



# **CS-473: System programming for Systems on Chip**

Practical work 4

## **Caches**

**Version:**  
1.0

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Prerequisites . . . . .	1
1.2	Objectives . . . . .	1
<b>2</b>	<b>Exercises</b>	<b>2</b>
2.1	Structure Memory Layout and Caches . . . . .	2
2.2	Programming Tasks . . . . .	4

# 1 Introduction

## 1.1 Prerequisites

On moodle you find the file `pw4.zip`; it contains the complete source code template for this practical work. Download this file, and unzip it in the directory `programs`.

## 1.2 Objectives

In the scope of this practical work, you will learn about:

1. memory layouts of structures in C,
2. optimizing the memory layout for caches,
3. designing cache-friendly algorithms,
4. the interaction between I/O devices and caches.

## 2 Exercises

### 2.1 Structure Memory Layout and Caches



Please read “The Lost Art of Structure Packing” by Eric S. Raymond to answer the questions in this section. It is available at: <http://www.catb.org/esr/structure-packing/>.

Project `cache_sweep` under `cache/sweep/` executes a C program summarized in Figure 2.1. The program contains a struct `item_t` consisting of an ID and data, an array of `item_t`s, and function `items_find` that finds an `item_t` with a given ID. The full source code of the program is available under `cache/sweep/src/datapoint/datapoint.c`. The executable contains multiple versions of this program for different configurations described by the following parameters:

1. `PARAM_PACKED`: `item_t` has the `__packed` modifier or not.
2. `PARAM_DATALEN`: length of field `data` of struct `item_t`.
3. `PARAM_COUNT`: length of the array.

We aim to analyze how these parameters affect (1) the code generated by the compiler and (2) the number of data cache misses. Both plots in Figure 2.2 sweep `PARAM_DATALEN`. Figure 2.2a plots `sizeof(item_t)`. Figure 2.2b plots the number of data cache misses for function `items_find`. For all the data points in Figure 2.2b, function `items_find` accesses all the array elements.

#### Questions:

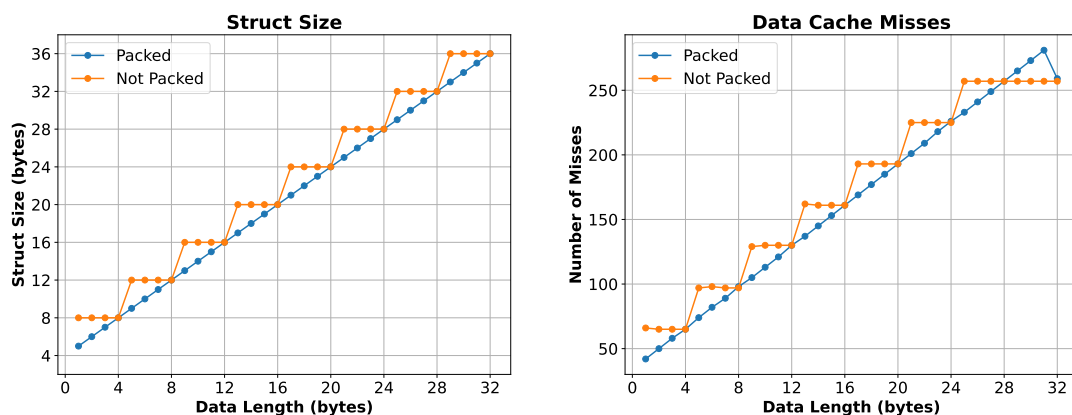
1. What does the `__packed` modifier do? Note that `__packed` is defined in Figure 2.1. Online resources and the compiler manual might help you with this part of the practical work.
2. Explain the memory layout of `item_t` with and without `__packed` and for `PARAM_DATALEN`  $\in \{1, 2, 3, 4\}$ . Specify (1) the ordering of the fields, (2) the size of each struct field, and (3) padding bytes (the processor has a 32-bit alignment). You do not need to make a table; however, be clear and concise!
3. Repeat the previous question for the memory layout of the array. Consider the first two array elements.
4. How does `__packed` affect the number of cache misses? Why?
5. Consider the data series in Figure 2.2 that do not use `__packed` (i.e., the orange lines). What is the cache line size for the data cache? Explain your approach.
6. How does `__packed` affect the generated assembly code size? Why? What does the compiler do differently? Generated assembly file for each data point is available under `cache/sweep/build-release-or1300/src/datapoint/datapoint_<COUNT>_<LENGTH>_<PACKED>.s`. **Important:** to have these files available you have to add `-save-temps` to line 7 in your makefile, it should like: `CFLAGS += -save-temps`

```

2  /** @note defined in 'defs.h'. */
3  #define __packed __attribute__((packed))
4
5  /** @brief Item struct. Relates an 'id' to a piece of 'data'. */
6  typedef struct __packed /* or nothing, depending on the configuration */ {
7      uint32_t id;
8      char data[PARAM_DATALEN];
9  } item_t;
10
11 /** @brief An array of items. */
12 static item_t items[PARAM_COUNT];
13
14 /**
15  * @brief Searches for an item matching the 'id'.
16  *
17  * @param id
18  * @return item_t* Pointer to the found item.
19  */
20 item_t* items_find(uint32_t id) {
21     for (size_t i = 0; i < PARAM_COUNT; ++i) {
22         if (items[i].id == id)
23             return &items[i];
24     }
25     return NULL;
26 }

