Raja R. Sambasivan

123 Sewall Ave. Apt. 3G
Brookline, MA 02446
412-983-1701
rrs AT bu.edu
http://www.rajasambasivan.com

Boston University
Harriri Institute for Computing
111 Cummington Mall, Room 172
Boston, MA 02215

Objective

A tenure-track position in the Computer Science or Computer Engineering Department of a leading university.

Summary

I am a systems and networking researcher, broadly interested in designing and building evolvable multi-party systems for the cloud ecosystem and the tools needed to diagnose problems in them. I am published in top conferences, including ATC, FAST, InfoVis, NSDI, SoCC, SIGCOMM, and SIGMETRICS (~770 citations total). I have also developed and taught a graduate-level class on cloud computing.

Education

Ph.D., Electrical & Computer Engineering, May 2013
Carnegie Mellon University, Pittsburgh, PA
Advisor: Greg Ganger
Dissertation: Diagnosing performance changes in distributed systems by comparing request flows

M.S., Electrical & Computer Engineering, May 2004
Carnegie Mellon University, Pittsburgh, PA

B.S., Electrical & Computer Engineering w/minor in Computer Science, May 2003
Carnegie Mellon University, Pittsburgh, PA

Honours & Awards

PI on a NSF-small grant, *A just in-time, cross-layer instrumentation framework for diagnosing performance problems in distributed applications*. $460,294. October 2018 to October 2021.

Best poster, EMC University Day 2012 (*Diagnosing performance changes by comparing request flows*).

Best paper, SIGMETRICS 2007 (*Modeling the relative fitness of storage*).

Featured in Piled Higher & Deeper. (PhDComics), February 14^{th}, 2007. strip.

Best paper, FAST 2005 (*Ursa Minor: versatile cluster-based storage*).

Research Summaries

Dec. 2017 Present

Automatic, just-in-time, cross-layer instrumentation: Diagnosing problems in distributed applications is extremely challenging. To help, distributed applications and the stack layers the run atop of (e.g., virtualization, OS) are instrumented heavily. But, it is difficult to know a priori where to add instrumentation (e.g., in which application node or lower stack layer) and what instrumentation to add (e.g., specific trace points or logging points) to help diagnose problems that may occur far in the future. In this research, we are creating a continuously running instrumentation framework that continuously searches the space of possible instrumentation choices to enable that needed to help engineers diagnose a new problem.

March 2018 Present

Higher-level abstractions for diagnosis Developers use powerful API-based abstractions to mitigate the amount of complexity they must deal with at any one time when building distributed applications. These abstractions have enabled developers to create extremely large-scale, complex applications. But, to diagnose problems observed in deployed versions of these applications, engineers are often forced to use no abstractions whatsoever (i.e., use flat logs or traces of low-level events). We are exploring the utility of a novel abstraction, the workflow motif, instances of which describe frequent or important processing patterns observed in the workflows of distributed applications' runtime behavior. We expect that workflow motifs will improve the efficacy of existing diagnosis tools and enable a wide range of new use cases.

June 2017 – Present

Flexibility for datacenter networks: As more organizations move their diverse applications to the cloud, the requirements placed on datacenters' networks continue to grow. To help, the networking community has proposed a variety of new protocols and technologies suited to these applications' needs. But, today's datacenters are architected to include only a single "one-size-fits-all" network, setting an extremely high bar for these new proposals. Even if datacenters could accommodate multiple technologies, applications only have coarse-grained options for expressing their networking requirements. To help, we are developing a networking interface that will both allow applications to express their requirements and allow for multiple datacenter networks (physical or virtual). We expect this interface will improve resource utilization within existing networks, enable deployment of new protocols and technologies, and enable a marketplace of networking providers within datacenters.

Aug. 2014 – Aug. 2017

Evolvability for inter-domain routing: BGP, the Internet's inter-domain routing protocol, is a foundational part of the Internet's architecture. It is responsible for connecting all the services we hold dear. It is also riddled with critical flaws and limitations. Numerous critical fixes and replacement protocols have been proposed, but almost none have been deployed because BGP is architecturally rigid and cannot facilitate the introduction of new protocols. In this research, we identified the features needed within any inter-domain routing protocol to allow it to bootstrap evolution to new inter-domain routing protocols—i.e., facilitate their deployment and gradually deprecate itself in favor of one or more new protocols. We created a version of BGP (called D-BGP) with these features and showed that it can support a rich Internet composed of any recently-proposed protocols. We demonstrated that D-BGP accelerates the rate at which early adopters see the benefits of several types of new protocols.

