Human-directed Approaches to Computer Systems Problems

Syllabus

Web Page

http://www.cs.northwestern.edu/~pdinda/hdsys

Instructor

Peter A. Dinda
Technological Institute, Room L463
847-467-7859
pdinda@northwestern.edu
Office hours: Thursdays, 2-4pm or by appointment
I also respond quickly via email
I am often available via Google Talk (pdinda at gmail dot com)
There is a course newsgroup (see below)

Location and Time

We will meet in TECH MG28 for the first week

For the next 9 weeks, we will meet in FORD 3-340 (conference room)

Tuesdays and Thursdays, 9:30-10:50am

Prerequisites

There are no formal prerequisites for this course at this time, although the following are noted:

- Basic knowledge of computer systems, to the level of EECS 213 is helpful. We will review the hardware/software stack in the first week of classes
- Some experience with human-computer interaction, such as in EECS 330, is helpful.
- Programming skills in an imperative language such as C, C++, Perl, or Python will helpful. Programming will be needed for some of the possible projects.
- Some familiarity with GUI programming on Windows, Unix/Linux, Mac, Java, or a toolkit is helpful. GUI programming will be needed for some of the possible projects.
- Familiarity with basic statistics is assumed.

As this is the first iteration of this course, I will try to adapt the course to the skills of the students.

Readings

Readings for the course will be in the form of research papers and other materials. For the most part, these materials are available online. I will hand out paper copies of those that are not. A separate reading list is provided. It also lists several reference texts that may be helpful.

Objectives, framework, philosophy, and caveats

Recently, human-centered approaches have made significant inroads well outside of the traditional area of human-computer interaction. Examples include CAPTCHAs in computer security, and "distributed human computation" for solving AI-complete problems such as in the Google Image Labeling Game.

This research-oriented course, suitable for graduate students and advanced undergraduates, explores the application of human-centered approaches to solving problems in systems software, networks, and architecture. The design and engineering decisions made at these layers of computing are typically considered quite independently of end-users. However, these layers have dramatic effects on the user experience. Work at Northwestern and elsewhere demonstrates there is massive potential in applying human-centered approaches to computer systems problems. This course will explain the state of the art, how to evaluate it, and will give students the opportunity to expand it.

The overall thinking and motivations behind this course are given in much more depth in the following position paper:

- P. Dinda, G. Memik, R. Dick, B. Lin, A. Mallik, A. Gupta, S. Rossoff, The User In Experimental Computer Systems Research, Proceedings of the Workshop on Experimental Computer Science (ExpCS 2007), June, 2007.

We will explore the following topics in this course:

- Review of computer systems software and hardware, and their impact on the user experience
- Human-centered computing as applied to systems
- Results outside of systems: CAPTCHAs, distributed human computation, games with a purpose, etc
- Results within systems: user-driven power management, user-driven scheduling, speculative remote display, user comfort with resource borrowing, distributed scheduling games, others
- How interface issues in systems are different from those in applications
- User-driven specification and user-driven search

- Role of learning
- Open applications of concepts
- Open research questions

When a student completes this course, he or she should be able to:

- Understand the concept of human-centered computing as applied to computer systems
- Understand the aspects of computer systems that impact the user experience
- Formalize computer systems problems that may be aided through human-centered computing
- Give examples of user-driven specification and user-driven search
- Understand the issues of interface design that arise in systems that are different from or in addition to those in applications
- Design and develop human-centered approaches to systems problems
- Understand how to evaluate, and how to judge the evaluation of, human-centered approaches to systems problems
- Have familiarity with the existing literature on the application of human-centered computing to computer systems, and how to evaluate it
- Have experience in a research project in the area, and how to document it in an effective conference paper and presentation.

This course is structured as a graduate research seminar, with undergraduates welcome. It combines extensive reading, presentation, and discussion of original research papers and a quarter-long research project in the area that leads to a conference-style paper and presentation.

Almost all of the readings for the course will be in the form of research papers, and similar materials. We will generally read several papers or equivalent materials each week, 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. Before each class, each student is expected to write up his/her comments on the course discussion group for access by all class members.

Student projects, which can be done individually or in groups, should strive to be original research, conducted with guidance from me. Each project will lead to a conference paper and a presentation. I want students to be very ambitious!

Course Discussion Group

The course web page includes a sign-up form and URL for the course discussion group. This is a Google Group that is accessible only to the members of the class.

