

The University of Texas at Austin
Department of Electrical and Computer Engineering

Course Descriptor: EE 382N.19 - Microarchitecture, Unique Number 16875
Spring semester, 2020

Instructor: Yale N. Patt (<http://www.ece.utexas.edu/~patt>)

Office: EER 5.802

Phone: 512-471-4085

Email Address: patt@ece.utexas.edu

Office Hours: MW, 6:30pm to 7:30pm, and by appointment

Course Title: Microarchitecture.

Objectives of the course – EE 382N.19 is intended to provide a serious introduction to microarchitecture at the graduate level.

At one time, I thought EE 382N was only for the serious graduate student who is interested either in PhD research in microarchitecture or in an industrial position on a leading edge microarchitecture project. However, I have consulted on enough application-specific chips over the past several years that I now believe this course can be useful to the serious student of embedded processors, vlsi circuit design, and hardware/software codesign, to name a few disciplines that are not microarchitecture-centric. In fact, most people are coming around to the notion that the microprocessor of the future will have lots of accelerators – the specification of those accelerators will benefit from the collaboration of the microarchitect and one specifically knowledgeable about the task to be accelerated.

EE382N.19 is known for being a lot of work, so before registering, you should check with your research advisor to make sure it is right for you, and check with graduate students who have previously taken the course to be sure you are aware of the amount of work involved.

As in everything else I teach, the emphasis in EE 382N.19 is on deep understanding of fundamental principles, so that students will be better prepared to carry on original research or take their places on design teams charged with pushing the frontiers of what one can do in (or with) a microprocessor. We expect to accomplish this objective in two ways:

(1) Each student will participate as a member of a design team to complete a substantial design of a cpu for a subset of a commercially available modern microprocessor. We will use Intel's x86 ISA as our challenging starting point. In the prerequisite course, EE 460N (formerly EE 360N), students complete the implementation of a simple microprocessor, the LC-3b. But, in that project, the baseline datapath, state machine, and control signal definitions were all provided. This time you get to start with a clean sheet of paper and design the data path, microsequencer, microprogrammed or hardwired control, microcode or logic, as appropriate, from scratch. You will also design an interface to memory I/O, and select and interconnect all the parts to implement all of the above. The design will be done at the logic gate level, in structural-level Verilog where the design will be concerned with timing issues (propagation delay,

cycle time). You will do all this as a member of a small design team. The design may be an aggressive pipeline, or a more conservative microarchitecture; it is up to you! Our expectation is that each student will come out of this experience more fully appreciating the problems that come up in designing the microarchitecture for a general purpose ISA.

(2) Lectures, in addition to dealing with design issues relevant to the project, will discuss in depth many of the important current hot topics in high performance microarchitecture, and an awareness and appreciation of the field of computer architecture, particularly alternative design styles and implementation tradeoffs. For example, although we will certainly discuss the key current hot topics that make up the latest superscalar microprocessors, we will also spend some time on other microarchitecture ideas that make sense when superscalar is not the right answer. One instance of this is the GPU, which lately is being heralded as "the answer," (regardless of the question) ...and for certain applications it is. We will discuss why that is the case. We will also deal with problems involving instruction supply, data supply, and instruction processing, compile-time/run-time tradeoffs, very aggressive branch prediction, wide-issue processors, in-order vs. out-of-order execution, instruction retirement, etc. Case studies will be taken mostly from current and near-term future microprocessors, although we will examine where relevant some aspects of classical older implementations. Now that just about all chips are multi-core, we will examine various diametrically different approaches as to what a multicore chip should look like. For example, a GPU from Nvidia looks nothing like Cell from IBM, and even less like the heterogeneous chips being proposed today. We will discuss the various design points, in particular the notion that energy consumption is a first class constraint, mollifying what we are allowed to aspire to in the way of performance. Finally, we will study the challenges presented by a chip containing fifty billion transistors operating at a clock frequency of 10 GHz, ...which we expect to be available within the next several years, before Moore's Law finally rests in peace.

Relevance of EE 382N.19 to careers in microarchitecture – This course provides a fundamental body of knowledge useful to graduate students who plan to do PhD research in microarchitecture or plan to seek employment in the microprocessor industry upon completion of their degree.

With respect to PhD research, several major IEEE and ACM conferences deal specifically with research results from this field, including ISCA, MICRO, HPCA, ASPLOS, PACT, and ICS. Several prestigious journals publish research based on the foundation material taught in this course. Notwithstanding some ill-advised pronouncements by some who should know better that computer architecture is dead, there does not appear to be any lessening of interest in this material in the research community.

