

David Plonka

Akamai Technologies

email: plonka@akamai.com or dave@plonka.us

<http://www.cs.wisc.edu/~plonka/>

Madison, Wisconsin

Primary Research & Development Interests

Internet measurement, traffic classification, analytics, and anomaly detection ♦ Internet Protocol version 6 ♦ Global Internet application performance ♦ managing the IoT - Internet of Things ♦ Security vulnerabilities

Education

- Ph.D. Computer Sciences** 2013
University of Wisconsin - Madison, Madison, WI
Dissertation: “Rendezvous-based Measurement, Traffic Classification and Host Profiling”
Advisor: Paul Barford
- M.S. Computer Sciences** 2008
University of Wisconsin - Madison
- B.S. Computer Science and Artificial Intelligence majors, *magna cum laude*** 1991
Carroll College (now Carroll University), Waukesha, WI

Professional Experience

- Akamai Technologies, Inc.**, Cambridge, MA 2013 - present
Senior Research Scientist, Research Scientist (2013 - 2016)
- Computer Sciences Department, University of Wisconsin**, Madison, WI 2006 - 2013
Fellow (2009 - 2010), Research Assistant
- Google, Inc.**, Madison, WI Summer, 2010
Platforms Software Engineering Intern
- Division of Information Technology, University of Wisconsin** 2006 - 2013
Network Services Consultant
- Computer Sciences Department, University of Wisconsin** 2001 - 2006
Co-manager of the Wisconsin Advanced Internet Laboratory (WAIL)
- Division of Information Technology, University of Wisconsin** 1997 - 2006
Senior Systems Programmer / Network Engineer, Systems Programmer
- McHugh Freeman (became RedPrairie, merged with JDA Software)**, Waukesha, WI 1991 - 1997
Lead Systems Programmer, Systems Programmer
- Computer Science Department, Carroll College**, Waukesha, WI 1989 - 1991
Lab Assistant

Consulting Experience

- External Advisory Board member - MAMI 2015 - present
“Measurement and Architecture for a Middleboxed Internet” (an EU Horizon 2020 project)
- Independent Consultant 1997 - 1999, 2009

Awards and Honors

ACM Internet Measurement Conference (IMC 2016), Best Paper Award	2016
Google Engineering Intern Scholarship	2010
Lawrence H. Landweber NCR Fellowship in Distributed Systems	2009
USENIX 21st Large Installation System Administration Conference (LISA '07), Best Paper Award	2007
Carroll College Computer Science Department, Outstanding Computer Science Student	1991

Professional Activities

Chair

Measurement and Analysis for Protocols (MAP) Research Group co-chair, Internet Research Task Force (IRTF)	2016 - present
Internet Protocol Flow Information eXport (IPFIX) Working Group co-chair, Internet Engineering Task Force (IETF)	2001 - 2006

Program Committees

IEEE/IFIP Workshop on Mobile Network Measurement (MNM'17)	2017
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

United States Department of Energy (DOE)	2015
National Science Foundation (NSF)	2009

Reviewer

ACM SIGCOMM Conference (Special Interest Group on Data Communications)	2013
Internet2 Network Research Review Committee (NRRC)	2009 - 2010

Memberships

Internet Research Steering Group (IRSG), Internet Research Task Force (IRTF)	2016 - present
Luxurient Flowing Hair Club for Scientists (LFHCfS)	2015 - present
Union of Concerned Scientists (UCS)	2014 - present
IEEE	2008 - present
Association for Computing Machinery (ACM)	2002 - , 2008 - present

Patents

1. Paul R. Barford, Jeffery Thomas Kline, Sangnam Nam, David J. Plonka, and Amos Ron. "Method and Apparatus for Network Anomaly Detection", United States Patent 9680693, 2017.
2. Paul R. Barford and David J. Plonka. "Apparatus and Method for Classifying Network Packet Data", United States Patent 7907543, 2011.
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 via Wisconsin Alumni Research Foundation (WARF).

