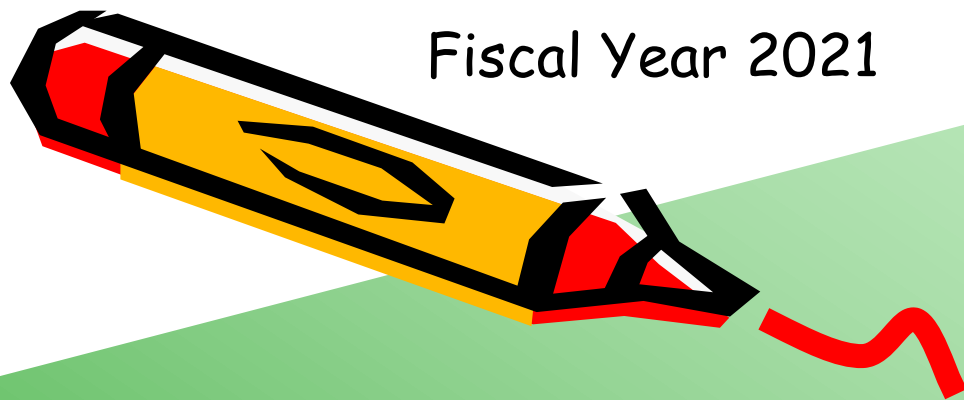Fiscal Year 2021
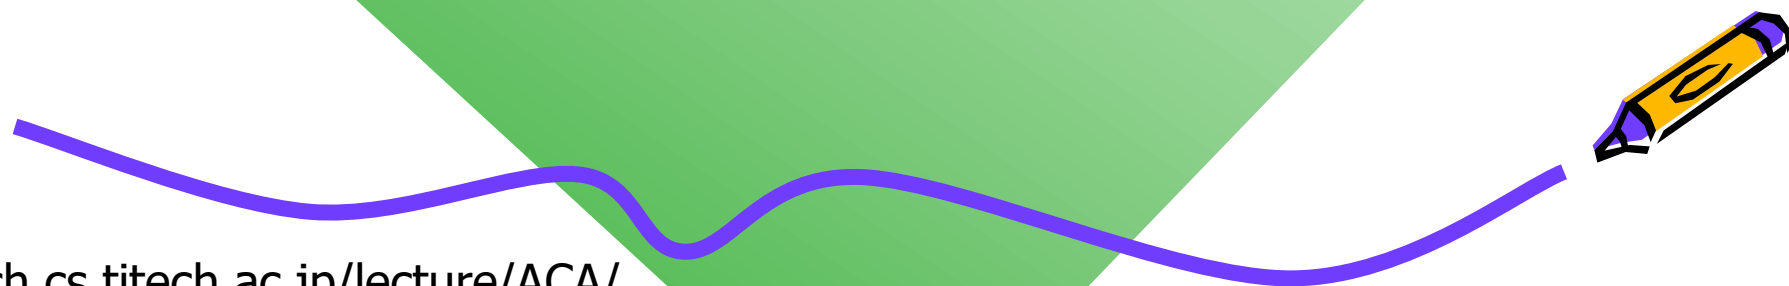
Course number: CSC.T433
School of Computing,
Graduate major in Computer Science

