

Lecture II: Memory management

- simple base/limit MMU
- address translation
- paging / TLB
- swapping

Defn state is safe if \exists an order in which processes can be run completely (assuming they acquire max allocation, then release all resources)

Bankers algorithm: maintain invariant that state is safe by blocking when resources are requested.

- when process requests resource, block until granting request would be safe.

| | current | | | max | | |
|------|---------|---|---|-----|---|---|
| | A | B | C | A | B | C |
| P1 | 1 | | | 1 | 1 | |
| P2 | | 1 | | | 1 | 1 |
| P3 | | | | | 1 | 1 |
| Free | | 1 | | 1 | 1 | 1 |

Can't run P1 fully
 Can't run P3 fully
 could do P2

⇒

| | A | B | C |
|---------------|---|---|---|
| P1 | 1 | | |
| P2 | | | |
| P3 | | | |
| F | | 1 | 1 |

then P1
 then P3.

P3 requests 1 resource C.

if granted:

| | A | B | C |
|----|---|---|---|
| P1 | 1 | | |
| P2 | | 1 | |
| P3 | | | 1 |
| F | 0 | 0 | 0 |

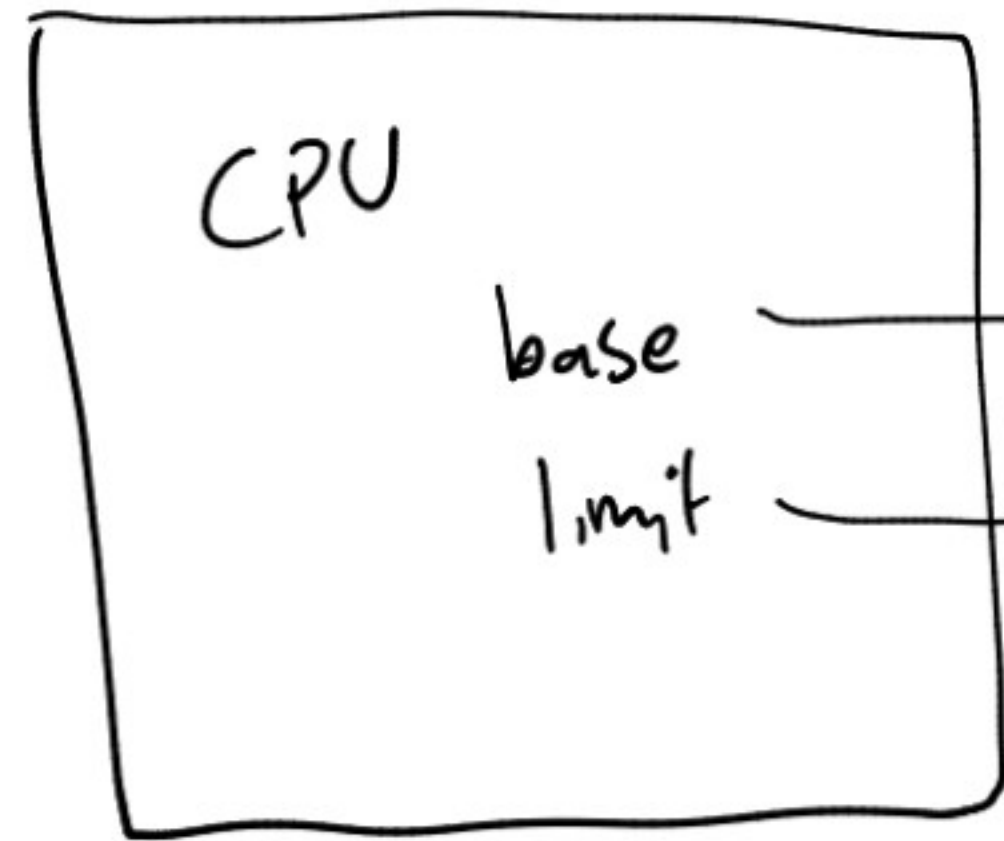
Not safe
 block.

