

Problem-Set #2 – EE8205

Accelerator-based Embedded Systems

Q-1. You are designing an embedded system using an Intel Xeon as a host. Does it make sense to add an accelerator to implement the function: $z = ax + by + c$? Explain.

Q-2. You are designing an embedded system 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-3. You are designing an embedded system 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-4. You are designing a bus based accelerated system that performs the following function as its main task:

```
for (i = 0; i < M; i++)
  for (j = 0; j < N; j++) {
    MAXVAL = max_of_5 (pix[i][j - 1], pix[i - 1][j], pix[i][j], pix[i + 1][j], pix[i][j + 1]) ;
    f[i][j] = (pix[i][j - 1] + pix[i - 1][j] + pix[i][j] + pix[i + 1][j] + pix[i][j + 1]) / (5 * MAXVAL);
  }
```

Assume that the accelerator has the entire pix and f arrays in its internal memory during the entire computation of pix is read into the accelerator before the operations begin and f is written out after all computations have been completed.

a) Show a system schedule for the host, accelerator, and bus assuming that the accelerator is inactive during all data transfers. (All data are sent to the accelerator before it starts and data are read from the accelerator after the computations are finished.)

b) Show a system schedule for the host, accelerator, and bus assuming that the accelerator has enough memory for two pix and f arrays and that the host can transfer data for one set of computations while another set is being performed.

Q-5. Compare and contrast a co-processor and an accelerator.

Q-6.

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-7.

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

Q-8.

Which is more important in an embedded computer system: throughput or latency? Explain your answer.

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 the opcodes are equally probable.


Q-10.

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 in a distinct module (no ALUs).

a)



b)



Q-11.

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?