Advanced Computer Architecture

3. Memory Hierarchy Design

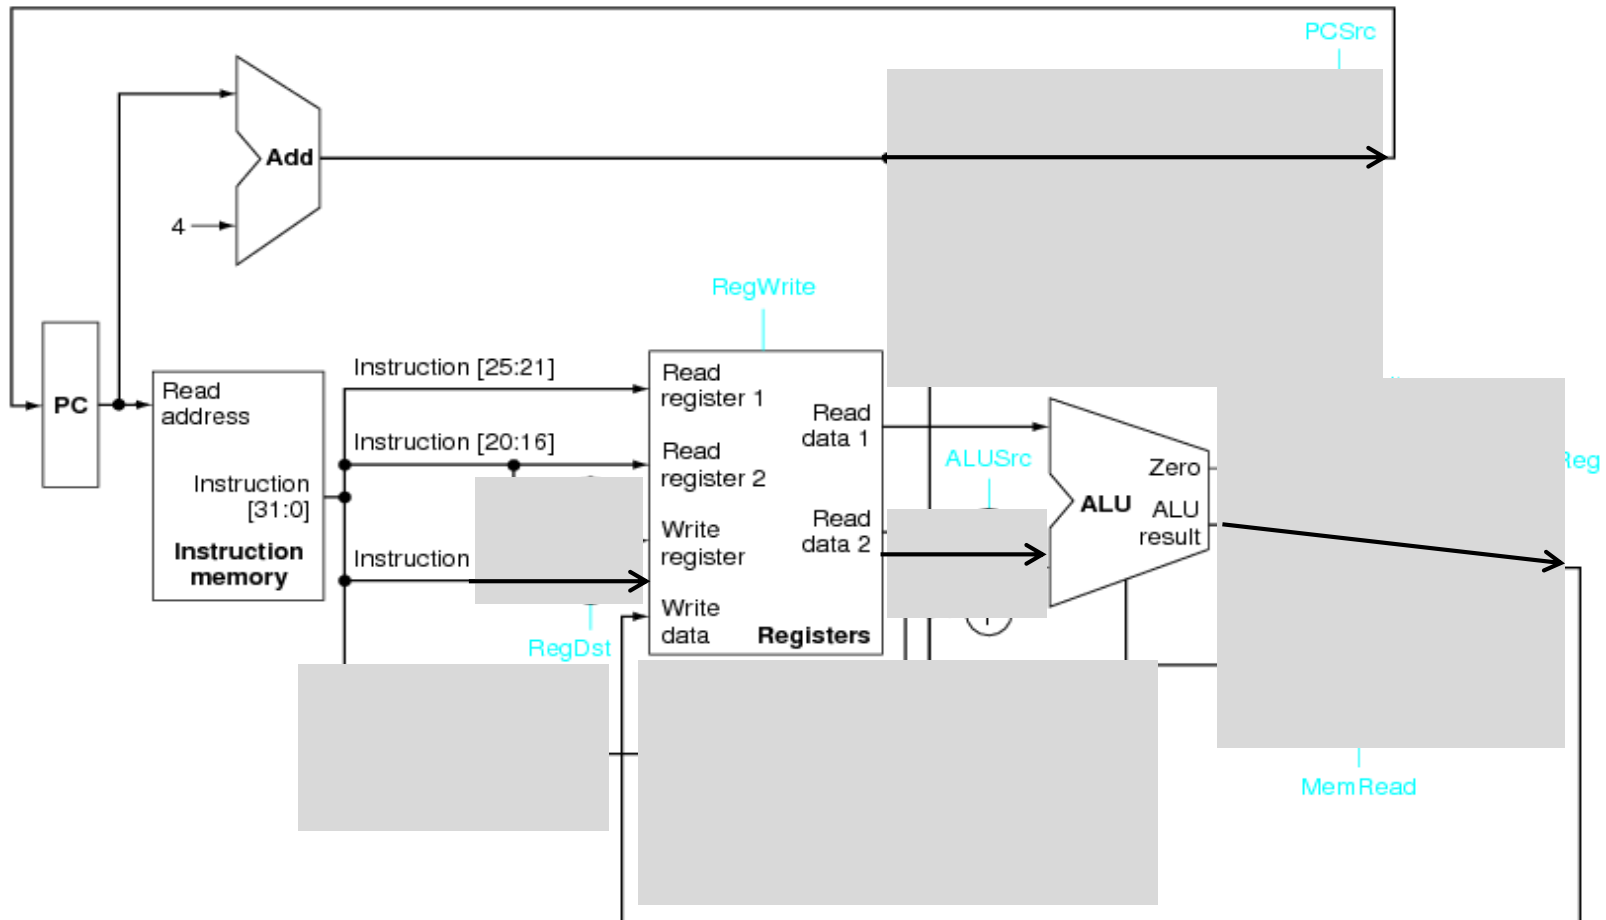www.arch.cs.titech.ac.jp/lecture/ACA/

Room No.W936

Mon 14:20-16:00, Thr 14:20-16:00

Kenji Kise, Department of Computer Science
kise _at_ c.titech.ac.jp

Datapath of single-cycle processor supporting ADD

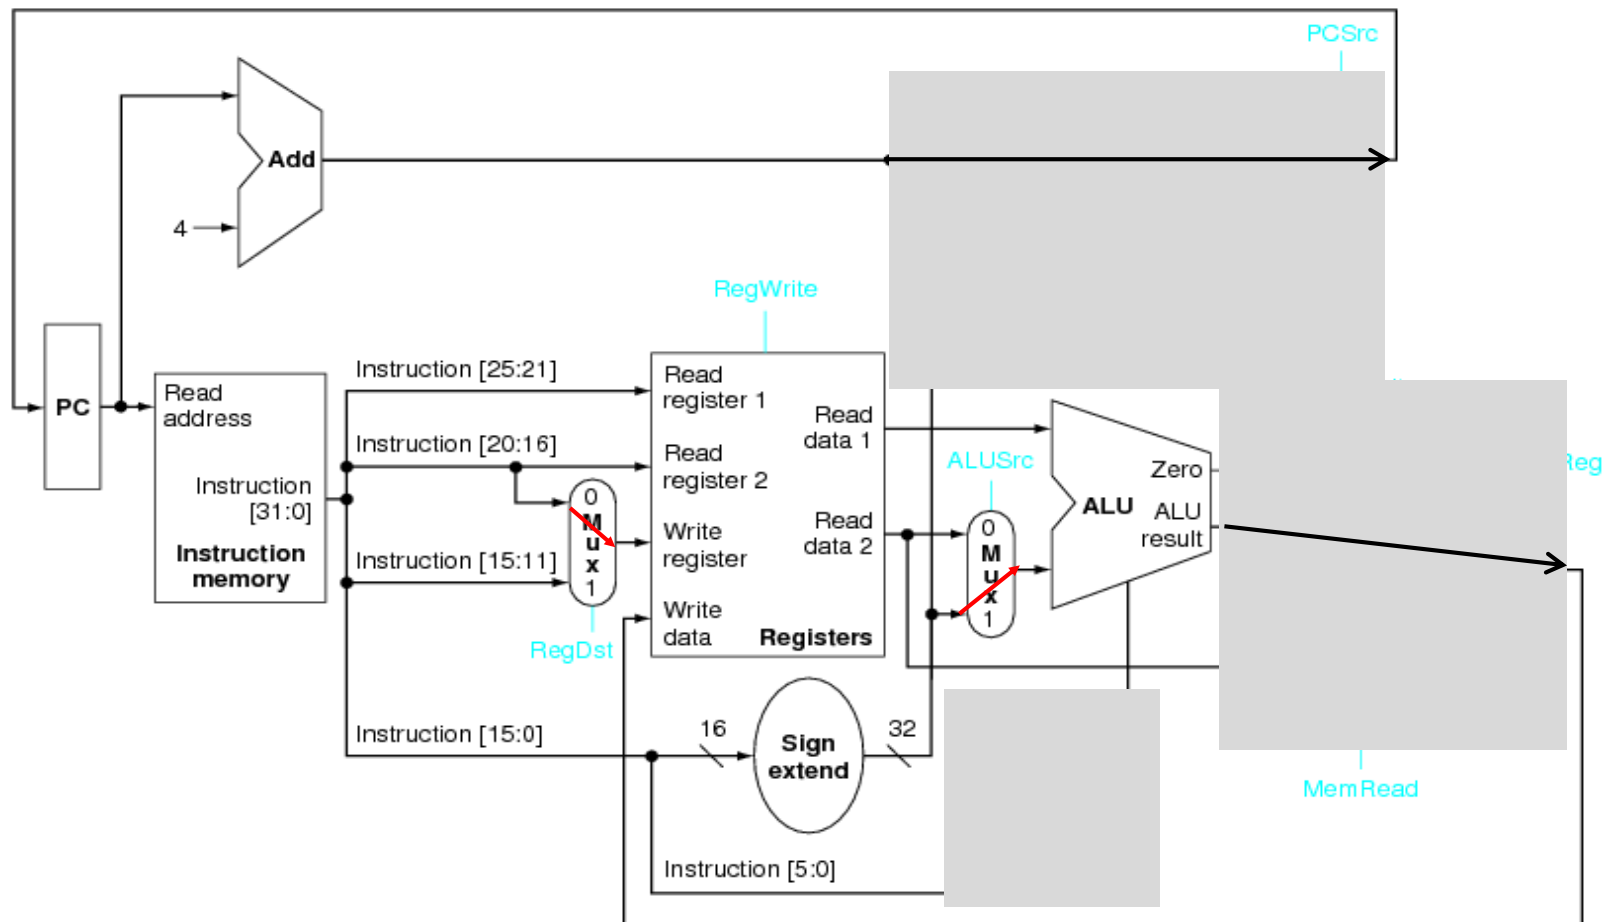| op	rs	rt	rd	shamt	funct

IR[25:21]   IR[20:16]   IR[15:11]

