

WEEK 4


Memory management and virtual memory

Stallings, Chapters 7 & 8.1


1

Fixed partitions

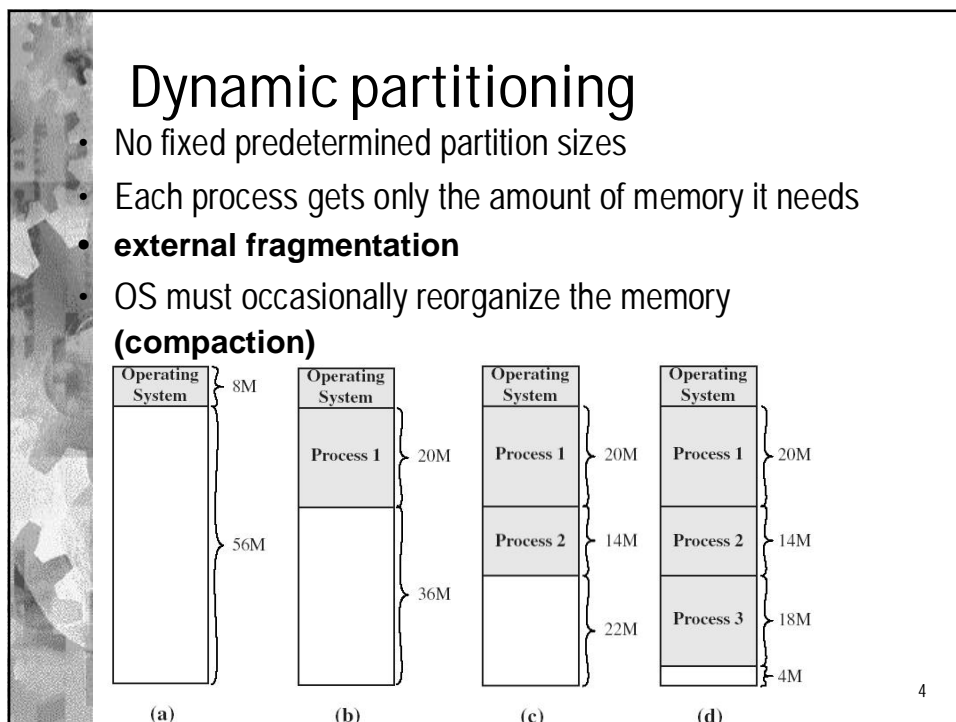
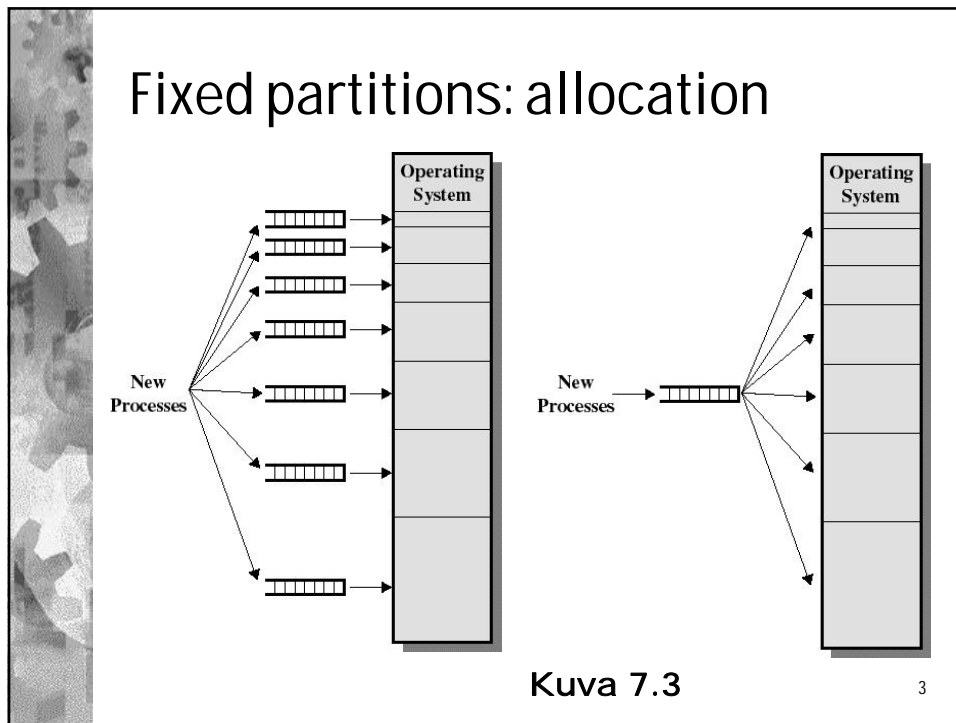
- OS splits the memory to partitions
- Whole process must fit to a partition
- Program sizes must be at most the size of the partitions
- Running larger programs required special effort from the programmer
- SWAPPING out to make place for another process
- internal fragmentation