Peer-Reviewed Conference and Workshop Publications

1. David Plonka and Arthur Berger. “kIP: a Measured Approach to IPv6 Address Anonymization”, *Manuscript under review*, May, 2017.
2. Paweł Foremski, David Plonka, Arthur Berger. “Entropy/IP: Uncovering Structure in IPv6 Addresses”, In *Proceedings of the ACM Sixteenth Internet Measurement Conference (IMC 2016)*, Santa Monica, California, November, 2016. 25% accept rate.
3. Philipp Richter, Georgios Smaragdakis, David Plonka, Arthur Berger. “Beyond Counting: New Perspectives on the Active IPv4 Address Space”, In *Proceedings of the ACM Sixteenth Internet Measurement Conference (IMC 2016)*, Santa Monica, California, November, 2016. 25% accept rate.
Best Paper Award winner.
4. David Plonka and Arthur Berger. “Temporal and Spatial Classification of Active IPv6 Addresses”, In *Proceedings of the ACM Fifteenth Internet Measurement Conference (IMC 2015)*, Tokyo, October, 2015. 26% accept rate.
5. David Plonka and Paul Barford. “Assessing Performance of Internet Services on IPv6”, In *Proceedings of the Eighteenth IEEE Symposium on Computers and Communications (ISCC 2013)*, Split, Croatia, July, 2013. 49% accept rate.
6. 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.
7. 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.
8. 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.
9. 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.
10. 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.**
11. 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.
12. 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.
13. 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.
14. 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 Elisa Boschi, “The Internet of Things Old and Unmanaged”, Internet Architecture Board (IAB) Internet of Things Software Update Workshop (IoTSU), Trinity College Dublin, Ireland, June, 2016.
2. 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.
3. 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.

Journals and Periodicals

1. David Plonka. “The Internet of Things Unchecked”, *IETF Journal* Volume 12, Issue 2, November, 2016.
2. Dave Plonka. “An Analysis of Napster and Other IP Flow Sizes”, *Network Analysis Times*, April, 2001.
3. Dave Plonka. “How I Recovered Data I Thought I’d Lost”, *Sys Admin - the Journal for UNIX Systems Administrators*, August, 2000.
4. Dave Plonka. “Sys Admin File Revision Control with RCS”, *Sys Admin - the Journal for UNIX Systems Administrators*, December, 1998.
5. Dave Plonka. “Managing System Administration Tasks Using GNATS”, *Sys Admin - the Journal for Unix Systems Administrators*, February, 1997.

Online Technical Reports

1. Aaron Bergstralh, 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: “An Akamai Update on IPv6 Address Activity”, Federal IPv6 Face-to-Face Meeting, McLean, Virginia, October, 2017.
2. Invited talk: “kIP: a Measured Approach to IPv6 Address Anonymization”, Measurement and Analysis for Protocols Research Group (maprg), 99th IETF meeting, Prague, July, 2017.
3. Invited talk, WiP: “A New MAP: Measurement and Analysis for Protocols Research Group”, ACM Sixteenth Internet Measurement Conference (IMC 2016), Santa Monica, California, November, 2016.
4. Invited talk, WiP: “A Continuing Study of the Active IPv6 Address Space”, ACM Sixteenth Internet Measurement Conference (IMC 2016), Santa Monica, California, November, 2016.
5. Invited talk, “Privacy Negotiation for TLS - Selectable SNI or SNO: Server Name Omission”, TCP Increased Security (tcpinc) Working Group, 96th IETF meeting, Berlin, July, 2016.