Jan. 2014 – Jan. 2016

Resource-aware routing: Content sources (e.g., CDNs), transit ISPs, and eyeball networks have a vested interest in sharing resource demands because of the large amount of traffic they exchange (e.g., in 2009, over 50% of all Internet traffic originated from less than 30 content sources). Since this information is not communicated by BGP, content sources can shift vast amounts of traffic onto unsuspecting ISPs as they search for resource-abundant paths that allow them to meet their goals (e.g., deliver high-quality video). At the same time, ISPs may only advertise resource-constrained paths to content sources, preventing them from meeting their goals. In this work, we explored a new inter-domain routing protocol that allows ISPs and eyeballs to create routing paths with knowledge of content sources' resource desires.

May 2012 – Oct. 2016

Principled workflow-centric tracing of distributed systems: Workflow-centric tracing (**also called end-to-end tracing**) captures the workflow of how individual requests are serviced within and among the components of a distributed system. It is an extremely powerful substrate on which to layer future automated-management tools (e.g., for diagnosis or resource provisioning). However, today, there is a lack of clarity with regards to when tracing infrastructure designs need to be different. This has led to both a proliferation of tracing designs and a mistaken belief that a single tracing design can be a "one-size-fits-all" solution for a variety of management tasks. For this research, we used our experiences building workflow-centric tracing infrastructures and building tools on top of them to distill the key design axes of tracing. We described which management tasks require different design choices for these axes.

Sept. 2006 – Sept. 2013

Diagnosing performance changes by comparing request flows: The first step of engineers' diagnosis efforts is to localize the source of a newly-observed problem from many entities that may contain the root cause (application nodes, switches, etc.) to a few potential culprits. This is extremely challenging. My dissertation focused on a technique, called request-flow comparison, that automatically localizes the source of a performance degradation from the many application nodes that could be involved in its processing to just a culprits. We implemented request-flow comparison in a tool called **Spectroscope** and demonstrated its effectiveness by using it to diagnose real performance problems observed in a distributed storage application called Ursa Minor and in certain Google infrastructure applications. Via a 26-person user study, we found that different visualizations for presenting request-flow comparison's results are best for different problem types. Request-flow comparison builds on workflow-centric tracing.

Jun. 2005 – May 2010

Creating a transparently scalable metadata service for distributed storage: Many object-based storage services, such as the Google File System, are scalable only with regard to data, not metadata. In this work, we explored simple techniques for building a scalable metadata service for Ursa Minor. The key challenge we addressed was how to minimize the complexity of handling multi-server operations, such as RENAMES. Our

approach was twofold. We first minimized such occurrences by automatically co-locating metadata for files that resided in similar parts of the filesystem namespace. In cases where the metadata involved resided on multiple servers, we used the migration functionality already present in most systems to move one item to the other item's server. This is a heavy-handed approach, but we found that it worked with little performance impact for many common workloads.

Jul. 2007 – Sept. 2007 **Improving the accuracy of query-progress indicators for data warehouses**: Queries in business intelligence workloads are often resource intensive and slow to complete. For these queries, accurate query progress estimation is paramount to help operators decide whether to kill a currently long-running query in favor of other more important queries. While an intern at HP Labs, I explored the use of statistical methods for creating a progress indicator for Neoview, HP's data warehousing solution.

Jul. 2005 – Dec. 2006 **//trace: Enabling accurate trace replay for parallel applications**: Since company security policies often prohibit explicit application sharing, it is difficult for storage vendors to evaluate their storage systems under a potential client's expected workload. One alternative is trace replay, but the many inter-node dependencies exhibited by most parallel applications make them hard to replay accurately. To address this problem, I helped create //trace, a program for extracting and replaying traces of parallel applications to recreate their I/O behavior. Its tracing engine automatically discovers inter-node data dependencies and inter-I/O compute times for each node in a parallel application. //trace embeds this information in per-node annotated I/O traces, allowing its parallel replayer to closely mimic the behavior of a traced application.

