

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

47

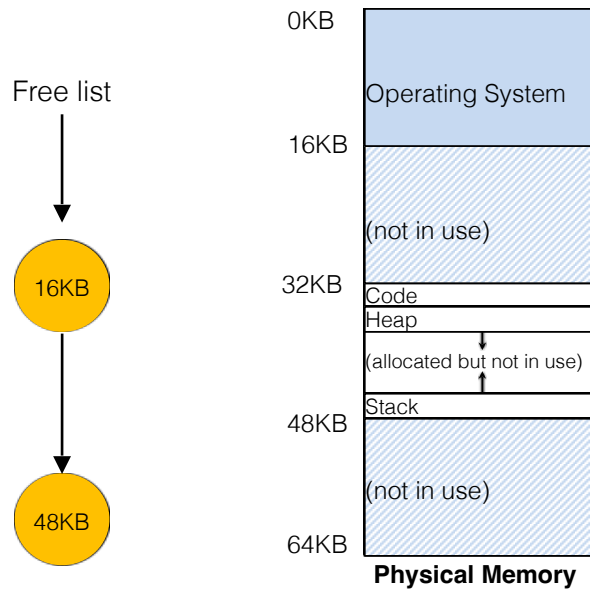
Issues (still Base and Bounds)

- 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

48

OS Issues: When a Process Starts Running


- The OS must find a room for a new address space.
 - free list : a list of unused ranges of physical memory



49

OS Issues: Process Termination

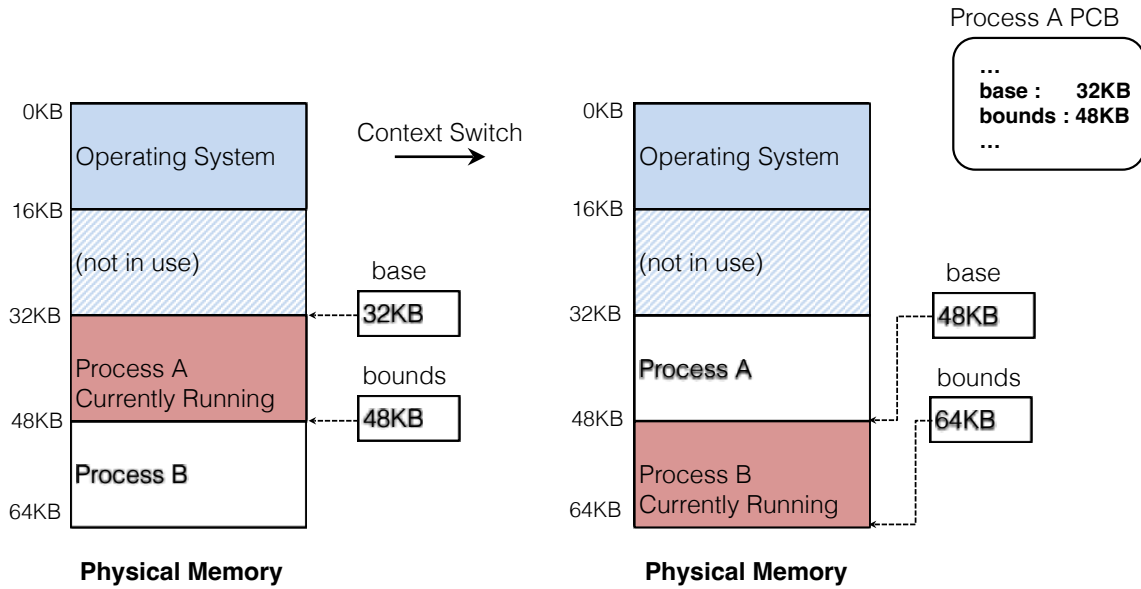
- OS must put memory back on the free list.



