Analysis and Prediction of the Dynamic Behavior of [Users,] Applications, Hosts, and Networks

Reading List

Note: We will not read all of these papers in class. They are included so that you can see a broad range of work. The syllabus is the final word on the specific papers that we shall read in class.

Most of these papers are available from the web (use http://www.google.com and http://citeseer.nj.nec.com to find them. I will make photocopies of the older, non-web papers available as needed.)

Books and Collections

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

- This book covers most common areas of performance analysis. It is perhaps the one performance analysis book that belongs on everyone's bookshelf.
- This is a required book for this course

Larry Golnick, et al, *The Cartoon Guide To Statistics*, 1994.

- A very readable introduction to basic probability theory and classic parametric statistics.
- If you've never seen stats/probability before, this is a good place to start, but you should talk to me about whether this course is appropriate for you.

StatSoft, Inc, *The StatSoft On-line Statistics Textbook*, http://www.statsoft.com/textbook/stathome.html, 2000.

- An excellent reference book and introduction to many different areas of modern statistics.

Alan V. Oppenheim, et al, *Signals and Systems*, 1983.

- Good introduction to linear systems theory
- You will read portions of this book.

Alan V. Oppenheim, et al, *Discrete-time Signal Processing*, 1993.

- Good book on this topic.

Benjamin Kuo, *Control Systems*, 1988.

- Good introduction to control systems theory.

G.E.P. Box, et al, *Time Series Analysis: Forecasting and Control*, 1994.

- The classic text on linear time series analysis.

Leonard Kleinrock, *Queuing Systems, Volumes I and II*, 1976.

- The classic text on queuing theory.

Henry Abarbanel, *Analysis of Observed Chaotic Data*, 1996.

- How to use concepts from chaotic dynamics to study data and systems.

Benoit Mandelbrot, *The Fractal Geometry of Nature*, 1988.

- The seminal book on this topic

.

Hosts: Process Behavior

1. W. Leland, and T. Ott, *Load-balancing heuristics and process behavior*, SIGMETRICS ’86.
2. D. Eager, et al, *The limited performance benefits of migrating active processes for load sharing*, SIGMETRICS '88.
3. M. Devarakonda and R. Iyer, *Predictability of process resource usage: a measurement-based study on UNIX*, IEEE Transactions on Software Engineering, 15:12, 1989.
4. M. Harchol-Balter, A. Downey, *Exploiting process lifetime distributions for dynamic load balancing*, SIGMETRICS '96.
5. S. Kleban, et al, *Hierarchical Dynamics, Interarrival Times, and Performance*, SC 2003.

Hosts: Availability, Load, and Power

6. M. Mutka and M Livny, *The available capacity of a privately owned workstation environment*, Performance Evaluation 12:4, July 1991.
7. P. Dinda, *The statistical properties of host load*, Scientific Programming, 7:3,4, 1999. (Also available as CMU Technical Report CMU-CS-TR-98-175.)
8. R. Wolski, et al, *Predicting the CPU availability of time-shared Unix syst*ems, HPDC ‘99.
9. P. Dinda and D. O’Hallaron, *Host load prediction using linear models*, HPDC ’99 (journal version appears in Cluster Computing, summary in SIGMETRICS 2001)
10. P. Dinda, *Online Prediction of the Running Time of Tasks*, HPDC 2001, (journal version appears in Cluster Computing.)

11. M. Knop, et al, *Windows Performance Monitoring and Data Reduction Using Argus*, SHAMAN 2002 (**This paper got its start in as a project in this course)**
12. T. Li, et al, *Run-time Modeling and Estimation of Operating System Power Consumption*, SIGMETRICS 2003.
13. L. Yang, et al, *Conservative Scheduling: Using Predicted Variance to Improve Scheduling Decisions in Dynamic Environments*, SC 2003.