0x800   add $t0, $s1, $s2   [add $8, $17, $18]

$17 = 3
$18 = 4

Datapath of processor supporting ADD and ADDI

op	rs	rt	16 bit immediate

IR[25:21]   IR[20:16]

I format

0x804 addi $t1, $t0, 3   [addi $9, $8, 3]

$8 = 7

Waveform of proc02
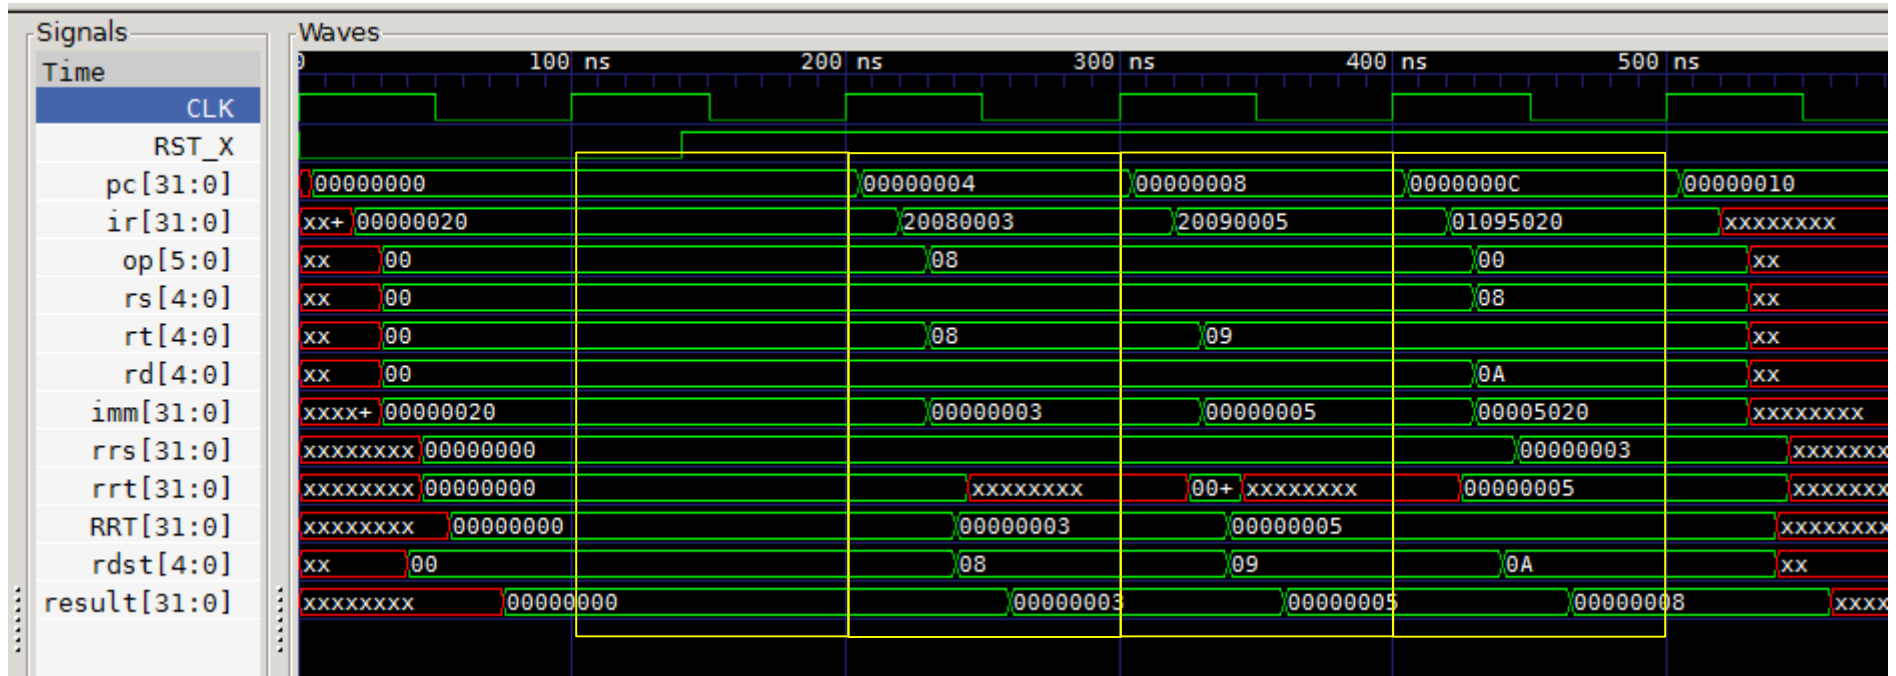

```
add  $0,  $0,  $0    # NOP {6'h0, 5'd0, 5'd0, 5'd0, 5'd0, 6'h20}
addi $t0, $zero, 3  # {6'h8, 5'd0, 5'd8, 16'd3}
addi $t1, $zero, 5  # {6'h8, 5'd0, 5'd9, 16'd5}
add  $t2, $t0, $t1  # {6'h0, 5'd8, 5'd9, 5'd10, 5'd0, 6'h20}
```

MIPS Memory Access Instructions

- MIPS has two basic data transfer instructions for accessing memory

```
lw  $t0, 4($s3)   #load word from memory
sw  $t0, 8($s3)   #store word to memory
```

- The data is loaded into (lw) or stored from (sw) a register in the register file – a 5 bit address

- The memory address – a 32 bit address – is formed by adding the contents of the base address register to the offset value

  - A 16-bit field meaning access is limited to memory locations within a region of $\pm 2^{13}$ or 8,192 words ($\pm 2^{15}$ or 32,768 bytes) of the address in the base register

