

Zoran Dimitrijević

Curriculum Vitae

zorand@gmail.com

+1 (805) 259-5262

<http://3opan.net/~zorand>

Technical Interests

Design and implementation of large-scale storage systems, quality of service, parallel and cluster-based computing, multimedia systems, and large-scale search engines. Security, privacy and encryption. News and citizen journalism. Open Source.

Experience Summary

2019-present	Independent consultant.
2017-2019	SAP Labs, Palo Alto, California.
2014-2017	Altiscale Inc. acquired by SAP Labs, Palo Alto, California.
2004-2013	Google Inc., Mountain View, California.
1999-2004	Teaching and Research Assistant, CS and ECE Department, UCSB, California.
Summer 2003	Instructor (Teaching Associate), CS Department, UCSB, California.
Summer 2001	Research Scientist, Sony 550 Digital Media Ventures, San Francisco, California.
Summer 2000	Research Intern, HP Labs, Palo Alto, California.
1998-1999	Research Assistant, CE&CS Department, University of Belgrade, Serbia.
Fall 1997	Research Intern, ECE Department, University of Campinas, Brazil.

Education

June 2004	Ph.D. in Computer Science , University of California, Santa Barbara (UCSB). Dissertation title: “Quality-of-Service Scheduling in Storage Systems” Advisor: Prof. Edward Y. Chang.
2003	M.S. in Computer Science , University of California, Santa Barbara. GPA: 4.00/4.00.
1999	Dipl.Ing. in Electrical Engineering , School of Electrical Engineering, University of Belgrade, Serbia, Yugoslavia. GPA: 8.95/10.00.

Experience

Principal Software Engineer, SAP Labs, Palo Alto, California, 2017–2019.

- Tech lead for Hadoop team (open-source components of SAP Cloud Platform Big Data Services including Apache Hadoop, Hive, Ranger and Tez). Service owner for internal infrastructure for Altiscale CA and x509 certificates, data encryption-at-rest, Kerberos, Hadoop build and deploy pipelines, and data transfer tools. Tech lead for extending Apache Ranger for SAP Big Data Services (<https://github.com/altiscale/ranger>). Launched Apache Ranger for SAP BDS in production. Launched SSL support for Altiscale HttpFs (and fixed a bug in open-source for swebhdfs tokens). Prototyped Altiscale Hadoop and Spark running on Kubernetes. Prototyped autoscaling Spark when running on Apache Yarn (on Kubernetes and on AWS). Implemented templates for microservices in Go with client-cert authentication and ported Altiscale encryption-at-rest key management service in Go.

Principal Software Engineer, Altiscale Inc., Palo Alto, California, 2014–2017 (acquired by SAP in September 2016).

- Tech lead for Data transfer tools. Improved open-source Hadoop distcp by multithreading file listings which were slow for large number of S3 objects (in some cases improvement from 60 minutes to 2 minutes). Implemented java-based tool for rsync between HDFS and S3 (and Glacier). Refactored Altiscale ruby-based client copytohdfs for high-volume enterprise customers. Wrote chef-server recipes for Altiscale disk partitioning and mkfs supporting dockerized environments. Designed and implemented Altiscale encryption-at-rest for HDFS based on Linux dmccrypt+LUKS block-device encryption and client-cert based secure-hash infrastructure.

- Tech lead for Hadoop team. Owned and maintained several production chef-server recipes and Jenkins pipelines including core Hadoop and internal x509 CA services. Various performance evaluation studies, especially for containerized environments and encrypted disks and SSDs. Designed and implemented open-source tool for multiplexing HttpFs transfers via multiple ssh tunnels: <https://github.com/altiscale/transfer-accelerator>. Helped certify Altiscale as SAP acquired-to-ship product, especially for data-transfer tools and data encryption-at-rest. Helped with compliance certifications.