With respect to the microprocessor industry, companies seek graduates who have the insights acquired from this course. Many major employers of our graduates (Intel, ARM, Apple, Nvidia, IBM, AMD, and Oracle, for example) have an increasing need for graduates who have these insights.

Where I am coming from – Some of the Topics I hope to discuss this semester include the following:

1. Introduction and Focus.
2. ISA tradeoffs.
3. uarch tradeoffs.
4. System tradeoffs.
5. Single thread.
6. Run-time optimizations.
7. Compile-time optimizations.
8. Branch Prediction.
9. Multiple threads.
10. Integer Arithmetic.
11. Floating Point Arithmetic.
12. Cache Coherency.
13. Memory consistency.
14. Measurement methodology and abuses.
15. RISC: A retrospective.
16. Case studies: the microarchitectures of existing processors.
17. Some Current specific issues in Microarchitecture.
18. Security.
19. Machine Learning.
20. A few words on quantum accelerators (an accelerator is not a computer.)
21. Multi-core, Mega-Nonsense.
22. My sense as to critical requirements for the future.
23. A few guest lectures from local industry (more than in the past.)

I will also reserve the last class meeting, at popular request, for what I call a free for all!

We will probably not get to all of these topics for several reasons: (1) there is too much to cover in one semester. (2) Simply "covering" the material is not something I particularly aspire to. Furthermore, we will probably not even cover the topics in the order I have shown, regardless how much I plan to, today.

My objective in our class meetings will be to explore ideas that will be useful to your future research and/or your future work in industry. (Incidentally, my view of research is that if you know the outcome before you start the project, then I am not interested in calling the work "research.") I suspect that many of our class meetings will follow some unintended path as we explore dynamically some issue that comes up.

I have included on the course webpage a link to our computer architecture seminar schedule. I encourage you to attend these seminars, and if a topic surfaces in any of them that we find ourselves interested in, we may end up spending an entire unintended class period discussing that topic. I want you to think critically about what you read, and explore creatively what might be possible. If that causes us to spend three times as long on a topic as we might otherwise if we covered the topic from my notes, it will not make me unhappy. If we get the material from my notes to yours without going through the brains of either of us, that will make me very unhappy.

Lest anyone think this is intended to encourage wild-eyed departures from fundamental knowledge and reasoning, let me assure you that the one thing we always try to do is tie everything we discuss to our grounding in the fundamentals. My hope is to encourage you to combine mastery of the fundamentals, critical reading and analysis, and creative thinking.

CAD Tools – For the project, we will be using a modern set of CAD design tools, provided by Synopsys, which use the Verilog design language. We will provide sufficient introductory material and examples to help you get started with these tools. Mastery of the tools is not an end in itself; on the contrary, the tools are expected to be a means to enhance your productivity in completing the project. **You are enthusiastically encouraged to help each other master the tools**, so that we can all get on with the business of carrying out our designs.

An important caveat about the design project – My experience from teaching this course has been that the design project requires a much larger amount of time to complete than most students expect to be the case at the beginning of the semester. If this semester goes as the ones before it, you will be pleased with what you have accomplished after the term is over. But during the semester, sometimes after a few consecutive sleepless nights, you may wonder what lapse in sanity caused you to sign up. Please consider the time requirement as you organize your workload for the semester.

Meeting Info: The course will meet for three hours of lecture plus one optional 1 1/2 hour discussion session each week. Lectures will be MW from 5 to 6:30pm in ECJ 1.306. The optional discussion session will be conducted by the TA every Thursday until Spring Break, from 5pm until 6:30pm, starting tomorrow. You will find that the discussion sessions are important to your successful completion of the course. The TA is a member of my research group, and has himself taken this course two years ago. If you plan on taking this course, I encourage you to take advantage of the help he will provide. Although it is always the case that all class meetings, both lecture and discussion sessions, are optional (we never take attendance), you are advised that it is probably not a good idea to sign up for this course unless you plan to attend both lectures and discussion sessions. If you want to attend, but have a conflict with the time, please let me know.

TA: Aniket Deshmukh (a.deshmukh@utexas.edu).

Course Home Page: <http://www.ece.utexas.edu/~patt/20s.382N>

Textbook: There is no required text. References will be suggested where appropriate, depending on the topic. I expect to provide handouts on additional material when I feel that is useful. Also, some of the lectures will use transparencies. In those cases, you will be provided with copies of the transparencies. From time to time, relevant material will be available for downloading on the course home page.

