

### Tietokoneen rakenne

# Internal Memory, Cache

Stallings: Ch 4, Ch 5

- Key Characteristics
- Locality
- Cache
- Main Memory

From Computer Outstrip Encyclopedia © 1999 The Computer Language Co. Inc.

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 1

### Key Characteristics of Memories / Storage

<b>Location</b>	<b>Performance</b>
Processor	Access time
Internal (main)	Cycle time
External (secondary)	Transfer rate
<b>Capacity</b>	<b>Physical Type</b>
Word size	Semiconductor
Number of words	Magnetic
<b>Unit of Transfer</b>	Optical
Word	Magneto-Optical
Block	<b>Physical Characteristics</b>
<b>Access Method</b>	Volatile/nonvolatile
Sequential	Erasable/nonerasable
Direct	<b>Organization</b>
Random	
Associative	

(Sta06 Table 4.1)

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 2

### 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

cache

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

virtual memory

HW help for SW solution

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 3

### Memory Hierarchy

- Most often needed data kept close
- Access to small data sets can be made fast
  - simpler circuits
- 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

(Sta06 Fig 4.1)

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 4

### 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)

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 5

### 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] [233] [71] [8] [305] [63] [91] [2]

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 6

### Tietokoneen rakenne

# Cache

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 7

### Cache Memory (välimuisti)

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

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 8

### Cache

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 9

### Cache Read

```

    graph TD
      START([START]) --> RA[Receive address RA from CPU]
      RA --> Q{Is block containing RA in cache?}
      Q -- Miss --> MM[Access main memory for block containing RA]
      Q -- Hit --> FC[Fetch RA word and deliver to CPU]
      MM --> AL[Allocate cache line for main memory block]
      AL --> LML[Load main memory block into cache line]
      LML --> FC
      AL --> DR[Deliver RA word to CPU]
      FC --> DONE([DONE])
      DR --> DONE
  
```

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 10

### Cache Organization

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 11

### Cache Design

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

- n **Size**
  - u Many blocks help for temporal locality
  - u Large blocks help for spatial locality
  - u Multi-level cache

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 12

### Mapping

- Which block contains the memory location?
- Is the block in cache?
- Where is it located?

**Solutions**

- direct mapping (suora kuvaus)
- fully associative mapping (joukkoassosiatiivinen)
- set associative mapping (täysin assosiatiivinen)

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 13

### Direct Mapping (4)

- Each block has only one possible location (line) in cache
  - determined by index field bits
- Several blocks may map into same cache line
  - identified with tag field bits

Block number (in memory) → tag (21 bits) | index (8 bits) | offset (5 bits)

34 bit address (byte address)

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

Unique bits that are different for each block, Stored into cache line = Fixed location in cache =  $2^8 = 256$  blocks = 8 KB

Sta06 Fig 4.7  
PaHe98 Fig 7.10

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 14

### Direct Mapping Example (5)

Word = 4B (here)      Block size =  $2^3 = 8$  bytes = 64 bits  
Cache line size

ReadW I2, 0xA4

8 bit address (byte address): tag (2 bits) | index (3 bits) | offset (3 bits) → 10 | 100 | 100

tag	block, 64b
000:	
001:	
010:	
011:	01 54 A7 00 91 23 66 32 11
100:	11 77 55 55 66 66 22 44 22
101:	01 65 43 21 98 76 65 43 32
110:	
111:	

compare → No match

Read new memory block from memory address 0xA0=1010 0000 to cache location 100, update tag, and then continue with data access

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 15

### Direct Mapping Example 2 (5)

ReadW I2, 0xB4

8 bit address (byte address): tag (2 bits) | index (3 bits) | offset (3 bits) → 10 | 110 | 100

tag	block
000:	
001:	
010:	
011:	01 54 A7 00 91 23 66 32 11
100:	11 77 55 55 66 66 22 44 22
101:	01 65 43 21 98 76 65 43 32
110:	10 00 11 22 33 44 55 66 77
111:	

compare → Match

Start with 4th byte

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 16

### Fully Associative Mapping (6)

- Each block can be in any cache line
  - tag must be complete block number

Alpha A XP uses 34 bit memory addresses

Block number (in memory) → tag (29 bits) | offset (5 bits)

34 bit address (byte address)

Offset from the beginning of the block (in bytes)

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

Unique bits that are different for each block

Each block can be anywhere Cache size can be any number of blocks

Sta06 Fig 4.9

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 17

### Fully Associative Example (4)

ReadW I2, 0xB4

34 bit address (byte address): tag (5 bits) | offset (3 bits) → 10110 | 100

tag	block
000:	11011 12 34 56 78 9A 01 23 45
001:	10111 87 00 32 89 65 A1 B2 00
010:	00011 87 54 00 89 65 A1 B2 00
011:	10100 54 A7 00 91 23 66 32 11
100:	00111 77 55 55 66 66 22 44 22
101:	10100 65 43 21 98 76 65 43 32
110:	10110 00 11 22 33 44 55 66 77
111:	10011 87 54 32 89 65 A1 B2 00

Parallel! → Match

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 18

### Fully Associative Mapping

- lots of circuits
  - tag fields are long - wasted space?
  - each cache line tag must be compared parallelly with the memory address tag
    - lots of wires, comparison circuits
    - large surface area on chip
- Final comparison "or" has large gate delay
  - did any of these 64 comparisons match?
    - $\log_2(64) = 6$  levels of binary OR-gates
  - how about 262144 comparisons?
    - 18 levels?

