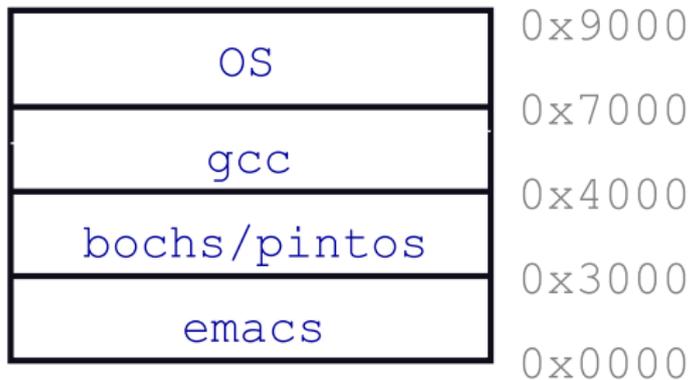


# Want processes to co-exist



- **Consider multiprogramming on physical memory**
  - What happens if pintos needs to expand?
  - If emacs needs more memory than is on the machine??
  - If pintos has an error and writes to address 0x7100?
  - When does gcc have to know it will run at 0x4000?
  - What if emacs isn't using its memory?

# Issues in sharing physical memory

- **Protection**

- A bug in one process can corrupt memory in another
- Must somehow prevent process *A* from trashing *B*'s memory
- Also prevent *A* from even observing *B*'s memory (ssh-agent)

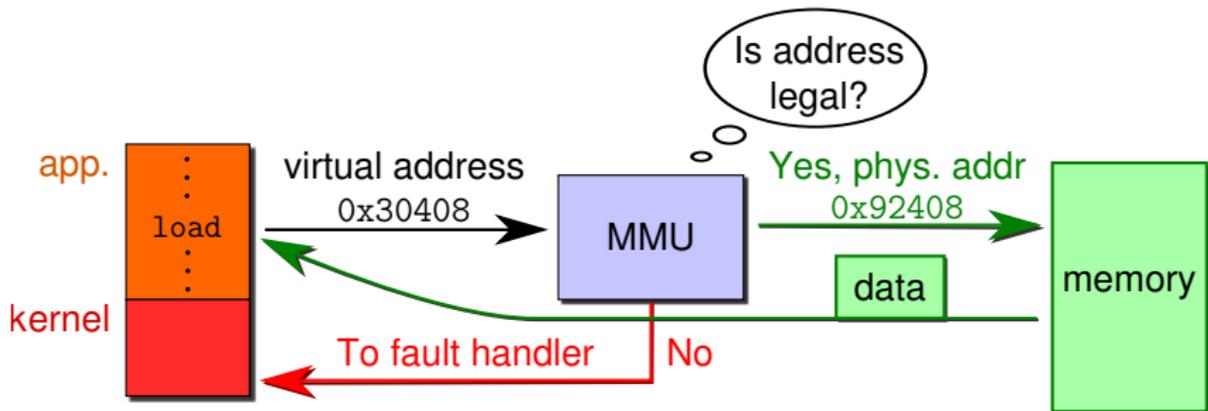
- **Transparency**

- A process shouldn't require particular physical memory bits
- Yes processes often require large amounts of contiguous memory (for stack, large data structures, etc.)

- **Resource exhaustion**

- Programmers typically assume machine has "enough" memory
- Sum of sizes of all processes often greater than physical memory

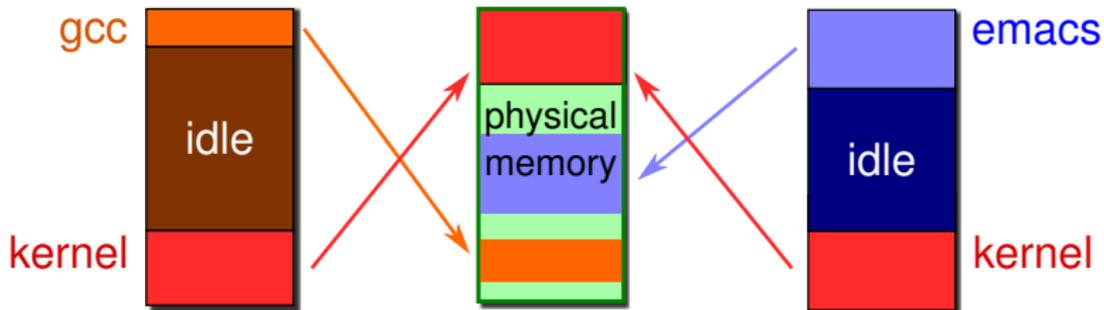
# Virtual memory goals



- **Give each program its own “virtual” address space**
  - At run time, Memory-Management Unit relocates each load, store to actual memory... App doesn't see physical memory
- **Also enforce protection**
  - Prevent one app from messing with another's memory
- **And allow programs to see more memory than exists**
  - Somehow relocate some memory accesses to disk