Before each class, students are expected to read the papers, and post a summary and comments to the course discussion group. These posts, which will be accessible to everyone in the class, are required, and will form part of your grade.

Tools

In addition to standard resources (TLab, VLab, Wilkinson Lab, Pocket PCs), students will have access to other tools from my research lab on an as-needed basis, and I will purchase other tools if needed.

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 conference-style paper, similar to those you'll be reading, and a 20-25 minute presentation. If you're an undergraduate or a new graduate student, this process will expose you to doing research, and documenting its results in a professional way. It is not uncommon for good work in a graduate class like this to lead to a paper submission to a high quality workshop or conference, and/or to a longer-term research project.

Students can work individually or in groups. I will expect a project proposal, and weekly progress reports in addition to the final paper and presentation. Ideally, we will meet once a week to discuss your project.

Exams

There will be no exams

Grading

40 % Project, **including proposal and weekly progress reports**
20 % Project paper and presentation
20 % In-class paper presentations of papers
20 % General classroom participation **and comments on discussion group**

Schedule

Week 1 Introduction

Tuesday, 1/8 Introduction
Review of how computer systems work

Efficient Reading of Papers in Science and Technology:
http://www.cs.columbia.edu/~hgs/netbib/efficientReading.pdf (very short!)

EECS 213 textbook (Bryant and O'Hallaron, *Computer Systems: A Programmer's Perspective)* may be handy if you've never taken a systems course before

J. Salzer, et al, *End-to-end Arguments in System Design*, ACM TOCS 2:4, 1984.

Also take a look at Lampson, *Hints for Computer System Design* if you haven't read it before.

Thursday, 1/10 Introduction - the overall ideas in the class

P. Dinda, G. Memik, R. Dick, B. Lin, A. Mallik, A. Gupta, S. Rossoff, *The User In Experimental Computer Systems Research,* Proceedings of the Workshop on Experimental Computer Science (ExpCS 2007), June, 2007.

Chapter 1 of B. Lin, *Human-driven Optimization*, Doctoral Dissertation, Technical Report NWU-EECS-07-04, Department of Electrical Engineering and Computer Science, Northwestern University, July, 2007

Week 2 Outside of Systems

Tuesday, 1/15 CAPTCHAs

L. von Ahn, M. Blum, N. Hopper, J. Langford, *CAPTCHA: Using Hard AI Problems for Security*, Eurocrypt 2003.

Also be sure to check out captcha.net and recaptcha.net

Thursday, 1/17 Human Computation

L. von Ahn, *Games With A Purpose*, IEEE Computer, 39:6, June 2006.

L. von Ahn, L. Dabbish, *Labeling Images with a Computer Game*, CHI 2004.

L. von Ahn, *Google Talk on Human Computation*, Google Video (in class, probably)

Week 3 Measuring the User

Tuesday, 1/22 Measuring the Desktop User

A. Gupta, B. Lin, P. Dinda, *Measuring And Understanding User Comfort With Resource Borrowing*, Proceedings of the 13th IEEE International Symposium on High Performance Distributed Computing (HPDC 2004)

C. Reynolds, *The sensing and measurement of frustration with computers*, Master's thesis, MIT Media Lab, 2001 (skim)

A. Komatsubara, *Psychological upper and lower limits of system response time and user's preference on skill level*, HCI International 1997

PROJECT PROPOSAL DUE AT END OF THIRD WEEK

Thursday, 1/24 Measuring the Web User

D. Olshefski, *Inferring Client Response Time at the Web Server*, SIGMETRICS 2002.

E. Kiciman, B. Livshits, *AjaxScope: A Platform for Remotely Monitoring the Client-side Behavior of Web 2.0 Applications*, SOSP 2007.

Week 4 Modeling the User

Tuesday, 1/29 Workload modeling and prediction

S. Bhola, M. Ahamad, *Workload modeling for highly interactive applications*, SIGMETRICS 1999. Extended version is Technical Report GIT-CC-99-2, College of Computing, Georgia Tech. Read the extended version.

B. Davison, *Learning Web Request Patterns*, Book Chapter in *Web Dynamics: Adapting to Change in Content, Size, Topology and Use*, 2004.

Thursday, 1/31 Workload modeling and prediction

B. Davison, H. Hirsh, *Predicting Sequences of User Actions*, AAAI-98/ICML-98 Workshop on Predicting the Future: AI Approaches to Time Series Analysis.