50


OS Issues: Context Switches

- OS must save and restore base-and-bounds pair.
 - In process structure or process control block(PCB)



51

Not Always Efficient



- Need big chunk of "free" space
 - physically consecutive memory
- Cannot run when address space **does not fit** into physical memory

52

Memory


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

53

Segmentation

- Segment is a *contiguous* portion of the address space:
 - Several types: code, stack, heap, ...
- Segments can be placed anywhere in physical memory.
 - Slightly modified base and bounds per segment

Segment	Base	Size
Code	32K	2K
Heap	34K	2K
Stack	28K	2K




54

Address Translation for Segments

$$\text{physical address} = \text{offset} + \text{base}$$

- The `offset` of below virtual address 100 is 100.
 - The code segment **starts at virtual address 0** in address space.




55

Address Translation for Segments

$$\text{Offset} + \text{base, not virtual address} + \text{base}$$


- The `offset` of virtual address 4200 is 104.
 - The heap segment **starts at virtual address 4096** in address space.



56

Segment Descriptors

- Explicit approach
 - Chop up the address space into segments based on the **top few bits** of virtual address.
- Example: virtual address 4200 (01000001101000)



Segment	bits
Code	00
Heap	01
Stack	10
-	11

57

Segment Descriptors


- Bits
 - SEG_MASK = 0x3000 (1100000000000000)
 - SEG_SHIFT = 12
 - OFFSET_MASK = 0xFFF (0011111111111111)

```
1 // get top 2 bits of 14-bit VA
2 Segment = (VirtualAddress & SEG_MASK) >> SEG_SHIFT
3
4 // now get offset
5 Offset = VirtualAddress & OFFSET_MASK
6 if (Offset >= Bounds[Segment])
7     RaiseException(PROTECTION_FAULT)
8 else
9     PhysAddr = Base[Segment] + Offset
10    Register = AccessMemory(PhysAddr)
```

58

Referring to Stack Segment

- Stack grows *backward*.
- *Extra hardware support needed*.
 - The hardware checks which way the segment grows.
 - 1: positive direction, 0: negative direction



Segment Register(with Negative-Growth Support)

Segment	Base	Size	Grows Positive?
Code	32K	2K	1
Heap	34K	2K	1
Stack	28K	2K	0

$$\text{address} = \text{base} + \text{offset} - \text{sizeof}(\text{stack})$$

59

Support for Sharing

- Segments can be shared between address spaces
 - Code sharing still used
- Need hardware support in form of *protection bits*.
 - Bits indicate read, write and execute permissions.

Segment Register Values(with Protection)

Segment	Base	Size	Grows Positive?	Protection
Code	32K	2K	1	Read-Execute
Heap	34K	2K	1	Read-Write
Stack	28K	2K	0	Read-Write

60

Fine-Grained and Coarse-Grained

- Coarse-Grained is small number of segments
 - e.g., code, heap, stack.
- Fine-Grained segmentation allows more flexibility
 - Hardware-supported segment tables

