

HELSINGIN YLIOPISTO
HELSINGFORS UNIVERSITET
UNIVERSITY OF HELSINKI

Lecture 4

Internal Memory Cache

Stallings: Ch 4, Ch 5

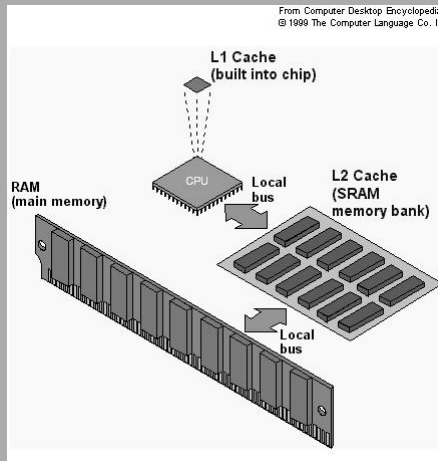
Key Characteristics


Locality

Cache

Main Memory

From Computer Desktop Encyclopedia
© 1999 The Computer Language Co., Inc






Key Characteristics of Memories / Storage

<p>Location</p> <ul style="list-style-type: none"> Processor Internal (main) External (secondary) <p>Capacity</p> <ul style="list-style-type: none"> Word size Number of words <p>Unit of Transfer</p> <ul style="list-style-type: none"> Word Block <p>Access Method</p> <ul style="list-style-type: none"> Sequential Direct Random Associative 	<p>Performance</p> <ul style="list-style-type: none"> Access time Cycle time Transfer rate <p>Physical Type</p> <ul style="list-style-type: none"> Semiconductor Magnetic Optical Magneto-Optical <p>Physical Characteristics</p> <ul style="list-style-type: none"> Volatile/nonvolatile Erasable/nonerasable <p>Organization</p>
----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

(Sta06 Table 4.1)



Goals

- I want my memory lightning fast
- I want my memory to be gigantic in size

- Register access viewpoint
 - data access as fast as HW register
 - **data size as large as memory**


- Memory access viewpoint
 - data access as fast as memory
 - **data size as large as disk**

cache

HW solution

virtual memory

HW help for SW solution

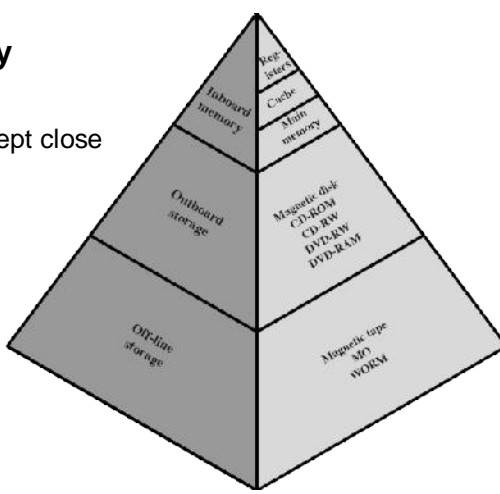


Memory Hierarchy

- Most often needed data kept close
- Access to small data sets can be made fast
 - simpler circuits
 - smaller gate delays
- Faster ~ more expensive
- Large can be bigger and cheaper (per B)

up: smaller, faster, more expensive, more frequent access

down: bigger, slower, less expensive, less frequent access



The pyramid is divided into four horizontal layers. From top to bottom: 1. Register, Cache, Main memory; 2. Inboard memory; 3. Outboard storage; 4. Off-line storage. To the right of the pyramid, storage technologies are listed: Magnetic disk, CD-ROM, CD-RW, DVD-RW, DVD-RAM, and Magnetic tape, MFM, WORM.

(Sta06 Fig 4.1)

Principle of locality (*paikallisuus*)

- In any given time period, memory references occur only to a small subset of the whole address space
= The reason why memory hierarchies work

Prob (small data set) = 99%	“Cost” (small data set) = 2 μs
Prob (the rest) = 1%	“Cost” (the rest) = 20 μs

Aver cost = 99% * 2 μs + 1% * 20 μs = 2.2 μs

- Average cost is close to the cost of small data set
- How to determine data for that small set?
- How to keep track of it?

Sta06 Fig 4.2

Principle of locality

- In any given time period
 - memory references occur only to a small subset of the whole address space
- Temporal locality (*ajallinen*)
 - it is likely that a data item referenced a short time ago will be referenced again soon
- Spatial locality (*alueellinen*)
 - it is likely that a data items close to the one referenced a short time ago will be referenced soon

