

## Part II CST: SoC D/M: Quick exercises (examples sheet) (preliminary draft copy 1st May 2013).

This sheet contains short exercises for quick revision. Please also look at past exam questions and/or try some of the longer exercises from the second sheet.

Unless otherwise stated, both here and in exam questions, you can answer using any mixture of RTL, SystemC or block or circuit diagrams. Syntax details are not important. Similarly, precise accuracy with C, C++ or any particular assembly language is not important.

### SG 1 – Register Transfer Language (RTL)

SG1a.1. Give an RTL structural netlist for an RS latch.

SG1a.2. Give a fragment of RTL that uses non-blocking assignments (in Verilog the `<=` operator) to swap a pair of bitwise registers on a clock edge when an externally-provided enable signal holds. Sketch the resultant circuit (or gate-level, structural netlist) that would arise when your fragment is compiled to gates.

SG1a.3. Give a fragment of RTL that uses blocking assignments (in Verilog the `=` operator) and give the resultant circuit (or structural netlist) showing clearly the difference that would arise if non-blocking assignments were used instead.

SG1a.4. The following fragment of RTL produces a repeating pattern on a one-bit output. Assuming it is accepted by some gate-generating synthesis tool, give the resultant circuit (or structural netlist) that it might sensibly generate.

```
always begin
    @(posedge clk) a<= 1;
    @(posedge clk) a<= 0;
    @(posedge clk) a<= 0;
end
```

SG1a.5. Consider the following fragment in a fictional BlueSPEC-like language. This is to be synthesised into a standard Verilog (or VHDL) RTL for further synthesis to gates using standard RTL tools. Assume that the `write` and `read` methods each add the following wires to the signature of the generated component: a request input and a ready output. Give the resultant circuit (or structural netlist) that the tool might sensibly generate. The handshake convention is that method call takes place on a positive clock edge where both request and ready hold.

```
module InvertingRegister( ... )

    reg [7:0] contents; // Basic RTL register (unguarded).

    method write(d) { contents <= d; }

    method read() { return ~d; }
```

Suppose there is a side condition that read and write cannot occur in the same clock cycle: modify your answer to reflect this. *Hint: the difference would be in the generation of one of the ready outputs.*

### SG 1 – Simulation:

- SG1b.1. Tabulate the events inserted in an event list when simulating one clock cycle of your net-level hardware design from question SG1a.2.. Assume each gate has a delay of 100 picoseconds.
- SG1b.2. Modify your net-level hardware design from SG1a.2. by adding an inverter so that one of the bits is swapped in polarity when moving between one register to the other. What extra events would arise in the net-level simulation. What has happened to the maximum clock frequency of the design?

### SG 1c – Hazards:

- SG1c.1. Consider a FIFO component that accepts 16-bit writes and provides 8-bit reads. Internally it is implemented with a 128-byte SRAM with organisation 128x8. A revised implementation uses an SRAM with organisation 64x16. Which implementation offers better throughput and which (if any) suffers from structural a hazard ?
- SG1c.2. The writeback pipeline stage in a simple RISC processor was once described as ‘a dummy stage whose only purpose is to avoid a structural hazard’. Is this a fair comment and why? How does data cache cycle time affect your answer?

### SG1 – Folding, Retiming & Recoding:

- SG1d.1. The critical path in a clock domain consists of three AND2 gates each with delay 35 ps and starts and ends at flip-flops with the following parameters: set-up time 4 ps, hold time 2ps, clock to Q time 15 ps. What is the maximum clock frequency ?
- SG1d.2. Time and space can be exchanged by design refactoring. Compare the algorithm on page 20 of the notes to a naive repeated addition approach to multiplication. Which has more time and which uses more space (if either)?
- SG1d.3. Booth’s radix-4 algorithm for long multiplication is embodied in the following ML function:

```
fun booth(x, y, c, carry) = // Inputs are x and y. Other args initially clear.
  if(x=0 andalso carry=0) then c else
  let val x' = x div 4
      val y' = y * 4
      val n = (x mod 4) + carry
      val (carry', c') = case (n) of
        (0) => (0, c)
      | (1) => (0, c+y)
      | (2) => (0, c+2*y)
      | (3) => (1, c-y)
      | (4) => (1, c)
      in booth(x', y', c', carry')
  end
```


How many adders does it require in its natural hardware implementation (no special fold/unfold applied)? Sketch the hardware datapath and list the inputs and outputs to the controlling FSM. How many clock cycles does it use ? *Perhaps this should be on the long exercise sheet.*

## SG 1 – Protocol and Interface:

- SG1e.1. What is the difference between a synchronous and an asynchronous interface?
- SG1e.2. What is the difference in throughput between the two-phase and the four-phase handshake ?
- SG1e.3. Why might an interface be ‘always ready’ ?
- SG1e.4. How can data from two receiving interfaces be transmitted over a single outgoing interface ? Sketch RTL or a circuit that implements this for a suitable handshaking protocol. Perhaps add an extra feature (such as an extra wire in a parallel bus) where the data is somehow tagged with the originating interface as it conveyed to the output. (NB: This is not a short/quick exercise for those not used to hardware design.)