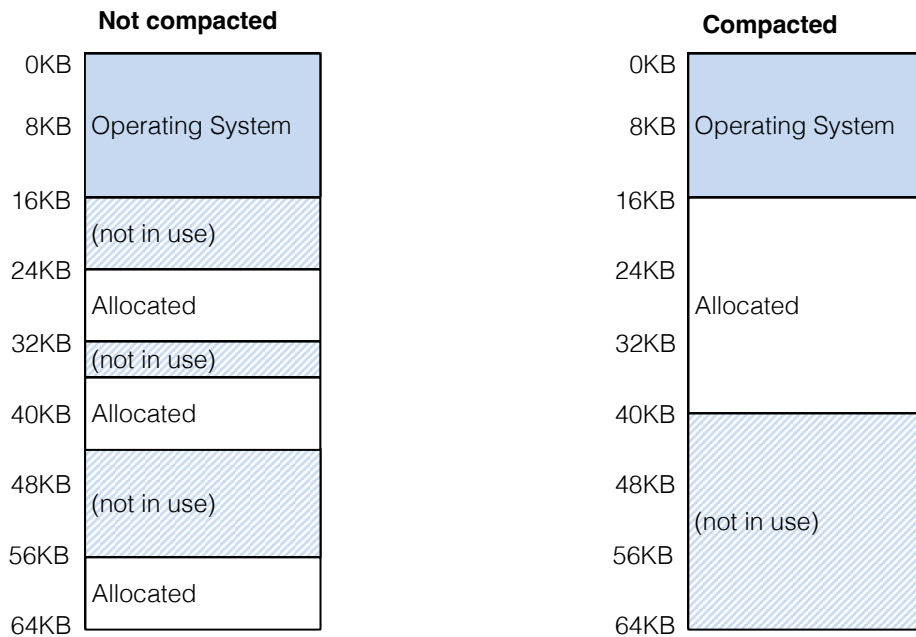
61

OS support: Fragmentation

- External Fragmentation:
 - Distinct runs of free space in physical memory
 - Might be **24KB free**, but **not in one contiguous** segment.
 - The OS **cannot** satisfy the **20KB request**.
- Compaction: consolidating existing segments in physical memory.
 - Compaction is **costly**.
 - **Stop** running process.
 - **Copy** data to somewhere.
 - **Change** segment register value.

62

Memory Compaction



63

GeekOS

- segmented memory addresses
 - 16-bit “segment selector”, 32-bit offset
 - segment selector has:
 - 1 bit: GDT or LDT
 - 13 bits: index into GDT or LDT
 - 2 bits: protection level of segment
 - segment descriptor (from table) has:
 - linear *base* physical address of segment: 32 bits
 - limit (size) of segment: 20 bits
 - descriptor privilege level (dpl): 2 bits
 - type of segment (data, code, system, tss, gate): 4 bits
 - present (in-memory): 1 bit
 - etc.

64

GeekOS

- GDT
 - entries point to kernel segments, optionally user segments
 - entry 0 (null selector) is not used to access memory
 - `gdt r` register points to the GDT
- LDT similar, but
 - points to segments of a single process
 - entry 0 can be used
 - any number of LDTs can be in memory
 - `LDTR` register points (via GDT) to currently used LDT

