VIRTUALIZATION: CPU TO MEMORY

Shivaram Venkataraman

CS 537, Spring 2020

ADMINISTRIVIA

- Project 1a: DONE!?
- How to use slip days? (Piazza)

- Project 1b is out, due Feb 5th (Next Wednesday)
- Discussion section
  - xv6 code walk through
  - How to use gdb

AGENDA / LEARNING OUTCOMES

CPU virtualization

- Recap of scheduling policies
- Work through problems

Memory virtualization

- What is the need for memory virtualization?
- How to virtualize memory?

RECAP: CPU VIRTUALIZATION

RECAP: SCHEDULING MECHANISM

Process: Abstraction to virtualize CPU

Use time-sharing in OS to switch between processes

Limited Direct Execution

- Use system calls to run access devices etc. from user mode
- Context-switch using interrupts for multi-tasking

RECAP: METRICS → POLICIES

Turnaround time = *completion_time - arrival_time*

FIFO: First come, first served

SJF: Shortest job first

SCTF: Shortest completion time first

RECAP: METRICS → POLICIES

Response time = *first_run_time - arrival_time*

RR: Round robin with time slice

Minimizes response time but could increase turnaround?

MULTI-LEVEL FEEDBACK QUEUE

MLFQ EXAMPLE

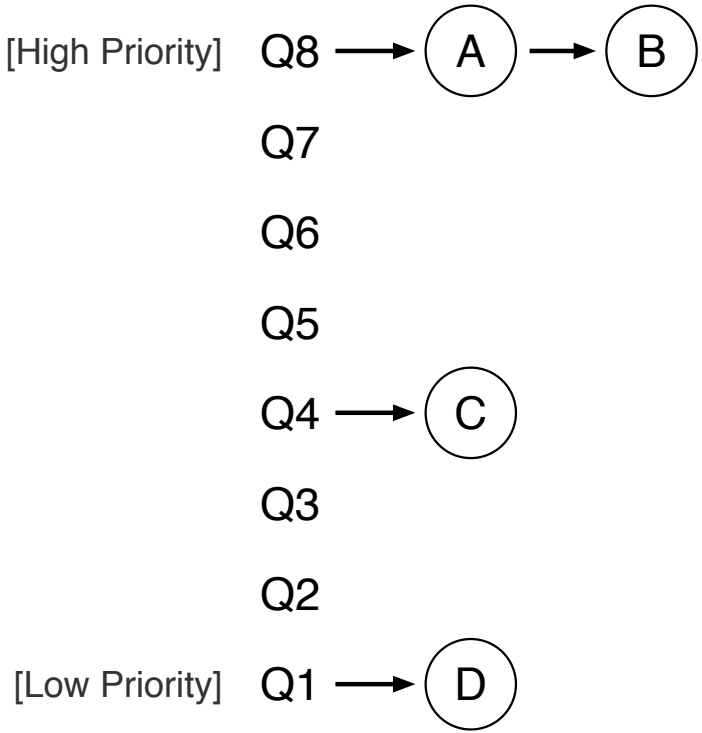

Rules for MLFQ

Rule 1: If priority(A) > Priority(B)
A runs

Rule 2: If priority(A) == Priority(B),
A & B run in RR

Rule 3: Processes start at top priority

Rule 4: If job uses whole slice, demote process.
If not stay at level

ONE LONG JOB

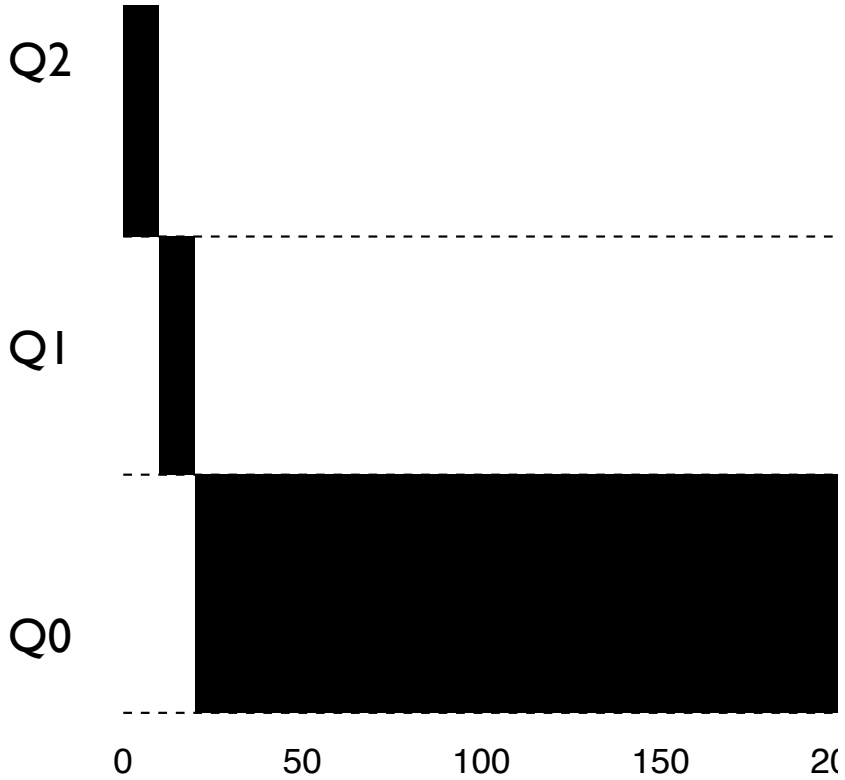

INTERACTIVE PROCESS JOINS

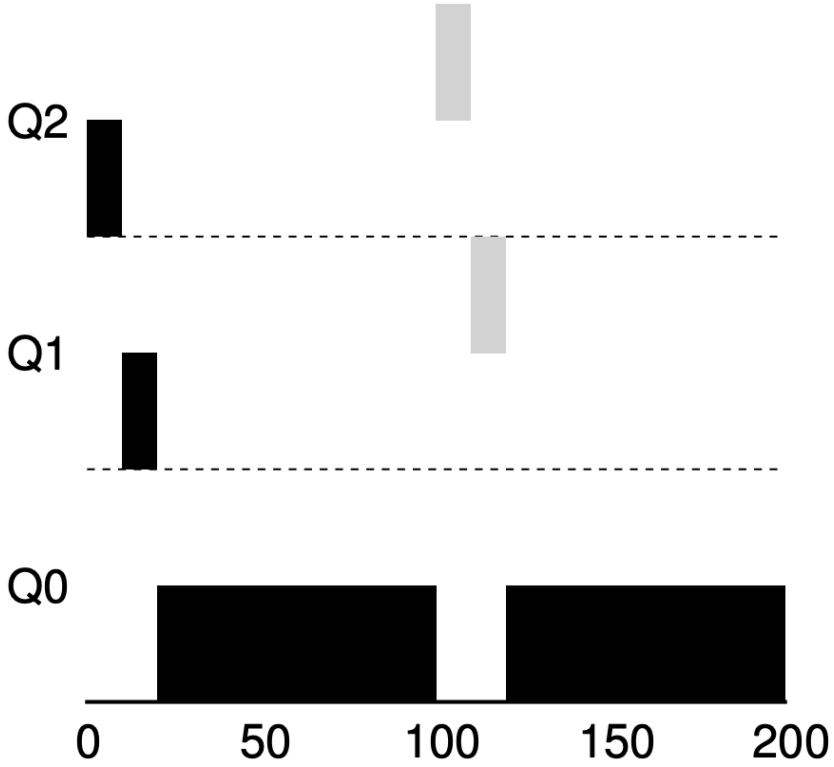

MLFQ PROBLEMS?

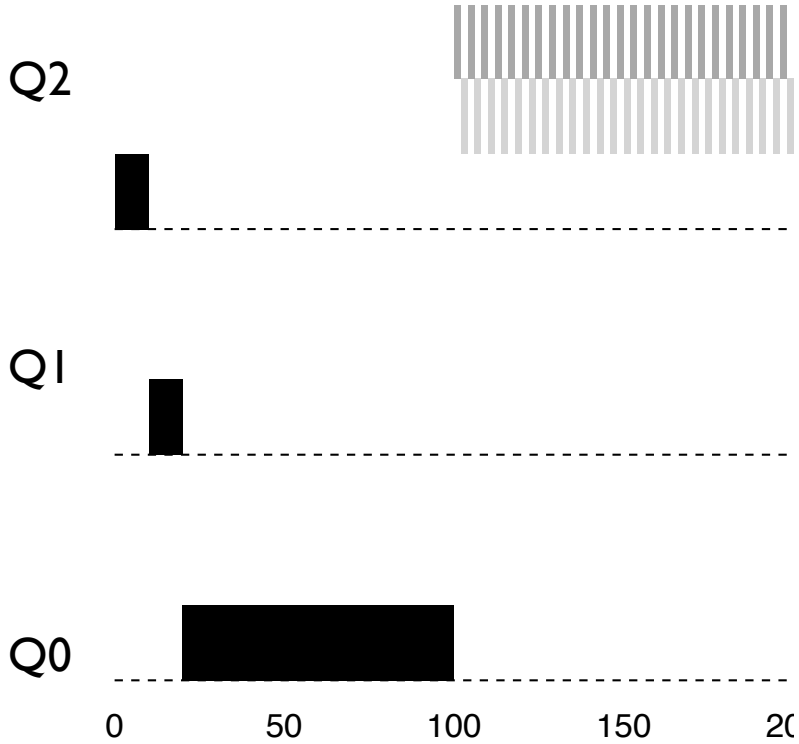

What is the problem with this schedule ?

AVOIDING STARVATION

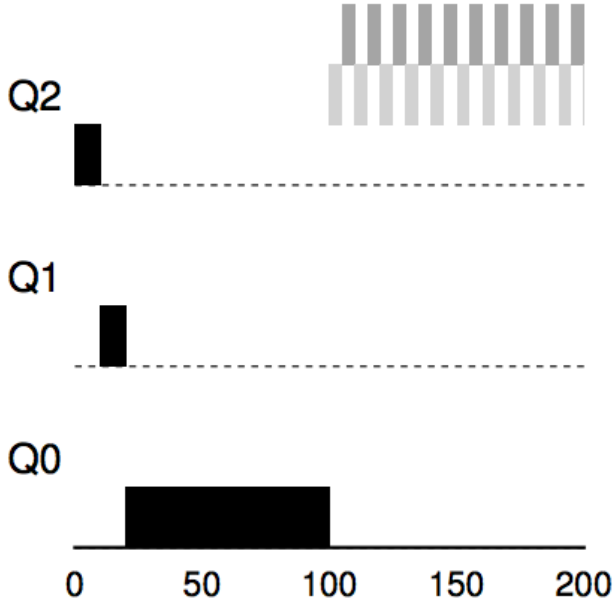

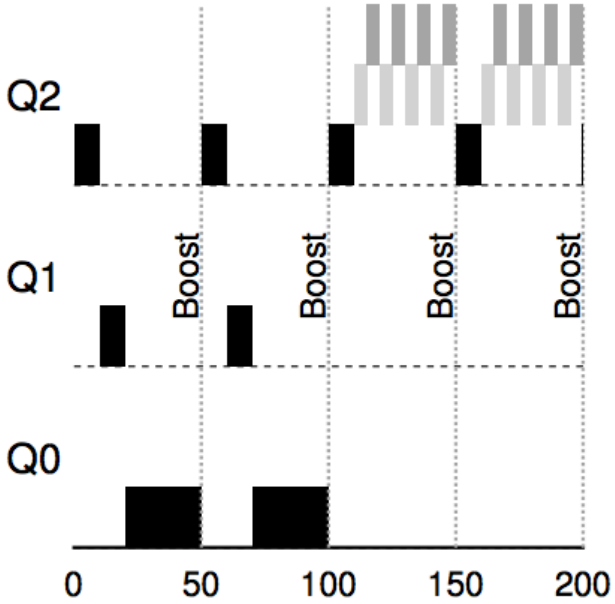

Rule 5: After some time period S, move all the jobs in the system to the topmost queue.

GAMING THE SCHEDULER ?