P2 requests C

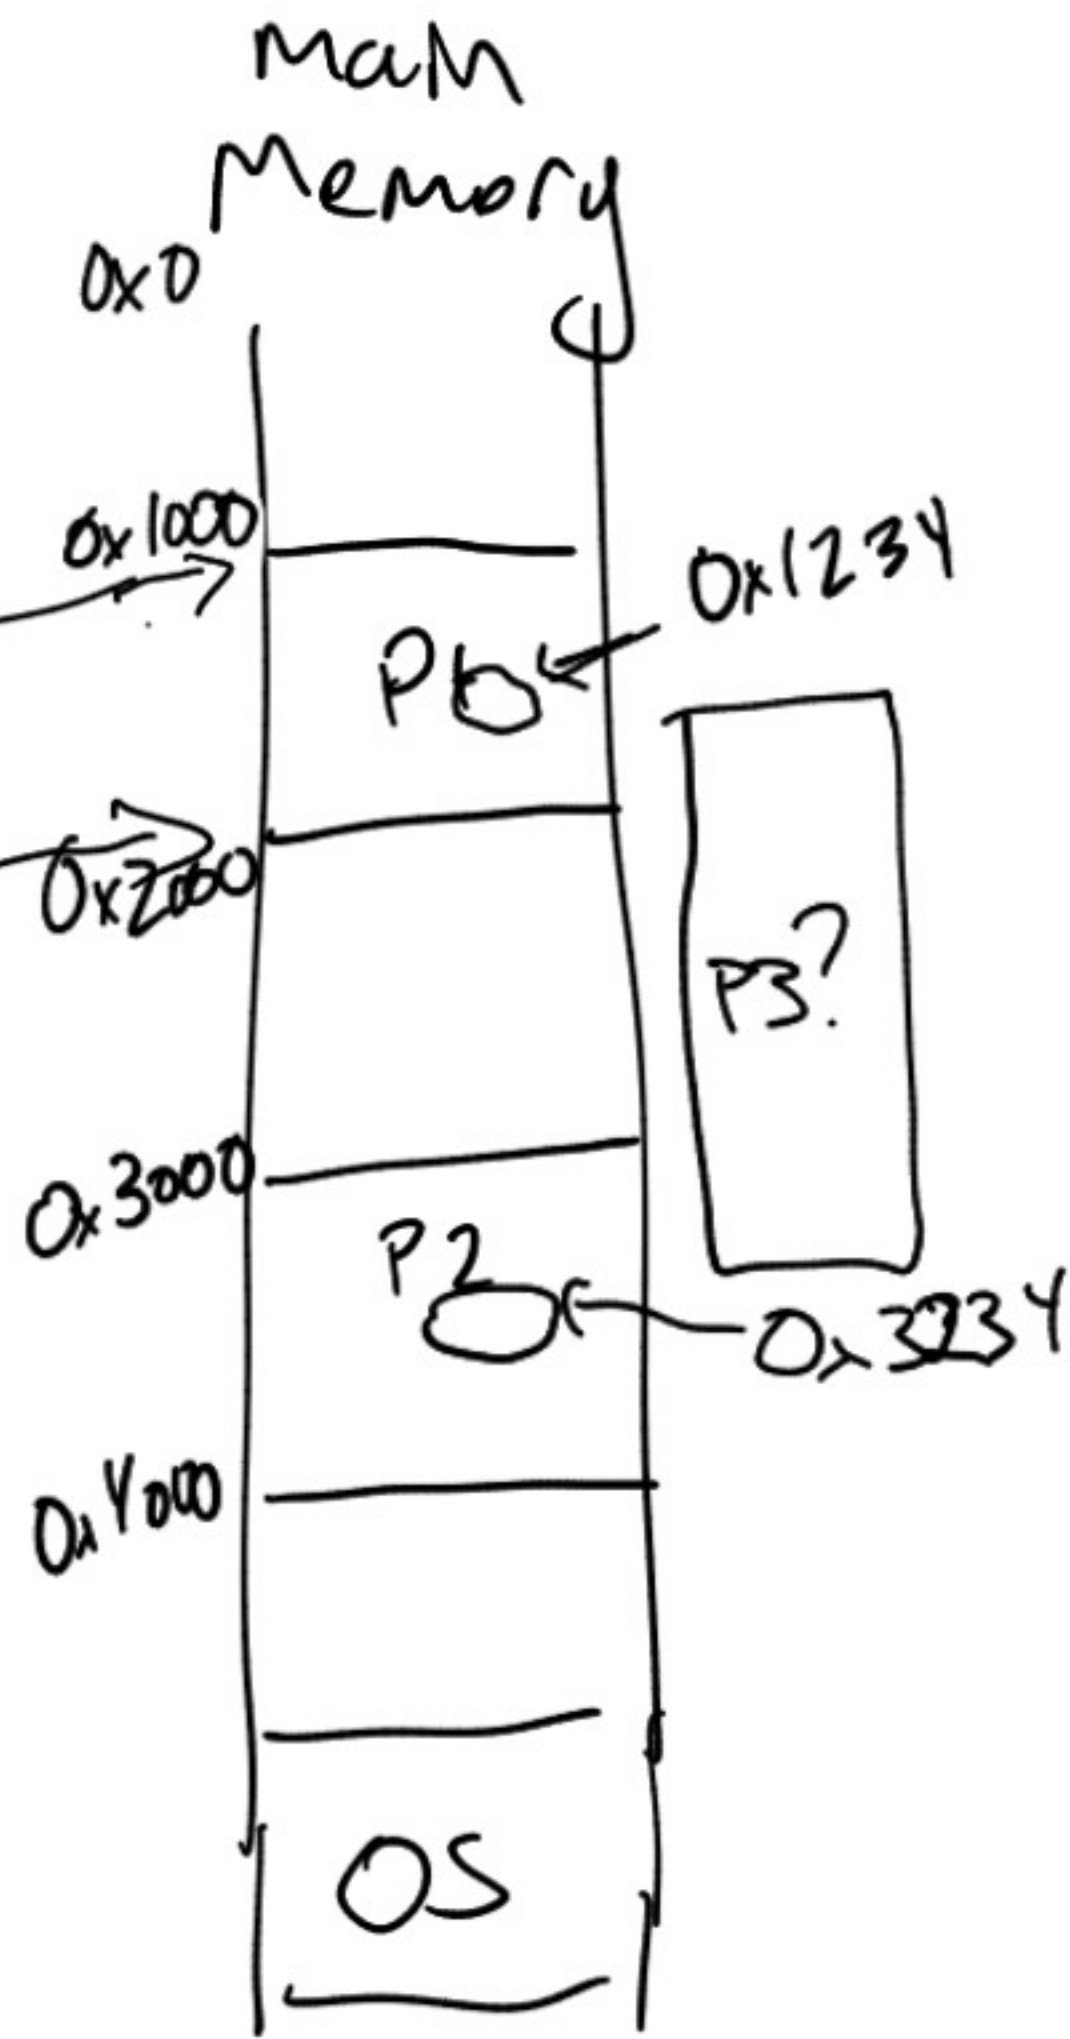
if granted:

| | A | B | C |
|----|---|---|---|
| P1 | 1 | | |
| P2 | | 1 | 1 |
| P3 | | | |
| F | 0 | 0 | 0 |

Safe:
 run P2, P1, P3



- on load, store, IR address is out of base, limit range, cause an exception.
- on ctxt switch, change base & limit to point to new prog.





Problems

- o Allocation: ensure processes don't overlap.
 - fragmentation enough free memory to use, but not contiguous.
 - o Procs need to adapt to their location.
- Not actually problems
- o how do we know base & limit are?

