


Memory Hierarchy

The background of the slide features a photograph of a statue, likely the 'The Spirit of the Law' statue at the University of Illinois, set in a park-like environment with trees. The entire image is overlaid with a semi-transparent orange color.

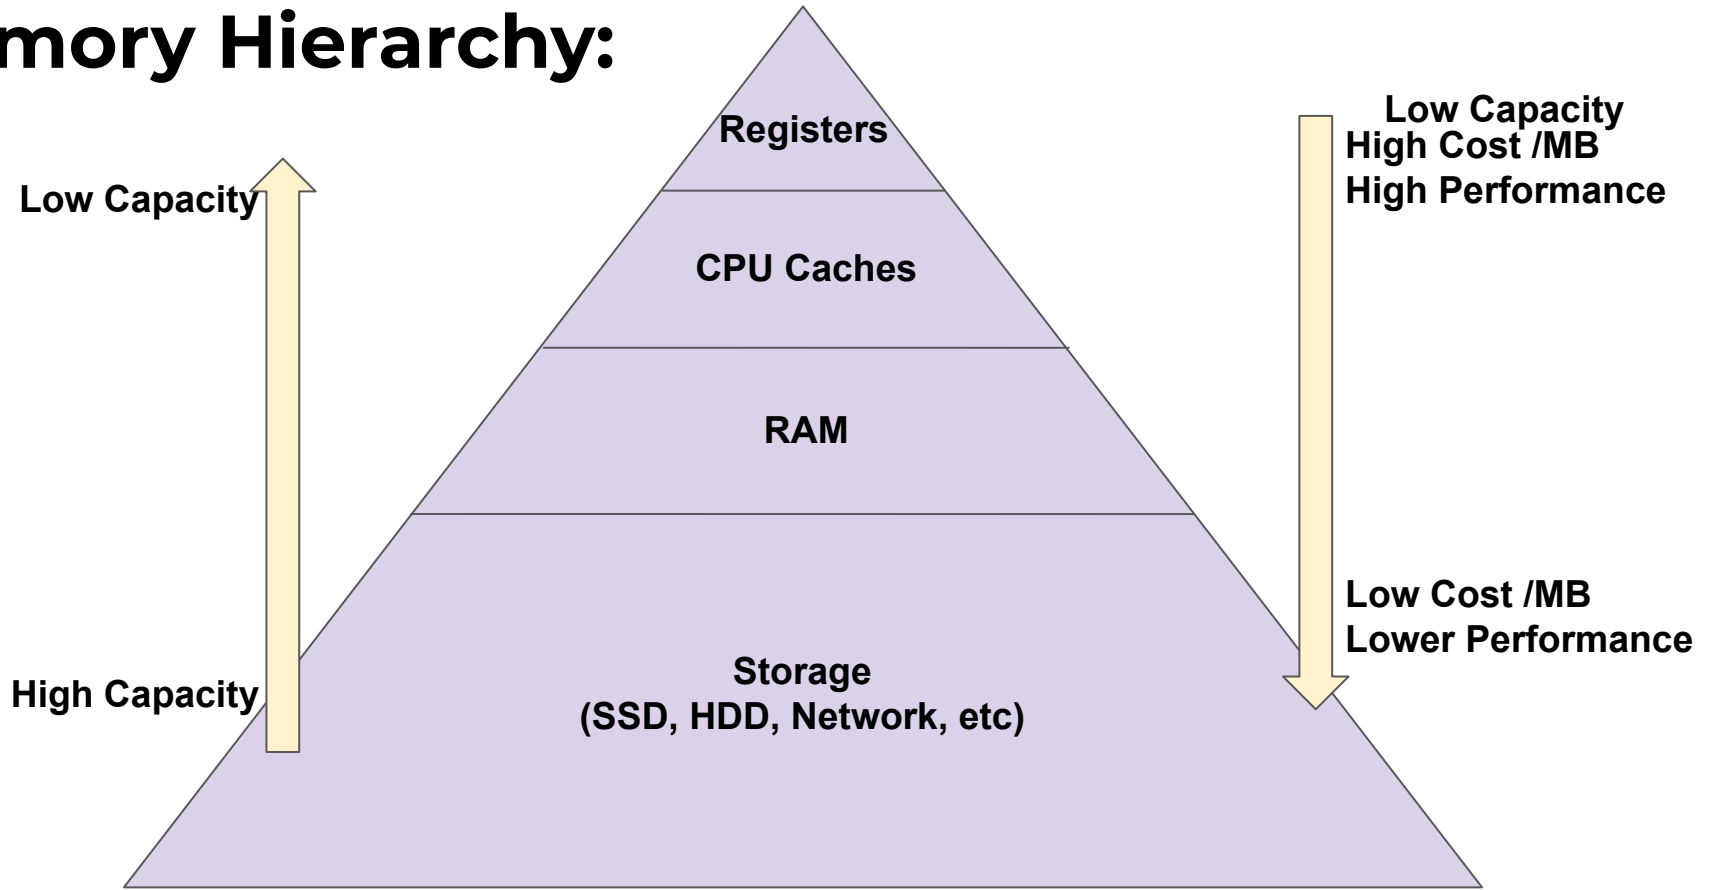
CS 423 - The University of Illinois

Wade Fagen-Ulmschneider

Memory Hierarchy:



Memory Hierarchy:



Memory Considerations

- ★ We have a **limited amount of fast** resources.
- ★ We have an **abundance of slow** resources.
- ★ How do we create an **illusion** of an **abundance of fast** resources?



ALMA MATER

TO THE UNIVERSITY OF CALIFORNIA

1907