Jun. 2003 – Sept. 2005 **Evaluating replication policies for layered clustering of NFS servers**: Layered clustering (or NAS aggregation) offers cluster-like load balancing for unmodified NFS servers. Read requests sent to a busy server can be offloaded to other servers holding replicas of the accessed files. In this research, we explored a key design question for this approach: which files should be replicated? By conducting a trace-based study, we found that the popular policy of replicating read-only files offers little benefit. A policy that replicates read-only portions of read-mostly files, however, implicitly coordinates with client cache invalidations and thereby allows almost all read operations to be offloaded. In a read-heavy trace, we found that 75% of all operations and 52% of all data transfers can be offloaded from a busy server.

Professional Experience

Research Scientist **November 2016 – Present**
Boston University, Mass Open Cloud & Red Hat Collaboratory

- *Host*: Orran Krieger
- Helping grow a new research lab by mentoring students and suggesting research directions
- Collaborating with Red Hat to find research topics of interest to industry and academia
- Started / leading a research group that focuses on automating problem diagnosis in clouds
- Mentoring students on topics: problem diagnosis, big-data frameworks, datacenter networking
- Organizer for the Red Hat Collaboratory Colloquium
- Guest lecturer for select topics in BU's undergraduate cloud computing class

Postdoctoral Researcher **June 2013 – September 2016**
Carnegie Mellon University, eXpressive Internet Architecture (XIA) Group

- *Host*: Peter Steenkiste
- Explored mechanisms to enable evolvability for inter-domain routing
- Explored resource-aware routing protocols for use by CDNs and clouds
- Helped write a NSF NeTS large proposal (*Design principles for a future-proof Internet control plane*)
- Co-wrote NSF annual reports
- Co-developed and co-taught CMU's inaugral Master's/PhD-level cloud-computing class

Consultant **November 2014 – December 2014**
Huawei, storage group

- Identified best practices for problem diagnosis in storage systems
- Provided advice on how to implement request-flow comparison for use in select products
- Helped design a workflow-centric tracing infrastructure for use in select products

Graduate Student **June 2006 – May 2013**
Carnegie Mellon University, Parallel Data Lab

- *Advisor*: Greg Ganger
- Developed techniques for tracing distributed systems and automatically localizing problems within them
- Collaborated with fellow graduate students on the following research topics:
  - Building scalable metadata services for distributed storage systems
  - Building an accurate MPI-based trace replayer for HPC applications
  - Applying machine learning to predict workload performance
- Developed and maintained Ursa Minor's NFS server and workflow-centric tracing infrastructure

Software Engineering Intern **May 2010 – December 2010**
Google

- *Mentor*: Michael De Rosa
- Implemented Spectroscope on top of Dapper, Google's workflow-centric tracing infrastructure
- Demonstrated Spectroscope's utility in helping diagnose real performance degradations
- Helped create a Dapper visualization that shows service inter-dependencies annotated with perf. metrics

Research Intern **July 2007 – March 2008**
HP Labs

- *Mentor*: Kivanc Ozonat
- Explored how to create a statistical query-progress indicator for Neoview, HP's enterprise data warehouse
- Created tools for visualizing the complex query execution plans created by Neoview

Systems Programmer **June 2004 – May 2006**
Carnegie Mellon University, Parallel Data Lab

- Core member of the Ursa Minor development team
  - Responsible for the development and maintenance of Ursa Minor's NFS server
  - Led an effort to ensure SPEC SFS compatibility with Ursa Minor
  - Developed tools for visualizing Ursa Minor's performance on key benchmarks over time
  - Co-led an effort to improve small-file performance via intelligent metadata prefetching
- Implemented a tee that compared the functionality of a NFSv3/v4 server to that of a reference impl
- Created a NFS trace-based simulator for evaluating NAS-replication policies

Teaching Experience

Course Co-developer & co-instructor **Fall 2013**
CS 15-719, Advanced Cloud Computing

This class provides an overview of cloud-computing concepts via a curriculum that emphasizes reading research papers, lectures, projects, and exams. It usually consists of 30-40 Master's students and a few PhD students.

