

Operating Systems 412

Pete Keleher
(slides mostly from Youjip Won)

1

Memory

- 14 - Memory API
- 17 - Free Space Management
- 13 - Address Spaces
- 15 - Address Translation
- 16 - Segmentation
- 18 - Paging
- 19 - Translation Lookaside Buffers
- 20 - Advanced Paging
- 21 - Swapping
- 22 - Swapping Policy

2

Memory API: malloc()

```
#include <stdlib.h>

void* malloc(size_t size)
```

- Allocate a memory region on the heap.
 - Argument
 - `size_t size`: size of the memory block(in bytes)
 - `size_t` is an unsigned integer type.
 - Return
 - Success : a void type pointer to the memory block allocated by `malloc`
 - Fail : a null pointer

3

sizeof()

- Routines and macros are utilized for `size` in `malloc` instead typing in a number directly.
- Careful of `sizeof`!

```
int *x = malloc(10 * sizeof(int));
printf("%d\n", sizeof(x));
```

4

```
int x[10];
printf("%d\n", sizeof(x));
```

40

4

Memory API: free()

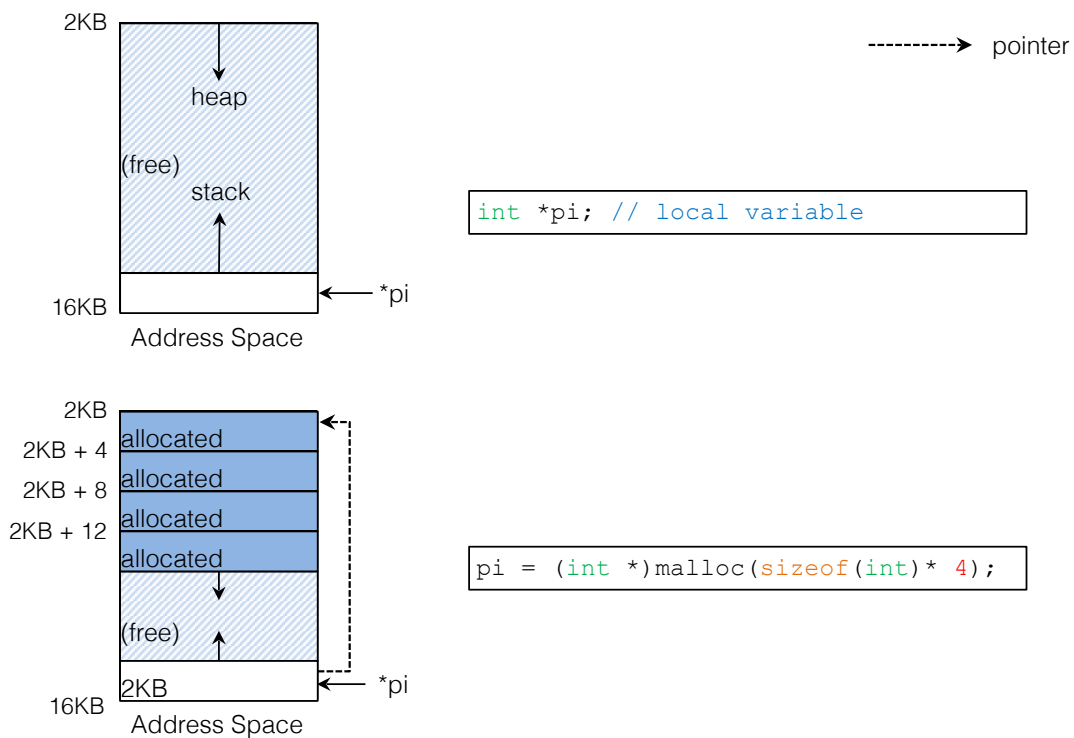
```
#include <stdlib.h>

void free(void* ptr)
```

- Free memory region allocated by a call to `malloc`.
 - Argument
 - `void *ptr`: a pointer to a memory block allocated with `malloc`
 - Return
 - none

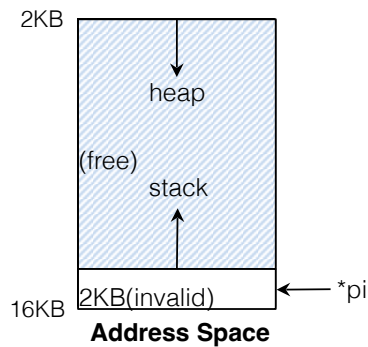
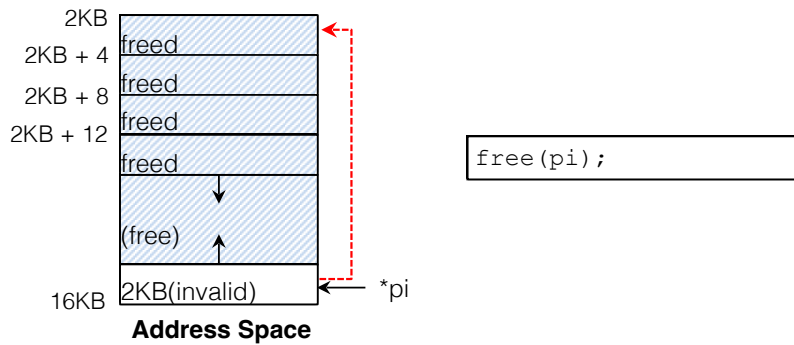
5