MEM: 345 23 71 8 305 63 91 2

Computer Organization II

Cache

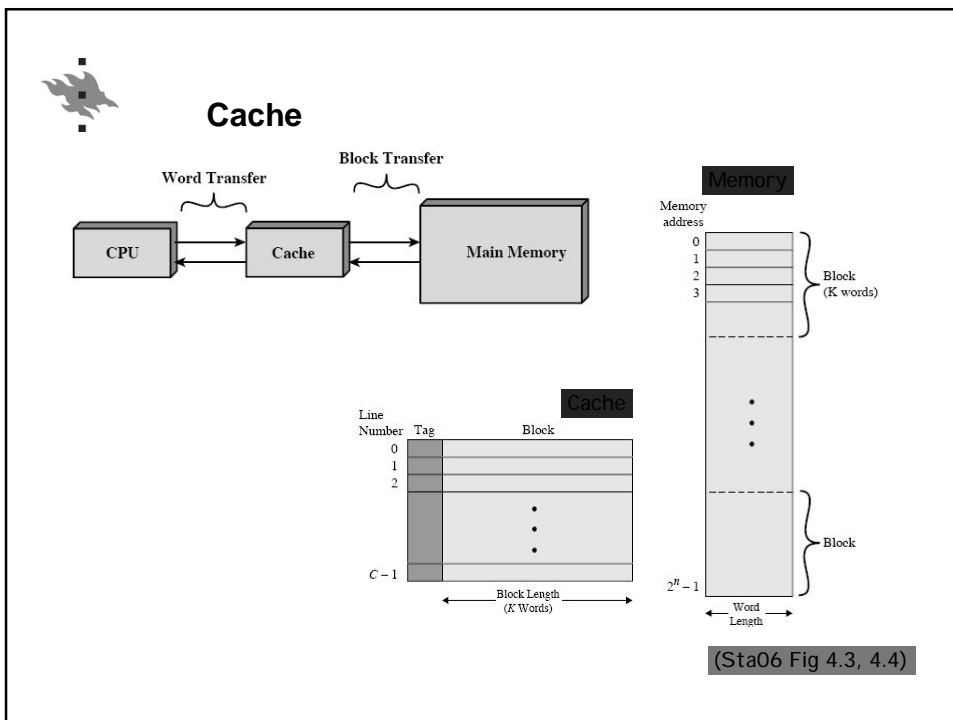
Teemu's Cheesecake

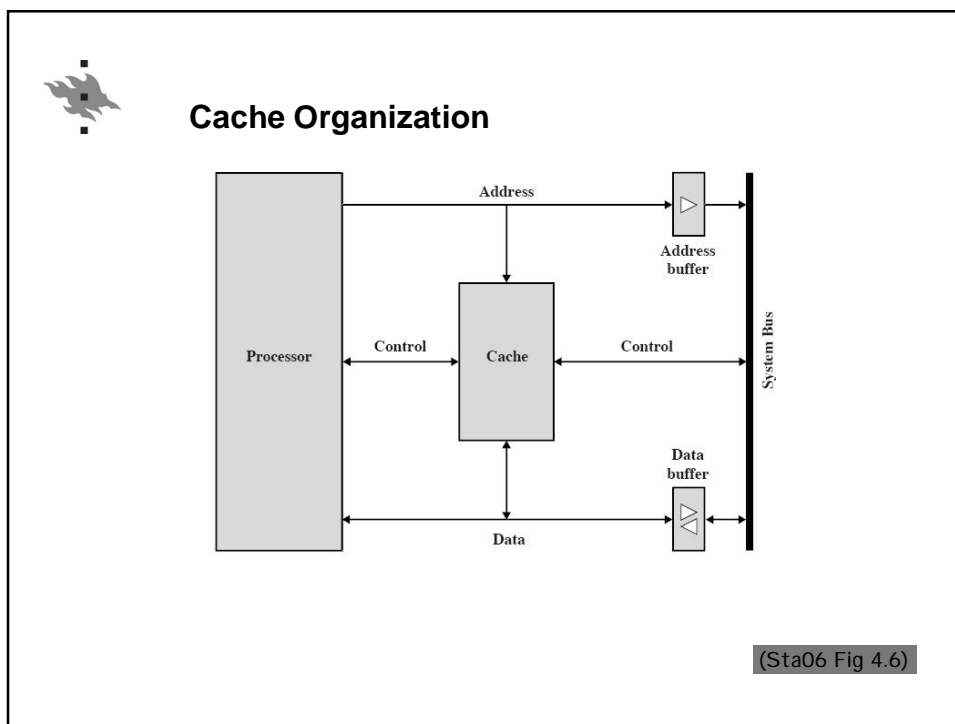
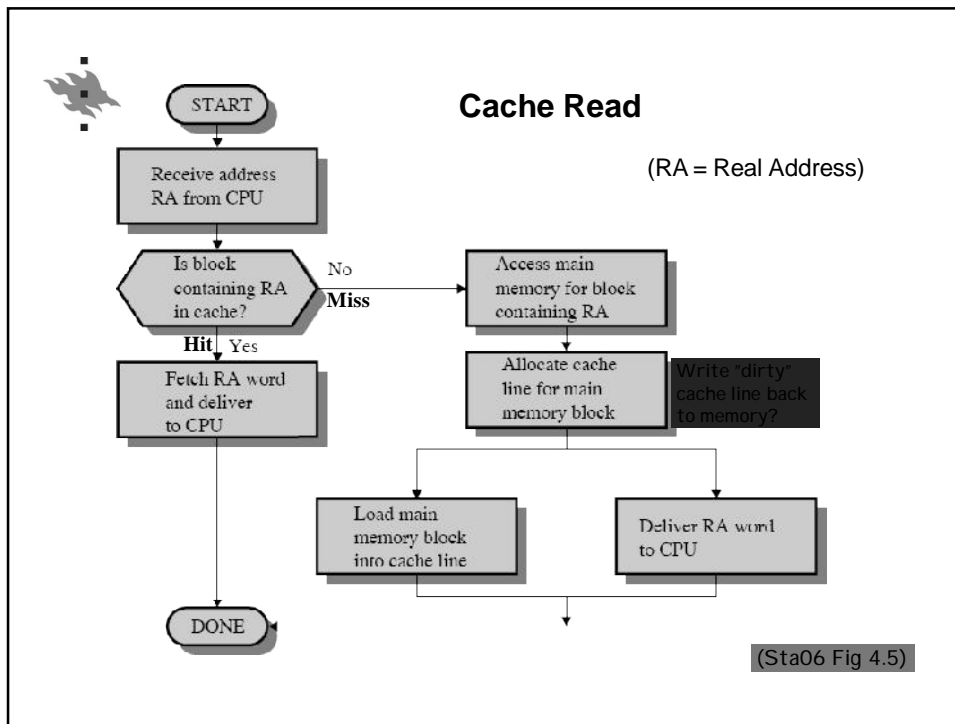
Register, on-chip cache, memory, disk, and tape speeds relative to times locating cheese for the cheese cake you are baking...