- *Co-instructors*: Garth Gibson, Majd Sakr, Greg Ganger
- Drove design of course syllabus
- Created and delivered lectures, created exams, and held regular office hours
- Created and supported two projects involving running Hadoop jobs within AWS and implementing a load balancer for OpenStack

Teaching Assistant **Fall 2005 & Spring 2010**
ECE 18-746, Storage Systems

This class covers a broad range of material, including hard-disk architecture, file-system design and debugging, RAID, and object-based storage. It usually consists of 40-60 Master's students and PhD students.

- Created and supported a project that required students to build a RAID-5 controller for iSCSI
- Created and graded both tests and projects and held regular office hours

Professional service

Program committee member:

- Transactions on Services Computing 2020, 2015
- Transactions on Software Engineering 2017
- HotStorage 2014, 2013

External reviewer: SIGCOMM CCR 2018, EuroSys 2017

Panel member: NSF Cloud Review 2017, NSF 2016

Session chair: NeNS 2017, SoCC 2016

Reading groups: BU diagnosis (2017-present), CMU network diagnosis (2013-2014), CMU diagnosis (2011-2013)

Outreach

Letter Writer, Letters to a Pre-Scientist middle-school outreach program (2018)

Research Mentor, MIT Primes high-school research program (2018, 2017)

- Mentees named Siemens Competition semi-finalists (2017)
- Mentees selected to present at Red Hat Developers' conference Devconf.us (2018)

CS Grand Awards Judge, Intel Science & Engineering Fair Finals (2015, 2012)

Presenter, Carnegie Science Center Buhl Planetarium (2010)

Mentoring

Current students:

- Mania Abdi, PhD, North Eastern University
  The workflow motif: a powerful, widely-useful abstraction for diagnosing distributed-application problems
  Instrumenting Ceph with workflow-centric tracing
- Emre Ates, PhD, Boston University
  Pythia: A just-in-time, cross-layer instrumentation framework for distributed applications
- Golsana Ghaemi, PhD, Boston University
  The workflow motif: a powerful, widely-useful abstraction for diagnosing distributed-application problems
  Instrumenting Ceph with workflow-centric tracing
- Emine Ugur Kaynar, PhD, Boston University
  D3N: A multi-layer cache for improving big-data applications' performance
- Lily Sturmann, MS, Harvard Extension School
  Pythia: A just-in-time, cross-layer instrumentation framework for distributed applications
- Mert Toslali, PhD, Boston University
  Pythia: A just-in-time, cross-layer instrumentation framework for distributed applications

Past students

- David Tran-Lam, PhD, University of Wisconsin-Madison
  Bootstrapping evolvability for inter-domain routing with D-BGP
- Harshal Sheth, High School'18, Westford Academy
  Skua: Support for workflow-centric tracing in the Linux Kernel
  Tarpan: A router that supports evolvability
- Andrew Sun, High School'18, Westford Academy
  Skua: Support for workflow-centric tracing in the Linux Kernel
  Tarpan: A router that supports evolvability
- Jethro Sun, MS'18, Boston University
  FlexNet: Enabling flexibility in cloud networks
- William Wang, MS'11, Carnegie Mellon University
  Adding workflow-centric tracing to HDFS
- Da Yu, PhD, Brown University
  FlexNet: Enabling flexibility in cloud networks

Raja R. Sambasivan

Funding

A just-in-time, cross-layer instrumentation framework for diagnosing performance problems in distributed applications. Raja R. Sambasivan, Ayse K. Coskun, Orran Krieger. $460,249. Award # CNS-1815323. October 2018 to October 2021.

Refereed Publications

Bootstrapping evolvability for inter-domain routing with D-BGP. Raja R. Sambasivan, David Tran-Lam, Aditya Akella, Peter Steenkiste. In proceedings of the ACM 2017 SIGCOMM Conference (SIGCOMM'17). August 21^{th} to August 25^{th}, 2017. Los Angeles, CA, USA.

Principled workflow-centric tracing of distributed systems. Raja R. Sambasivan, Ilari Shafer, Jonathan Mace, Rodrigo Fonseca, Gregory R. Ganger. In proceedings of the 7^{th} ACM Symposium on Cloud Computing (SoCC'16). October 5^{th} to October 7^{th}, 2016. Santa Clara, CA, USA.

