



CS162 - Operating Systems and Systems Programming

Address Translation - Continued

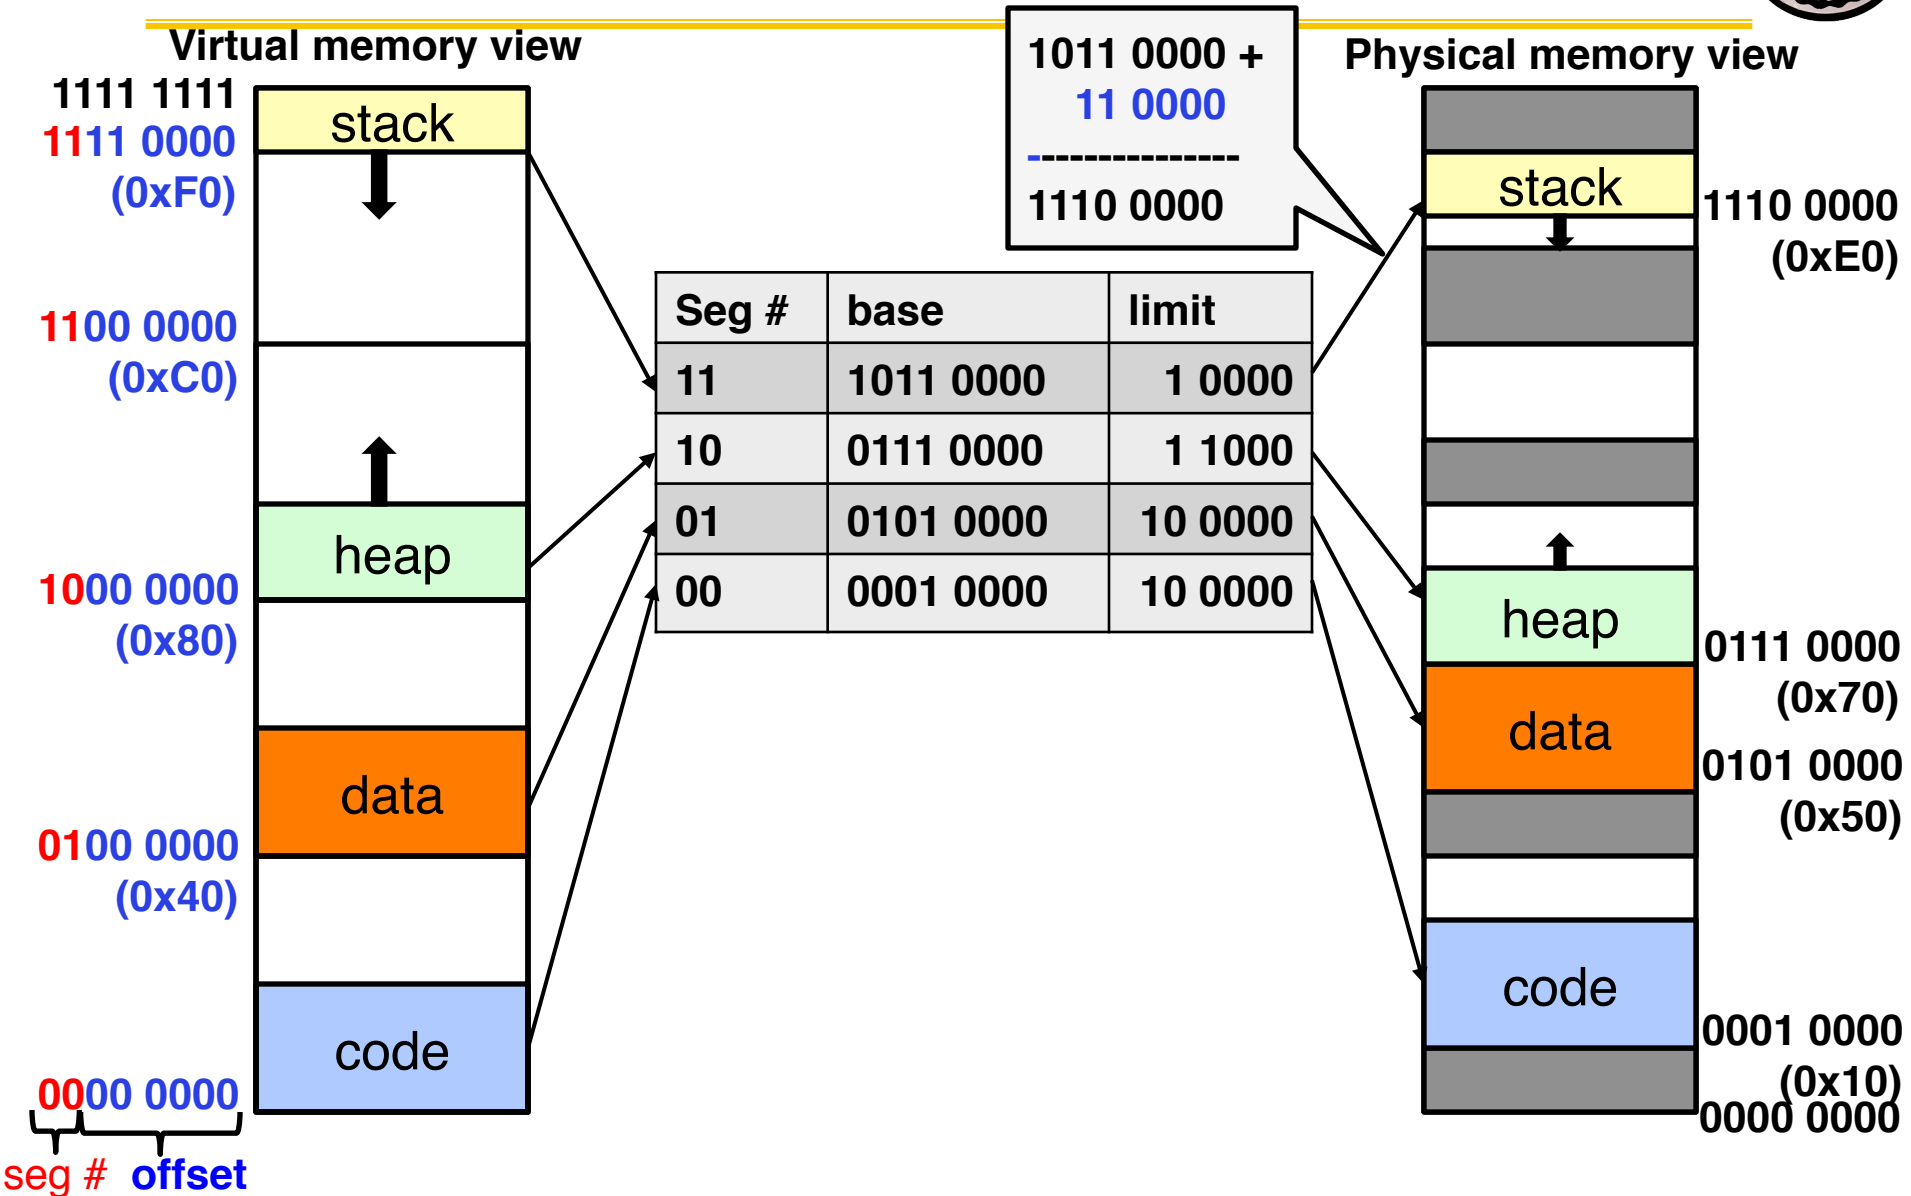
David E. Culler

<http://cs162.eecs.berkeley.edu/>

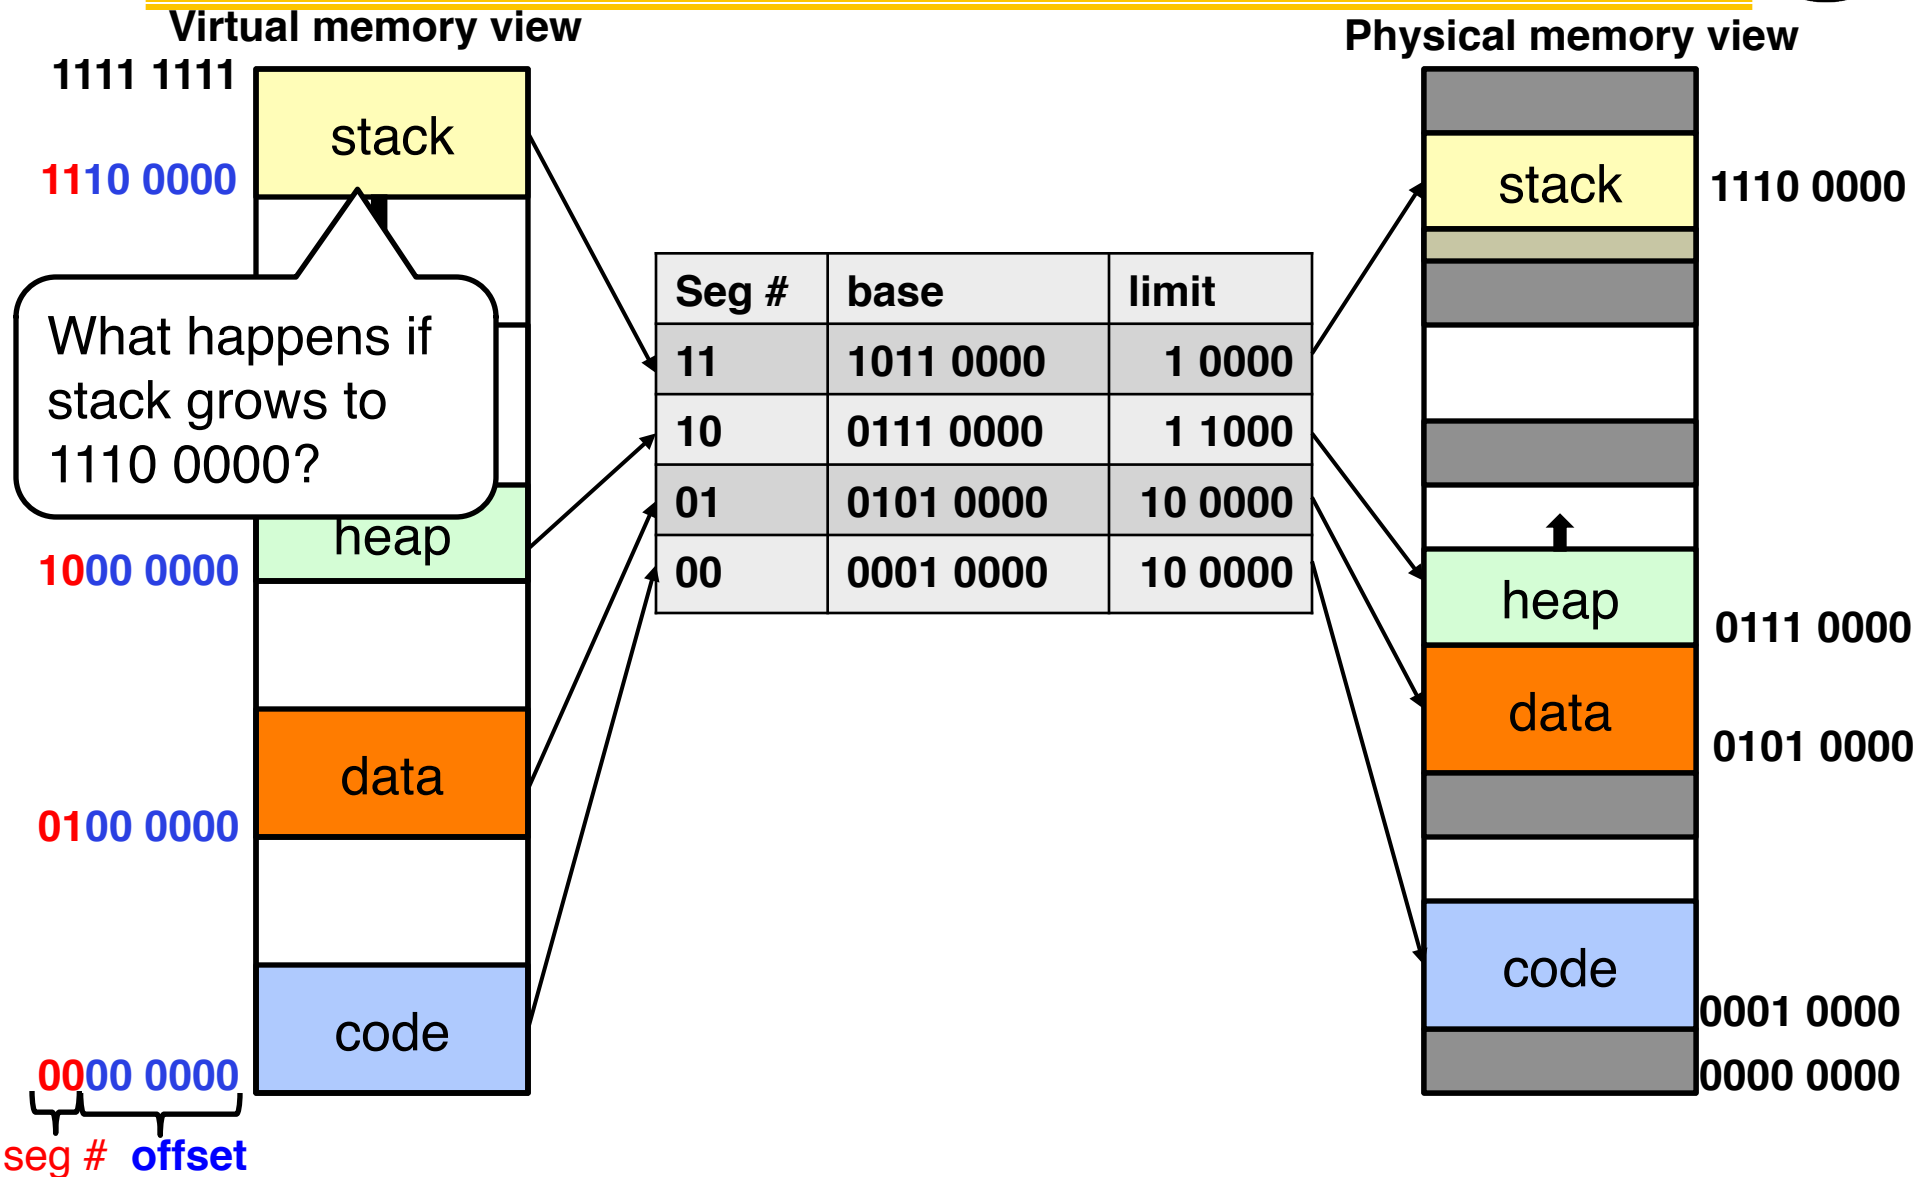
Lecture #15

Oct 1, 2014

Summary: Address Segmentation



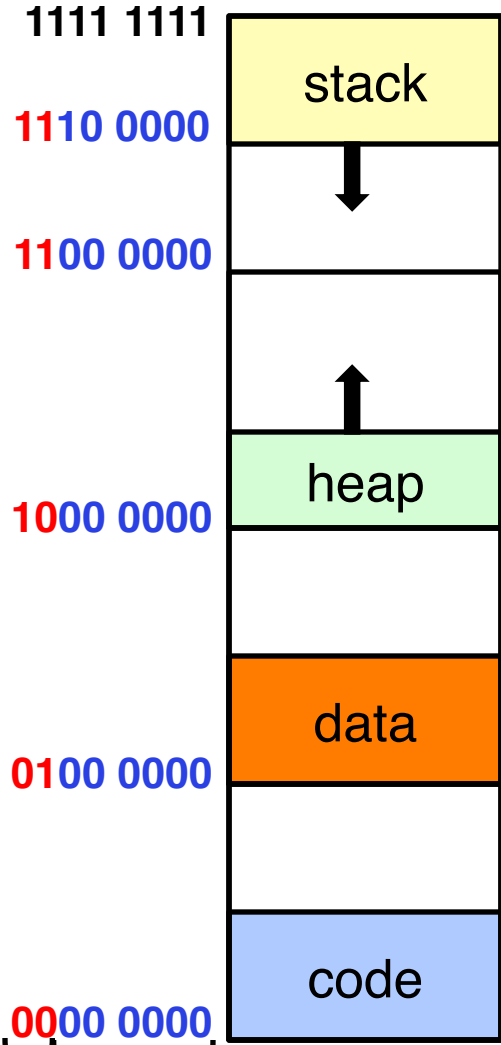
Summary: Address Segmentation



Recap: Address Segmentation




Virtual memory view



| Seg # | base | limit |
|-------|-----------|---------|
| 11 | 1011 0000 | 1 0000 |
| 10 | 0111 0000 | 1 1000 |
| 01 | 0101 0000 | 10 0000 |
| 00 | 0001 0000 | 10 0000 |

Physical memory view



No room to grow!!
 Buffer overflow error or
 resize segment **and**
 move segments around
 to make room

seg # offset




How do we run more programs than fit in memory ?



Schematic View of “Swapping”

- Q: What if not all processes fit in memory?
- A: Swapping: Extreme form of Context Switch
 - In order to make room for next process, some or all of the previous process is moved to disk
 - This greatly increases the cost of context-switching



- Desirable alternative?
 - Some way to keep only active portions of a process in memory at any one time
 - Need finer granularity control over physical memory

Problems with Segmentation



- Must fit variable-sized chunks into physical memory
- May move processes multiple times to fit everything
- Limited options for swapping to disk
- **Fragmentation**: wasted space
 - **External**: free gaps between allocated chunks
 - **Internal**: don't need all memory within allocated chunks


Paging: Physical Memory in Fixed Size Chunks



