

Luento 4

Tietokoneen rakenne

Internal Memory, Cache (välimuisti)

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

The diagram illustrates the memory hierarchy. At the top is the CPU, which has an L1 Cache built into its chip. Below the CPU is the L2 Cache, which is an SRAM memory bank. Both the CPU and the L2 Cache are connected to a Local bus. The Local bus is also connected to the RAM (main memory), which is represented by a long strip of memory modules.

Stallings: Ch 4, Ch 5

- n Key Characteristics
- n Locality
- n Cache
- n Main Memory

Tietokoneen rakenne / 2007 / Liisa Marttinen
11/9/2007
Luento 4 - 1

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)

Tietokoneen rakenne / 2007 / Liisa Marttinen
11/9/2007
Luento 4 - 2

Goals

- n I want my memory lightning fast
- n I want my memory to be gigantic in size
- n Register access viewpoint
 - u data access as fast as HW register
 - u data size as large as memory
- n Memory access viewpoint
 - u data access as fast as memory
 - u data size as large as disk

cache

HW solution

virtual memory

HW help for SW solution

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 3

Memory Hierarchy

- n Most often needed data kept close
- n Access to small data sets can be made fast
 - u simpler circuits
 - u smaller gate delays
- n Faster ~ more expensive
- n 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 / 2007 / Liisa Marttinen 11/9/2007 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 μ s
 Prob (the rest) = 1% "Cost" (the rest) = 20 μ s



Aver cost = 99% * 2 μ s + 1% * 20 μ s = 2.2 μ s

- 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 / 2007 / Liisa Marttinen 11/9/2007 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	23	71	8	305	63	91	2
-----	----	----	---	-----	----	----	---

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 6

Tietokoneen rakenne

Cache

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 7

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 Type	Access Time
hand	0.5 sec (register)
table	1 sec (cache)
refrigerator	10 sec (memory)
moon	12 days (disk)
Europa (Jupiter)	4 years (tape)

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 8

Cache Memory (välimuisti)

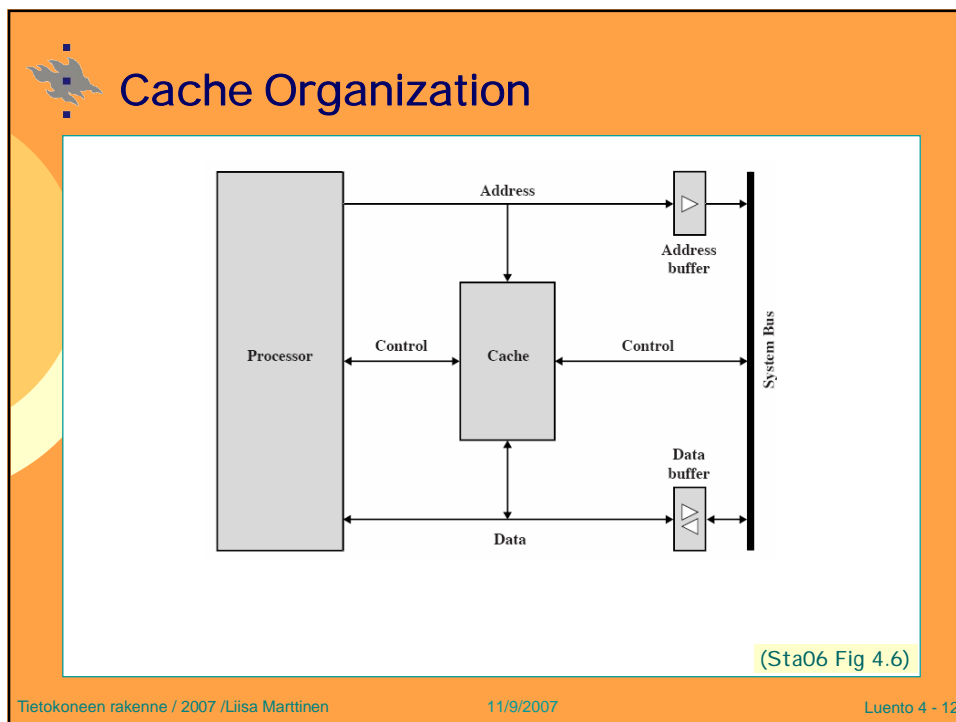
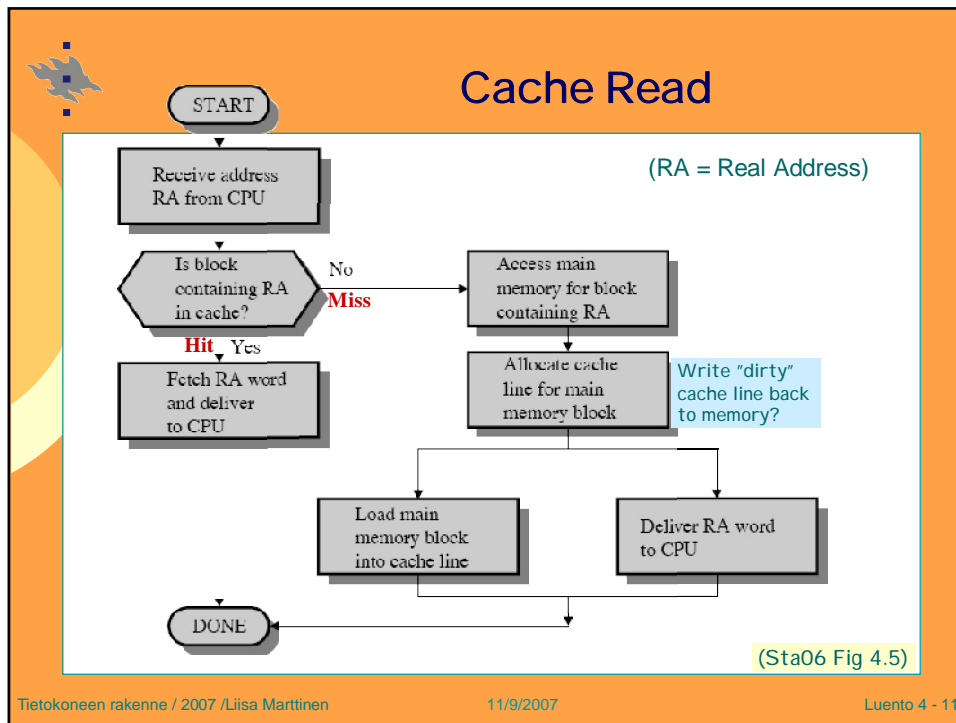
- n How to access main memory as fast as registers?
- n Locality → Use (CPU) cache!
 - u Keep most probably referenced data in fast cache close to processor, and rest in memory
 - u Most of data accesses only to cache
 - § hit ratio 0.9-0.99
 - u Cache is much smaller than main memory
 - u Cache is (much)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