Storage Technology	Relative Time to Locate Cheese
Hand (register)	0.5 sec
Table (cache)	1 sec
Refrigerator (memory)	10 sec
Moon (disk)	12 days
Europa (Jupiter) (tape)	4 years

Cache Memory (*välimuisti*)

- How to access main memory as fast as registers?
- Locality → Use (CPU) cache!
 - Keep most probably referenced data in fast cache close to processor, and rest in memory
 - Most of data accesses only to cache
 - hit ratio 0.9-0.99
 - Cache is much smaller than main memory
 - Cache is (much) more expensive (per byte) than memory








Cache Design

Cache Size	Write Policy
Mapping Function	Write through
Direct	Write back
Associative	Write once
Set Associative	Line Size
Replacement Algorithm	Number of caches
Least recently used (LRU)	Single or two level
First in first out (FIFO)	Unified or split
Least frequently used (LFU)	
Random	

- **Cache Size & Line Size**
 - Many blocks help for temporal locality
 - Large blocks help for spatial locality
 - Larger cache is slower
 - Multi-level cache

Typical sizes:
 L1: 8 KB – 64 KB
 L2: 256KB - 8 MB

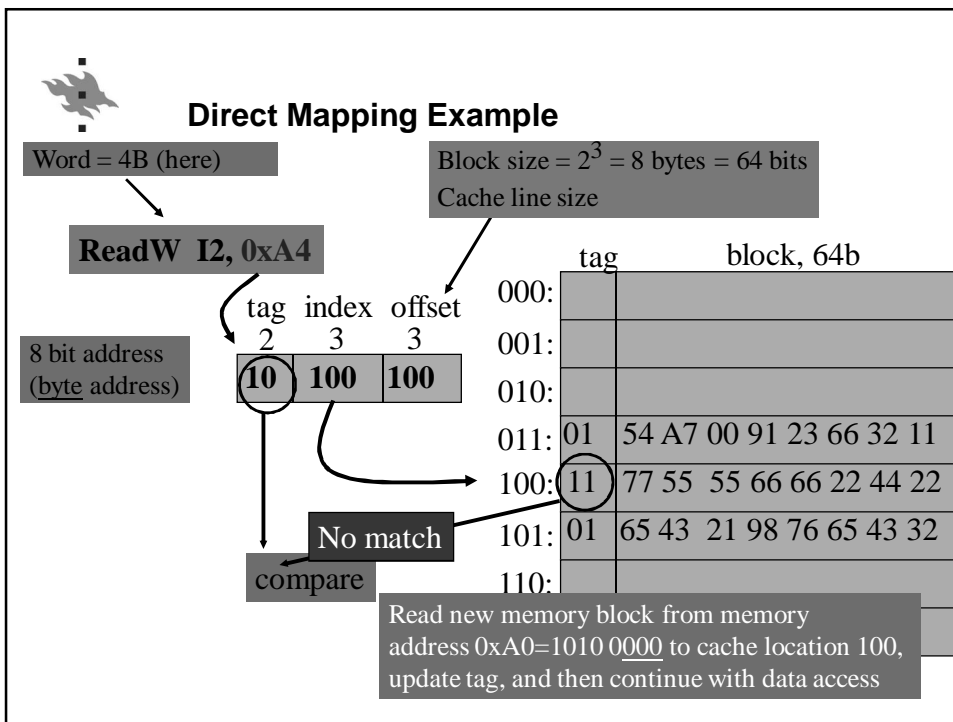
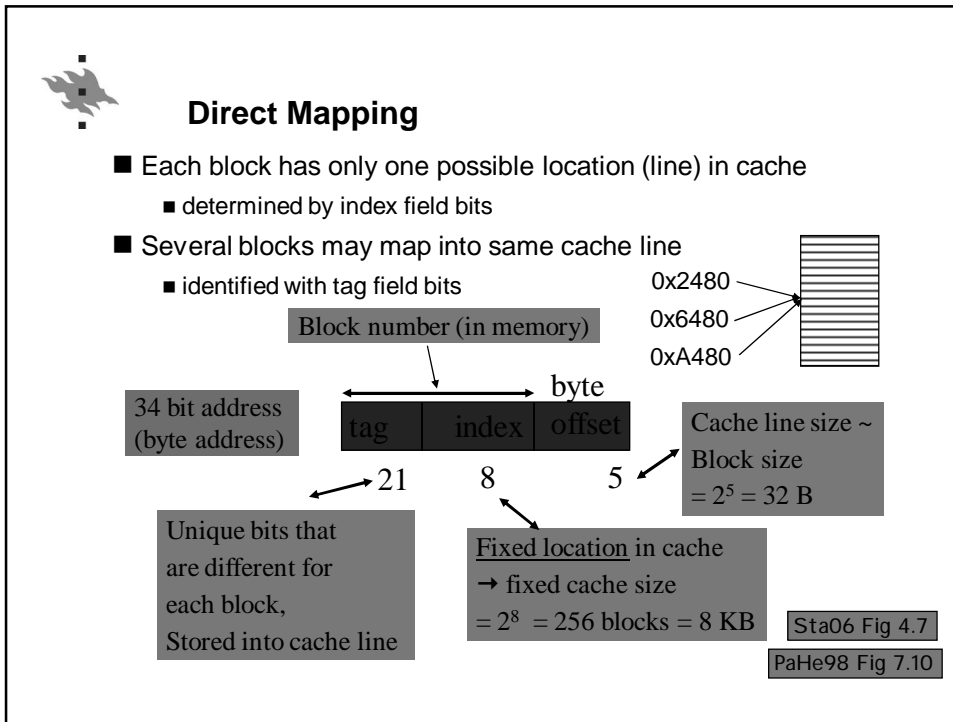
(Sta06 Table 4.2)