- Solution to fragmentation from segments?
 - Allocate physical memory in fixed size chunks (“pages”)
 - Every chunk of physical memory is equivalent
 - » Can use simple vector of bits to handle allocation:
00110001110001101 ... 110010
 - » Each bit represents page of physical memory
1⇒allocated, 0⇒free
- Should pages be as big as our previous segments?
 - No: Can lead to lots of internal fragmentation
 - » Typically have small pages (1K-16K)
 - Consequently: need multiple pages/segment



How to Implement Paging?



- Page Table (One per process)
 - Resides in physical memory
 - Contains physical page and permission for each virtual page
 - » Permissions include: Valid bits, Read, Write, etc
- Virtual address mapping
 - Offset from Virtual address copied to Physical Address
 - » Example: 10 bit offset ⇒ 1024-byte pages
 - Virtual page # is all remaining bits
 - » Example for 32-bits: 32-10 = 22 bits, i.e. 4 million entries
 - » Physical page # copied from table into physical address
 - Check Page Table bounds and permissions



What about Sharing?

Virtual Address
(Process A):




PageTablePtrA

| | |
|---------|-------|
| page #0 | V,R |
| page #1 | V,R |
| page #2 | V,R,W |
| page #3 | V,R,W |
| page #4 | N |
| page #5 | V,R,W |

PageTablePtrB

| | |
|---------|-------|
| page #0 | V,R |
| page #1 | N |
| page #2 | V,R,W |
| page #3 | N |
| page #4 | V,R |
| page #5 | V,R,W |

Virtual Address
(Process B):

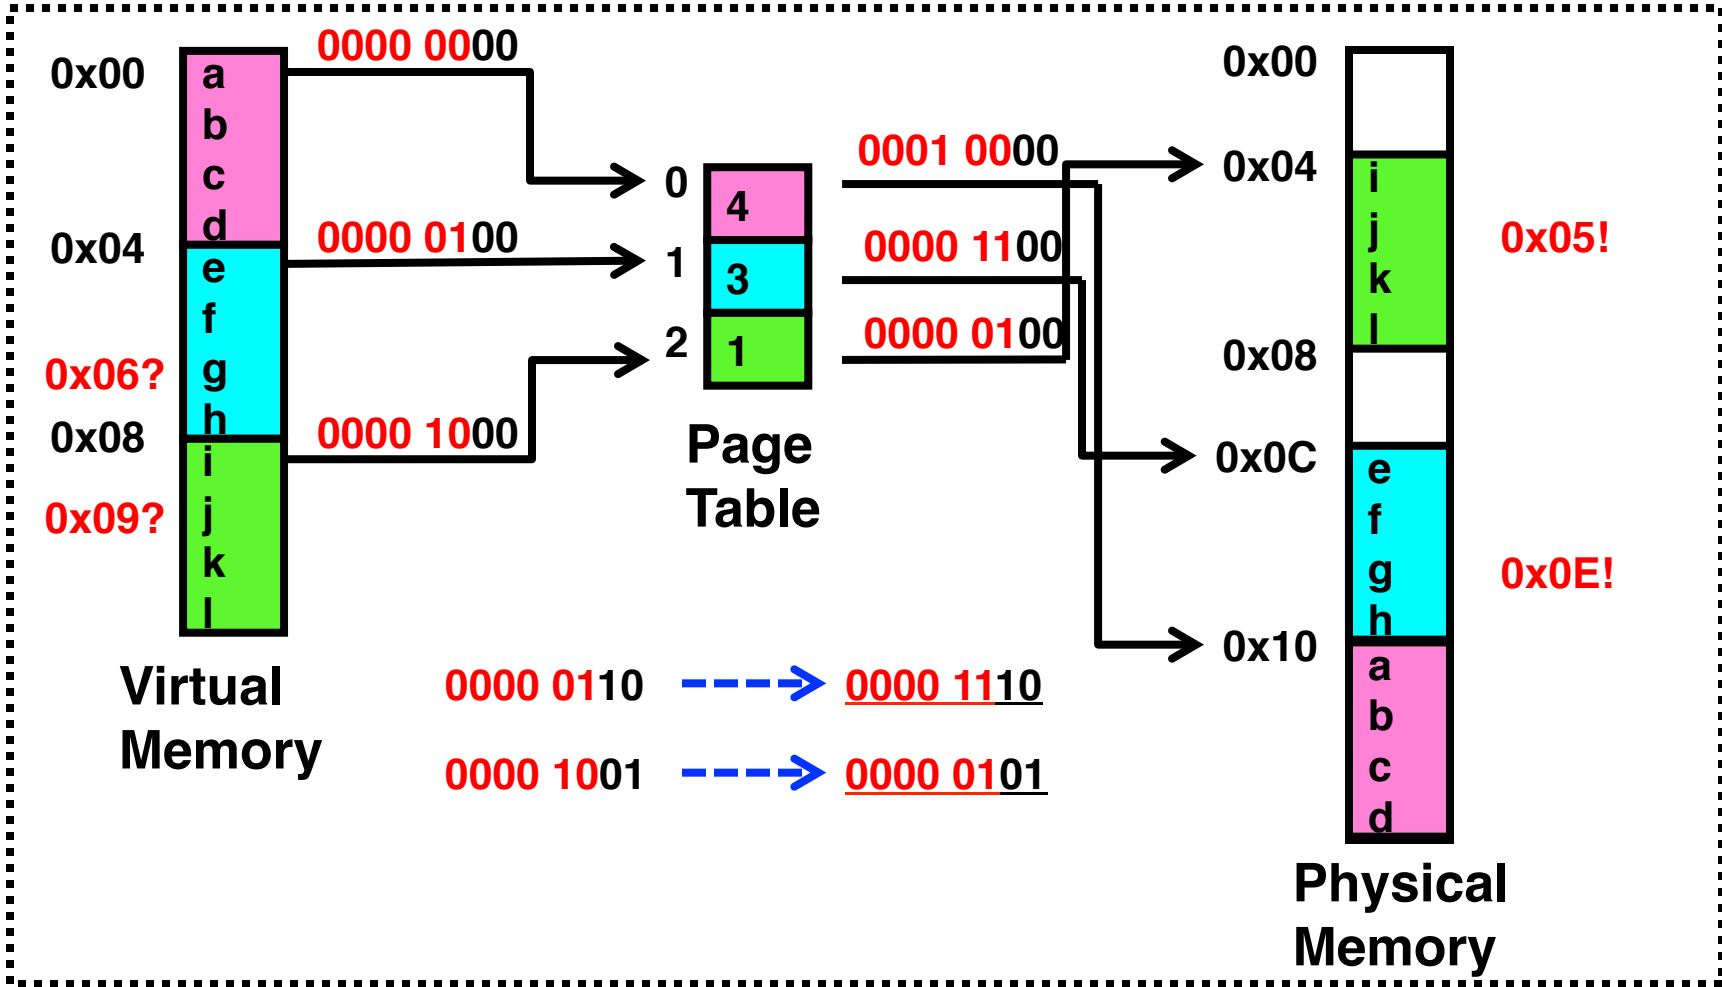


This physical page appears in address space of both processes



Simple Page Table Example

Example (4 byte pages)