```

Figure 2.1: A simplified version of the measured program. Each data point consists of PARAM\_DATALEN, PARAM\_COUNT, and PARAM\_PACKED. The program is compiled and executed for each data point. PARAM\_PACKED determines if the struct has \_\_packed modifier. For each data point, we record sizeof(item\_t) and the number of data cache misses.



(a) Structure size for varying data length.

(b) Number of data cache misses for varying data length.

Figure 2.2: Both figures vary the data length (PARAM\_DATALEN). Figure (b) is drawn for PARAM\_COUNT=256. The number of memory accesses is same as PARAM\_COUNT.

## 2.2 Programming Tasks



Please read Section 4 of “Cache-Conscious Data Structures” from Microsoft Research to answer the questions in this section. It is available at: <https://www.microsoft.com/en-us/research/wp-content/uploads/2016/12/ccds.pdf>.

In this part of the assignment, you will modify the source of project `cache_tasks` under `cache/tasks/`. Take a look at the directory structure and understand it. As a first step, please check your setup:

1. Compile the project as usual.
2. Upload the `.cmem` file and execute the program. Note that uploading might take up to 10 seconds.

The expected program output is the following (cache miss numbers might be slightly different):

```
task1_main
2 sizeof(node_t) = 64:
  struct layout for node_t:
4 member      -> offset
  node.id     -> 0x000
6 node.data   -> 0x004
  node.next   -> 0x038
8 node.prev   -> 0x03c
#0 -> #11 -> #6 -> #12 -> #9 -> #2 -> #4 -> #8 -> #1 -> #3 -> #7 -> #15 -> #14 ->
  #13 -> #10 -> #5 -> Done.
10 Task 1: dcache misses:      42
task2_main
12 Item ID = 15, data = random data: 11
  Task 2: dcache misses:      16
14 task3_main
  Task 3: dcache misses:      65893
16 out_vector verification successful!
task4_main
18 original address: 0x50000800
  new address: 0x08000800
```

## Task 1: Optimize node\_t

Consider the following source files:

- ▶ cache/tasks/src/task1.c
- ▶ cache/tasks/src/node.c
- ▶ cache/tasks/include/node.h

struct node\_t is defined in node.h and function node\_count is defined in node.c:

```

1 #define NODE_DATALEN 52
2
3 typedef struct node_t node_t;
4
5 /**
6  * @brief Defines a node.
7  * @note **Do not remove** any fields from this structure.
8  */
9
10 struct node_t {
11     /** @brief Node ID. */
12     unsigned id;
13
14     /** @brief Node data. */
15     char data[NODE_DATALEN];
16
17     /** @brief The next node. */
18     node_t* next;
19
20     /** @brief The previous node. */
21     node_t* prev;
22 };
23
24 uint32_t node_count(node_t* node) {
25     // YOU ARE NOT SUPPOSED TO MODIFY THIS.
26
27     uint32_t result = 0;
28
29     while (node) {
30         printf("#%u -> ", node->id);
31         result++;
32         node = node->next;
33     }
34
35     printf("Done.\n");
36
37     return result;
38 }

```

### Questions:

7. What accesses in function node\_count cause cache misses? Why?
8. Propose a strategy to decrease the number data cache misses caused by function node\_count. Test your approach on the Gecko board. What is the new number of cache misses? Do not modify the functions or remove fields from struct node\_t.



Note that node\_count is unlikely to access node\_ts that are adjacent in memory.