*Hint:* An arbiter, will be needed. For a synchronous design it will be like

<http://www.cl.cam.ac.uk/teaching/1011/P35/arbiter/zhp711199d4d.html> and for asynchronous design it might be based on a violated RS latch as shown: It might be helpful to also think about a companion design that implements the reverse direction (a demultiplexor).



## SG 2 – SystemC Components

- SG2.1. Why does SystemC implement the (non-blocking) signal paradigm where the value written is not immediately readable?
- SG2.2. Why does SystemC distinguish between threads and methods and what is the difference ?
- SG2.3. Sketch SystemC code for design of question SG1e.4.. (If you answered that question using RTL perhaps just list the editor commands needed to make the change!)

## SG 3 – Basic SoC Components

- SG3a.1. What is meant by the address space and the address map of a computer ?
- SG3a.2. Explain why I/O devices and other items in the memory map might appear more than once ?
- SG3a.3. What problem is encountered with memory-mapped I/O devices when the cache structure of a computer is misconfigured?
- SG3a.4. Why is it common practice for the interrupt signal to be deasserted by an I/O device when the status register is read?
- SG3a.5. Why is it necessary to be able to disable transmit interrupts ?
- SG3a.6. Why must the memory map of a computer be designed taking into account the reset and interrupt vectors hardwired into a processor core ?
- SG3a.7. A single LED is connected to a single GPIO pin of microcontroller. Sketch C or assembler code to make the LED flash at 1 Hertz.
- SG3a.8. Why is DMA generally used to service an Ethernet I/O device whereas it is not for a UART connected to an RS232 serial line ?
- SG3a.9. Why might the processor on a laptop computer adjust its clock own frequency ?
- SG3a.10. In the big/little approach to processor design, a simple core is used when there is not much processing to do. How might the switch between big and little cores be achieved ? *Note: there are many possible answers. May not be lectured.*
- SG3a.11. Why do devices have multiple clock domains? What precaution is needed when crossing clock domains? How does this influence the protocol or framing for messages crossing between domains ?

### SG 3 Bus questions:

- SG3b.1. A 64 bit bus is clocked at 200 MHz. A target always takes 4 clock cycles to respond to acknowledge a write. What is the write throughput for a transaction size of one word? How could write posting or larger transactions improve performance?
- SG3b.2. If there are two general-purpose ARM cores on a SoC, should each have its own bus and should there be two external DRAM banks?
- SG3b.3. What advantage does the BVCII bus have that enables it to tolerate pipeline delays in the bus structure.
- SG3b.4. In terms of throughput and number of elements, compare a full cross-bar made of 2x2 elements to interconnect N initiators with N targets with a single ring network. What assumptions have you made about traffic patterns and are they realistic?
- SG3b.5. Why might an ISS run faster than the real hardware ?
- SG3b.6. What difference does it make whether an ISS includes cache models or not ?

SG3b.7. Why is DRAM so much slower than SRAM in terms of latency and throughput ? Why then do we use it?

## SG 4 – ESL and TLM

### ESL – Electronic System Level Modelling.

- SG4.1. Why not wait for the silicon to be made before developing application software ?
- SG4.2. Describe an example where ‘functional modelling’ of an application to be implemented on a SoC does not reflect the final layout of data in memory (i.e. not memory-accurate modelling).
- SG4.3. ‘Cycle-accurate modelling is faster than event-level modelling’. If true, what sort of event is being considered ?
- SG4.4. To what extent can embedded software and firmware be tested by compiling it with the workstation compiler (e.g. to x86 code) compared with compiling it for the target architecture (e.g. ARM code) ?
- SG4.5. What advantage can be made of the fact that embedded firmware is often written in the same language as behavioural models of components (i.e. C++) ?
- SG4.6. How can RTL or net-level circuitry for an IP block be imported into an ESL simulation and how might the result differ from using a high-level model of that block?

### SG 4 – Transactional Level Modelling (TLM):

- SG4b.1. How can TLM modelling avoid the need to model handshake wires?
- SG4b.2. Does device driver software have to be changed when using an ESL model ?
- SG4b.3. Transactor code for four-phase handshake is included in the lecture notes. Sketch equivalent code for two-phase handshake, where data is transferred on both edges of the request signal. Say whether your implementation is synchronous or asynchronous.
- SG4b.4. What is the advantage of carrying the delay parameter in the transaction signature and what would be the more-simplistic alternative ?