Memory Overlays



CS 423 - The University of Illinois

Wade Fagen-Ulmschneider

Overlays

- ★ **Memory overlays** are fixed-sized segments of data used when a program exceeds the available memory.
- ★ Simple, minimal complexity; implemented at compile-time.
- ★ Still used in embedded systems.

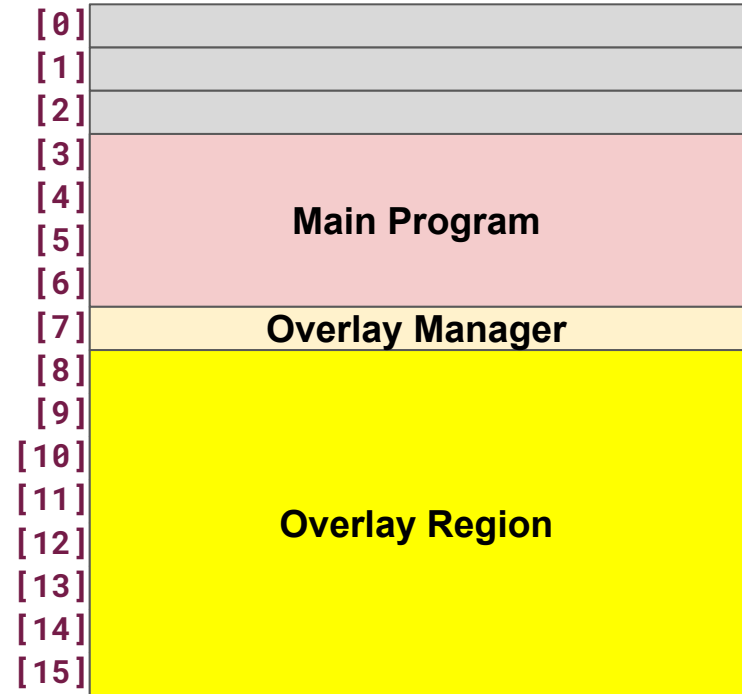
Physical Memory Pages:

[0]	
[1]	
[2]	
[3]	
[4]	
[5]	
[6]	
[7]	
[8]	
[9]	
[10]	
[11]	
[12]	
[13]	
[14]	
[15]	

Overlays

- ★ **Memory overlays** are fixed-sized segments of data used when a program exceeds the available memory.
- ★ Simple, minimal complexity; implemented at compile-time.
- ★ Still used in embedded systems.