Networks: Topology and Routing

14. V. Paxson, *End-to-end routing behavior in the Internet*, IEEE/ACM Transactions on networking, 5:5, 1997.
15. M. Faloutsos, el al, *On power-law relationships of the Internet topology*, SIGCOMM '99.
16. N. Duffield and M. Grossglauser, *Trajectory sampling for direct traffic observation*, SIGCOMM '00.
17. H. Tangmunarankit, et al, *Network Topology Generators: Degree-based versus Structural*, SIGCOMM '02.
18. O, Maennel and A. Feldmann, *Realistic BGP Traffic for Test Labs*, SIGCOMM '02.
19. Q. Chen, et al, *The Origin of Power-Laws in Internet Topologies Revisited*, INFOCOM '02.
20. N. Spring, et al, *Measuring ISP Topologies With Rocketfuel*, SIGCOMM '02
21. M. Coates, et al, *Maximum Likelihood Network Topology Identification from Edge-based Unicast Measurements*, SIGMETRICS '02.
22. C. Gkantsidis, et al, *Spectral Analysis of Internet Topologies*, INFOCOM '03.
23. C. Gkantsidis, et al, *Conductance and Congestion in Power Law Graphs*, SIGMETRICS 2003.
24. N. Spring, et al, *The Causes of Path Inflation*, SIGCOMM 2003.
25. L. Li, *A First-principles Approach to Understanding the Internet's Router-level Topology*, SIGCOMM 2004. (*)

Networks: Links, Paths, And Their Traffic

26. V. Paxson, and S. Floyd, *Wide-area traffic: The failure of Poisson modeling*. {IEEE/ACM} Transactions on Networking. 3:3, June 1995.
27. W. Willinger, et al, *Self-similarity in high-speed packet traffic: Analysis and modeling of ethernet traffic measurements*, Statistical Science 10:1, January 1995.
28. W. Willinger, et al, *Self-similarity through high-variability: Statistical analysis of ethernet lan traffic at the source level*, SIGCOMM '95.
29. S. Basu, et al, *Time series models for Internet traffic*. Tech. Rep. GIT-CC-95-27, College of Computing, Georgia Institute of Technology, February 1995.
30. D. Eckhardt and P. Steenkiste, *Measurement and Analysis of the Error Characteristics of an In-Building Wireless Network*, SIGCOMM '96.

31. H. Balakrishnan, et al, *Analyzing stability in wide area network performance*, SIGMETRICS '97.
32. A. Feldmann, et al, *Data Networks as Cascades: Investigating the Multifractal Nature of Internet WAN Traffic*, SIGCOMM '98.
33. A. Feldmann, et al, *Dynamics of IP traffic: a study of the role of variability and the impact of control*, SIGCOMM '99.
34. V. Ribeiro, et al, *Simulation of non-Gaussian long-range-dependent traffic using wavelets*, SIGMETRICS '99.
35. A. Sang and S. Li, *A Predictability Analysis of Network Traffic*, INFOCOM 2000.
36. K. Lai and M. Baker, *Measuring link bandwidths using a deterministic model of packet delay*, SIGCOMM '00.
37. J. Cao, et al, *On the Nonstationarity of Internet Traffic*, SIGMETRICS 2001.
38. A. Downey, *Using pathchar to estimate Internet link characteristics*, SIGCOMM '99.
39. M. Allman, and V. Paxson, *On estimating end-to-end network path properties*, SIGCOMM '99.
40. A. Medina, et al, *Traffic Matrix Estimation: Existing Techniques and New Directions*, SIGCOMM '02.
41. D. Schwab, et al, *Characterizing the Use of a Campus Wireless Network*, INFOCOM 2004.
42. T. Karagiannis, et al, *A Nonstationary Poisson View of Internet Traffic*, INFOCOM 2004.
43. A. Kakhina, et al, *Structural Analysis of Network Traffic Flows*, SIGMETRICS 2004
44. Y. Qiao, et al, *An Empirical Study of the Multiscale Predictability of Network Traffic*, HPDC 2004. (**This paper got its start as a project in this course**)
45. D. Aguayo, *Link-level Measurements From an 802.11b Mesh Network*, SIGCOMM 2004.
46. Y. Chen, et al, *An Algebraic Approach to Practical and Sclable Overlay Network Monitoring*, SIGCOM 2004.
47. H. Jiang, et al, *Why is the Internet Traffic Bursty in Short (sub-RTT) Time Scales?*, SIGMETRICS 2005
48. K. Xu, et al, *Profiling Internet Backbone Traffic: Behavior Models and Applications*, SIGCOMM 2005.

