



High Performance Embedded Systems EEE4120F



FINAL EXAM 10 June 2024

Venue: SNAPE TS2B Time starting: 08h00 3 hours

Examination Prepared by: Simon Winberg Last Modified: 08-Jun-2024

NB! NB! FIRST ENTER YOUR NAME AND STUDENT NUMBER HERE:

Name (first and last name)	Student Number

REGULATIONS

This is a closed-book exam. Scan through the questions quickly before starting, so that you can plan your strategy for answering the questions. If you are caught cheating, you will be referred to University Court for expulsion procedures. Answer on the answer sheets provided.

Make sure to **put your student number on any separate answer sheets!** Also **put your name and student number** and course code **EEE4120F** and title **Final Exam** in your UCT exam answer book (this answer book is mainly for rough work, you need to return this question paper containing your answers).

FOLLOW THE EXAMINATION PROCEDURES ON THE PROCEDURES PAGE ON YOUR DESK

DO NOT TURN OVER UNTIL YOU ARE TOLD TO

**NB!!! ANSWER QUESTIONS ON THIS PAPER, USE EXAM BOOKLET
FOR ADDITIONAL SPACE IF NEEDED**

Exam Structure

Marked out of 100 marks / 180 minutes. Time per mark = 1min 48sec

<u>Section 1</u>	<u>Section 2</u>	<u>Section 3</u>	<u>Appendices</u>
Short Answers (2 x 10-mark questions) [20 marks] pg 2	Multiple choice (10 x 5-mark questions) [50 marks] pg 6	Long Answers (2x question design/code) [30 marks] pg 9	A: Verilog cheatsheet (can detach last page) pg 13

(hint: consider taking a quick look at Q3.1 and Q3.2 to help you plan your answering)

RULES

- NB!**
- You must write your name and student number on each answer script book you hand in.
 - Write the question numbers attempted on the cover of each answer book that you use and hand in.
 - Make sure that you cross out material you do not want marked. Your first attempt at any question will be marked if two answers are found.
 - Answer all questions, and note that the time for each question relates to the marks allocated to it.

Section 1: Short Answers [20 marks]

Question 1.1 [10 marks]

Large clusters can include possibly hundreds of networked computers. It is now becoming more common to find clusters with heterogeneous nodes, for instance each node of the cluster may be a computer that also has FPGA-based accelerators, if not a graphics processing unit (GPU) as well. Programming such systems may use a combination of Pthreads and the Message Passing Interface (MPI), including specialized compilers and tools such as OpenCL for the GPU and Xilinx Vivado for the FPGA. Answer the following:

Q1.1(a) Why might it be useful to have a cluster with nodes comprising heterogeneous architectures? Why could this provide an overall performance better than nodes of homogeneous architectures? [2 marks]

Q1.1(b) What are some of the difficulties and expenses concerned in developing applications for heterogeneous architectures? Explain using comparisons to traditional clusters of SMP processor architectures, and also motivate reasons why SMP-only clusters are sometimes favored rather than heterogeneous architectures. [4 marks]

Q1.1(c) Pthreads provides a particular, almost intuitive, approach to developing a parallel program for deployment on a computer where the processors are largely all the same. Consider the following situation:

- We have global array variables: `int x[N], y[N]`; and global `int sum=0`, where `N` is a constant positive value under 1024.
- We spawn `N` threads where thread IDs, a variable `i` in the thread function, ranges from 0 to `N-1`
- Within each thread instance, i.e. in thread ID `i`, the operation `sum+= x[i]*y[i]` is performed

But there are problems with the above program design, which likely causes the wrong sum value to be calculated when run with more than one thread due to a race condition on the shared sum variable. Briefly explain what this problem is and how it can lead to incorrect results in this case [2 marks].

In addition, explain why, when run using multiple threads, this code actually ends up slower than using a single-threaded for-loop to sequentially add up the $x[i]*y[i]$ terms. Consider factors such as thread management overhead, synchronization overhead, and memory contention in your explanation [2 marks].

[illegible]

[Q1.1c total = 4 marks]

Question 1.2 [10 marks]

Answer the following questions and sub-questions concerning computation approaches...

- (a) Reconfigurable Computing (RC) platforms are constructed using FPGAs among other processing components. They often also include a PLA or CPLD in addition to an FPGA on the platform.

Explain the following...