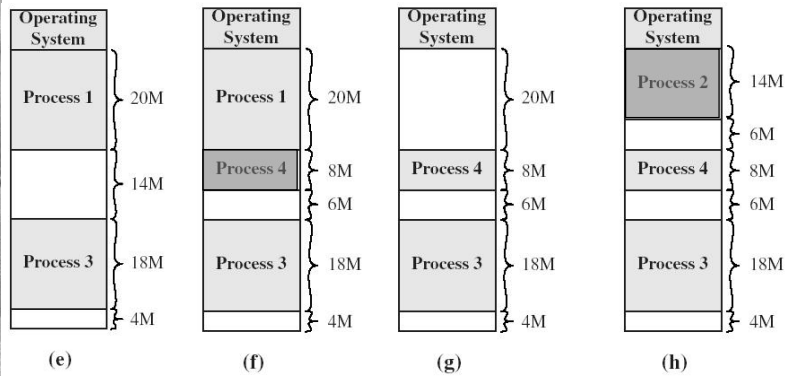
Equal-size partitions



Unequal-size partitions



Dynamic partitions




- Needed space for process 4 -> swap out process 2
- External fragments: 6M + 6M + 4M = 14M

5

Allocation Kuva 7.5

Where to place the new process?

- Goal: avoid external fragmentation and compation
- Some alternatives:
- Best-fit
- First-fit
- Next-fit




Example Memory Configuration Before and After Allocation of 16 Mbyte Block

Buddy System

- Compromise: Fixed partition sizes, dynamic splitting
- Partition sizes are 2:n potenssissa
 - Some fixed smallest allocation size
- Allocation and combining efficient
- List for each separate size
- Allocation
 - If correct size not available, split a larger one
- Release
 - When two adjacent same sized areas free combine them

