

CmpSci 635
Modern Computer Architecture
Course Syllabus
Spring, 2006

Chip Weems
Office: CS-342
Phone: (413)545-3163
weems@cs.umass.edu

Office Hours:
Wednesday: 10:30 - 11:30
and by appointment

Office Hour Appointments: Send me e-mail to make an appointment. Feel free to drop in whenever I'm in the office. If I'm really too busy to see you, I'll ask you to make an appointment for a later time.

Text: Computer Architecture A Quantitative Approach, 3rd Ed.
John Hennessy and David Patterson, Morgan Kaufmann 2003
Copies available at Jeffrey Amherst Bookstore, or online.

Lecture Notes: My lecture notes are available on line via my personal web page (<http://www.cs.umass.edu/~weems/courses.html>). The previous semester's notes are there already, and I will be updating some of them to reflect recent advances. The CSCF has requested that these not be printed on the Edlab printers, due to the excessive load that they place on those already overburdened machines, and the unbudgeted additional paper cost. Videotapes of the lectures will be on reserve at the main library, usually a few days after the lecture.

Course Goals: The goal of this class is for you to develop an understanding of the factors influencing modern microprocessor and parallel processor designs, and a working knowledge of the specific choices made in some current designs. You will also gain an understanding of some of the acceleration mechanisms employed in modern processors, such as superpipelining, branch prediction, out-of-order superscalar issue, predication, speculation, multithreading, and chip multiprocessing. By the end of the course you should be able to read between the lines of a newly-announced architectural specification and describe some of the tradeoffs that influenced the design and the implications for software. You should also be able to make well-reasoned predictions of system performance and limitations for specific application scenarios.

Grading:

Midterm :	10%	March 27 (Monday after Spring Break)
Final:	10%	
Homework:	20%	
Mini-essays	10%	
Class participation	30%	
Project::	20%	

Exams: Exams are open book and open notes. This is so that you can look up formulas or data from the text or lecture notes -- I don't expect you to memorize this sort of information for

exams. But I do expect you to have read the book, attended class, and to know where to look up the information that you need. There won't be time in the exams for you to read through material you're unfamiliar with. The purpose of the exams is for you to demonstrate that you have attained an operational level of understanding of the material. Bring a calculator capable of exponential functions to all exams.

Homework:

Number: At most eight homework assignments will be given out.

Due Dates: They are each due at the end of class one week later unless otherwise indicated.

Late Policy: You must do your work on time because we'll be correcting it in class. If you know that you have a specific time conflict, make arrangements with me in advance for a separate assignment for late submission. The answers will usually be given out in class on the due date, so credit cannot be given for regular homework turned in after that.

Format: All homework and in-class assignments will be done in the style of an engineering notebook, which is a permanent record of a conceptual flow. Please use an **11x14** spiral-bound field sketch book with at least 80 pages of high quality, non-lined paper (e.g., 65lb acid-free). **DO NOT** use 8.5 x 11 or other smaller sizes of notebook. The campus store usually has a good supply of such books, as does Hastings in Amherst, and Michael's in Hadley. A hardcover version is preferred, as it is more durable. Do your work in pen (color choice is up to you). Please do not use pencil, with the exception of colored art pencils, which may be used for illustrations. Do not remove pages from the book. Only paste other work in if I specifically indicate that you should. Neatly cross out mistakes. This is standard practice for engineering notebooks used in support of patent applications – it is thus a good habit to develop. Think before writing and plan drawings out in advance so that you keep a neat and presentable book that you can be proud of. If you keep your notes in the same book, I will review them when you submit it for homework.

Collaboration: Do not collaborate on the actual solution of the homework. The purpose of the homework is for each person to gauge their comprehension of material from class, and to practice answering questions similar to those in the exams. It is thus best for you to do the work alone. If someone else asks you for help, take the role of a teacher and help them find the solution on their own; don't just give them the answer. If I feel that people have submitted answers that are merely copies of each other, I will grade the one solution and divide the credit for it equally among the copies. You may think that it helps your team to split up the homework and share answers, but if your team-mates then do poorly on the exams because they lack the experience of solving similar problems in the homework, your team will actually suffer more.