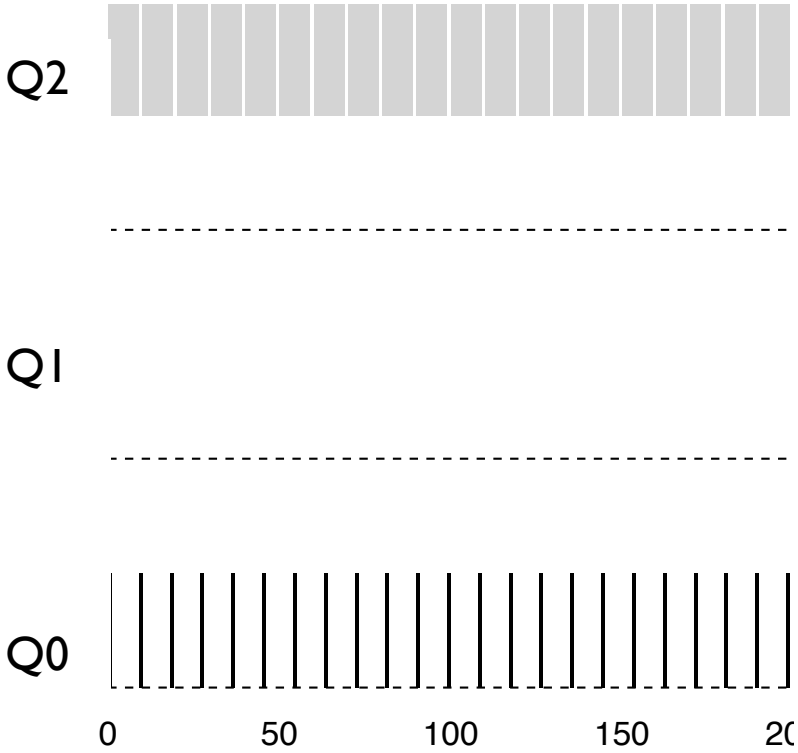

Job could trick scheduler by doing I/O just before time-slice end

GAMING THE SCHEDULER ?

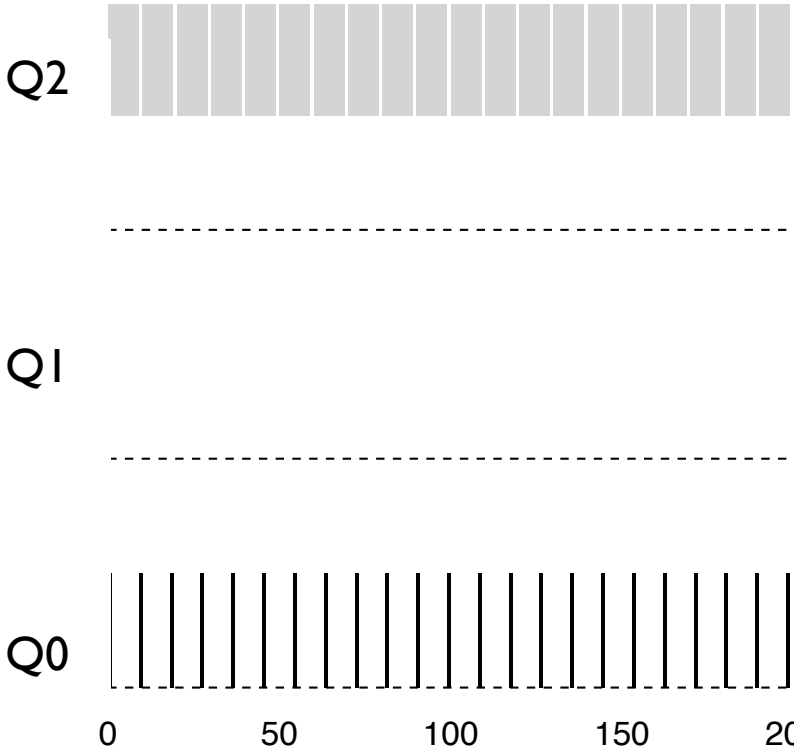

Job could trick scheduler by doing I/O just before time-slice end

Rule 4*: Once a job uses up its time allotment at a given level (regardless of how many times it has given up the CPU), its priority is reduced

QUIZ 5

https://tinyurl.com/cs537-sp20-quiz5

Jobs	Runtime	Arrival Time
Job A	100	0
Job B	10	50

Jobs	Runtime	Arrival Time
Job A	100	0
Job B	10	50
Job C	20	70

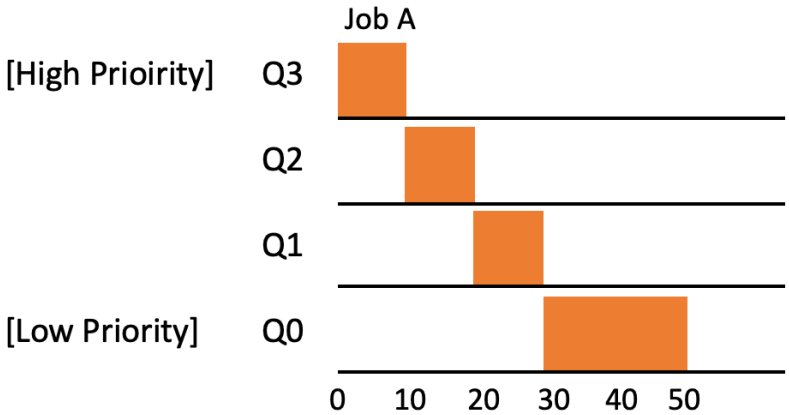

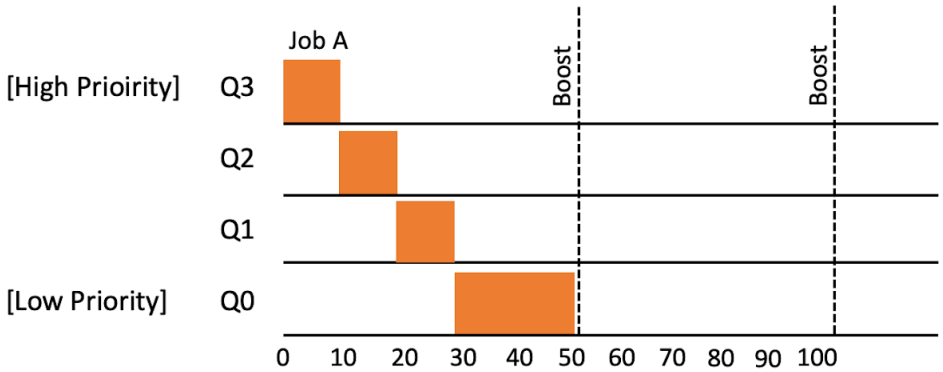

CPU SUMMARY

Mechanism

- Process abstraction
- System call for protection
- Context switch to time-share

Policy

- Metrics: turnaround time, response time
- Balance using MLFQ

VIRTUALIZING MEMORY

BACK IN THE DAY...