  - Note that the offset can be positive or negative

Machine Language - Load Instruction

- Load/Store Instruction Format (I format):

$$\texttt{lw} \quad \texttt{\$t0,} \quad \texttt{24(\$s2)}$$

op	rs	rt	16 bit offset

Memory

$24_{10} + \$s2 =$

```
  ...0001 1000
+ ...1001 0100
  ...1010 1100 = 0x120040ac
```

$t0 ←   0x120040ac

$s2 →   0x12004094

0xf f f f f f f f

0x0000000c
0x00000008
0x00000004
0x00000000

data   word address (hex)

Datapath of processor supporting ADD, ADDI, LW

	IR[25:21]	IR[20:16]		
op	rs	rt	16 bit immediate	I format

0x808   lw $t2, 4($t0)   [lw $10, 4($8)]

$8 = 0x10

mem[0x14] = 3

A Typical Memory Hierarchy

❑ By taking advantage of the principle of locality

- Present much memory in the cheapest technology
- at the speed of fastest technology

On-Chip Components

	Control					
Datapath	RegFile	ITLB	Instr Cache	Second Level Cache (SRAM)	Main Memory (DRAM)	Secondary Memory (Disk)
		DTLB	Data Cache			
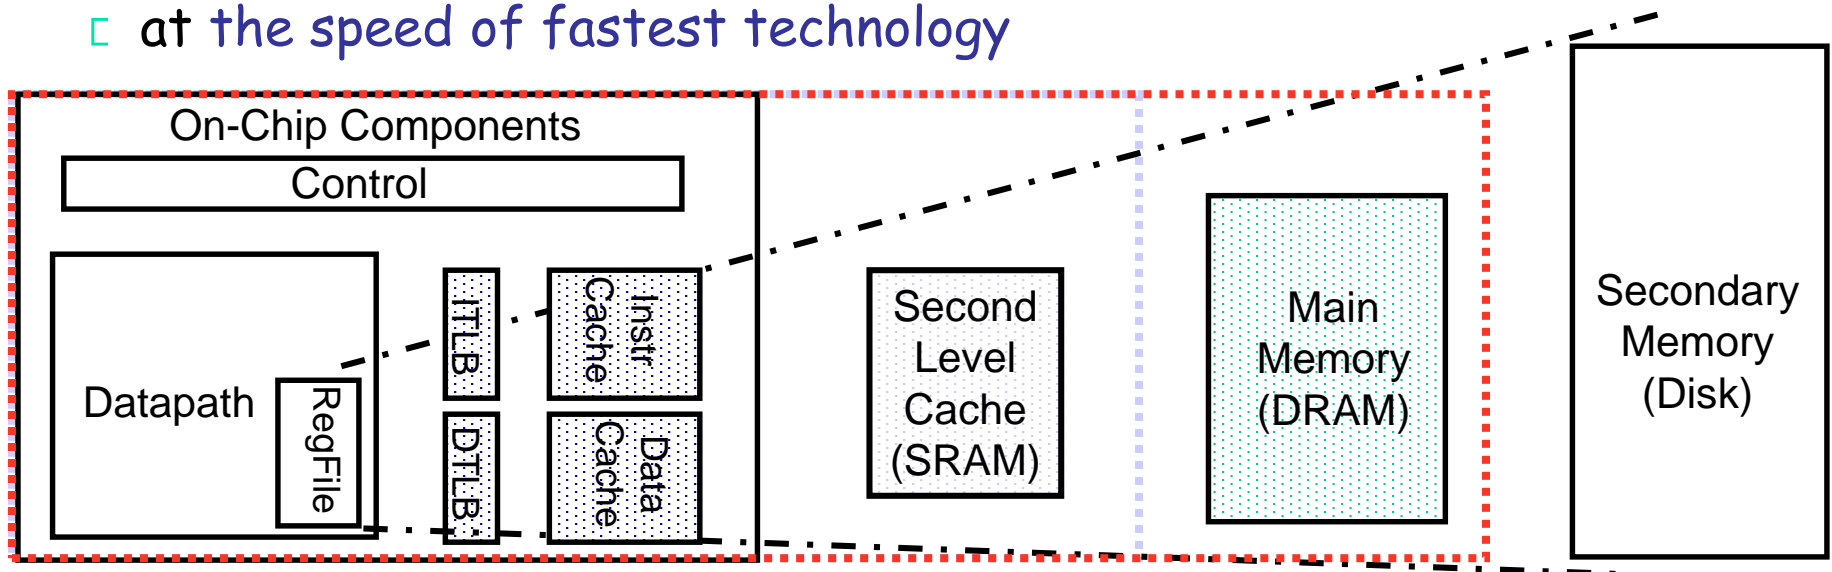
Speed (%cycles):	½'s	1's	10's	100's	1,000's
Size (bytes):	100's	K's	10K's	M's	G's to T's
Cost:	highest				lowest

TLB: Translation Lookaside Buffer

MIPS Direct Mapped Cache Example

- One word/block, cache size = 1K words (4KB)

What kind of locality are we taking advantage of?

Multiword Block Direct Mapped Cache

- Four words/block, cache size = 1K words (4KB)

Hit

Data

31 30 . . . 13 12 11 . . . 4 3 2 1 0

Byte offset

Tag

20

Index

8

Block offset

Data (4 word)

Index Valid Tag

0
1
2
.
.
.
.
253
254
255

20

=

32

What kind of locality are we taking advantage of?

