

memory (virtualization)

~ dive into xv6!

recap: the process abstraction

- + xv6 riscv git
- + xv6 book
- + xv6 source code listing

decouple program/logic
encoding / development
from
program execution

CPU
virtualization
(via context save & restore
& scheduling).

memory

process-view

- ~ data objects / variables
- + code / functions
- + stack ptr.

~ an (independent) / explicit
CPU for execution

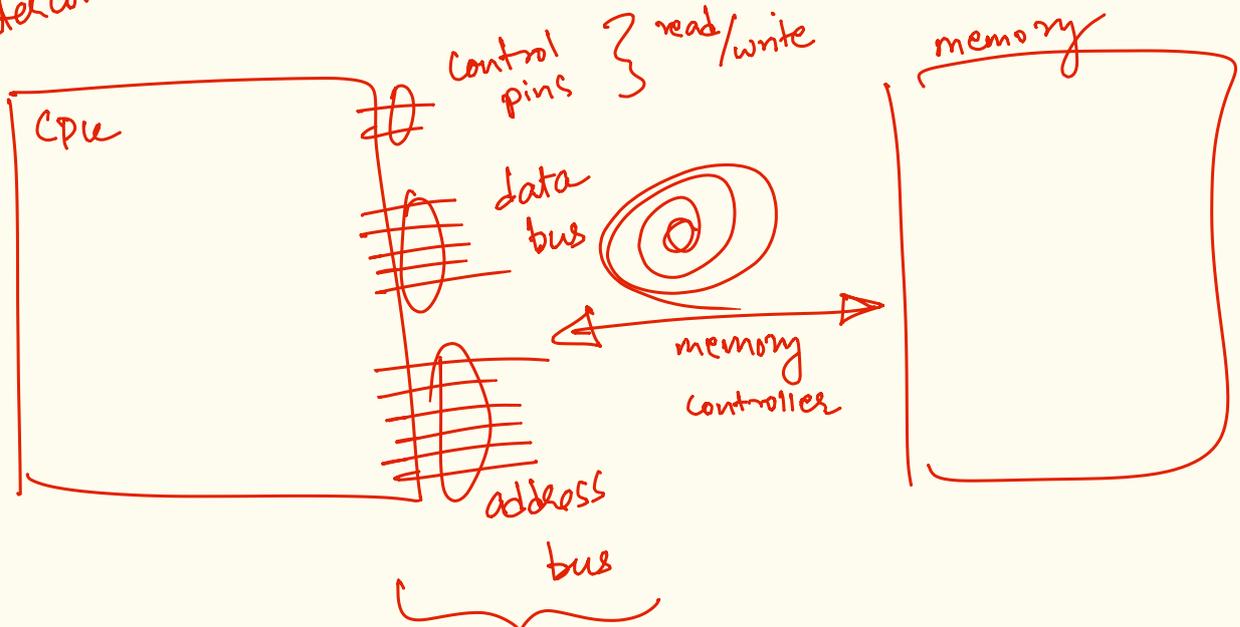
(virtualization)

mgmt.

addressing
mechanisms
for allocation

- + how much?
- + where to allocate?
- + eviction policy?

memory access / interconnect view



- identifying / addressing memory
- address-width - 8-bits, 16, 32, 64 ...
- address range ~ how much memory can the CPU address.

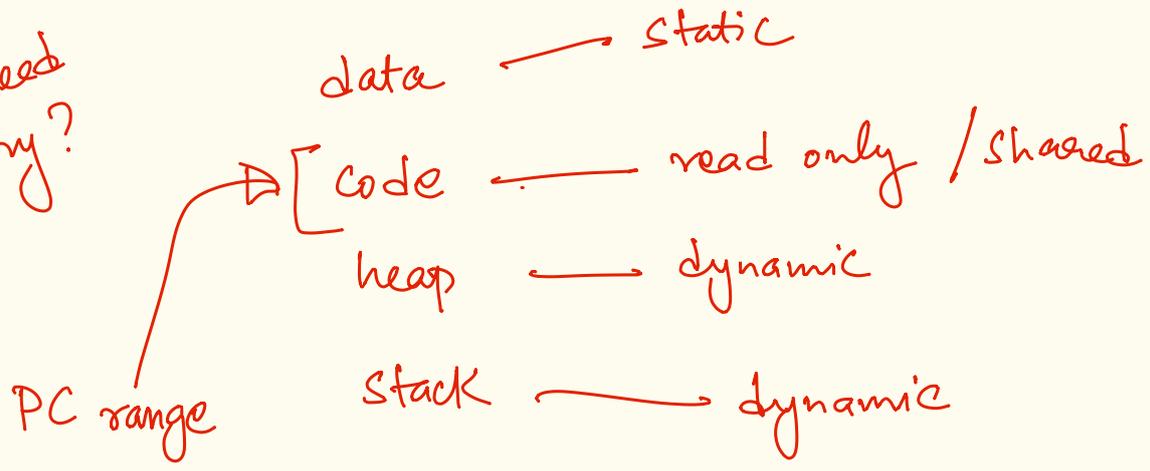
address: 0/1 encoding on the address-bus.

0000 0000 0000 1111

0x 000f ~ 16-bit address!