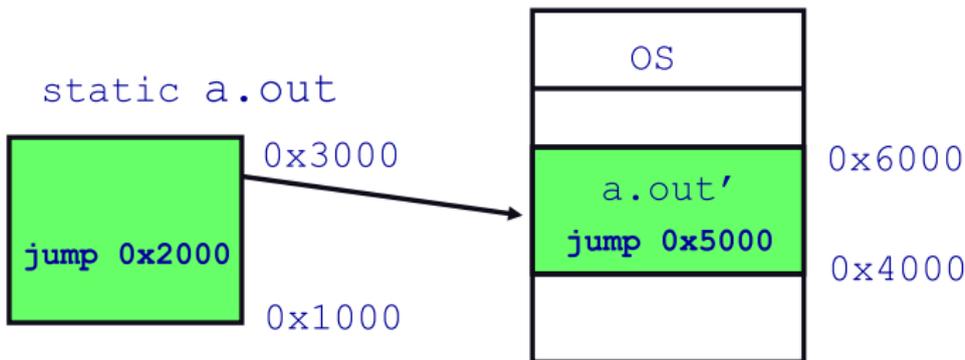
# Virtual memory advantages

- **Can re-locate program while running**
  - Run partially in memory, partially on disk
- **Most of a process's memory will be idle (80/20 rule).**



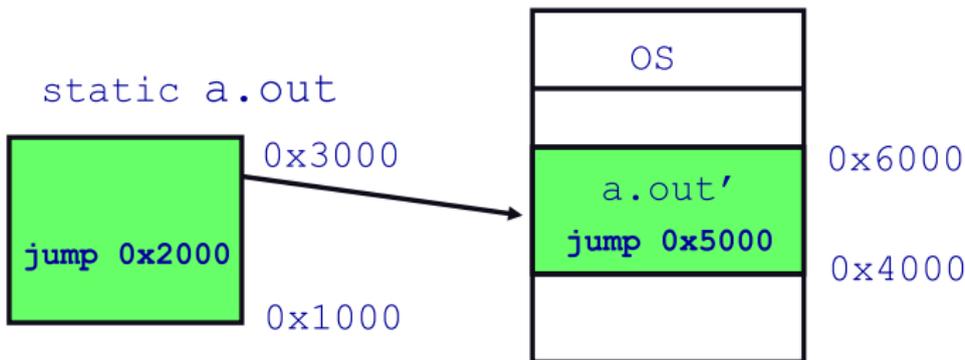
- Write idle parts to disk until needed
  - Let other processes use memory of idle part
  - Like CPU virtualization: when process not using CPU, switch (Not using a memory region? switch it to another process)
- **Challenge: VM = extra layer, could be slow**

# Idea 1: load-time linking



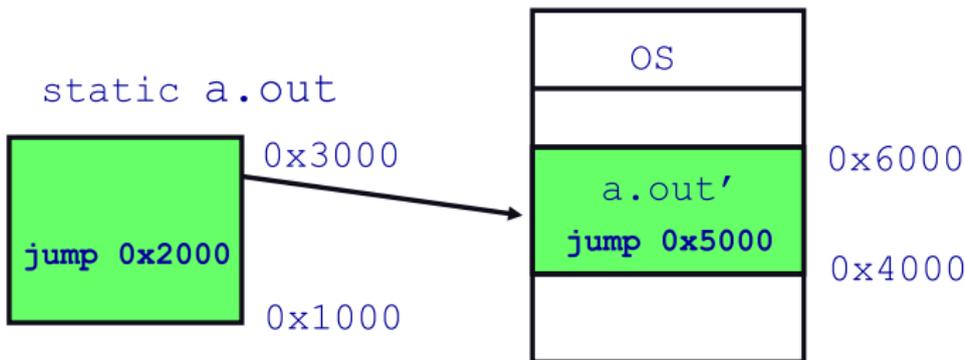
- *Linker* patches addresses of symbols like printf
- **Idea: link when process executed, not at compile time**
  - Determine where process will reside in memory
  - Adjust all references within program (using addition)
- **Problems?**

