

ECE 18-845: Internet Services

David R. O'Hallaron

Dept of Electrical and Computer Engineering
and School of Computer Science
Carnegie Mellon University
Spring 2002

1 Organization

Instructors

David R. O'Hallaron, WeH 8125, x8-8199, droh@cs.cmu.edu

TA: Saumitra Das, TBD, TBD, saumi@cmu.edu

Administrative Assistant

Barb Grandillo, WeH 8212, x8-7550, bag@cs.cmu.edu

Class

Tue Thu 12:30-2:30, Doherty Hall A317.

Resources

Web page: www.ece.cmu.edu/~ece845

Newsgroup: cmu.ece.class.ee845

Course directory: </afs/ece/class/ece845>

2 Objectives

As recently as 10 years ago, a busy network server might be expected to handle thousands of requests. And if it went down for awhile, few people were seriously inconvenienced. Today it's a much different story.

With the explosion of the World Wide Web, busy servers are expected to handle millions of requests a day, 24 hours a day, 7 days a week. Slow response times and downtimes cost money, jobs, and business.

Our aim in 18-845 is to understand the research issues involved in providing Internet services that are scalable, secure, and highly available. Examples of such services include Web servers, mail servers, search engines, proxy caches, and online auction systems. We will study the techniques for building scalable, secure, and highly available services, such as concurrent servers, clustering, and caching, trying to gain a deep understanding of the advantages and disadvantages of the various approaches. We will place particular emphasis on how to measure and evaluate the effectiveness of various approaches.

3 Textbook

There is no required textbook for 18-845. The basic aspects of network programming will be covered in lecture, and complete documentation is available in the Unix man pages. There are also numerous tutorials on the Web. However, if you have a serious interest in network programming, then you will want to acquire the latest version of the classic network programming text by Stevens. The coverage is comprehensive and motivated with lots of useful examples. It is an excellent reference for your course projects and your future work.

W. Richard Stevens, *Unix Network Programming: Networking APIs*, Volume 1 (Second Edition), Prentice Hall PTR, 1998.

W. Richard Stevens, *Unix Network Programming: Interprocess Communications*, Volume 2 (Second Edition), Prentice Hall PTR, 1999.

We'll also supplement with some draft chapters on systems programming from the 15-213 textbook:

R. Bryant and D. O'Hallaron, *Computer Systems: A Programmer's Perspective*, Prentice Hall, to appear in 2002.

4 Course Organization

18-845 is a computer systems research course, and we will approach the material as researchers. This means diligent investigation of existing work, striving for deep understanding of ideas and their advantages and disadvantages, careful measurement and quantitative evaluation, and sharing our ideas with others in discussions, papers, and oral presentations.

Your participation in the course will involve five forms of activity. The percentage of your final grade is shown in parentheses. Note that there are no exams.

1. Attending and participating in the classroom discussions (5%).
2. A 3-week individual project (15%).

3. Leading classroom discussions (20%).
4. Research paper critiques (20%).
5. An 8-week group project paper and poster presentation (40%).

Classroom participation

A research course like this is only as good as the intellectual energy and involvement of the students. It is not the kind of course where you can show up sporadically, cram for a few exams, and expect to learn much. **You are expected to attend every class and participate actively in the discussions.**

Individual project

The first part of the course will be some review classes taught in a traditional lecture format by your instructor, covering basic concepts in network programming such as IP, internetworking, TCP and UDP, basic sockets, TCP and UDP sockets, processes, threads, signals, shared memory, record and file locking, synchronization, server design alternatives, HTTP servers, DNS, and ICMP. How much of this material we cover depends on the experience level of the class.

This part of the course culminates with a 3-week individual project, assigned by your instructors. You will work by yourself on this project. The purpose is to get everyone up to the same basic technical level.

Leading classroom discussions

The bulk of the course will be devoted to reading and discussing research papers. Students will take turns leading the discussions on particular papers and topics. In the default case, your instructors will identify the topic and an initial set of papers for that topic and will assign a student to lead the discussion on the topic. If you have a particularly interesting topic of your own to talk about, then we can arrange that as well.

As the discussion leader, your job is provide any necessary technical background (your instructor can help you here), summarize the problem being addressed and the key ideas that address the problem, compare and contrast the different approaches, and generally lead the discussion with the rest of the class.

Paper critiques

Students who are not leading the discussion that day should prepare a 1-page critique of the paper denoted by a "*" on the course home page. The critique should include two parts:

- A summary of the key ideas in the paper.
- 5 discussion questions.

Examples of good discussion questions include aspects of the paper you need clarification on, questions about claims made by the authors, or "what if" questions about situations that aren't handled well by the

ideas in the paper. The purpose of the critiques is to ensure that you have carefully read and actively thought about one of the papers being presented that day, so that you are armed with good questions to stimulate the discussion.

Handing in your critiques: Bring a hardcopy of your critique (no email) to class and give it to Saumitra. He will grade it and return it the next class.

Group project

This part of the course will culminate in a group project of your choice. You may work in groups of 1 or 2. There is a lot of leeway with the choice of group projects. Your instructor will work with you individually to come up with a suitable project.

At the end of the term, each group will hand in a paper (10 pages or so) describing their project, and will present your results to the rest of the class in a poster session.

5 Getting help

You are always welcome to visit or phone your instructors if you have questions, problems, comments, or just want to chat.

Saumitra Das, the course TA, is your primary point of contact for questions related to specific programming and systems issues.

Prof. O'Hallaron normally works with his office door open, and welcomes visits from students whenever his door is open. However, if his door is closed, then he is busy with a meeting or a phone call and should not be disturbed.

We will use the Web as the central repository for all information about the class. We have also set up a news group for this class. This group will be used by members of the teaching staff to post announcements and clarifications. You may also post to this group to make queries.

6 Cheating

In general, the only thing you are not allowed to do is copy code from another student or group. If you do this, you will automatically fail the course. Otherwise, we encourage to talk to and help each other.

7 Class Schedule

The up-to-date course schedule will be maintained on the Web page.