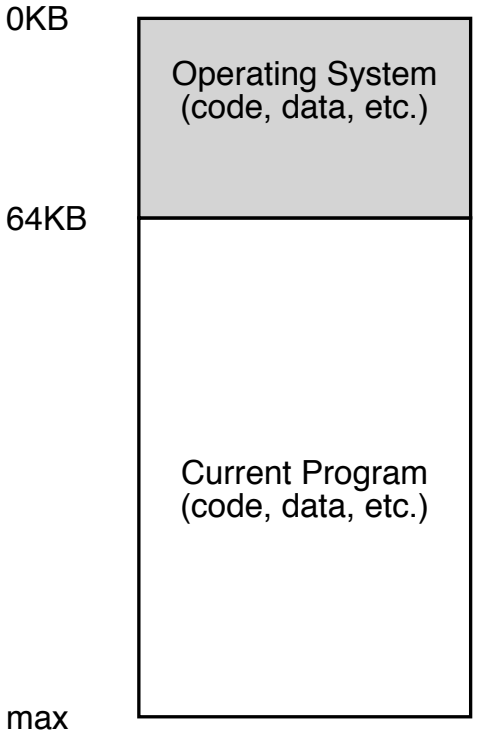

Uniprogramming: One process runs at a time

MULTIPROGRAMMING GOALS

Transparency: Process is unaware of sharing

Protection: Cannot corrupt OS or other process memory

Efficiency: Do not waste memory or slow down processes

Sharing: Enable sharing between cooperating processes

ABSTRACTION: ADDRESS SPACE

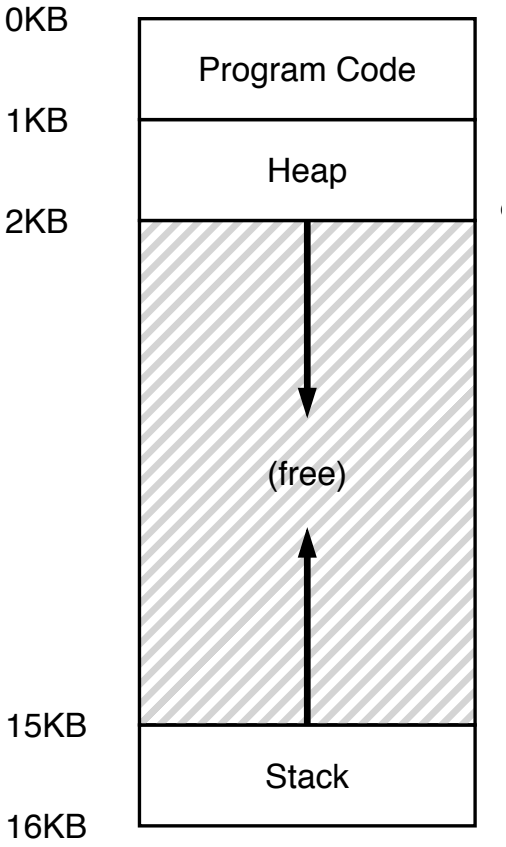

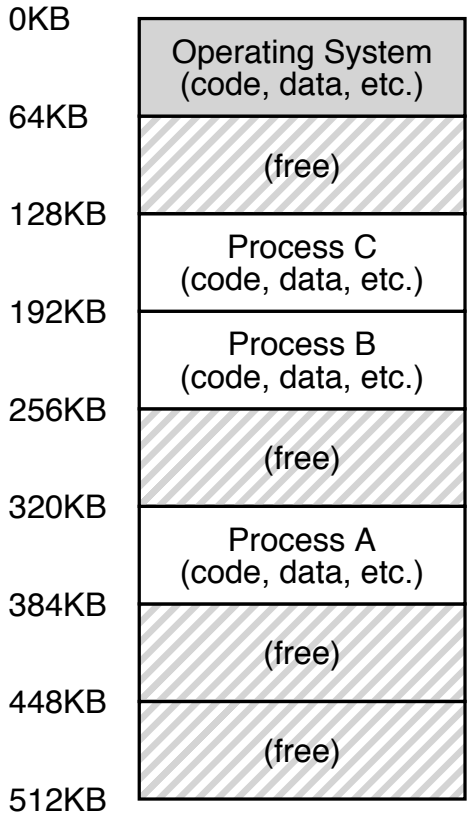

WHAT IS IN ADDRESS SPACE?

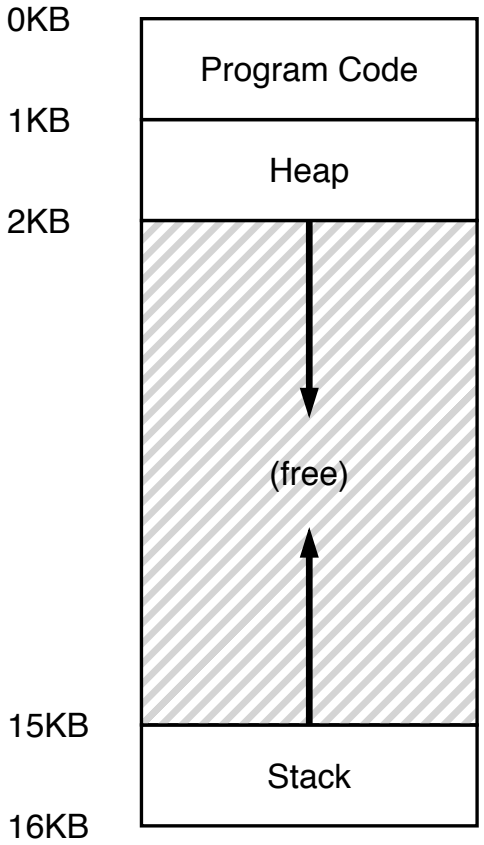

the code segment:
where instructions live

the heap segment:
contains malloc'd data
dynamic data structures
(it grows downward)

(it grows upward)
the stack segment:
contains local variables
arguments to routines,
return values, etc.

Static: Code and some global variables

Dynamic: Stack and Heap

STACK ORGANIZATION

```
alloc(A);
alloc(B);
alloc(C);
free(C);
alloc(D);
free(D);
free(B);
free(A);
```

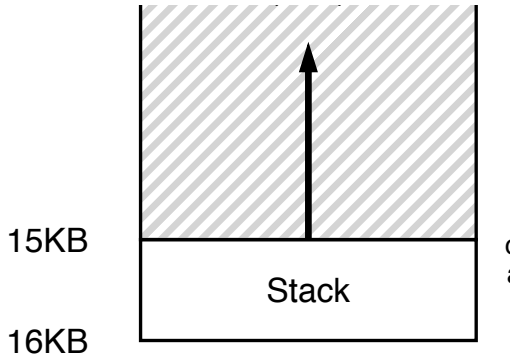

Pointer between allocated and free space

- Allocate: Increment pointer
- Free: Decrement pointer

No fragmentation!