Grading: Homework will be corrected as an exercise in class. I will have teams exchange homework and mark each other's work. Be considerate of your fellow classmates by writing legibly and organizing your answers in a clear manner. You will mark each answer for correctness and provide corrections as necessary. I will collect all of the homework to review and assign grades. An overall numerical grade will be given -- some number of points, with the maximum number indicated. Each of the homework assignments makes up the same percentage of your final grade. You should mark answers that are obviously correct or wrong, and comment on answers that are not clear-cut. I will review your marks and make the final assignment of credit for the answers.

Mini-essays

In most lectures, a mini-essay assignment will be given out, due a week later. The purpose of these is to encourage you to keep up in the reading. The questions should be fairly easy, once you've read the associated material and given it some thought. They are intended to be answerable in one or two written paragraphs. Be succinct in your answer. Writing more won't necessarily gain you more credit, and may give the impression that you don't have a clear understanding of the reading. Essays must be on either a 5x8 index card (preferred) or a half-sheet of 8.5 x 11 paper (write on one side only). The standard grade is a check, indicating that your answer shows that you read the material and understood it reasonably well. A minus indicates that your answer doesn't show a solid understanding of the reading (or of the question), and you should reread the material. On occasion, if an answer is especially insightful, I'll give a plus grade.

Class Participation

A large component of your learning takes place in class. It is where you have me as a resource. I will lecture to you less than in a typical class. The actual concepts of modern architectures are fairly simple, although their implementation is often complicated by real-world constraints. Thus, I tend to give short lectures to explain these concepts, and pose questions for discussion that are meant to draw out these implications. I will guide discussion, and add information here and there as necessary to carry the discussion forward or to lead it into a digression that adds depth in a different direction. I will frequently have in-class exercises that you will do as teams. Thus, it is very important that you attend class regularly. I will keep attendance throughout the semester. Please let me know in advance of any scheduled absences. If you are going to be out due to illness, please let your team know so that they can properly prepare for class.

It is very important that we focus our attention during the limited time we have together in class. Each of us comes in filled with thoughts from outside of the classroom. Thus, each day we will begin promptly with a brief mind-clearing exercise, followed by a focusing exercise. During the mind-clearing exercise we will sit in silence and focus as completely as possible on our breathing, or a sound, or a verse. To then focus our attention within the classroom, we will recite some words in unison (these may be a piece of a poem, some nonsense words, song lyrics, a short piece of prose, etc.). The recitation also serves to warm up our voices. Because it is important that we not be distracted while doing these exercises, I will close the door promptly at the start of class. **If you arrive late and the door is already closed, wait outside until I reopen it and invite you in.** I will record late arrivals on the day's attendance, and chronic lateness will count against your class participation grade (three late arrivals equals one absence).

In addition to attendance, your class participation grade depends on how much you contribute to discussions and the in-class exercises, your record of in-class exercises in your sketch-book, my sense of how well you are prepared for class, and the overall attitude and quality of participation that you display. As a team, it is important for you to draw each other out. I will especially notice teams that are dominated by one or two members, with other members remaining silent. Don't let this happen in your team!

Learning Process

It is well documented that we learn in different ways with the two hemispheres of our brains. Our right hemisphere is oriented toward the artistic and creative side of things while our left hemisphere is oriented toward the logical and analytical. Learning is more effective when both hemispheres are engaged and reinforcing one another. Thus, you will find that I sometimes assign unusual homework and class work that is designed to stimulate this kind of holistic learning. I may ask you to write a poem about a technical topic, or to draw an artistic representation of a concept. I won't evaluate these works on their artistic merit, but I will look for evidence of both personal creativity and knowledge being employed. Keep in mind they will also be reviewed by other students in the class.

