

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

Reading List

Note: We will not read all of these papers in class. 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.)

Host Characterization and Prediction

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. Mutka and M Livny, *The available capacity of a privately owned workstation environment*, Performance Evaluation 12:4, July 1991.
4. M. Harchol-Balter, A. Downey, *Exploiting process lifetime distributions for dynamic load balancing*, SIGMETRICS '96.
5. 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.)
6. R. Wolski, et al, *Predicting the CPU availability of time-shared Unix systems*, HPDC '99.
7. P. Dinda and D. O'Hallaron, *Host load prediction using linear models*, HPDC '99 (journal version appears in Cluster Computing, summary in SIGMETRICS 2001)
8. P. Dinda, *Online Prediction of the Running Time of Tasks*, HPDC 2001, (journal version appears in Cluster Computing.)

Network Characterization and Prediction

9. R. Caceres, et al, *Characteristics of wide-area TCP/IP conversations*, SIGCOMM '91.
10. V. Paxson, and S. Floyd, *Wide-area traffic: The failure of Poisson modeling*. {IEEE/ACM} Transactions on Networking. 3:3, June 1995.
11. W. Willinger, et al, *Self-similarity in high-speed packet traffic: Analysis and modeling of ethernet traffic measurements*, Statistical Science 10:1, January 1995.

12. W. Willinger, et al, *Self-similarity through high-variability: Statistical analysis of ethernet lan traffic at the source level*, SIGCOMM '95.
13. 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.
14. V. Paxson, *End-to-end routing behavior in the Internet*, IEEE/ACM Transactions on networking, 5:5, 1997.
15. 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).
16. H. Balakrishnan, et al, *Analyzing stability in wide area network performance*, SIGMETRICS '97.
17. M. Faloutsos, et al, *On power-law relationships of the Internet topology*, SIGCOMM '99.
18. A. Feldmann, et al, *Dynamics of IP traffic: a study of the role of variability and the impact of control*, SIGCOMM '99.
19. V. Ribeiro, et al, *Simulation of non-Gaussian long-range-dependent traffic using wavelets*, SIGMETRICS '99.
20. N. Duffield and M. Grossglauser, *Trajectory sampling for direct traffic observation*, SIGCOMM '00.
21. K. Lai and M. Baker, *Measuring link bandwidths using a deterministic model of packet delay*, SIGCOMM '00.
22. J. Bolliger, et al, *Bandwidth Modelling for Network-Aware Applications*, INFOCOM '99.
23. A. Sang and S. Li, *A Predictability Analysis of Network Traffic*, INFOCOM 2000.
24. J. Cao, et al, *On the Nonstationarity of Internet Traffic*, SIGMETRICS 2001.
25. A. Balachandran, *Characterizing User Behavior and Network Performance in a Public Wireless LAN*, SIGMETRICS 2001.
26. A. Downey, *Using pathchar to estimate Internet link characteristics*, SIGCOMM '99.
27. M. Allman, and V. Paxson, *On estimating end-to-end network path properties*, SIGCOMM '99.
28. P. Barford and D. Plonka, *Characteristics of Network Traffic Flow Anomalies*, Internet Measurement Workshop.
29. W. Feng and P. Tinnakornsrisuphap, *The Failure of TCP in High-Performance Computational Grids*, Supercomputing 2000.

Web, Video, and Audio

30. M. Garrett and W. Willinger, *Analysis, modeling and generation of self-similar {VBR} video traffic*, SIGCOMM '94.
31. M. Krunz, et al, *On the Characterization of VBR MPEG Streams*, SIGMETRICS 97.
32. A. Myers, et al, *Performance characteristics of mirror servers on the Internet*, INFOCOM '99. .