6. Invited talk, “Measurement and Analysis for the Internet of Things”, Measurement and Analysis for Protocols (proposed) Research Group (maprg), 96th IETF meeting, Berlin, July, 2016.
7. Invited participant, Future Internet Dialogue, Berlin, Germany, July, 2016.
8. Invited talk, “An Akamai View of Global IPv6 Address Activity”, Trinity College Dublin, Ireland, June, 2016.
9. Invited participant, Internet Architecture Board (IAB) Internet of Things Software Update Workshop (IoTSU), Trinity College Dublin, Ireland, June, 2016.
10. Invited talk, “An Akamai View of Global IPv6 Address Activity”, Networking @Scale Conference, Menlo Park, California, May, 2016. **Attendee-rated Most Informative Talk.**
11. Invited talk, “IPv6 Prefix Intelligence”, Measurement and Analysis for Protocols (proposed) Research Group (maprg), 95th IETF meeting, Buenos Aires, Argentina, April, 2016.
12. Invited talk, “Exploring Global IPv6 Address Activity”, PROTOSEC 2016, Buenos Aires, Argentina, April, 2016.
13. Invited talk, “Monitoring and Classification of Active IPv6 Addresses”, FloCon 2016 Conference, Daytona Beach, Florida, January, 2016.
14. Invited talk, “Temporal and Spatial Classification of Active IPv6 Addresses”, Massachusetts Institute of Technology, Cambridge, Massachusetts, November, 2015.
15. Invited talk, “Temporal and Spatial Classification of Active IPv6 Addresses”, IPv6 Operations (v6ops) Working Group, 94th IETF meeting, Yokohama, November, 2015.
16. Invited participant, 26th Messaging/Malware/Mobile Anti-Abuse Working Group (M3AAWG/MAAWG) General Meeting, Baltimore, October, 2012.
17. Invited talk, “Rendezvous-based Traffic Classification, Measurement, and Analysis”, ISC/CAIDA Data Collaboration Workshop, Baltimore, October, 2012.
18. Invited talk, “A Rendezvous-based Paradigm for Analysis of Solicited and Unsolicited Traffic”, 1st International Workshop on Darkspace and Unsolicited Traffic Analysis (DUST 2012), San Diego, May, 2012.
19. Invited talk, “Rendezvous-based Network Traffic Analysis”, Lockdown 2009 Conference, UW-Madison, Madison, Wisconsin, July, 2009.
20. Invited talk, “Rendezvous-based Network Traffic Analysis”, Massachusetts Institute of Technology, Cambridge, Massachusetts, February, 2009.
21. Invited talk, “Rendezvous-based Network Traffic Analysis”, Boston University, Boston, February, 2009.
22. Invited talk, WiP: “Network Admins are Programmers: An Analysis of Network Management Artifacts”, LISA 2008, San Diego, November, 2008.
23. Invited talk, BoF: “DNS and Network Traffic Analysis with TreeTop”, LISA 2008, San Diego, November, 2008.
24. Invited participant and session moderator, “NSF NeTS Future Internet Design (FIND) Student Meeting”, Seattle, August, 2008.
25. 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.
26. Invited participant, “Wisconsin Entrepreneurial Bootcamp”, UW-Madison, Madison, Wisconsin, July/August, 2007.