Memory Allocating



6

Freeing Memory

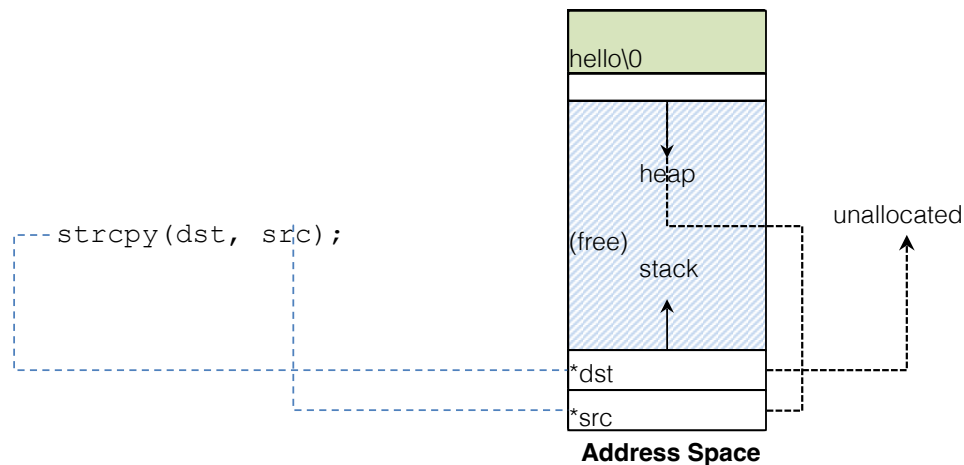


7

Forgetting To Allocate Memory

- Incorrect code

```
char *src = "hello"; //character string constant
char *dst;           //unallocated
strcpy(dst, src);    //segfault and die
```

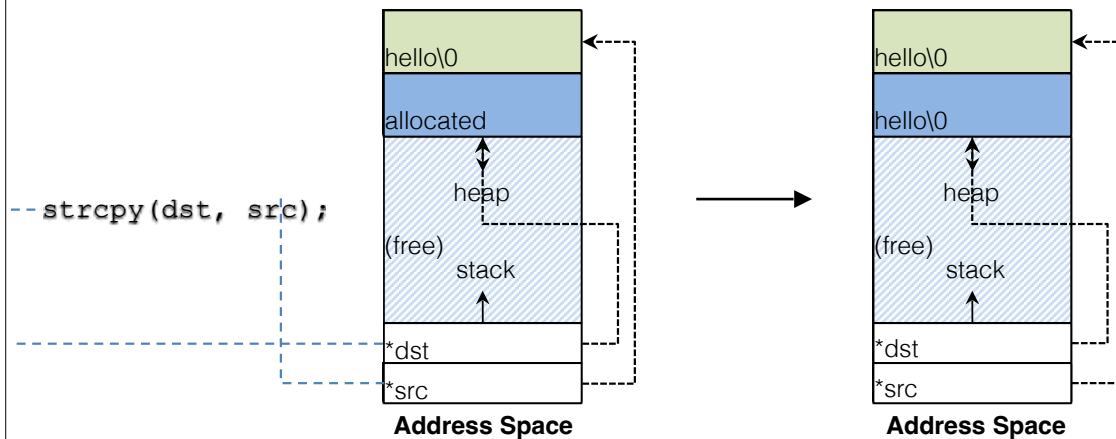


8

Forgetting To Allocate Memory(Cont.)

- Correct code

```
char *src = "hello"; //character string constant
char *dst (char *)malloc(strlen(src) + 1 ); // allocated
strcpy(dst, src); //work properly
```

