SoC D/M Exercise Sheet(s), 2016/2017

This sheet contains exercises of various lengths, but the very short ones are in a separate PDF file called soc-damquick.pdf. Many exercises are nominally allocated marks at Tripos examination level (i.e. with 20 marks making a full exam question).

There is some repetition of material between the exercises, so a suitable target is to solve approximately one third of them. Exercises marked with a ♡ form a recommended core. Some sections contain additional preference instructions.

Supervisors are recommended to mainly just set the exercises marked with the ♡ symbol, leaving other parts of the sheets for discussion or for self-study revision purposes. Example answers are available to supervisors.

The course has been rejigged a few times in the last ten years and, despite my attempts, it is likely that one or two of the questions here are now no-longer-relevant or else under the wrong heading. SystemC is no longer being lectured for net-level modelling: it is just presented as a basis for TLM modelling, so where a question suggests giving a SystemC descriptons instead of RTL, that route will not be available unless you have separately studied low-level SystemC coding.

SP1: SoC Components and Bus Structure Exercises

0.1 SP1a: SoC Components and Bus Structure Exercises - Full Questions

SoC-SP1a.1 What is meant by polled I/O and how does it compare with interrupt driven I/O ? [4 Marks]

SoC-SP1a.2 ♡ Swapping bit order in a word requires an expensive sequence of instructions on many processors. However, it can be done instantly in wiring. It is proposed to provide an I/O device with three registers respectively of width 8, 16, 32 bits. Any word written to such a register can then be immediately read back with its bits swapped.

a) Give the programming model for such an I/O device and sketch RTL or the circuit diagram for the actual device. [4 Marks]

b) How should such a device be connected to the processor such that it actually saves time or energy compared with the sequence of instructions approach? Discuss various wiring possibilities and their pros and cons [6 Marks].

For your reference, here is some C code (from the ARM ISS), one quarter of this meets one of third of our specification:

```
switch(size_code)
  {
  case 0:  //Reverse byte order in a word.
    d = ((d>>24) & 0xFF) | ((d>>8) & 0xFF00) | ((d<<8) & 0xFF0000) | ((d<<24) & 0xFF000000);
    break;
  case 1: // Reverse byte order in each halfword independently.
    d = ((d>>8) & 0x00FF00FF) | ((d<<8) & 0xFF00FF00);
    break;
  case 2://   RBIT - (a full butterfly) reverse the bit order in a 32-bit word.
    d = ((d>>16) & 0x0000FFFF) | ((d<<16) & 0xFFFF0000);
    d = ((d>>8) & 0x00FF00FF)  | ((d<<8)  & 0xFF00FF00);
    d = ((d>>4) & 0x0F0F0F0F0) | ((d<<4)  & 0xF0F0F0F0);
    d = ((d>>2) & 0x333333333) | ((d<<2)  & 0xCCCCCCCC);
    d = ((d>>1) & 0x555555555) | ((d<<1)  & 0xAAAAAAAA);
    break;
  case 3: //REVSH      Reverse byte order in the bottom halfword, and sign extend to 32 bits.
    d = ((d >> 8) & 0xFF) | ((d<<8) & 0xFF00);
    if (d & 0x8000) d |= 0xFFFF0000;
    break;
  }
```

SoC-SP1a.3 ♡ Sketch a set of typical macro definitions in C suitable for making low-level hardware access to the device of the previous question. Also give the 'device driver' which is three methods that user code can call. (This is not a typical device driver since the whole unit should appear stateless and does not do any I/O.)

What must be done about thread saftey?

SoC-SP1a.4 ♡ Show how to wire up a push button to a GPIO pin and describe the action of (or sketch out the code for) a device driver that sets up the GPIO device and returns how many times it has so far been pressed. Describe firstly polled code. Then write a few sentences about how the interrupt driven solution would be implemented and when it should be used instead. (No credit is allocated for proper debouncing so feel free to neglect that aspect). [10 Marks]

SoC-SP1a.5 A bus arbiter decides which initiator next makes access to the set of shared targets connected to that bus. In simple systems these have fixed behaviour and so do form part of the programming model. Here we consider an arbiter with programmable policy.

a) Sketch the circuit, RTL or SystemC code for a programmable bus arbiter that manages two initiators. Start by specifying the behaviour and then defining a suitable programming model. Then list the net-level connections to the component to draw a schematic symbol before going on to implementation details.

b) Should the programmed I/O access to the bus arbiter itself be subject to arbitration and if this were to be a problem, in what ways could higher-priority access to the arbiter's register file be granted?

c) What are the advantages and disadvantages of having an arbiter programmable? What method might a network-on-chip use to get much richer programmable arbitration?

