DECEMBER 2025

JOHN DUNAGAN

Redmond, WA 98052| dunaganj@amazon.com | jdunagan@hotmail.com

CAREER

Amazon Web Services

Principal Research Scientist / Principal Algorithm Engineer / Senior Principal (as of April 2020) **2012-present**

Amazon Web Services provides a public cloud for hosting computation and data. During my time at Amazon Web Services, I have worked on Amazon Kinesis, AWS Fraud Prevention, EC2 Security, the EC2 Control Plane, EC2 Capacity Management, and AWS Supply Chain Automation. Since August 2023, I work on the AWS Commerce Platform. As of December 2025, I am the senior-most engineer in an organization of ~1,700, and I have recently served as the Executive Sponsor for one of AWS's customers.

Microsoft High Performance Computing Team

Principal Architect / Principal Program Manager **2009-2012**

During my time on the High Performance Computing team, the team was first part of the Technical Computing Group and then later part of Windows Azure. The Technical Computing Group focused on building software that helps scientists and engineers to do modeling and simulation. Windows Azure provides a public cloud for hosting computation and data.

On the High Performance Computing team, I was responsible for the productization of DSC, Dryad, and DryadLINQ, a suite of technologies related to the Big Data trend.

Towards the end of my time on the High Performance Computing team, I changed from being a Principal Architect to being a Principal Program Manager.

Microsoft Research

Researcher **2002-2009**

I worked in the Networking Research Group, the Distributed Systems and Security Group, the Systems Management Group, the Systems and Networking Group, and the Theory Group.

My publication record and product group collaborations (described in more detail below) occupied most of my time in research, though I also put in some long hours helping Microsoft meet its EU antitrust obligations.

EDUCATION

Massachusetts Institute of Technology

PhD in Mathematics **1998-2002**

Thesis: A Geometric Theory of Outliers and Perturbation

Massachusetts Institute of Technology

BS in Mathematics with Computer Science **1994-1998**

AWARDS AND SERVICE

Technical Program Committee for USENIX Annual Technical Conference **2009**

Technical Program Committee for International Workshop on Peer-to-Peer Systems (IPTPS) **2008**

Attendee to Department of Energy Workshop for Cyber Security Research Needs for Open Science **2007**

Member of Organizing Committee for National Academy of Engineering's
US Conference on Frontiers of Engineering **2007**

Attendee to National Academy of Engineering's US Conference on Frontiers of Engineering **2006 and 2005**

Charles and Holly Housman Award for Excellence in Teaching
Massachusetts Institute of Technology **2002**

Lester Wolfe Fellowship
Massachusetts Institute of Technology **2002**

LANGUAGES

English – native language
Spanish – speak, read and write with basic competence

MEMBERSHIPS

ACM, IEEE, Usenix

PATENTS

39 filed within US; 2 filed outside US.

INVITED TALKS

"Datacenter Research at Microsoft"
Institute for Defense Analysis, Virginia **2009**

PRODUCT GROUP COLLABORATIONS WHILE IN MICROSOFT RESEARCH

Forefront Credential Management: collaborated on development of technologies to prevent escalation of privilege attacks using credentials.

Live Mesh Partitioning: collaborated on development of a partitioning and recovery engine that supports multiple services within Microsoft's Live Mesh offering.

MSN Distributed Cache: collaborated on development of a distributed cache incorporating load balancing and admission control.

MSN Logging and Monitoring Infrastructure: collaborated on development of a logging and monitoring system that uses application-level failures to adapt the data it collects from across a distributed system.

PUBLICATIONS

- **Machine Learning to Diagnose Failures in Server Tiers:**
  Bilinear Logistic Regression for Factored Diagnosis Problems.
  Appeared in Workshop on Managing Large-Scale Systems via the Analysis of System Logs and the Application of Machine Learning Techniques (SLAML) 2011; later appeared in the journal Operating Systems Review (OSR).
  Joint work with Sumit Basu, Kevin Duh and Kiran-Kumar Munuswamy-Reddy.

