

GEORGE AMVROSIADIS

Assistant Research Professor

Collaborative Innovation Center
5000 Forbes Ave, Office 2311
Carnegie Mellon University, Pittsburgh, PA 15213

EMAIL: gamvrosi@cmu.edu

PHONE: +1 (412) 268 7814

EDUCATION

- 2009 - 2016 **Doctor of Philosophy** in Computer Science, **University of Toronto**, Canada
AREA: Computer Systems & Networks
THESIS: “Unobtrusive Storage Maintenance” [9]
ADVISORS: Angela Demke Brown, Ashvin Goel, Bianca Schroeder
- 2005 - 2009 **Bachelor of Science** in Computer Science, **University of Ioannina**, Greece
THESIS: “Namespace Management in Federated Filesystems” [17]
ADVISOR: Stergios Anastasiadis
GPA: 9.15 / 10.0 (Graduated 2nd in class, Top 1%)

PROFESSIONAL POSITIONS

- | | | |
|-----------------------|--|----------------------------|
| JAN 2018 - PRESENT | Assistant Research Professor
Depts. of Electrical and Computer Engineering, Computer Science (by courtesy)
Carnegie Mellon University (<i>Pittsburgh, PA</i>)
→ Member of the Parallel Data Lab
→ Co-instructing graduate courses on Storage Systems, Cloud Computing | |
| SEPT 2016 - DEC 2017 | Postdoctoral Fellow
Dept. of Computer Science
Carnegie Mellon University (<i>Pittsburgh, PA</i>)
SUPERVISOR: Garth Gibson | Position length: 16 months |
| SEPT 2009 - AUG 2016 | Graduate Research Assistant
Systems and Networks Lab, Dept. of Computer Science
University of Toronto (<i>Toronto, ON</i>)
SUPERVISORS: Angela Demke Brown, Ashvin Goel, Bianca Schroeder | Position length: 7 years |
| JULY 2015 - JULY 2016 | Ph.D. Intern
Veritas Labs
Veritas World Headquarters (<i>Mountain View, CA</i>)
MANAGERS: Medha Bhadkamkar, Henry Aloysius | Position length: 12 months |
| OCT 2014 - JUNE 2015 | Ph.D. Intern
Symantec Research Labs
Symantec World Headquarters (<i>Mountain View, CA</i>)
MANAGERS: Medha Bhadkamkar, Sanjay Sawhney
→ Received two internal performance awards
→ Awarded Symantec Research Labs Graduate Fellowship (3 winners worldwide) | Position length: 9 months |
| SEPT 2013 - DEC 2013 | Student Intern
Storage Technologies Group
IBM Research – Zurich Lab (<i>Zurich, Switzerland</i>)
MANAGERS: Christian Cachin, Thomas D. Weigold | Position length: 4 months |

SEPT 2008 - JULY 2009

Assistant Administrator

Systems Support Group, Dept. of Computer Science
University of Ioannina (*Ioannina, Greece*)

ADMINISTRATOR: Spyros Melissovas

| *Position length: 11 months*

PUBLICATIONS

- [1] **“File Systems Unfit as Distributed Storage Backends: Lessons from 10 Years of Ceph Evolution”**
Abutalib Aghayev, Sage Weil, Michael Kuchnik, Mark Nelson, Gregory R. Ganger, George Amvrosiadis.
To Appear in the 27th Symposium on Operating Systems Principles (SOSP), 2019. (Acc.Rate: 13.7%)
- [2] **“Compact Filter Structures for Fast Data Partitioning”**
Qing Zheng, Charles D. Cranor, Ankush Jain, Gregory R. Ganger, Garth A. Gibson, George Amvrosiadis, Bradley W. Settlemyer, Gary A. Grider.
To Appear in the International Conference on Cluster Computing (CLUSTER), 2019.
- [3] **“Data Storage Research Vision 2025”**
George Amvrosiadis, Ali R. Butt, Vasily Tarasov, Erez Zadok, Ming Zhao, et al.
Report on NSF Visioning Workshop held May 30 – June 1, 2018.
- [4] **“The Atlas Cluster Trace Repository”**
George Amvrosiadis, Michael Kuchnik, Jun Woo Park, Chuck Cranor, Greg Ganger, Elisabeth Moore, Nathan DeBardeleben.
In USENIX ;login;, Volume 43, Number 4, Winter 2018.
- [5] **“Scaling Embedded In-Situ Indexing with DeltaFS”**
Qing Zheng, Charles D. Cranor, Gregory R. Ganger, George Amvrosiadis, Garth A. Gibson, Bradley W. Settlemyer, Gary Grider.
In Proc. of Int’l Conference for High Performance Computing, Networking, Storage & Analysis (SC18). (Acc.Rate: 19.0%)
- [6] **“On the diversity of cluster workloads and its impact on research results”**
George Amvrosiadis, Jun Woo Park, Gregory R. Ganger, Garth A. Gibson, Elisabeth Baseman, Nathan DeBardeleben.
In Proc. of the 2018 USENIX Annual Technical Conference (ATC). (Acc.Rate: 20.0%)
- [7] **“Bigger, Longer, Fewer: what do cluster jobs look like outside Google?”**
George Amvrosiadis, Jun Woo Park, Gregory R. Ganger, Garth A. Gibson, Elisabeth Baseman, Nathan DeBardeleben.
Tech report CMU-PDL-17-104.
- [8] **“Software-Defined Storage for Fast Trajectory Queries using a DeltaFS Indexed Massive Directory”**
Qing Zheng, George Amvrosiadis, Saurabh Kadekodi, Garth A. Gibson, Charles D. Cranor, Bradley W. Settlemyer, Gary Grider, Fan Guo.
In Proc. of the 2nd Joint International Workshop on Parallel Data Storage & Data Intensive Scalable Computing Systems (PDSW-DISCS), 2017.
- [9] **“Unobtrusive Storage Maintenance”**
Doctoral Dissertation, University of Toronto, Canada.
Advisors: Angela Demke Brown, Ashvin Goel.
- [10] **“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%)
- [11] **“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.
- [12] **“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%)
- [13] **“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%)