[`http://www.cl.cam.ac.uk/research/srg/han/ACS-P35/zynq/Zynq-7000-TRM.pdf` If you look at the Zynq 7000 technical reference manual you will see a number of Quality-of-Service blocks instantiated for traffic control. An overview is given in section 5.2.2 They are programmable with maximum number of outstanding transactions, AW and AR channel peak rates, burstiness, average rates.]

SoC-SP1a.6 Fully describe (using diagrams and perhaps sketches of RTL (or SystemC) code) an interrupt controller that stores eight vectors with individual interrupt enable flags. The controller monitors eight interrupt inputs and presents the highest-priority, non-masked interrupt vector to the processor when the processor asserts an interrupt acknowledge signal or otherwise reads the device. *Start by defining a suitable programming model and then list the net-level connections to the component to draw a schematic symbol.*

The ARM Generic Interrupt Controller `http://www.cl.cam.ac.uk/research/srg/han/ACS-P35/zynq/arm_gic_architecture_specification.pdf` *supports multiple cores, so is more complex than needed. It also has additional complexity in that the ARM itself does not have a vectored interrupt capability, so the GIC needs to assist the ARM rapidly branch to the appropriate ISR, whereas a vectored CPU will support an interrupt jump table in hardware.*

As always, syntactic accuracy is not required in SoC D/M examination answers. [15 Marks]

SoC-SP1a.7 How does the processor set up the interrupt distributor device of SoC-SP1a.6 and what must it do after servicing an interrupt ? [4 Marks]

SoC-SP1a.8 How would you make an interrupt distributor that shares work over two CPUs? What are the pros and cons of dynamic distribution of interrupts ? [6 Marks]

SoC-SP1a.9 Bus Bridge.

a) What is the function of a bus bridge in a SoC ? [2 Marks]

b) What typical and possible address translation semantics might a bus bridge implement ? [4 Marks]

c) How might internal queue structure vary between bus bridge designs ? [3 Marks]

d) How might arbitration policy vary between bus bridge designs ? [3 Marks]

SoC-SP1a.10 Input and Output to a Network Controller or (Streaming Media Device)

a) Sketch the structural schematic symbol for a generic network block that is bus target only, giving full details and descriptions of the signals used to connect to a typical system bus. *The network type or internal structure does not matter, it could be Ethernet, USB, Firewire etc..* [6 Marks]

b) Assuming the device cannot be a bus master, give skeleton code for a typical interrupt-driven device driver. What are the essential differences and similarity between an Ethernet and and a Sound Card? Work out a typical CPU and interrupt load for a sound card that does not use DMA. [8 Marks]

c) What advantages are there to giving the controller the capability of being a bus master? [2 Marks]

d Describe the additional signals needed to make the controller a bus master. [4 Marks]

e) Assuming the device cannot be a bus master, give skeleton code for a typical device driver. [6 Marks]

SoC-SP1a.11 Define a feasible protocol for a simplex serial interface that is used between a sound controller (that is connected to a system bus) and an audio output DAC (digital-to-analog convertor) chip. The interface conveys a pair of stereo channels of 16 bit precision at 44.1 ksps. The physical interface that caries the protocol has three nets used respectively for clock, serial data and synchronisation. Just the PCM data needs be conveyed: no volume or other control is needed. [4 Marks]

SoC-SP1a.12 Sketch the block diagram or RTL for a simple audio output controller that uses DMA to send a serial audio data-stream to a DAC. Include the full programmers' model. [12 Marks]

SoC-SP1a.13 Clock Domain Crossing.

a) List basic principles used in the design of a reliable clock-domain crossing bridge to avoid metastability problems and achieve reliable transfer of data ? [6 Marks]

b) Sketch the RTL or block diagram for a simplex clock crossing bridge that internally uses one parallel data bus and four-phase handshake ? *If giving RTL, only the receiving side logic is needed.* [6 Marks]

c) What constraints exist for simplex protocols that cross clock domains ? [6 Marks]

d) What constraints exists for duplex protocols that span clock domains ? [2 Marks]

SoC-SP1a.14 Exercise: sketch RTL code for a non-preemptive version of the 3-input arbiter given on this page `https://www.cl.cam.ac.uk/teaching/1011/SysOnChip/slides/sp3socparts/zhp711199d4d.html` Alternatively, provide RTL code for a round-robin, non-preemptive version of the 3-input arbiter. *An asynchronous implementation is quite tricky unless you are experienced at logic designing with transparent latches and other level-sensitive latches, so feel free to present a synchronous design, which is just a finite-state machine.*