WHAT GOES ON STACK?

```
main () {
    int A = 0;
    foo(A);
    printf(“A: %d\n”, A);
}
void foo (int Z) {
    int A = 2;
    Z = 5;
    printf(“A: %d Z: %d\n”, A, Z);
}
```

HEAP ORGANIZATION

Allocate from any random location: malloc(), new() etc.

- Heap memory consists of allocated and free areas (holes)
- Order of allocation and free is unpredictable

MEMORY ACCESS

```
#include <stdio.h>
#include <stdlib.h>

int main(int argc, char *argv[]) {
  int x;
  x = x + 3;
}
```

```
0x10: movl  0x8(%rbp), %edi
0x13: addl  $0x3, %edi
0x19: movl  %edi, 0x8(%rbp)
```

%**rbp** is the base pointer:
points to base of current stack frame

MEMORY ACCESS

Initial %rip = 0x10
%rbp = 0x200

```
0x10: movl  0x8(%rbp), %edi
0x13: addl  $0x3, %edi
0x19: movl  %edi, 0x8(%rbp)
```

%**rbp** is the base pointer:
points to base of current stack frame

%rip is instruction pointer (or program counter)

MEMORY ACCESS

Initial %rip = 0x10
%rbp = 0x200

```
0x10: movl  0x8(%rbp), %edi
0x13: addl  $0x3, %edi
0x19: movl  %edi, 0x8(%rbp)
```

%rbp is the base pointer:
points to base of current stack frame

%rip is instruction pointer (or program counter)

Fetch instruction at addr 0x10
Exec:
load from addr 0x208

Fetch instruction at addr 0x13
Exec:
no memory access

Fetch instruction at addr 0x19
Exec:
store to addr 0x208

QUIZ 6

https://tinyurl.com/cs537-sp20-quiz6

```
int x;
int main(int argc, char *argv[]) {
  int y;
  int* z = malloc(sizeof(int)););
}
```

Possible locations:
static data/code, stack, heap

Address	Location
x	
main	
y	
z	
*z	

HOW TO VIRTUALIZE MEMORY

Problem: How to run multiple processes simultaneously?

Addresses are "hardcoded" into process binaries

How to avoid collisions?

Possible Solutions for Mechanisms (covered today):

1. Time Sharing
2. Static Relocation
3. Base
4. Base+Bounds

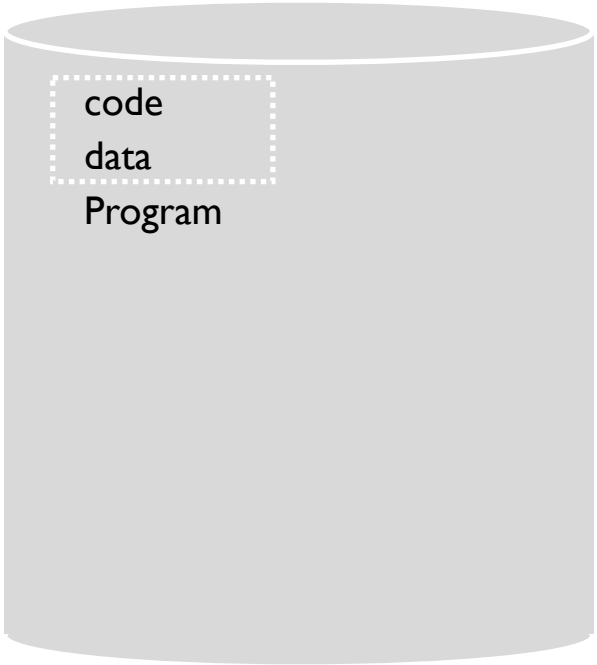

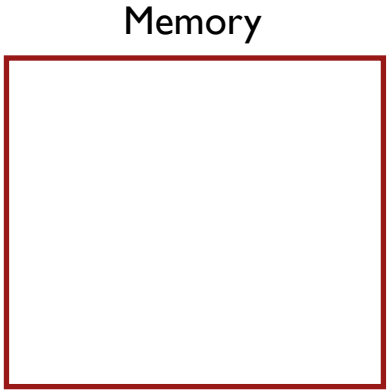

TIME SHARE MEMORY: EXAMPLE

PROBLEMS WITH TIME SHARING?

Ridiculously poor performance

Better Alternative: space sharing!

- At same time, space of memory is divided across processes
- Remainder of solutions all use space sharing