Bootstrapping evolvability for inter-domain routing. Raja R. Sambasivan, David Tran-Lam, Aditya Akella, Peter Steenkiste. In proceedings of the 14^{th} ACM Workshop on Hot Topics in Networks (HotNets'15). November 16^{th} to November 17^{th}, 2015. Philadelphia, PA, USA.

Visualizing request-flow comparison to aid performance diagnosis in distributed systems. Raja R. Sambasivan, Ilari Shafer, Michelle L. Mazurek, Gregory R. Ganger. IEEE Transactions on Visualization and Computer Graphics 19(12), December 2013. In proceedings of Information Visualization 2013.

Specialized storage for big numeric time series. Ilari Shafer, Raja R. Sambasivan, Anthony Rowe, Gregory R. Ganger. In proceedings of the 5^{th} USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage'13). June 27^{th} to June 28^{th}, 2013. San Jose, CA, USA.

Automated diagnosis without predictability is a recipe for failure. Raja R. Sambasivan, Gregory R. Ganger. In proceedings of the 4^{th} USENIX Workshop on Hot Topics in Cloud Computing (HotCloud'12). June 12^{th} to June 13^{th}, 2012. Boston, MA, USA.

Diagnosing performance changes by comparing request flows. Raja R. Sambasivan, Alice X. Zheng, Michael De Rosa, Elie Krevat, Spencer Whitman, Michael Stroucken, William Wang, Lianghong Xu, Gregory R. Ganger. In proceedings of the 8^{th} USENIX Symposium on Network Systems Design and Implementation (NSDI'11). March 30^{th} to April 1^{st}, 2011. Boston, MA, USA.

A transparently-scalable metadata service for the Ursa Minor storage system. Shafeeq Sinnamohideen, Raja R. Sambasivan, Likun Liu, James Hendricks, Gregory R. Ganger. In proceedings of the 2010 USENIX Annual Technical Conference (USENIX ATC'10). June 23^{rd} to 25^{th}, 2010. Boston, MA, USA.

Categorizing and differencing system behaviours. Raja R. Sambasivan, Alice X. Zheng, Eno Thereska, Gregory R. Ganger. Appears in the proceedings of the 2^{nd} International Workshop on Hot Topics in Autonomic Computing (HotAC II). June 15^{th}, 2007. Jacksonville, Florida, USA.

Modeling the relative fitness of storage. Michael Mesnier, Matthew Wachs, Raja R. Sambasivan, Alice X. Zheng, Gregory R. Ganger. In proceedings of the International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS'07). June 12^{th} to 16^{th}, 2007. San Diego, CA, USA.

//Trace**: parallel trace replay with approximate causal events**. Michael Mesnier, Matthew Wachs, Raja R. Sambasivan, Julio Lopez, James Hendricks, Gregory R. Ganger. In proceedings of the 5^{th} conference on File and Storage Technologies (FAST'07). February 13^{th} to 16^{th}, 2007. San Jose, CA, USA.

Ursa Minor: versatile cluster-based storage. Michael Abd-El-Malek, William V. Courtright II, Chuck Cranor, Gregory R. Ganger, James Hendricks, Andrew J. Klosterman, Michael Mesnier, Manish Prasad, Brandon Salmon,

Raja R. Sambasivan, Shafeeq Sinnamohideen, John D. Strunk, Eno Thereska, Matthew Wachs, Jay J. Wylie. In the proceedings of the 4^{th} USENIX conference on File and Storage Technologies (FAST'05). December 13^{th} to 16^{th}, 2005. San Francisco, CA, USA.

Replication policies for layered clustering of NFS servers. Raja R. Sambasivan, Andrew J. Klosterman, Gregory R. Ganger. Appears in the proceedings of the 13^{th} Annual Meeting of the IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS'05). September 27^{th} to 29^{th}, 2005. Atlanta, Georgia, USA.

Journal Publications

Relative fitness modeling. Michael Mesnier, Matthew Wachs, Raja R. Sambasivan, Alice Zheng, Raja R. Sambasivan, Gregory R. Ganger Research Highlights, Communications of the ACM. April 2009.

Early experiences on the journey towards self-* storage. Michael Abd-El-Malek, William V. Courtright II, Chuck Cranor, Gregory R. Ganger, James Hendricks, Andrew J. Klosterman, Michael Mesnier, Manish Prasad, Brandon Salmon, Raja R. Sambasivan, Shafeeq Sinnamohideen, John D. Strunk, Eno Thereska, Matthew Wachs, Jay J. Wylie. In the Bulletin of the IEEE Computer Society Technical Committee on Data Engineering 29(3). Special issue on self-managing database systems. September 2006.