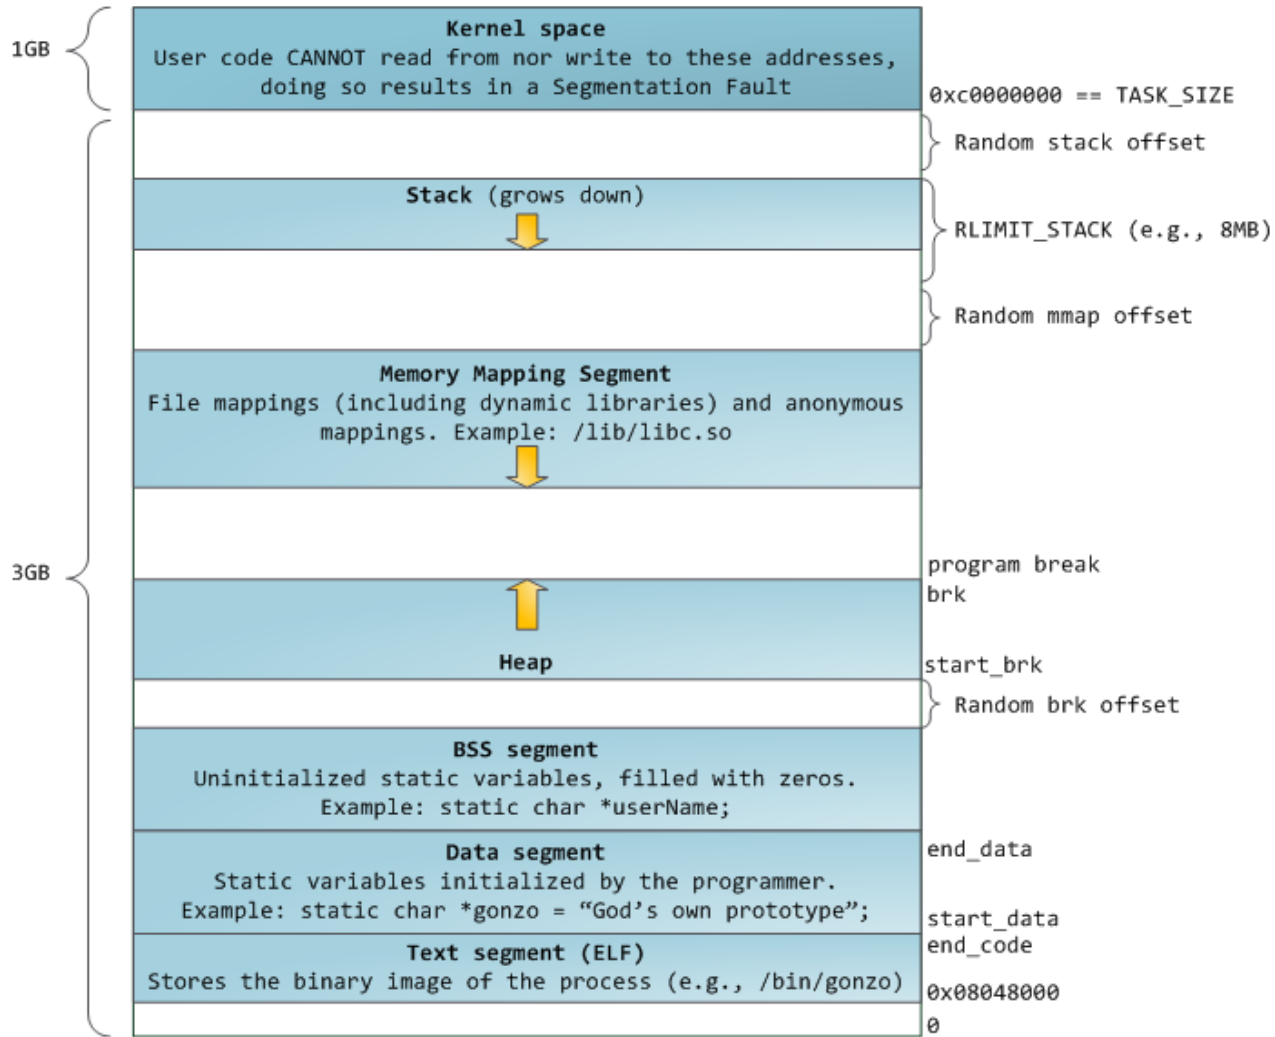


Page Table Discussion



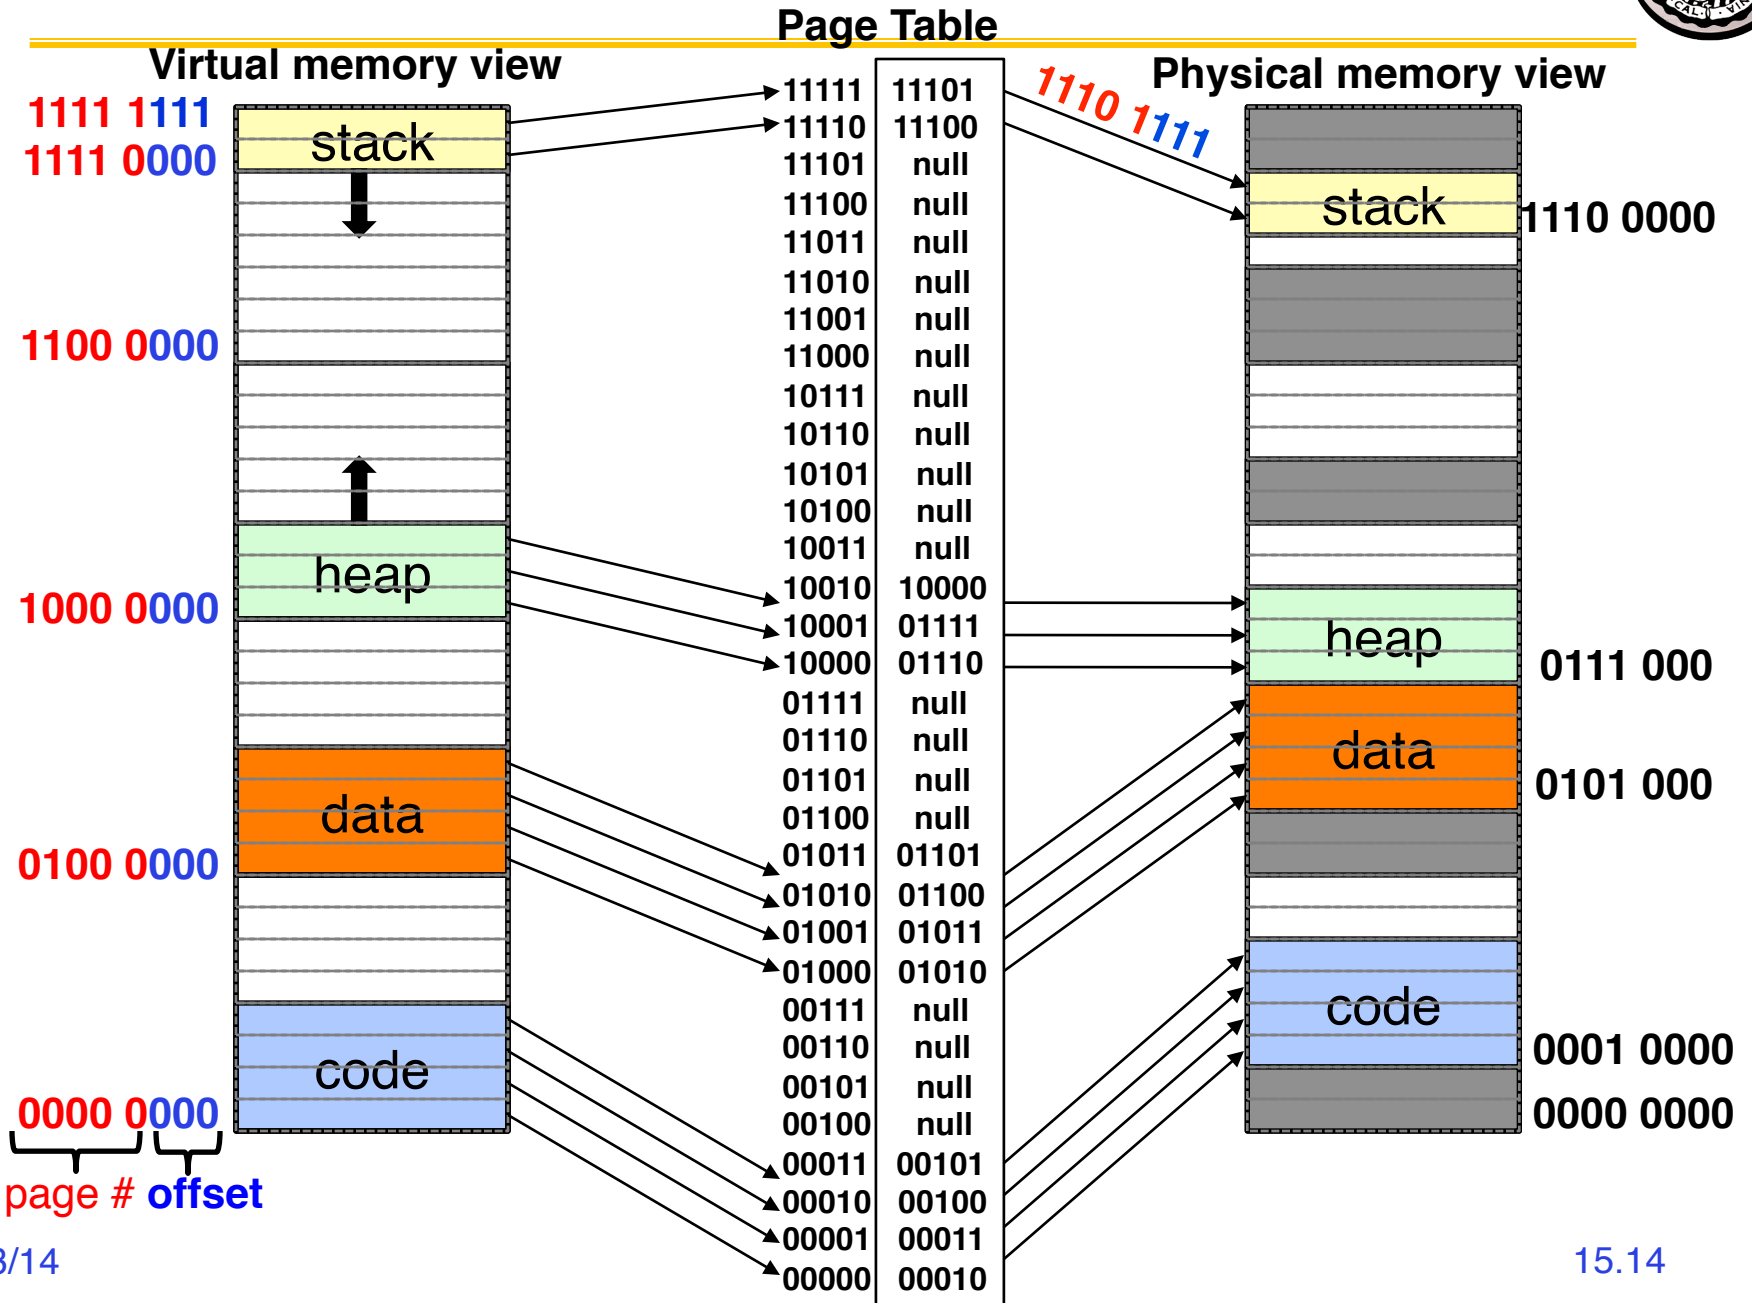
- What needs to be switched on a context switch?
 - Page table pointer and limit
- Analysis
 - Pros
 - » Simple memory allocation
 - » Easy to Share
 - Con: What if address space is sparse?
 - » E.g. on UNIX, code starts at 0, stack starts at $(2^{31}-1)$.
 - » With 1K pages, need 2 million page table entries!
 - Con: What if table really big?
 - » Not all pages used all the time \Rightarrow would be nice to have working set of page table in memory
- How about combining paging and segmentation?