2) STATIC RELOCATION

Idea: OS rewrites each program before loading it as a process in memory

Each rewrite for different process uses different addresses and pointers

Change jumps, loads of static data

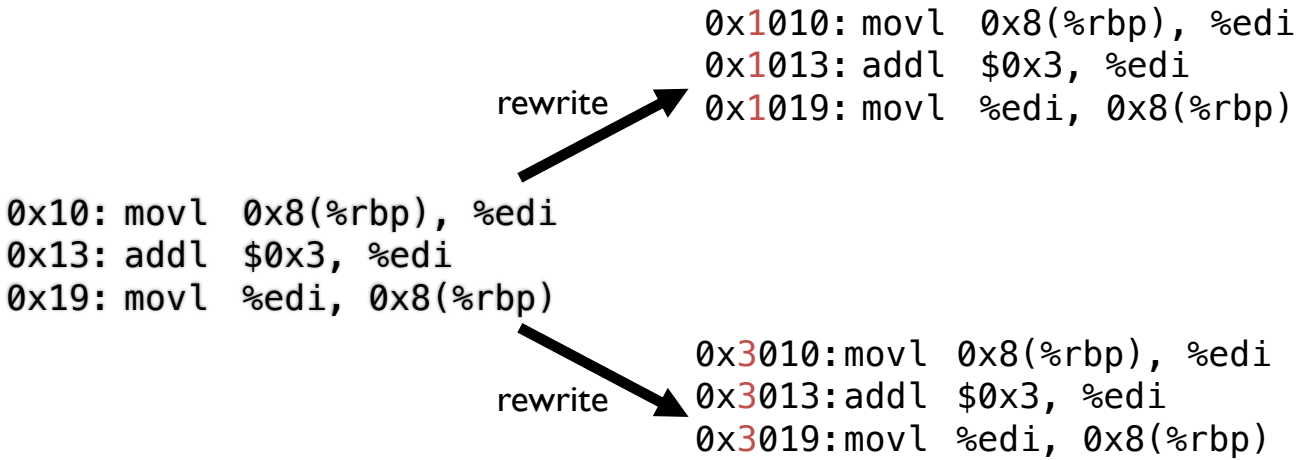

STATIC: LAYOUT IN MEMORY

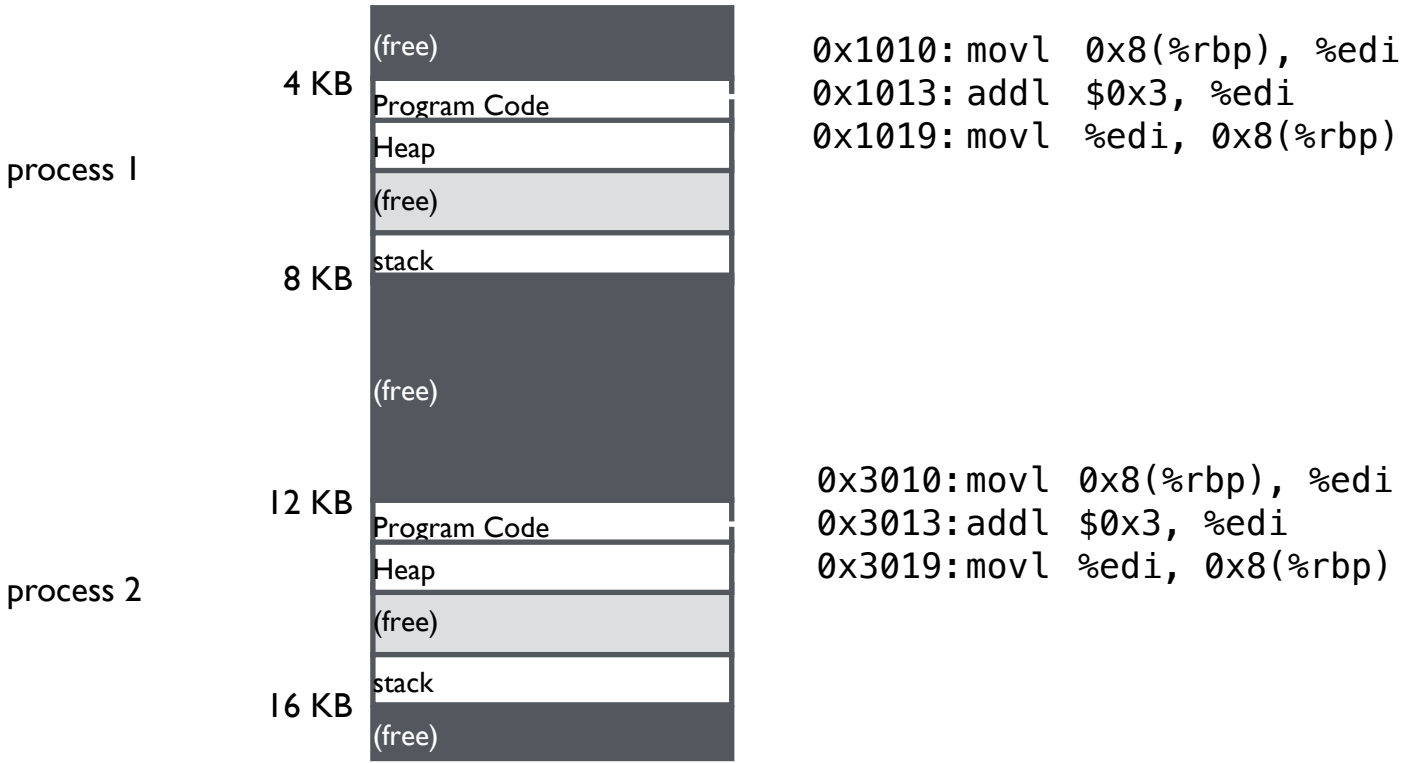


```
0x1010: movl  0x8(%rbp), %edi
0x1013: addl  $0x3, %edi
0x1019: movl  %edi, 0x8(%rbp)
```

```
0x3010: movl  0x8(%rbp), %edi
0x3013: addl  $0x3, %edi
0x3019: movl  %edi, 0x8(%rbp)
```

STATIC RELOCATION: DISADVANTAGES

No protection

- Process can destroy OS or other processes
- No privacy

Cannot move address space after it has been placed

- May not be able to allocate new process

3) DYNAMIC RELOCATION

Goal: Protect processes from one another

Requires hardware support

- Memory Management Unit (MMU)

MMU dynamically changes process address at every memory reference

- Process generates logical or virtual addresses (in their address space)
- Memory hardware uses physical or real addresses

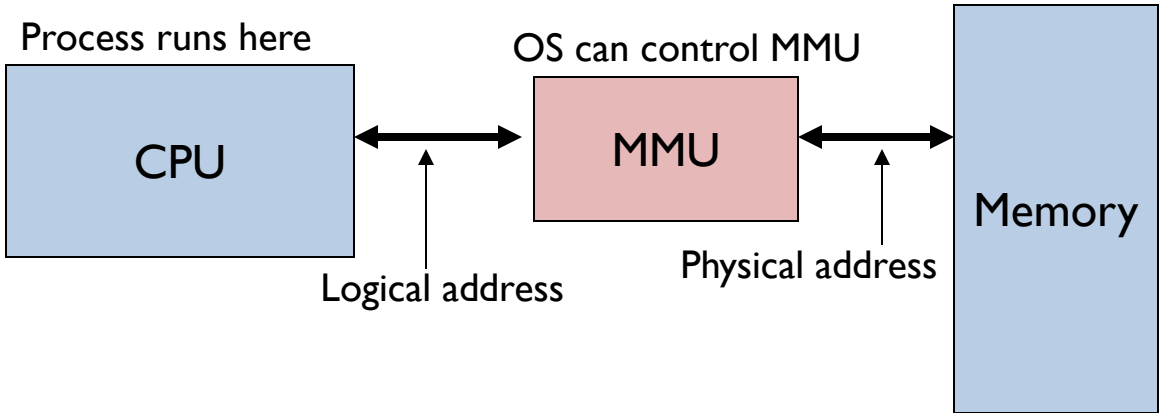

HARDWARE SUPPORT FOR DYNAMIC RELOCATION

Privileged (protected, kernel) mode: OS runs

- When enter OS (trap, system calls, interrupts, exceptions)
- Allows certain instructions to be executed
  (Can manipulate contents of MMU)
- Allows OS to access all of physical memory