Four-Way Set Associative Cache

- 2^8 = **256 sets** each with four ways (each with one block)

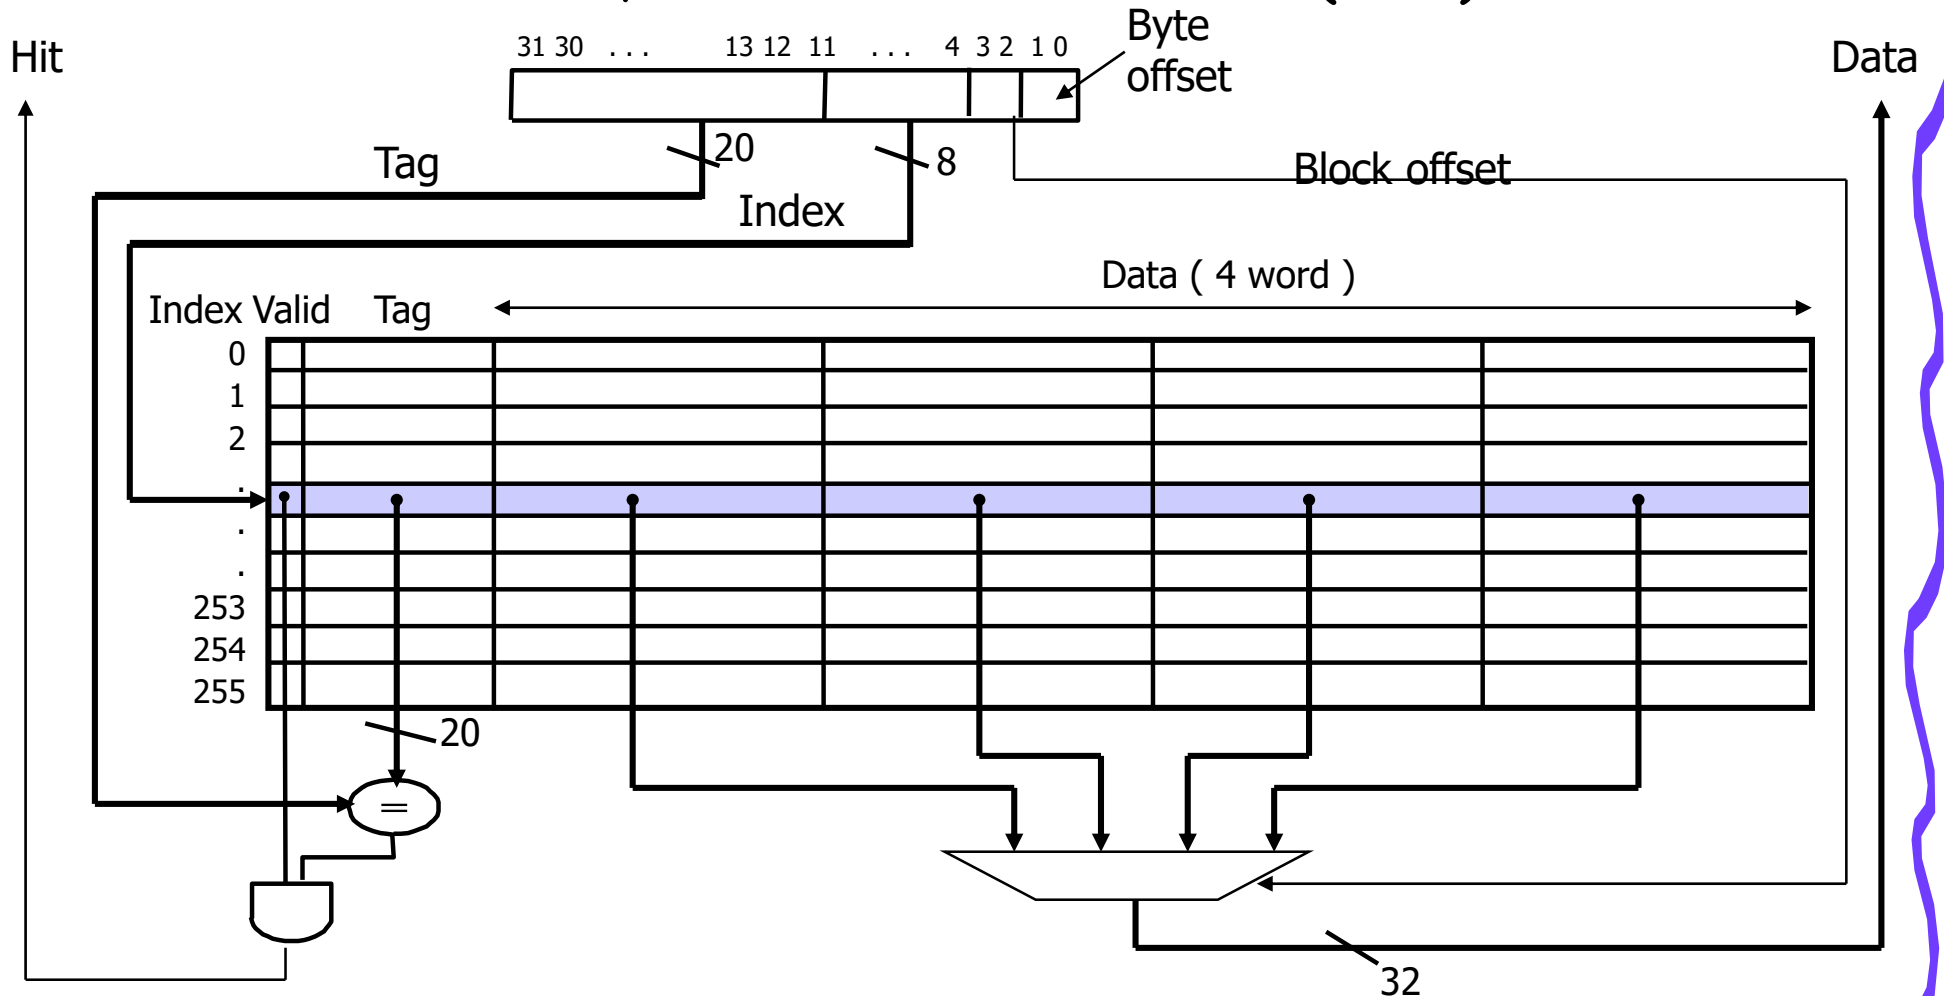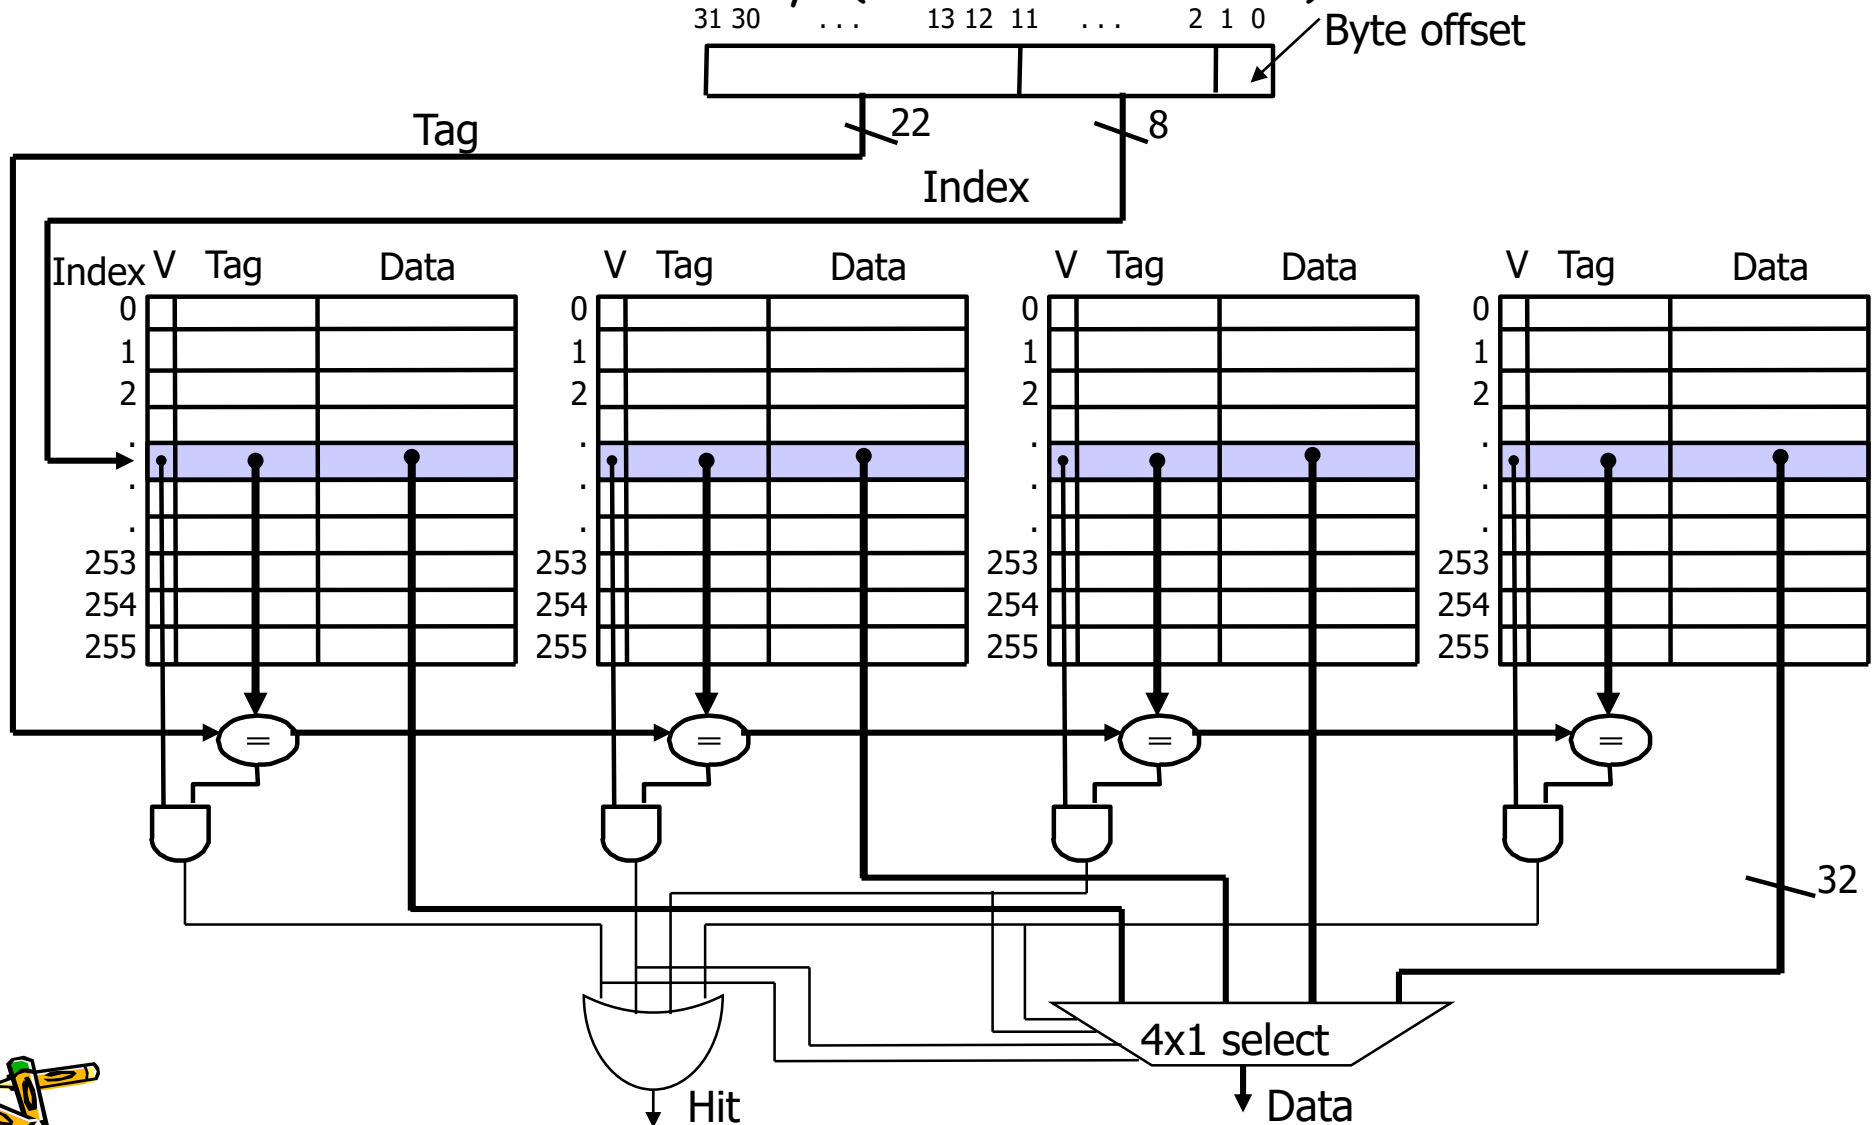31 30 . . . 13 12 11 . . . 2 1 0