65

Virtual Memory

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

66

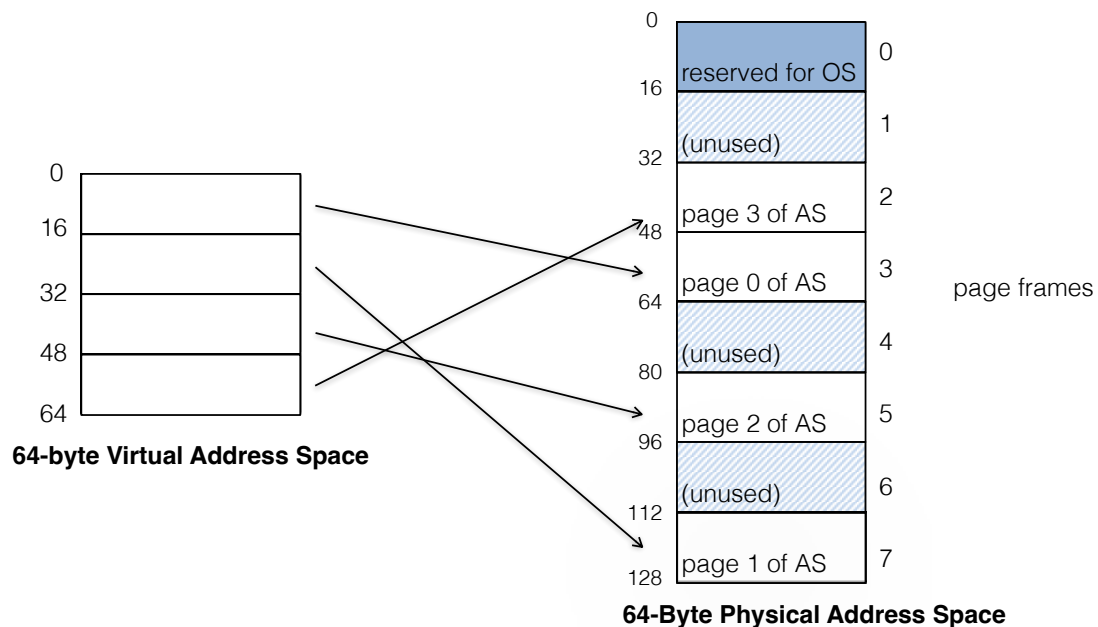
Paging

- Paging splits address space into fixed-size *pages*.
 - vs segmentation: variable size of logical segments
- Physical memory holding a page is the *page frame*
- Page table per process
 - translates virtual address to physical address.
- Flexibility:
 - No assumptions on how heap and stack grow or are used
- Simplicity: ease of free-space management
 - All pages and page frames are the same size
 - Free lists are easy...

67

Paging Example


- 128-byte physical memory with eight 16-byte page frames
- 64-byte address space with 16-byte pages




68

Address Translation

- Two components in the virtual address
 - VPN: virtual page number
 - Offset: offset within the page



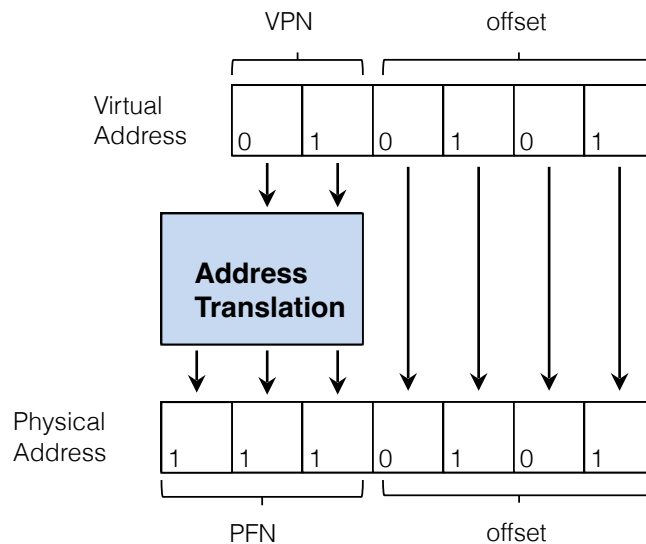
- Example: virtual address 21 in 64-byte address space



69

Example: Address Translation

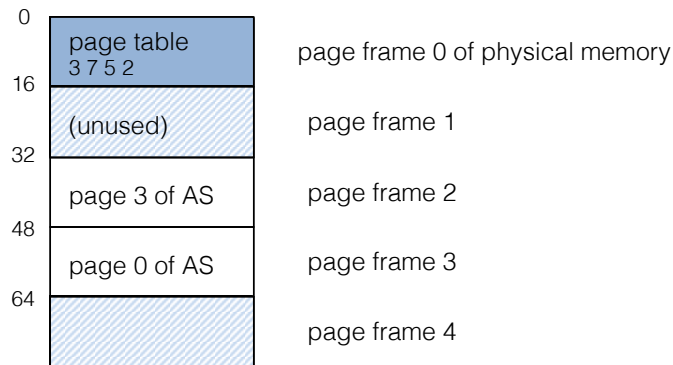
- The virtual address 21 in 64-byte address space



70

Where Are Page Tables Stored?

- Page tables can be large...
 - 32-bit address space with 4-KB pages, 20 bits for VPN
 - assume entry is 4 bytes:
 - page table size is $2^{20} * 4 = 2^{22} = 4\text{MB}$ of space
- Page tables for each process are stored in memory...




71

What Is In The Page Table?

- A page table is just a **data structure** that is used to map the virtual address to physical address.
 - Simplest form: a linear page table, an array
- The OS/hardware accesses a page-table entry by indexing into the array by virtual page-number
- Common bits:
 - Valid Bit: whether the particular translation is valid.
 - Protection Bit: read, write, execute
 - Present Bit: in physical memory or swapped out
 - Dirty Bit: page modified since it brought into memory
 - Reference Bit(Accessed Bit): page has been accessed

72

Example: x86 Page Table Entry



- P: present
- R/W: read/write bit
- U/S: supervisor
- A: accessed bit
- D: dirty bit
- PFN: the page frame number
- Others: mostly caching directives

73

Paging: Too Slow

- To find a location of the desired PTE, the starting location of the page table is needed.
- For every memory reference, paging requires the OS to perform one extra memory reference.

