

Part II CST: SoC D/M: Quick exercises (examples sheet) 2015 (rev a).

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 main 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.

This sheet will be extended to include further topics as lectures progress.

SP1Q: Basic SoC Components

- SPQ1.1. What is meant by the address space and the address map of a computer ?
- SPQ1.2. Explain why I/O devices and other items in the memory map might appear more than once ?
- SPQ1.3. What problem is encountered with memory-mapped I/O devices when the cache structure of a computer is misconfigured?
- SPQ1.4. Why is it common practice for the interrupt signal to be deasserted by an I/O device when the status register is read?
- SPQ1.5. Why is it necessary to be able to disable transmit interrupts ?
- SPQ1.6. Why must the memory map of a computer be designed taking into account the reset and interrupt vectors hardwired into a processor core ?
- SPQ1.7. 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.
- SPQ1.8. 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 ?
- SPQ1.9. 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 ?
- SPQ1.10. 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?
- SPQ1.11. 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?

- SPQ1.12. What advantage does the BPCI bus have that enables it to tolerate pipeline delays in the bus structure?
- SPQ1.13. In terms of throughput and number of elements, compare a full cross-bar made of 2×2 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?
- SPQ1.14. Why is DRAM so much slower than SRAM in terms of latency and throughput ? Why then do we use it?

SP3Q: Design Partition and Architectural Exploration:

- SPQ2.1. What is the difference between a co-processor and a peripheral ?
- SPQ2.2. What is the advantage of full-custom VLSI design and why is semi-custom generally used?
- SPQ2.3. 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 ?
- SPQ2.4. What, in theoretical terms, might be the minimum content for a FPGA configurable logic block ?
- SPQ2.5. 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?
- SPQ2.6. 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.