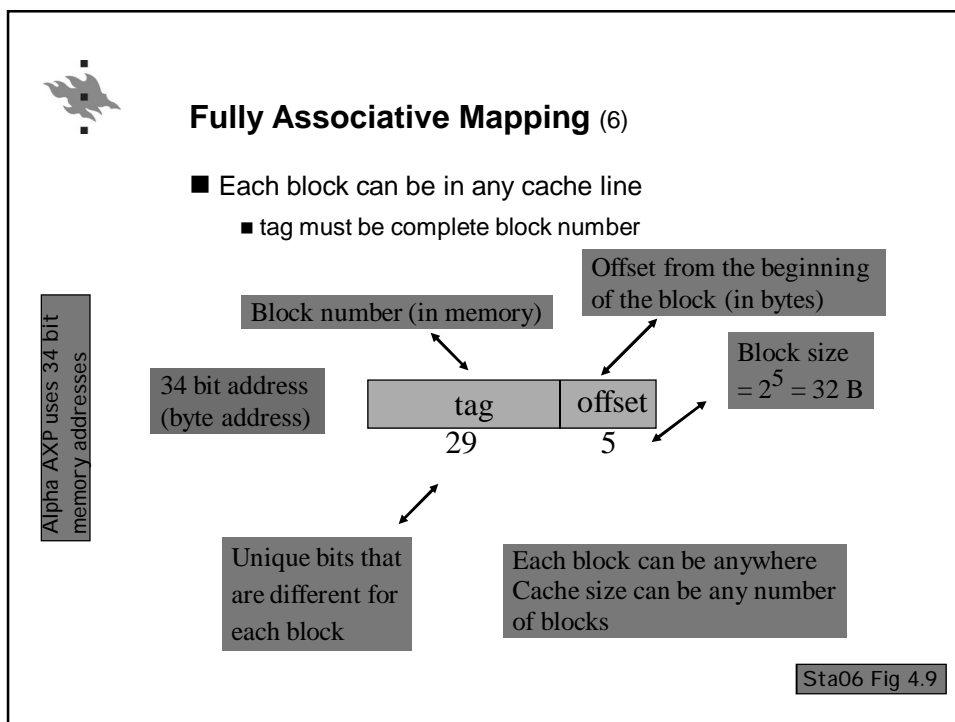
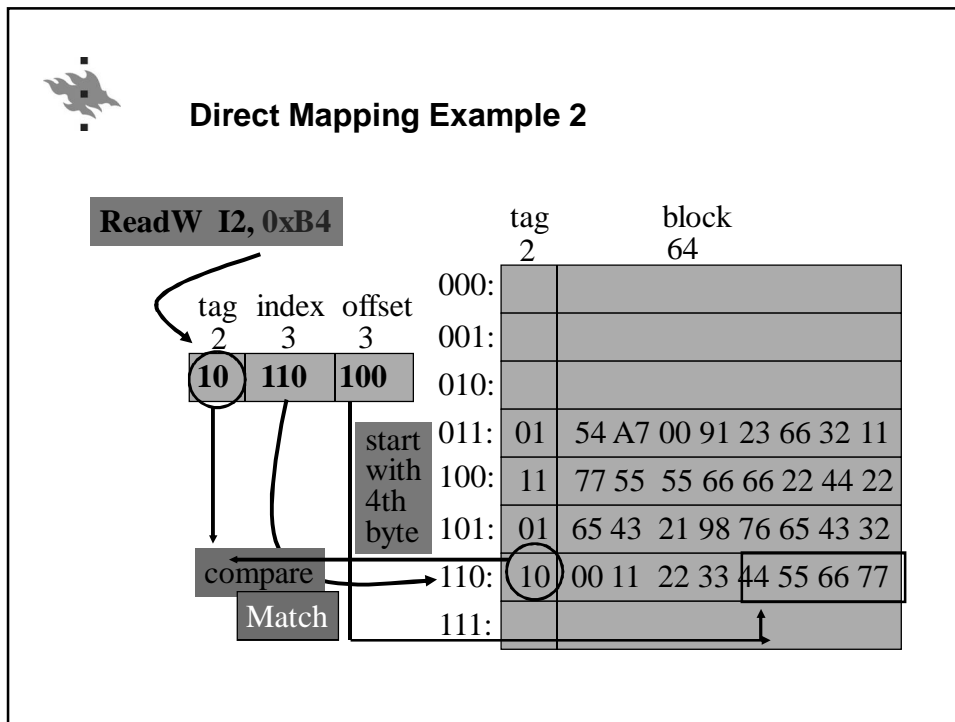