0.2 SP1b: SoC Busses, Partition and Technology - Full Questions

SOC-SP1b.1. ♡ : SoC Structure.

a) Sketch the block diagram for a SoC with one processor, one SRAM, one ROM, one Counter/Timer block and one PIO section, all connected to a single bus without any bus bridges. [5 Marks]

b) List the (main) net-level interface signals needed for a bus port that enables multiple bus operations to be in flight at once (such as the BVCI port lectured, or an IP block interface of similar functionality) and explain the protocol. [6 Marks]

c) Is it appropriate for DMA to be supported or used in the SoC of part *a)* ? [3 Marks]

d) How are interrupt signals routed in the SoC of part *a)* ? [3 Marks]

e) What modifications are needed if a second processor core were to be added ? Is a second bus a good idea ? [3 Marks]

SOC-SP1b.2. : Multiple Busses With Bridges.

a) In SoC terms, what is a bus and how does it compare with the 1980's concept of a motherboard bus (such as the ISA or PCI bus) ? [2 Marks]

b) How might the destination port for a transaction over such a bus be decided ? [2 Marks]

c) What is a bus bridge, what transactions might it support and what internal operations might it implement ? [4 Marks]

d) If a SoC is designed with a number of bridged busses, what are the main aspects that determine the allocation of initiators and targets to the busses ? [3 Marks]

e) Is there no real difference between a Network On Chip and a set of bus bridges ? [3 Marks]

f) What form of bus protocol is needed for good performance on a SoC that uses a number of bridges busses or clock domains ? [3 Marks]

g) How is contention for destinations handled in a SoC that uses a number of bridges busses compared with a NoC (network on chip) ? [3 Marks]

SOC-SP1b.3. : Network-On-Chip (NoC).

a) What is meant by the term Network-on-Chip and what are the main differences between using a number of bus bridges and a network fabric? [5 Marks]

b) Describe two buffering techniques that might be used in a NoC ? [2 Marks]

c) Describe a flow control technique used in a NoC ? [2 Marks]

SOC-SP1b.4. : DRAM and Cache.

a) What are the main features of DRAM and why was it not commonly integrated as part of a SoC ? [5 Marks]

b) Why should out-of-order read responses ideally be supported by a SoC Bus or NoC ? [5 Marks]

c) Using a system clock of 400 MHz, a 32 bit MIPS/ARM-like CPU is served without a cache by a 16-bit DRAM system with the following parameters

Operation	Clock cycles	Function
RAS	3	Sending row address,
CAS	1	Read or write 16 bits in current row,
RAS precharge	2	Write back time when finished with row.

Making some assumptions about the pattern of access that the processor will make of the memory, estimate its performance in terms of instructions per second. [5 Marks]

d) If all instructions for inner loops are copied to a 32-bit wide on-chip SRAM (that provides true random access at 400 MHz) at code start, what is the performance now. [5 Marks]

e) If a cache structure with 98 percent instruction and 80 percent data hit rate is applied, what processor performance is now achieved ? You may consider in-order and out-of-order processors but full credit awarded for either. [5 Marks]

SOC-SP1b.5. : JTAG Port and Test Modes. *(In lectures JTAG was only briefly mentioned in relationship to debugging and GDB remote serial protocol (RSP)). This question is suitable for discussion in supervisions but the material was not lectured.*

a) Why do ASICs commonly support special test modes? [4 Marks]

b) Define and compare boundary scan with full scan test path [4 Marks]

c) Briefly describe the structure and operation of the JTAG test port used on many chips. [4 Marks]

d) How can JTAG ports be combined and is this a good idea within a single SoC ? [4 Marks]

e) What other uses can the JTAG port frequently be put to ? [4 Marks]

SOC-SP1b.6. : Cell Library.

a) Give a short list of the types of logic cells typically found in a standard cell library. [5 Marks]

b) List five types of information that should be stored about each cell. [5 Marks]

c) Name several illustrative, specialist VLSI structures or components that cannot readily be made out of standard logic cells and explain why custom design is needed. [5 Marks]

SP2: Silicon Energy, Power and Technology

1.3 SP2: Silicon Energy, Power and Technology - Full Questions

Power Estimation and Control questions:

POWER.1. Dynamic Power Gating.

a) Why might a subsystem on a SoC never be turned on in that SoC's lifetime ? [2 Marks]

b) What low-level and high-level mechanims implement such permanent power gating? [4 Marks]

POWER.2. Dynamic Clock Gating.

a) What is dynamic clock gating and why is it used ? [4 Marks]

b) Describe some common clock-gate insertion transformations. [6 Marks]

c) Compare dynamic clock gating with power isolation islands in terms of automation, scale and functionality. [6 Marks]

d) Consider the pros and cons of coarse-grained manual clock gating logic insertion with the use of an automatic tool/compiler that inserts clock gating. [4 Marks]

POWER.3. ♡ : VLSI Energy Use. For this question, use the following figures and assume values for or look up values of any other information that you feel you need. Credit is awarded for the method and not for the numerical results.