## Task 2: Optimize item\_t

Consider the following source files:

- ▶ cache/tasks/src/task2.c
- ▶ cache/tasks/src/item.c
- ▶ cache/tasks/include/item.h

struct item\_t is defined in item.h; functions items\_find and item\_init are defined in item.c:

```

1 #define ITEM_DATALEN 32
2
3 typedef struct item_t item_t;
4
5 /**
6  * @brief An item connects an 'id' to 'data'.
7  *
8  */
9 struct item_t {
10     /** @brief Item ID. */
11     unsigned id;
12
13     /** @brief Item data. */
14     char data[ITEM_DATALEN];
15 };
16
17 item_t* items_find(item_t* items, size_t log2n, unsigned id) {
18     // YOU ARE NOT SUPPOSED TO MODIFY THIS.
19     for (size_t i = 0; i < (1 << log2n); ++i) {
20         if (items[i].id == id)
21             return &items[i];
22     }
23     return NULL;
24 }
25
26 void item_init(item_t* item, uint32_t id, const char* data) {
27     // YOU CAN MODIFY THIS.
28     item->id = id;
29
30     if (data != NULL)
31         memcpy(item->data, data, ITEM_DATALEN);
32     else
33         memset(item->data, 0, ITEM_DATALEN);
34 }

```

### Questions:

9. What is the source of data cache misses in function items\_find? Explain.
10. Propose a strategy to decrease the number data cache misses caused by function items\_find. Test your approach on the Gecko board. What is the new number of cache misses? You can modify only function item\_init. You can also modify struct item\_t as long as it holds the ID-data relationship.



Notice that function items\_find accesses items\_t objects that are adjacent in memory.

### Task 3: Optimize Matrix-Vector Multiplication

Consider `cache/tasks/src/task3.c`. In this file, function `multiply` is defined as:

```
2  /**  
3  * @brief Multiplies the matrix with the input vector.  
4  * Writes to the output vector.  
5  */  
6  static void multiply() {  
7      // YOU CAN MODIFY THIS.  
8  
9      for (int j = 0; j < MATRIX_N; ++j) {  
10         for (int i = 0; i < MATRIX_N; ++i) {  
11             out_vector[i] += matrix[i][j] * in_vector[j];  
12         }  
13     }  
14 }
```

#### Questions:

11. Explain the terms row-major and column-major order. Which approach is used in C?
12. What accesses cause cache misses in function `multiply`?
13. Propose a strategy to decrease the number data cache misses caused by function `multiply`.

## Task 4: Bouncing Ball Example

Consider `cache/tasks/src/task4.c`. In this file, functions `bouncing_ball` and `init_dcachel` are defined as:

```

2 // old base address
3 #define LEDS_OLD_BASE 0x50000800ull
4
5 // new base address
6 #define LEDS_NEW_BASE 0x08000800ull
7
8 // pixel-based control of LEDs
9 #define LEDS_LEDS_OFFSET 0x400ull
10
11 // the base address is configurable.
12 // defines the offset of the base address register.
13 #define LEDS_BASEADDR_OFFSET 0x7FCull
14
15 void init_dcachel() {
16     // YOU CAN MODIFY THIS.
17     dcachel_enable(0);
18     dcachel_write_cfg(CACHE_FOUR_WAY | CACHE_SIZE_4K | CACHE_REPLACE_LRU |
19     CACHE_WRITE_BACK);
20     dcachel_enable(1);
21 }
22
23 void bouncing_ball() {
24     // YOU CAN MODIFY THIS.
25     int xdir, ydir, xpos, ypos, index;
26     volatile unsigned int* leds = (unsigned int*)(LEDS_NEW_BASE + LEDS_LEDS_OFFSET);
27
28     xdir = ydir = 1;
29     xpos = ypos = 5;
30     while (1) {
31         index = ypos * 12 + xpos;
32         leds[index] = 0;
33         if (ypos == 8)
34             ydir = -1;
35         if (ypos == 0)
36             ydir = 1;
37         if (xpos == 11)
38             xdir = -1;
39         if (xpos == 0)
40             xdir = 1;
41         ypos += ydir;
42         xpos += xdir;
43         index = ypos * 12 + xpos;
44         leds[index] = swap_u32(2);
45         for (volatile long i = 0; i < 100000; i++)
46             ;
47     }
48 }

```

### Questions:

14. Why does not the bouncing ball work as expected? You should be observing no change in the LEDs. The expected behavior is a single pixel moving and reflects as it bounces at the edges.
15. Propose two approaches to fix the behavior.