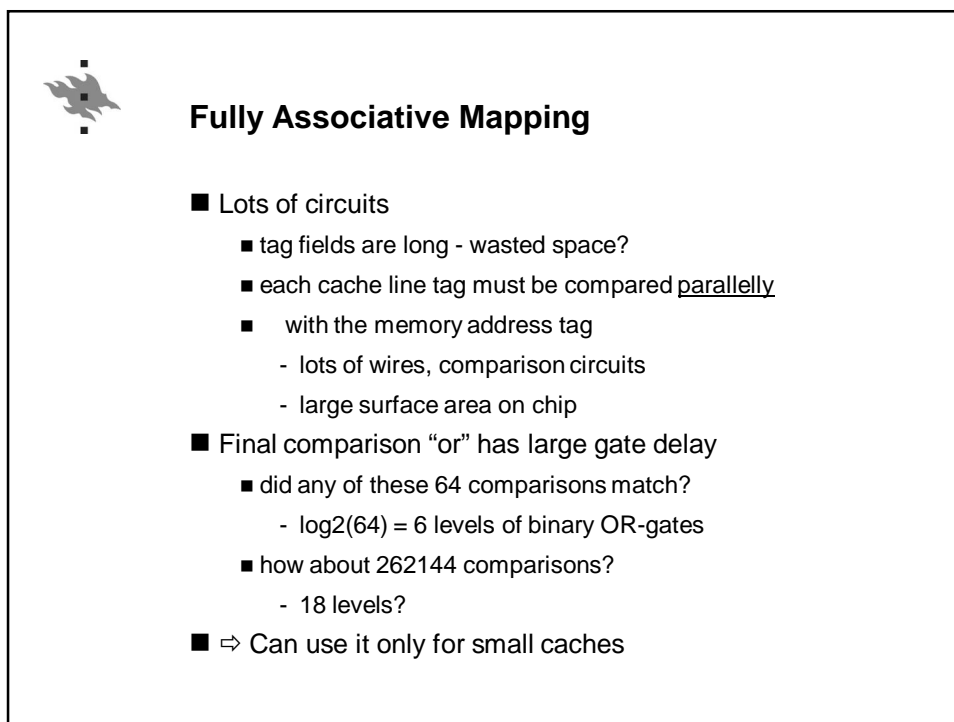
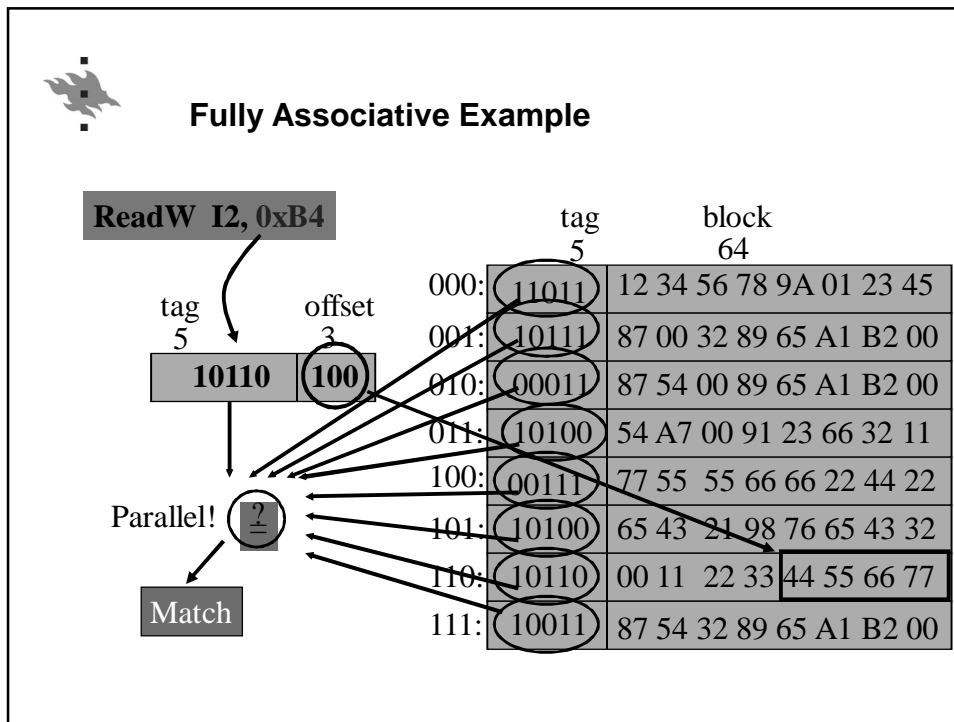
Mapping

