A1. Architecture (25 points) - Shared-memory Multiprocessor

(a) (10 points) Describe the basic architecture of a centralized shared-memory multiprocessor in terms of number and type of processors, memory, and network interconnect.

(b) (5 points) In symmetric shared-memory multiprocessors, what is meant by the term "*uniform memory access*"?

(c) (10 points) Indicate the major disadvantage for sharing a centralized memory when the interconnection means is a bus. In other words, when is it that sharing a centralized memory becomes less attractive with this interconnection means?

A2. Architecture (25 points) - Distributed-memory Multiprocessor

(a) (10 points) Describe the basic architecture of a distributed-memory multiprocessor in terms of number and type of processors, memory, and network interconnect.

(b) (7 points) List the major advantages of using distributed-memory multiprocessors.

(c) (8 points) Briefly indicate what communication mechanism is used in:

- a multiprocessor with a shared address space
- a multiprocessor with multiple address spaces

A3. Architecture (25 points) - Maintaining Coherence for Multiprocessors

(a) (12 points) There are two classes of protocols to maintain cache coherence for multiprocessors, called *directory based* and *snooping*. Briefly describe the two classes, comparing and contrasting them in terms of the techniques used to track the sharing status of a block of physical memory.

(b) (8 points) There are two ways to maintain the coherence requirements in snooping protocols, i.e., *write invalidate protocol* and *write broadcast protocol*. Briefly describe the two protocols in terms of access to data items in cache.

(c) (5 points) Briefly explain why all recent multiprocessors have opted for *write invalidate protocols* rather than *write broadcast protocols*.

A4. Architecture (25 points) Techniques to Exploit more Instruction-level Parallelism

(a) (9 points) Hardware-based speculation extends dynamic scheduling. Briefly explain how *hardware-based speculation* deals with (a.1) branch predictions, (a.2) control dependencies, and (a.3) scheduling of different combinations of basic blocks.

(b) (16 points) Consider the Tomasulo algorithm extended to support hardware speculation. List and briefly described the four steps involved in instruction execution for this architecture. Clearly describe in which step and how RAW hazards are checked and solved.