Physical Memory Pages:



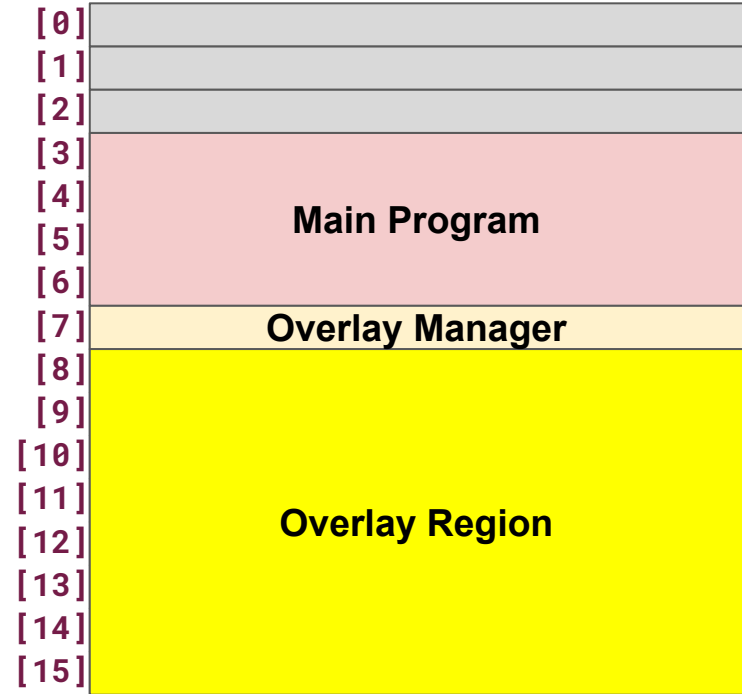
Overlays

Program's Overlays:

(Stored in secondary storage)



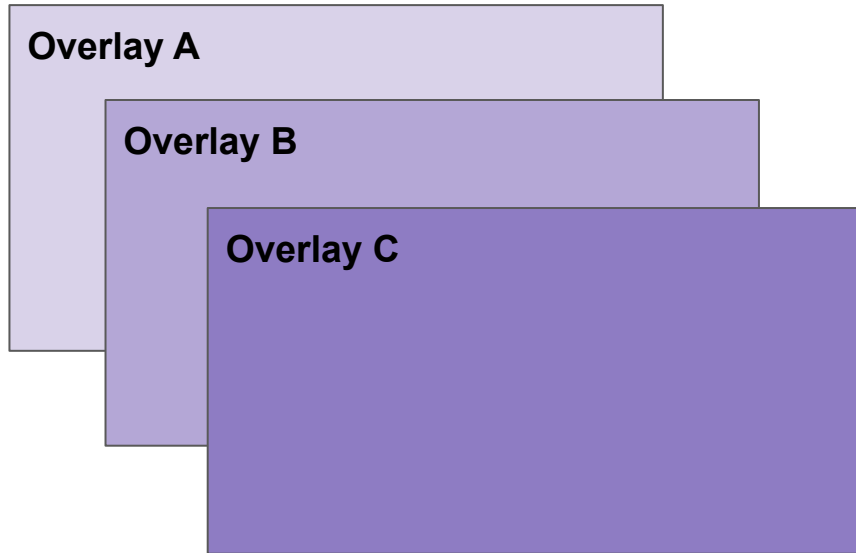
Physical Memory Pages:



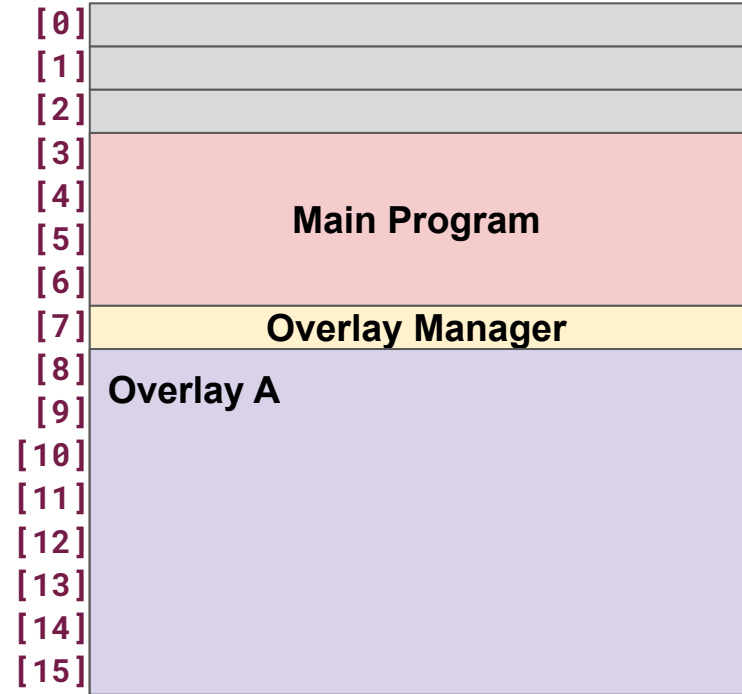
Overlays

Program's Overlays:

(Stored in secondary storage)



Physical Memory Pages:



Overlays

- ★ Systems may have multiple overlays and overlays are loaded in before they're required by the program code.
 - All modern compilers/linkers support overlays.
 - Compiled code target a specific overlay (ex: 2 x64 KB overlays).
- ★ **Disadvantages:**
 - Fixed size segments (ex: 64 KB),
 - Platform-specific (must compile for different segment sizes),
 - Raw access to RAM; limited process isolation.