Technical Reports

Bootstrapping evolvability for inter-domain routing with D-BGP. Raja R. Sambasivan, David Tran-Lam, Aditya Akella, Peter Steenkiste. Carnegie Mellon Computer Science Technical Report CMU-CS-16-117. June 2016.

So, you want to trace your distributed system? Key design insights from years of practical experience. Raja R. Sambasivan, Rodrigo Fonseca, Ilari Shafer, Gregory R. Ganger. Carnegie Mellon University Parallel Data Laboratory Technical Report CMU-PDL-14-102. April 2014.

Visualizing request-flow comparison to aid performance diagnosis in distributed systems. Raja R. Sambasivan, Ilari Shafer, Michelle L. Mazurek. Carnegie Mellon University Parallel Data Laboratory Technical Report CMU-PDL-13-104. May 2013. Supersedes CMU-PDL-12-102.

Visualizing request-flow comparison to aid performance diagnosis in distributed systems. Raja R. Sambasivan, Ilari Shafer, Michelle L. Mazurek. Carnegie Mellon University Parallel Data Laboratory Technical Report CMU-PDL-12-102. May 2012.

Automation without predictability is a recipe for failure. Raja R. Sambasivan, Gregory R. Ganger. Carnegie Mellon University Parallel Data Laboratory Technical Report CMU-PDL-11-101. January 2011.

Diagnosing performance changes by comparing system behaviours. Raja R. Sambasivan, Alice X. Zheng, Elie Krevat, Spencer Whitman, Michael Stroucken, William Wang, Lianghong Xu, Gregory R. Ganger. Carnegie Mellon University Parallel Data Laboratory Technical Report CMU-PDL-10-107. July 2010. Supersedes CMU-PDL-10-103.

A transparently-scalable metadata service for the Ursa Minor storage system. Shafeeq Sinnamohideen, Raja R. Sambasivan, James Hendricks, Likun Liu, Gregory R. Ganger. Carnegie Mellon University Parallel Data Laboratory Technical Report CMU-PDL-10-102. March 2010.

Diagnosing performance problems by visualizing and comparing system behaviours. Raja R. Sambasivan, Alice X. Zheng, Elie Krevat, Spencer Whitman, Gregory R. Ganger. Carnegie Mellon University Parallel Data Lab Technical Report CMU-PDL-10-103. February 2010.

Eliminating cross-server operations in scalable file systems. James Hendricks, Shafeeq Sinnamohideen, Raja R. Sambasivan, Gregory R. Ganger. Carnegie Mellon University Parallel Data Lab Technical Report CMU-PDL-06-105. May 2006.

Improving small file performance in object-based storage. James Hendricks, Raja R. Sambasivan, Shafeeq Sinnamohideen, Gregory R. Ganger. Carnegie Mellon University Parallel Data Lab Technical Report CMU-PDL-06-104. May 2006.

Selected project reports, Spring 2005 Advanced OS & Distributed Systems (15-712). Garth A. Gibson and Hyang-Ah Kim, Editors. Jangwoo Kim, Eriko Nurvitadhi, Eric Chung; Alex Nizhner, Andrew Biggadike, Jad Chamcham; Srinath Sridhar, Jeffrey Stylos, Noam Zeilberger; Gregg Economou, Raja R. Sambasivan, Terrence Wong; Elaine Shi, Yong Lu, Matt Reid; Amber Palekar, Rahul Iyer. Carnegie Mellon Computer Science Technical Report CMU-CS-05-138. May 2005.

Ursa Minor: Versatile cluster-based storage. Michael Abd-El-Malek, William V. Courtright II, Chuck Cranor, Gregory R. Ganger, James Hendricks, Andrew J. Klosterman, Michael Mesnier, Manish Prasad, Brandon Salmon, Raja R. Sambasivan, Shafeeq Sinnamohideen, John D. Strunk, Eno Thereska, Matthew Wachs, Jay J. Wylie. Carnegie Mellon University Parallel Data Laboratory Technical Report CMU-PDL-05-104. April 2005.