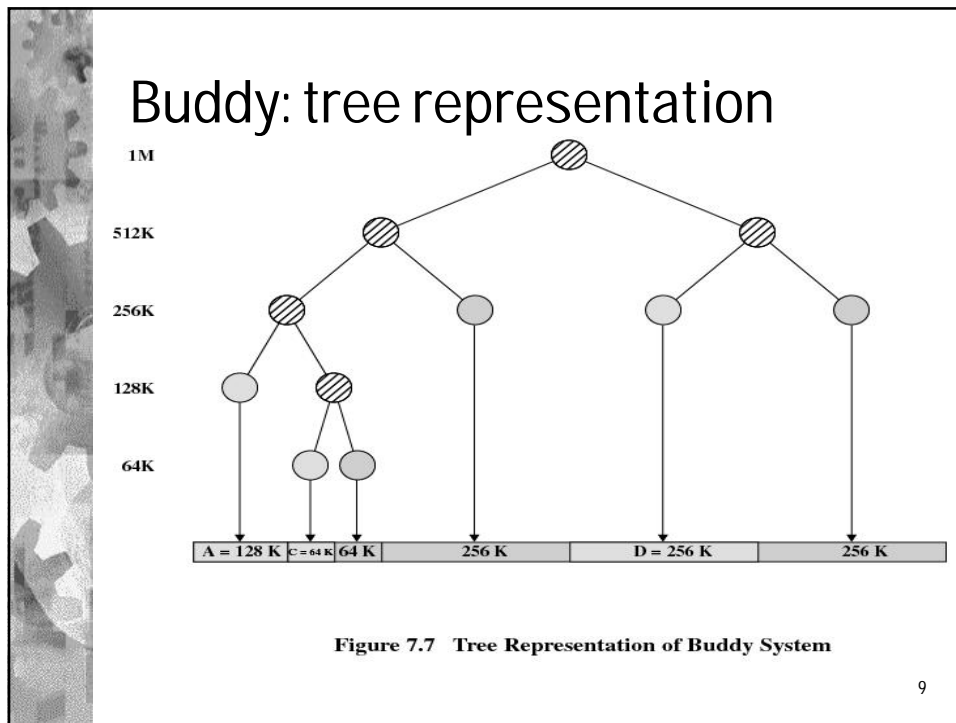
7

Buddy System

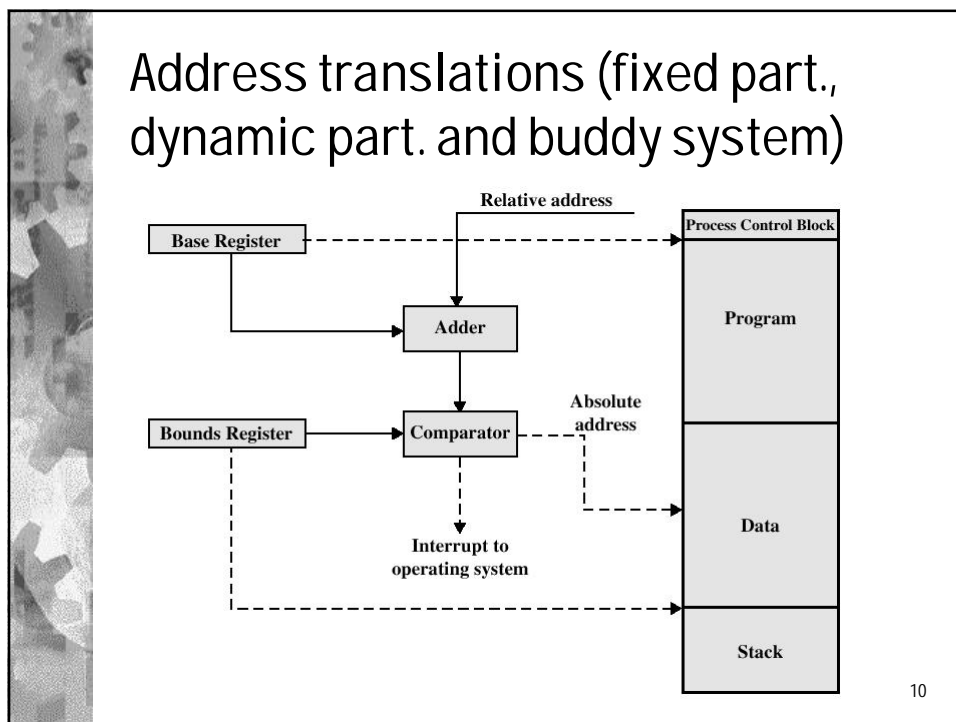


Kuva 7.6

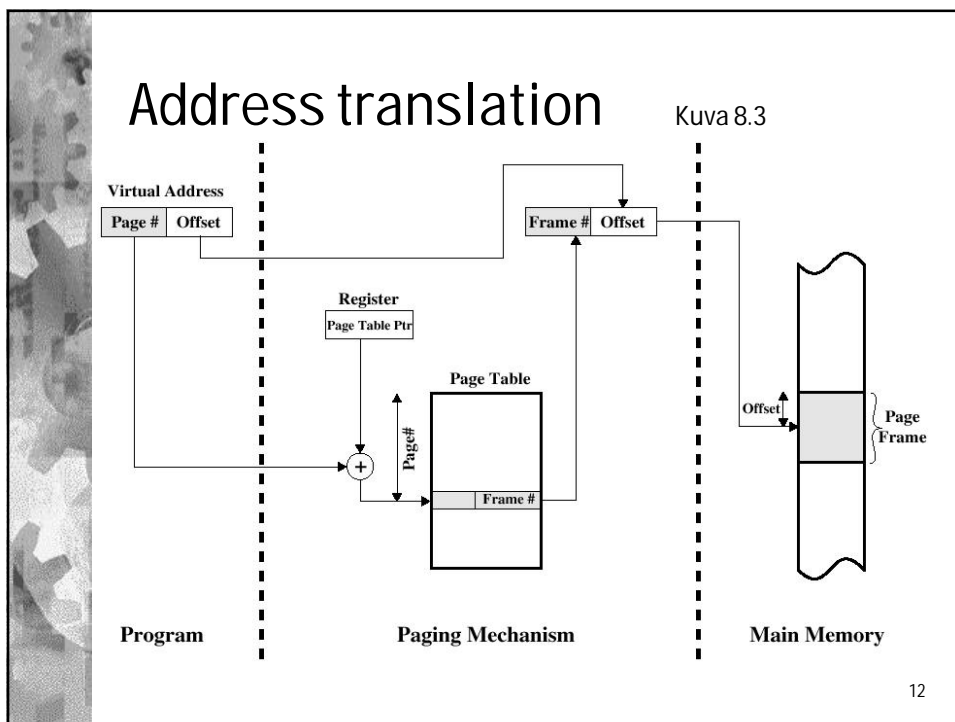

8



9

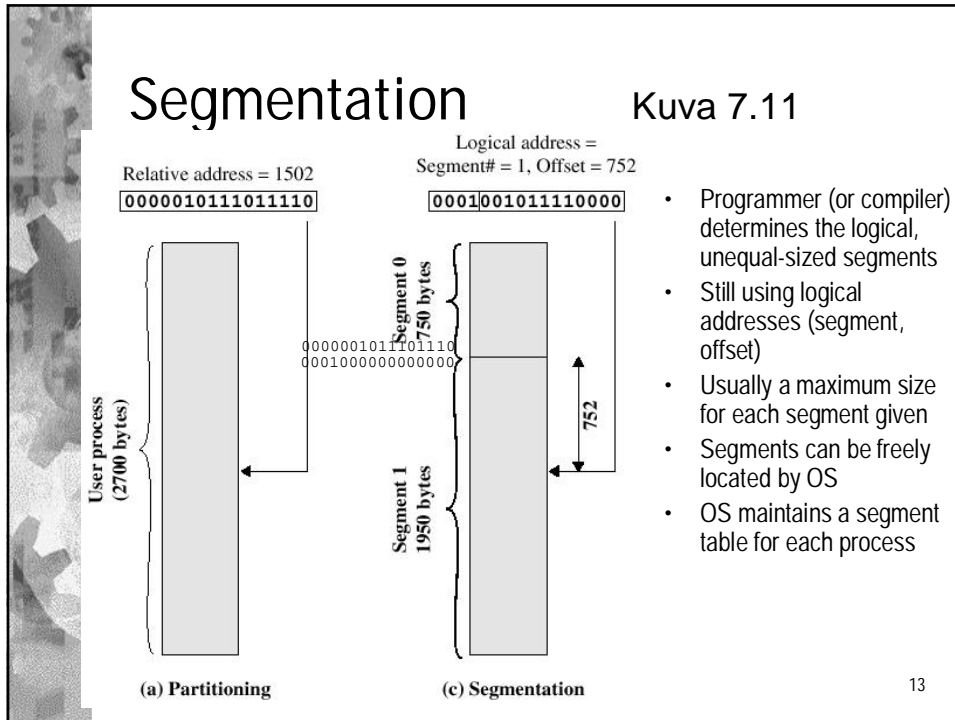


10



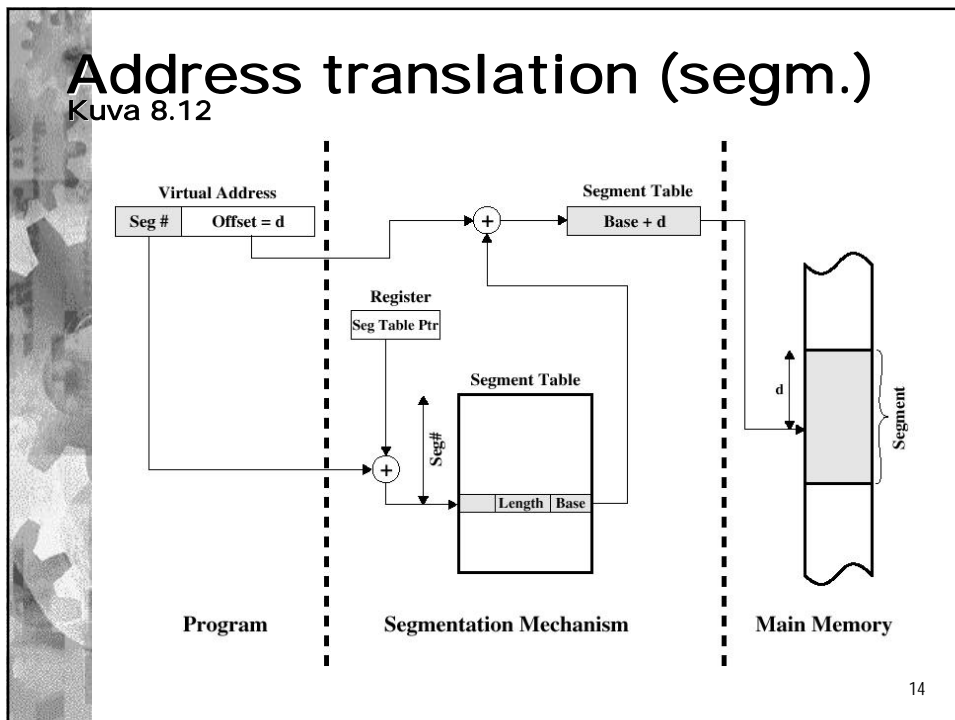
Segmentation

Kuva 7.11



Address translation (segm.)

Kuva 8.12



WEEK 4

Virtual memory


Stallings, Chapter 8

15

Page table

Kuva 8.2a

Virtual Address



P= present bit
M = Modified bit

Page Table Entry



- Each process has its own page table
- Each entry has a present bit, since not all pages need to be in the memory all the time -> page faults
- Remember the locality principle
- Logical address space can be much larger than the physical

16

Structure of page table

- Large virtual address space
 - Logical address could be 32- or 64-bits
- Each process has a large page table
 - Using 32-bit address and frame size 4KB (12 bit offset), means $2^{20} = 1\text{M}$ of page tables entries for a single process
 - Each entry requires several bytes (lets say 4 bytes), so the final size of page table could be for example 4 MB
- Thus the page table is divided to several pages also and part of it can be on the disk
 - We only need the part of the page table in the memory that covers the pages used currently in the execution of the process

