

# Processes

- What are they / why needed?
- How does dispatcher run a process?
- How does system create a process?

"Virtual CPU"

"Activity of some kind"  
=> Program, I/O, state

Process: def

Execution stream, in the context of a process state  
more simply

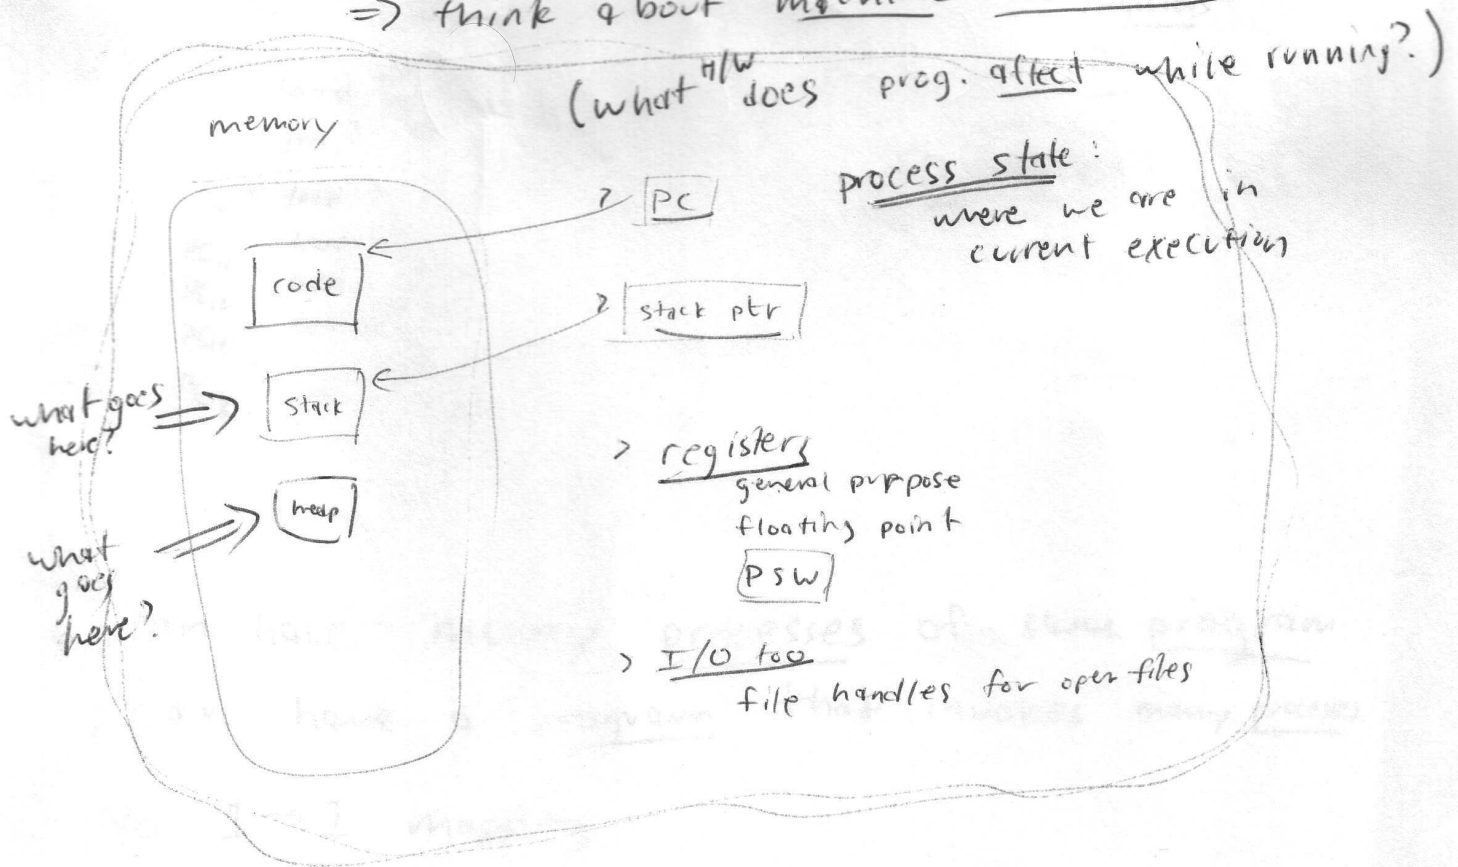
"Running program + stuff it needs/uses"

"thread of control"

can affect, be affected by

What things can code affect / be affected by?

=> think about machine architecture



# Program vs Process

②

③

Different!

Program : static code + static data

## Program

```
int  
main ( int argc, char * argv ) {  
    f(x);  
}  
  
void  
f ( int x ) {  
    ...  
}
```

sequence  
of instructions

```
load x, reg  
jmp "f"
```

```
f.  
load  
load  
add  
store  
ret
```

## Process

dynamic flow of instructions  
"prog in execution"  
+ "process state"

- heap
- stack
  - main
  - foo
- registers
- PC

PC →

PC <sub>1</sub>	load
PC <sub>2</sub>	jmp 10
PC <sub>10</sub>	load
PC <sub>11</sub>	load
PC <sub>12</sub>	add
PC <sub>13</sub>	return
PC <sub>3</sub>	—

Dynamic

Thus,

{ can have many processes of same program  
can have a program that invokes many processes

No 1 → 1 mapping

# Threads


## Address Space:

> will discuss in detail in mem mgmt

- > "virtual" address space
  - ⇒ <sup>really</sup> large memory to use
- > linear array of bytes:  $[0 \dots N-1]$

$$N \sim \underline{\underline{2^{32}}} \quad \underline{\underline{2^{64}}}$$

def: all logical entities used by a process + address by which they are referenced



one-to-one mapping: ~~all~~ Process  $\Leftrightarrow$  AS  
⇒ protection boundary


# Threads

Process  $\neq$  Thread

Threads: Many execution streams w/ in one process!