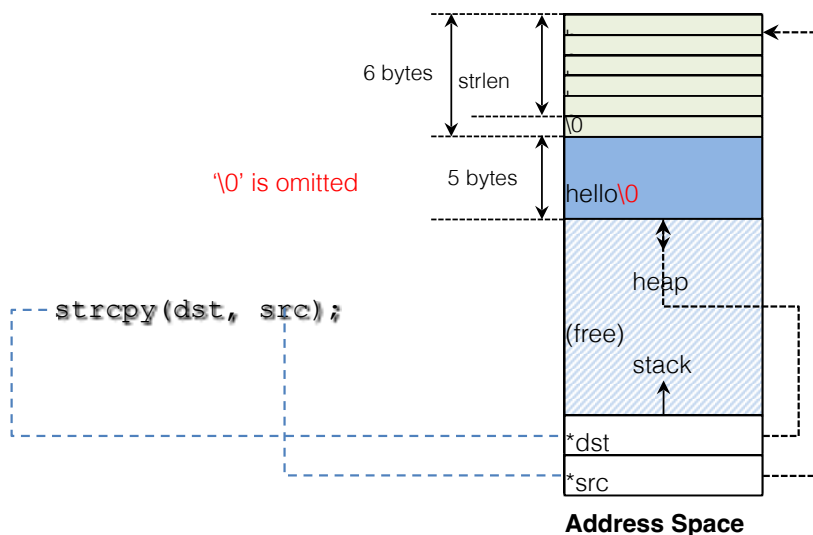


9

Not Allocating Enough Memory

- Incorrect code, but sometimes works

```
char *src = "hello"; //character string constant
char *dst (char *)malloc(strlen(src)); // too small
strcpy(dst, src); //work properly
```

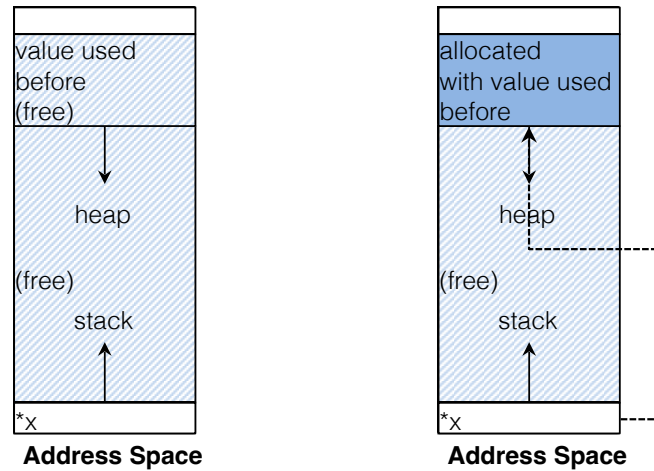


10

Forgetting to Initialize

- Encounter an uninitialized read

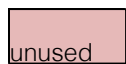
```
int *x = (int *)malloc(sizeof(int)); // allocated
printf("**x = %d\n", *x); // uninitialized memory access
```

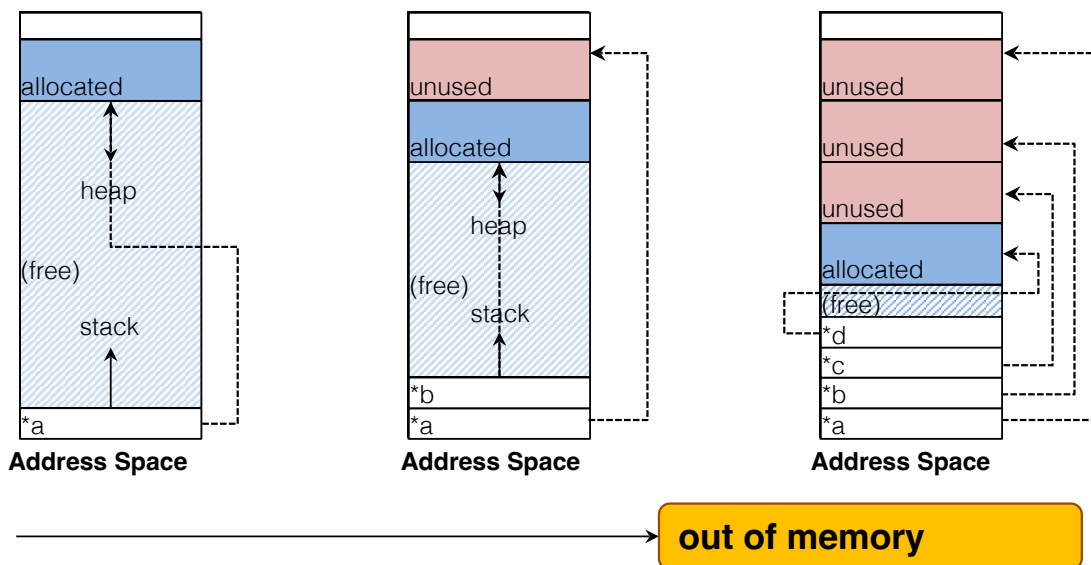


11

Memory Leak

- A program runs out of memory and eventually dies.