17

Two-level hierarchical page table

- Top most level in one page and always in the memory


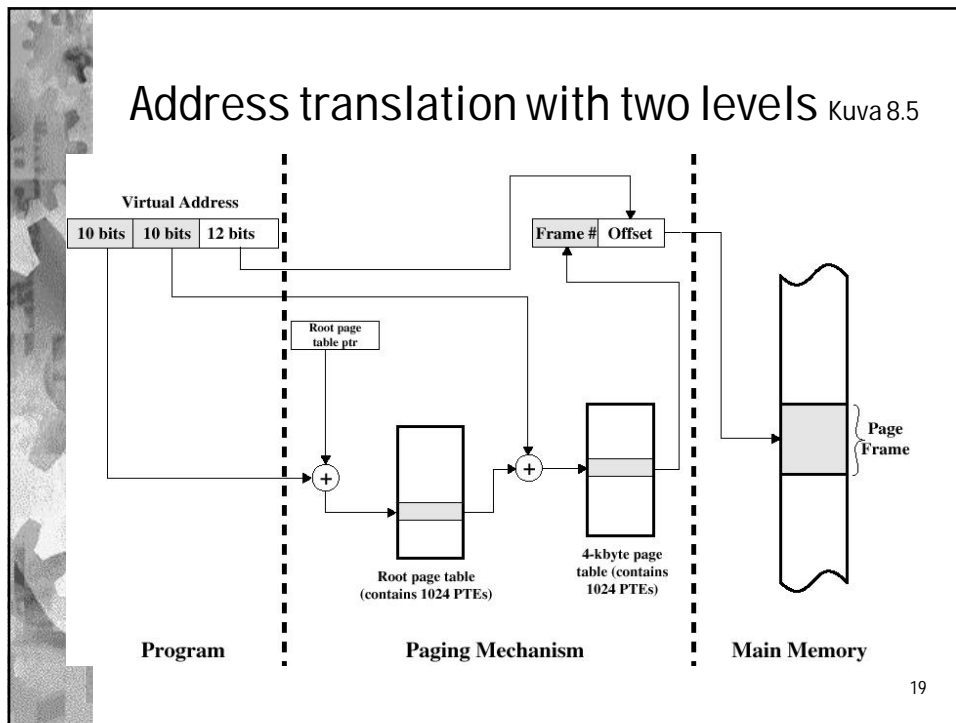


Figure 8.4 A Two-Level Hierarchical Page Table

8



Inverted page table

20

Inverted page table

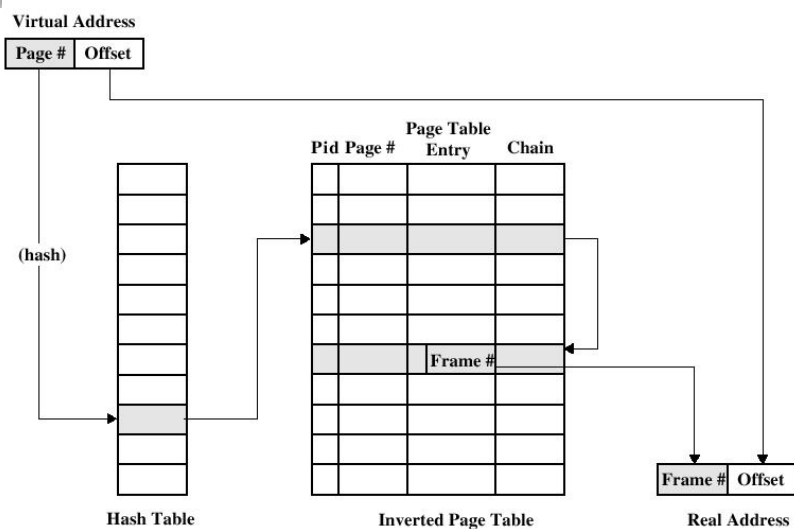
- Physical memory often smaller than the virtual address space of processes
- Invert booking: Store for each page frame what page (of which process) is stored there
 - Only one global table (inverted page table), one entry for each page frame.
- Search for the page based on the content of the table
 - Inefficient, if done sequentially,
 - Use hash to calculate the location, start search from there
- If page not found, page fault
- Useful, only if TLB is large