ALMA MATER

TO THE UNIVERSITY OF CALIFORNIA

1907

Fixed Partitions

The background of the slide is a photograph of a statue in a park, overlaid with a semi-transparent red filter. The statue is a large, classical-style figure with outstretched arms, standing on a pedestal. The surrounding area is filled with the bare branches of trees, suggesting a winter or late autumn setting. The overall aesthetic is academic and formal.

CS 423 - The University of Illinois

Wade Fagen-Ulmschneider

Fixed Partitions

- ★ **Fixed Partitions** allocate a fixed amount of physical RAM to every process in a fixed location.

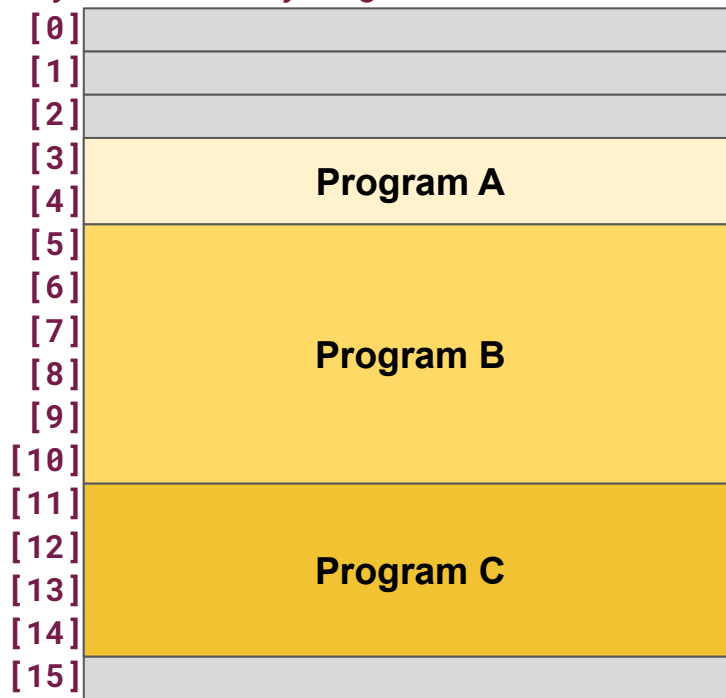
Physical Memory Pages:

[0]	
[1]	
[2]	
[3]	
[4]	
[5]	
[6]	
[7]	
[8]	
[9]	
[10]	
[11]	
[12]	
[13]	
[14]	
[15]	

Fixed Partitions

- ★ **Fixed Partitions** allocate a fixed amount of physical RAM to every process in a fixed location.
- ★ On creation, each process declares the **maximum** memory space it may need.
 - OS allocates a sequential amount of space for the process.

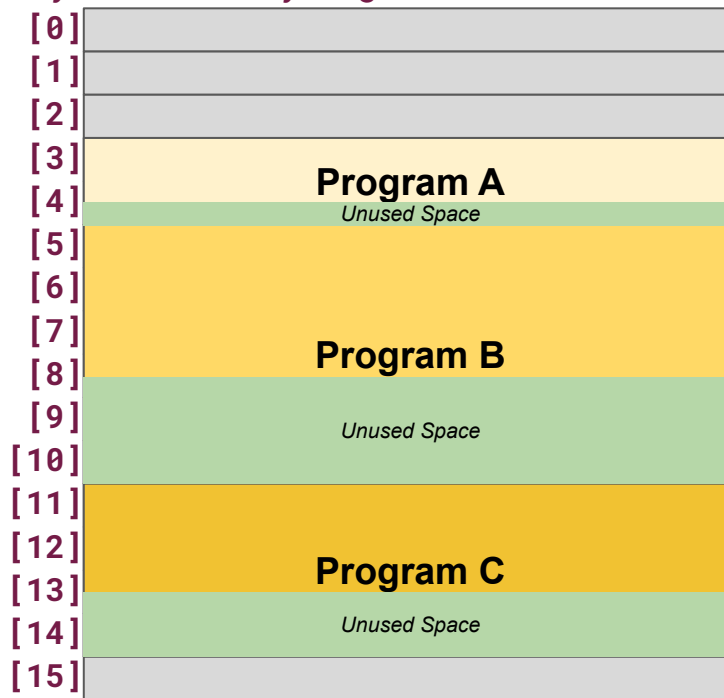
Physical Memory Pages:



