COMP32211

Two hours

Question ONE is COMPULSORY

UNIVERSITY OF MANCHESTER
SCHOOL OF COMPUTER SCIENCE

Implementing System-on-Chip Designs

Date: Monday 18th January 2016

Time: 14:00 - 16:00

Please answer Question 1 and also
TWO other questions from the remaining FOUR Questions provided

---

This is a CLOSED book examination

The use of electronic calculators is NOT permitted

[PTO]

Section A

This question is compulsory.
Answer any *ten* of the subsections.
Each subsection carries two marks.

1. a) What is ‘clock skew’? (2 marks)

b) Estimate (to about one significant figure) the memory bandwidth needed to read out data for a ‘high-resolution’ colour graphics display. You may define an appropriate resolution, as long as it is reasonable but you must clearly state the parameters you choose and the units for your answer. (2 marks)

c) During the ASIC ‘place and route’ process, in an initial placement of standard cells the desired gates will be interspersed with inactive ‘spacer’ cells. Why is this? (2 marks)

d) Why is ‘leakage’ in CMOS transistors more a concern now than it was 10 years ago? (2 marks)

e) Write down the truth table for the transistor circuit shown in figure 1. (2 marks)

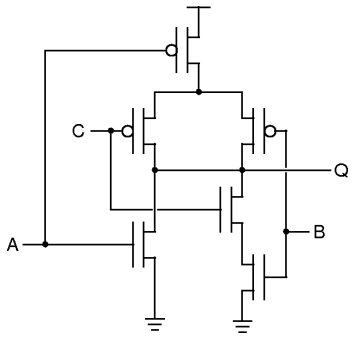

Figure 1:

f) What is the difference in definition between *static* and *dynamic* power dissipation in a CMOS ASIC? Which of these should be the more important design concern in a chip for a utility supplier’s gas ‘smartmeter’ – and why? (2 marks)

g) ‘JTAG’ can provide serial test access to parts of an SoC. Why is a *serial* interface convenient? What else is the interface sometimes useful for? (2 marks)

h) When implementing an SoC, what is meant by ‘technology mapping’? (2 marks)

i) What is meant by the ‘length’ and ‘width’ of a CMOS transistor? (A sketch may aid your description.) Which of these dimensions will often occur in a range of options for the same *logical* gate in a standard cell library, and why? (2 marks)

j) With reference to the following Verilog code:

```
always @ (posedge clk)
  begin
  #3;
  a <= #1 a + 1;
  b <= #2 a - 1;
  end
```

if variable `a` starts with the value 4 and `clk` makes a 0 → 1 transition at time ‘10’, what do the variables `a` and `b` become and at what simulated times? (2 marks)

k) Write a description of an 8-bit, 4:1 multiplexer module in Verilog. (2 marks)

l) In a given CMOS technology, why might you expect a 2-input AND gate to have a greater t_{pd} than a 2-input NAND gate? (2 marks)

Section B

Answer any *two* questions from this section.

2. a) What differences are appropriate in test patterns used for production testing of SoCs as opposed to the patterns a designer might apply to verify the correctness of the design? (2 marks)

   b) Imagine you are responsible for the *post-production test* of a new ASIC. What feature(s) might you include in the design to make the job easier? Describe how this/these operate and state, *qualitatively*, any associated costs. (6 marks)

   c) Modern SoCs are likely to contain several processor 'cores' which are software programmable. How might these be exploited to assist with testing the device? (4 marks)

   d) Give *two different* ways of *providing the software* for an on-chip processor to perform chip tests. Briefly discuss the pros and cons of each. (8 marks)

3. a) Explain what is meant by ‘GALS’ interconnection. (2 marks)

b) Suggest a reason why different parts of an SoC might have independent clock frequencies. (2 marks)

c) Figure 2 shows a synchronisation loop between two independently clocked domains. The logic ‘clouds’ are combinatorial logic only and conceal any decision making. Explain why, in each clock domain, the input from the other domain has been passed through *two* serial flip-flops with no logic in between. (3 marks)