Thus, in addition to a sketch book, you will need appropriate art supplies for some of your assignments. I recommend Stockmar beeswax art crayons as they have excellent color, are easy to use in a variety of creative ways to achieve varied effects, are non-toxic, and do not require additional tools to use. They can be hard to find, but I can get them for you from a local source. You are welcome to use other media, but keep in mind that it must be appropriate to the paper in your book – so oils, watercolors and pastels are going to be problematic. Colored art pencils and other brands of art crayons are likely to be good choices. Inexpensive children's crayons and colored felt pens are difficult to use with subtlety and I strongly recommend against them. You should have at least eight different colors in your set. Red, green, blue, yellow, black, and sometimes purple, are used in some VLSI layout exercises that we will do.

Teams

Some time soon after the end of the add-drop period, I will divide the class into four-person teams. Your grade for the semester depends on your performance as a team. Only under the most exceptional cases will I assign a different grade to the members of a team. Thus, you are responsible not only for your own grade, but for that of your team partners. Your grade thus depends equally on how well your partners do in the class. If one of your partners is having trouble with some concept, you need to work to ensure that he or she understands it before an exam or homework assignment. This sort of teaching within teams by team members leverages my teaching efforts, ensures that no one falls behind, and helps to deepen understanding of the material for those who are helping others. Here are some ways that you can help to insure that your team does well:

Meet regularly for study sessions. Quiz each other on concepts. Work together on problems from textbooks, or on problems that you make up which are similar to assigned homework. Try doing some of the latter before you tackle the actual homework assignment, so that doing the assignment comes easily. Get together and review homework assignments prior to each exam. Make up mock exams for each other and then grade them.

Do not help each other on the actual homework, but use it as a mini-test between yourselves. Review each other's homework after it is written and discuss how it might be improved. Then make corrections separately. You won't be able to collaborate on exams, and you need this sort of practice to be well-prepared.

Make sure that you have each other's schedules, phone numbers, e-mail addresses, etc. Establish a clear policy of when it is OK to contact one another. Let each other know when you are going to be unavailable with enough advance warning for good contingency planning.

Be realistic with regard to your abilities and time commitments when you divide up project work. Look at the long-term project workload – you can vary the workload over the semester as well as on each portion of the project.

Come to class on time, well prepared. Meet before class to review your preparations; identify, discuss and clarify issues from the reading, and formulate any questions you need to ask in class. Meet briefly after class to debrief each other on what you learned. Each person in a class takes away a different perspective from each lecture and you can potentially quadruple your understanding of what was covered by sharing your perspectives among your team.

Be very explicit in communicating your expectations of each other. Write them down and sign them. Communicate! Communicate! Communicate! Tell each other what you are thinking. Don't keep thoughts inside. Remember to praise each other for jobs well done. If you feel the need to criticize, use statements that start with "I think...." or "I feel..." Avoid criticism that starts with an accusatory "You did...", or "You always...". Praise others in the presence of the team when possible, and criticize them privately.

Think of yourselves as a team. Develop some team spirit. Name your team or your project. Develop a unique style for your homework and demos that you share in common (e.g., team colors, a logo, a border pattern that appears on each page of your sketch book). Get to know each other as individuals. Do lunch. Discuss hobbies, career goals. If you each take the time to care about the other, then you'll do what has to be done to excel as a team.

Write down a list of your individual goals in the class at one of your first meetings. Identify the goals that you have in common, and acknowledge the goals that differ. Make lists of what each of you is willing to contribute to the team effort, and a list of the limits that you won't exceed. Try to reach a consensus on the group's expectations of each other.

