

David Plonka

Computer Sciences Department
University of Wisconsin - Madison
1210 West Dayton Street
Madison, WI 53706
email: plonka@cs.wisc.edu
<http://www.cs.wisc.edu/~plonka/>

Primary Research & Development Interests

◊ Large scale Internet applications performance and optimization ◊ Network performance, measurement, and anomaly detection ◊ Internet protocols, network management, security vulnerabilities, and software design.

Education

- Ph.D. Candidate, Computer Sciences** expected by May, 2012
University of Wisconsin - Madison, Madison, WI
Advisor: Paul Barford
- M.S. Computer Sciences** December, 2008
University of Wisconsin - Madison
- B.S. Computer Science and Artificial Intelligence majors, *magna cum laude*** May, 1991
Carroll College (now Carroll University), Waukesha, WI

Professional Skills

◊ Technical leadership in research and development of solutions to open problems ◊ Multi-layer troubleshooting and diagnosis of performance problems across physical network, ethernet, TCP/IP, and application layers ◊ Relational database application design, implementation, SQL, and performance optimization ◊ Network management and Internet operation, network systems programming and software design ◊ Unix and Linux systems programming in C, perl, and Korn/Bourne Again/POSIX shells.

Professional Experience

- Computer Sciences Department, University of Wisconsin**, Madison, WI 2006 - present
Research Assistant, Fellow (2009 - 2010)
- ◊ Collaborated with researchers to invent and implement techniques for Internet traffic classification and measurement, for detection of malicious traffic and network abuse, and for network anomaly detection
◊ Solved Linux systems scalability and performance problems ◊ Performed novel studies of the Domain Name System (DNS), longitudinal studies of Internet email delivery performance, and studies of network configuration artifacts and network operations practice.
- Google Inc.**, Madison, WI Summer, 2010
Platforms Software Engineering Intern
- ◊ Worked within an advanced development team for Google Platforms, *i.e.*, the computing and networking infrastructure on which Internet-scale Google applications run ◊ Implemented measurement infrastructure for testbeds, performed controlled experiments, and made performance recommendations.
- Division of Information Technology, University of Wisconsin** 2006 - present
Network Services Consultant
- ◊ Counseled Network Services staff and assisted with maintenance of custom software systems to manage and operate the high performance production campus network comprising 110,000 ethernet ports, of which 98,500 are edge ports, serving a campus with approximately 45,000 users.

Computer Sciences Department, University of Wisconsin 2001 - 2006
Co-manager of the Wisconsin Advanced Internet Laboratory

◊ Performed design, equipment selection, configuration, and operation of a unique multi-million dollar network research facility focusing on simulation and emulation of end-to-end Internet communications at large scale.

Division of Information Technology, University of Wisconsin 1997 - 2006
Senior Systems Programmer / Network Engineer, Systems Programmer

◊ Introduced configuration management for the campus and state-wide networks (WiscNet), authored network measurement, management, and traffic visualization software for network operation and Internet usage trend analysis ◊ Worked within teams to develop, manage, and troubleshoot redesigns and upgrades of campus and state-wide networks utilizing equipment from multiple vendors ◊ Developed Internet standards within the IETF in the areas of Internet traffic measurement (the IPFIX protocol), global Internet operations, and the Network Time Protocol (NTP).

McHugh Freeman (now RedPrairie Corporation), Waukesha, WI 1991 - 1997
Lead Systems Programmer, Systems Programmer

◊ Maintained and managed proprietary library software releases ◊ Developed APIs and programming guidelines and trained employees ◊ Performed advanced portable Unix and VMS systems programming in the logistics industry, *i.e.*, warehouse and distribution automation, and developed database application software for the Fortune 500 and other corporate customers.

Consulting Experience

Independent Consultant, Madison, WI 1997 - 1999, 2009

◊ Consulting in logistics software and computer network technologies.

Awards and Honors

Google Engineering Intern Scholarship 2010

Lawrence H. Landweber NCR Fellowship in Distributed Systems 2009

Student Travel Grants: USENIX LISA and ACM IMC Conferences 2007, 2008, 2009

USENIX 21st Large Installation System Administration Conference (LISA '07), **Best Paper Award** 2007

Carroll College Computer Science Department, **Outstanding Computer Science Student** 1991

Patents

1. Paul R. Barford and David J. Plonka. "Apparatus and Method for Classifying Network Packet Data", United States Patent 7907543, 2011.
2. Amos Ron, Paul R. Barford, Jeffery Thomas Kline, Sangnam Nam, and David John Plonka. "Method and Apparatus for Network Anomaly Detection", U.S. Patent pending, 2007.
3. Vinod T. Yegneswaran, Paul R. Barford and David J. Plonka. "Scalable Monitor of Malicious Network Traffic", United States Patent 8015605, 2011.

