

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

## Project Ideas

### Comments

The following are various ideas that I think could lead to cool projects in this course. You are not obligated to use any of these ideas. The most important thing is that the project you choose will hold your interest for at least the rest of the quarter.

In any case, you should choose or design a project in consultation with Peter. You must formally propose your project in a one page document. Over the course of the quarter, you must also hand in two one page project updates. At the end of the quarter, you must hand in a 10 page (SIGMETRICS 2 column format) paper documenting your project and its results, and give a 20 minute public presentation of your work.

### Example Ideas

1. Efficient monitoring of virtual machines. Given the monitoring information in the host operating system (e.g., /proc on Linux or the performance counters in Windows), can you reconstruct the monitoring information in a guest operating system?
2. Study user irritation. Implement the Panic Button or QoS joystick idea described here: <http://www.cs.northwestern.edu/~knop/wtpb.html>. The implementation would include network, disk, and memory resources as well as CPU and would provide other parameterized patterns of usage. Do a user study to derive “irritation curves” for actual end users as a function of the parameters.
3. Study intrusion detection based on statistical signal processing of performance data.
4. Develop techniques to use Diffusion (<http://www.cs.northwestern.edu/~plab/Diffusion>) to better measure a network and disseminate results.
5. Read about non-periodic sampling techniques and how signals can be reconstructed. Implement one technique.
6. Characterize the workload of an interactive application such as a game or a media authoring tool.
7. Characterize the workload of an interactive multiplayer web game. Rob Thomas has offered Settlers of Catan for this purpose.
8. Apply nonlinear time series analysis to predict network traces.
9. Study the errors introduced by “average over a window” sampling of networks.
10. Develop a mirror selection algorithm and evaluate it against the Myers web traces or other traces.

11. Write a survey paper on chaotic dynamic approaches to signal prediction. (survey papers shall be considerably more in-depth than project reports)
12. Try applying genetic programming to time series prediction.
13. Develop a Kahlman prediction filter for RPS and evaluate its performance.
14. Develop an Innovations algorithm-based implementation of RPS's prediction filters and evaluate its performance.
15. Study my host load traces using a new technique, such as wavelets or Abarbanel's methodology for chaotic dynamics
16. Instrument an interactive windows program and perform a user study to determine how its resource demand varies over time.
17. Study web client traces to characterize how users interact with the web.
18. Collect LAN or wireless network traces within the CS department and analyze them
19. Study file access patterns using the Kroeger CODA traces from UCSC
20. Write a sensor for the Linux /proc filesystem and integrate it into RPS
21. Model the data generated from Linux /proc as a multivariate time series
22. Model the data from the Windows performance counters as a multivariate time series.
23. Create a plausible generative model for host characteristics (memory, disk, cpu) based on location in the network, etc.
24. ...
- 25.