Can use it only for small caches

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 19

### 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 22 set 7 offset 5

Block size =  $2^5 = 32$  B

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

Nr of sets =  $v = 2^7 = 128$  blocks = 4 KB

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

Sta06 Fig 4.11  
PaHe98 Fig 7.19

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 20

### 2-way Set Associative Cache

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

3 2 3 tag set offset 8 bit address (byte address)

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

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 21

### 2-way Set Assoc. Cache Example (5)

ReadW 12, 0xB4

tag set offset 101 10 100

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

Parallel! ? Match

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 22

### 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)
    - fully associative mapping
  - Minimum (1)
    - direct mapping

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 23

### 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?

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 24

### Cache Write Policy – memory writes?

- n **Write through** (läpikirjoittava)
  - u Each write goes always to cache and memory
  - u Each write is a cache miss!
- n **Write back** (lopuksi/takaisin kirjoittava)
  - u Each write goes only to cache
  - u Write cache block back to memory only when it is replaced in cache
  - u Memory may have stale (old) data
  - o 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)

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 25

### Cache Line Size

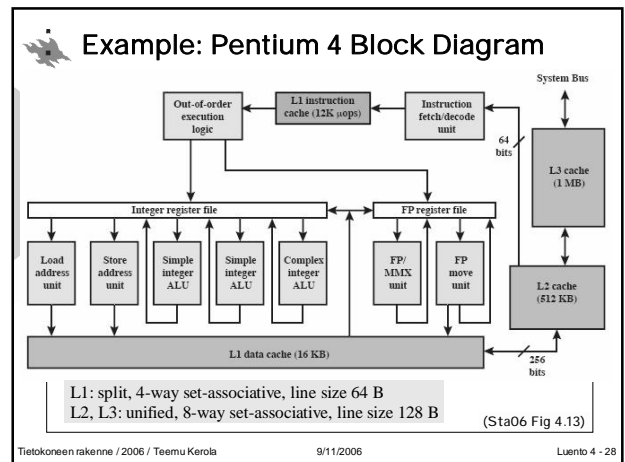
- n **How big cache line?**
- n **Optimise for temporal or spatial locality?**
  - u bigger cache line  $\bar{o}$  better for spatial locality
  - u more cache lines  $\bar{o}$  better for temporal locality
- n **Best size varies with program or program phase?**
- n **Best size different with code and data?**
- n **2-8 words?**
  - u word = 1 float??

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 26

### Types and Number of Caches

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

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 27



### Tietokoneen rakenne

# Main Memory

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 29

### 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		
Programmable ROM (PROM)			Read-mostly memory	Electrically, block-level	Electrically
Erasable PROM (EPROM)	UV light, chip-level				
Electrically Erasable PROM (EEPROM)	Electrically, byte-level				
Flash memory					

(Sta06 Table 5.1)

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

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 30

## RAM

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

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 31

## DRAM Access, 16 Mb DRAM (4M x 4)

22 bit address

- n row access select (RAS)
- n column access select (CAS)
- n interleaved on 11 address pins

(Sta06 Fig 5.3)

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 32

## 256-KB DRAM Memory Organization

- n Simultaneous access to 256K 8-bit word memory chip to access larger data items
- n Access 64-bit words in parallel? Need 8 chips.

(Sta06 Fig 5.5)

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 33

## SDRAM (Synchronous DRAM)

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

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 34

## Rambus DRAM (RDRAM)

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

(Sta06 Fig 5.14)

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 35

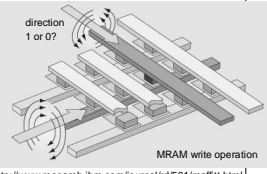
## Flash memory

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

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 36

## MRAM

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



direction  
1 or 0?

MRAM write operation

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

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 37

## Kertauskysymyksiä

- n Muistihierarkia ja paikallisuus?
- n Millä tavoin paikallisuutta huomioidaan välimuistitratkaisussa?
- n Assosiativisen ja joukkoassosiativisen kuvauksen erot?
- n Miksi käskyille oma välimuisti ja datalle oma?

Tietokoneen rakenne / 2006 / Teemu Kerola 9/11/2006 Luento 4 - 38