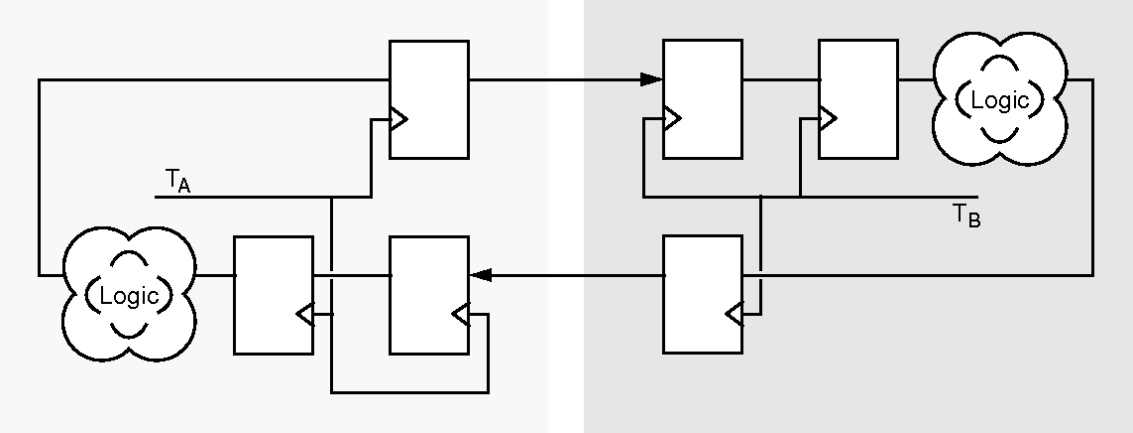

Figure 2: An inter-block synchronisation loop.

d) Referring to the circuit in figure 2, if a single data word is to be sent from a clock domain *A*, with a clock period of T_A, to a clock domain *B*, with a clock period of T_B, derive a formula for the *average* time taken from the start of one transfer to the start of the next, assuming that transfers take place as rapidly as possible.
State any assumptions about the circuit you make. (3 marks)

e) A colleague is worried that the circuit in figure 2 will not be reliable enough when synthesized onto an SoC. Suggest a mechanism for improving the reliability, justifying why you believe it would help. (2 marks)

f) The circuit from figure 2 is proposed to control and validate the flow of 32-bit words between clock domains on a SoC. However it is observed that there is a need to transfer large (˜kilobytes) data packets and that the *data throughput* of the resulting interface will be too low for the application. The clock frequencies are fixed by other constraints.
Discuss, in *qualitative* terms, how you might increase the throughput of the interface. Include a *short* commentary on any merits and demerits of your mechanism(s).
A text description is adequate here; a circuit is not *required.* (8 marks)

4. a) What is meant by the ‘threshold’ of a CMOS transistor? (2 marks)

   b) A manufacturer offers gates made from transistors with ‘high’ and ‘low’ thresholds which can be mixed on a single SoC. What characteristics might you expect from the two different families? Where might it be appropriate for each type to be used? (4 marks)

   c) Briefly discuss the concept of *power gating* a subcircuit on an SoC. Your answer should, at least, attempt to address the following points:

   - What the term means.
   - What (if any) benefit(s) are there?
   - To what types of circuits can it be applied relatively easily?
   - To what types of circuits can it not be applied without additional consequences? and why not?
   - Under what dynamic conditions is it feasible? Why?

   (8 marks)

   d) Sketch the transistor-level schematic diagram for a high-performance *power gated* 3-input NAND gate. Clearly indicate what (i.e. ‘high’ or ‘low’) the threshold of each transistor is. (6 marks)

5. a) What is meant by “pipelining” a data processing logic function? (2 marks)

   b) Explain why pipelining a logic function can increase both its *throughput* and its *latency*. (4 marks)

   c) What is meant if a data processing pipeline is said to be “well balanced”? (2 marks)

   d) Give *two* different reasons why pipelining is an appropriate way of introducing parallelism to accelerate the execution path of a ‘general purpose’ microprocessor. (4 marks)

   e) In a particular design it turns out that one particular pipeline stage has a critical path which is significantly longer than its neighbouring stages. *The logic within this stage cannot be further pipelined.*
   Outline *two* different techniques which you might employ to improve the throughput, noting any particular drawbacks in each case. (8 marks)

END OF EXAMINATION