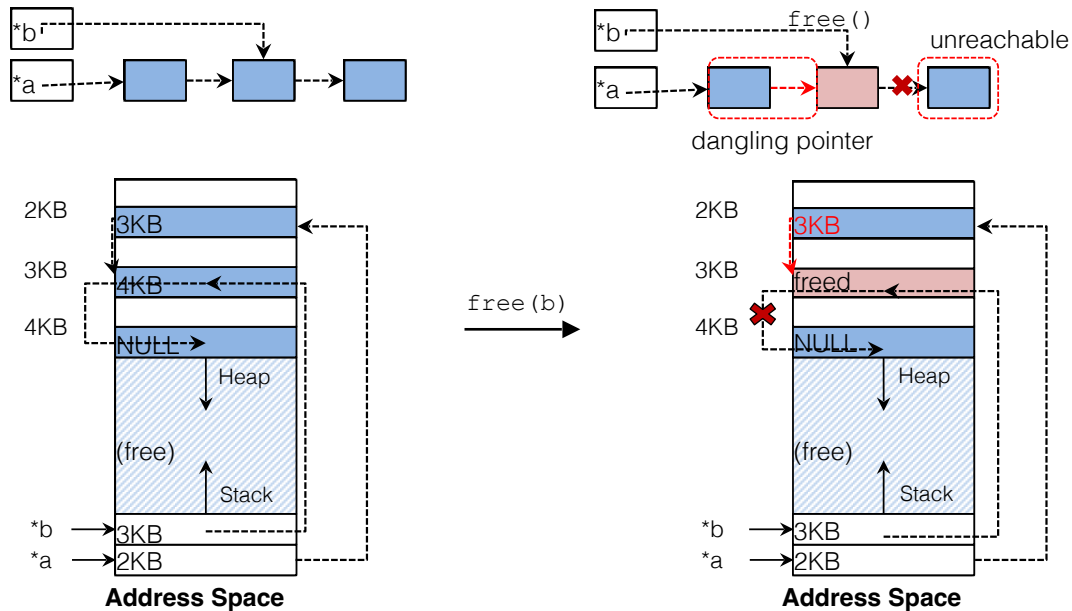
 : unused, but not freed



12

Dangling Pointer

- Use after memory freed: invalid pointer



13

Other Memory APIs: calloc()

```
#include <stdlib.h>

void *calloc(size_t num, size_t size)
```


- Allocate memory on the heap and zeroes it before returning.
 - Argument
 - `size_t num` : number of blocks to allocate
 - `size_t size` : size of each block(in bytes)
 - Return
 - Success : a void type pointer to the memory block allocated by `calloc`
 - Fail : a null pointer

14

Double Free

- Free memory that was freed already.

```
int *x = (int *)malloc(sizeof(int)); // allocated
free(x); // free memory
free(x); // free repeatedly
```



15

Other Memory APIs: realloc()

```
#include <stdlib.h>
void *realloc(void *ptr, size_t size)
```

- Change the size of memory block.
 - A pointer returned by `realloc` may be either the same as `ptr` or a new.
 - Argument
 - `void *ptr`: Pointer to memory block allocated with `malloc`, `calloc` or `realloc`
 - `size_t size`: New size for the memory block(in bytes)
 - Return
 - Success: Void type pointer to the memory block
 - Fail : Null pointer

16

System Calls

```
#include <unistd.h>

int  brk(void *addr)
void *sbrk(intptr_t increment);
```

- malloc library uses the brk system call
 - brk is called to expand the program's break.
 - break: The location of the end of the heap in address space
 - sbrk is an additional call similar with brk.
 - Programmers should never directly call either brk or sbrk.

17

System Calls(Cont.)

```
#include <sys/mman.h>

void *mmap(void *ptr, size_t length, int port, int flags, int
fd, off_t offset)
```

18


Memory

- 14 - Memory API
- 17 - Free Space Management
- 13 - Address Spaces
- 15 - Address Translation
- 16 - Segmentation
- 18 - Paging
- 19 - Translation Lookaside Buffers
- 20 - Advanced Paging
- 21 - Swapping
- 22 - Swapping Policy

19

Splitting


- Finding a free chunk of memory that can satisfy the request and splitting it into two.
 - When request for memory allocation is **smaller** than the size of free chunks.



20

Splitting(Cont.)


- Two 10-byte free segments with a **1-byte request**



21

Coalescing

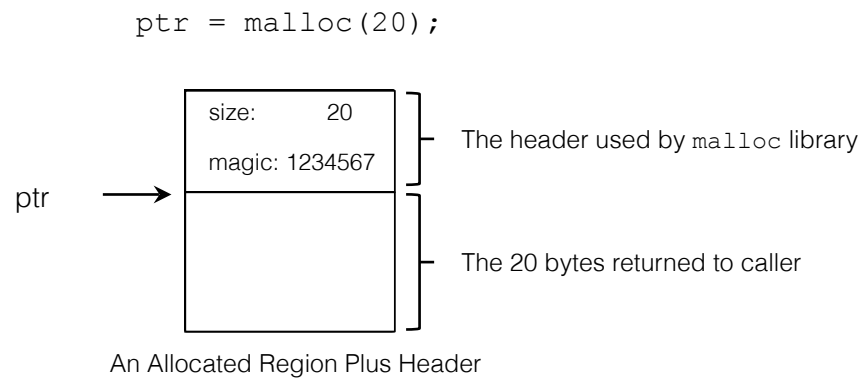
- If no large enough chunk exists, need to *coalesce*:



22

Tracking The Size of Allocated Regions

- `free(void *ptr)` does *not take a size parameter*.
 - How does the library **know the size** at dealloc time?
- Most allocators store extra information in a header block
 - size
 - additional pointers to speed up deallocation
 - magic number for integrity checking



23

Embedding A Free List

- The memory-allocation library initializes the heap and puts the first element of the free list in the free space.
- Description of a node of the list

```
typedef struct __node_t {  
    int size;  
    struct __node_t *next;  
} nodet_t;
```


- Building the heap and free list
 - Assume that the heap is built via `mmap()` system call.

```
// mmap() returns a pointer to a chunk of free space  
node_t *head = mmap(NULL, 4096, PROT_READ|PROT_WRITE,  
    MAP_ANON|MAP_PRIVATE, -1, 0);  
head->size = 4096 - sizeof(node_t);  
head->next = NULL;
```

24

A Heap With One Free Chunk

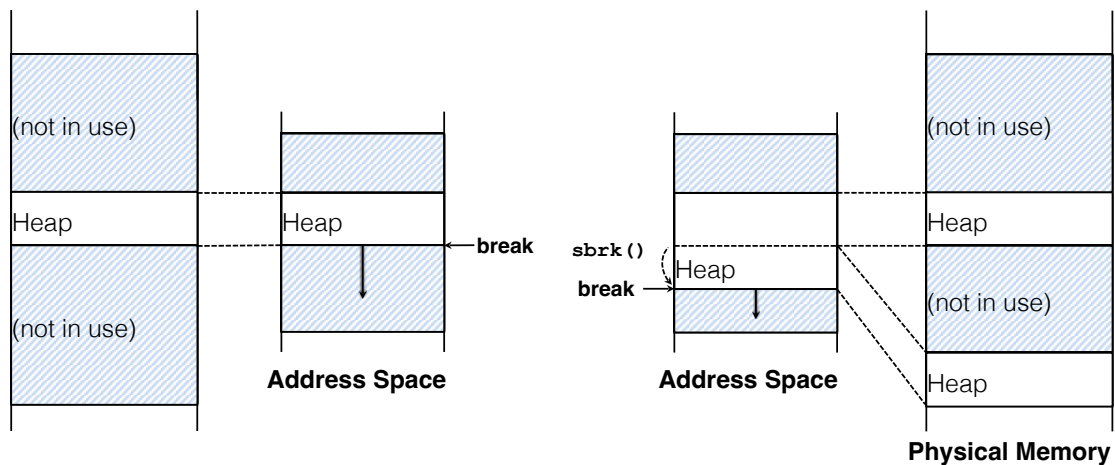
```
// mmap() returns a pointer to a chunk of free space
node_t *head = mmap(NULL, 4096, PROT_READ|PROT_WRITE,
                    MAP_ANON|MAP_PRIVATE, -1, 0);
head->size = 4096 - sizeof(node_t);
head->next = NULL;
```



25

Growing The Heap

- Most allocators **start with a small-sized heap** and then **request more** memory from the OS when they run out.
 - e.g., `sbrk()`, `brk()` in most UNIX systems.



26


Managing Free Space: Basic Strategies

- Best Fit:
 - Finding free chunks that are **big or bigger than the request**
 - Returning the **one of smallest** in the chunks **in the group** of candidates
- Worst Fit:
 - Finding the **largest free chunks** and allocation the amount of the request
 - **Keeping the remaining chunk** on the free list.


27

Examples of Basic Strategies


- Allocation Request Size 15



- Best-fit:



- Worst-fit:



28


Other Approaches: Segregated List

- **Segregated List:**
 - Keeping separate free lists for popular requests.
 - New complication:
 - How much memory should dedicate to the pool of memory that serves specialized requests of a given size?
 - Slab allocator handles this issue.
- **Slab Allocator**
 - Allocate a number of caches for popular specific sizes at system boot.
 - e.g., locks, file-system inodes, etc.
 - Request a new slab from general memory allocator (size multiple of page size * object size) when a given cache is running low.

29

Buddy Allocation

- **Binary Buddy Allocation**
 - The allocator divides free space by two until a block that is big enough to accommodate the request is found.



64KB free space for 7KB request

- **Internal fragmentation dealt with through coalescing:**
 - When block b is freed, coalesce w/ buddy if also free, etc.