## SG 5 – ABD - Assertion-Based Design

SG5.1. Classify the following as safety or liveness assertions:

- (a) ‘Either the inner or outer airlock doors is always closed’
- (b) ‘The sun never rises in the north.’
- (c) ‘When the big hand has just passed twelve the little hand is always on a number.’
- (d) ‘However far you drive along the M25 there will always be more road ahead.’
- (e) ‘Whichever way you branch you will always come to dead end.’

SG5.2. Give a fragment of RTL that will always eventually do something that no practical simulation run would ever reveal.

SG5.3. Write an assertion for a formal tool to check concerning your answer to SG5.2..

SG5.4. ‘*Sequential logic equivalence implies combinational logic equivalence*’. Discuss.

SG5.5. A simplistic generator algorithm, for use in directed/constrained random stimulus generation, will run very slowly. Give an example and explain why.

SG5.6. Why are regular expressions a rather limited language for expressing assertions over a hardware subsystem?

## SG 6 – SoC Architecture Exploration and Tools

### SG 6 – Engineering and Associated Tools:

SG6a.1. Design a small circuit containing just flip-flops and 2-input multiplexors where a static timing analyser would produce the wrong result since the longest path never carries an event from end to end. (This pattern often arises when an IP block has a test mode controlled by a test input, so perhaps consider an external input called test that feeds your circuit.)

SG6a.2. Design a scan path flip-flop (i.e. list its terminals and sketch its internal logic). This is a normal flip-flop that can become part of a shift register during testing. *May not have been lectured.*

SG6a.3. Give a list of test vectors that will reveal all stuck-at faults in a direct hardware translation of the following RTL fragment: *May not have been lectured.*

```
assign p = a && b;  
assign q = p || r;
```

SG6a.4. A BIST unit for a 32Kx4 static RAM operates by writing a pseudo-random pattern to the device and then checks whether 1024 other pseudo-random writes corrupt any location. If there are five such RAMs on a SoC, each with their own BIST unit, and the wafer probe tester supplies a clock of 32 MHz (lower than normal operation), give a lower bound on the time to wafer probe each chip.

## **SG 6 – Architectural Design Exploration:**

- SG6b.1. What is the difference between a co-processor and a peripheral ?
- SG6b.2. What is the main difference between a standard cell design and a gate array? Which can be mask programmed and which can be field programmed ?
- SG6b.3. What are the main two benefits of making transistors smaller in VLSI and what disadvantage is now becoming predominant below 45nm feature size ?
- SG6b.4. A ‘respin’ occurs when a chip has a design fault that means new masks are needed before it can go into production. Estimate (or Google for) the cost of a respin to the nearest million dollars!
- SG6b.5. The company V-Tech makes children’s toys which normally consist of a battery, a loudspeaker, one or more motors, LEDs, switches, a number of pretty plastic moldings, the odd bit of artificial fur and a small circuit board. The company is very successful, with production runs of a given toy often exceeding 300,000 units. What technology should they use on their circuit board?
- SG6b.6. Give two major reasons why chips should increasingly be made with the intention that most/much of their logic is switched off or not in use much/most of the time.

## **SG 7 – Silicon Power and Technology**

- SG7.1. A logic gate with intrinsic delay of 190 picoseconds feeds four other logic gates using a total net length of 0.05 mm on metal layer 2 of a chip. The capacitance per unit length on layer 2 is 0.3 pF per mm. The logic gates being driven each have an input capacitance of  $4 \times 10^{-15}$  F. The gate derates as 10 picoseconds per load unit where a load unit is also  $4 \times 10^{-15}$ . What gate delay should be modelled?
- SG7.2. The clock network within a subsystem has total length 3 mm. Its frequency is 200 MHz. The supply voltage is 1.2 Volts. The wiring capacitance is 0.3 pF/mm. How much power is dissipated by this clock net when in operation?
- SG7.3. How can dynamic voltage scaling lead to a cubic relationship between clock frequency and power dissipation?
- SG7.4. Why might the power consumption be different for different programs run on the same core even if they are all highly CPU-intensive ?

SG7.5. Why can dynamic clock gating respond more quickly than dynamic power gating ?

SG7.6. There are numerous potential mechanisms for frequency, power and clock gating control, including:

- (a) a dedicated input pad (or pads) on the chip,
- (b) additional logic inserted by the Verilog compiler as it generates a netlist from RTL,
- (c) programmed I/O by a core writing to control registers.

Suggest a use case, such as an IP block, subsystem, application or device that might use each of these three techniques.