33. L. Breslau, et al, *Web caching and Zipf-like distributions: evidence and implications*, INFOCOM '99.
34. M. Crovella and A. Bestavros, *Self-similarity in world wide web traffic*, SIGMETRICS '96.
35. M. Arlitt and C. Williamson, *Web server workload characterization: the search for invariants*, SIGMETRICS '96.
36. A. Bavier, et al, *Predicting MPEG execution times*, SIGMETRICS '98.
37. P. Barford and M. Crovella, *Generating representative web workloads for network and server performance evaluation*, SIGMETRICS '98.
38. A. Mena and J. Heidemann, *An Empirical Study of Real Audio Traffic*, INFOCOM 2000.
39. S. Dykes, et al, *An Empirical Evaluation of Client-side Server Selection Algorithms*, INFOCOM 2000.
40. F. Smith, et al, *What TCP/IP Protocol Headers Can Tell Us About the Web*, SIGMETRICS 2001.
41. D. Olsheski, *Inferring Client Response Time at the Web Server*, SIGMETRICS 2002.

Games and Interactive Applications

42. 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.
43. A. Abdelkhalek, et al, *Behavior and Performance of Interactive Multi-player Game Servers*, ISPASS 2001.
44. 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.
45. 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.
46. D. Cavitt, et al, *A Performance Monitoring Application for Distributed Interactive Simulations (DIS)*, Winter Simulation Conference, 1997.
47. DIS Steering Committee, *The DIS Vision, A Map to the Future of Distributed Simulation*. Orlando, Florida, Institute for Simulation and Training, 1994.

Other Applications

48. M. Devarakonda and R. Iyer, *Predictability of process resource usage: a measurement-based study on UNIX*, IEEE Transactions on Software Engineering, 15:12, 1989.
49. R. Arpaci-Dusseau, et al. *The Interaction of Parallel and Sequential Workloads on a Network of Workstations*. SIGMETRICS '95.

50. T. Kroeger and D. Long, Predicting file system actions from prior events, USENIX '96.
51. S. Gribble, et al, *Self-similarity in file systems*, SIGMETRICS '98.
52. P. Dinda, et al, *The measured network traffic of compiler-parallelized programs*, ICPP 2001.
53. N. Kapadia, et al, *Predictive application-performance modeling in a computational grid environment*, HPDC '99.
54. J. Douver and W. Bolosky, *A Large-Scale Study of File-System Contents*, SIGMETRICS '99.
55. J. Vetter, *Dynamic Statistical Profiling of Communication Activity in Distributed Applications*, SIGMETRICS 2001.
56. A. Adya, et al, *Analyzing the Browse Patterns of Mobile Clients*, Internet Measurement Workshop.
57. S. Hofmeyer, et al, *Intrusion detection using sequences of system calls*, Journal of Computer Security, volume 6, pp 151-180, 1998.
58. J. Subhlok, et al, *Impact of Job Mix on Optimizations for Space Sharing Schedulers*, Supercomputing '96.

Measurement and Prediction Systems

59. B. Lowekamp, et al, *A resource monitoring system for network-aware applications*, HPDC '98.
60. B. Lowekamp, et al, *Direct queries for discovering network resource properties in a distributed environment*, HPDC '99.
61. P. Dinda, and D. O'Hallaron, *An extensible toolkit for resource prediction in distributed systems*, Tech. Rep. CMU-CS-99-138, Carnegie Mellon University, July 1999.
62. 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.)
63. M. Stemm, et al, *A Network Measurement Architecture for Adaptive Applications*, INFOCOM 2000.
64. Grid Performance Working Group (www.gridforum.org)
65. Grid Information Services Working Group (www.gridforum.org)

Measurement and Sampling Theory

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

Data Analysis Techniques

68. J. Bassingthwaight, 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.
69. J. Vetter, and D. Reed, *Managing performance analysis with dynamic statistical projection pursuit*, Supercomputing '99.
70. Time Series Analysis (see Statsoft Guide, Box's Time Series Analysis,
71. Wavelets (<http://www.mathsoft.com/wavelets.html>)
72. Statistics and Probability Intro (see Jain's *Art of Computer Systems Performance Analysis*, Statsoft *Guide*, S-Plus *Guide*, Golnick's *Cartoon Guide*, all below)
73. Signal processing and Fourier (see A Oppenheim, et al, *Signals and Systems*, below)
74. H. Abarbanel, et al, *Obtaining order in a world of chaos*, IEEE Signal Processing Magazine, May, 1998.
75. D. Dasgupta, and S. Forrest, *Novelty detection in time series data using ideas from immunology*, International Conference on Intelligent Systems, 1999.