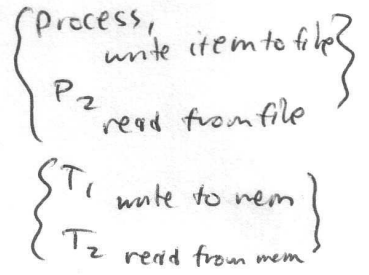
"lightweight process"

"shared address space"



> why do this?

- > efficient communication if cooperating to do something
- > express concurrency + ↑



# Why use Processes?

(4)

1) Divide and Conquer

Take large problem  $\Rightarrow$  many small ones  
(smaller  $\Rightarrow$  easier)

2) ~~Simple way to express~~ concurrency  
lots of users, devices, etc.

3) Easy to reason about processes  
sequential activity

## System Classifications

Uniprogramming: One process @ a time

early PCs

Why? simple

Not? inconvenient, poor perf.

Multiprogramming: many @ time

All modern OSs

! multiprocessing

Why? perf, easier

$\Rightarrow$  machine w/  
 $> 1$  processor

Not? complex

# Multiprogramming

## OS requirements

- Policy : when to schedule A, B, ...
- Mechanism : How to switch between processes
- How to protect them from one another

## Separation of P/M

- Policy : depends on ① workload, other assumptions about env.
- Mechanism : basic functionality for doing stuff

w/ processes

P = scheduling (later)

M = dispatching (today)

Java/C


Project

#1 out  
we on Tuesday

## Dispatch Mechanism

Data structure : list of processes (per-process state)

Process : 3 nodes



### Dispatcher

How to desched A, run B

> "gain control", save state of A, load state of B, run B

How?

what state?

context switch

# How to gain control?


6

general

User Processes

"user mode"

limited access to h/w



(system)

"kernel mode"

can access anything

how: h/w support

H/w slw interface

Problem:

only one CPU


CPU is running user program

dispatcher: how to run??

Two ways:

Traps: internal event in user process  
e.g. syscall, <sup>(illegal inst, 0)</sup> error, <sup>(page fault)</sup> fault

Interrupt: external event  
character typed, disk I/O done



DISP: gaining control

Cooperative approach

trust process to give up CPU (yield)

why bad? OS trusts apps!

[Alto => Mac]

Non-cooperative

OS trusts no one

Timer interrupt => give OS control (dispatcher)  
(every 10 ms)

Can count "ticks", decide to switch to diff job (policy)

key: user can't turn off interrupts  
(user/sys mode)

=> Building OS is about being paranoid!

H/W Support

# What state must be saved?

7

OS: must track state of  $P_{procs}$  when not running

on trap/interrupt,

save state to PCB (Process Control Block)

## Info in PCB

execution state : registers, status word, PC, stack ptr, etc)

i/o state : open files

sched info : state, priority

accounting : owner, pid

What needs to be saved? (on switch)

no protection (a) Trust: early PCs, macs


(b) All of memory [A to]

Slow: 1 GB of mem  $\Rightarrow$  disk?

(b) Protect memory  
[later]

trust

non-trusting



## Context-switch: Implementation

Machine-dependent (assembly)  
save/restore registers



### Why hard?

code needs registers to run

$\Rightarrow$  H/W support

CISC : instr. to save regs  $\Rightarrow$  stack

RISC : convention: don't use  $R_1, R_2$

H/W

S/W

interface

### How expensive?

many loads/stores

problem: copy can be expensive (esp. if  $exec()$  is next)

# Process Creation

8

2 ways:


{ from scratch  
by imitation (cloning) }

## Scratch:

Load code/data  
Create stack, PCB  
Process  $\Rightarrow$  ready list

## Clone

Stop current, save state  
Copy state  
Add process  $\Rightarrow$  ready list



Unix: "clone approach"

fork() clones process

But: just copies same process

exec() Overlays new image on calling process

shell  
shell.c

```
cmd = get-command();  
int rc = fork();  
if (rc < 0) {  
    // error  
} else if (rc == 0) {  
    // child  
    execv(cmd, args);  
    printf("failed");  
} else {  
    // parent  
    int pid = rc;  
    wait(pid);  
}
```

1) not an exact copy: why?

2) problem: copy can be expensive (esp. if exec() is next)

skip



# Example : Process Creation in Windows NT

9

Create Process ( AppName, CmdLine, Process Attributes,  
Thread Attributes, InheritHandles,  
Creation Flags, Environment,  
Current Directory, Startup Info,  
Process Info );

History  
VAX/VMS  $\Rightarrow$  NT  
Mach

As David Korn says

"There is a single primitive, named  
CreateProcess(), that takes ten  
arguments, yet still cannot perform  
the simple operation of overlaying  
the current process with a new  
program as execve() requires"