

Problem-Set #2 – COE838

Accelerator-based System-on-Chip and Basics of SoCs

Q-1. You are designing an embedded SoC using a CPU core with no floating-point support as host. Does it make sense to add an accelerator to implement the floating-point function? Explain.

Q-2. You are designing an embedded SoC using a high-performance embedded processor with floating point. Does it make sense to add an accelerator to implement the floating-point function? Explain.

Q-3. Compare and contrast a co-processor and an accelerator.

Q-4. What factors determine the time required for two processes to communicate? Does your analysis depend on whether the processes are implemented in hardware or software?

Q-5. Which is better suited to implementation in an accelerator: Viterbi decoding or discrete cosine transform?

Q-6. Which is more important in an embedded SoC: throughput or latency? Explain your answer.


Q-7.

A video compressor performs motion estimation on 16×16 macroblocks; the search field is 31 pixels vertically and 41 pixels horizontally.

- a. If we search every point in the search area, how many SAD operations must we perform to find the motion vector for one macroblock?
- b. If we search 16 points in the search area, how many SAD operations must we perform to find the motion vector for one macroblock?

Q-8.

Estimate the execution time and required hardware units for each dataflow graph. Assume that one operator executes in one clock cycle and that each operator type is implemented as a distinct module (no ALUs).



Q-9.

Use a Huffman code to encode these five-bit opcodes: 00000, 00001, 10010, 10001, 00011. Show the Huffman coding tree and the codes for each opcode. Assume that all opcodes are equally probable.

Basic of ICs and SoCs

Q-10.

A four - segment pipeline implements a function and has the following delays for each segment ($b = 0.2$):

<u>Segment #</u>	<u>Maximum delay *</u>	* Excludes clock overhead of 0.2 ns.
1	1.7 ns	
2	1.5 ns	
3	1.9 ns	
<u>4</u>	<u>1.4 ns</u>	

- a) What is the cycle time that maximizes performance without allocating multiple cycles to a segment?
- b) What is the total time to execute the function (through all stages)?

Q-11.

Repeat Question 10 if there is a 0.1 ns clock skew (uncertainty of ± 0.1 ns) in the arrival of each clock pulse.

Q-12.

A processor die ($1.4 \text{ cm} \times 1.4 \text{ cm}$) will be produced for five years. Over this period, defect densities are expected to drop linearly from 0.5-to-0.1 defects/cm². The cost of 20 cm wafer production will fall linearly from \$5,000 to \$3,000, and the cost of 30 cm wafer production will fall linearly from \$10,000 to \$6,000. Assume production of good devices is constant in each year. Which production process should be chosen?