[14] "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.

[15] "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%)

[16] "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%)

[17] "Namespace Management in Federated Filesystems"

Undergraduate Dissertation, University of Ioannina, Greece.

Advisor: Stergios Anastasiadis.

TEACHING EXPERIENCE

FALL 2016	Storage Systems (18-746/15-746) Instructor (WITH GREG GANGER) <i>Electrical and Computer Engineering Dept., Carnegie Mellon University, USA</i> <ul style="list-style-type: none">• Taught lectures, generated fraction of syllabus and course material• Managed 6 graduate student TAs• Extended course projects, migrated student auto-marking infrastructure to AWS → Typical class size and level: 100 M.Sc. students	<i>Total semesters:</i> 4
FALL 2017		
FALL 2018		
FALL 2019		
SPRING 2018	Advanced Cloud Computing (18-709/15-719) Instructor (WITH GREG GANGER, MAJD SAKR) <i>Electrical and Computer Engineering Dept., Carnegie Mellon University, USA</i> → Typical class size and level: 150 M.Sc. students	<i>Total semesters:</i> 3
SPRING 2019		
SPRING 2020		
SUMMER 2011	Computer Science Help Centre (HLP101) Teaching Assistant <i>Dept. of Computer Science, University of Toronto, Canada</i> <ul style="list-style-type: none">• Tutor students and student groups on a drop-in basis (4 hours/week)• Maintain familiarity with <i>all</i> undergraduate courses• Communicate feedback to Chair of Undergraduate Studies, course instructors → Typical attendance: 15 students per TA hour	<i>Total semesters:</i> 11
FALL 2011		
SPRING 2012		
FALL 2012		
SPRING 2013		
SUMMER 2013		
SPRING 2014		
SUMMER 2014		
SPRING 2015		
FALL 2015		
SPRING 2016		
FALL 2009	Operating Systems (CSC369) Teaching Assistant <i>Dept. of Computer Science, University of Toronto, Canada</i> <ul style="list-style-type: none">• Taught weekly student tutorials• Helped construct and mark coding assignments and exams• Conducted student interviews for progress monitoring and held office hours → Typical class size and level: 100 junior undergraduate students	<i>Total semesters:</i> 5
SPRING 2010		
FALL 2010		
SPRING 2011		
SPRING 2013		

FALL 2010
SPRING 2011
SPRING 2013
SPRING 2014

Software Tools & Systems Programming (CSC209)

Teaching Assistant

Dept. of Computer Science, University of Toronto, Canada

| *Total semesters: 4*

- Managed weekly workshop sessions
 - Taught weekly student tutorials, filled in for lectures
 - Helped mark coding assignments and exams
- **Typical class size and level:** 200 sophomore undergraduate students

STUDENT SUPERVISION

• Doctoral Students

1. Michael Kuchnik → Ph.D. student, Carnegie Mellon University
2. Abutalib Aghayev → Ph.D. student, Carnegie Mellon University
3. Ankush Jain → Ph.D. student, Carnegie Mellon University
4. Qing Zheng → Ph.D. student, Carnegie Mellon University

• Master's Students

1. Chengze Fan (Spring 2019, Fall 2019) → M.Sc. student, Carnegie Mellon University
2. Minghua Deng (Spring 2019, Fall 2019) → M.Sc. student, Carnegie Mellon University
3. Mengyang Lyu (Spring 2019, Fall 2019) → M.Sc. student, Carnegie Mellon University
4. Anush Baskaran (Spring 2019) → M.Eng. student, Carnegie Mellon University
5. Zi Liang (Summer 2018) → M.Sc. student, Carnegie Mellon University
6. Saksham Jain (Spring 2018, Summer 2018) → M.Eng. student, Carnegie Mellon University
7. Venkata Parvathi Bhogaraju (Summer 2017, Fall 2017) → M.Sc. student, Carnegie Mellon University