Intellectual property licensed by Nemean Networks, LLC, acquired by Qualys, Inc., September, 2010.

Professional Activities

Chair

Internet Protocol Flow Information eXport (IPFIX) Working Group co-chair, 2001 - 2006
Internet Engineering Task Force (IETF)

Program Committees

USENIX Large Installation System Administration Conference (LISA) 2008, 2009
ACM Workshop on Network Data Anonymization (NDA) 2008
IEEE/IFIP Workshop on End-to-End Monitoring Techniques and Services (E2EMON) 2006

Panelist

National Science Foundation 2009

Other Service

Internet2 Network Research Review Committee (NRRC) 2009 - present

Memberships

Association for Computing Machinery (ACM) 2002 - , 2008 - present
IEEE 2008 - present

Refereed Conference and Workshop Publications

1. David Plonka and Paul Barford. “Flexible Traffic and Host Profiling via DNS Rendezvous”, In *Proceedings of the 1st Securing and Trusting Internet Names Workshop (SATIN 2011)*, Teddington, UK, April, 2011. 73% accept rate.
2. David Plonka and Andres Jaan Tack. “An Analysis of Network Configuration Artifacts”, In *Proceedings of the USENIX 23rd Large Installation System Administration Conference (LISA '09)*, Baltimore, November, 2009. 35% accept rate.
3. David Plonka and Paul Barford. “Context-aware Clustering of DNS Query Traffic”, In *Proceedings of the ACM SIGCOMM / USENIX Eighth Internet Measurement Conference (IMC 2008)*, Vouliagmeni, Greece, October, 2008. 17% accept rate.
4. Jeff Kline, Sangnam Nam, Paul Barford, David Plonka, Amos Ron. “Traffic Anomaly Detection at Fine Time Scales with Bayes Nets”, In *Proceedings of the IEEE Third International Conference on Internet Monitoring and Protection (ICIMP 2008)*, Bucharest, Romania, June/July, 2008. 23% accept rate.
5. David Plonka, Archit Gupta, Dale Carder. “Application Buffer-Cache Management for Performance: Running the World’s Largest MRTG”, In *Proceedings of the USENIX 21st Large Installation System Administration Conference (LISA '07)*, Dallas, November, 2007. **Best Paper Award winner.**
6. David Mills, Judah Levine, Richard Schmidt, David Plonka. “Coping with Overload on the Network Time Protocol Public Servers”, In *Proceedings of the 36th Annual Precise Time and Time Interval Systems and Applications Meeting (PTTI 2004)*, Washington, D.C., December, 2004.
7. Vinod Yegneswaran, Paul Barford and David Plonka. “On the Design and Use of Internet Sinks for Network Abuse Monitoring”, In *Proceedings of the Symposium on Recent Advances in Intrusion Detection (RAID '04)*, pp. 146-165, Sophia Antipolis, France, September, 2004. 20% accept rate.

8. Paul Barford, Jeff Kline, David Plonka, Amos Ron. “A Signal Analysis of Network Traffic Anomalies”, In *Proceedings of ACM SIGCOMM Internet Measurement Workshop '02*, pp. 71-82, Marseilles, France, November, 2002. 42% accept rate.
9. Paul Barford and David Plonka. “Characteristics of Network Traffic Flow Anomalies”, In *Proceedings of ACM SIGCOMM Internet Measurement Workshop '01*, pp. 69-73, San Francisco, November, 2001. 26% accept rate.
10. Dave Plonka. “FlowScan: A Network Traffic Flow Reporting and Visualization Tool”, In *Proceedings of the USENIX 14th Systems Administration Conference (LISA 2000)*, New Orleans, December, 2000.

Other Publications

Internet RFC Documents

1. David Plonka. Best Current Practice 105, RFC 4085: “Embedding Globally-Routable Internet Addresses Considered Harmful”, Internet Engineering Task Force, June, 2005.

Invited Papers

1. David Plonka and Paul Barford. “Network Anomaly Confirmation, Diagnosis and Remediation”, In *Proceedings of the 47th Annual Allerton Conference on Communication, Control, and Computing (Allerton Conference 2009)*, Monticello, Illinois, September/October, 2009.

Journals and Periodicals

1. Dave Plonka. “An Analysis of Napster and Other IP Flow Sizes”, *Network Analysis Times*, April, 2001.
2. Dave Plonka. “How I Recovered Data I Thought I’d Lost”, *Sys Admin - the Journal for UNIX Systems Administrators*, August, 2000.
3. Dave Plonka. “Sys Admin File Revision Control with RCS”, *Sys Admin - the Journal for UNIX Systems Administrators*, December, 1998.
4. Dave Plonka. “Managing System Administration Tasks Using GNATS”, *Sys Admin - the Journal for Unix Systems Administrators*, February, 1997.