30

Memory

- 14 - Memory API
- 17 - Free Space Management
- 13 - Address Spaces
- 15 - Address Translation
- 16 - Segmentation
- 18 - Paging
- 19 - Translation Lookaside Buffers
- 20 - Advanced Paging
- 21 - Swapping
- 22 - Swapping Policy

31

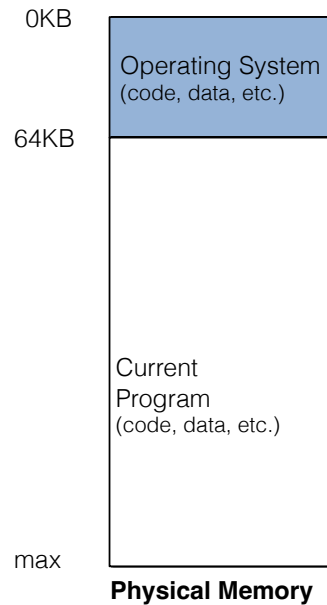
Memory Virtualization

- What is memory virtualization?
 - OS virtualizes its physical memory.
 - OS provides a [virtual address space](#) for each process.
 - Illusion of [each process using the entire physical memory](#) .
- Why?
 - Ease of use in programming
 - Memory efficiency in [time](#) and [space](#)
 - *Isolation* for processes as well as OS
 - Protection from [errant accesses](#) of other processes

32

Early Operating Systems


- Load only one process in memory.
 - Poor utilization and efficiency



33

Multiprogramming and Time Sharing


- Load multiple processes in memory
 - Execute one for a short while.
 - Switch processes between them in memory.
 - Increase utilization and efficiency.
- But what about protection?
 - Errant memory accesses from other processes



34

Address Space

- An abstraction of physical memory:



- **Code**
 - Where instructions live
- **Heap**
 - Dynamically allocate memory.
 - `malloc` in C
 - `new` in object-oriented languages
- **Stack**
 - Store return addresses or values.
 - Contain local variables arguments to routines.

35

Virtual Addresses

- Every address in a running program is virtual.

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

int main(int argc, char *argv[]){

    printf("location of code : %p\n", (void *) main);
    printf("location of heap : %p\n", (void *) malloc(1));
    int x = 3;
    printf("location of stack : %p\n", (void *) &x);

    return x;
}
```

- OS uses hardware to translate virtual addresses to physical

36

Virtual Addresses

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

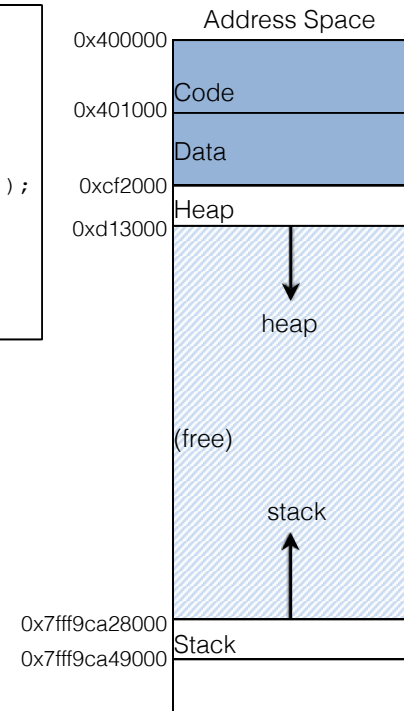
int main(int argc, char *argv[]){

    printf("location of code : %p\n", (void *) main);
    printf("location of heap : %p\n", (void *) malloc(1));
    int x = 3;
    printf("location of stack : %p\n", (void *) &x);

    return x;
}
```

Output in 64-bit Linux machine:

```
location of code : 0x40057d
location of heap : 0xcf2010
location of stack : 0x7fff9ca45fcc
```



37

Memory

- 14 - Memory API
- 17 - Free Space Management
- 13 - Address Spaces
- 15 - Address Translation
- 16 - Segmentation
- 18 - Paging
- 19 - Translation Lookaside Buffers
- 20 - Advanced Paging
- 21 - Swapping
- 22 - Swapping Policy

38

Need Efficiency, and Control...

- Limited direct execution (LDE)
 - Programs run directly (not emulated)
 - Memory virtualizing, efficiency, control maintained by **hardware support**.
 - e.g., registers, TLBs (Translation Look-aside Buffers), page-tables
- Hardware transforms virtual addresses to physical addresses
 - Memory only addressed with physical addresses
- The OS sets up the hardware.
 - Hardware raises interrupts when needed.

39

Example: Address Translation

```
void func()  
    int x;  
    ...  
    x = x + 3; // this is the line of code we are interested in
```


- **Load** a value from memory
 - **Increment** by three
 - **Store** the value back into memory
- **Assembly**

```
128 : movl 0x0(%ebx), %eax      ; load 0+ebx into eax  
132 : addl $0x03, %eax        ; add 3 to eax register  
135 : movl %eax, 0x0(%ebx)    ; store eax back to mem
```

- Assume address of 'x' in ebx register.
- **Load** the value at that address into eax register.
- **Add 3** to eax register.
- **Store** the value in eax back into memory.

