- 1. Instruction level parallelism (ILP) can be exploited through two primary approaches: static scheduling approach and dynamic scheduling.
  - (a) Explain the differences between these two approaches.
  - (b) In which cases can the dynamic scheduling approach exploit more ILP than the static scheduling approach? Support your answer with examples.

(10/10 points)

- 2. The Branch Target Buffer (BTB) is a technique used in modern processors to predict the target address of a branch instruction before it's executed.
  - (a) Explain how predicting a branch's target address helps the processor reduce pipeline stalls and improve performance?
  - (b) What cases may cause BTB misses?

(10/10 points)

3. Amdahl's Law is a principle that helps to understand the potential performance improvements in a system when only a portion of the system is improved. It is often used in the context of parallel computing to predict the theoretical maximum speedup of a program as a result of enhancing a specific part of it. Amdahl's Law can be expressed mathematically as:

$$S = \frac{1}{(1-P) + P/N}$$

Where:

*S* is the overall speedup of the system.

*P* is the proportion of the program that can be improved.

*N* is the factor by which that portion can be improved.

If a program is 60% parallelizable, what is the maximum speedup we can obtain by using a 4-core processor to execute this program? Explain your answer using Amdahl's Law.

(10 points)

- 4. The cache performance is determined by temporal locality and spatial locality in a program.
  - (a) Describe an approach that optimizes caching for temporal locality, and explain how

it works.

(b) Describe an approach that optimizes caching for spatial locality, and explain how it works.

(10/10 points)

- 5. Answer "True" or "False" for the following statements, and explain WHY.
  - (a) Larger cache blocks exploit spatial locality. Therefore, increasing the block size always decreases the miss rate.
  - (b) Cutting the instruction datapath into many small pipeline stages can increase the processor clock rate. Therefore, deeper pipelines (i.e., more pipeline stages) will always have higher performance.
  - (c) An 8-core processor always has higher performance than a 4-core processor.

(10/10/10 points)