Semester Project:

The project is a team effort to explore some aspect of modern computer architecture in an area of particular interest to the team members. Some reasonable projects include:

- designing a new high performance architecture, simulating it, and analyzing its performance
- conceiving a new idea for enhancing an existing architecture, and modifying existing simulation tools to analyze the performance effects of these changes
- measuring the performance of an application on a sample population of processors in order to develop a detailed analysis of the architectural bottlenecks and then propose enhancements which can be measured using software emulation

- using a trace analysis tool to study in detail the execution of some application, and determine the impact of various hardware changes (such as alternate branch prediction schemes, new instructions, different memory hierarchy, etc.) on its performance
- developing a benchmark that tests an aspect of performance that is not adequately covered by existing benchmarks, ensuring its portability to various systems, and evaluating the performance of different processors
- developing a compiler optimization or a new algorithm that takes advantage of some feature of a modern architecture, or of a feature that you add to a simulation of an architecture
- exploration of architecture-dependent performance optimizations for a parallel processor, such as profiling execution of a parallel application on the nodes of a cluster and identifying architectural bottlenecks in different nodes that motivate adjustment of the load distribution
- experimenting with a nontraditional architecture, such as a network interface processor, router core, extreme low power processor, graphics processor, or vector processor. Then develop a novel application or show how an existing application can benefit from the special features of the architecture. Or explore performance issues with the architecture.
- studying security or reliability issues in a processor, identifying architectural issues that may contribute to (or degrade) security or reliability, and proposing ways to enhance security and reliability.

I would like to see you propose an idea of your own that interests you -- the purpose of this effort is for you to gain an in-depth appreciation of the impact of architecture on some area of computer science that is important to you. You will need to develop a written project proposal with a schedule of milestones and a division of labor among the team members. Several classes will be devoted to brief project status reports from teams. If you would like to use the video projector for your presentation, please bring a laptop. If your presentation is available on the web, you may use the machine in the studio. Your status reports and final presentation should be a properly prepared research-style talk of 5 minutes or less. These presentations have the purpose of teaching the rest of the class about your work -- make it interesting and informative. A final written report will be due at the end of finals week.

Past projects have occasionally resulted in publications, synthesis projects, or useful tools. I am happy to discuss such opportunities with your team. It is best to do so in the early planning stages, so that I can help you design your project to produce results that will have the necessary quality.

Reading Assignments -- Read by Date Shown

1.	Feb. 1	W	First class, no reading	Historical perspective
2.	Feb. 6	M	Sect. 1.1 - 1.3, Lecture 1	Processor Technology
3.	Feb. 8	W	Sect. 1.4	Memory Technology
4.	Feb. 13	M	Sect. 1.5, Lecture 2	VLSI layout and cost models
5.	Feb. 15	W	Sect. 1.6 - 1.11, Lecture 3	Measuring performance
	Feb. 20	M	Holiday	
6.	Feb. 21	Tu	Sect. 2.1 - 2.8, Lecture 4	Instruction sets and addressing
7.	Feb. 22	W	Sect. 2.9 - 2.14, Lecture 5	Control flow, <u>proposal presentations</u>
8.	Feb. 27	M	Sect. A.1 - A.3, 3.1, Lecture 6	Pipelining and hazards
9.	Mar. 1	W	Sect. A.4 - A.5, 3.2 - 3.3	Implementation issues
10.	Mar. 6	M	Sect. A.6 - A.9, Lecture 7	Implementation example
11.	Mar. 8	W	Sect. 3.4 - 3.5	Branch prediction
12.	Mar. 13	M	Sect. 3.6 - 3.7, Lecture 8	Superscalar and speculative issue
13.	Mar. 15	W	Sect. 3.8 - 3.15	Limits of ILP, <u>project status reports</u>
	Mar. 20	M	Spring Break	
	Mar. 22	W	Spring Break	
14.	Mar. 27	M	Midterm Exam	Chapters 1 - 3, Appendix A
15.	Mar. 29	W	Sect. 5.1 - 5.3, Lecture 9	Cache basics
16.	Apr. 3	M	Sect. 5.4 - 5.6, Lecture 10	Reducing miss effects
17.	Apr. 5	W	Sect. 5.7 - 5.9	Accelerating access, specialized caches
18.	Apr. 10	M	Sect. 5.10 - 5.17, Lecture 11	Virtual memory and examples
19.	Apr. 12	W	Chapter 4	Software ILP
	Apr. 18	M	Holiday	
20.	Apr. 20	W	Sect. 7.1 - 7.7, Lecture 12	Buses, I/O, and disk, <u>project status reports</u>
21.	Apr. 24	M	Sect. 6.1 - 6.4, Lecture 15	Shared memory multiprocessors (G)
22.	Apr. 26	W	Sect. 6.5 - 6.8	Cache coherence schemes (G)
23.	May. 1	M	Sect. 8.1 - 8.4, Lecture 13	Multiprocessor networks
24.	May. 3	W	Sect. 8.5 - 8.9	Topologies and practical issues
25.	May 8	M	Sect. 8.10 - 8.11, Lecture 14	Distributed memory multiprocessors
26.	May 10	W	Lecture 16, Appendix G	Vector and array processors
27.	May 15	M	Sect. 6.9 - 6.14, Lecture 17	Multithreading, VLIW
28.	May 17	W	No reading	<u>Project final reports</u>