Networks: Connections And Their Behavior

49. R. Caceres, et al, *Characteristics of wide-area TCP/IP conversations*, SIGCOMM '91.
50. R. Wolski, *Forecasting network performance to support dynamic scheduling using the network weather service*, HPDC '97 (Extended version available as UCSD Technical Report TR-CS96-494.

51. J. Bolliger, et al, *Bandwidth Modeling for Network-Aware Applications*, INFOCOM '99.
52. W. Feng and P. Tinnakornsrisuphap, *The Failure of TCP in High-Performance Computational Grids*, Supercomputing 2000.
53. L. Guo and I. Matta, *The War Between Mice and Elephants*, ICNP 2001.
54. M. Allman, *Measuring End-to-end Bulk Transfer Capacity*, IMW 2001.
55. A. Akella, et al, *Selfish Behavior and Stability of the Internet: A Game-theoretic Analysis of TCP*, SIGCOMM '02
56. M. Jain and C. Dovrolis, *End-to-end Available Bandwidth: Measurement Methodology, Dynamics, and Relation with TCP Throughput*, SIGCOMM '02
57. Q. He, *On The Predictability of Large Transfer TCP Throughput*, SIGCOMM 2005.
58. D. Lu, et al, *Characterizing and Predicting TCP Throughput on the Wide Area Network*, ICDCS 2005
59. D. Lu, et al, *Modeling and Taming Parallel TCP on the Wide Area Network*, IPDPS 2005.

Applications: Intrusion Detection

60. S. Hofmeyer, et al, *Intrusion detection using sequences of system calls*, Journal of Computer Security, volume 6, pp 151-180, 1998.
61. P. Barford, et al, *A Signal Analysis of Network Traffic Anomalies*, IMW 2002.
62. D. Moore, et al, *Code-Red: A Case Study on the Spread and Victims of an Internet Worm*, IMW 2002.
63. V. Yegneswaran, et al, *Internet Intrusions: Global Characteristics and Prevalence*, SIGMETRICS 2003.
64. C. Estan, et al, *Automatically Inferring Patterns of Resource Consumption In Network Traffic*, SIGCOMM 2003.
65. S. Saroiu, et al, *Measurement and Analysis of Spyware in a University Environment*, NSDI 2004.
66. A. Moore, et al, *Internet Traffic Classification Using Bayesian Analysis Techniques*, SIGMETRICS 2005.

Applications: P2P

67. M. Ripeanu, et al, *Mapping the Gnutella Network: Macroscopic Properties of Large-Scale Peer-to-Peer Systems*, IPTPS 2002.
68. S. Saroiu, et al, *A Measurement Study of Peer-to-Peer File Sharing Systems*, MCN 2002.
69. R. Bhagwan, et al, *Understanding Availability*, IPTPS 2003
70. F. Bustamante, et al, *Friendships that Last: Peer Lifespan and Its Role in P2P Protocols*, WCCD 2003.
71. Y. Qiao, et al *Looking at the Server Side of Peer-to-Peer Systems*, LCR 2004.
72. A. Iamnitchi, et al, *Small World File-sharing Communities*, INFOCOM 2004

73. A. Klemm, et al, *Characterizing the Query Behavior in Peer-to-Peer File Sharing Systems,* IMC 2004.
74. D. Leonard, et al, *On Lifetime-based Node Failure and Stochastic Resilience of Decentralized Peer-to-Peer Networks,* SIGMETRICS 2005.
75. S. Kirshnamurthy, et al, *A Statistical Theory of Chord Under Churn,* IPTPS 2005.
76. M. Yang, et al, *An Empirical Study of Free-Riding Behavior in the Maze P2P File-Sharing System,* IPTPS 2005.
77. J. Pouwelse, et al, *The Bittorrent P2P File-Sharing System: Measurements and Analysis,* IPTPS 2005.
78. L. Guo, *Measurmenets, Analysis, and Modeling of BitTorrent-like Systems,* IMC 2005.
79. Y. Qiao, et al, *Structured and Unstructured Overlays Under the Microscope: A Measurement-based View of Two P2P Systems That People Use,* USENIX 2006.