0x 0000
 |
 0x ffff } 2^{16} addresses

⊗ why do process need memory?



memory allocation/access for processes requirements

(i) isolation

(ii) "full" addressability

& zero-starting address range

(iii) transparent, correctness

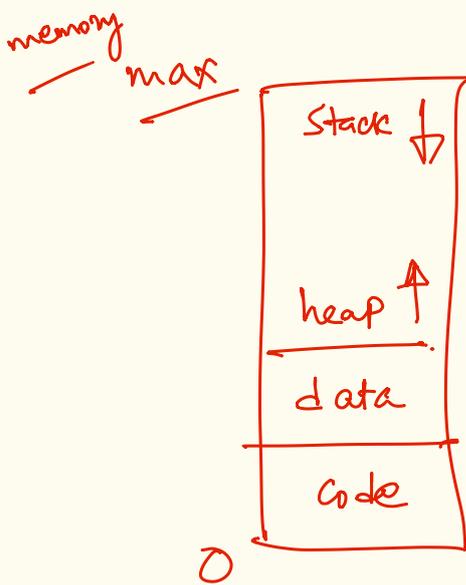
(iv) efficient.

design 0

~ similar to the context switch idea

~ ^{extend} save & restore to memory!

~ all memory available to single executing process!



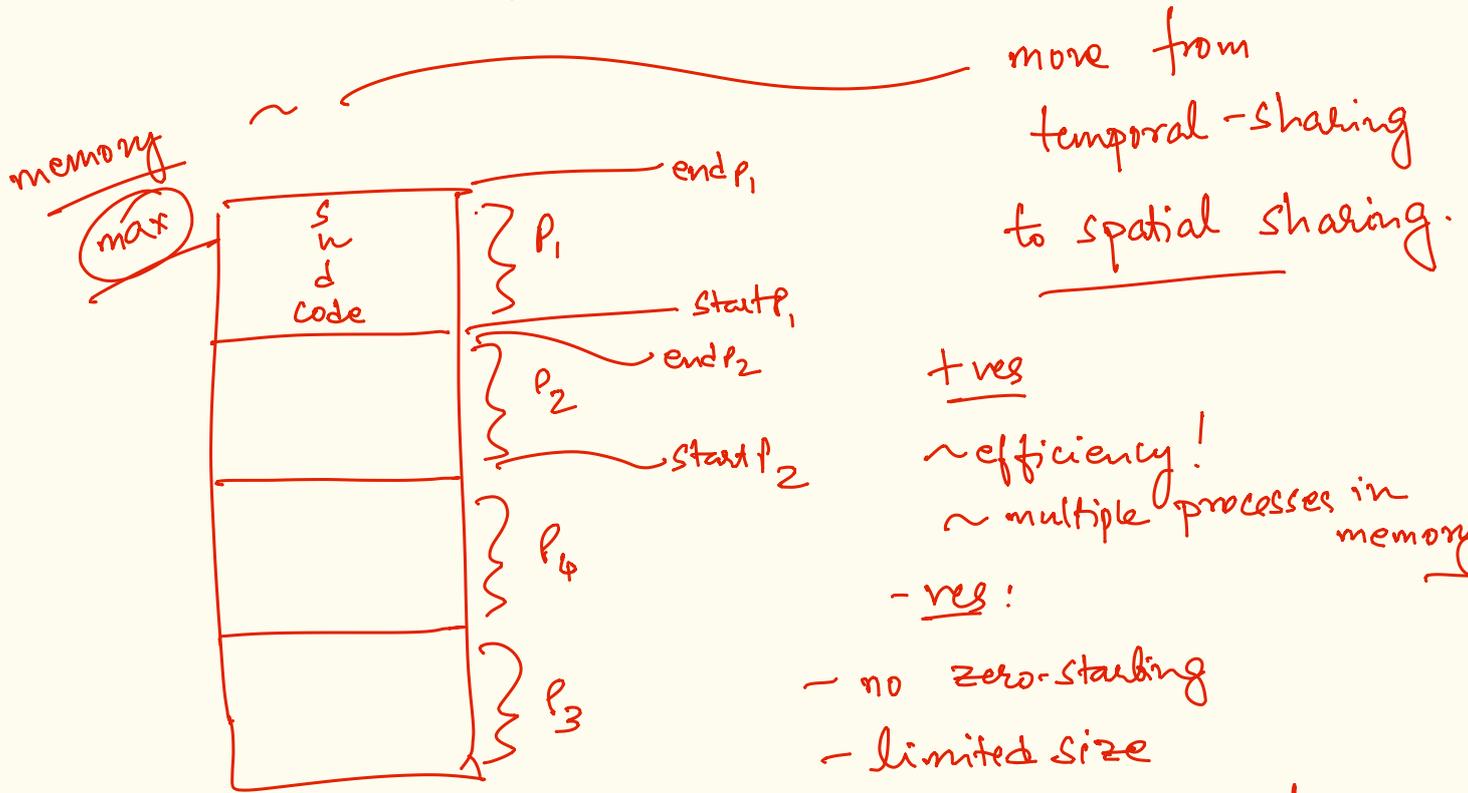
+ves: zero-starting
full range

-ves: ~ potential offload destⁿ
disk
is slow, slow, slow

~ size of context
is all memory!

design 1:

~ not all memory may be used at all times!



+ves

~ efficiency!

~ multiple processes in memory.

-ves:

- no zero-starting

- limited size

- no dynamic size change possibility.

design 2:

~ decoupling of addressing from allocation
via ISA feature.