E.g., Linux 32-bit

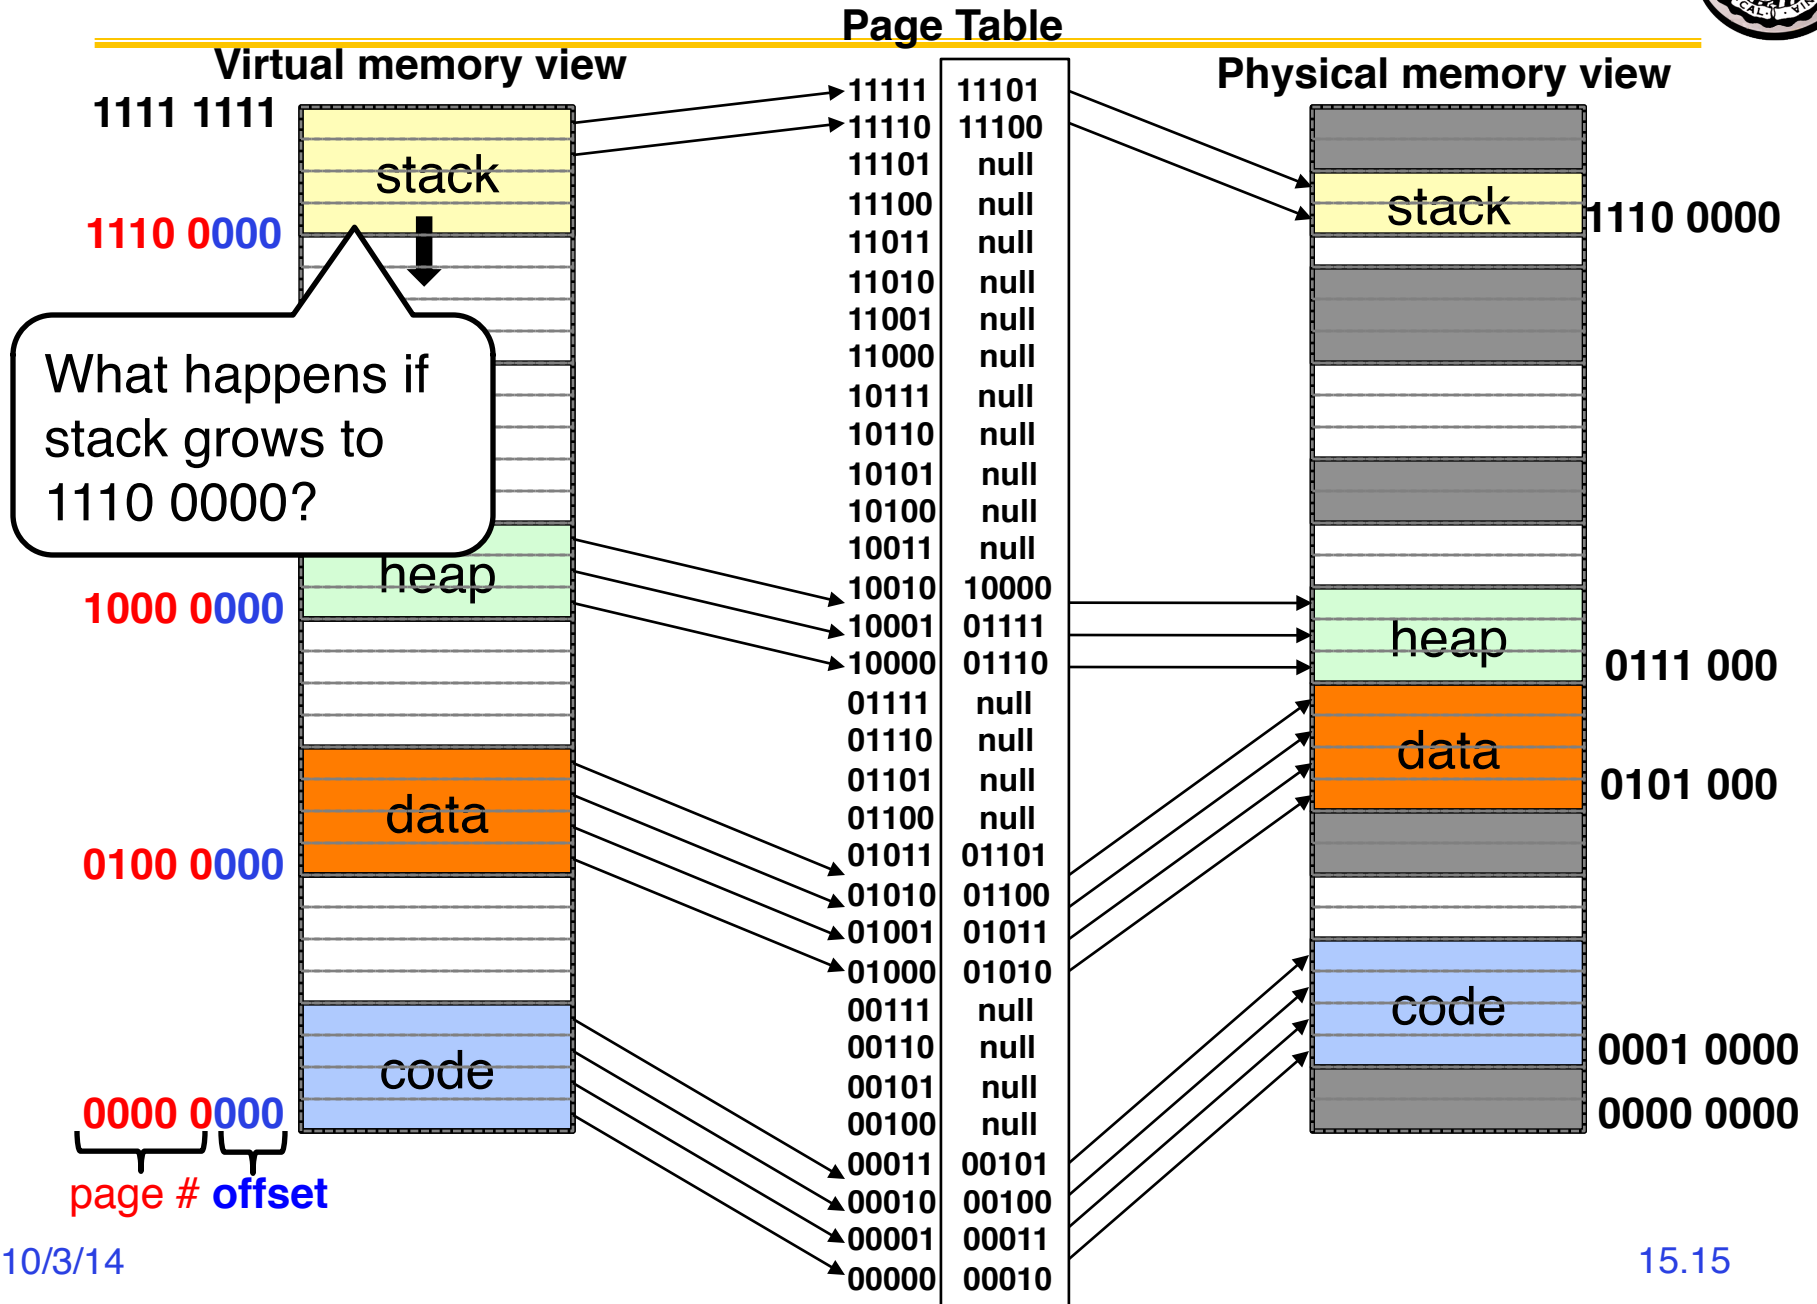


<http://static.duartes.org/img/blogPosts/linuxFlexibleAddressSpaceLayout.png>

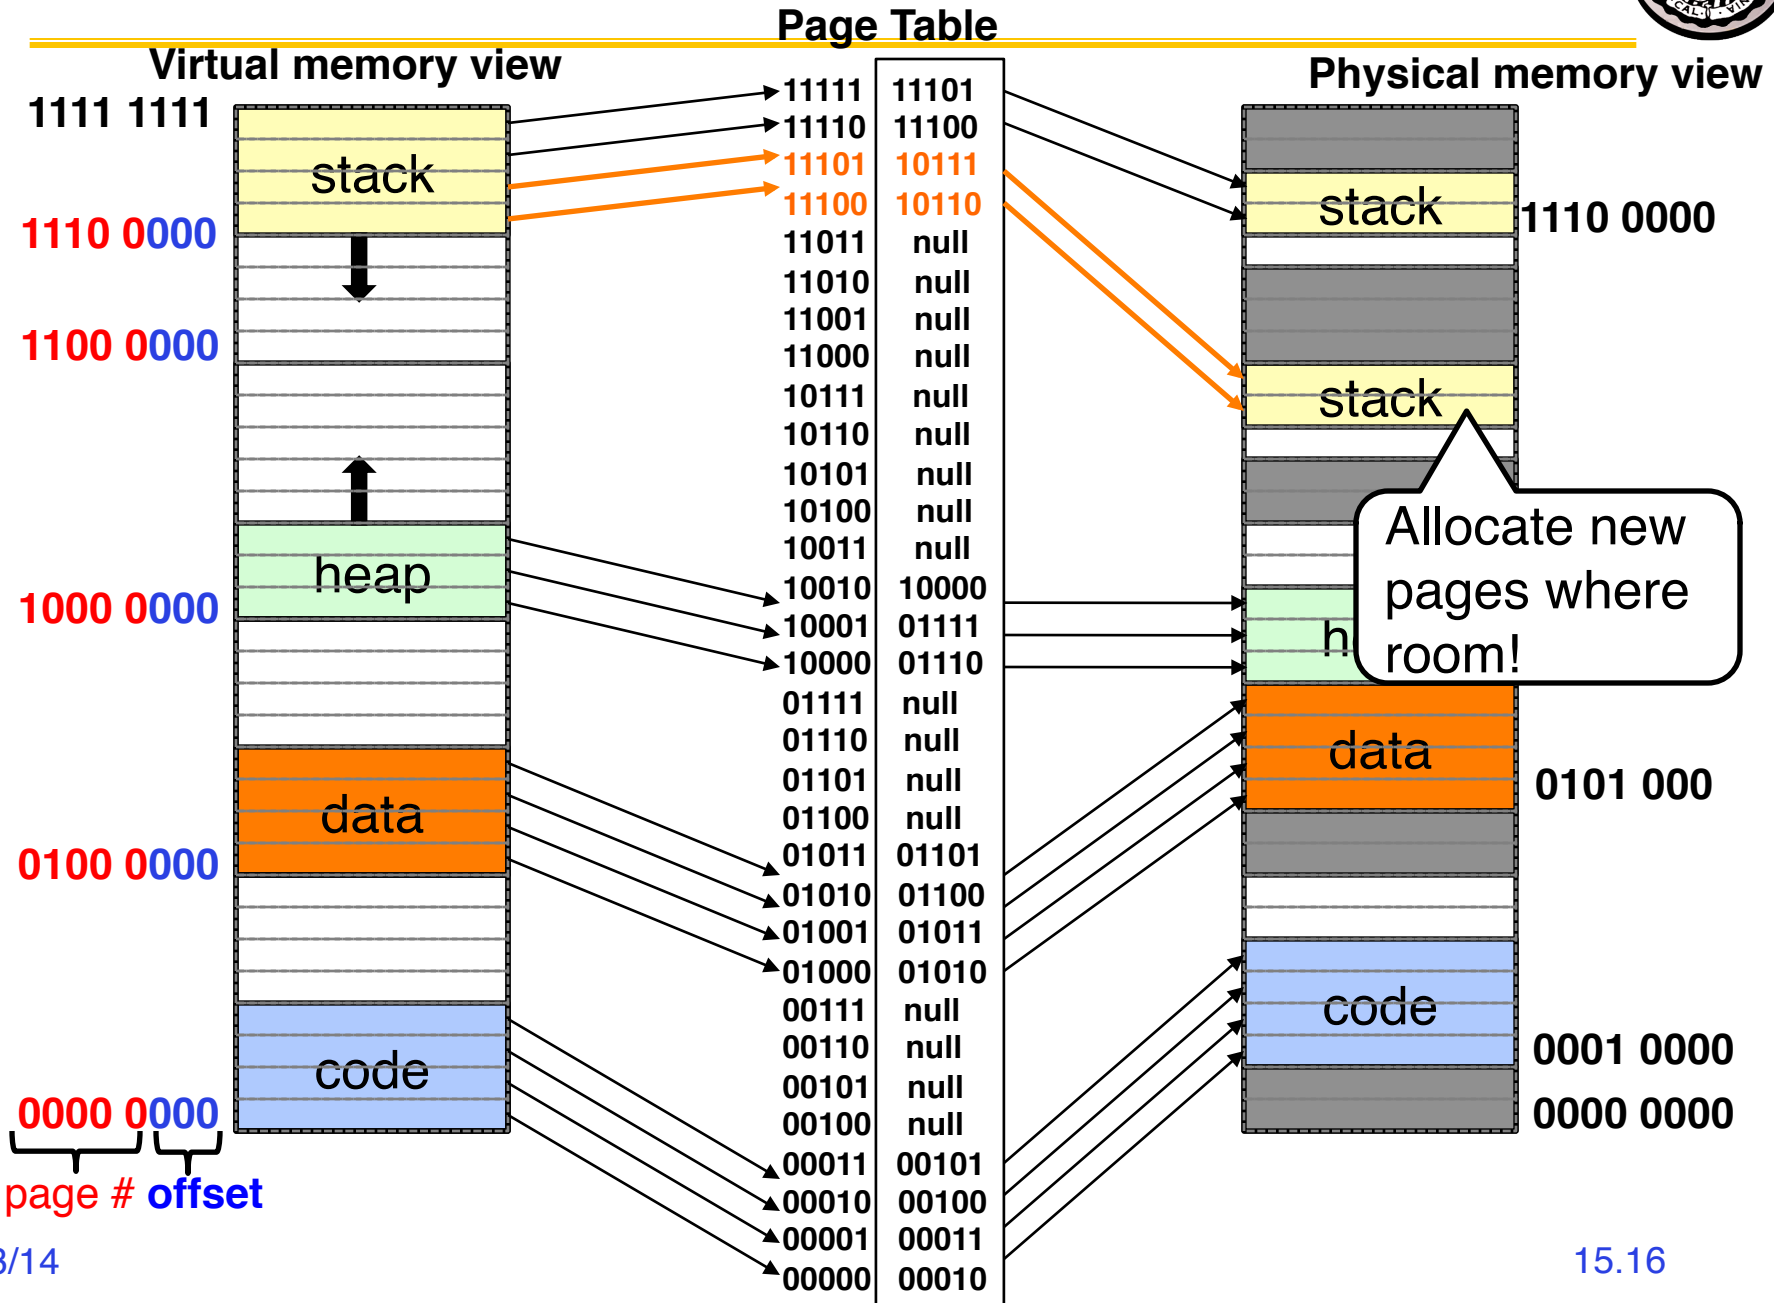
Summary: Paging



Summary: Paging



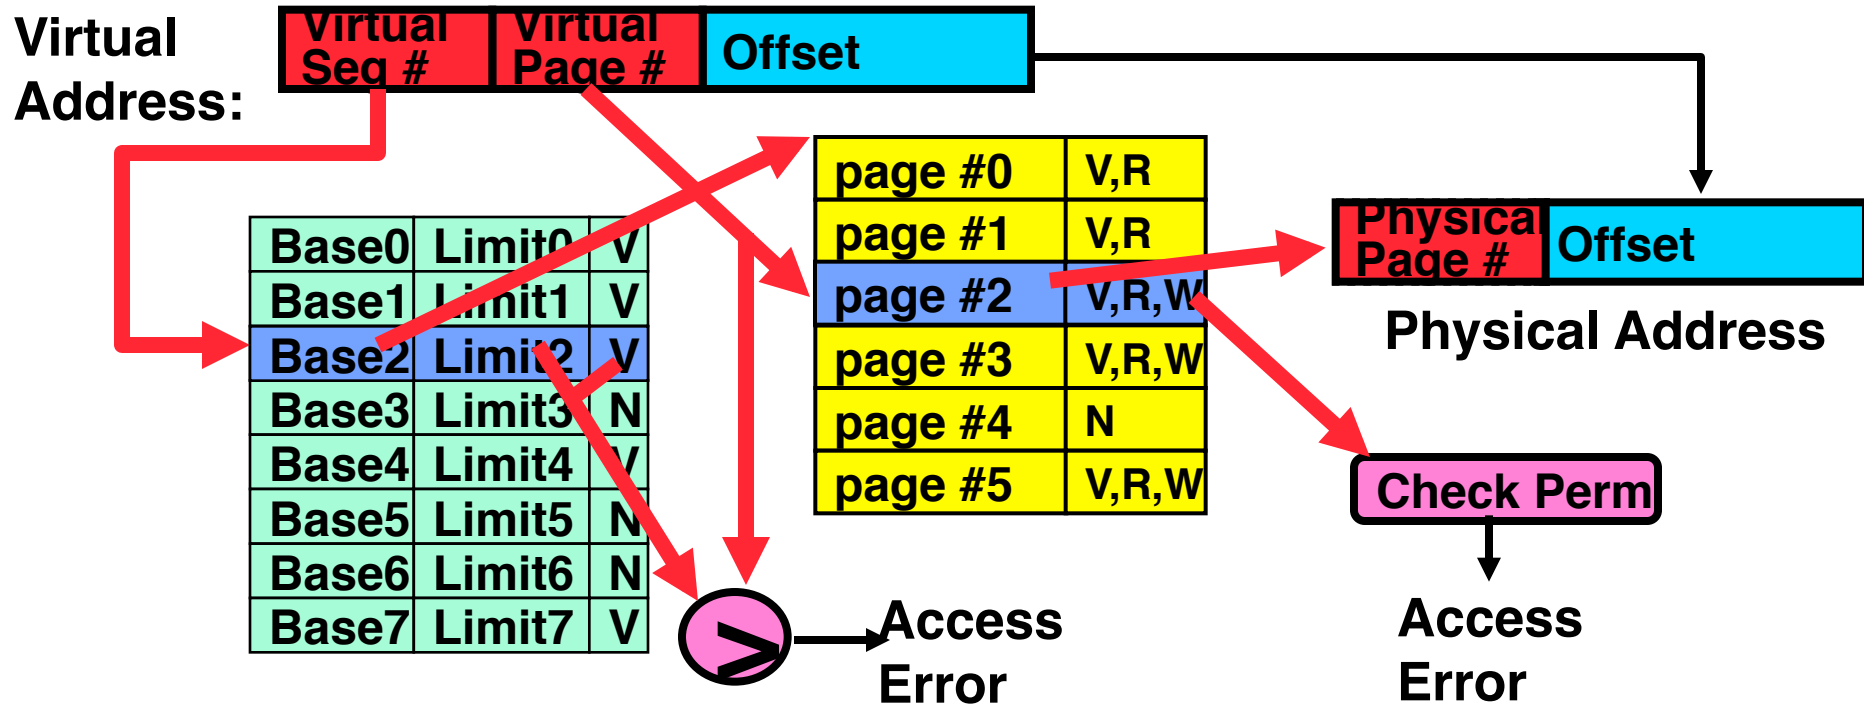