**Software Engineer and Senior Software Engineer,
Google Inc., Mountain View, California, 2004–2013.**

- 2004–2005: Systems Infrastructure, Google File System (GFS) team. Worked with Howard Gobioff as my TL in Bill Coughran’s team. Added background checksumming in GFS Chunkserver (C++) which run on most Google machines in production until migration from GFS to Collosus. Worked on improving integration testing infrastructure for GFS (injecting faults, Python). As a 0% project launched www.google.co.yu. Launched www.google.cs. Launched www.google.rs after adding support to Google Geo-IP for my former country which kept changing names and borders. As a 20% project, designed and prototyped Google Video Streaming Service (buffering layer for videos stored on GFS) which was later used as a starting architecture for YouTube streaming infrastructure.
- 2005–2008: Systems Infrastructure, Google News Archive Search. As part of a small team with Yanghua Chu and Anurag Acharya (and later Dale Neal) designed and implemented Google News Archive Search. Implemented Google premium crawler. Implemented C++ frontend for News Archive Search (arfe). We launched News Archives in 2006 as part of news.google.com. Wrote all arfe-related borgcfg and related GFE and rpc-balancer rules. Worked with SRE teams to get required production Borg resources. Worked with PM teams to get all UI approved. Implemented most UI and Google-data-push (GDP) infrastructure for News Archives. Implemented all borgmon and related varz interfaces for arfe. Implemented backend APIs for Google one-box and news.google.com in arfe. Handled most of google.com traffic as Google News backend. Added support for pre-1970 dates for Google Search. Implemented and launched News Archive Search in several European languages. Launched Google Timeline. Implemented iGoogle module based on Archive Search. Worked with a team in India to include OCR documents from microfilms to News Archive Search.
- 2009–2011: Google News. Technical owner for videos in Google News. Added user generated content from YouTube into Google News index. Designed, implemented and lunched this feature in production. Added news-video pubsub stream for Fresh Docs (Google real-time index at that time). Re-implemented search result pages for Google News back in GWS as part of One Google project (Google-wide project to provide unified infrastructure for search-result page rendering). On-call in production for Google News and News Archives. Launched Google News in several new languages including Serbian news.google.rs.
- 2012: YouTube. Joined YouTube to start a new experimental project YouTube News.

**Graduate Research Assistant for Prof. Edward Y. Chang,
CS and ECE Department, UCSB, Santa Barbara, California, 2001–2004.**

- Proposed preemptible disk scheduling algorithms. Investigated preemptibility of disk IOs. Designed and implemented *Semi-preemptible IO* prototype. Proposed and designed preemptive RAID scheduling algorithms. Designed IO preemption and resumption criteria. Implemented a simulator for preemptible RAID systems (*PraidSim*). Research in the areas of disk profiling, modeling, and data storage. Designed and implemented QoS extensions for Linux disk scheduling (*UCSB-IO*).
- Co-designed MEMS-based disk buffer for streaming media servers. Co-designed the analytical framework, admission-control criteria, and data-placement algorithms.
- Design and implementation of SfinX video surveillance system. Proposed the hardware and software architecture for SfinX. Co-designed and co-implemented Xstream video streaming and storage system. Investigated video streaming over wireless networks.

**Research Scientist (Internship) for Dr. Sean Varah,
Sony 550 Digital Media Ventures, San Francisco, California, Summer 2001.**

- Research in the design and implementation of real-time storage systems for streaming media.

Research Intern for Dr. Dejan Milošević, HP Labs, Palo Alto, California, Summer 2000.

- Researched the susceptibility of operating systems and software to soft errors.

Publication Highlights

- [1] Zoran Dimitrijević. **Quality of service scheduling in storage systems**. Ph.D. Dissertation (ISBN 0-496-84077-0), Department of Computer Science, University of California, Santa Barbara, June 2004.

Journal Papers

- [2] R. Rangaswami, Z. Dimitrijević, E. Chang, and K. Schauer. **Building MEMS-based Storage Systems for Streaming Media**. *ACM Transactions on Storage*, Volume 3 Issue 2, June 2007.
- [3] Z. Dimitrijević, R. Rangaswami, and E. Chang. **Systems Support for Preemptive Disk Scheduling**. *IEEE Transactions on Computers*, pages 1314–1326, October 2005.
- [4] A. Messer, P. Bernadat, G. Fu, D. Chen, Z. Dimitrijević, D. Lie, D. Mannaru, A. Riska, and D. Milojičić. **Susceptibility of commodity systems and software to memory soft errors**. *IEEE Transactions on Computers*, pages 1557–1568, December 2004.
- [5] R. Rangaswami, Z. Dimitrijević, E. Chang, and S.-H. G. Chan. **Fine-grained device management in an interactive media server**. *IEEE Transactions on Multimedia*, pages 558–569, December 2003.

Conference Papers

- [6] Z. Dimitrijević, C. Sahin, C. Tinnefeld, and J. Patvarczki. **Importance of Application-level Resource Management in Multi-cloud Deployments (invited)**. *Accepted to IEEE International Conference on Cloud Engineering (IC2E)*, June 2019.
- [7] B. Liu, R. Rangaswami, and Z. Dimitrijević. **Thwarting virtual bottlenecks in multi-bitrate streaming servers**. *Proceedings of IEEE Real-Time Systems Symposium (RTSS)*, December 2005.
- [8] Z. Dimitrijević, R. Rangaswami, and E. Chang. **Design and implementation of Semi-preemptible IO**. *Proceeding of the Second Usenix File and Storage Technology (FAST)*, pages 145–158, San Francisco, California, March 2003.
- [9] R. Rangaswami, Z. Dimitrijević, E. Chang, and K. E. Schauer. **MEMS-based disk buffer for streaming media servers**. *Proceedings of the 19th IEEE International Conference on Data Engineering (ICDE)*, pages 619–630, Bangalore, India, March 2003.
- [10] D. Chen, A. Messer, P. Bernadat, G. Fu, Z. Dimitrijević, D. Lie, D. Mannaru, A. Riska, and D. Milojičić. **JVM susceptibility to memory errors**. *Proceedings of Usenix Java[tm] Virtual Machine Research and Technology Symposium*, pages 67–77, Monterey, California, April 2001.