Fixed Partitions

- ★ At any moment in time, a program may use only a part of its partition.

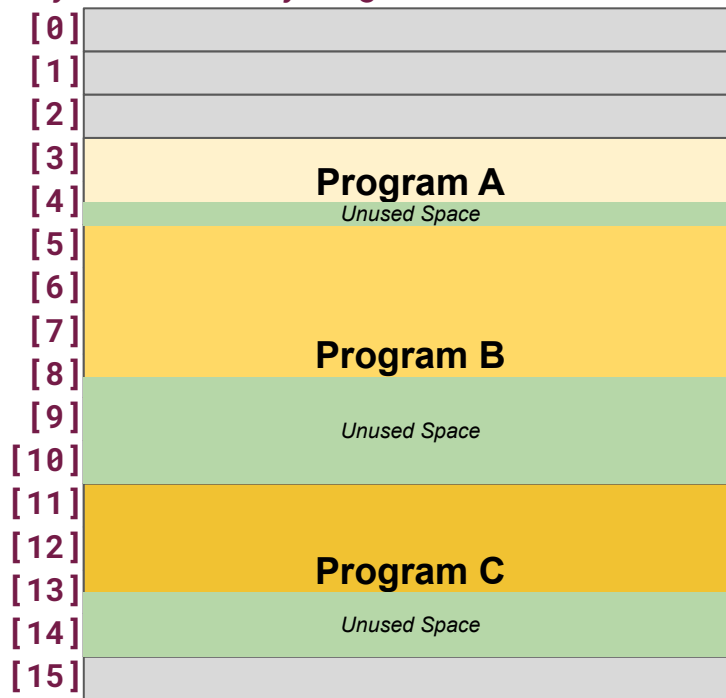
Physical Memory Pages:



Fixed Partitions

- ★ At any moment in time, a program may use only a part of its partition.
- ★ The unused space is **internal fragmentation** -- OS allocated the space, but process does not utilize it fully.

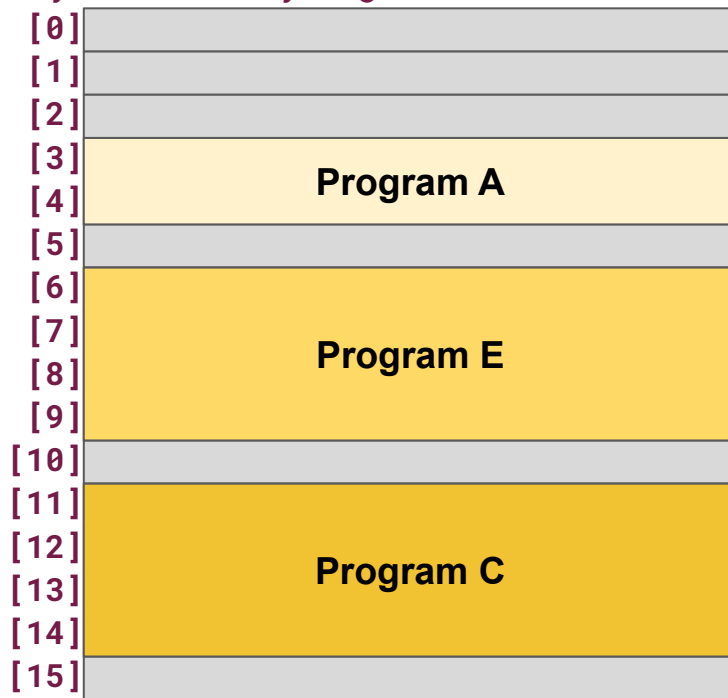
Physical Memory Pages:



Fixed Partitions

- ★ Additionally, the RAM will become fragmented with various sized holes as processes enter/exit.
 - *Some processes creation may be blocked until a partition is available.*
- ★ Raw access to RAM; limited process isolation.