21

Inverted page table

Kuva 8.6

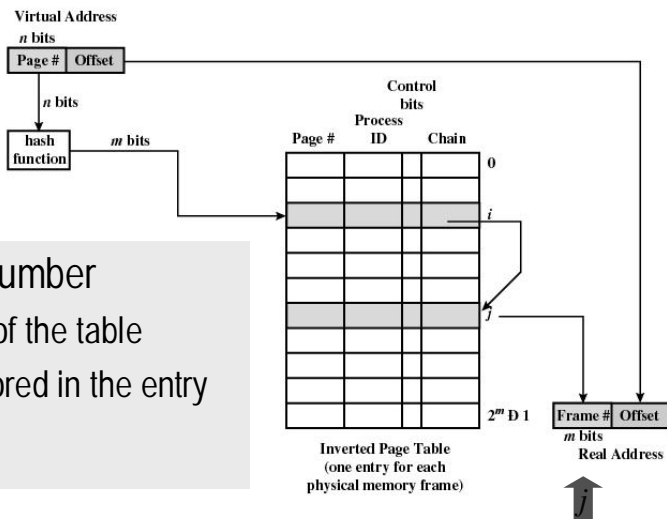
4. ed



22

5. ed

Inverted page table v.2



- Frame number
- Index of the table
- Not stored in the entry

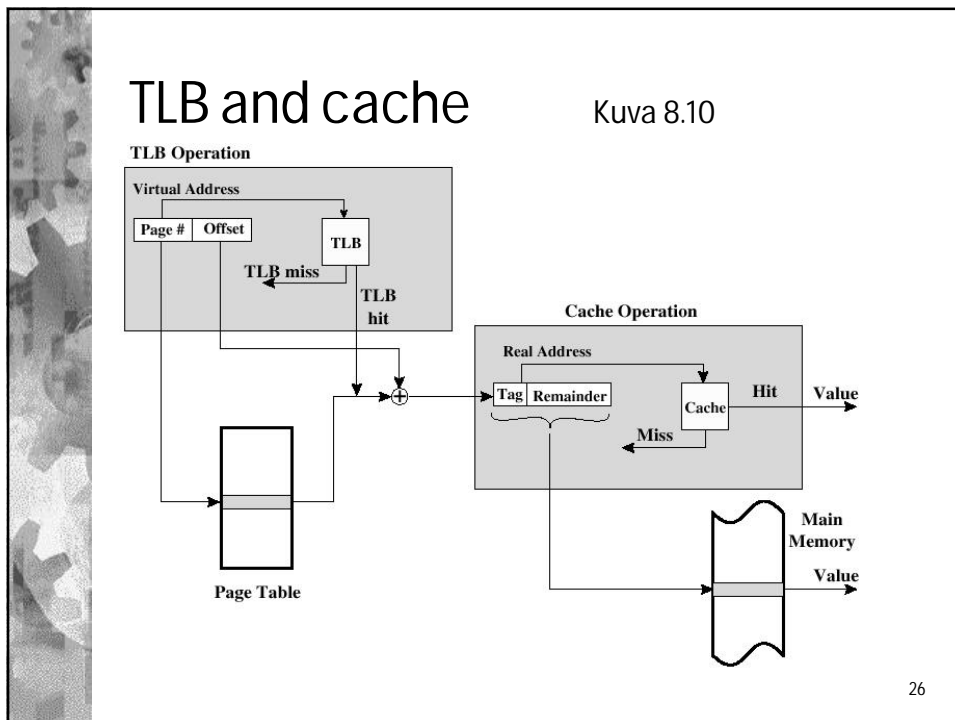
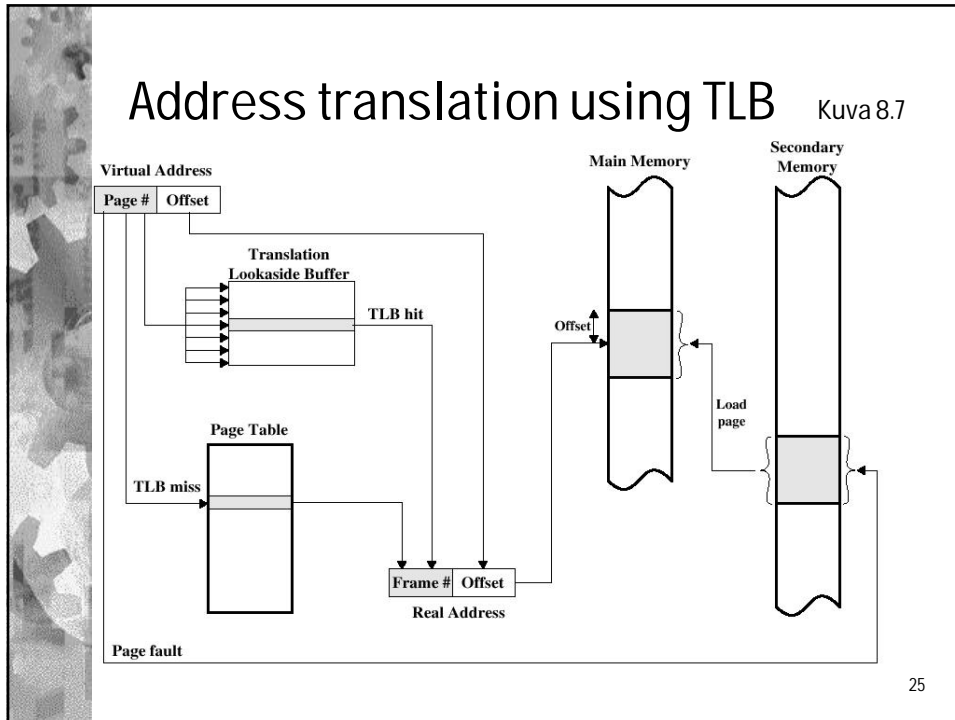
Figure 8.6 Inverted Page Table Structure