Summary: Paging





Multi-level Translation

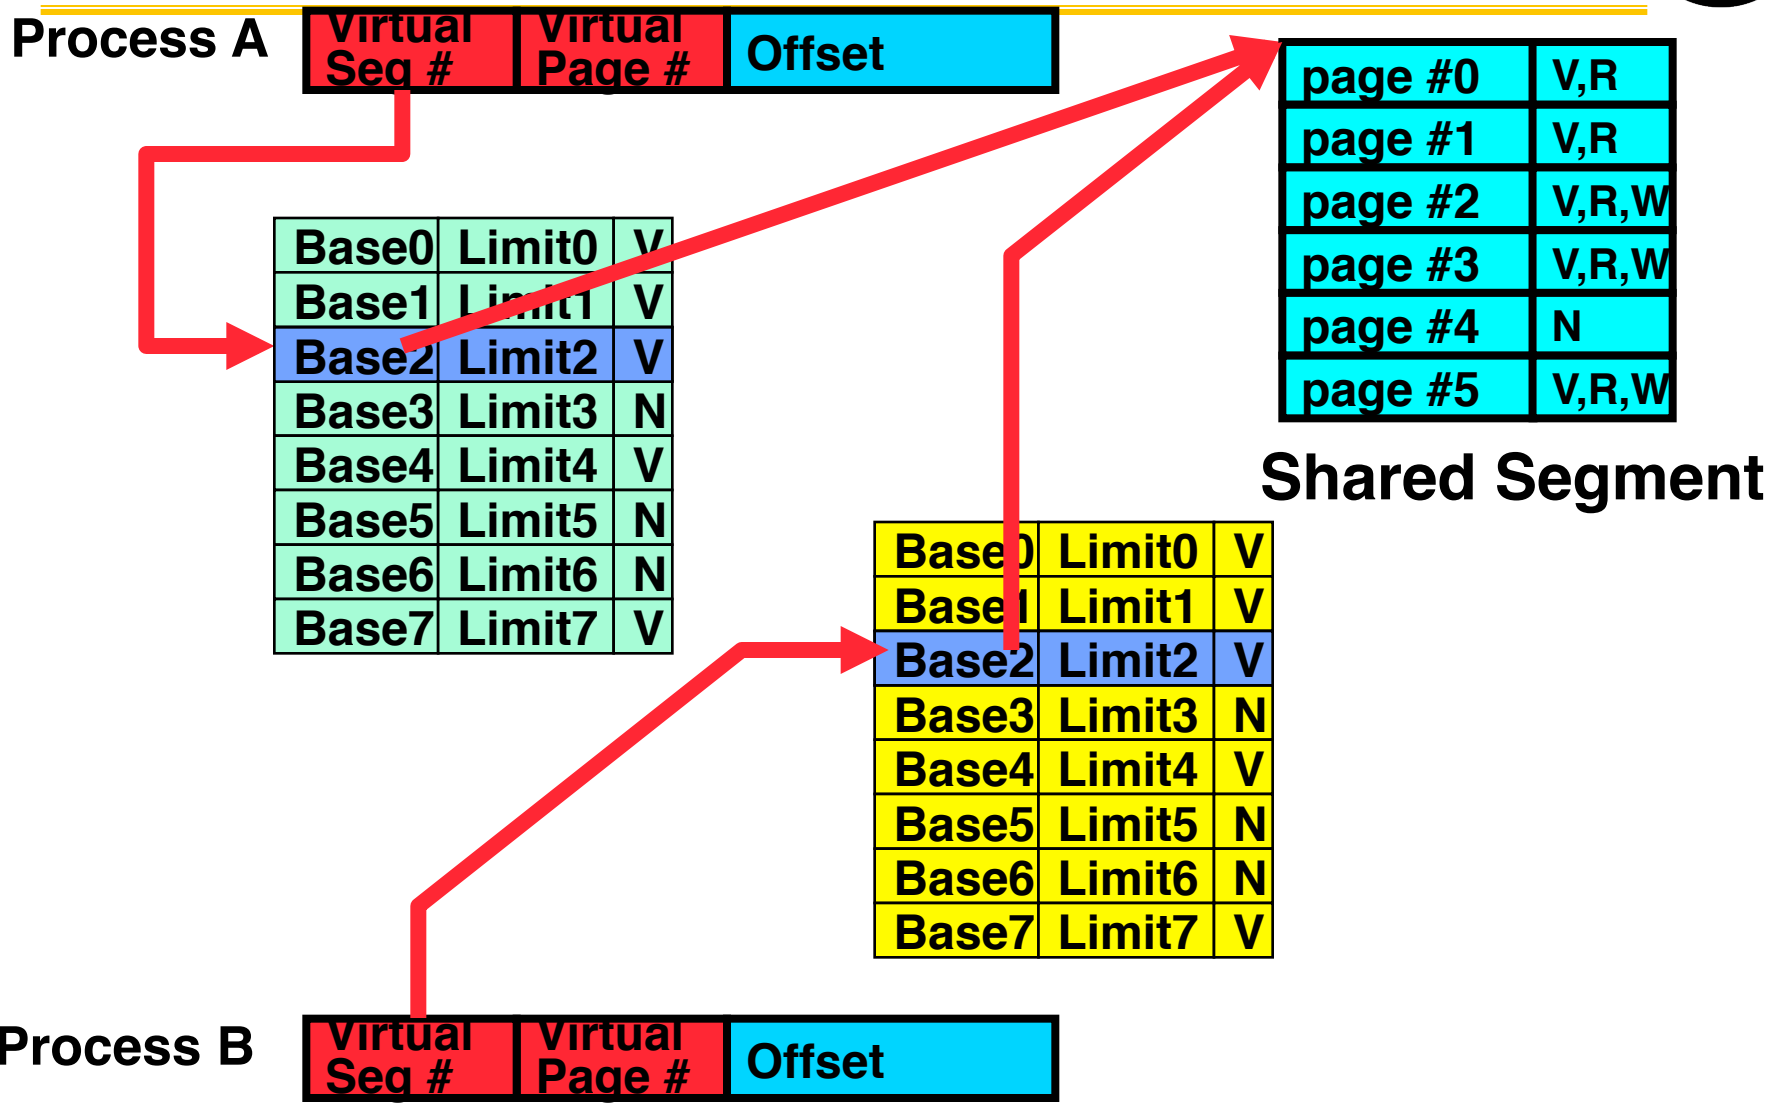
- What about a tree of tables?
 - Lowest level page table \Rightarrow memory still allocated with bitmap
 - Higher levels often segmented
- Could have any number of levels. Example (top segment):