Video Class Addendum

For students in the video section, some aspects of the class must be done differently. Specifically:

Exams: These will be mailed shortly after we do them. You must arrange for a proctor to supervise and keep time. They are open book and open notes. You will need a calculator and colored pencils. Because the midterm solution will be discussed in the next class, you won't receive the next tape until you've sent in the exam.

Homework: Email these as an attached Adobe Acrobat pdf file or as a single Word file. **DO NOT** send multiple files. Spreadsheets and drawings must be inserted into the homework file. It is OK to insert a scan of a hand-drawn figure into the homework document. Homework will be graded by the teaching assistant, who will send it back to you with comments electronically embedded. Homework is due one week later than the date given in class, to account for shipping delays. Immediately after the due date, the TA will send out the solution, so late submissions **cannot** be accepted. It is better to turn in a partial homework than to miss the deadline and receive no credit. Homework should be sent directly to the TA, and you can copy the instructor.

Mini Essays: These should be sent as plain text email messages to the instructor. Keep them short (< 1/4 page of typing). They are due one week later than the reading assignment date in the syllabus, to account for shipping delays. Please keep up with the reading schedule!

Participation: I assume that you watch all of the videos, so you automatically get perfect attendance. You should try to do the in-class exercises along with us. One advantage of the video presentation is that you can stop the playback if you need extra time. Pausing will probably be necessary to compensate for the lack of having a team with which to discuss the problem. Feel free to email me questions or comments, ideas you would have contributed to a discussion, etc.

Learning Process: We will hold our minute of silence while the color-bars and tone are being placed on the tape, so you will see me begin by announcing the lecture number, and then asking the class to stand for our focus exercise. You may skip this if you wish, since I assume you are watching the video at a time that you have set aside to focus on the class. You're also welcome to do this mental stretching work with us – it does help to stir up a new level of brain activity.

Teams: Most video students do not form teams unless they are taking the course together with others in their company, for example. I have had virtual teams of people at different locations, who have split the work of a project. However, many companies have intellectual property policies that make this difficult. Thus, you will most likely be working separately.

Semester Project: Unless you are in a team, or already happen to have access to simulation facilities or an ongoing research project, you will probably do a less ambitious project. For many students, it takes a large time commitment simply to get a simulator, such as SimpleScalar, operational and to be able to use it. There is then not much time left to do an in-depth research project. Thus, if you have access to special facilities that could be the basis of a project, I urge you to explore using them. Some examples: A student working for a database company did a

benchmarking study of the company product running on proprietary data sets that were more realistic than typical benchmarks. A student who had access to a parallel processor studied its performance on some applications. A student working in a lab that had five generations each of several different processor architectures did a study of performance improvement over time, for each architecture, and compared them. Another student did some benchmarking of an embedded processor that was being used in company products. Some students have been able to leverage work projects into also serving as semester projects, by exploring some way that their work project depends on aspects of architecture.