Applications: Web

80. M. Arlitt and C. Williamson, *Web server workload characterization: the search for invariants,* SIGMETRICS ’96.
81. M. Crovella and A. Bestavros, *Self-similarity in world wide web traffic,* SIGMETRICS ’96.
82. P. Barford and M. Crovella, *Generating representative web workloads for network and server performance evaluation,* SIGMETRICS ’98.
83. A. Myers, et al, *Performance characteristics of mirror servers on the Internet,* INFOCOM '99.
84. L. Breslau, et al, *Web caching and Zipf-like distributions: evidence and implications,* INFOCOM ’99.
85. S. Dykes, et al, *An Empirical Evaluation of Client-side Server Selection Algorithms,* INFOCOM 2000.
86. F. Smith, et al, *What TCP/IP Protocol Headers Can Tell Us About the Web,* SIGMETRICS 2001.
87. A. Adya, et al, *Analyzing the Browse Patterns of Mobile Clients,* IMC 2002.
88. M. Harchol-Balter, et al, *Size-based Scheduling to Improve Web Performance,* ACM TOCS 21:2, May 2003.
89. D. Lu, et al, *Size-based Scheduling Policies With Inaccurate Scheduling Information,* MASCOTS 2004.
90. D. Lu, et al, *Effects and Implications of File Size/Service Time Correlation on Web Server Scheduling Policies,* MASCOTS 2005.

Applications: Video and Audio

91. M. Garrett and W. Willinger, *Analysis, modeling and generation of self-similar {VBR} video traffic,* SIGCOMM '94.
92. M. Krunz, et al, *On the Characterization of VBR MPEG Streams,* SIGMETRICS 97.

93. A. Bavier, et al, *Predicting MPEG execution times*, SIGMETRICS '98.
94. A. Mena and J. Heidemann, *An Empirical Study of Real Audio Traffic*, INFOCOM 2000.
95. Z. Su, et al, *A Prediction System for Multimedia Prefetching*, ACM Multimedia 2000.
96. D. Loguinov and H. Radha, *Measurement Study of Low-bitrate Internet Video Streaming*, IMW 2001.
97. D. Loguinov and H. Radha, *End-to-end Internet Video Traffic Dynamics: Statistical Study and Analysis*, INFOCOM '02.
98. K. Sripanidkulchai, et al, *An Analysis of Live Streaming Workloads On The Internet*, IMC 2004.

Applications: Databases

99. K. Keeton and D. Patterson, *Towards A Simplified Database Workload For Computer Architecture Evaluations*, Chapter 3 of Workload Characterization for Computer System Design, edited by L. John and A. Maynard, Kluwer, 2000.
100. TPC benchmarks.

Applications: Games and Interactive Applications

101. DIS Steering Committee, *The DIS Vision, A Map to the Future of Distributed Simulation*. Orlando, Florida, Institute for Simulation and Training, 1994.
102. D. Cavitt, et al, *A Performance Monitoring Application for Distributed Interactive Simulations (DIS)*, Winter Simulation Conference, 1997.
103. T. Mitra, T. Chiueh, *Dynamic 3D Graphics Workload Characterization and the Architectural Implications*, 32nd ACM/IEEE International Symposium on Microarchitecture, November 1999. Also available as SUNY Stony Brook Experimental Systems Lab Technical Report TR-61.
104. B. Schmidt, et al, *The Interactive Performance of SLIM: A Stateless Thin Client Architecture*, SOSP 1999.
105. A. Abdelkhalek, et al, *Behavior and Performance of Interactive Multi-player Game Servers*, ISPASS 2001.
106. A. Lai and J. Nieh, *Limits of Wide-Area Thin-client Computing*, SIGMETRICS '02.
107. C. Chambers, *Measurement-based Characterization of a Collection of On-line Games*, IMC 2005.