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 9

Cache

(Sta06 Fig 4.3, 4.4)

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 10



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	

n **Cache Size & Line Size**

- u Many blocks help for temporal locality
- u Large blocks help for spatial locality
- u Larger cache is slower
- u Multi-level cache

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

(Sta06 Table 4.2)

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 13

Mapping

n Which block contains the memory location?

n Is the block in cache?

n Where is it located?

n **Solutions**

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

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

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 14

Direct Mapping (4)

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

Block number (in memory)

0x2480
0x6480
0xA480

34 bit address
(byte address)

tag

index

offset

byte

21

8

5

Unique bits that are different for each block. Stored into cache line. ~2 million different!

Fixed location in cache
 Z fixed cache size
 = $2^8 = 256$ blocks = 8 KB

Cache line size ~ Block size = $2^5 = 32$ B

Sta06 Fig 4.7
PaHe98 Fig 7.10

Tietokoneen rakenne / 2007 / Liisa Marttinen
11/9/2007
Luento 4 - 15

Main memory block (a' 32 B) groups

0

256 * 32 = 8192 B

8192

16324

...

Cache lines

Tag	Data	
		0
		1
wich?		2
		3
		4
		...
		255

34 b address = memory size ~ 17 GB

Tietokoneen rakenne / 2007 / Liisa Marttinen
11/9/2007
Luento 4 - 16

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	index	offset
10	100	100

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

	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:		

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 17

Direct Mapping Example 2 (5)

ReadW I2, 0xB4

tag	index	offset
10	110	100

compare

Match

start with 4th byte

	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:		

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 18

Fully Associative Mapping (6)

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

Alpha AXP uses 34 bit memory addresses

34 bit address (byte address)

tag 29	offset 5
-----------	-------------

Block number (in memory) → tag
 Offset from the beginning of the block (in bytes) → offset
 Block size = $2^5 = 32$ 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 / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 19

Fully Associative Example (4)

ReadW I2, 0xB4

	tag 5	block 64
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

tag 5: 10110

offset 3: 100

Parallel! ?

Match

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 20

Fully Associative Mapping

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

☹ Can use it only for small caches

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 21

Set Associative Mapping

- n With set size $k=2$, each cache entry contain 2 blocks
 - u Use set (set index) field to find the cache entry
 - u Use tag to determine if the block belongs to the set
 - u 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 / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 22

2-way Set Associative Cache

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

3	2	3
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

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 23

2-way Set Assoc. Cache Example (5)

ReadW I2, 0xB4

tag	set	offset
101	10	100

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

Parallel!

?