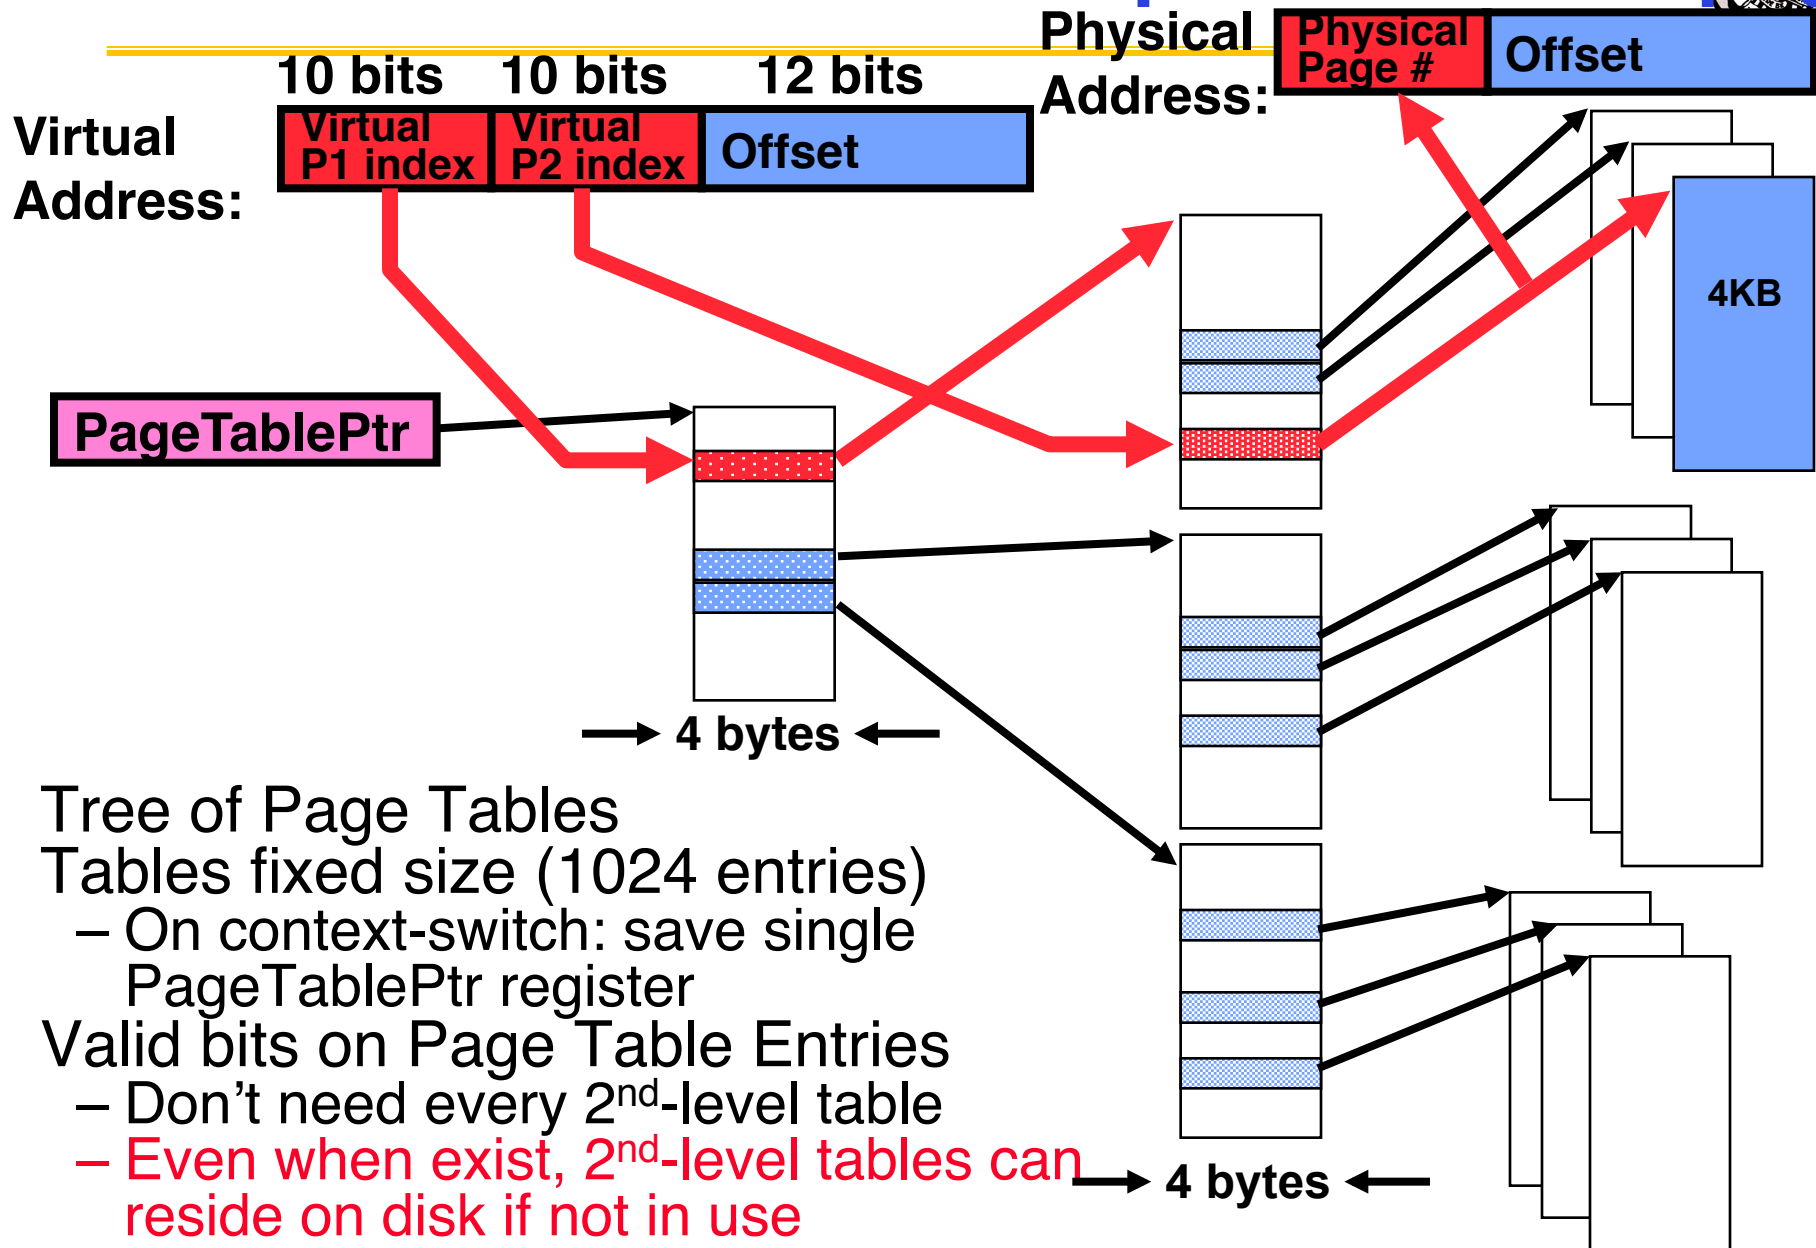
- What must be saved/restored on context switch?
 - Contents of top-level segment registers (for this example)
 - Pointer to top-level table (page table)



What about Sharing (Complete Segment)?