Applications: File Systems

108. T. Kroeger and D. Long, Predicting file system actions from prior events, USENIX '96.

109. S. Gribble, et al, *Self-similarity in file systems*, SIGMETRICS ’98.
110. J. Douver and W. Bolosky, *A Large-Scale Study of File-System Contents,* SIGMETRICS ‘99.

Applications: Scientific and Parallel Applications

111. R. Arpaci-Dusseau, et al. *The Interaction of Parallel and Sequential Workloads on a Network of Workstations*. SIGMETRICS ’95.
112. N. Kapadia, et al, *Predictive application-performance modeling in a computational grid environment,* HPDC '99.
113. J. Subhlok, et al, *Impact of Job Mix on Optimizations for Space Sharing Schedulers,* Supercomputing ’96.
114. P. Dinda, et al, *The measured network traffic of compiler-parallelized programs*, ICPP 2001.
115. J. Vetter, et al, *An Empirical Performance Evaluation of Scalable Scientific Applications,* SC 2002.
116. J. Vetter, *Dynamic Statistical Profiling of Communication Activity in Distributed Applications,* SIGMETRICS ’02.
117. J. Vetter, et al, *Communication Characteristics of Large-Scale Scientific Applications for Contemporary Cluster Architectures,* IPDPS 2002.
118. V. Taylor, et al, *Using Kernel Couplings to Predict Parallel Application Performance*, HPDC 2002.
119. D. Thain, et al, *Pipeline and Batch Sharing in Grid Workloads,* HPDC 2003
120. S. Kleban, et al, *Quelling Queue Storms,* HPDC 2003.
121. G. Marin, et al, *Cross-Architecture Performance Predictions for Scientific Applications Using Parameterized Models,* SIGMETRICS 2004.
122. D. England, *A New Metric For Robustness With Application To Job Scheduling*, HPDC 2005.
123. L. Yang, et al, *Cross-platform Performance Prediction of Parallel Applications Using Partial Execution,* SC 2005
124. U. Srinivasan, et al, *Characterization and Analysis of HMMER and SVM-RFE Parallel Bioinformatics Applications*, IISWC 2005.
125. S. Sodhi, et al, *Automatic Construction and Evaluation of Performance Skeletons,* IPDPS 2005.

Users

126. W. Tetzlaff, *State Sampling of Interactive VM/370 Users*, IBM Systems Journal 18(1), 1979.
127. D. Embley, et al, *Behaviorial Aspects of Text Editors,* ACM Computing Surveys 13:1, January, 1981.

128. B. Chen, et al, *The Measured Performance of Personal Computer Operating Systems*, ACM TOCS 14:1, 1996.
129. Y. Endo, et al, *Using Latency to Evaluate Interactive System Performance*, OSDI 1996.
130. A. Komatsubara, et al, *Psychological Upper and Lower Limits Of System Response Time and the User's Preference on Skill Level*, HCI 1997.
131. S. Bhola and M. Ahamad, *Workload Modeling for Highly Interactive Distributed Applications*, Technical Report GIT-CC-99-2, College of Computing, Georgia Institute of Technology, 1999.
132. J. Klein, *Computer Response to User Frustration*, Masters Thesis, MIT, 1999.
133. C. Reynolds, *The Sensing and Measurement of Frustration With Computers*, Masters Thesis, MIT, 2001.
134. T. Henderson and S Bhatti, *Modeling User Behavior in Network Games*, ACM Multimedia 2001.
135. A. Balachandran, *Characterizing User Behavior and Network Performance in a Public Wireless LAN*, SIGMETRICS 2001.
136. D. Olshefski, *Inferring Client Response Time at the Web Server*, SIGMETRICS 2002.
137. C. Dewes, et al, *An Analysis of Internet Chat Systems*, IMC 2003.
138. J-D Ruvini, *Adapting to the User's Internet Search Strategy*, UM 2003.
139. T. Zhu, et al, *Learning a Model of a Web User's Interests*, UM 2003.
140. A. Gupta, et al, *Measuring and Understanding User Comfort With Resource Borrowing*, HPDC 2004. (**This paper got its start as a project in this course**)
141. B. Davison, *Learning Web Request Patterns*, Chapter in *Web Dynamics: Adapting to Change in Content, Size, Topology, and Use*, Levene and Poulovassilis, editors, Springer, 2004.
142. H. Liu, et al, *Client Behavior and Feed Characteristics of RSS, A Publish-Subscribe System for Web Micronews*, IMC 2005.
143. R. Balan, et al, *Simplifying Cyber Foraging For Mobile Devices*, Technical Report CMU-CS-05-157R, Carnegie Mellon.
144. J. Sousa, et al, *Giving Users the Steering Wheel For Guiding Resource-Adaptive Systems*, Technical Report CMU-CS-05-198, Carnegie Mellon.
145. N. Hine, et al, *Modeling the Behavior of Elderly People as a Means of Monitoring Well Being*, UM 2005.
146. *Putting the User in Direct Control of CPU Scheduling*, preprint
147. *Process- and User-driven Dynamic Voltage and Frequency Scaling*, preprint
148. Workshop on Adaptive Systems and User Modeling on the World Wide Web: http://wwwis.win.tue.nl/asum99/