Parameter	Value	Unit
Drawn Gate Length	0.08	μm
Metal Layers	6 to 9	layers
Gate Density	400K	gates/mm^2
Track Width	0.25	μm
Track Spacing	0.25	μm
Gate Output Capacitance	0.06	fF
Gate Input Capacitance	0.03	fF
Tracking Capacitance	1	fF/mm
Core Supply Voltage	0.9 to 1.4	V
FO4 Delay	51	ps
Leakage current	21	nA/gate

A processor core in the above technology uses 200k gates, excluding cache memories. It has two operating conditions: 100 MHz at 0.9 volts or 400 MHz at 1.4 volts. The average net activity ratio during halt is negligible and 0.3 when running.

Give all working and intermediate results. State any additional assumptions you need or use.

a) Estimate the area of the processor. [2 Marks]

b) Compute the power consumed per gate at each operating condition when driving a tracks of 0 mm and 1 mm. [2 Marks]

c) Estimate the power consumption of the processor core when halted and running for each operating condition. [3 Marks]

d) Compared with having the processor running at full performance all the time, how much power is saved just by halting the processor when it is idle ? [2 Marks]

e) How much power is saved by dynamic frequency scaling ? [2 Marks]

f) How does dynamic frequency scaling compare with halting ? [2 Marks]

g) How much power is saved by combined dynamic voltage and frequency scaling ? [2 Marks]

h) How much power might be saved by power gating (i.e. power isolation) ? [2 Marks]

i) Estimate the relative costs of performing a 32 bit addition and sending the 32 bit result 1 mm over the chip [3 Marks]

POWER.4. : Dynamic Voltage and Frequency Scaling.

a) Give a formula for the power dissipation associated with a net on a silicon chip. [3 Marks]

b) What is meant by dynamic clock gating and compare this to a technique where software writes to a control register to turn off a clock generator ? [3 Marks]

c) For a fixed supply voltage, quantify the power benefits of frequency scaling. In other words, compare computing quickly and halting with computing more-slowly and finishing just in time. [3 Marks]

d) Give two ways that the supply voltage to a region may be varied? [3 Marks]

e) Using variable supply voltages, quantify the power benefits of frequency scaling. [3 Marks]

f) In supervisions, discuss the architecture of an ASIC (or part of) that uses all of these techniques. [5 Marks]

POWER.5. : Power Consumption

This question is primarily for discussion in supervisions.

a) What are the main components of power consumption in a laptop computer? [5 Marks]

b) How does clock frequency affect power consumption ? [5 Marks]

c) How might clock frequency be controller in a laptop and for what reasons ? [5 Marks]

d) When viewing a DVD (including moving video and audio) on a laptop, what is the best clock frequency policy? [5 Marks]

POWER.6. Technology/Scaling.

This question is primarily for discussion in supervisions.

a) What is meant by the term *feature size* in VLSI ? Give typical values. [5 Marks]

b) What are the main consequences of moving to a smaller feature size in VLSI fabrication ? [5 Marks]

c) What happens to the relative costs of computation and communication as features get smaller ? [5 Marks]

d) Why has parallel computation become more important than ever before ? [5 Marks]

POWER.7. ♡ Battery Life Planning

a) Data can be compressed before being sent by a hand-held unit over a radio link or later on at a central server. What are the advantages and disadvantages? [4 Marks]

b) How might you estimate the battery life of a sensor device that has a 3 Watt radio transmitter, no LEDs or displays and needs to run about five minutes worth of CPU cycles per day? Would solar cells serve it well? [4 Marks]

c) In the last decade or so, digital sensors with radio (wireless) backlinks have been installed in all rooms in all public buildings. What have been the main forward steps in technology that have faciliated this? Does the price of copper make a difference? [3 Marks]

POWER.8. Dark Silicon

a) What is the penalty for putting more and more transistors on a chip? (Or is there none?).

b) How can yield problems be mitigated for regular structures ?

c) When/why might we choose a high-leakage transistor technology over a low-leakage one (for supervision discussion)?

d) What is the Power Wall ?

 END OF DOCUMENT.