Byte offset

Tag

22

8

Index

Index	V	Tag	Data
0			
1			
2			
.			
253			
254			
255			

	V	Tag	Data
0			
1			
2			
.			
253			
254			
255			

	V	Tag	Data
0			
1			
2			
.			
253			
254			
255			

	V	Tag	Data
0			
1			
2			
.			
253			
254			
255			

= = = =

32

4x1 select

Hit

Data

Cache Associativity & Replacement Policy

Book

Bookshelf

A B C D

Desk

Costs of Set Associative Caches

- N-way set associative cache costs
  - N comparators (delay and area)
  - MUX delay (set selection) before data is available
  - Data available after set selection and Hit/Miss decision.
- When a miss occurs,
  which way's block do we pick for replacement ?
  - Least Recently Used (LRU):
    the block replaced is the one that has been unused for the longest time
    - Must have hardware to keep track of when each way's block was used
    - For 2-way set associative, takes one bit per set →
      set the bit when a block is referenced
   (and reset the other way's bit)
  - Random

Recommended Reading

- Emulating Optimal Replacement with a Shepherd Cache
    - Kaushik Rajan, Govindarajan Ramaswamy, Indian Institute of Science
    - MICRO-40, pp. 445-454, 2007
   - Session 8: Cache Replacement Policies

- A quote:
    "The inherent temporal locality in memory accesses is filtered out by the L1 cache. As a consequence, an L2 cache with LRU replacement incurs significantly higher misses than the optimal replacement policy (OPT). We propose to narrow this gap through a novel replacement strategy that mimics the replacement decisions of OPT."

Memory Hierarchy Design

Memory Hierarchy

L2 and lower caches

- Objective : Need to reduce expensive memory accesses
- Design : Large size, Higher associativity, Complex design
- Problem : Do not interact with program directly and observe filtered temporal locality

- High Associativity $\implies$ replacement policy crucial to performance
- L1 cache services temporal accesses $\implies$ Lack of temporal accesses at L2 $\implies$ LRU replacement inefficient
- Replacement decisions are taken off the processor critical path

Emulating Optimal Replacement with a Shepherd Cache, MICRO-2007

LRU has room for improvement

LRU vs OPT

MPKI values shown above bars: 87.5, 81.5, 73.9, 58.4, 80, 79.7, 72.6, 65.9

Legend: 512KB-lru16, 512KB-lruFA, 256KB-opt8, 512KB-opt16

Benchmarks (x-axis): art, mcf, gcc, luca, swim, appl, ammp, twolf, vpr, facer, mgrid, apsi, avg26

13%

MPKI for SPEC2000 suite, Benchmarks with MPKI < 5 not plotted but count towards average

Huge performance gap between LRU and OPT
OPT at half the size preferable to LRU at double the size

MPKI: Miss Per Kilo Instructions

OPT: Optimal Replacement Policy

The Optimal Replacement Policy

1. **Replacement Candidates** : On a miss any replacement policy could either choose to replace any of the lines in the cache or choose not to place the miss causing line in the cache at all.

2. **Self Replacement** : The latter choice is referred to as a self-replacement or a cache bypass

Optimal Replacement Policy

On a miss replace the candidate to which an access is least imminent [Belady1966,Mattson1970,McFarling-thesis]

3. **Lookahead Window** : Window of accesses between miss causing access and the access to the least imminent replacement candidate. Single pass simulation of OPT make use of lookahead windows to identify replacement candidates and modify current cache state [Sugumar-SIGMETRICS1993]

Emulating Optimal Replacement with a Shepherd Cache, MICRO-2007

Example of Optimal Replacement Policy

Understanding OPT

Access Sequence	A_5	A_1	A_6	A_3	A_1	A_4	A_5	A_2	A_5	A_7	A_6	A_8
OPT order for A_5		0		1		2	3	4				
OPT order for A_6				0	1	2	3				4	

- Consider 4 way associative cache with one set initially containing lines (A_1, A_2, A_3, A_4), consider the access stream shown in table
- Access A_5 misses, replacement decision proceeds as follows
  1. Identify replacement candidates : $(A_1, A_2, A_3, A_4, A_5)$
  2. Lookahead and gather imminence order : shown in table, lookahead window circled
  3. Make replacement decision : A_5 replaces A_2
- A_6 self-replaces, lookahead window and imminence order in table

Shepherd Cache emulation OPT

Emulating OPT with a Shepherd Cache

- Split the cache into two logical parts
  - Main Cache (MC) for which optimal replacement is emulated
  - Shepherd Cache (SC) used to provide a lookahead and guide replacements from MC towards OPT
- Operation
  1. Buffer lines temporarily in SC before moving them to MC, SC acts as a FIFO buffer
  2. While in SC, gather imminence information and emulate lookahead
  3. When forced out of SC, make an MC replacement based on the gathered imminence order