Measurement and Prediction Tools and Systems

(note: systems are also described in the other sections. these papers are *predominantly* about the systems)

149. B. Lowekamp, et al, *A resource monitoring system for network-aware applications*, HPDC '98.
150. K. Obraczka, et al, *The Performance of A Service For Network-Aware Applications*, SPDT 98.
151. B. Lowekamp, et al, *Direct queries for discovering network resource properties in a distributed environment*, HPDC '99.
152. R. Wolski, et al, *The network weather service: A distributed resource performance forecasting system*, Journal of Future Generation Computing Systems, 1999, (A version is also available as UC-San Diego technical report number TR-CS98-599. Initial work in HPDC '97.)
153. M. Stemm, et al, *A Network Measurement Architecture for Adaptive Applications*, INFOCOM 2000.
154. D. Gunter, et al, *NetLogger: A Toolkit for Distributed System Performance Analysis*, MASCOTS 2000.
155. P. Dinda, *Design, Implementation, and Performance of an Extensible Toolkit for Resource Prediction In Distributed Systems*, IEEE TPDS 17:2, February, 2006.
156. Other Network Measurement Tools: NLANR List: http://dast.nlanr.net/NPMT/, CAIDA List: http://www.caida.org/tools

Additional Measurement Principles

157. PASTA Principle (see V. Paxson, *End-to-end routing behavior in the Internet*, above)
158. Nyquist Criterion and Sampling Theory (see A. Oppenheim, et al, *Signals and Systems*, below)

Additional Modeling and Analysis Techniques

159. Arrival processes (see Jain on M/M/1 and M/G/1)
160. Time Series Analysis (see Statsoft Guide and Box's Time Series Analysis)
161. J. Bassingthwaighte, et al, *Fractal structures and processes*, Chaos and the Changing Nature of Science and Medicine: An Introduction, D. Herbert, Ed., no.376 in AIP Conference Proceedings, American Institute of Physics, pp. 54—79, April 1995.
162. J. Vetter, and D. Reed, *Managing performance analysis with dynamic statistical projection pursuit*, Supercomputing '99.
163. J. Skicewicz, et al, *Multi-resolution Resource Behavior Queries Using Wavelets*, HPDC 2001. (**This paper got its start as a project in this course**)
164. Statistics and Probability Intro (see Jain's *Art of Computer Systems Performance Analysis*, Statsoft *Guide*, S-Plus *Guid*e, Golnick's *Cartoon Guide*, all above)
165. Signal processing and Fourier (see A Oppenheim, et al, *Signals and Systems*, above)

166. H. Abarbanel, et al, *Obtaining order in a world of chaos*, IEEE Signal Processing Magazine, May, 1998.
167. D. Dasgupta, and S. Forrest, *Novelty detection in time series data using ideas from immunology*, International Conference on Intelligent Systems, 1999.
168. A. Arpaci-Dusseau and R. Arpaci-Dusseau, *Information and Control in Gray-box systems*, SOSP 2001