- **Machine Learning to Partition Graphs:**
  Active Graph Reachability Reduction for Network Security and Software Engineering.
  Appeared in International Joint Conference on Artificial Intelligence (IJCAI)2011.
  Joint work with Alice Zheng and Ashish Kapoor.

- **Scaling out Middle Tiers:**
  Centrifuge: Integrated Lease Management and Partitioning for Cloud Services.
  Appeared in Symposium on Networked Systems Design and Implementation (NSDI) 2010.
  Joint work with Atul Adya and Alec Wolman.

- **Data Placement for Cloud Services:**
  Volley: Automated Data Placement for Geo-Distributed Cloud Services.
  Appeared in Symposium on Networked Systems Design and Implementation (NSDI) 2010.
  Joint work with Sharad Agarwal, Alec Wolman, Navendu Jain, Stefan Saroiu and Harbinder Bhogan.

- **Connecting Middle Tiers to Scale-out Storage:**
  Stout: Performance, Consistency, and Agility across Scalable Cloud Stores
  Appeared in USENIX Annual Technical Conference 2010.
  Joint work with Alec Wolman, Alec Snoeren and John McCullough.

- **Managing Enterprise Security Configuration:**
  Heat-ray: Combating identity snowball attacks using machine learning, combinatorial optimization and attack graphs
  Appeared in *Symposium on Operating Systems Principles (SOSP) 2009*.
  Joint work with Alice Zheng and Dan Simon.

- **Characterizing Botnets from Spam:**
  Characterizing Botnets from Email Spam Records
  Appeared in *Workshop on Large-Scale Exploits and Emergent Threats (LEET) 2008.*
  Joint work with Li Zhuang, Daniel R. Simon, Helen J. Wang, Ivan Osipkov, Geoff Hulten and J. D. Tygar.

- **Investigating PC Unresponsiveness:**
  Why Did My PC Suddenly Slow Down?
  Appeared in *Workshop on Tackling Computer Systems Problems with Machine Learning Techniques (SysML) 2007.*
  Joint work with Sumit Basu and Greg Smith.

- **A New Algorithm for Solving Systems of Linear Equations:**
  Iteratively Constructing Preconditioners via the Conjugate Gradient Method.
  Appeared in *Symposium on Theory of Computing (STOC) 2007.*
  Joint work with Nicholas J. A. Harvey.

- **Infrastructure for Analyzing Network Protocols:**
  Generic Application-Level protocol Analyzer and its Language.
  Appeared in *Network and Distributed System Security (NDSS) 2007.*
  Joint work with David Brumley, Nikita Borisov, Helen Wang, Pallavi Joshi and Chuanxiong Guo.

- **Protecting Web Browsers using JavaScript Virtualization:**
  BrowserShield: Vulnerability-Driven Filtering of Dynamic HTML.
  Appeared in *Operating Systems Design and Implementation (OSDI) 2006.*
  Joint work with Charlie Reis, Helen Wang, Opher Dubrovsky, and Saher Esmeir.

- **Measuring Shellcode Polymorphism:**
  Finding Diversity in Remote Code Injection Exploits.
  Appeared in *Internet Measurement Conference (IMC) 2006.*
  Joint work with Justin Ma, Helen Wang, Stefan Savage, and Geoffrey Voelker.

- **Applying Machine Learning to Analyzing Network Protocols:**
  Automatically Extracting Fields from Unknown Network Protocols.
  Appeared in *Systems and Machine Learning Workshop (SysML) 2006*
  Joint work with Karthik Gopalratnam, Sumit Basu, and Helen Wang.

- **Security Access Tracing:**
  A Black-Box Tracing Technique to Identify Causes of Least-Privilege Incompatibilities.
  Appeared in *Network and Distributed System Security Symposium (NDSS) 2005.*
  Joint work with Shuo Chen, Chad Verbowski, and Yi-Min Wang.