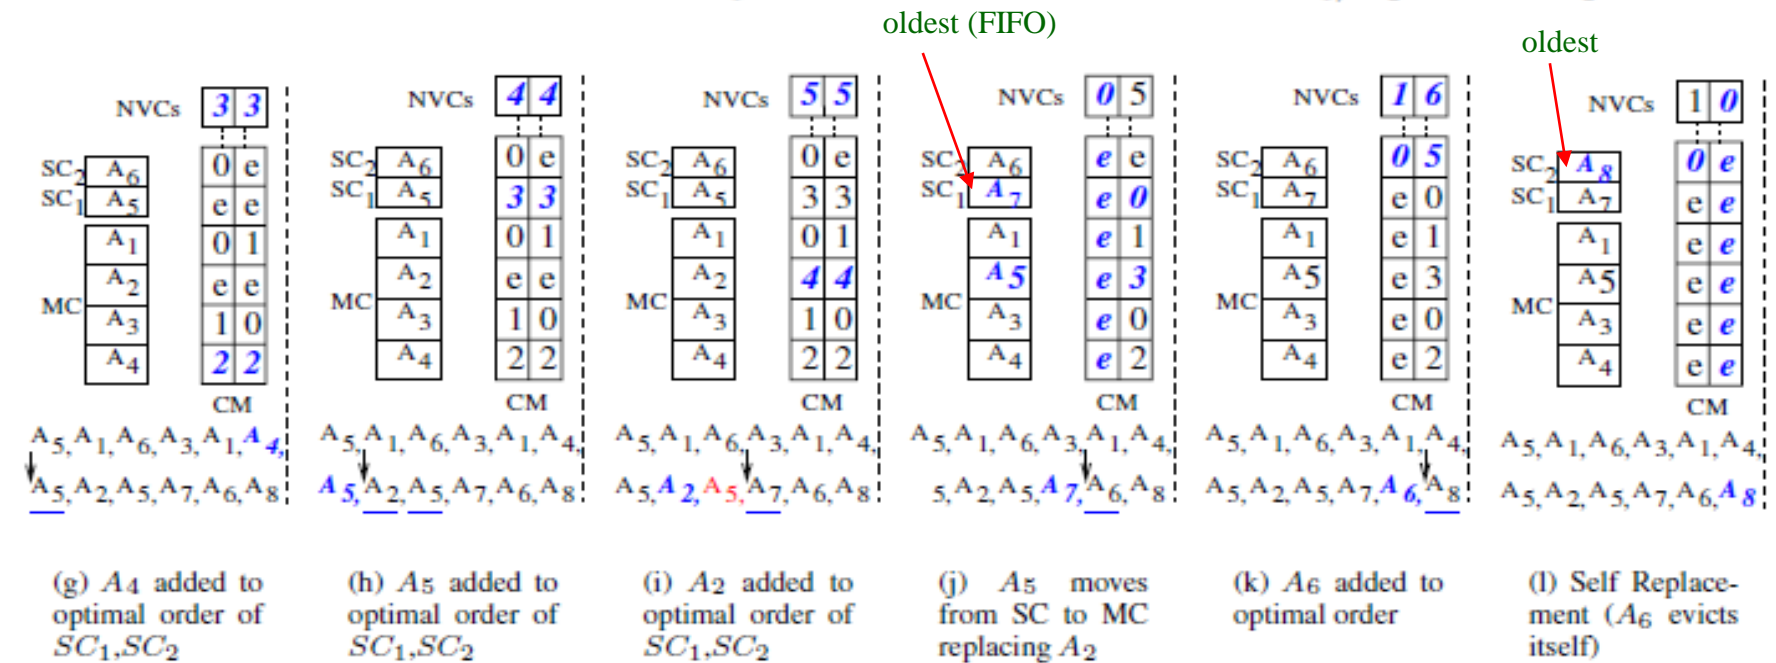Emulating Optimal Replacement with a Shepherd Cache, MICRO-2007

Shepherd Cache Overview

Overview of Shepherd Caching
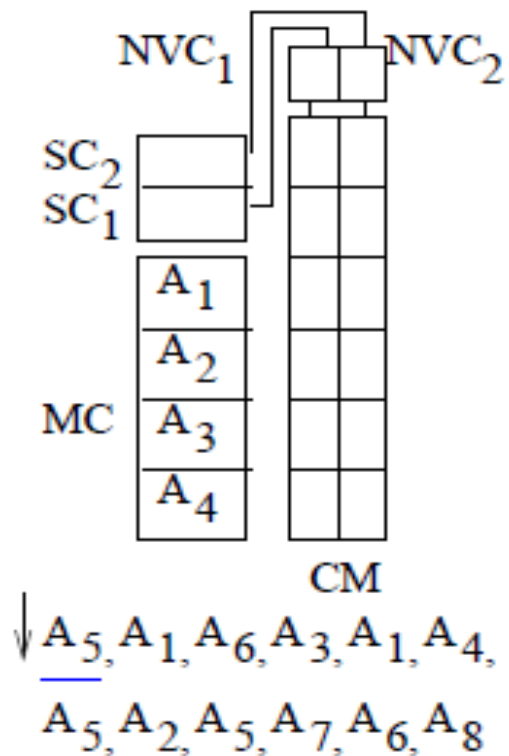
- To emulate MC with 4 ways per set and 2 SC ways per set
- To gather imminence order add a counter matrix (CM)
- CM has one column per SC way to track imminence order w.r.t to it
- CM has one row per SC and MC line as any of them can be a replacement candidate
- Each column has one Next Value Counter (NVC) to track the next value to assign along column

Emulating Optimal Replacement with a Shepherd Cache, MICRO-2007

(a) Initial State

(b) A_5 inserted at SC_1

(c) A_1 added to the optimal order of SC_1

(d) A_6 inserted at SC_2

(e) A_3 added to the optimal order of SC_1, SC_2

(f) A_1 added to optimal order of SC_2

(g) A_4 added to optimal order of SC_1, SC_2

(h) A_5 added to optimal order of SC_1, SC_2

(i) A_2 added to optimal order of SC_1, SC_2

(j) A_5 moves from SC to MC replacing A_2

(k) A_6 added to optimal order

(l) Self Replacement (A_6 evicts itself)

Shepherd cache bridges 32 – 52% of the gap

Bridging the performance gap

Bridging the LRU-OPT gap

- *SC-4 bridges 32-52% of gap*
- SC moves closer to OPT as cache size increases

Legend:
- lru-16 (0B)
- sc-12 (136B)
- sc-8 (92B)
- sc-6 (71B)
- sc-4 (48B)
- sc-2 (25B)
- opt-16 (-)

Average MPKI (y-axis), cache sizes 512KB, 1MB, 2MB, 4MB (x-axis)

Avg MPKI over SPEC2000 suite

MPKI: Miss Per Kilo Instructions

Emulating Optimal Replacement with a Shepherd Cache, MICRO-2007

Assignment 2

1. Design a single-cycle processor supporting MIPS add, addi, lw and sw instructions in Verilog HDL. Please download proc03.v from the support page and refer it.

2. Verify the behavior of designed processor using following assembly code

   - `add  $0,  $0,  $0   # NOP {6'h0, 5'd0, 5'd0, 5'd0, 5'd0, 6'h20}`
   - `addi $t0, $zero, 8  # {6'h8, 5'd0, 5'd8, 16'd8}`
   - `sw   $t0, 4($t0)    # {6'h2b,5'd8, 5'd8, 16'd4}`
   - `lw   $t1, 4($t0)    # {6'h23,5'd8, 5'd9, 16'd4}`
   - `addi $t2, $t1, 6    # {6'h8, 5'd9, 5'd10,16'h6}`

3. Submit your report in a PDF file via E-mail **by the next Thursday**.

   - The report should include a block diagram, a source code in Verilog HDL, and obtained waveforms of your design.
   - E-mail address : report@arch.cs.titech.ac.jp
   - E-mail title: Assignment of Advanced Computer Architecture