Computer Design Exercises

This set of suggested exercises represents the minimum I would expect completed for supervisions. Supervisors are encouraged to set additional material. Suggestions of additional questions will be gratefully received by the lecturer, as will corrections.

Lectures 1 to 4 – ECAD

- Teach yourself Verilog using the Interactive Verilog Compiler (IVC). IVC teaches you the Verilog you need to complete the ECAD assessed exercises as painlessly as possible. A number of self-tests check your progress through learning the basics of the language. The tutorial explains what you need to do to write Verilog programs. Link to IVC: http://www.cl.cam.ac.uk/Teaching/current/ECAD+Arch/

Lecture 5 — Historical Computer Architecture

- What was Leo and what was it used for?
- Where and when was Whirlwind developed?
- What did TRADIC first demonstrate?
- If memory capacity continues to double every 18 months, how much main memory will a typical desktop PC have in 20 years time?

Lecture 6 — Early instruction set architecture

- Read and discuss chapters 1 & 2 from "Computer Architecture — A Quantitative Approach" (2nd or 3rd editions, not the most recent), particularly the "fallacies and pitfalls" sections.

Lecture 7 — Build your first computer

- Download the Baby examples on the ECAD+Arch web page and try them out in simulation and on the DE2 board.
- Have a go at writing your own Baby program and simulating it.

Lecture 8 — RISC Processor Design & the MIPS Instruction Set

- Given that the MIPS processor has a branch delay slot, what will the following contrived piece of code do?

```
foo:  slti   $t0, $t0, 5
      beq    $t0, $zero, foo
      addi   $t0, $t0, -1
```

- Write a loop that copies a region of memory from the address in $a0 to the address in $a1 for the number of words specified in $a2. You may assume that the regions of memory are none overlapping.
- How might you improve the performance of your code in (b) by copying more than one word on each iteration of the loop?

Lecture 9 — MIPS Tools & Examples

- Answer past exam. question 2, paper 5, 2004: `http://www.cl.cam.ac.uk/tripos/y2004p5q2.pdf`

Lecture 10 — CISC machines and the Intel IA32 Instruction Set

- How did AMD extend Intel's 32-bit instruction set (IA32) to the 64-bit version AMD64?
- How does AMD64 differ from IA64?

Lecture 11 — Java Virtual Machine

- Write an iterative version of Fibonacci in Java and figure out what the disassembled code means. Run through the code for a fib(4).

Lecture 12 — Memory Hierarchy

- Read and discuss "Memory Hierarchy Design" from "Computer Architecture — A Quantitative Approach" (2nd or 3rd editions, not the most recent), particularly the "fallacies and pitfalls" sections.

Lecture 13 — Simulating a MIPS Processor

- Add the XOR instruction to the simulation and write some MIPS assembler to create a (rather poor!) pseudo random bit generator using a linear feedback shift register. Verilog simulation code describing this algorithm:

```
module random(
        output reg clk,
        output random);

   reg [3:0]  s=4'h1; // shift register
   initial begin
          clk = 0;
          #200 $stop;
   end

   always #5 clk = !clk;

   always @(posedge clk)
         s <= {s[0]^s[1],s[3:1]};

   assign random=s[0];

   always @(negedge clk)
         $display("at time %3d random=%1d",$time,random);

endmodule // random
```

Lecture 14 — Simplified MIPS in SystemVerilog

- Repeat the exercise for lecture 13 but for the SystemVerilog version of the simplified MIPS processor

Lecture 15 — Pipelining

- Read about pipelining and resolution of hazards, e.g. in Chapter 7 of Harris & Harris, *Digital Design and Computer Architecture*, 2007.
- Attempt past exam. question 2007 P6 Q2

Lecture 16 — Communication on and off chip

- What is the difference between serial and parallel communication?
- What is the difference between latency and bandwidth?
- What is the difference between a bus and a switched communication network?
- Why is it difficult to communicate data in parallel at GHz frequencies?

Lecture 17 — Manycore

- What is the difference between instruction-level parallelism and thread-level parallelism?
- How does SMT exploit both instruction-level parallelism and thread-level parallelism?
- What is companion scheduling?
- What is the difference between manycore and CMP?

Lecture 18 — Data-flow

- What is the difference between data-flow and control-flow?
- Do RISC processors execute instructions in a data-flow or control-flow manner?
- When can modern SMT processors exhibit any data driven behaviour?