40

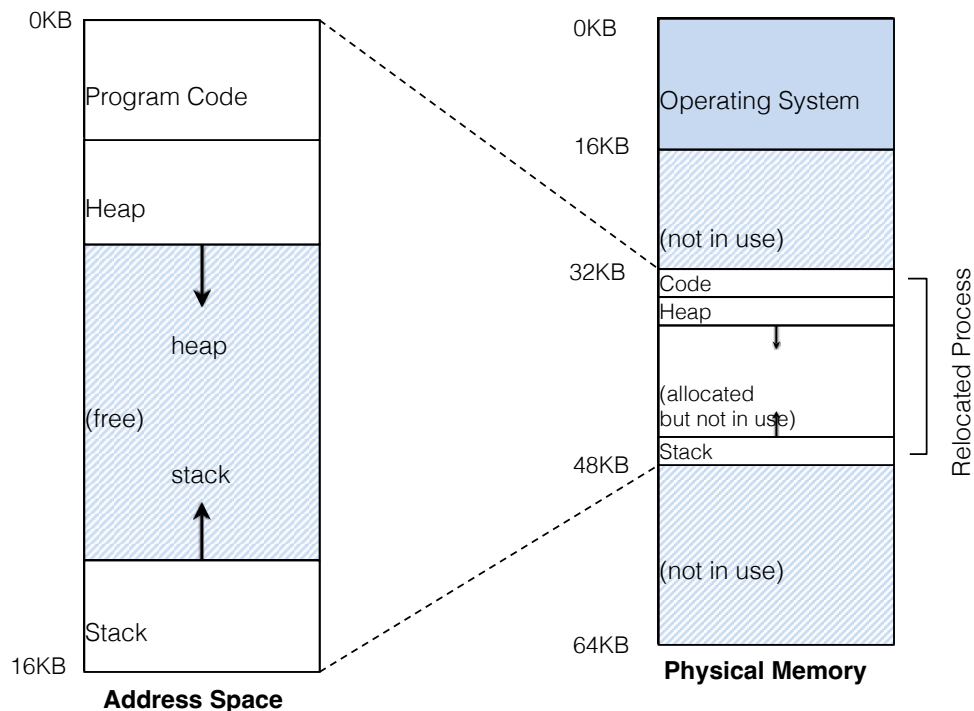
Example: Address Translation



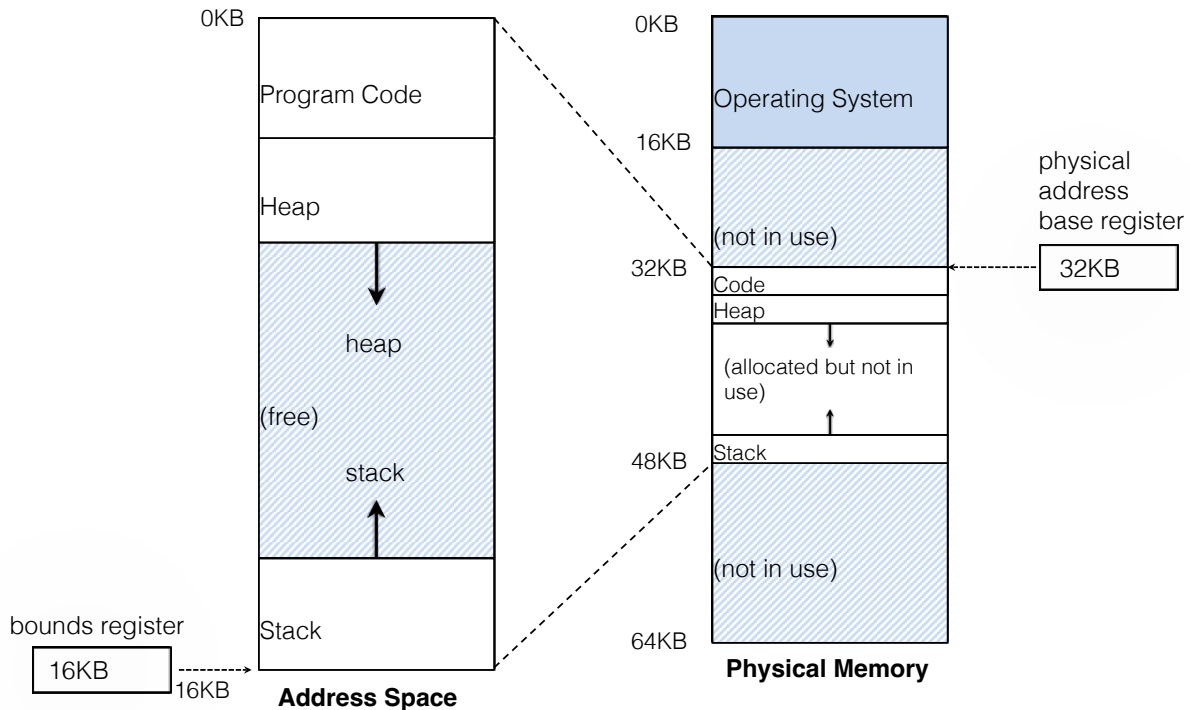
- Fetch instruction at address 128
- Execute instruction (load from address 15KB)
- Fetch instruction at address 132
- Execute instruction (no memory reference)
- Fetch the instruction at address 135
- Execute instruction (store to address 15 KB)