User mode: User processes run

- Perform translation of logical address to physical address

IMPLEMENTATION OF DYNAMIC RELOCATION: BASE REG

Translation on every memory access of user process

MMU adds base register to logical address to form physical address

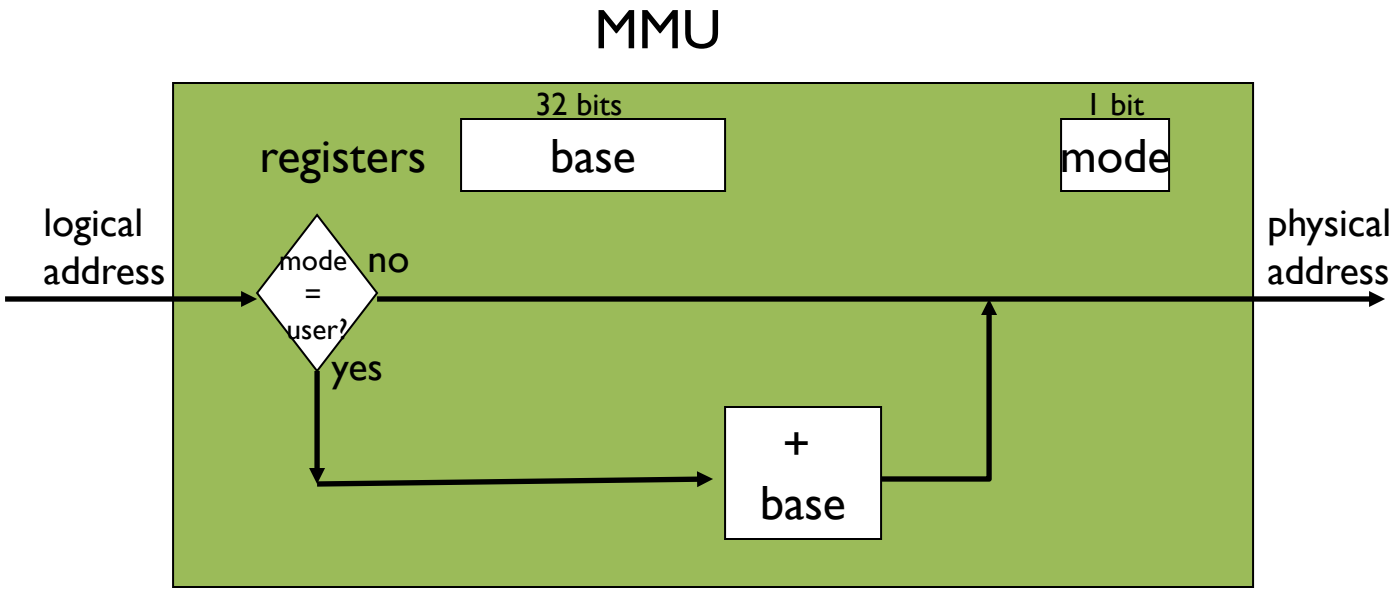

DYNAMIC RELOCATION WITH BASE REGISTER

Translate virtual addresses to physical by adding a fixed offset each time.

Store offset in base register

Each process has different value in base register

Dynamic relocation by changing value of base register!

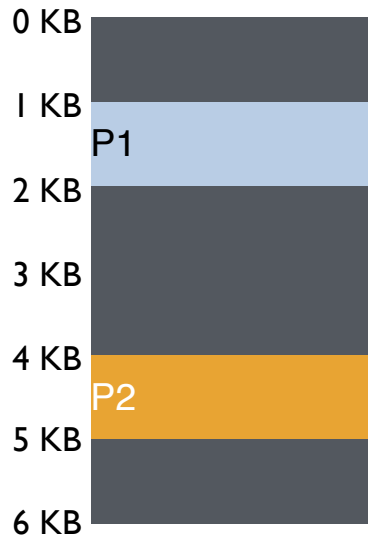

VISUAL EXAMPLE OF DYNAMIC RELOCATION: BASE REGISTER

Virtual

P1: load 100, R1

P2: load 100, R1

P2: load 1000, R1

P1: load 100, R1

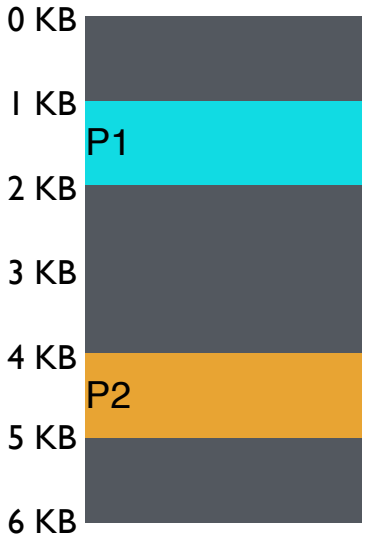

Virtual	Physical
P1: load 100, R1	load 1124, R1
P2: load 100, R1	load 4196, R1
P2: load 1000, R1	load 5096, R1
P1: load 1000, R1	load 2024, R1

Can P2 hurt P1?
Can P1 hurt P2?

How well does dynamic relocation do with base register for protection?

Virtual	Physical	
P1: load 100, R1	load 1124, R1	
P2: load 100, R1	load 4196, R1	
P2: load 1000, R1	load 5096, R1	
P1: load 100, R1	load 2024, R1	
P1: store 3072, R1	store 4096, R1	(3072 + 1024)

How well does dynamic relocation do with base register for protection?

