

GEORGE AMVROSIADIS

B.Sc., Ph.D. Candidate ~ ACM, IEEE, Usenix Member

Bahen Centre of Informational Technology
40 St. George Street, Room BA5214 D06
University of Toronto, Toronto, ON M5S 2EA, Canada

EMAIL: gamvrosi@gmail.com, cs.toronto.edu }

PHONE: +1 (416) 618 0426

RESEARCH OBJECTIVE

We have transitioned to an era where data plays an integral role in business continuity, and data analytics is often used to inform research and development practices. In this context, storage systems are faced with multiple tasks trying to access the same large body of data, either for data analytics, or to perform maintenance operations that will ensure data reliability, availability, performance, and security. Such tasks may range from MapReduce jobs to backup, virus scanning, and layout reorganization, among others. My current research objective is to understand how these tasks operate in the field, and enable them to collaborate by adjusting their data access patterns in a manner reactive to other tasks running in the system. This way, they can leverage data availability in system caches, and collectively reduce the I/O needed to reach their goals, which in turn reduces their runtime.

My **research interests** lie in the areas of Storage Reliability and Performance, File Systems, Operating and Distributed Systems, Systems modeling and simulation, Databases, and Networks.

EDUCATION

- SEPT. 2009 - **Doctor of Philosophy** in Computer Science, **University of Toronto**, Canada
PRESENT Area: Storage Performance and Reliability, Computer Systems & Networks
CURRENT GPA: A | Advisors: Prof. Angela DEMKE-BROWN, Prof. Ashvin GOEL
- SEPT. 2005 - **Bachelor of Science** in Computer Science, **University of Ioannina**, Greece
JULY 2009 Thesis: "Namespace Management in Federated Filesystems" | Advisor: Prof. Stergios ANASTASIADIS
My undergraduate thesis [4] focused on designing and evaluating flexible namespace management for federated file systems. We created a level of abstraction that allows administrators to transparently relocate information over a global, distributed namespace, through the use of a distributed LDAP-based directory that was used to map static NFSv4 mount points to unique keys.
GPA: 9.15 / 10.0 ("Excellent") | Graduated 2nd in class, Top 1%

PUBLICATIONS

- [1] **"Getting back up: Understanding how enterprise data backups fail"**,
George Amvrosiadis, Medha Bhadkamkar.
In Proc. of the 2016 USENIX Annual Technical Conference (ATC). (Acc.Rate: 19.0%)
- [2] **"Quartet: Harmonizing task scheduling and caching for cluster computing"**,
Francis Deslauriers, Peter McCormick, George Amvrosiadis, Ashvin Goel, Angela Demke Brown.
In Proc. of the 8th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage), 2016.
- [3] **"Opportunistic Storage Maintenance"**,
George Amvrosiadis, Angela Demke Brown, Ashvin Goel.
In Proc. of the 25th Symposium on Operating Systems Principles (SOSP), 2015. (Acc.Rate: 16.5%)
- [4] **"Identifying Trends in Enterprise Data Protection Systems"**,
George Amvrosiadis, Medha Bhadkamkar.
In Proc. of the 2015 USENIX Annual Technical Conference (ATC). (Acc.Rate: 21.2%)
- [5] **"Temperature Management in Datacenters: Cranking Up the Thermostat Without Feeling the Heat"**,
Nosayba El-Sayed, Ioan Stefanovici, George Amvrosiadis, Andy Hwang, Bianca Schroeder.
In USENIX ;login:, Volume 38, Number 1, February 2013.

- [6] **“Practical Scrubbing: Getting to the bad sector at the right time”**,
George Amvrosiadis, Alina Oprea, Bianca Schroeder.
In Proc. of the 42nd Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), Dependable Computing and Communication Symposium (DCCS), 2012. **Highest rated in conference!** (Acc.Rate: 17.3%)
- [7] **“Temperature Management in Data Centers: Why Some (Might) Like It Hot”**,
Nosayba El-Sayed, Ioan Stefanovici, George Amvrosiadis, Andy Hwang, Bianca Schroeder.
In Proc. of the Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS/Performance), 2012. **Best Paper Award!** (Acc.Rate: 15.3%)
- [8] **“Namespace Management in Federated Filesystems”**
Undergraduate Thesis, University of Ioannina, Greece, 2009. Advisor: Prof. Stergios Anastasiadis.