- i. What is the main difference between an FPGA and a PLA? [1 marks]
- ii. Explain why a PLA (or a CPLD) in a reconfigurable computing system is often needed in the configuration architecture / glue logic whereas FPGA(s) are more commonly used for the actual computation (e.g., signal processing) operations. [2 marks]

- (b) It is often stated that an FPGA gives around a tenth of the performance (if that) of an ASIC. But what is an ASIC? [1 mark] Explain why an FPGA is generally not going to be as fast as an ASIC implementation; you can consider a case of comparing a softcore RISC-V processor (implemented in HDL running on a FPGA) to a RISC-V processor implemented in silicon that are both running the same processing tasks. [3 marks] [total = 4 marks]

(question 1.2 continues on next page)

- (c) Contrast the benefit of using parallel code over sequential code for an embedded microprocessor-based or microcontroller-based solution. [3 marks] (marking: 1 mark for explaining a significant benefit, and 2 marks for a clear motivation for why you consider it a significant benefit).

(exam continues next page)

Section 2: Multiple Choice [50 marks]

NOTE: Choose only one option i.e. **circle the letter A., B., C. or D.** in the column left of the answer option to select your answer choice for Q2.1 to Q2.9. Questions are 5 marks each.

Q2.1 An embedded system development company often encounters two general types of costs related to their product manufacture. Which one of these is an example of “non-recurring costs”? (choose one option):

- A. Battery purchase, as this has become a dominant factor in the production of new systems.
- B. Maintenance costs.
- C. Postage on returns.
- D. Research and development costs.

[5 marks]

Q2.2 You might remember some discussion about the term baud rate. Sentence below most accurately describes what baud rate is? Consider each sentence starts with The baud rate is ...

- A. ... the number of bits transferred in the communication channel each second.
- B. ... the speed that a packet of data is encoded for sending over the communication channel.
- C. ... the symbol rate used for sending information over the communication channel.
- D. ... the time period after which particles of the conveying medium decay and lose information.

[5 marks]

Q2.3 In the first term, the paper “The Landscape of Parallel Computing Research: A view from Berkeley” was presented. It was highlighted that for processor cores “small is beautiful”. Which of these statements accurately reflects the article’s advice on multi-core processor architecture design?

- A. A small number of big cores should be used as they are much better at doing many small tasks.
- B. Many small cores tend to give better performance per unit area than a few big cores taking an equivalent area.
- C. A small processor gives better performance, so processor design should always be kept simple.
- D. A big cluster of small cores will almost always run faster, for high granularity problems, than a small cluster of big cores.

[5 marks]

Q2.4 Which one of the following memory partitioning models is used in the program shown below? (Choose just one letter for your choice). Assume an image gets successfully loaded into the image object.

- A. Partitioned
- B. Interleaved
- C. Haphazard
- D. Interlaced

[5 marks]

(... code on next page ...)

```

#include <thread1> // thread library
#include <Image>    // image library

using namespace std;

JPG_Image image; // image object of image.size_x by image.size_y pixels
const int nth = 4;

void threadfn ( int arg ) {
    int y = arg; // arg indicates starting y coord
    while (y<image.size_y) {
        for (int x=0; x<(image.size_x-1)/2; x++) {
            Pixel t = image.pixel[x][y];
            image.pixel[x][y] = image.pixel[image.size_x-x-1][y];
            image.pixel[image.size_x-x-1][y] = t;
        }
        y = y + nth;
    }
} // end function

int main () {
    int i;
    thread threads[nth]; // array of threads
    image.load("happy.jpg"); // load an image from file
    for (i = 0; i < nth; ++i)
        threads[i] = create_thread(threadfn,i);
    for (i = 0; i < nth; ++i) join_thread(threads[i]);
    showimage(image);
    return 0; // program completed
}

```

Q2.5 The concept of load balancing is often discussed in terms of parallel computing. But what is load balancing in relation to parallel computing? Select the best description of load balancing below:

- A. It is ensuring precise synchronization between tasks to ensure task order is maintained.
- B. It is distributing work to minimize idle time.
- C. It is maintaining fair access to shared resources to avoid deadlock or excessive wait times.
- D. It is regulating equal processor allocation between threads.

[5 marks]

Q2.6 Consider that it has been decided to implement a state machine using one-hot encoding for the state index. But what does this mean? Select the best option below explains this encoding method.

- A. One-hot encoding is potentially wasteful on bits, but it can enhance robustness (e.g. tolerance to radiation) whereby valid state indexes have only one bit set and all other bits clear.
- B. It is an encoding scheme by which the state index will not have more than one bit change when changing from one state to another state.
- C. One-hot encoding is a scheme by which the sum of all bits in a state adds up to a prime number (if this is not the case, then the state is invalid).
- D. It is a state machine design by which only one state does any processing or result generation.

