Computer Science 654
Advanced Computer Architecture
Spring 2009

Instructor: Peter Kemper, kemper@cs.wm.edu, ph: 1-3462
Time/Place: MWF 9:00-9:50pm, Washington Hall 307
Office hours: MW 2:004:00pm, McGlothlin Street 104A.
Textbook: J. Hennessy and D. Patterson, Computer Architecture: a quantitative approach, Fourth Edition, 2006
Webpage: http://www.cs.wm.edu/~kemper/cs654

Description:
The course covers the principles, characteristics, and technology trends of computer architecture design at a graduate level. Architecture provides a holistic approach to designing computer systems, so it has to include both hardware and software. Our goal in this class is to learn more about the insides of the box and where these designs are headed, so that we are able to develop software that makes more effective use of current and future computers. The course will cover many of the topics of the textbook, but at varying lengths. As noted in the latest edition of the textbook, computers are undergoing a sea of change, with parallelism and multiprocessing being central in all topics. Whenever possible, the lectures will reflect that.

Fundamentals of Computer Design. We will overview basic architectural concepts, computer abstractions, control and data path, as well as study the technology trends. The rapidly changing nature of this area makes the prediction of future directions necessary. We will also study the principles of measuring, evaluating, and comparing computer systems.

Instruction Sets. This is an immediate interface to the computer hardware. Its design directly affects the performance of a computer system. We will classify instruction sets to RISC and CISC and study the merits and limits of each.

Pipelining. We will study this important architecture performance enhancement design for overlapping computations and communications.

Memory Hierarchy. One of the hurdles in extracting performance from todays (and future) architectures has been (and will be) the dramatically increasing disparity between memory speed and processor speed. Building a memory hierarchy is an effective way to narrow this gap. We will study the issues involved in cache and virtual memory design.

Storage Systems. Data intensive applications, such as commercial data bases, rely mainly on high performance Input/Output devices. We will overview the functions of I/O devices and their interfaces to computer systems.

Interconnection networks and Multiprocessors. Computers with multiple CPUs have now forcefully entered the mainstream. Moreover, low-cost but powerful computers can be interconnected together to build computing clusters that rival the performance even of high end supercomputers. Although parallelism material will be present in all the above topics, multiprocessor systems and networks of PCs/workstations deserve a closer look. GPU gen- eral purpose processing will be addressed as well.

Coursework: Your grade will be determined as follows:

1. Homework: 40% Three or four homework sets.
2. Project: 30% Topics to be given out in March. Final presentation and writeup required.
3. Final test: 30% A three hour, in class test.

Grading policy: The Honor Code applies on all assignments as follows. You are encouraged to discuss the homework problems with your fellow students in 654, but not with students outside the class. The write up, however must be yours. For the pro ject, students can work in couples, handing in one pro ject and receiving a common grade. For the pro ject, students may NOT discuss with members of other groups.
Students with disabilities: If you have a disability that may affect your participation in this course and wish to discuss academic accommodations, please contact me as soon as possible.