CSE 490/590 – Quiz 3 – April 18, 2025 – «Version»

«Last_Name», «First_Name» «OrgDefinedId» «Username» SEAT «Seat»

Instructions:

- 1. This Quiz is **closed book and closed notes.**
- 2. You may have pens, pencils, erasers, a calculator, and a water bottle/drink.
- 3. Electronic devices are **NOT ALLOWED** during this Quiz, with the exception of a calculator. **You may NOT share a calculator during the Quiz.**
- 4. When you are finished, remain in your seat and raise your hand and we will come and collect your Quiz. You must not talk to anyone in the room until your Exam has been collected and you have left the room.
- 5. Any form of cheating/academic integrity violation, including (but not limited to) violation of the rules above will result in an automatic 0 on the Exam.
- 6. Answer based on your own understanding.

Please fill your name in the blank and sign the statement below:

I,

have read, and acknowledge that I will adhere to the instructions above and if not followed, I will accept the penalty given by the instructor.

Signature: _____

STOP!! DO NOT START THIS QUIZ

UNTIL YOU ARE TOLD TO DO SO

FOR GRADERS ONLY:

Q1:	/6	Q4:/10
Q2:	/8	Q5:/18
Q3:	/8	Total:/50

[Question 1] (6 Points)

Referring to Flynn's Taxonomy in Computer architecture, explain each, their benefits and possible applications.

a. SISD

SISD: One processor executes one instruction on one data item at a time. Simple, used in traditional CPUs.

b. SIMD

SIMD: One instruction operates on multiple data items in parallel. Efficient for vector/matrix processing, used in GPUs.

c. MIMD

MIMD: Multiple processors execute different instructions on different data simultaneously. Supports multitasking and parallelism in multicore systems.

[Question 2] (8 Points)

What are centralized shared memory multiprocessors (SMP) and Distributed Memory multiprocessors (DSM)? Mention their key features and differences.

SMP (Centralized Shared Memory Multiprocessors): All processors share a single memory. Easy to program, but less scalable due to memory bottlenecks.

DSM (Distributed Shared Memory Multiprocessors): Each processor has its own local memory, but logically shared. Better scalability, but complex memory access and coherence.

[Question 3] (8 Points 6+2)

Explain the concept of gather-scatter operation. (6 points)

Gather: Collects non-contiguous data from memory into a single vector register.

Scatter: Distributes elements from a vector register to non-contiguous memory addresses. Useful for sparse matrix or irregular memory accesses.

Explain Stride. (2 points)

Stride: Distance between successive data elements accessed (e.g., stride-2 means every 2nd element is accessed).

[Question 4] (10 Points)

Considering the VMIPS structure below, answer the following questions.

a. What is a Vector Architecture? Discuss its application(s). (4 Points)

Vector Architecture: Uses vector registers to process data in parallel. Efficient for scientific computing, image/video processing.

b. Explain the difference between vector registers and scalar registers. (2 Points)

Vector vs Scalar Registers: Vector holds multiple data elements; scalar holds one.

c. Explain how vector load differs from scalar load. (2 Points)

Vector Load vs Scalar Load: Vector load moves multiple data items to vector registers; scalar load moves one to a scalar register.

d. Explain how vector store differs from scalar store. (2 Points)

Vector Store vs Scalar Store: Vector store writes a whole vector to memory; scalar store writes one item.



[Question 5] (18 Points)

a. Briefly explain how Snoopy-based and Directory-based cache coherence protocols work. (4 Points)

Snoopy-based Protocol: Caches listen ("snoop") on a bus to detect and react to memory writes. Works well for SMPs with a shared bus.

Directory-based Protocol: A central directory tracks which caches have copies of a block and coordinates updates. Scales better for DSM systems.

- b. Compare the above two cache coherence in terms of:
 - i. Scalability (3 Points) Snoopy is limited by bus traffic; Directory scales better as it reduces global communication.
 - ii. Use cases (3 Points) Snoopy for small SMPs; Directory for large-scale DSM systems or clusters.

- c. Explain the three block states tracked by the directory in a Directory-based cache coherence protocol. (4 points)
 - Uncached no cache holds the block.
 - Shared –One or more nodes have the block cached, value in memory

is up-to-date, Set of node IDs

• Modified/Exclusive - Exactly one node has a copy of the cache block, value in

memory is out-of-date, Owner node ID

d. In a directory-based cache coherence protocol, how are read and write misses handled in a shared block? (4 points)

• Read miss – The requesting node is sent the requested data from memory, node is added to sharing set

• Write miss – The requesting node is sent the value, all nodes in the sharing set are sent invalidate messages, sharing set only contains requesting node, block is now exclusive