Strawman: Address translation

Process always uses addresses from 0... site



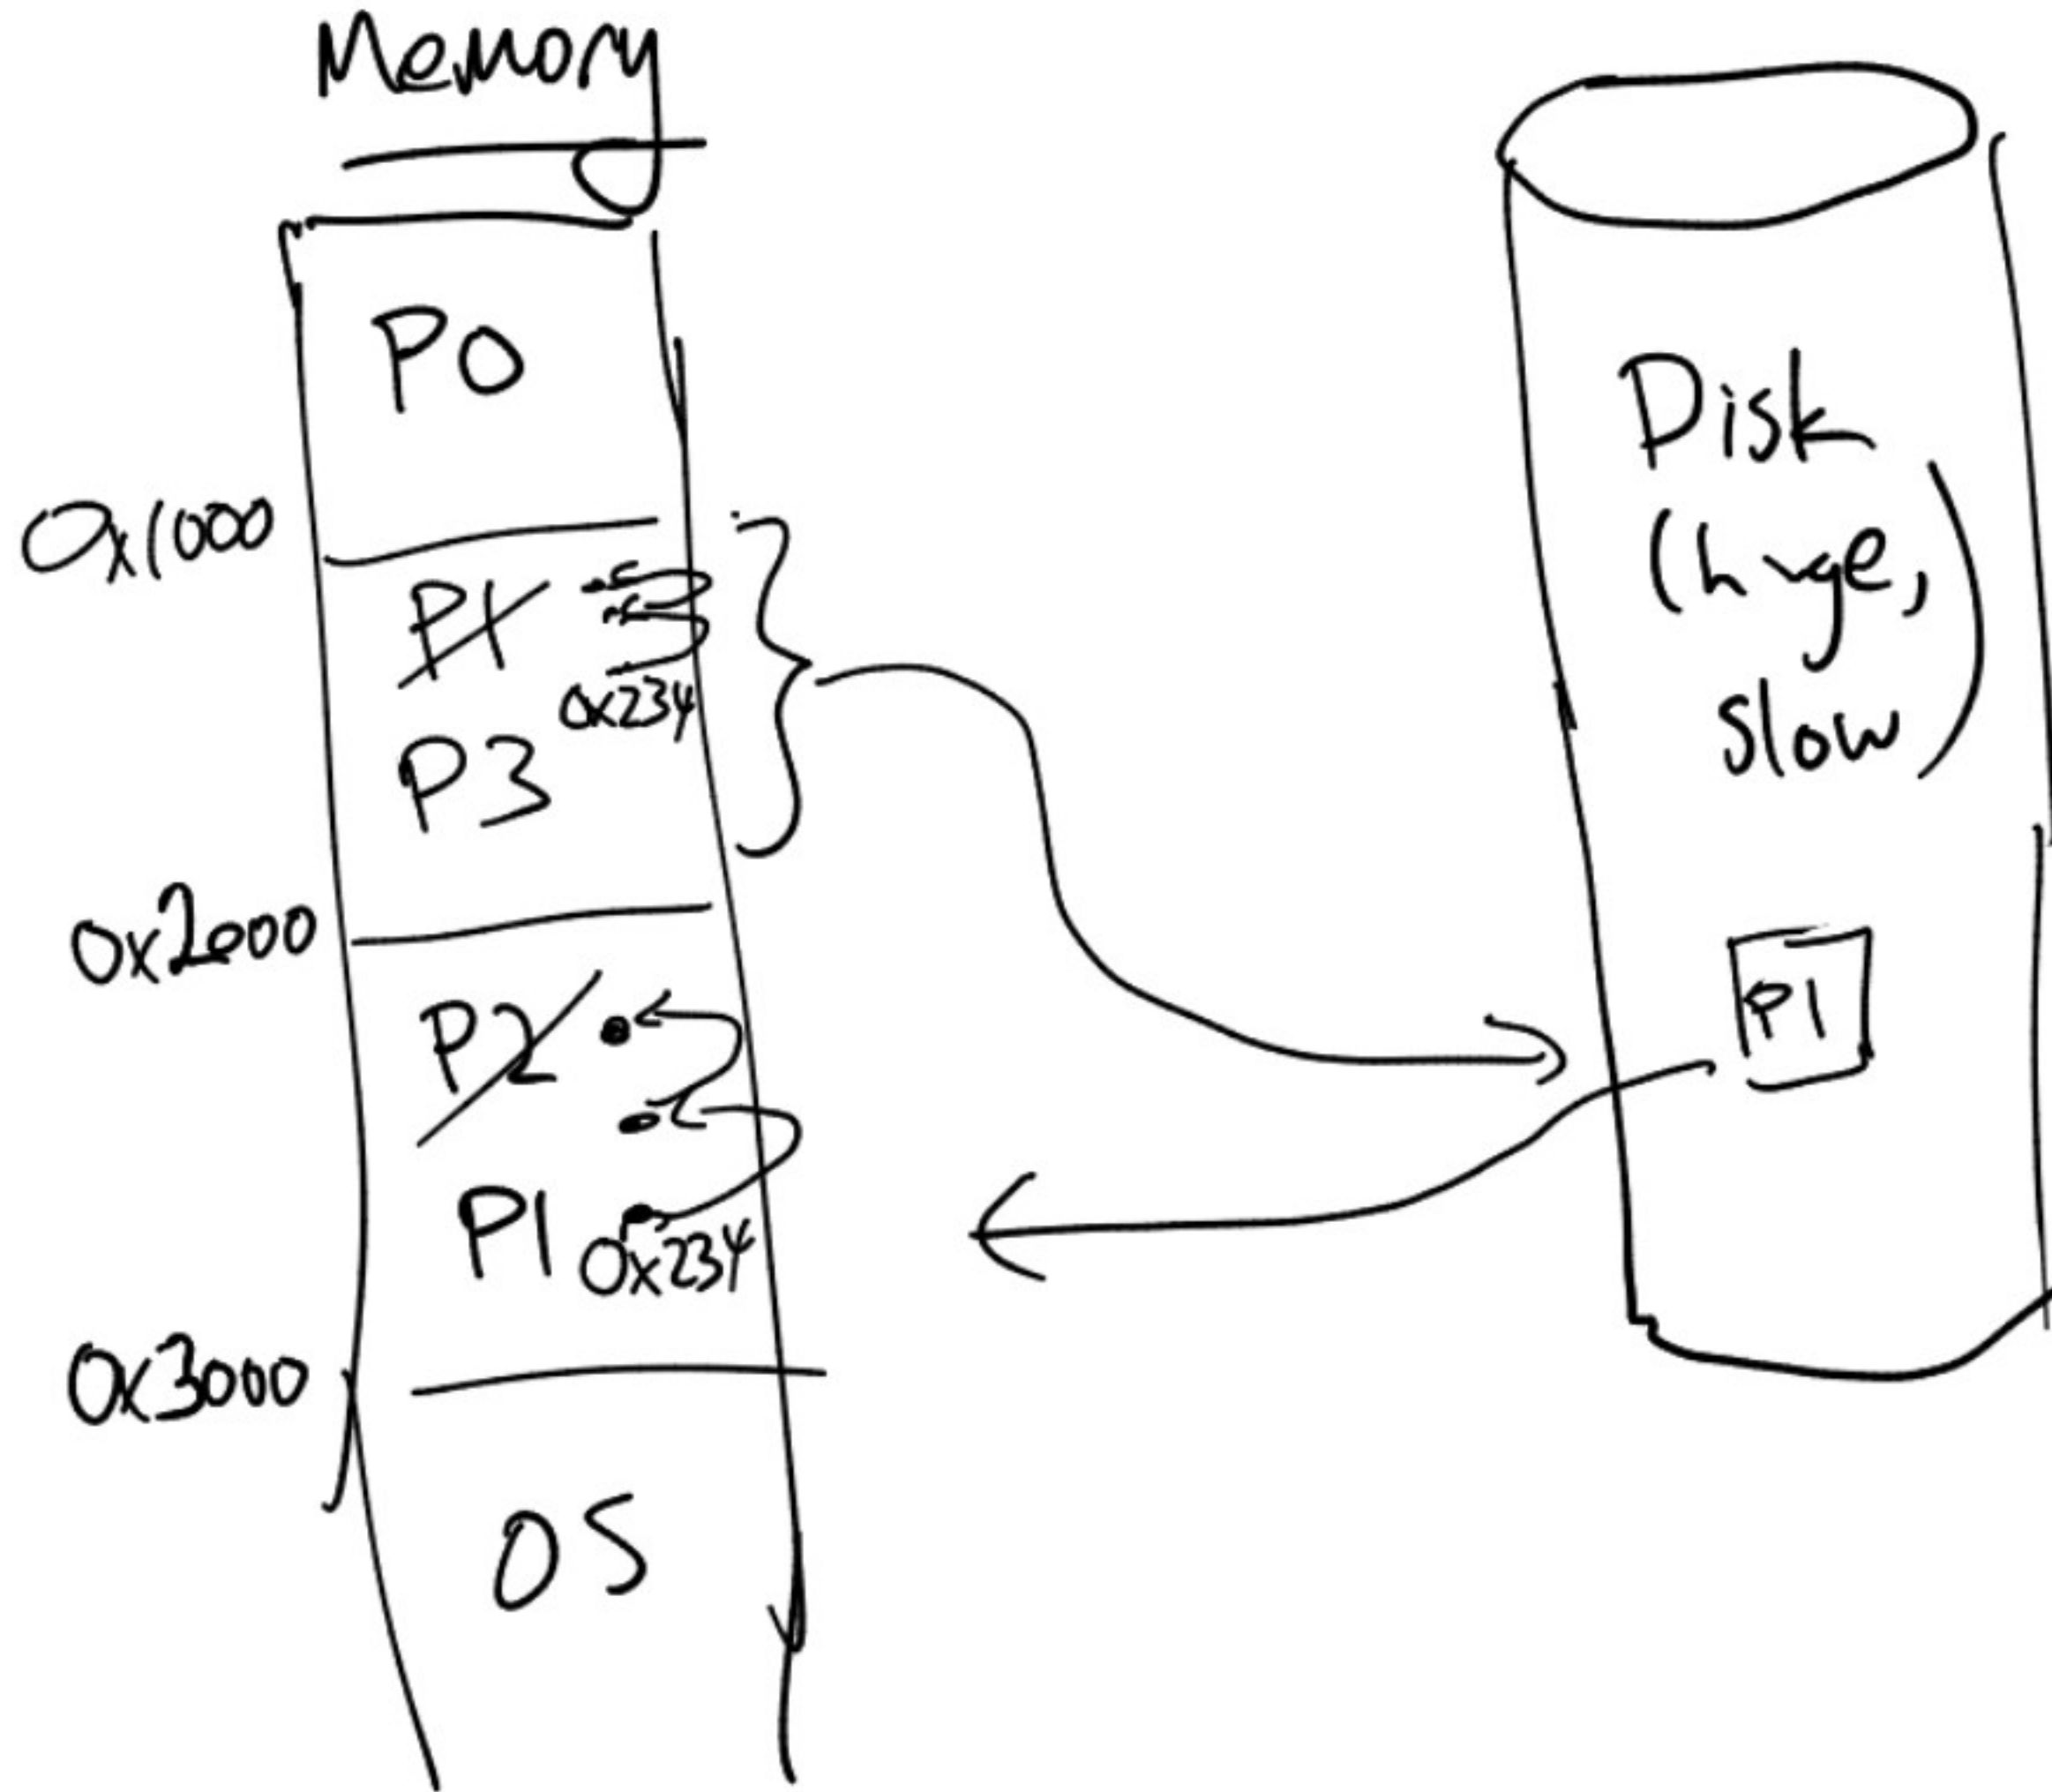