(Full papers are available on-line at <http://3opan.net/~zorand/publications.html>.)

Teaching Experience

- Summer 2003 **Instructor (Teaching Associate)**, CS Department, Summer Sessions, UCSB.
CS170 Operating Systems (upper-division undergraduate course).
- Winter 2003 **Teaching Assistant for Prof. Tao Yang**, CS Department, UCSB.
CS240B High Performance Computing Systems and Applications (graduate-level course).
- Winter 2001 **Teaching Assistant for Prof. Peter Cappello**, CS Department, UCSB.
CS172 Software Engineering (upper-division undergraduate course).
- Fall 2000 **Teaching Assistant for Prof. Alan Konheim**, CS Department, UCSB.
CS176A Computer Networks (upper-division undergraduate course)
- Spring 2000 **Teaching Assistant for Prof. Klaus Schauer**, CS Department, UCSB.
CS290I Scalable Internet Services and Systems (graduate-level course).
- Winter 2000 **Teaching Assistant for Prof. Amr El Abbadi**, CS Department, UCSB.
CS130A Data Structures and Algorithms I (upper-division undergraduate course).
- Fall 1999 **Teaching Assistant for Prof. Anurag Acharya**, CS Department, UCSB.
CS170 Operating Systems (upper-division undergraduate course).

Selected Talks

- [1] **YouTube News: Broadcasting the World.** Industry Keynote, UCSB Graduate Student Workshop, October 2011.
- [2] **Google: organizing the world's information... and loving it.** Google recruiting tech talk, University of Colorado at Boulder, October 2005.
- [3] **Google: a computer scientist's playground.** Invited talk at Computer Science Department, FIU, Miami, April 2005.
- [4] **Two Pieces of Google Core Infrastructure: GFS and MapReduce.** Invited lecture for Prof. Rich Wolski's CS290B Advanced Operating Systems course, UCSB, March 2005.
- [5] **Google: a computer scientist's playground.** Invited talk at Innovative Computing Laboratory, University of Tennessee at Knoxville, October 2004.
- [6] **High-performance preemptible and MEMS-based IOs.** Invited talk at IBM Almaden Research Center, Almaden, California, November 2003.
- [7] **Design and implementation of Semi-preemptible IO.** Conference talk at Usenix FAST'03, San Francisco, California, April 2003.
- [8] **Multiprocessor systems.** Guest lecture for ECE154 Computer Architecture, UCSB, Winter 2003.
- [9] **I/O storage.** Guest lecture for ECE154 Computer Architecture, UCSB, Winter 2003.
- [10] **Virtual IO: Preemptible disk access.** Research seminar for Prof. Richard Muntz's group, UCLA, August 2002.

Awards, Fellowships, Honors

2018	SAP Catalyst.
2015	IEEE Senior Member.
2015-present	Member of IEEE Computer Society Industry Advisory Board.
2010	Google OC Award for Google File System.
2006–2013	Member of Google Hiring Committee for new-grad PhDs.
2003	President's Work Study Award for 2003-2004, UCSB.
2002	Graduate Student Travel Grant, Graduate Division, UCSB.
1999	Dean's Fellowship, College of Engineering, UCSB.
1999	Teaching Assistantship and Tuition Fellowship for 1999-2003, CS Department, UCSB.
1993	Second place on the classification exam in Physics, University of Belgrade.
1988-1992	First place awards on Serbian state Physics competitions in 1988, 1990, and 1992.

Refereeing

Journals: *ACM Transactions on Storage*, *IEEE Transactions on Parallel and Distributed Systems*, *The Computer Journal* (published by Oxford University Press on behalf of British Computer Society), *Software Practice and Experience* (published by John Wiley & Sons, Ltd.), *IEEE Transactions on Computers*, *IEEE Transactions on Cloud Computing*, and *Applied Computing and Informatics*.

Conferences: *IEEE Infocom 2005*, *International Conference on Database Systems for Advanced Applications (DASFAA 2004)*, *IEEE Conference on Multimedia and Expo (ICME2004, ICME2005)*, *IEEE MMM 2006*, *ACM Multimedia 2006*, *IEEE (MSST2010) Symposium on Massive Storage Systems and Technologies*, *ACM/IFIP/USENIX Middleware 2016*, *UCSB Computer Science Capstone 2019*, *IEEE International Conference on Cloud Engineering (IC2E) 2019*.