Match

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 24

Set Associative Mapping

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

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 25

Cache Replacement Algorithm

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

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 26

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 A bit set
 - 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)

Coherence problems:

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

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 27

Cache Line Size

- n How big cache line?
- n Optimise for temporal or spatial locality
 - u bigger cache line →  better for spatial locality
 - u more cache lines →  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 / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 28

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

Trend towards split caches: Pentium, Power PC,... (instruction pipelining)

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 29

Example: Pentium 4 Block Diagram

The diagram illustrates the Pentium 4 architecture. It shows the System Bus connected to an L3 cache (1 MB) and an L2 cache (512 KB). The L2 cache is connected to the L1 instruction cache (12K μops) and the L1 data cache (16 KB). The L1 instruction cache feeds into the Instruction fetch/decode unit, which then feeds into the Out-of-order execution logic. The Out-of-order execution logic feeds into the Integer register file and the FP register file. The Integer register file feeds into the Load address unit, Store address unit, Simple integer ALU, and Complex integer ALU. The FP register file feeds into the FP/MMX unit and the FP move unit. The L1 data cache feeds into the Load address unit, Store address unit, Simple integer ALU, and Complex integer ALU. The FP/MMX unit and FP move unit feed into the FP register file. The L2 cache is connected to the L1 data cache via a 256-bit bus. The L3 cache is connected to the L2 cache via a 64-bit bus.

L1: split, 4-way set-associative, line size 64 B
 L2, L3: unified, 8-way set-associative, line size 128 B

(Sta06 Fig 4.13)

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 30

Tietokoneen rakenne

Main Memory

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 31

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	Read-mostly memory	Electrically, block-level		

(Sta06 Table 5.1)

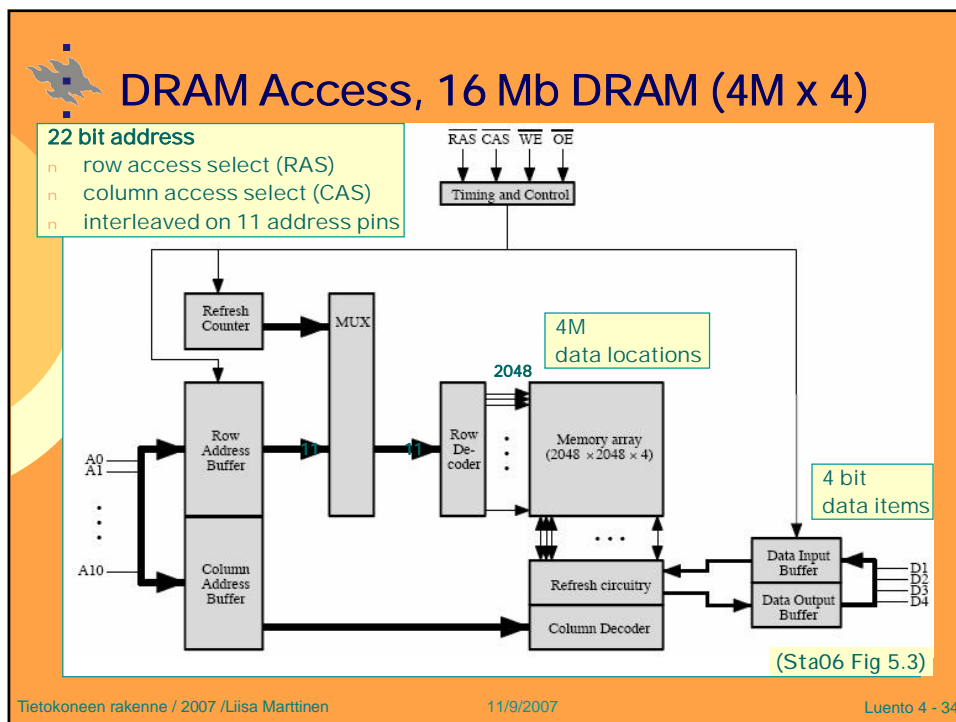
- Random access semiconductor memory
 - Direct access to each memory cell
 - Access time same for all cells

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 32

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 / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 33



256-KB DRAM Memory Organization

The diagram illustrates the memory organization for a 256-KB DRAM. It features a Memory address register (MBR) with 9 row and 9 column inputs. These inputs are connected to two memory chips, each labeled 'Decode 1 of 512' and '512 words by 512 bits Chip #1'. The output of each chip is connected to a Memory buffer register (MBR) with 8 outputs, numbered 1 through 8. The diagram is labeled '(Sta06 Fig 5.5)'.

- 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.

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 35

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 / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 36

Rambus DRAM (RDRAM)

(Sta06 Fig 5.14)

- 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!)

STI Cell processor

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 37

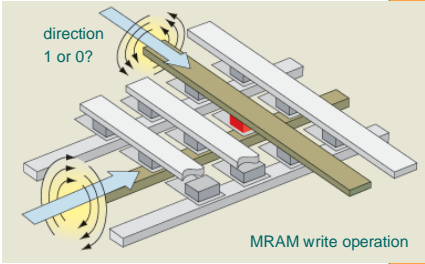
Flash memory

- 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 FlashBI OS in PC's, USB-memory
 - u In phones, digital cameras, hand-held devices,....

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 38

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 (Flash has upper limit)
 - 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 / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 39

Kertauskysymyksiä

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

Tietokoneen rakenne / 2007 / Liisa Marttinen 11/9/2007 Luento 4 - 40