But not all programs can be at location 0

A Single Relocated Process



Base and Bounds Register



43

Dynamic(Hardware base) Relocation

- OS decides where in physical memory a process is loaded.
 - Set the **base** register:
 $\text{physical address} = \text{virtual address} + \text{base}$
 - Virtual addresses must **not be greater than bound** or **negative**:
 $0 \leq \text{virtual address} < \text{bound}$

44

Relocation and Address Translation

128 : `movl 0x0(%ebx), %eax`

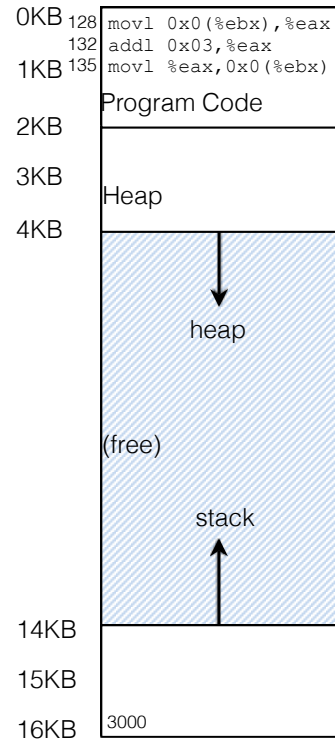
- **Fetch** instruction at address 128

$$32896 = 128 + 32KB(base)$$

- **Execute** this instruction

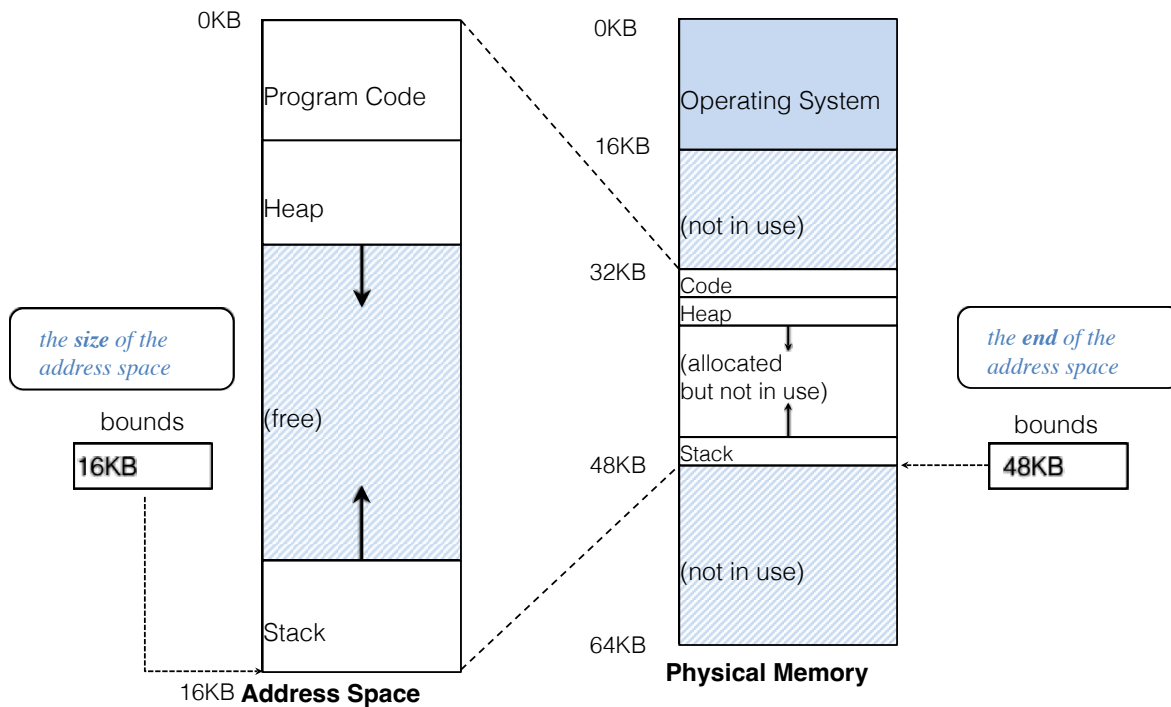
- Load from address 15KB

$$47KB = 15KB + 32KB(base)$$



virtually addressed

Two ways to Use the Bounds Register



OS Issues for Memory Virtualizing

- OS intervenes at three critical junctures:
 - When a process starts running:
 - find space for address space in physical memory
 - When a process is terminated:
 - reclaims the memory for use
 - When context switch occurs:
 - Save and store the base-and-bounds pair