TEACHING EXPERIENCE

| | |
|------------------------|--|
| MAY 2011 - APRIL 2016 | <p>Help Centre Teaching Assistant, COMPUTER SCIENCE HELP CENTRE, <i>Dept. of Computer Science, St. George Campus, University of Toronto, Canada</i></p> <p>Staffed a Help Centre for Computer Science students, providing help on a drop-in basis for all undergraduate courses. Part of my responsibilities was maintaining familiarity with current assignments in key courses, and tutoring students or student groups on a variety of concepts. <i>Total semesters: 11</i></p> |
| SEPT 2009 - APRIL 2013 | <p>Teaching Assistant, OPERATING SYSTEMS, <i>Dept. of Computer Science, St. George Campus, University of Toronto, Canada</i></p> <p>Worked with instructors: Karen Reid, Prof. Angela Demke-Brown, Olga Irzak. Part of my responsibilities was to teach tutorials for students on a weekly basis, mark assignments and exams, and conduct student interviews for progress monitoring. <i>Total semesters: 5</i></p> |
| JAN 2011 - APRIL 2014 | <p>Teaching Assistant, SOFTWARE TOOLS & SYSTEMS PROGRAMMING, <i>Dept. of Computer Science, St. George Campus, University of Toronto, Canada</i></p> <p>Worked with instructors: Karen Reid, Alan J. Rosenthal, Arnold Rosenbloom. This is a class with high enrolment, and my responsibilities included: administering and teaching student labs on a weekly basis, filling in for course lectures, and marking assignments and exams. <i>Total semesters: 3</i></p> |

WORK EXPERIENCE

| | |
|-----------------------|--|
| JULY 2015 - PRESENT | <p>Intern, VERITAS LABS, <i>Veritas World Headquarters, Mountain View, California</i></p> <p>Working with Medha Bhadkamkar on designing models to auto-configure backup systems for efficiency, to meet predefined SLAs, and as a means for avoiding misconfigurations [1].</p> |
| OCT 2014 - JUNE 2015 | <p>Intern, SYMANTEC RESEARCH LABS, <i>Symantec World Headquarters, Mountain View, California</i></p> <p>Worked with Medha Bhadkamkar on identifying common practices in the configuration of backup systems [4]. Received two internal performance awards in the duration of the internship.</p> |
| SEPT 2013 - DEC 2013 | <p>Student Intern, STORAGE TECHNOLOGIES GROUP, <i>IBM Research – Zurich Lab, Switzerland</i></p> <p>Worked with Christian Cachin, Alessandro Sorniotti, Anil Kurmus on the design and implementation of a comprehensive security model for an enterprise cluster file system.</p> |
| SEPT 2008 - JULY 2009 | <p>Assistant Administrator, SYSTEMS SUPPORT GROUP, <i>Dept. of Computer Science, University of Ioannina, Greece</i></p> <p>As an assistant to the department’s administrator, my responsibilities included: introduction and integration of new technologies in production environments, routine system audits and maintenance, software modifications and hardware configuration.</p> |

STUDENT SUPERVISION

Through my advisors, Prof. Angela Demke Brown and Prof. Ashvin Goel, I have been fortunate to work closely with the following (brilliant!) undergraduate students, on a number of on-going research projects.

| | |
|---|--|
| KASRA KYANZADEH WILLIAM CUI | Realistic benchmarking with statistical file system aging, WINTER/SUMMER 2014 The goal of this project is to analyze the data layout in aged file systems. Using this information, we can then generate file systems with the same performance characteristics, for the purposes of realistic benchmarking. Kasra's work focused on extending a tool that extracts metadata for the analysis of existing aged file systems. William worked on creating a tool that uses statistical information from an existing aged file system, to generate another with the same performance characteristics. |
| ABDI DAHIR | Towards more realistic file system benchmarks, WINTER 2014 To improve the realism of existing file system benchmarks, Abdi analyzed real file system traces for a set of characteristics, such as access patterns and request inter-arrival rates. Using these measurements, he then extended the filebench benchmark to use them. Part of this work was used in the evaluation section of [3]. |
| MAX HOLDEN SUBANJAN MUKHERJEE PRANAY JAIN | Opportunistic log-structured file system segment cleaning, FALL 2014/WINTER 2015 Max, Suvanjan, and Pranay used the framework described in [3], to develop an opportunistic segment cleaner for the f2fs file system. The augmented segment cleaner was made more efficient, by taking into account data available in memory. |
| WILLIAM KINGSFORD PATRICK J. PAYNE | Enhancing file synchronization applications using Duet, FALL 2015/WINTER 2016 Project in progress. |