address translation

- processes work with "logical" addresses
ex: 0x0234

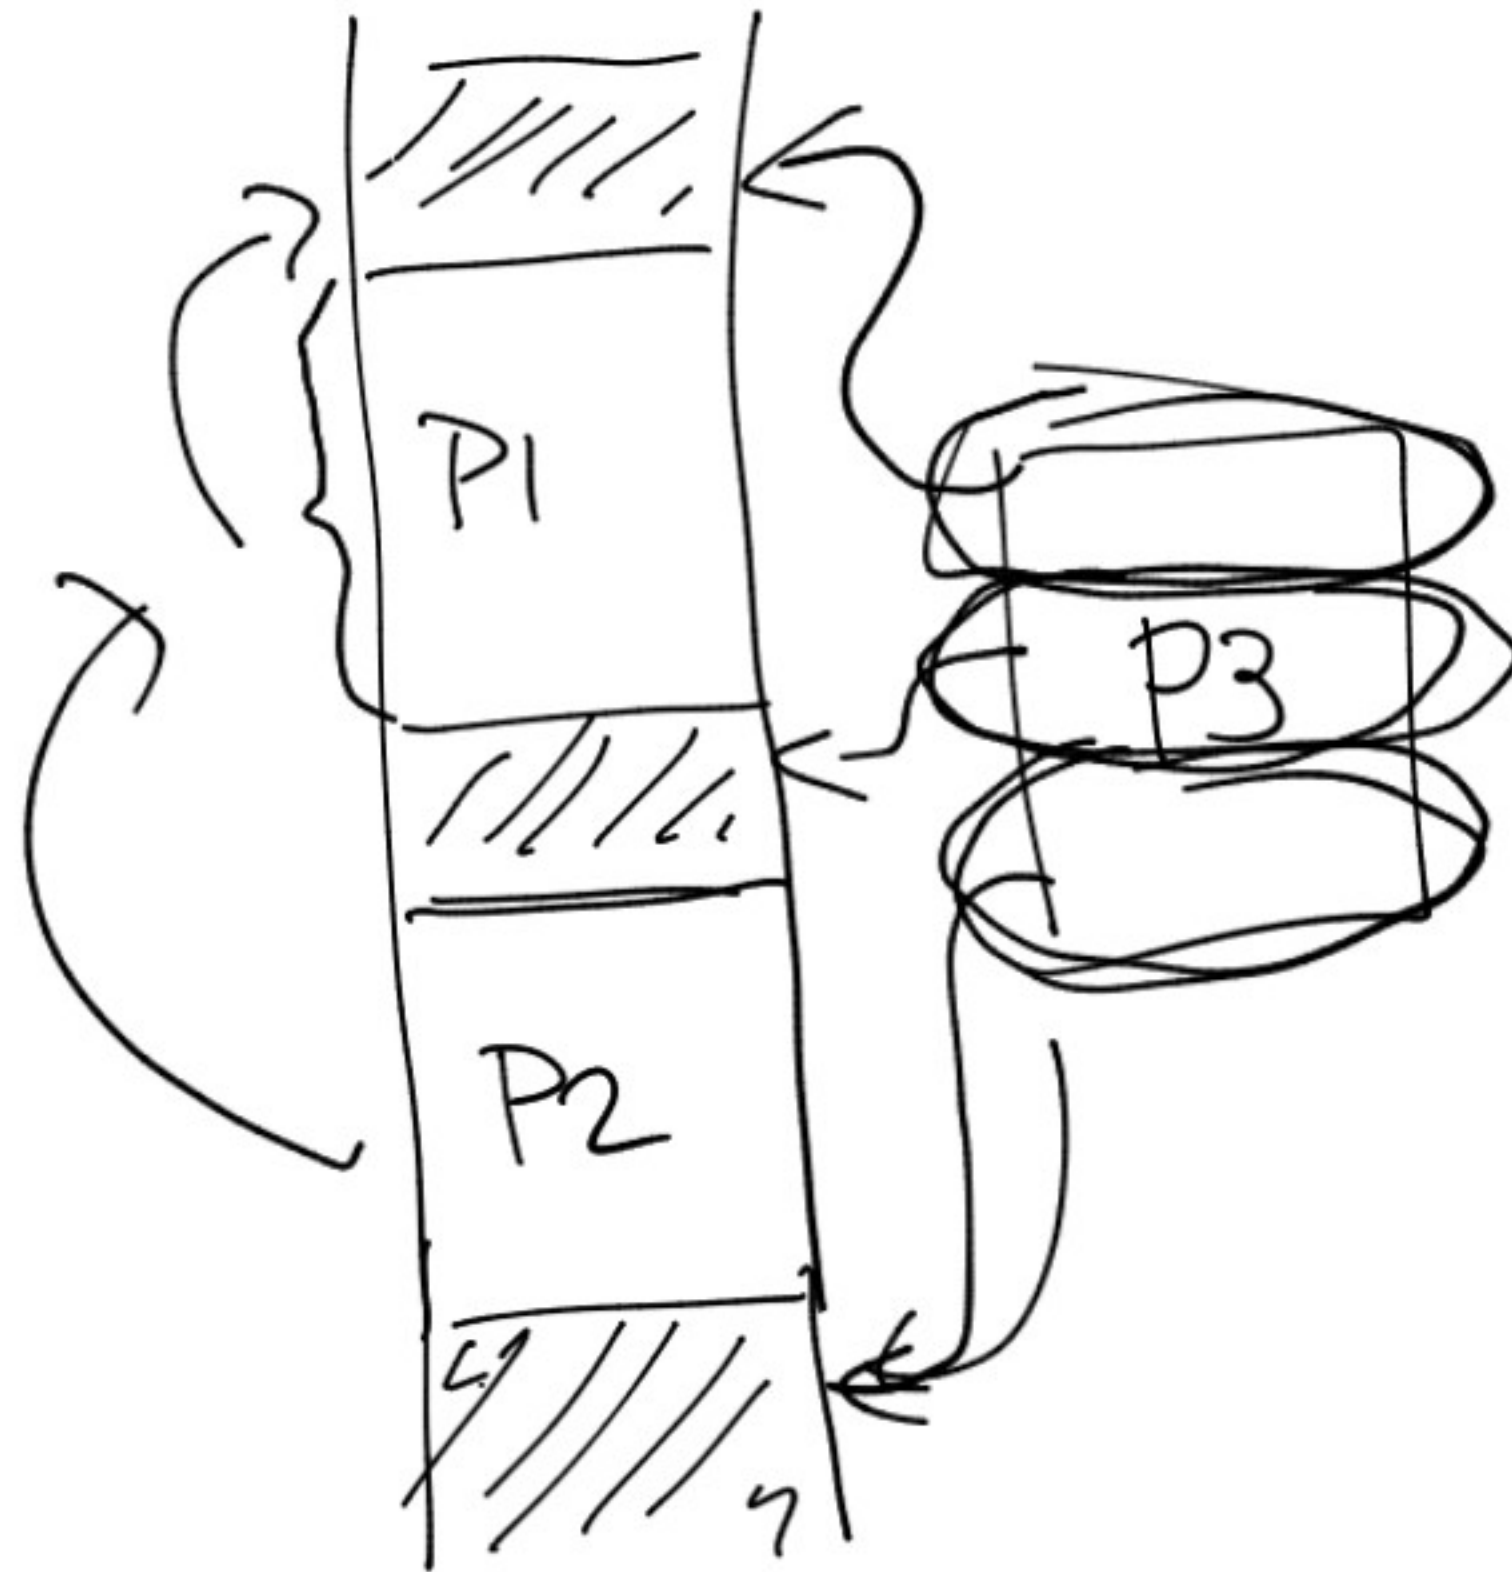
- memory management unit (MMU) translates logical address to physical addr.
e.x: 0x3234