[5 marks]

Q2.7 Which one of the following algorithms is the most course-grained when parallel processing a NxN pixels 8-bit grayscale image that is sitting in shared memory accessible by multiple processors?

- A. Perform a horizontal image flip (i.e. for $x=1:N$ for $y=1:N$ tmp = img[x,y]; img[x,y]=img[N-x+1,y]; img[N-x+1,y]=tmp; end; end;)
- B. A mean image filter that, for each pixel, calculates the average grey scale value from that pixel from the surrounding 5x5 pixels.
- C. Find the location of the brightest pixel on each row and see if that pixel is also the brightest pixel on the corresponding column.
- D. Reduce the value of each pixel in an image if its value is above 10.

[5 marks]

Q2.8 The acronym DDR-SDRAM might have been used to describe a type of memory or memory module for an embedded system? What does DDR-SDRAM refer to if it is actually a valid term.

- A. This is not a valid term.
- B. Double-Data Rate Synchronous Dynamic Random Access Memory.
- C. Data-Debounced Regulated Standard Dynamic Random Access Memory.
- D. Dual-Data Regulated Symmetric Dynamic Random Access Memory

[5 marks]

Q2.9 Which one of the following statements is correct? (choose only one answer).

- A. A softcore processor cannot be placed on an FPGA, as the propagation delays on FPGAs are excessively lengthy and will lead to instructions not being able to complete on time.
- B. Daisy chaining is a bus architecture for capturing time of data transfers.
- C. Fine-grained (compared to course-grained) parallel tasks require a great deal of intercommunication relative to the amount of computation they do.
- D. It is necessary to use the same number of blocking assignments as nonblocking assignments within the same Verilog 'always block'.

[5 marks]

Q2.10 Select True (T) or False (F) for the following statements to indicate if the statement is true or false?

#	Statement	ANSWER T/F
i	We used OpenCL in the pracs; the acronym OpenCL stands for "Open Compile Language"	
ii	In Verilog, non-blocking assignments ('<=') are for sequential logic, ensuring all assignments within the same 'always@' block occur one after the other.	
iii	FPGAs incorporate a routing network for connecting up PLBs.	
iv	Access to network resources (e.g. database on different machine) is considered a software bottleneck in many modern high-performance systems.	
v	The term 'embarrassingly parallel' is used to refer to a very course-grained processing operation that is easy to make parallel.	

[5 marks]

Section 3: Long Answers [30 marks]

This section involves one rather lengthy answer and one not-so-lengthy answer

Huge Calculator With Tiny Battery

(an assessment of your thinking of parallel systems... without you needing to do any programming)

Imagine that you are an engineer working for a development company called Amalgamated Constructors of Meaningless Equipment (or just ACME for short). The company produces a variety of widgets – for which there is high demand, mainly as corporate gifts. These widgets include: huge desk calculator (with tiny non-standard built-in batteries designed to go flat after a year), colorful USB hubs (version 1.0 USB naturally), pens with chrome parts and dried out ink, and laptop carrier bags that are just a few millimeters too small for a full-size laptop. Business has been booming! It is imperative that the production lines are enhanced. You, with your knowledge of parallel programming and electronics (since they could not find a systems engineer), have been tasked to streamline part of the Huge Calculator (HC) with tiny non-standard battery production line.



You need to think about the Verilog coding and testing of it and respond to the questions that follow.

Question 3.1 Verilog Module Coding [20 marks]

The ACME Company is planning a prototype version of their Huge Calculator. This simplified form is not going to implement any of the fancy functions such as square-root or divide. It will just do basic operations using 32-bit integers (i.e. no decimals). How the system works is that one punches in a number (e.g. 10) and then presses one of the math function keys (only ‘+’ ‘-’ ‘x’ and ‘=’ works). Each time you press a function key it displays the result of the previous operation. A reset is needed, that resets things, i.e. sets regs to 0.

Example scenarios of operation and what is displayed are as follows:

<i>Keypresses</i>	<i>Display</i>
1 0	1 0
+	1 0
2	2
=	1 2

Thus, when a math key is pressed the last entered value (LEV register) and the accumulator (ACC register) are applied to the math function and the result is saved to the ACC and displayed on the screen, by copying the value of register ACC to the display register (DSP). (i.e. assume the DSP register is already connected to the LCD display so that the LCD will display whatever is in the DSP register).