Prerequisites: Satisfactory completion of courses covering the material of EE 316 and 460N (or 360N) with a grade of B, or consent of the instructor. If you are concerned about whether you are prepared to take 382N.19, please do not hesitate to meet with me privately.

Homework policy: Homework consists of three types: (1) problem sets, scheduled at the beginning of the semester to get the student ready for the major term project of the course, (2) individual problems assigned in class from time to time as I feel appropriate to test your comprehension of a point or to get you to take the point we are discussing to the next level, and (3) the major design project.

Quizzes and Exams: There will be two exams, a written exam in class on March 25, and an oral exam on April 2 or 3. There will be no final exam in this course. Also, I do not expect to give a make-up exam except in rare and well-documented circumstances.

The written exam will be closed book, with two exceptions: (1) A student may bring into the exam three sheets of paper on which he/she may have written anything he/she wishes. All three sheets must be original sheets (not printed nor xeroxed) in the student's own handwriting. (2) Students may also bring into the exam any handouts that have been expressly permitted by the instructor prior to the exam.

The oral exam will be an individual exam in my office, EER 5.802, at a time suitable to fit each student's schedule. It will be 30 minutes in length. Students will be free to bring whatever reference material they wish into the exam.

Important dates:

Problem set 1a is due at the beginning of class, January 29.
Problem set 1b is due at the beginning of class, February 5.
Problem set 2 is due at the beginning of class, February 19.
Problem set 4 is due at the beginning of class, February 26.
Initial Design Reviews with each Project Team, March 11,12,13.
March 16-21, no class. Spring break.
Written exam, (exam 1) in class, March 25.
Oral exams (exam2) in EER 5.802, April 2,3.
No class, April 29.
Final project design reviews, May 7,8.
Final project report due in EER 5.802, 10pm, May 15.

Grading mechanics: Three major items will contribute to your grade in this course: the major design project, scores on the two mid-term exams, and homework and problem sets. You will note below that I have allocated 6% of your grade to "other" to give the TA and me some flexibility to include our subjective evaluation of your performance in the assignment of a final grade.

The items will be weighted, approximately as follows:

- exams, 42%
- major design project, 42%
- homework, problem sets, etc., 10%
- other, 6%

Final Exam: There will be no final exam in this course. In lieu of a final exam, you will have a final design review of your design project, and you will submit a final design project report. The final design project report will be due in EER 5.802 on May 15, 2018, which is the date specified for a final exam in this course if we were to have a final exam.

Cheating: Students are encouraged to help each other master the Synopsis tools. Students are also encouraged to work together BEFORE both exams to study for the exams. In fact, forming study groups is a good way to better master the material discussed in class. However, in the case of the oral exam, students are NOT permitted to discuss anything about the oral exam with ANYONE once the oral exam period has started (April 2), until the last student has taken the oral exam. In the case of the written exam, no aid may be given or received during the exam. In the case of the design project, students are encouraged to work together to learn/understand the design tools, but not to work together on their individual design assignments (Problem sets 1,2, and 4). Students must work together within their design project group to complete the term project. If you need help, you are welcome to see the Instructor or the TA. Receiving help from any source other than the Instructor, TA, or someone designated by the Instructor or TA is not permitted.

Any violation of the above paragraph constitutes academic dishonesty and will be directed to the Dean of Students for Disciplinary Action. If you are not clear as to what is permissible and what is not, please ask the instructor or the TA. Failure to ask ahead of time, and later invoking the statement, "I thought it was ok to do" does not constitute an acceptable excuse. If you cheat, you violate the soul of the University, which I take very seriously, and will not compromise.

Course evaluation: The MEC Common Evaluation form will be used to evaluate the instructor in this course.

Additional details:

I am asked to remind you to consult the University policy on the deadlines for adding/dropping courses. If you need help with this, please check with me or the TA.

Allegations of Scholastic Dishonesty will be dealt with according to the procedures outlined in Appendix C, Chapter 11, of the General Information Bulletin, <http://www.utexas.edu/student/registrar/catalogs/>.

The University of Texas at Austin provides, upon request, appropriate academic adjustments for qualified students with disabilities. For more information, contact the Office of the Dean of Students at 471-6259, 471-4241 TDD, or the College of Engineering Director of Students with Disabilities, 471-4321.

Finally, for those of you who decide to continue in this course, Good Luck. I hope you find the experience an important part of your computer engineering education. I also hope you have a good time doing it.