- Which block contains the memory location?
- Is the block in cache?
- Where is it located?
- **Solutions**
 - direct mapping (suora kuvaus)
 - fully associative mapping (täysin assosiatiivinen)
 - set associative mapping (joukkoassosiatiivinen)

Cache simulation tools:
<http://www.ecs.umass.edu/ece/koren/architecture/Cache/frame0.htm>







Set Associative Mapping

- With set size $k=2$, each cache entry contain 2 blocks
 - Use set (set index) field to find the cache entry
 - Use tag to determine if the block belongs to the set
 - Use offset to find the proper byte in the block

34 bit address
(byte address)

tag

set

offset

Block size
 $= 2^5 = 32 \text{ B}$

Unique bits that are different for each block, stored with block

Nr of sets = $v = 2^7 = 128 \text{ blocks} = 4 \text{ KB}$

Total cache size = $k*v = 2*4 \text{ KB} = 8 \text{ KB}$
(without tag bits!)

Sta06 Fig 4.11

PaHe98 Fig 7.19

2-way Set Associative Cache

- $k=2 \rightarrow$ Two blocks in each set (= in one cache entry)
- 4 sets \rightarrow 2 bits for set index
- 2 words in a block = 8 Bytes \rightarrow 3 bits for byte offset
- 3 bits for tag


tag

set

offset

8 bit address
(byte address)

set	tag	block	tag	block
00:	110	12 34 56 78 9A 01 23 45	011	77 55 55 66 66 22 44 22
01:	110	87 00 32 89 65 A1 B2 00	101	65 43 21 98 76 65 43 32
10:	100	87 54 00 89 65 A1 B2 00	101	00 11 22 33 44 55 66 77
11:	101	54 A7 00 91 23 66 32 11	111	00 11 22 33 44 55 66 77
	3	64	3	64



- ### Set Associative Mapping
- Set associative cache with set size $k=2$
= 2-way cache (common)
 - Degree of associativity = nbr of blocks in a set = v
 - Large degree of associativity?
 - More data items in one set
 - Less "collisions" within set
 - Final comparison (matching tags?) gate delay?
 - Maximum (nr of cache lines) Whole cache is one set!
⇒ fully associative mapping
 - Minimum (1) Each cache line is a set!
⇒ direct mapping



Cache Replacement Algorithm

- Which cache block to replace
to make room for new block from memory?
- Direct mapping: trivial
- First-In-First-Out (FIFO)?
- Least-Frequently-Used (LFU)?
- Random?
- Which one is best / possible?
 - Chip area?
 - Fast? Easy to implement?



Cache Write Policy – memory writes?

- Write through (*läpikirjoittava*)
 - Each write goes always to cache and memory
 - Each write is a cache miss!
- Write back (*lopuksi/takaisin kirjoittava*)
 - Each write goes only to cache
 - Write cache block back to memory
 - only when it is replaced in cache A bit set
 - Memory may have stale (old) data
 - cache coherence problem (*yhdenmukaisuus, yhtäpitävyys*)
- Write once (*"vain kerran kirjoittava?"*)
 - Write-invalidate Snoopy-cache coherence protocol for multiprocessors
 - Write invalidates data in other caches
 - Write to memory at replacement time, or when some other cache needs it (has read/write miss)


Coherence problems:

- More users of the same data: memory valid? cache valid?
- multiple processors with own caches