# Idea 1: load-time linking



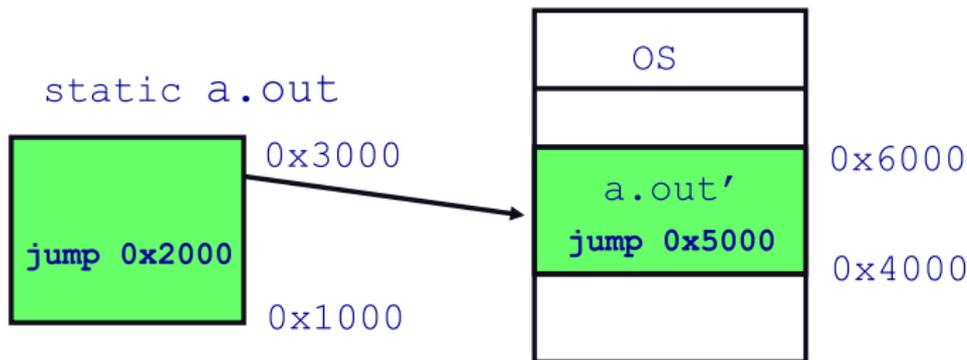
- *Linker* patches addresses of symbols like `printf`
- **Idea: link when process executed, not at compile time**
  - Determine where process will reside in memory
  - Adjust all references within program (using addition)
- **Problems?**
  - How to enforce protection
  - How to move once in memory (Consider: data pointers)
  - What if no contiguous free region fits program?

## Idea 2: base + bound register



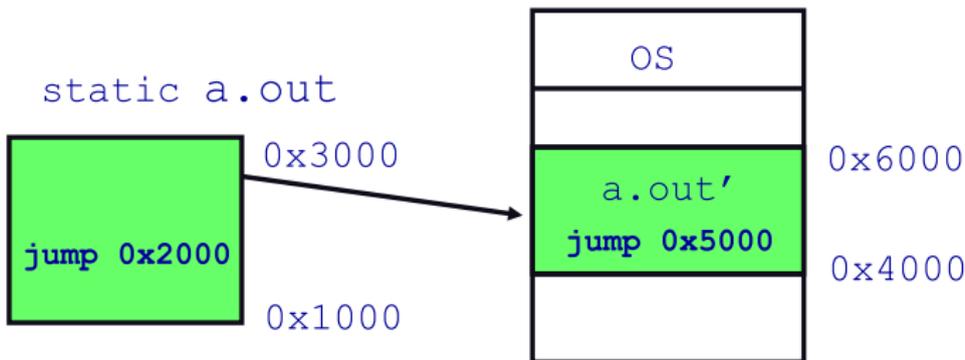
- Two special privileged registers: **base** and **bound**
- On each load/store:
  - Physical address = virtual address + **base**
  - Check  $0 \leq \text{virtual address} < \text{bound}$ , else trap to kernel
- How to move process in memory?
- What happens on context switch?

## Idea 2: base + bound register



- Two special privileged registers: **base** and **bound**
- On each load/store:
  - Physical address = virtual address + **base**
  - Check  $0 \leq \text{virtual address} < \text{bound}$ , else trap to kernel
- How to move process in memory?
  - Change **base** register
- What happens on context switch?