Another common example: two-level page



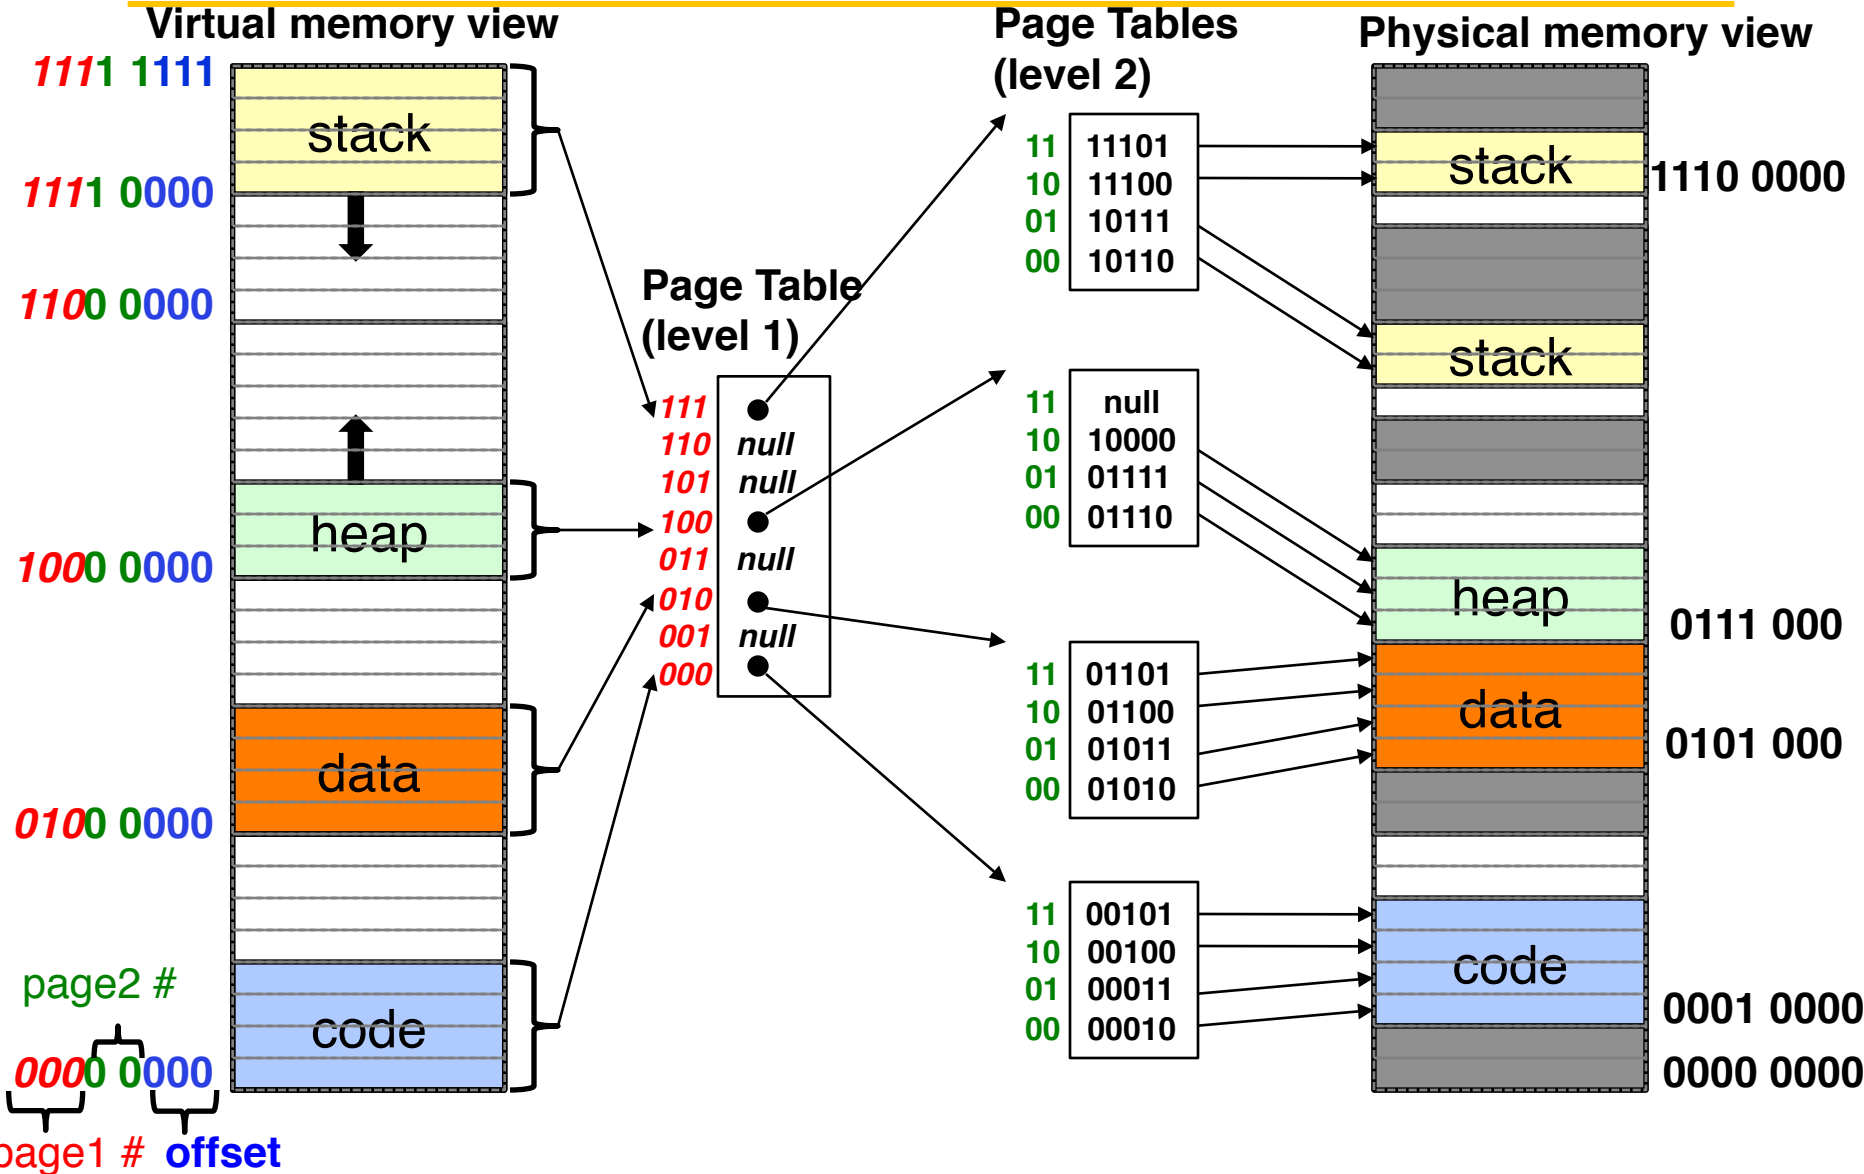
- Tree of Page Tables
- Tables fixed size (1024 entries)
 - On context-switch: save single PageTablePtr register
- Valid bits on Page Table Entries
 - Don't need every 2nd-level table
 - Even when exist, 2nd-level tables can reside on disk if not in use

Multi-level Translation Analysis

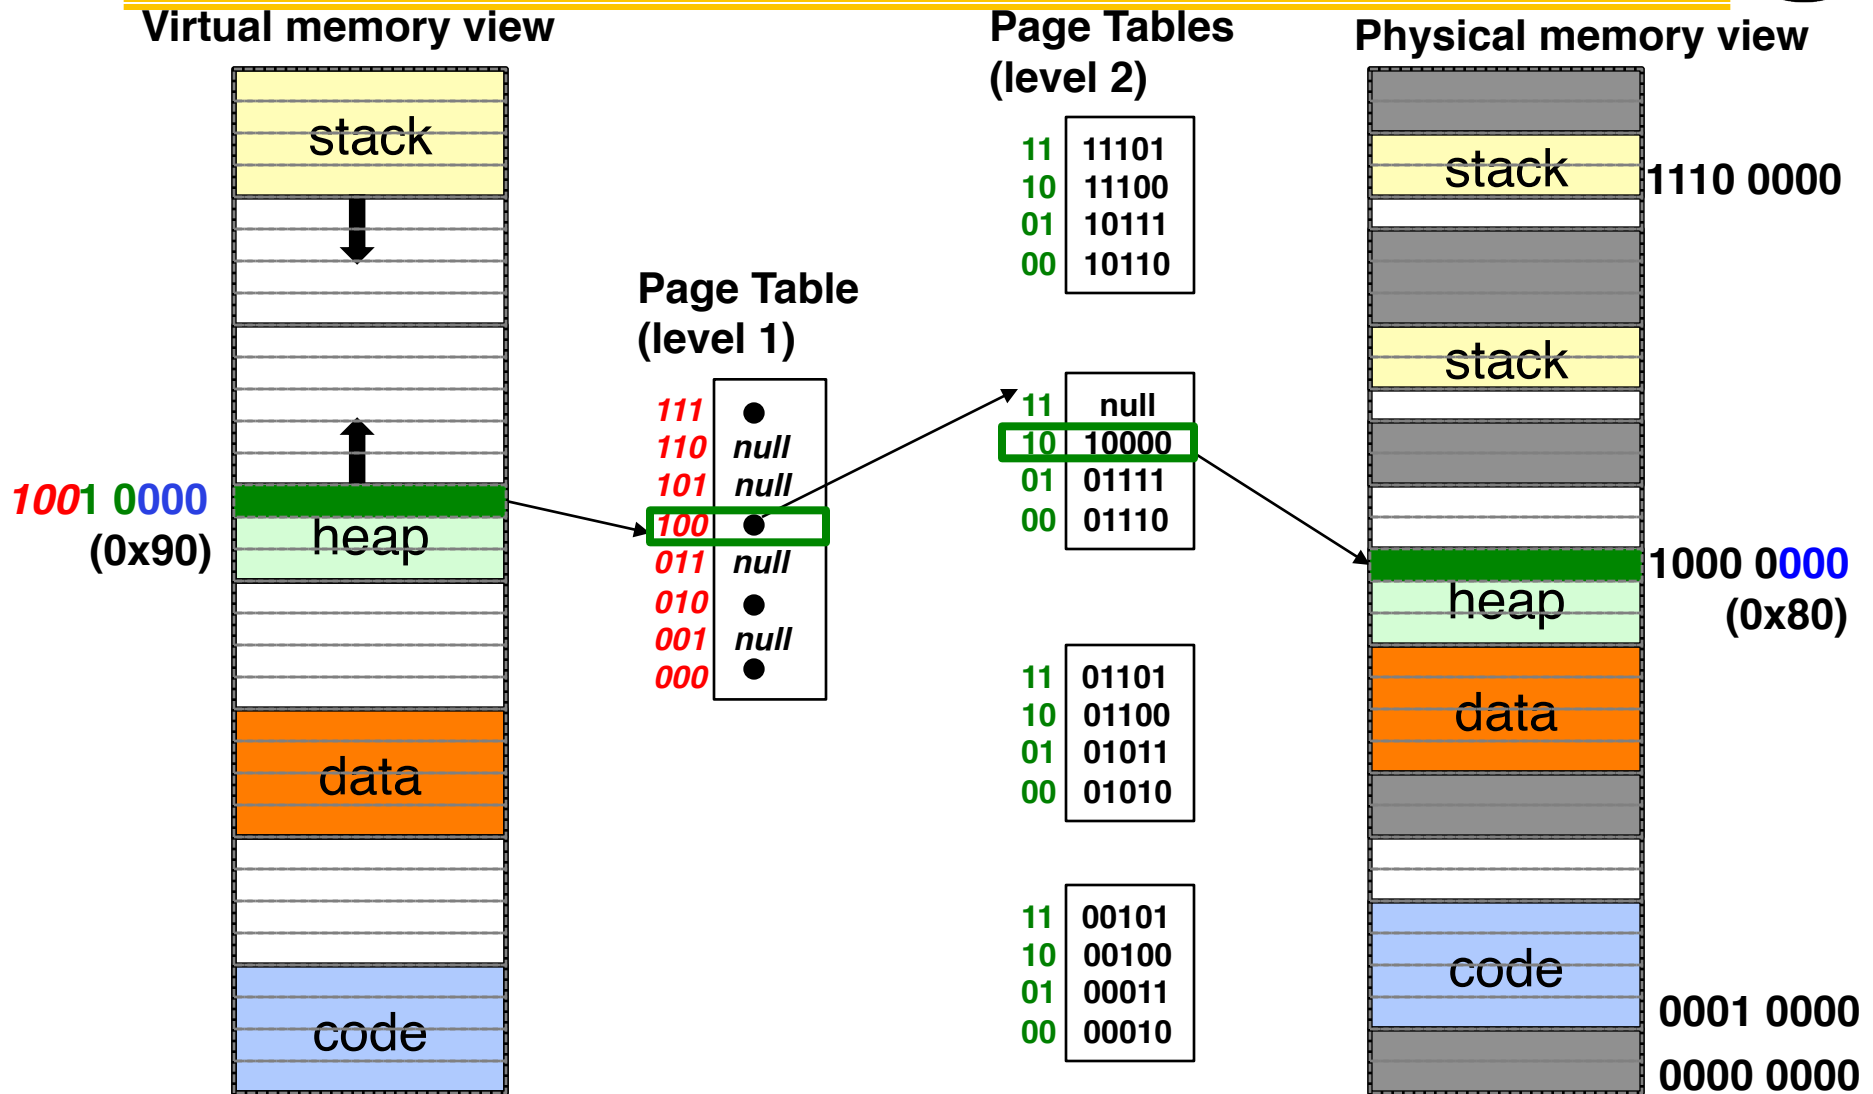


- Pros:
 - Only need to allocate as many page table entries as we need for application – size is proportional to usage
 - » In other words, sparse address spaces are easy
 - Easy memory allocation
 - Easy Sharing
 - » Share at segment or page level (need additional reference counting)
- Cons:
 - One pointer per page (typically 4K – 16K pages today)
 - Page tables need to be contiguous
 - » However, previous example keeps tables to exactly one page in size
 - Two (or more, if >2 levels) lookups per reference
 - » Seems very expensive!