Online Technical Reports

1. Aaron Bergstrahl, Erik Paulson, and David Plonka. “Bill-Pay Control: An Interactive User Interface to Select IP Service Quality”, December, 2006.
2. Dave Plonka. “Flawed Routers Flood University of Wisconsin Internet Time Server”, August, 2003.
3. Dave Plonka. “UW-Madison Napster Traffic Measurement”, March, 2000.

Invited Talks and Workshops

1. Invited talk, “Rendezvous-based Network Traffic Analysis”, Lockdown 2009 Conference, UW-Madison, Madison, Wisconsin, July, 2009.
2. Invited talk, “Rendezvous-based Network Traffic Analysis”, Massachusetts Institute of Technology, Cambridge, Massachusetts, February, 2009.
3. Invited talk, “Rendezvous-based Network Traffic Analysis”, Boston University, Boston, February, 2009.
4. Invited talk, WiP: “Network Admins are Programmers: An Analysis of Network Management Artifacts”, LISA 2008, San Diego, November, 2008.
5. Invited talk, BoF: “DNS and Network Traffic Analysis with TreeTop”, LISA 2008, San Diego, November, 2008.

6. Invited participant and session moderator, “NSF NeTS Future Internet Design (FIND) Student Meeting”, Seattle, August, 2008.
7. Invited talk, “Application Buffer-Cache Management for Performance: Running the World’s Largest MRTG”, Madison Chapter of The League of Professional System Administrators (LOPSA-Madison) Meeting, Madison, Wisconsin, January, 2008.
8. Invited participant, “Wisconsin Entrepreneurial Bootcamp”, UW-Madison, Madison, Wisconsin, July/August, 2007.
9. Invited talk, “Anomalous Traffic from Internet Consumer Products”, ESCC/Internet2 Joint Techs Workshop, Madison, Wisconsin, July, 2006.
10. Invited talk, “Get Your FIX: Flow Information eXport Analysis and Visualization”, ESCC/Internet2 Joint Techs Workshop, Madison, Wisconsin, July, 2006.
11. Invited panelist, “12 Hot Research Topics in Monitoring and Measurements”, Fourth IEEE/IFIP Workshop on End-to-End Monitoring Techniques and Services (E2EMON’06), Vancouver, Canada, April, 2006.
12. Invited talk, “AANTS: Web-Based Tools for Cooperative Campus Network Administration”, Fall 2005 Internet2 Member Meeting, Philadelphia, September, 2005.
13. Invited participant, Sampling 2005 Workshop, Paris, July, 2005.
14. Invited talk, “Exposing Abuse in Internet Traffic Measurements”, Lockdown 2005 Conference, UW-Madison, Madison, Wisconsin, July, 2005.
15. Invited talk, “Flow-based Measurements and Other Means to Expose Internet Pathology”, Institut Eurecom, Sophia Antipolis, France, December, 2004.
16. Invited talk, “Bare-Bones Measurement Data Archiving”, ISMA Data Catalog 2004 Workshop, CAIDA/SDSC, La Jolla, California, June, 2004.
17. Invited talk, “Embedding Globally-Routable Internet Addresses Considered Harmful”, Global Routing Operations (GROW) Working Group, 58th IETF meeting, Minneapolis, November, 2003.
18. Invited talk, “A Case Study in Internet Pathology: Flawed Routers Flood University’s Network”, LISA 2003, San Diego, October, 2003.
19. Invited talk, “Flawed Routers Flood University of Wisconsin Internet Time Server”, NANOG 29, Chicago, October, 2003.
20. Invited talk, “IP Flow Measurement and Analysis with FlowScan”, Duke University, Durham, North Carolina, September, 2003.
21. Invited talk, “A Case Study in Internet Pathology: Flawed Routers Flood University’s Network”, Duke University, Durham, North Carolina, September, 2003.
22. Invited talk, “Flawed Routers Flood University of Wisconsin Internet Time Server”, MAD-SAGE meeting, Madison, Wisconsin, August, 2003.
23. Invited talk, “Internet Sink Deployments”, IEPG meeting, Vienna, Austria, July, 2003.
24. Invited talk, “IP Flow Measurement and Analysis with FlowScan”, IPAM Mini-Workshop on “Internet-wide Measurements”, Los Angeles, March, 2002.
25. Invited talk, “FlowScan at the University of Wisconsin-Madison”, EDUCAUSE 2001, Indianapolis, October, 2001.
26. Invited talk, “Practical Flow-based Passive Measurements and FlowScan”, Computer Sciences Department Network Seminar, UW-Madison, Madison, Wisconsin, June, 2001.