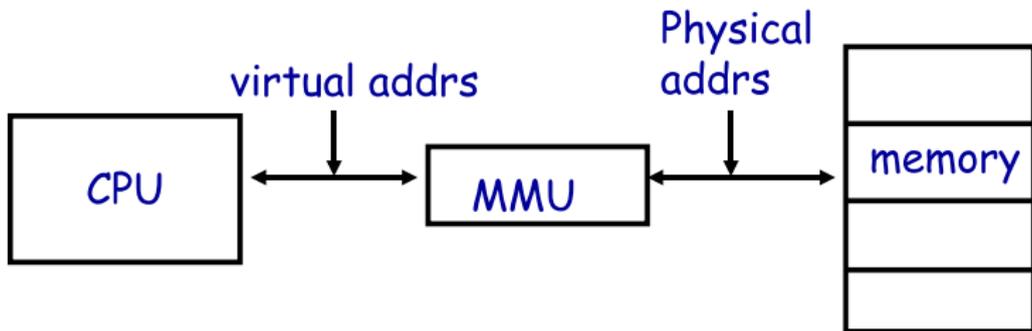
## Idea 2: base + bound register



- Two special privileged registers: **base** and **bound**
- On each load/store:
  - Physical address = virtual address + **base**
  - Check  $0 \leq \text{virtual address} < \text{bound}$ , else trap to kernel
- How to move process in memory?
  - Change **base** register
- What happens on context switch?
  - OS must re-load **base** and **bound** register

# Definitions

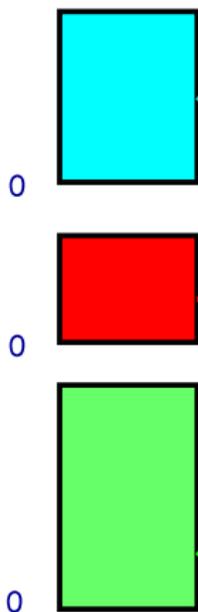
- Programs load/store to **virtual** (or **logical**) addresses
- Actual memory uses **physical** (or **real**) addresses
- VM Hardware is Memory Management Unit (**MMU**)



- Usually part of CPU
- Accessed w. privileged instructions (e.g., load bound reg)
- Translates from virtual to physical addresses
- Gives per-process view of memory called **address space**

# Address space

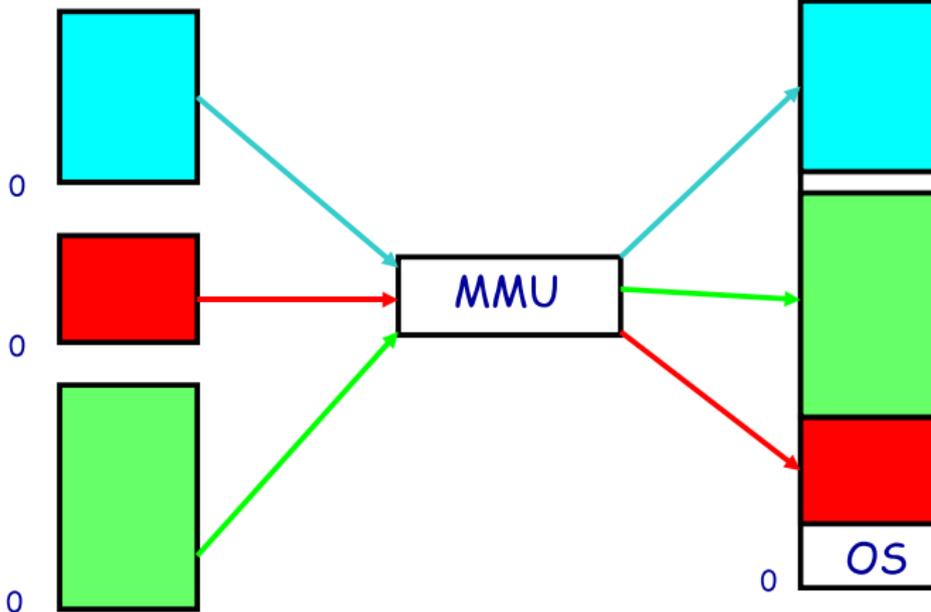
Virtual Address View



Physical Address View



MMU



# Base+bound trade-offs

- **Advantages**

- Cheap in terms of hardware: only two registers
- Cheap in terms of cycles: do add and compare in parallel
- Examples: Cray-1 used this scheme

- **Disadvantages**

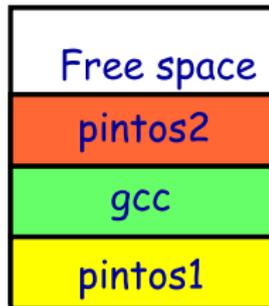
# Base+bound trade-offs

- **Advantages**

- Cheap in terms of hardware: only two registers
- Cheap in terms of cycles: do add and compare in parallel
- Examples: Cray-1 used this scheme