TALKS

| | |
|-----------|--|
| APR. 2016 | Data Loss: Worse than Zombie Apocalypse? CSC469/CSC2208: ADVANCED OPERATING SYSTEMS COURSE, GUEST LECTURE, COMPUTER SCIENCE, UNIVERSITY OF TORONTO. |
| NOV. 2015 | Opportunistic Storage Maintenance SYSTEMS DESIGN AND IMPLEMENTATION SEMINAR, PARALLEL DATA LAB, CARNEGIE MELLON UNIVERSITY, PITTSBURGH. |
| OCT. 2015 | Opportunistic Storage Maintenance ACM SOSP 2015, MONTEREY. |
| OCT. 2015 | Opportunistic Storage Maintenance VERITAS RESEARCH LABS, MOUNTAIN VIEW. |
| JUNE 2015 | Identifying Trends in Enterprise Data Protection Systems USENIX ATC 2015, SANTA CLARA. |
| JUNE 2015 | Identifying Trends in Enterprise Data Protection Systems VERITAS RESEARCH LABS, MOUNTAIN VIEW. |
| DEC 2014 | Towards Understanding Enterprise Backup Systems SYMANTEC RESEARCH LABS, MOUNTAIN VIEW. |
| NOV 2013 | Dealing with Failure: Detecting hard disk errors before it's too late IBM RESEARCH, ZURICH. |
| JULY 2012 | Practical Scrubbing: Getting to the bad sector at the right time EMC RSA LABS, BOSTON. |
| JULY 2012 | Practical Scrubbing: Getting to the bad sector at the right time IEEE DSN 2012, BOSTON. |

SELECTED AWARDS

| | |
|------------|--|
| MAR. 2016 | Doctoral Completion Award (2015-2016), UNIVERSITY OF TORONTO |
| FEB. 2016 | FAST '16 Conference Student Grant, USENIX |
| SEPT 2015 | SOSP '15 Conference Student Grant, ACM |
| JUNE 2015 | ATC '15 Conference Student Grant, USENIX |
| FEB. 2015 | Symantec Graduate Fellowship, SYMANTEC CORPORATION |
| AUG. 2014 | Doctoral Completion Award (2014-2015), UNIVERSITY OF TORONTO |
| JAN. 2014 | FAST '14 Conference Student Grant, USENIX |
| MAY 2012 | DSN 2012 Student Travel Grant, IEEE |
| APRIL 2012 | School of Graduate Studies Conference Grant, UNIVERSITY OF TORONTO |
| JAN. 2011 | FAST '11 Conference Student Grant, USENIX |
| SEPT. 2009 | Wolfond Fellowship, UNIVERSITY OF TORONTO |

SIDE PROJECTS

| | |
|------------------------------|--|
| FILEBENCH | Simplifying and extending the functionality of Filebench (with V. Tarasov), IN PROGRESS Filebench is a popular open-source benchmark used to evaluate different aspects of file systems performance. This project aims to simplify the codebase of Filebench, making it easier to maintain. In the process, important extensions are made to the functionality of the benchmark to allow it to simulate more realistic workloads. |
| INNODB BUFFER POOL EXTENSION | Extending MySQL InnoDB's Buffer Pool (with N. El-Sayed), FALL 2011 We implemented a device-agnostic staging area to function as a space extension for InnoDB's buffer pool. The staging area used lazy-checkpointing to transfer pages between the buffer pool and the disk, while guaranteeing quick recovery in the event of a failure. |
| TRASIM | Traffic Simulator (with F. Sitaras, K. Tziomakis), FALL 2009 TraSim simulates crossroad traffic light algorithms, supporting cars, pedestrians, and user interaction. The GUI and core codebase are written entirely in Java. TraSim is hosted by SourceForge.net |

PERSONAL TRIVIA

| | |
|-------------------|---|
| COMPUTER SKILLS | Preferred Languages: C/C++, R, Matlab, awk, Python; OSes: *nix, MacOS and Windows. |
| LANGUAGES | Greek (<i>native</i>), English (<i>fluent; IELTS, TOEFL, Univ. of Cambridge Lower and Proficiency, Univ. of Michigan Proficiency</i>), German (<i>Zertifikat Deutsch</i>) |
| MUSICAL EDUCATION | Music theory and Piano studies (1992 - 2004); Panorama secondary school orchestra: pianist (2000-2002), conductor (2002-2005) |
| COMMUNITY | Organization Committee Member, Computer Science Undergraduate Research Internship program (UTRECS/NSERC-USRA), <i>University of Toronto</i> , 2 TERMS (2012-2014). Research In Action Showcase Presenter, Computer Science, <i>University of Toronto</i> , 2016. Graduate Students' Union Representative for the Dept. of Computer Science, <i>University of Toronto</i> , 4 TERMS (2012-2016). |