Patents

Managing execution of database queries. Stefan Kompress, Harumi Anne Kuno, Umeshwar Dayal, Janet Wiener, Raja Sambasivan. U.S. Patent 9,910,892. March 6^{th}, 2018.

Conference Talks

Bootstrapping evolvability for inter-domain routing with D-BGP. Presented at the ACM 2017 SIGCOMM Conference (SIGCOMM'17). August, 2017.

Principled workflow-centric tracing of distributed systems. Presented at the 7^{th} ACM Symposium on Cloud Computing (SoCC'16). October 2016.

Bootstrapping evolvability for inter-domain routing. Presented at the 14^{th} ACM Workshop on Hot Topics in Networks (HotNets'15). November 2015.

Visualizing request-flow comparison to aid performance diagnosis in distributed systems. Presented at IEEE InfoVis 2013.

Automated diagnosis without predictability is a recipe for failure. Presented at the 4^{th} USENIX Workshop on Hot Topics in Cloud Computing (HotCloud'12). June 2012.

Generalizing request-flow comparison to more systems. WiP talk at 23^{rd} ACM Symposium on Operating Systems Principles (SOSP'11). October 2011.

Diagnosing performance changes by comparing request flows. Presented at the 8^{th} USENIX Symposium on Networked Systems Design and Implementation (NSDI'11). March 2011.

Spectroscope: a tool for categorizing and differencing system behaviours. Presented at the 2^{nd} International Workshop on Hot Topics in Autonomic Computing (HotACII). June 2007.

Replication policies for layered clustering of NFS servers. Presented at the 13^{th} Annual Meeting of the IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS'05). September 2005.

Invited Talks & Guest Lectures

Toward a diagnosis plane for cloud computing. Presented at LightStep (April 2018), Columbia University (March 2018), Facebook (February 2018), Brown University (February 2018).

Diagnosis and inter-domain support for an Internet of clouds. Presented at Tufts University (October 2017), Yale University (March 2017), AT&T Labs (May 2016), Intel Labs (April 2016), NYU (April 2016).

Diagnosing performance changes by comparing request flows. Presented at UCSD (April 2014), Brown

University (April 2012), NetApp RTP (September 2011), Google NYC (June 2011).

Networking at Google: B4 & Jupiter Rising. Guest lecture in BU CS 528 (March 2018, March 2017).

When the cloud fizzles: Outages and how to debug them. Guest lecture in NU CS 6620 (April 2018) and BU CS 528 (April 2017).

A case study of the AWS outage on April 21st, 2011. Guest lecture in CMU 15-719 (Fall 2015, Fall 2014).

Diagnosis via monitoring & tracing. Guest lecture in CMU 15-719 (Fall 2015, Fall 2014).

Raja R. Sambasivan

References

Academic

Orran Krieger
Professor of ECE
Director, Mass Open Cloud
Boston University
111 Cummington Mall
Boston, MA 02215
okrieg AT bu.edu

Peter Steenkiste
Professor of CS & ECE
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213
prs AT cs.cmu.edu

Gregory R. Ganger
Professor of ECE & CS
Director, Parallel Data Lab
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213
ganger AT ece.cmu.edu

Ayse K. Coskun
Associate Professor of ECE
Boston University
8 St. Mary's St.
Boston, MA 02215
acoskun AT bu.edu

Rodrigo Fonseca
Associate Professor of CS
Brown University
Box 1910
115 Waterman St
Providence, RI 02912
rfoneseca AT cs.brown.edu

Aditya Akella
Professor of CS
University of Wisconsin-Madison
1210 W. Dayton St.
Madison, WI 53706
akella AT cs.wisc.edu

Mark Crovella
Professor of CS
Boston University
140E Math & CS
111 Cummington Mall
Boston, MA 02215
crovella AT bu.edu

Industry

Peter Portante
Software Engineer
Red Hat
300 A St
Boston, MA 02110
pportant@redhat.com

Michael De Rosa
Software Engineer
Google
6425 Penn Ave. Suite 700
Pittsburgh, PA 15206
mderosa AT google.com

Kivanc Ozonat
Senior Research Scientist
HP Labs
1501 Page Mill Rd
Palo Alto, CA 94304
kivanc.ozonat AT hp.com