- **Disadvantages**

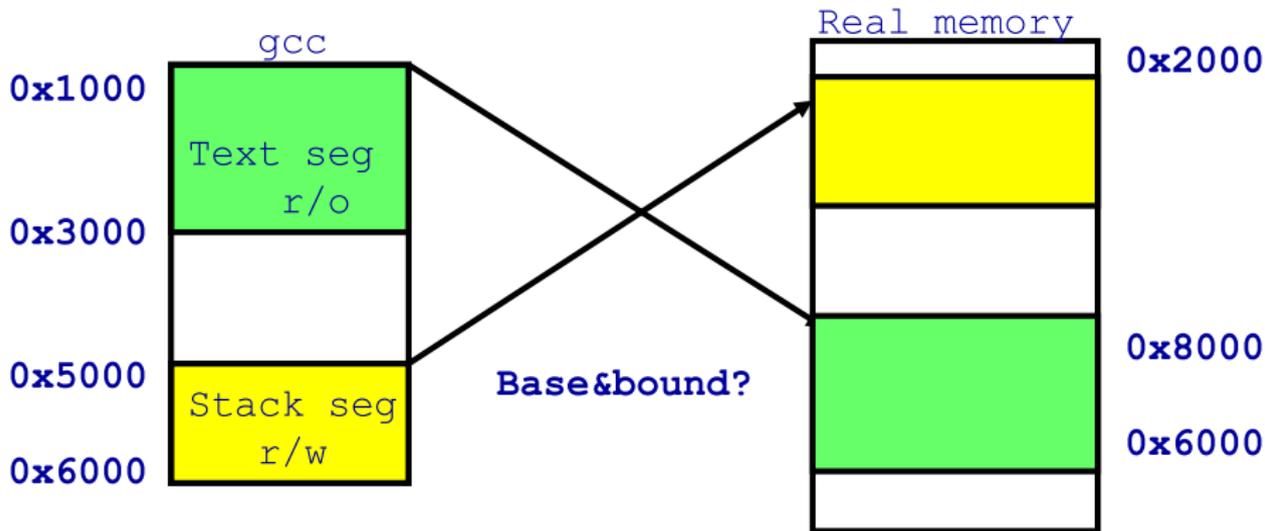
- Growing a process is expensive or impossible
- No way to share code or data (E.g., two copies of bochs, both running pintos)



- **One solution: Multiple segments**

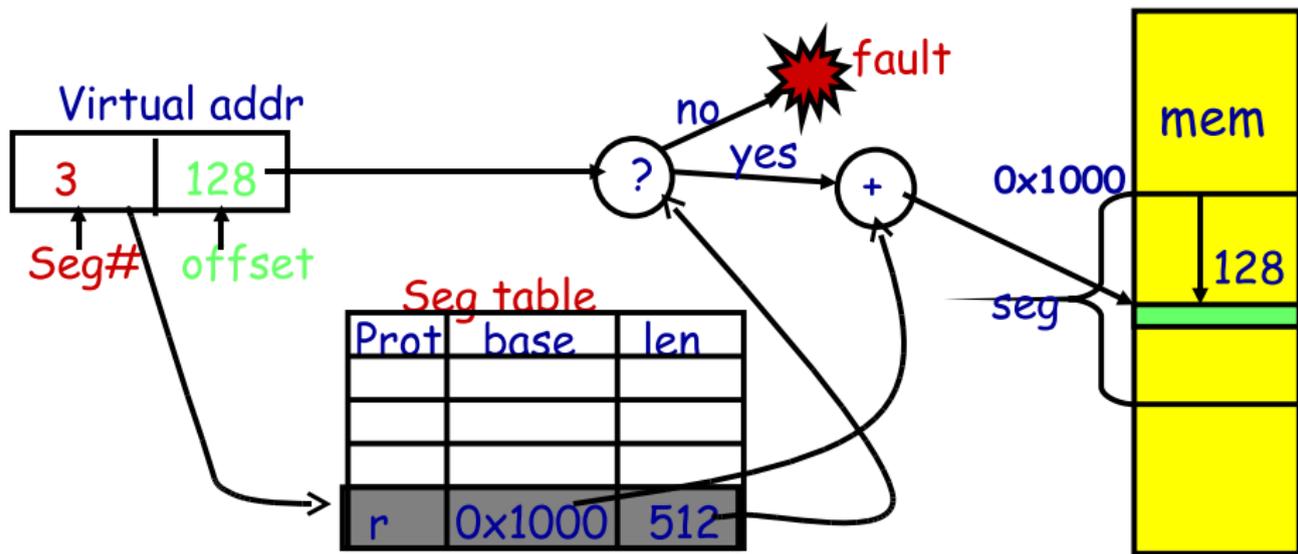
- E.g., separate code, stack, data segments
- Possibly multiple data segments