Added Materials

76. D. Eckhardt and P. Steenkiste, *Measurement and Analysis of the Error Characteristics of an In-Building Wireless Network*, SIGCOMM '96.
77. A. Feldmann, et al, *Data Networks as Cascades: Investigating the Multifractal Nature of Internet WAN Traffic*, SIGCOMM '98.
78. V. Taylor, et al, *Using Kernel Couplings to Predict Parallel Application Performance*, HPDC '02. (systems: prophesy)
79. D. Gunter, et al, *Dynamic Monitoring of High-performance Distributed Applications*, HPDC '02 (systems: netlogger)
80. O. Maennel and A. Feldmann, *Realistic BGP Traffic for Test Labs*, SIGCOMM '02.
81. A. Akella, et al, *Selfish Behavior and Stability of the Internet: A Game-theoretic Analysis of TCP*, SIGCOMM '02.
82. N. Spring, et al, *Measuring ISP Topologies With Rocketfuel*, SIGCOMM '02
83. H. Tangmunarankit, et al, *Network Topology Generators: Degree-based versus Structural*, SIGCOMM '02.
84. A. Medina, et al, *Traffic Matrix Estimation: Existing Techniques and New Directions*, SIGCOMM '02.
85. M. Jain and C. Dovrolis, *End-to-end Available Bandwidth: Measurement Methodology, Dynamics, and Relation with TCP Throughput*, SIGCOMM '02
86. M. Coates, et al, *Maximum Likelihood Network Topology Identification from Edge-based Unicast Measurements*, SIGMETRICS '02.
87. A. Lai and J. Nieh, *Limits of Wide-Area Thin-client Computing*, SIGMETRICS '02.

88. J. Vetter, *Dynamic Statistical Profiling of Communication Activity in Distributed Applications*, SIGMETRICS '02.
89. Q. Chen, et al, *The Origin of Power-Laws in Internet Topologies Revisited*, INFOCOM '02.
90. D. Loguinov and H. Radha, *End-to-end Internet Video Traffic Dynamics: Statistical Study and Analysis*, INFOCOM '02.
91. C. Gkantsidis, et al, *Spectral Analysis of Internet Topologies*, INFOCOM '03.
92. T. Henderson and S Bhatti, *Modeling User Behavior in Network Games*, ACM Multimedia 2001.
93. Z. Su, et al, *A Prediction System for Multimedia Prefetching*, ACM Multimedia 2000.
94. A. Arpaci-Dusseau and R. Arpaci-Dusseau, *Information and Control in Gray-box systems*, SOSP 2001.
95. M. Allman, *Measuring End-to-end Bulk Transfer Capacity*, IMW 2001.
96. D. Loguinov and H. Radha, *Measurement Study of Low-bitrate Internet Video Streaming*, IMW 2001.
97. P. Barford, et al, *A Signal Analysis of Network Traffic Anomalies*, IMW 2002.
98. D. Moore, et al, *Code-Red: A Case Study on the Spread and Victims of an Internet Worm*, IMW 2002.
99. L. Guo and I. Matta, *The War Between Mice and Elephants*, ICNP 2001.

Performance Analysis Books

- 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. However, none of its coverage is particularly deep, and the author makes undeclared assumptions.

General Books

- Larry Golnick, et al, *The Cartoon Guide To Statistics*, 1994.
- A very readable introduction to basic probability theory and classic parametric statistics.
- 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.
- Mathsoft, Inc, *The S-Plus Guide to Statistics, Volumes 1 and 2*,
- Good reference book with examples given in S-Plus.
- Alan V. Oppenheim, et al, *Signals and Systems*, 1983.
- Good introduction to linear systems theory

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