4) DYNAMIC WITH BASE+BOUNDS

Idea: limit the address space with a bounds register

Base register: smallest physical addr (or starting location)
Bounds register: size of this process's virtual address space

- Sometimes defined as largest physical address (base + size)

OS kills process if process loads/stores beyond bounds

IMPLEMENTATION OF BASE+BOUNDS

Translation on every memory access of user process

- MMU compares logical address to bounds register
  if logical address is greater, then generate error
- MMU adds base register to logical address to form physical address

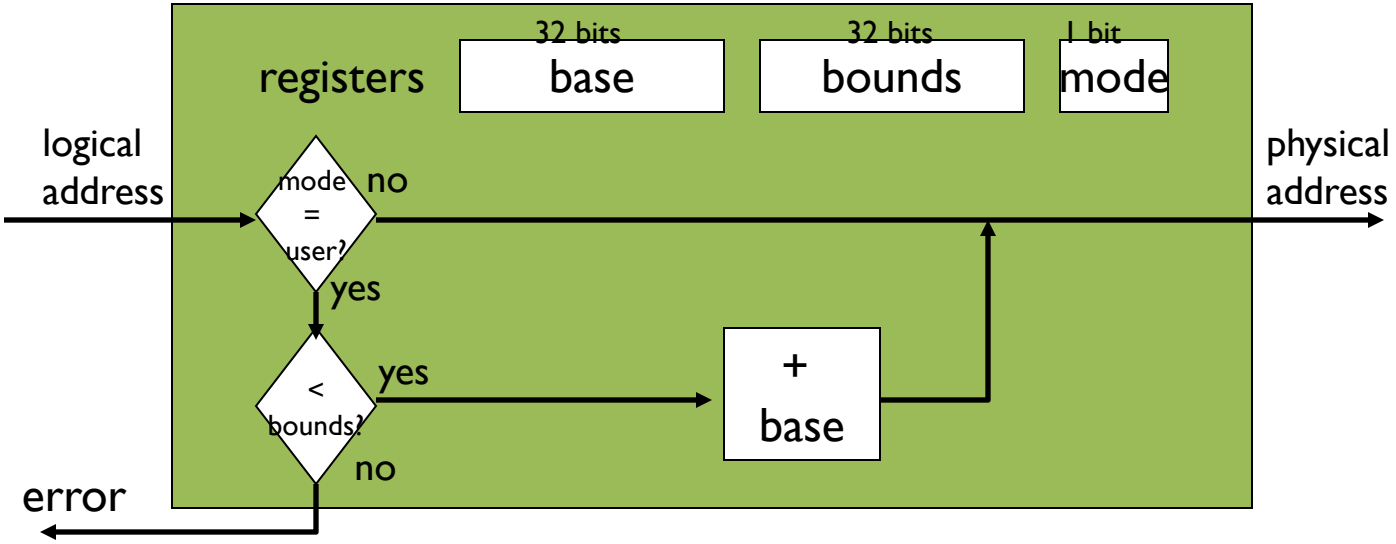

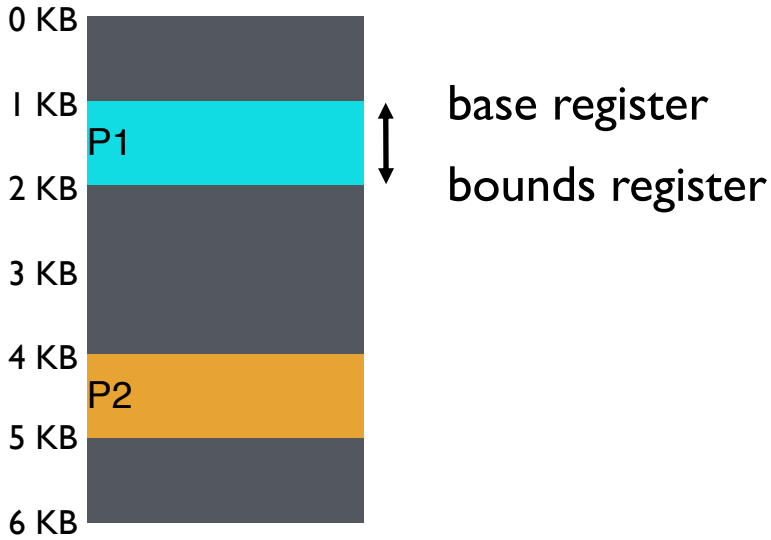
0 KB
1 KB
P1
2 KB
3 KB
4 KB
P2
5 KB
6 KB
base register
bounds register

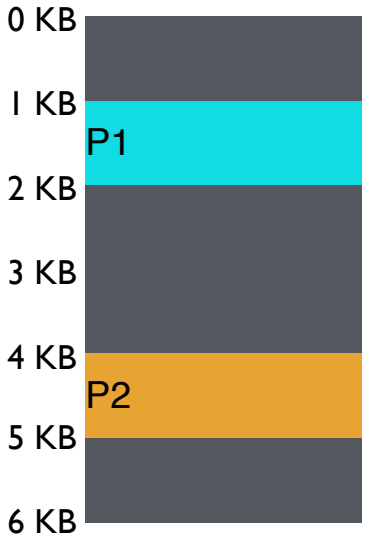

Virtual	Physical
P1: load 100, R1	load 1124, R1
P2: load 100, R1	load 4196, R1
P2: load 1000, R1	load 5196, R1
P1: load 100, R1	load 2024, R1
P1: store 3072, R1	

Can P1 hurt P2?

MANAGING PROCESSES WITH BASE AND BOUNDS

Context-switch: Add base and bounds registers to PCB

Steps

- Change to privileged mode
- Save base and bounds registers of old process
- Load base and bounds registers of new process
- Change to user mode and jump to new process

Protection requirement

- User process cannot change base and bounds registers
- User process cannot change to privileged mode

BASE AND BOUNDS ADVANTAGES

Provides protection (both read and write) across address spaces

Supports dynamic relocation

- Can place process at different locations initially and also move address spaces

Simple, inexpensive implementation: Few registers, little logic in MMU

Fast: Add and compare in parallel

BASE AND BOUNDS DISADVANTAGES

Disadvantages

- Each process must be allocated contiguously in physical memory
  Must allocate memory that may not be used by process

- No partial sharing: Cannot share parts of address space

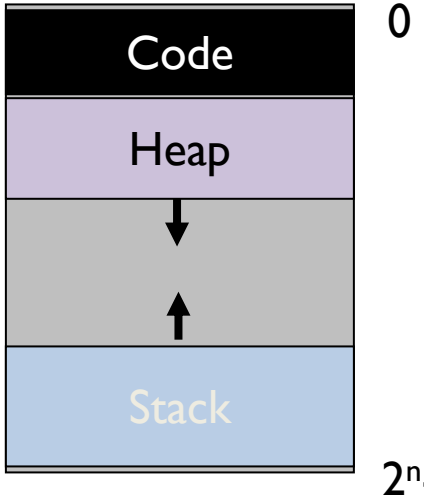

NEXT STEPS

Project 1b: Out now, due Feb 5th

Thursday discussion

- xv6 introduction, walk through
- Project 1b tips

Next week: Virtual memory segmentation, paging and more!