23

TLB – translation lookaside buffer

- Part of memory management unit (MMU)
 - Cache for used page table entries to avoid extra memory access during address translation
- Associative search
 - Compare with all elements at the same time (fast)
- Each TLB element contains: page number, page table entry, validity bit
- Each process uses the same page numbers $0, 1, 2, \dots$, stored on different page frames
 - TLB must be cleared during process switch
 - At least clear the validity bits (this is fast)

24







Segmentointi ja sivutus yhdistettynä

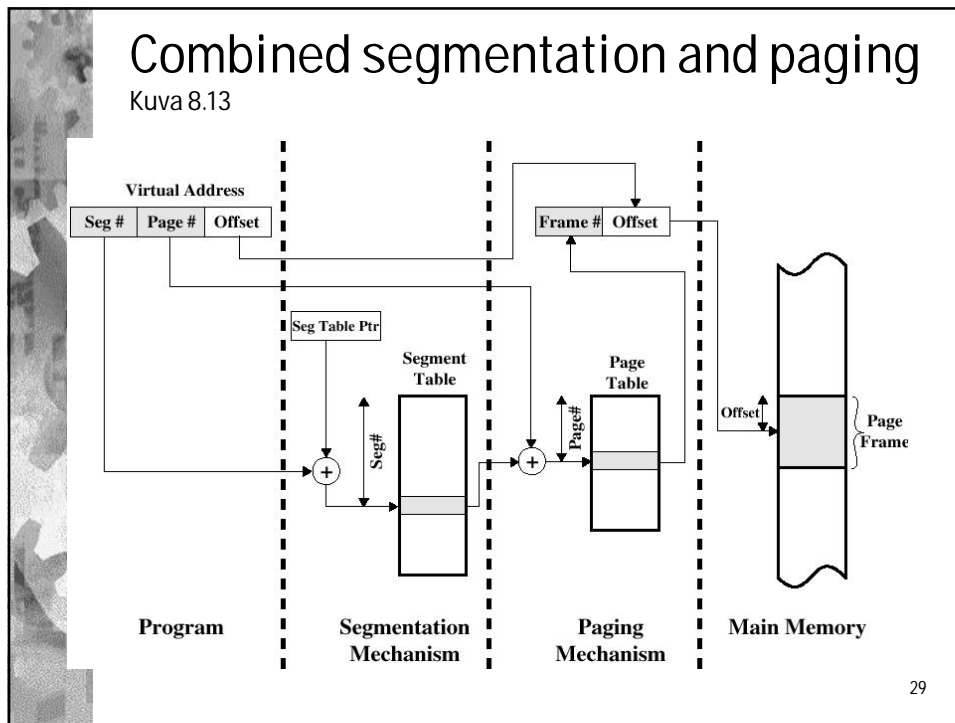
27

Combining segmentation and paging Tan01 4-39

PCB



28



Protection and sharing

No slides in English (sorry)

30