74

Accessing Memory With Paging

```
1 // Extract the VPN from the virtual address
2 VPN = (VirtualAddress & VPN_MASK) >> SHIFT
3
4 // Form the address of the page-table entry (PTE)
5 PTEAddr = PTBR + (VPN * sizeof(PTE))
6
7 // Fetch the PTE
8 PTE = AccessMemory(PTEAddr)
9
10 // Check if process can access the page
11 if (PTE.Valid == False)
12     RaiseException(SEGMENTATION_FAULT)
13 else if (CanAccess(PTE.ProtectBits) == False)
14     RaiseException(PROTECTION_FAULT)
15 else
16     // Access is OK: form physical address and fetch
17     offset = VirtualAddress & OFFSET_MASK
18     PhysAddr = (PTE.PFN << PFN_SHIFT) | offset
19     Register = AccessMemory(PhysAddr)
```

75

A Memory Trace

- Example: A Simple Memory Access

```
int array[1000];
...
for (i = 0; i < 1000; i++)
    array[i] = 0;
```

- Compile and execute

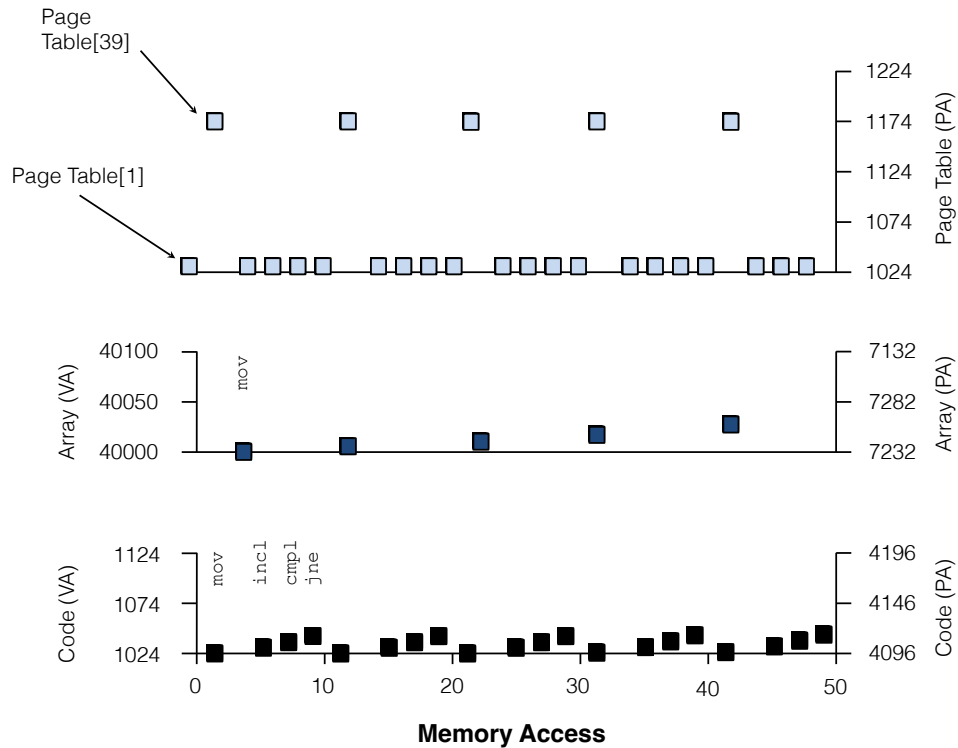
```
prompt> gcc -o array array.c -Wall -o
prompt> ./array
```

- Resulting Assembly code

```
0x1024 movl $0x0, (%edi,%eax,4)
0x1028 incl %eax
0x102c cmpl $0x03e8,%eax
0x1030 jne 0x1024
```