Swapping

create P3?




Fragmentation



external fragmentation:

enough free space, but
not contiguous.

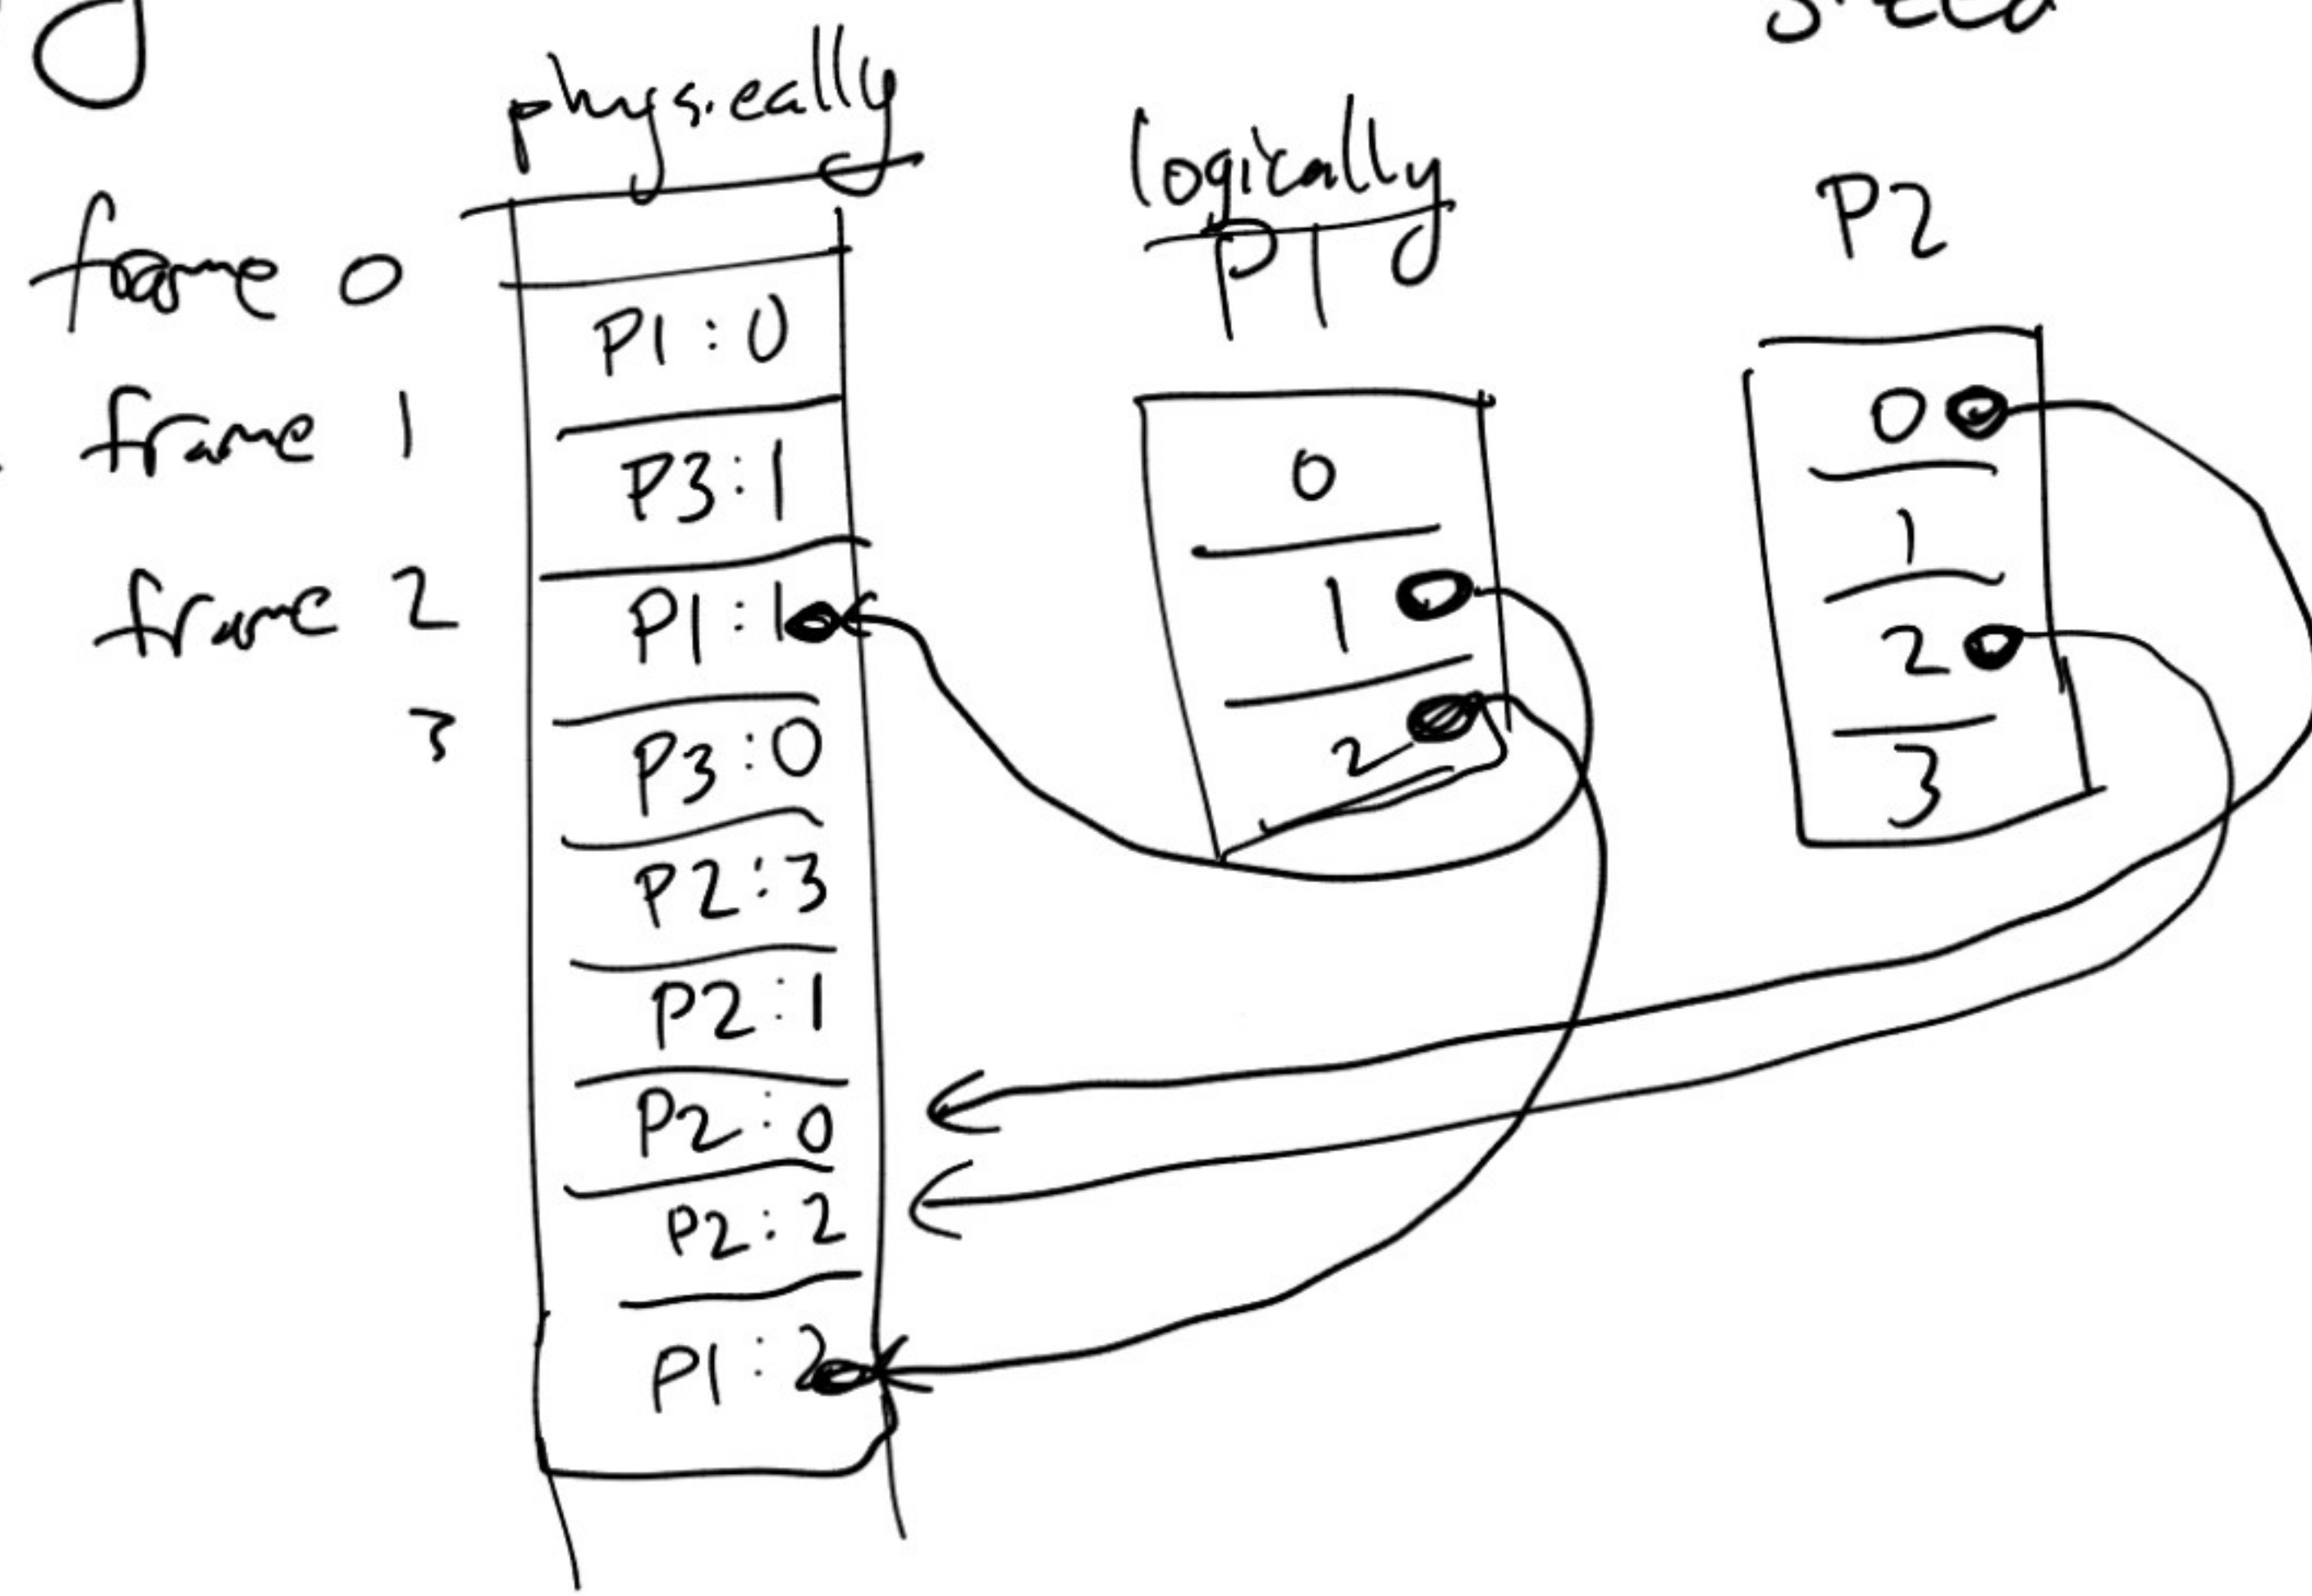
Could make all processes
same size: No external
frag.



