Analysis and Prediction of the Dynamic Behavior of Applications, Hosts, and Networks

Syllabus

Web Page

http://www.cs.northwestern.edu/~pdinda/predclass-s03

Instructor

Peter A. Dinda
1890 Maple Avenue, Room 338
847-467-7859
pdinda@cs.northwestern.edu
Office hours: 2-4pm on Thursdays or by appointment
I am always interested in talking to you about your projects and will make time as needed to help you with them.

Location and Time

1890 Maple Avenue, Computer Science Department,
Bliss Conference Room (340), WF 10:00-11:20

Prerequisites

Required	CS 311 or equivalent data structures course
Required	Calculus and some linear algebra
Recommended	Basic Statistics and Probability
Recommended	Familiarity with Matlab, Maple, Mathematica, or S-Plus

Permissions

CS Graduate Students	No special permissions required (see Caryn)
Others	Permission required (see me)

Objectives, framework, and caveats

Important, measurable properties of distributed computing environments, such as application workloads, network bandwidth, and host load, vary dynamically over time. These properties drastically affect the performance of applications running on these systems. However, if applications are made aware of this dynamicity, ideally in the form of statistically meaningful predictions, they often can adapt their behavior to nonetheless provide consistent high performance. In effect, such predictions enable them to exploit the degrees of freedom available in a distributed computing environment to ameliorate the dynamicity found there. In addition, the analyses upon which prediction is based can tell us valuable and fascinating things about the nature of computer systems and networks qua natural systems.

This course focuses on how we can measure, analyze and predict the dynamic behavior of distributed computing environments and their applications. Broadly speaking, this is an area of **performance analysis** that includes **workload characterization and modeling** with the goal of providing **online** modeling and prediction for the direct benefit of applications.

For the most part, we will use an approach based on probability, statistics, and signal processing, although we will also touch on queuing theory and other approaches. The course has three objectives. The first objective is for you to learn some of the theory behind measurement, analysis, and prediction. The second objective is for you to learn how this theory has been applied to computer systems in the past, and what fascinating new things were learned. The final objective is for you to become comfortable in applying the theory to his or her own data and systems, and in evaluating other methods for studying your data.

We will generally read about 2 papers or equivalent materials for each session, covering fundamental ideas and important recent results. Each paper will be formally presented to the group by a student and then discussed in a round-table manner. In parallel with the readings, students will be strongly encouraged to apply what they are learning by using analytical tools such as Matlab, Maple, Mathematica, S-Plus, Prophet, and others to study real data, ideally data that they themselves, as well as the instructor, are interested in. Students will also be encouraged to use and extend on-line measurement and prediction systems such as RPS, Netlogger, Remos, and NWS. Finally, each student will complete a quarter-long project in which they will apply what they learn to an area that interests them. The goal of these investigations will be to produce interesting new research results, perhaps even some that will lead to publications.

This is a graduate course and all students in it will be treated like graduate students. I will assume that you are interested in this material, that you can motivate yourself to learn about it, and that you will not be afraid to venture into uncharted territory (i.e., do research). The undergraduate section will differ primarily in that the expectations for the project will be *slightly* lower.

Reading

While there is no textbook for the course, the following book is required and is available at the campus store and via the web:

Raj Jain, *The Art of Computer Systems Performance Analysis*, Wiley, 1991.

This book provides background reference material for many of the tools we will use in this class, including measurement techniques and presentation, probability and statistics, experiment design, simulation, and basic queuing theory. While it's not perfect, it is a very useful reference to have on your shelf even beyond this course. I will provide photocopied supplements for signal processing, time-series analysis, and other areas.

The vast majority of the reading for the course will consist of original research papers that report on studies of real computer systems, networks, and their applications. In addition to these papers, there are a number of books, papers, and other resources that are very helpful in understanding the theoretical and statistical analysis techniques that were applied to produce their results. We will read some of this material in tandem with the research papers so that you will be able to generalize the techniques and learn to apply them to your own work. Most of what we will read is available on the web, and I will hand out photocopies of what is not.

Project

Over the course of the quarter, you will apply what you learn to a project of your choice, and then document your project in a high quality paper and open presentation. Project topics will be chosen in consultation with me. Projects may be done individually or in groups. Project complexity and expectations will be tied to group size. There are specific projects that I have in mind that would be well-geared to groups of two or three.

The expectation for graduate students is that the project will be quality work that the students would not be embarrassed to submit to a workshop. The two previous iterations of this course, totaling about 14 projects, have resulted in three publications, two publications in current submission, two long-term research projects, and several conference-level drafts that we have not had the time to push to publication.

The expectation for undergraduates is that the project be something they would be proud to list on their resumes, although students are encouraged to aim high.

All projects will be presented at a public colloquium.

Example project ideas are listed in a separate handout. Because of the high expectations placed on the project, it is vital that you choose to work on something that interests you deeply and that I can advise strongly.

Exams

There will be no exams

Grading

50 % Project

10 % Project paper and presentation
20 % In-class paper presentations of papers
20 % General classroom participation

Schedule

Lecture	Date	Topics	Theory Readings	Application Readings
1	4/2	CANCELLED		
2	4/4	Mechanics, motivation, overview, calibration, probability/stats review	Jain 1-4, 6, 10, 11	
3	4/9	Distributions, summaries, estimates, and their implications for the web	Jain 12-14	Myers (32), Arlitt (35), Smith (40)
4	4/11	Distributions, summaries, estimates, and their implications for hosts		Eager (2); Mutka (3),
5	4/16	Distributions, summaries, estimates, and their implications for hosts **PROJECT PROPOSAL DUE**		Harchol-balter (4), Leland (1)
6	4/18	Signal Analysis	Oppenheimer	Dinda (5)
7	4/23	Heavy tails, power laws, and self-similarity	Bassingthwaighte	Willinger (11), Willinger (12), Gao (99)
8	4/25	Self-similarity and non-stationarity		Paxson (10), Cao (24)
9	4/30	Self-similarity and power laws in network topology		Faloutsos (17), Chen (89), Gkantsidis (91)
10	5/2	Self-similarity in video and other applications **PROJECT UPDATE DUE**		Garrett (30), Loguinov (90)
11	5/7	Wavelets	Wavelet intro	Ribeiro (19)
12	5/9	Chaotic dynamics	Abarbanel (74)	
13	5/14	Application Behavior		Bavier (36), Kapadia (53), Lai (87)
14	5/16	Application Behavior		Abdelkhalek (43), Henderson (92)

15	5/21	Application Behavior		Bhola (45), Keeton (44)
16	5/23	Predicting hosts **PROJECT UPDATE DUE**	Box	Wolski (6), Dinda (7), Devarakonda (48)
17	5/28	Predicting networks		Wolski (15), Basu (13), Sang (23)
18	5/30	Anomaly detection		Dasgupta (75), Barford (97), Hofmeyer (57)
19	6/4	Zipf, file system prediction		Breslau (33), Kroeger (50), Douver (54)
20	6/6	Network paths and links		Spring (82), Coates (86)
Spares		Wireless networks		Balachandran (25), Eckhardt (76)
Spares		Systems		Lowekamp (59), Stemm (63), Dusseau (94)
Spares		Systems		Wolski (62), Dinda (61)
	6/9 – 6/13	**Project Papers Due, Project Presentations (TBA)**		