Physical Memory Pages:





Relocation and Variable Partitions



CS 423 - The University of Illinois

Wade Fagen-Ulmschneider

Relocation

- ★ **Reallocation** provides a translation from a “offset (logical) address” to the physical address through the reallocation register.

Physical Memory Pages:


[0]	
[1]	
[2]	
[3]	
[4]	
[5]	
[6]	
[7]	
[8]	
[9]	
[10]	
[11]	
[12]	
[13]	
[14]	
[15]	

Relocation

- ★ All programs will address their memory from $0x0 \Rightarrow 0x\{\text{MAX}\}$.
- ★ The “offset” or “logical” address would be translated into the physical address by relocating the request:

$0x\ 3ac \Rightarrow 0x\mathbf{3c3a}3ac$
 +Relocation
 Register

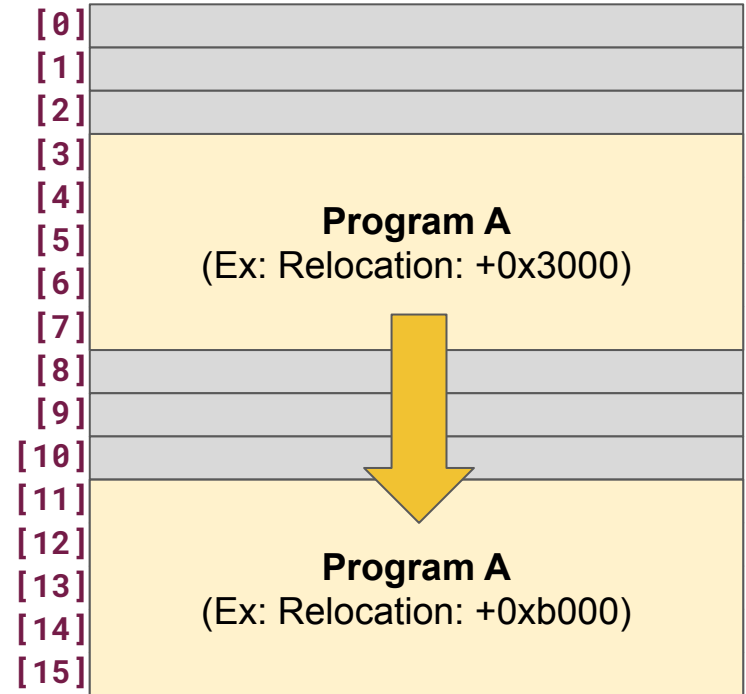
Physical Memory Pages:



Relocation

- ★ By changing the value of the relocation register, each process can now be moved around within RAM.

Physical Memory Pages:



Relocation

- ★ First system with a “translation” between a “logical address” and the “physical address” in RAM.
 - **Disadvantage:** Still requires sequential memory to be committed.
- ★ **Overhead:** Single offset is needed to translate the page; the offset can be adjusted by the OS as needed. **(Low overhead!)**



ALMA MATER

TO THE UNIVERSITY OF CALIFORNIA

1868

Paging and Virtual Memory



CS 423 - The University of Illinois

Wade Fagen-Ulmschneider

Paging

- ★ **Paging** is an extension of segmentation, where we divide all data on our system into **fixed-sized pages**.

Physical Memory Pages:

[0]	
[1]	
[2]	
[3]	
[4]	
[5]	
[6]	
[7]	
[8]	
[9]	
[10]	
[11]	
[12]	
[13]	
[14]	
[15]	

Paging

- ★ **Paging** is an extension of segmentation, where we divide all data on our system into **fixed-sized pages**.
 - Small enough to have minimal internal fragmentation.
 - Large enough to have minimal external fragmentation and overhead.