Wasted
space
inside
procs.

internal fragmentation

Paging: split process's address space up into fixed-sized "pages", put pages into different "frames" of phys. memory

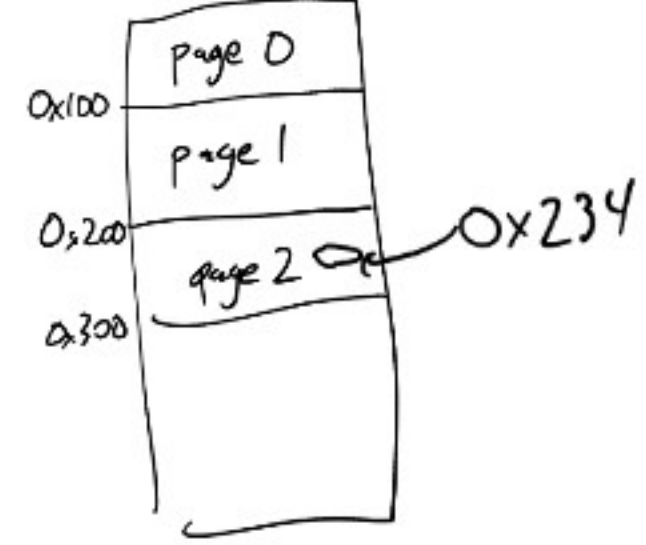



PI running

MMU for paging

logical address space

load
0x234



know which frame contains which page?

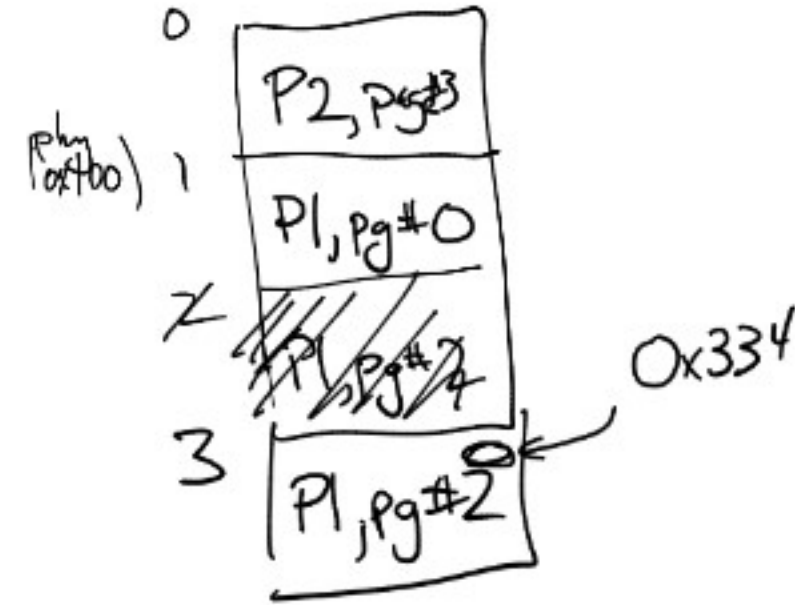
phys memory

TLB: translation lookaside buffer



page# | frame#

0
2
1

1
3
(nowhere)



log. address



physical address.

bits indicating that page is not present, raise HW exception if not

Segmentation

→ processes tell OS that certain regions of log. address space are protected (e.g. read only (constants) executable, not writable (code))

→ OS crashes process if program misbehaves

→ TLB has 3 permission bits for each entry (r, w, x), causes a HW exception on access to a page w/o permissions.

logical address space