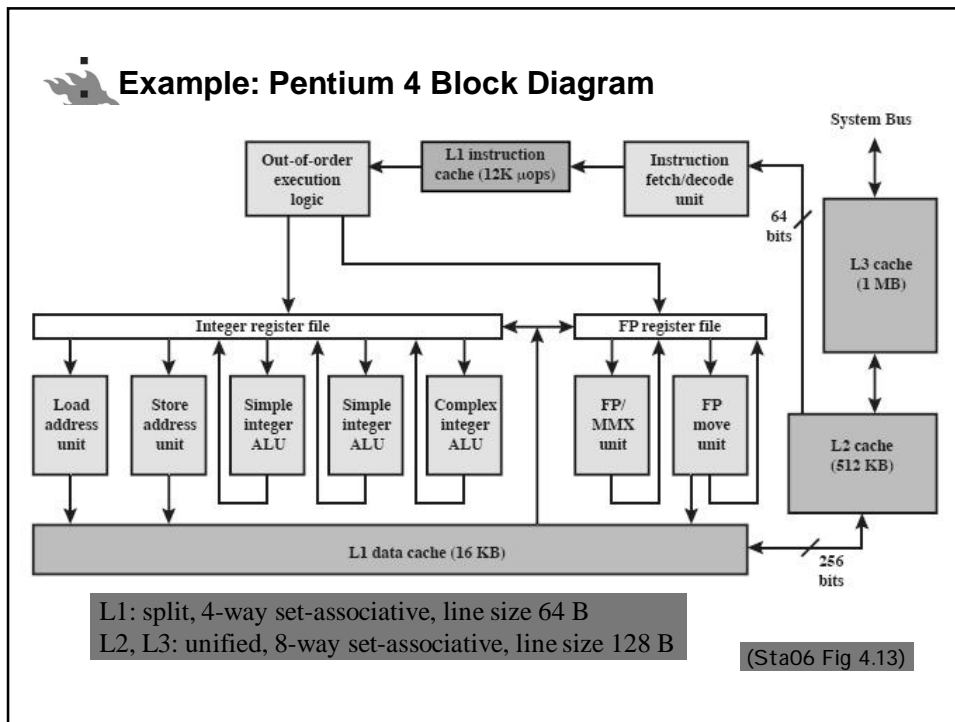
Cache Line Size


- How big cache line?
- Optimise for temporal or spatial locality?
 - bigger cache line
 - ⇒ better for spatial locality
 - more cache lines
 - ⇒ better for temporal locality
- Best size varies with program or program phase?
- Best size different with code and data?
- 2-8 words?
 - word = 1 float??



Types and Number of Caches

- Same cache for data and code, or not?
 - Data references and code references behave differently
- Unified vs. split cache (yhdistetty/erilliset)
 - split cache: can optimise structure separately for data and code
 - Trend towards split caches: Pentium, Power PC,... (instruction pipelining)
- One cache too large for best results
- Multiple levels of caches
 - L1 on same chip as CPU
 - L2 on same package or chip as CPU
 - older systems: same board
 - L3 on same board as CPU






Main Memory Types

Memory Type	Category	Erasure	Write Mechanism	Volatility
Random-access memory (RAM)	Read-write memory	Electrically, byte-level	Electrically	Volatile
Read-only memory (ROM)	Read-only memory	Not possible	Masks	Nonvolatile
Programmable ROM (PROM)			Electrically	
Erasable PROM (EPROM)	UV light, chip-level			
Electrically Erasable PROM (EEPROM)	Electrically, byte-level			
Flash memory	Electrically, block-level			

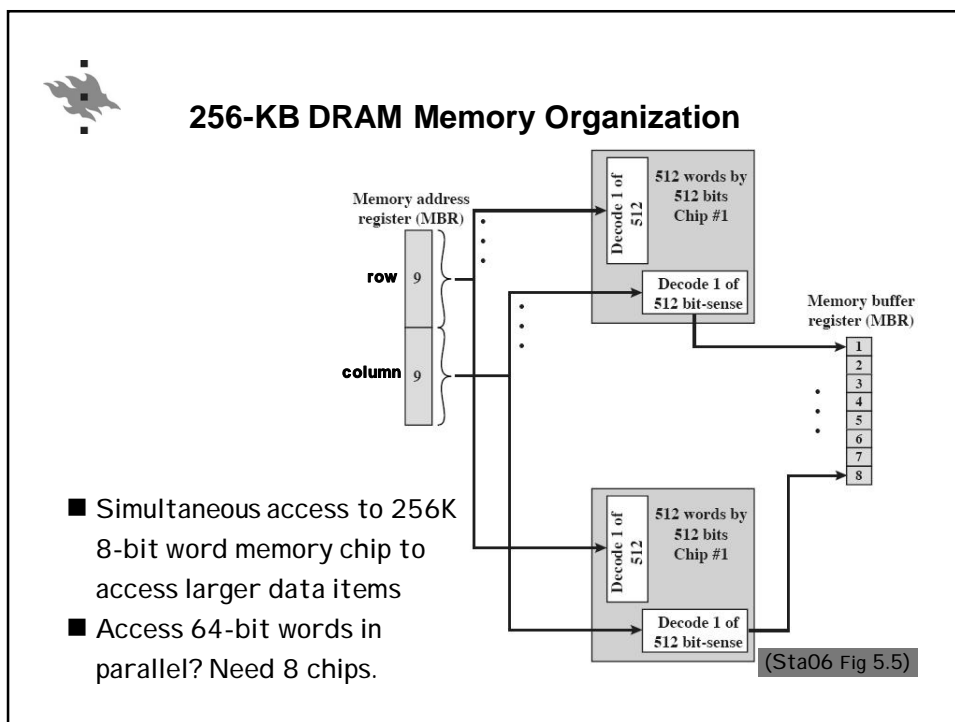
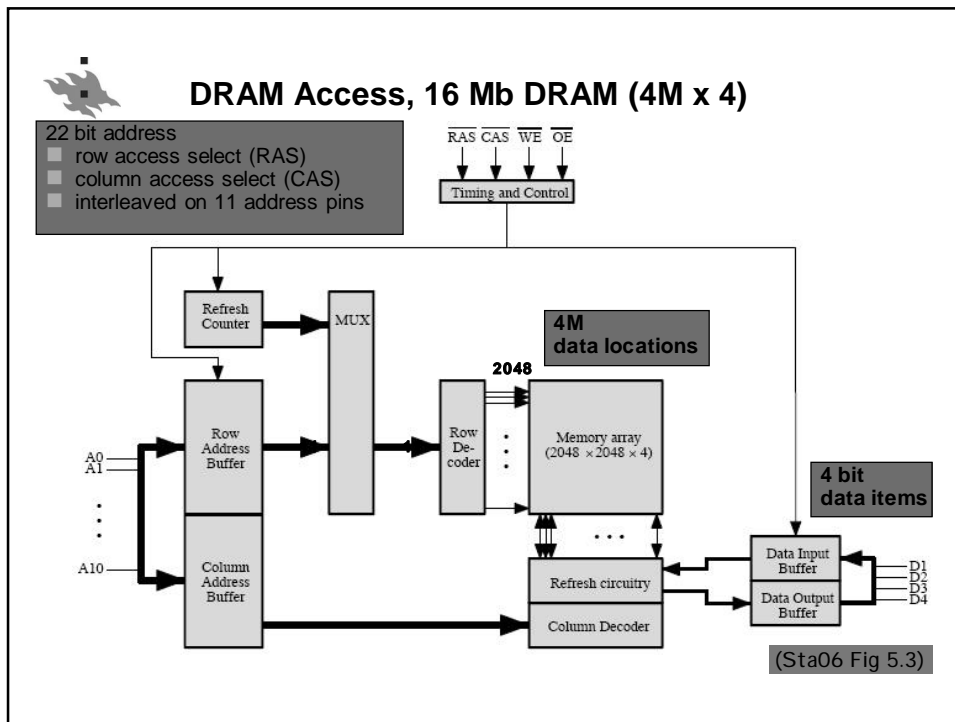
(Sta06 Table 5.1)