Note that you do not need to worry about capturing the number input, there is a separate module (called *levin*, which you need not worry about connecting up) that senses number key presses and updates the LEV register (so, if you press reset, LEV is set to 0; then pressing ‘1’ LEV is set to 1, then pressing ‘2’ causes LEV to be multiplied by 10 and have 2 added, resulting in LEV=12. And so on as per normal use).

The math operations are involved by pressing keys ‘+’, ‘-’, ‘x’, and ‘=’ which respectively results in signal lines add, sub, mul or eq being raised high for a few microseconds and returning to 0V. Assume hardware debounce is already dealt with so you don’t need to write code to accommodate this.

The function code below, which you need to add to, gives you a starting point; please complete it to implement the above functionality given the operation of the inputs and outputs as explained above.

(you can assume this is the starting point of your solution, add code in the blank lined spaces below. You don't have to rewrite this part, if you choose to answer in your exam answer booklet)

```

module hugecalc (
    lev, acc, dsp, add, sub, mul, eq    /* can add below */

);

/* FPGA prototyped version of the Huge Calculator */
input  reset;      /* reset the calculator */
input  lev[31:0]; /* last entered value */
inout  acc[31:0]; /* accumulator          */
output dsp[31:0]; /* display register    */
input  add,sub,mul,eq; /* the math function buttons */

/* add your code here: */

```

[illegible]

[illegible][illegible]

(more answer space available on the back on the next page, Appendix A, if you choose not to detach it)

END OF EXAMINATION

Appendix A: Verilog Cheat sheet

Numbers and constants

Example: 4-bit constant 10 in binary, hex and in decimal: 4'b1010 == 4'ha -- 4'd10

(numbers are unsigned by default)

Concatenation of bits using {}

4'b1011 == {2'b10, 2'b11}

Constants are declared using parameter:

parameter myparam = 51

Operators

Arithmetic: and (+), subtract (-), multiply (*), divide (/) and modulus (%) all provided.

Shift: left (<<), shift right (>>)

Relational ops: equal (==), not-equal (!=), less-than (<), less-than or equal (<=), greater-than (>), greater-than or equal (>=).

Bitwise ops: and (&), or (|), xor (^), not (~)

Logical operators: and (&&) or (||) not (!) note that these work as in C, e.g. (2 && 1) == 1

Bit reduction operators: [n] n=bit to extract

Conditional operator: ? to multiplex result

Example: (a==1)? funcif1 : funcif0

The above is equivalent to:

```
((a==1) && funcif1)
```

```
|| ((a!=1) && funcif0)
```

Registers and wires

Declaring a 4 bit wire with index starting at 0:

```
wire [3:0] w;
```

Declaring an 8 bit register:

```
reg [7:0] r;
```

Declaring a 32 element memory 8 bits wide:

```
reg [7:0] mem [0:31]
```

Bit extract example:

r[5:2] returns 4 bits between pos 2 to 5 inclusive

Assignment

Assignment to wires uses the assign primitive outside an always block, e.g.:

```
assign mywire = a & b
```

Registers are assigned to inside an always block which specifies where the clock comes from, e.g.:

```
always@(posedge myclock)
```

```
    cnt = cnt + 1;
```

Blocking vs. unblocking assignment <= vs. =

The <= assignment operator is non-blocking (i.e. if use in an always@(posedge) it will be performed on every positive edge. If you have many non-blocking assignments they will all updated in parallel. The <= operator must be used inside an always block – you can't use it in an assign statement.

The blocking assignment operator = can be used in either an assign block or an always block. But it causes assignments to be performed in sequential order. This tends to result in slower circuits, so avoid using it (especially for synthesized circuits) unless you have to.

Case and if statements

Case and if statements are used inside an always block to conditionally update state. e.g.:

```
always @(posedge clock)
    if (add1 && add2) r <= r+3;
    else if (add2) r <= r+2;
    else if (add1) r <= r+1;
```

Note that we don't need to specify what happens when add1 and add2 are both false since the default behavior is that r will not be updated. Equivalent function using a case statement:

```
always @(posedge clock)
    case({add2,add1})
        2'b11 : r <= r+3;
        2'b10 : r <= r+2;
        2'b01 : r <= r+1;
        default: r <= r;
    endcase
```

Module declarations

Modules pass inputs, outputs as wires by default.

```
module ModName (
    output reg [3:0] result, // register output
    input [1:0] bitsin, input clk, inout bidirectnl );
    ... code ...
endmodule
```

Verilog Simulation / ISIM commands

```
$display ("a string to display");
$monitor ("like printf. Vals: %d %b", decv,bitv);
#100 // wait 100ns or simulation moments
$finish // end simulation
```

(additional answer space, remember to indicate question numbers)