- **Failure Detection and Notification:**
  FUSE: Lightweight Guaranteed Distributed Failure Notification.
  Appeared in *Operating Systems Design and Implementation (OSDI) 2004.*
  Joint work with Nicholas J. A. Harvey, Michael B. Jones, Dejan Kostic, Marvin Theimer, and Alec Wolman.

- **Channel Hopping Strategies for Wireless Networks:**
  SSCH: Slotted Seeded Channel Hopping for Capacity Improvement in IEEE 802.11 Ad-Hoc Wireless Networks.
  Appeared in *International Conference on Mobile Computing and Networking (Mobicom) 2004.*
  Joint work with Ranveer Chandra and Victor Bahl.

- **Black-box Troubleshooting:**
  STRIDER: A Black-box, State-based Approach to Change and Configuration Management and Support.
  Appeared in *Usenix Large Installation System Administration Conference (LISA) 2003 (Best Paper).*
  Joint work with Yi-Min Wang, Chad Verbowski, Yu Chen, Helen J. Wang, Chun Yuan, and Zheng Zhang.

- **Patch Management:**
  Towards A Self-Managing Software Patching Process Using Black-Box Persistent State Manifests.
  Appeared in *International Conference on Autonomic Computing (ICAC) 2004.*
  Joint work with Roussi Roussev, Brad Daniels, Aaron Johnson, Chad Verbowski, and Yi-Min Wang.

- **Peer-to-peer Overlay Networks:**
  SkipNet: A Scalable Overlay Network with Practical Locality Properties.
  Appeared as *Microsoft Research Technical Report TR-2002-92.*
  Joint work with Nicholas J. A. Harvey, Michael B. Jones, Stefan Saroiu, Marvin Theimer, and Alec Wolman.

- **Overlay Multicast Trees:**
  Subscriber/Volunteer Trees: Polite, Efficient Overlay Multicast Trees.
  Appeared as *Microsoft Research Technical Report TR-2004-131.*
  Joint work with Nicholas J. A. Harvey, Michael B. Jones, Marvin Theimer, and Alec Wolman.

- **Experience Implementing a Peer-to-peer System**:
  Engineering Realities of Building a Working Peer-to-Peer System.
  Appeared as *Microsoft Research Technical Report TR-2004-54.*
  Joint work with Michael B. Jones.

- **Thesis:**
  A Geometric Theory of Outliers and Perturbation.
  The thesis contains most of the work in the paper on outliers and the two papers on smoothed analysis of linear programming plus some additional discussion.

- **A New Algorithm for Linear Programming:**
  A Polynomial-time Rescaling Algorithm for Solving Linear Programs.
  Appeared in *Symposium on Theory of Computing (STOC) 2004*, and then in the journal *Mathematical Programming: Series A.*
  Joint work with Santosh Vempala.

- **Understanding Linear Programming in the Presence of Noise:**
  Smoothed Analysis of the Perceptron Algorithm for Linear Programming.
  Joint work with Avrim Blum.
  Appeared in *Symposium on Discrete Algorithms (SODA) 2002,* and then in the *Journal of Algorithms*.

  Smoothed Analysis of the Renegar's Condition Number for Linear Programming.
  Appeared in *SIAM Conference on Optimization (SIOPT) 2002.*
  Joint work with Dan Spielman and Shang-Hua Teng.

- **Outliers:**
  Optimal Outlier Removal in High-Dimensional Space.
  Appeared in *Symposium on Theory of Computing (STOC) 2001, a*nd then in the *Journal of Computer and System Sciences (JCSS).*
  Joint work with Santosh Vempala.

- **Approximation Algorithms:**
  On Euclidean Embeddings and Bandwidth Minimization.
  Appeared in *Workshop on Randomization and Computation (RANDOM) 2001.*
  Joint work with Santosh Vempala.