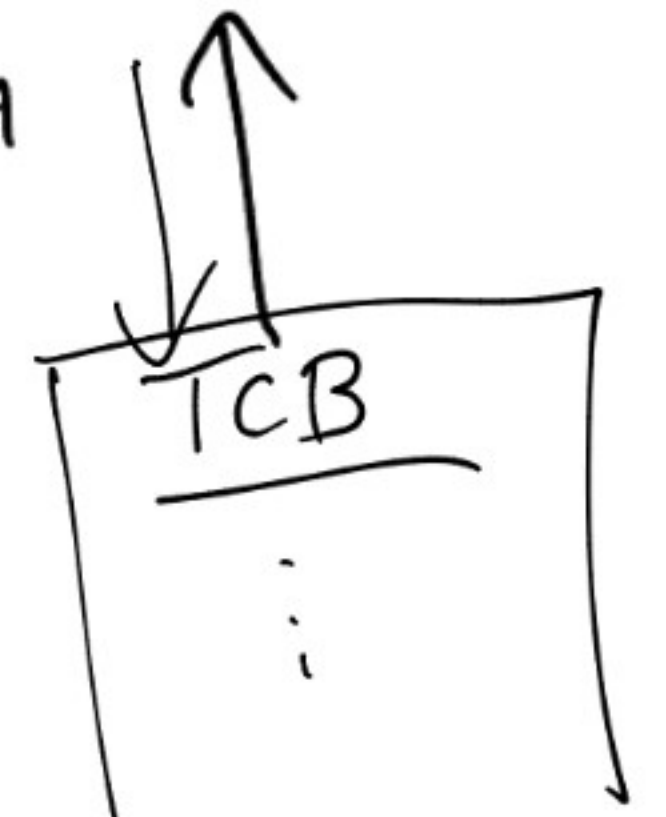
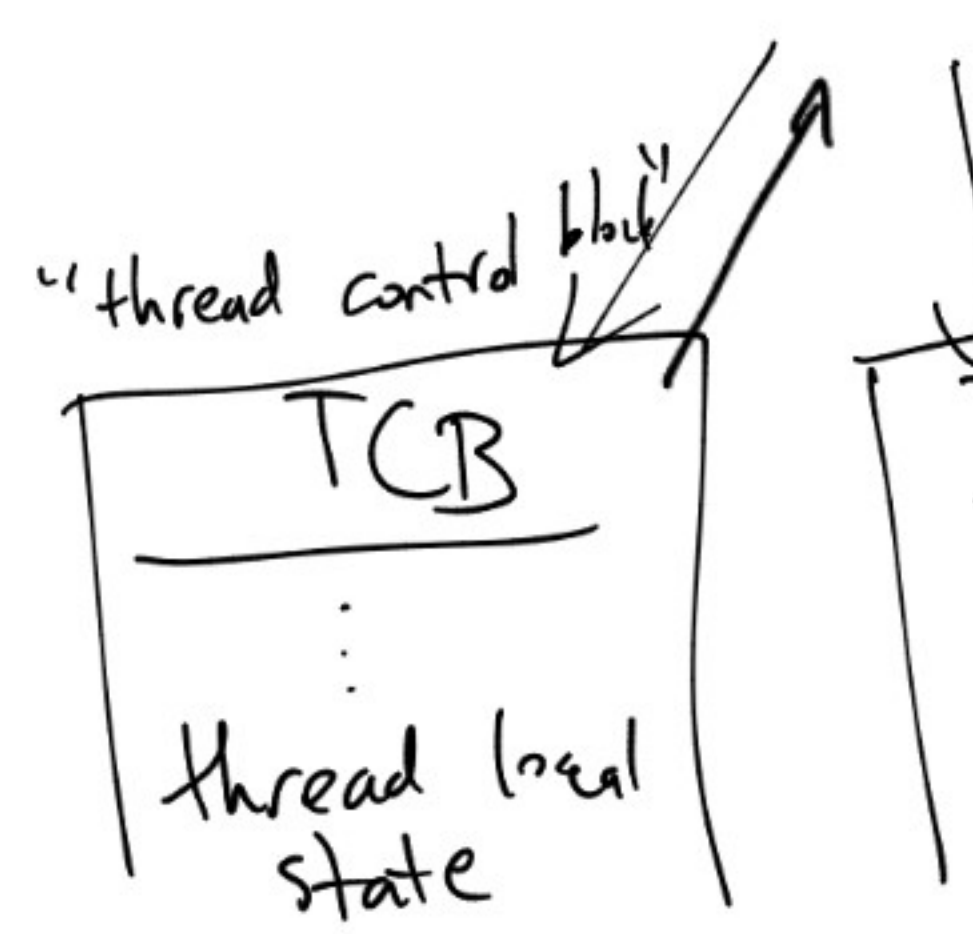
- o "Processes" that share address space are called threads.
 - have own local vars,
shared global variables

OS's memory

ready

running

writing
- dev1
- dev2



Scheduling

- when CPU becomes free, what to do next?

Criteria:

- Fairness

- Efficiency:

 - scheduler itself

 - avoid context switches

- Responsiveness:

 - respond quickly to user requests

- Priority
 - (specified by user)

- Simplicity.

FCFS (First come, first served)

- put incoming jobs in queue (at end)
- to schedule, select a job (from front)

Pros:

- Fair (if all jobs finish)
- simple.
- scheduler super efficient
- not many ctx switches.

Cons:

- No notion of priority
- average waiting time can be bad.

Preemption

- o set a timer (will cause interrupt after a delay)
 - when timer goes off, schedule another process.

Round robin: Preemptive FCFS

- schedule proc for fixed amount of time (quantum)
- go to end of line after timeout.