Comparative Architectures Supervision Problem (Set 1)

Wei Song 30/01/2016

The following problems are mostly open-ended. Diagrams or charts are not needed in the answers. Therefore it is highly recommended to type the answers. Please keep your answers organised and tidy. Most of the questions are taken from the supervision notes from the lecturer. If you are not sure about the answers, you could look up from the Hennessy and Patterson book or directly from the Internet. If you could not find the answer or would like a discussion during the supervision session, please label it in your answer.

Lecture 1. Introduction

Q1: Will we ever converge on a single optimal architecture for all computers? [Sup-work Q 1.2]

Q2: Why does power consumption now constrain processor design much more than it has historically? [Sup-work Q 1.3]

Q3: Why is MIPS (millions of instructions per second) a very poor measure of processor performance? [Sup-work Q 1.9]

Q4: (Hard) discuss the important metrics in evaluating the performance of computers targeting different markets, including mobile phones, servers, real-time systems, sensor networks. A non-exhaustive list of available metrics includes: power, energy, MIPS, power efficiency, cost, reliability, flexibility, responding time, etc.

Q5: (Hard) how might recent advances in die stacking help to improve microprocessor performance and reduce costs? (2015 Paper 7, Q 5)

Lecture 2. Fundamentals of Computer Design

Q6: (Hard) discuss the pros and cons of architectures with a 64-bit word size versus those with a 32-bit size. What applications are likely to benefit most? (2005 Paper 7, Q 1) What hardware extensions are needed for the 64-bit CPU to run a 32-bit application to ensure backward compatibility?

Q7: (Hard) If you were a processor architect targeting embedded applications where memory is a scarce resource, how might you design a RISC-like instruction set that will achieve efficient use of memory? (2005 Paper 7, Q 1)

Lecture 3 & 4. Advanced Pipelining

Q8: Provide a case when anti-dependence can cause actual data hazard in a processor.

Q9: What determines the optimal pipeline length for a microprocessor? (Sup-work Q 3.1)

Q10: How does the existence of multiple parallel pipelines make it more difficult for us to provide support for precise exceptions? (Sup-work Q 3.7)

Q11: (Hard) discuss the pros and cons of different address formats, from zero-address to three-address format (“specifying an instruction’s operands” in Lecture 2). (hint, what is the effect on structure and data hazards?)

Q12: Why are 2-bit saturating counters often used to predict a branch's direction in preference to single bit schemes? (Sup-work Q 4.3)

Q13: Why might a branch target buffer provide a poor prediction of procedure return addresses and what hardware solution may be employed to improve the accuracy of such predictions? (2011 Paper 8, Q 3(a))

Q14: (Hard) why might a single conditional move instruction be supported rather than the conditional (or predicated) execution of every instruction in the instruction set? (Sup-work Q 4.8)