Summary: Two-Level Paging




Summary: Two-Level Paging






Inverted Page Table

- With all previous examples (“Forward Page Tables”)
 - Size of page tables is at least as large as amount of virtual memory allocated to *ALL* processes
 - Physical memory may be much, much less
 - › Much of process’ space may be out on disk or not in use



- Answer: use a hash table
 - Called an “Inverted Page Table”
 - Size is independent of virtual address space
 - Directly related to amount of phy mem (1 entry per phy page)
 - Very attractive option for 64-bit address spaces (IA64, PowerPC, UltraSPARC)
- Cons: Complexity of managing hash chains in hardware

Summary: Inverted Table



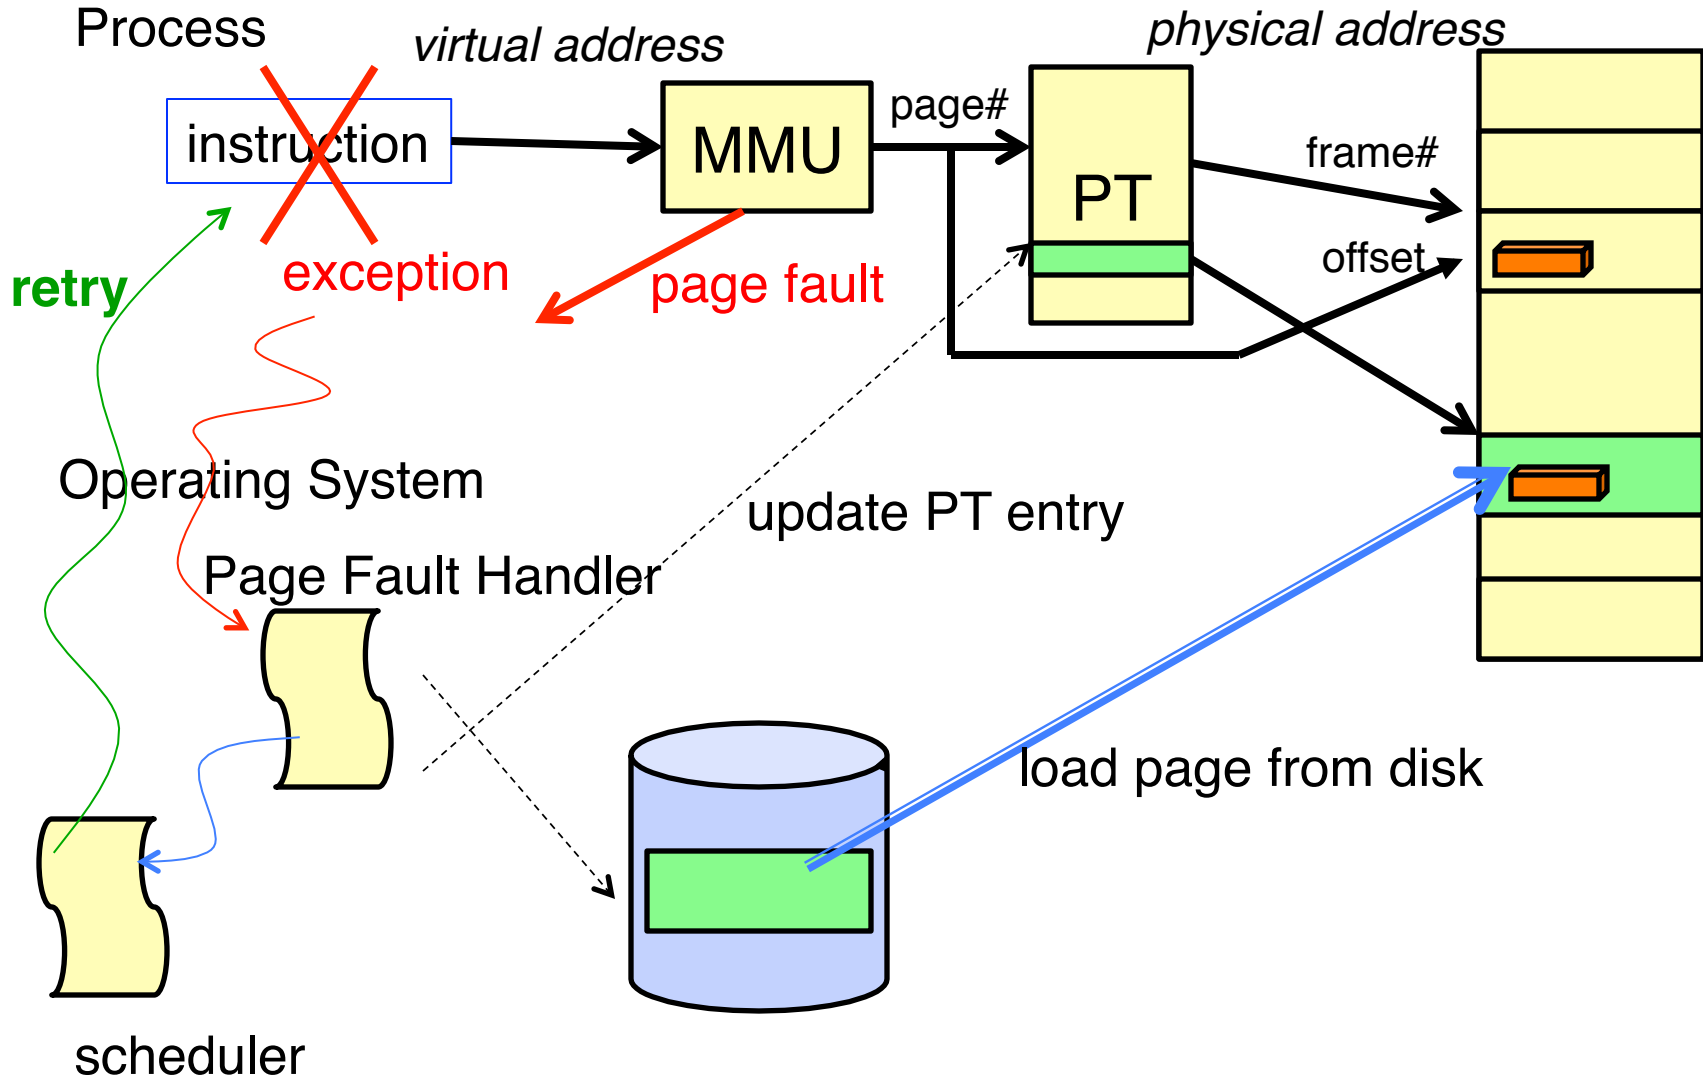
Address Translation Comparison




| | Advantages | Disadvantages |
|----------------------------|---|---|
| Segmentation | Fast context switching: Segment mapping maintained by CPU | External fragmentation |
| Paging (single-level page) | No external fragmentation, fast easy allocation | Large table size ~ virtual memory Internal fragmentation |
| Paged segmentation | Table size ~ # of pages in virtual memory , fast easy allocation | Multiple memory references per page access |
| Two-level pages | | |
| Inverted Table | Table size ~ # of pages in physical memory | Hash function more complex |



What happens when ...



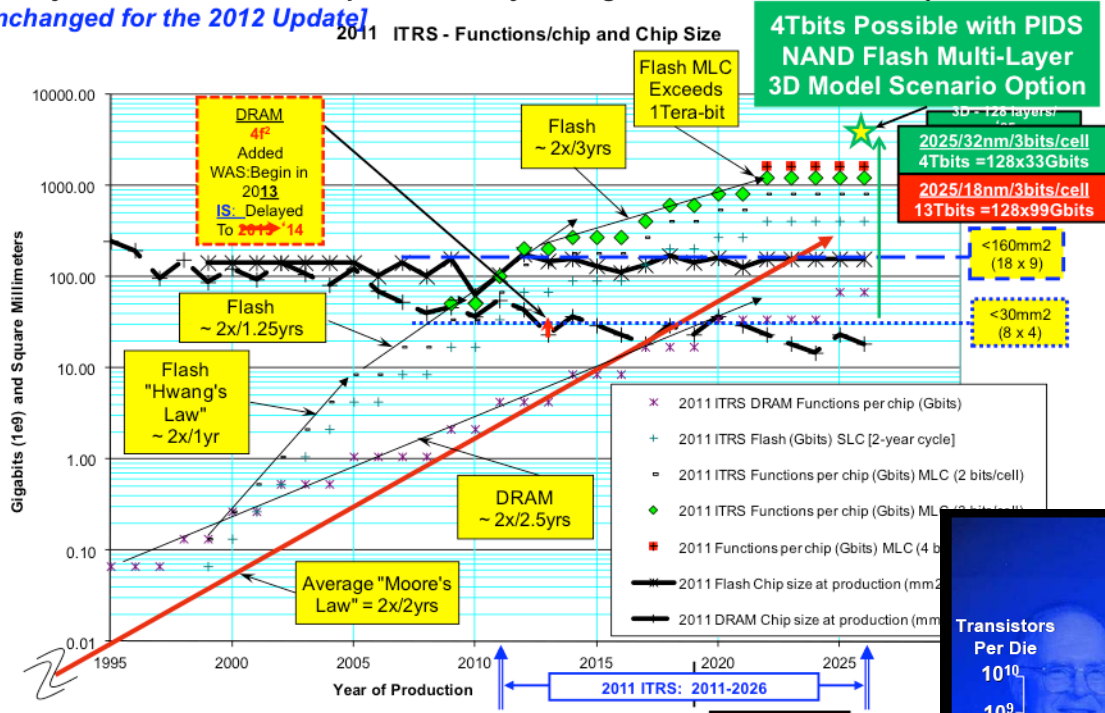
How has OS design choices been influenced by technological change?




RAM?




Figure 7 2011 ITRS Product Technology Trends: Memory Product Functions/Chip and Industry Average "Moore's Law" and Chip Size Trends [unchanged for the 2012 Update]



Source: 2011 ITRS - Executive Summary Fig 7





Summary



- Memory is a resource that must be multiplexed
 - Controlled Overlap: only shared when appropriate
 - Translation: Change virtual addresses into physical addresses
 - Protection: Prevent unauthorized sharing of resources
- Simple Protection through segmentation
 - Base + Limit registers restrict memory accessible to user
 - Can be used to translate as well
- Page Tables
 - Memory divided into fixed-sized chunks of memory
 - Offset of virtual address same as physical address
- Multi-Level Tables
 - Virtual address mapped to series of tables
 - Permit sparse population of address space
- Inverted page table: size of page table related to physical memory size