# Segmentation



- **Let processes have many base/bound regs**
  - Address space built from many segments
  - Can share/protect memory at segment granularity
- **Must specify segment as part of virtual address**

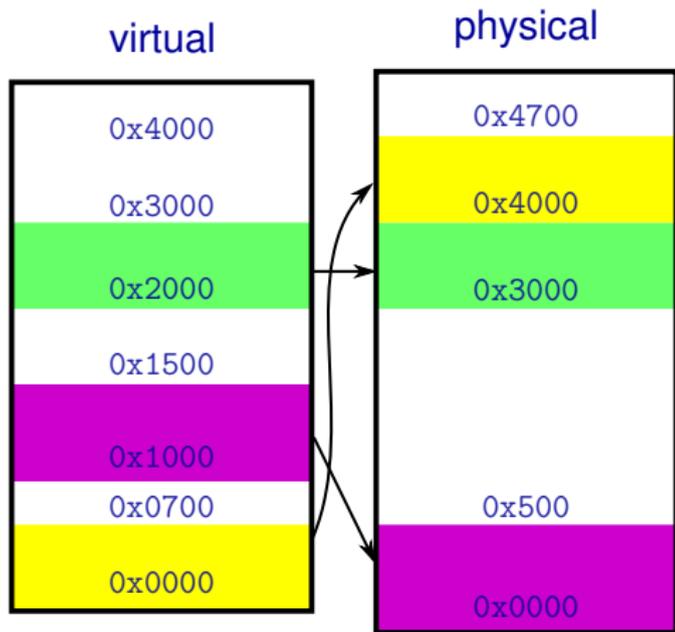
# Segmentation mechanics



- Each process has a segment table
- Each VA indicates a segment and offset:
  - Top bits of addr select segment, low bits select offset (PDP-10)
  - Or segment selected by instruction or operand (means you need wider "far" pointers to specify segment)

# Segmentation example

Seg	base	bounds	rw
0	0x4000	0x6fff	10
1	0x0000	0x4fff	11
2	0x3000	0xffff	11
3			00

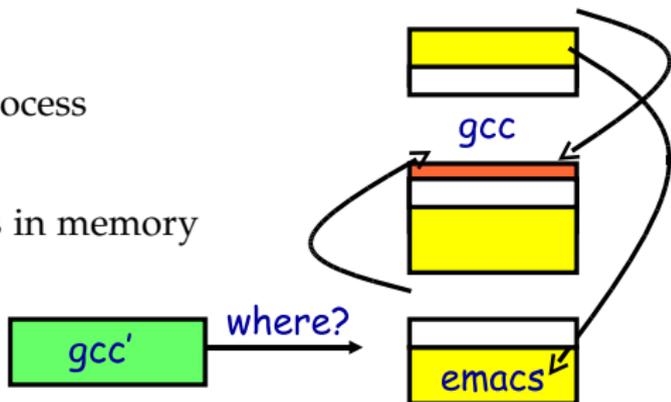


- **2-bit segment number (1st digit), 12 bit offset (last 3)**
  - Where is 0x0240? 0x1108? 0x265c? 0x3002? 0x1600?

# Segmentation trade-offs

- **Advantages**

- Multiple segments per process
- Allows sharing! (how?)
- Don't need entire process in memory



- **Disadvantages**

- Requires translation hardware, which could limit performance
- Segments not completely transparent to program (e.g., default segment faster or uses shorter instruction)
- $n$  byte segment needs  $n$  contiguous bytes of physical memory
- Makes *fragmentation* a real problem.

# Fragmentation

- **Fragmentation**  $\implies$  Inability to use free memory
- **Over time:**
  - Variable-sized pieces = many small holes (external fragmentation)
  - Fixed-sized pieces = no external holes, but force internal waste (internal fragmentation)