76

A Virtual(And Physical) Memory Trace



77

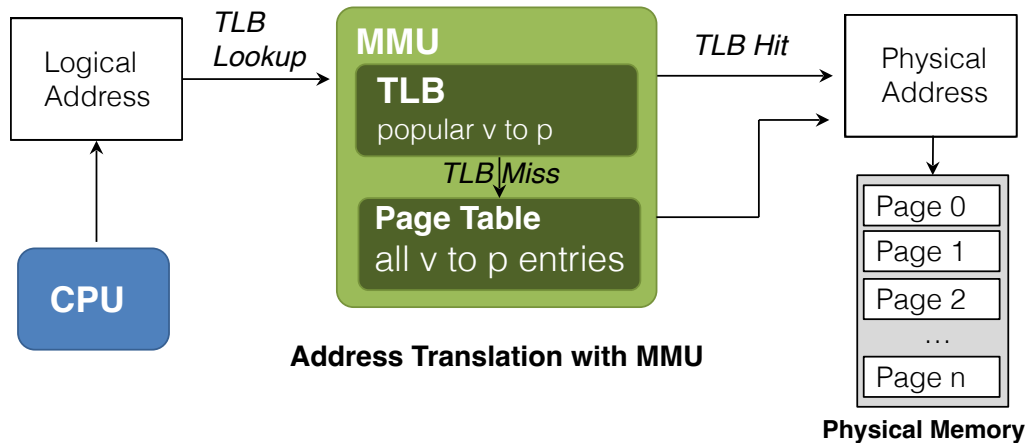
Virtual Memory

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

78

TLB

- Part of the chip's memory-management unit(MMU).
- A hardware cache of **popular** virtual-to-physical address translation.



79

TLB Basic Algorithms

- extract the virtual page number (VPN).
- check for hit in the the TLB
- extract page frame number from relevant TLB entry, form desired physical address, and access memory

```
1: VPN = (VirtualAddress & VPN_MASK ) >> SHIFT
2: (Success , TlbEntry) = TLB_Lookup(VPN)
3:   if (Success == True){ // TLB Hit
4:     if (CanAccess(TlbEntry.ProtectBit) == True ){
5:       offset = VirtualAddress & OFFSET_MASK
6:       PhysAddr = (TlbEntry.PFN << SHIFT) | Offset
7:       AccessMemory( PhysAddr )
8:     } else RaiseException(PROTECTION_ERROR)
```

80

TLB Basic Algorithms (Cont.)

- (11-12 lines) The hardware accesses the page table to find the translation.
- (16 lines) updates the TLB with the translation.

```
11:     } else { //TLB Miss
12:         PTEAddr = PTBR + (VPN * sizeof(PTE))
13:         PTE = AccessMemory(PTEAddr)
14:         (...)
15:     } else {
16:         TLB_Insert( VPN , PTE.PFN , PTE.ProtectBits)
17:         RetryInstruction()
18:     }
19: }
```

81

Example: Accessing An Array

- ▣ How a TLB can improve its performance.

	OFFSET			
	00	04	08	12
VPN = 00				
VPN = 01				
VPN = 03				
VPN = 04				
VPN = 05				
VPN = 06		a[0]	a[1]	a[2]
VPN = 07	a[3]	a[4]	a[5]	a[6]
VPN = 08	a[7]	a[8]	a[9]	
VPN = 09				
VPN = 10				
VPN = 11				
VPN = 12				
VPN = 13				
VPN = 14				

The TLB improves performance due to **spatial locality**

```
0:     int sum = 0 ;
1:     for( i=0; i<10; i++){
2:         sum+=a[i];
3:     }
```


3 TLB misses and 7 hits.
Thus **TLB hit rate** is 70%.

82

Locality

- Temporal Locality

- An instruction or data item that has been recently accessed will likely be re-accessed soon in the future.



- Spatial Locality

- If a program accesses memory at address x , it will likely soon access memory near x .