27. Invited talk, “Better Campus Network Health Through Instrumentation”, UW System Network Operators Meeting, Madison, Wisconsin, March, 2001.
28. Invited talk, “FlowScan - LIVE!”, NANOG 21, Atlanta, February, 2001.
29. Invited talk, “FlowScan - a Network Traffic Reporting and Visualization Tool”, MAD-SAGE meeting, Madison, Wisconsin, November 2000.
30. Invited talk, BoF: “Flow-based Passive Measurements”, CANARIE/NLANR/Internet2 Techs Workshop, Toronto, August, 2000.
31. Invited talk, “Unix Security”, UW-Madison Unix Fair, Madison, Wisconsin, August, 1999.
32. Invited participant, CAIDA ISMA Workshop on “Passive Measurement Data and Analysis”, La Jolla, California, January, 1999.
33. Invited talk, “Unix Security Basics and Tools - an Introduction”, UW-Unix meeting, Madison, Wisconsin, April, 1998.

Teaching Experience

Computer Sciences Department, University of Wisconsin, Madison, WI

CS 640 - Introduction to Computer Networks

Fall 2008, Fall 2009

Guest lectures on Ethernet, Wireless Ethernet, IP Multicast, TCP

Available Software

The following software packages are available at: <http://net.doit.wisc.edu/~plonka/packages.html>

fincore

fincore (“File IN CORE”) is a command that shows which pages (blocks) of a file are in core memory. It is particularly useful for determining the contents of the buffer-cache.

FlowScan

FlowScan is a system to analyze and report on flows exported by IP routers and collected using flow-tools, argus, cflowd, etc. It could also be considered a “front-end” for RRDTOOL.

RRGrapher

RRGrapher is a CGI graphing tool for RRDTOOL. It’s basically an interactive web graph construction set.

Net::Patricia

This is a perl API to a Patricia Trie data structure to perform fast lookups by IP address. Patricia Trie is the data structure used by the BSD kernel routing code where it is named “radix.”

junipoll

junipoll is a JUNIper router snmp POLLer. It is an mrtg-like utility that polls the counter values from the firewall filters configured on Juniper routers.

Spectrum::CLI

This is a perl module/class that provides a way to invoke the functionality of Spectrum Enterprise Manager’s Command Line Interface.

headers

This is a utility to identify C header files that should be included in C source files. You run it on your C source files, and it looks for your use of ANSI or POSIX identifiers and generates include directives for those source files.

Altoids

This is a distribution of my perl module that provides SNMP get/walk functionality with built-in name to OID translation alternatively by using .oid files.

Cflow

Cflow is a perl module for analyzing flow files written by flow-tools, argus, cflowd, etc.

IOSchat / getcnf

This is a distribution of my perl module that provides the ability to “chat” (ala expect(1), chat2.pl, or Comm.pl) with Cisco Internet Operating System routers. The getcnf utility, probably the useful component of this package, uses IOSchat to get the flash and/or running configuration. This output is suitable for storing in an RCS revision file, or to diff(1) to discover configurations changes.

NetCMS

NetCMS is a Network Configuration Management System for network devices. It supports various hardware including Cisco routers and switches.

NetTree

NetTree is a perl package to encapsulate the subnet allocations within a network. If you’re a hostmaster that manages a large number of subnets it may be useful to you to be sure that you don’t define overlapping subnets and such.

find_revisions

find_revisions is a reporting utility for use with RCS. This is the accompanying code for my article “Sys Admin File Revision Control with RCS” from the December 1998 issue of Sys Admin magazine.

ipwatch

ipwatch is a perl script that maintains a simple database of IP addresses that it has “seen.” It was written to help discover which IP addresses can be recovered by observing that they were not in use over some specified period of time - e.g. 1 month. (Once this tool is adopted, its data file serves as a rudimentary database to track assignment of IP addresses.)

ipwatch uses two methods to determine if a given IP address is in use: (1) by sending ICMP echoes (pings) and waiting for responses, and (2) by examining the ARP cache of specified routers or switches, via SNMP.

physaddrwatch

physaddrwatch is a perl script that maintains a simple database of physical media addresses (MAC or ATM) that it has “seen” and the IP or AppleTalk address to which they belong. It does this by walking various tables via SNMP on gateways (routers) that you specify.

iftop

iftop shows router top interfaces in real time, similarly to the way the Unix top command displays top CPU processes. This script is supplied with Simon Leinen’s “SNMP in Perl” distribution.

ip2anonip

ip2anonip is a perl script that can be used to anonymize or obfuscate IP addresses or translates IP addresses to hostnames, like ip2hostname. The anonymization employs the prefix-preserving technique implemented in tcpdpriv.

stpgraph

stpgraph discovers and graphs the topology of ethernet networks that utilize the spanning tree protocol.

Personal Information

Amateur radio station N9HZF operator.

Bicyclist and recreational inline skater.

Amateur artist, visual media.