Physical Memory Pages:

[0]	
[1]	
[2]	
[3]	
[4]	
[5]	
[6]	
[7]	
[8]	
[9]	
[10]	
[11]	
[12]	
[13]	
[14]	
[15]	

Paging

- ★ **Linux: `getconf PAGESIZE`**
 - Reports the size of the page on a system.
 - Most systems use 2^{12} , or 4096 B/page.
 - *I've started to see 2^{16} (64 KiB) used in the wild more and more.*

```
$ getconf PAGESIZE
4096
```

Physical Memory Pages:

[0]	
[1]	
[2]	
[3]	
[4]	
[5]	
[6]	
[7]	
[8]	
[9]	
[10]	
[11]	
[12]	
[13]	
[14]	
[15]	

Page Tables

P1 Page Table: RAM:

[0]		[0]	
[1]		[1]	
[2]		[2]	
[3]		[3]	
[4]		[4]	
[5]		[5]	
[6]		[6]	
[7]		[7]	
[8]		[8]	
[9]		[9]	
[10]		[10]	
[11]		[11]	
[12]		[12]	
[13]		[13]	
[14]		[14]	
[15]		[15]	

P2 Page Table:

[0]	
[1]	
[2]	
[3]	
[4]	
[5]	
[6]	
[7]	
[8]	
[9]	
[10]	
[11]	
[12]	
[13]	
[14]	
[15]	


P3 Page Table:

[0]	
[1]	
[2]	
[3]	
[4]	
[5]	
[6]	
[7]	
[8]	
[9]	
[10]	
[11]	
[12]	
[13]	
[14]	
[15]	

Page Table

- ★ Every process has its own **page table**.
 - Page table provides a translation between a “virtual address” used by the program and the “physical address” on RAM.
 - No user process will ever see the real physical address!

Page Tables



Page Table

- ★ Every “virtual address” now has two components:
 - **Page Offset:** Where, within the page, is the data I’m addressing?
 - **Page Number:** What index is our page within our virtual page table?

Page Table

- ★ If our page is 2^{12} bytes in size, the lowest 12 bits of a memory address is the **page offset**.
- ★ The remaining bits is the **page number**.

Page Table

- ★ If our page is 2^{12} bytes in size, the lowest 12 bits of a memory address is the **page offset**.
- ★ The remaining bits is the **page number**.

Memory Address: **0x 32ac 51c16**

Page Number: 0x 32ac51

Page Offset: 0x c16 (lowest 12 bits)


Page Table

Memory Address: **0x 32ac 51c16**

<p>Page Number: 0x 32ac51 Page Offset: 0x c16 (lowest 12 bits)</p>
--

- ★ The page table entry at **0x32ac51** will provide the translation to the physical address in RAM.

Page Tables



Page Table

★ Advantages:

- Processes can have the **allusion of sequential memory** even though the pages may be located in (*translated to*) various locations throughout RAM.
- **Pages do not have to always be “present” in RAM;** can point to data on storage and load it in RAM when needed.
 - Need a mechanism to load pages in when needed.
- Processes have no direct access to RAM; allows OS to provide protections to RAM.

Page Table

★ Disadvantages:

- Overhead:
 - Consider 4 GiB of RAM divided into 4 KiB pages:
 $4 \text{ GiB} / 4 \text{ KiB} == 1 \text{ MiB pages (!!)}$
...each process has its own 1 MiB pages!
- Complexity: Non-trivial to translate addresses.



ALMA MATER

TO THE UNIVERSITY OF CALIFORNIA

AND THE STATE OF CALIFORNIA

Page Faults and Page Evictions

The background of the slide features a photograph of a large, classical-style statue in a park setting, possibly the 'The Spirit of the Law' statue at the University of Illinois. The entire image is overlaid with a semi-transparent red filter. The statue is the central focus, with its arms outstretched. The surrounding area shows the silhouettes of trees and other figures in the distance.

CS 423 - The University of Illinois

Wade Fagen-Ulmschneider

Virtual Memory Analysis

What is the range of possible file sizes for img.png?

Disk Pages:

...
PC1
PC2
PC3
PC4
PC5
PC6
img.png
img.png
img.png
...


Virtual Memory Analysis

What is the range of possible file sizes for `./programCode` (“PC”)?

Disk Pages:

...
PC1
PC2
PC3
PC4
PC5
PC6
img.png
img.png
img.png
...


Page Eviction/Replacement Strategies:



Page Access: 17 33 40 17 43 8 99 33 99 17


Page Faults

Page Eviction/Replacement Strategies:




Page Access: 17 33 40 17 43 8 99 33 99 17

Page Eviction/Replacement Strategies:




Page Access: 17 33 40 17 43 8 99 33 99 17

Page Eviction/Replacement Strategies:



Page Access: 17 33 40 17 43 8 99 33 99 17

Page Eviction/Replacement Strategies:



Page Access: 17 33 40 17 43 8 99 33 99 17