- Random access semiconductor memory
 - Direct access to each memory cell
 - Access time same for all cells



RAM

- Dynamic RAM, DRAM Analog: Charge on capacitors
 - Periodic refreshing required
 - Refresh required after read
 - Simpler, slower, denser, bigger (bytes per chip)
 - Access time ~ 60 ns
 - Main memory? (early systems)
- Static RAM, SRAM Digital: flip-flop gates
 - No periodic refreshing needed
 - Data remains until power is lost
 - More complex (more chip area/byte), faster, smaller
 - Access time ~ 2-5 ns
 - Level 2 cache?



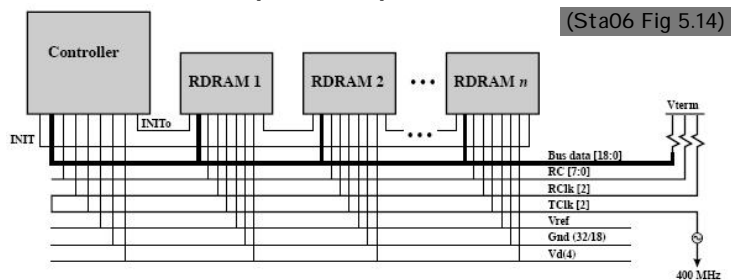


SDRAM (Synchronous DRAM)

- CPU clock synchronizes also the bus
 - Runs on higher clock speeds than ordinary DRAM
 - CPU knows how long it takes to make a reference,
 - can do other work while waiting
- 16 bits in parallel
 - Access 4 DRAMs (4 bits each) in parallel
 - Access time ~ 18 ns, transfer rate ~ 1.3 GB/s
- DDR SDRAM, double data rate
 - Current main memory technology
 - Supports transfers both on rising and falling edge of the clock cycle
 - Consumes less power
 - Access time ~ 12 ns, transfer rate ~ 3.2 GB/s



Rambus DRAM (RDRAM)



- Works with fast Rambus memory bus (800Mbps)
 - Controller + RDRAM modules
 - Access time ~ 12 ns, transfer rate ~ 4.8 GB/s
- Speed slows down with many memory modules
 - Serially connected on Rambus channel
 - Not good for servers with 1 GB memory (for now!)

ST1 Cell processor




Flash memory

- Based on transistors that are separated by a thin oxide layer
 - Flash cell is analog, not digital storage:
 - uses different charge levels to store 2 (or more) bits in each cell
- Non-volatile, data remains with power off
 - Electrical erasing in blocks = "flash"
 - Slow to write
 - Access time ~ 50 ns
- Used as a solid state storage
 - No moving parts
 - FlashBIOS in PC's, USB-memory
 - In phones, digital cameras, hand-held devices,....



MRAM

- Magnetoresistive Random Access Memory (MRAM)
 - Data stored with magnetic fields on two plates
 - Magnetic field directions determine bit value
- Non-volatile, data remains with power off
 - Fast to read/write
 - No upper limit for write counts (Flash has upper limit)
 - Access time comparable to DRAM
 - Almost as fast as SRAM
- Future open
 - Small market share now
 - Expensive now (2006: \$25 4Mbit)
 - Still under development
 - May replace flash in a few years
 - May replace SRAM later on
 - May replace DRAM and become "universal memory"



<http://www.research.ibm.com/journal/rd/501/maffitt.html>



Kertauskysymyksiä/Review questions

- Memory hierarchy and principle of locality?
- Different ways to use locality in cache solutions?
- Differences of associative and set associative mappings?
- Why to have separate caches for instructions and data ?

- Muistihierarkia ja paikallisuus?
- Millä tavoin paikallisuutta huomioidaan välimuistitarkaisussa?
- Assosiatiivisen ja joukkoassosiatiivisen kuvauksen erot?
- Miksi käskyille oma välimuisti ja datalle oma?