P. Gorniak, D. Poole, *Predicting Future User Actions by Observing Unmodified Applications*, AAAI 2000

Week 5 Speculation

Tuesday, 2/5 Speculative execution and display

C. Zilles, *Increasing Interactivity By Predicting User Actions,* ASPLOS 2004 "Wild and Crazy Ideas" Session (WACI)

J. Lange, P. Dinda, S. Rossoff, *Experiences with Client-based Speculative Remote Display*, DRAFT

Thursday, 2/7 (likely out of town at V3VEE board meeting)

SLACK TIME FOR NOW

Week 6 User-driven Scheduling

Tuesday, 2/12 Resource Management

J. Sousa, R. Balan, V. Poladian, D. Garlan, M. Satyanarayanan, *Giving users the steering wheel for guiding resource-adaptive systems*. Tech Rep. CMU-CS-05-198, Department of Computer Science, Carnegie Mellon University, December 2005.

B. Lin, P. Dinda, *Towards Scheduling Virtual Machines Based On Direct User Input*, Proceedings of the 1st International Workshop on Virtualization Technology in Distributed Computing (VTDC 2006), November, 2006. Read the full version, Northwestern Technical Report NWU-EECS-06-07.

Thursday, 2/14 Resource Management

Y. Endo, and M. Seltzer, *Improving Interactive Performance using TIPME*, SIGMETRICS 2000.

Chapter 7 of B. Lin, *Human-driven Optimization*, Doctoral Dissertation, Technical Report NWU-EECS-07-04, Department of Electrical Engineering and Computer Science, Northwestern University, July, 2007

Week 7 User-driven Power Management

Tuesday, 2/19 Power management

K. Flautner., T. Mudge, *Vertigo: Automatic Performance-setting for Linux*, OSDI 2002.

B. Lin, A. Mallik, P. Dinda, G. Memik, R. Dick, *Power Reduction Through Measurement and Modeling of Users and CPUs: Summary*, Proceedings of

ACM SIGMETRICS 2007, June, 2007. Read the full technical report instead: NWU-EECS-06-11.

Thursday, 2/21 Power managment

A. Mallik, J. Cosgrove, R. Dick, G. Memik, P. Dinda, *PICSEL: Measuring User-Perceived Performance to Control Dynamic Frequency Scaling*, ASPLOS 2008.

L. Yan, L. Zhong, N. Jha, *User-perceived Latency based Dynamic Voltage Scaling for Interactive Applications*, DAC 2005

Week 8 Evaluating User-based Systems

Tuesday, 2/26 Example user-oriented evaluations of systems

B. Schmidt, et al, *The Interactive Performance of SLIM: A Stateless Thin-client Architecture*, SOSP 1999.

Y. Endo, Z. Wang, J. Chen, M. Seltzer, *Using latency to evaluate interactive system performance*, OSDI 1996.

Thursday, 2/28 Challenges in user studies

Reading: TBD

Week 9 Security and Visualization

Tuesday, 3/4 Visualizing hashes and keys

R. Dhamija, *Hash Visualization in User Authentication*, CHI 2000.

A. Perrig, D. Song, *Hash Visualization: A New Technique to Improve Real-World Security*, CRYPTEC 1999.

Thursday, 3/6 Visualizing computer systems and security

P. Ren, Y. Gao, Z. Li, Y. Chen, B. Watson, *IDGraphs: Intrusion Detection and Analysis Using Stream Compositing*, IEEE Computer Graphics and Applications, 26:2, March-April, 2006. (whole issue is about visualization in cybersecurity)

Genetic Art for Intrusion Detection Project: http://ga-ids.cs.northwestern.edu

Week 10 Brainstorming/Crazy Ideas

Tuesday, 3/11
Thursday 3/13

Is there a grand unified abstraction for what users want?
Do emotions matter for empathic systems software?
Your questions here and paper suggestions here…

B. Sabata, et al, *Taxonomy for QoS Specifications*, WORDS 1997

M. Whang, *The Emotional Computer Adaptive To Human Emotion*, Probing Experience From Assessement of User Emotions and Behavior to Development of Products, J. Westerink, et al, eds, 2008.

R. Mandryk, M. Atkins, *A fuzzy physiological approach for continuously modeling emotion during interaction with play technologies*, Journal of Human-Computer Studies, 65 (2007), 329-347.

PROJECT DUE LAST DAY OF FINALS WEEK

Project Presentations In Finals Week Final Exam Slot