27. Invited talk, "Anomalous Traffic from Internet Consumer Products", ESCC/Internet2 Joint Techs Workshop, Madison, Wisconsin, July, 2006.
28. Invited talk, "Get Your FIX: Flow Information eXport Analysis and Visualization", ESCC/Internet2 Joint Techs Workshop, Madison, Wisconsin, July, 2006.
29. 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.
30. Invited talk, "AANTS: Web-Based Tools for Cooperative Campus Network Administration", Fall 2005 Internet2 Member Meeting, Philadelphia, September, 2005.
31. Invited participant, Sampling 2005 Workshop, Paris, July, 2005.
32. Invited talk, "Exposing Abuse in Internet Traffic Measurements", Lockdown 2005 Conference, UW-Madison, Madison, Wisconsin, July, 2005.
33. Invited talk, "Flow-based Measurements and Other Means to Expose Internet Pathology", Institut Eurecom, Sophia Antipolis, France, December, 2004.
34. Invited talk, "Bare-Bones Measurement Data Archiving", ISMA Data Catalog 2004 Workshop, CAIDA/SDSC, La Jolla, California, June, 2004.
35. Invited talk, "Embedding Globally-Routable Internet Addresses Considered Harmful", Global Routing Operations (GROW) Working Group, 58th IETF meeting, Minneapolis, November, 2003.
36. Invited talk, "A Case Study in Internet Pathology: Flawed Routers Flood University's Network", LISA 2003, San Diego, October, 2003.
37. Invited talk, "Flawed Routers Flood University of Wisconsin Internet Time Server", NANOG 29, Chicago, October, 2003.
38. Invited talk, "IP Flow Measurement and Analysis with FlowScan", Duke University, Durham, North Carolina, September, 2003.
39. Invited talk, "A Case Study in Internet Pathology: Flawed Routers Flood University's Network", Duke University, Durham, North Carolina, September, 2003.
40. Invited talk, "Flawed Routers Flood University of Wisconsin Internet Time Server", MAD-SAGE meeting, Madison, Wisconsin, August, 2003.
41. Invited talk, "Internet Sink Deployments", IEPG meeting, Vienna, Austria, July, 2003.
42. Invited talk, "IP Flow Measurement and Analysis with FlowScan", IPAM Mini-Workshop on "Internet-wide Measurements", Los Angeles, March, 2002.
43. Invited talk, "FlowScan at the University of Wisconsin-Madison", EDUCAUSE 2001, Indianapolis, October, 2001.
44. Invited talk, "Practical Flow-based Passive Measurements and FlowScan", Computer Sciences Department Network Seminar, UW-Madison, Madison, Wisconsin, June, 2001.
45. Invited talk, "Better Campus Network Health Through Instrumentation", UW System Network Operators Meeting, Madison, Wisconsin, March, 2001.
46. Invited talk, "FlowScan - LIVE!", NANOG 21, Atlanta, February, 2001.
47. Invited talk, "FlowScan - a Network Traffic Reporting and Visualization Tool", MAD-SAGE meeting, Madison, Wisconsin, November 2000.

48. Invited talk, BoF: “Flow-based Passive Measurements”, CANARIE/NLANR/Internet2 Techs Workshop, Toronto, August, 2000.
49. Invited talk, “Unix Security”, UW-Madison Unix Fair, Madison, Wisconsin, August, 1999.
50. Invited participant, CAIDA ISMA Workshop on “Passive Measurement Data and Analysis”, La Jolla, California, January, 1999.
51. 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, Fall 2011, Fall 2012
 Guest lecturer on Ethernet, Wireless Ethernet, IP Multicast, TCP, DNS

Student Mentoring

Anant Shah Intern, Akamai Technologies Ph.D. candidate, Colorado State University, expected 2017	Summer 2017
Eric W. Gaston Master of Science in Applied Cyber Operations, Naval Postgraduate School, 2017 Co-Advised Capstone Project: “High-Frequency Mapping of the IPv6 Internet Using Yarrp”	Fall 2016 - Spring 2017
Ramakrishna Padmanabhan Intern, Akamai Technologies Ph.D. student, University of Maryland, College Park	Summer 2016 - present
Philipp Richter Intern, Akamai Technologies Ph.D., Technische Universität Berlin, 2017	2015 - present
Paweł Foremski Intern, Akamai Technologies Ph.D. candidate, Institute of Theoretical and Applied Informatics, Polish Academy of Sciences	2015 - 2016
Ramakrishnan Durairajan Intern, Akamai Technologies Ph.D., Computer Sciences, University of Wisconsin-Madison, 2017	2015

Authored Software

timemail

The Timemail system consists of a stand-alone script, a database, and tools to analyze the structure and performance of the email delivery network over time.

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 most 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 encapsate 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 to translate 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.