LSU EE 7722 **Homework 1** **Due: 23 February 2018**

Read Sadasivam 19 Micro Mag: "IBM Power9 Processor Architecture" and answer the questions below based on this paper.

Problem 1: The lsli02-cores notes, `http://www.ece.lsu.edu/gp/2018/lsli02-cores.pdf`, has a page headed "Vector Units in Current Processors," which is on page 65 as of this writing. Here is the contents of that page:

Vector Units in Current Processors

Intel (i5, i7) Haswell

SSE: Vector Registers (`xmm`): 16×128 bits (4 SP, 2 DP).

AVX: Vector Registers (`ymm`): 16×256 bits (8 SP, 4 DP).

Eight-Way Superscalar (based on *microops*, etc.).

Two vector units, and so bandwidth is 2 vector insn / cycle.

Vector latency (including FMA), 5 cycles.

(*a*) Provide the same kind of information for the Power9 based on the paper and other sources. If other information sources are used please list them.

Problem 2: Figure 3 in the Sadasivam paper shows the pipeline stages in Power8 and Power9, emphasizing that Power9 has fewer stages. This figure is drawn to show the sequence of stages an instruction passes through (from top to bottom) in each microarchitecture. For this problem only consider Power9.

Figure 5 shows a block diagram of a core, which is drawn to show the major connections between storage devices (caches, buffers), instruction processing units (such as decode), functional units, etc.

(*a*) As best you can label the blocks in Figure 5 with the stages shown in Figure 3. When there are multiple units of the same type, such as AGEN just label one of them.

(*b*) In Figure 5 the diagram show which blocks are part of the front end (through which instructions flow in program order) and which are part of execution unit pipelines (into which instructions may enter out of program order).

Problem 3: The paper emphasizes the benefits of its *Execution Slice Microarchitecture*.

In the class's set-2 notes, `http://www.ece.lsu.edu/gp/2018/lsli02-cores.pdf`, typical supercalar and short-vector instruction organizations are described and contrasted. (Don't confuse modern short-vector instructions which we've talked about in class, such as those which are part of Intel AVX or Sun VIS with vector instructions in 20th century supercomputers such as Cray and an attempted revival described in the Stephens ARM paper.)

The class notes compared the flexibility and cost of providing a particular FP bandwidth with either a superscalar processor without a vector unit or with a scalar processor that included a vector unit. Explain whether the Execution Slice Microarchitecture as described in Sadasivam provides a good example of the comparisons described in the class notes or whether the Execution Slice Microarchitecture undermines the points made in the class notes.

Provide code examples and statements from the slides and along with them indicate material from the paper that confirms or undermines the points the slide material was making.