• Undergraduate Students

1. Ishaan Gupta (Summer 2018) → B.Eng. student, Carnegie Mellon University
2. Chelsea Chen (Spring 2018) → B.Eng. student, Carnegie Mellon University
3. Ian Stewart-Binks (Fall 2016) → B.Sc. student, University of Toronto
4. Hao Dino Liu (Summer 2016) → B.Sc. student, University of Toronto
5. Patrick J. Payne (Fall 2015, Spring 2016) → Microsoft, Redmont
6. William Kingsford (Summer 2015) → B.Sc. student, University of Toronto
7. Max Holden (Fall 2014, Spring 2015) → Apple, Cupertino
8. Suvanjan Mukherjee (Fall 2014, Spring 2015) → Intel, Toronto
9. Pranay U. Jain (Fall 2014, Spring 2015) → Amazon Web Services, Vancouver
10. William Cui (Summer 2014) → B.Eng. student, University of Toronto
11. Abdi Dahir (Spring 2014, Summer 2014) → Electronic Arts, Redwood City
12. Kasra Kyanzadeh (Spring 2014, Summer 2014) → Airtable, San Francisco

FUNDING

• ***A Software Defined Storage Approach to Exascale Storage Services***

Source of Support: Department of Energy

Period: 02/01/2016 through 01/31/2019 – **Award Amount:** \$675,000

- **LANL Institute for Reliable High Performance Information Technology (IRHPIT)**
Source of Support: Los Alamos National Security, LLC
Period: 10/05/2016 through 09/30/2021 – **Award Amount:** \$1,539,078
- **Versatile Data Management Services for Future DOE Science**
Source of Support: Los Alamos National Security, LLC
Period: 02/01/2019 through 01/01/2022 – **Award Amount:** \$561,000

TALKS

1. *“Storage Research is Alive and Well”*
 CMU Parallel Data Lab Spring Visit Day 2019
 PITTSBURGH, PA – MAY 2019
2. *“Data Storage Research Vision 2025”*
 USENIX Conference on File and Storage Technologies (FAST 2019)
 BOSTON, MA – FEBRUARY 2019
3. *“How I Learned to Stop Worrying and Love the Exascale”*
 Los Alamos National Laboratory
 LOS ALAMOS, NM – AUGUST 2018
4. *“Computing with exaflops and other scary stories”*
 Veritas Corporation
 ROSEVILLE, MN – JULY 2018
5. *“How I Learned to Stop Worrying and Love the Exascale”*
 Microsoft Research
 REDMOND, WA – JUNE 2018
6. *“Scale for Masochists: Overview of High-Performance Computing Research”*
 CMU Parallel Data Lab Spring Visit Day 2018
 PITTSBURGH, PA – MAY 2018
7. *“Paving the Way for Exascale Computing”*
 Hewlett-Packard Enterprise, Nimble Storage
 SAN JOSE, CA – FEBRUARY 2018
8. *“Bigger, Longer, Fewer: what do cluster jobs look like outside Google?”*
 CMU Parallel Data Lab Retreat 2017
 BEDFORD SPRINGS, PA – OCTOBER 2017
9. *“Unobtrusive Storage Maintenance”*
 Software Systems Lab Seminar, Columbia University
 NEW YORK CITY, NY – SEPTEMBER 2017
10. *“File Systems in the Era of Exascale Computing and Tiered Storage”*
 CMU Parallel Data Lab Retreat 2016
 BEDFORD SPRINGS, PA – OCTOBER 2016
11. *“Unobtrusive Storage Maintenance”*
 CMU Parallel Data Lab Retreat 2016
 BEDFORD SPRINGS, PA – OCTOBER 2016
12. *“Getting Back Up: Understanding how enterprise backups fail”*
 USENIX Annual Technical Conference 2016
 DENVER, CO – JUNE 2016
13. *“Data Loss: Worse than Zombie Apocalypse?”*
 CSC469/CSC2208: Advanced Operating Systems course, University of Toronto (Guest Lecture)
 TORONTO, CANADA – APRIL 2016
14. *“Opportunistic Storage Maintenance”*
 Systems Design and Implementation Seminar, Carnegie Mellon University
 PITTSBURGH, PA – NOVEMBER 2015

15. *"Opportunistic Storage Maintenance"*
ACM Symposium on Operating Systems Principles (SOSP) 2015
MONTEREY, CA – OCTOBER 2015
16. *"Opportunistic Storage Maintenance"*
Veritas Research Labs
MOUNTAIN VIEW, CA – OCTOBER 2015
17. *"Identifying Trends in Enterprise Data Protection Systems"*
USENIX Annual Technical Conference 2015
SANTA CLARA, CA – JUNE 2015
18. *"Identifying Trends in Enterprise Data Protection Systems"*
Symantec Research Labs
MOUNTAIN VIEW, CA – JUNE 2015
19. *"Towards Understanding Enterprise Backup Systems"*
Symantec Research Labs
MOUNTAIN VIEW, CA – DECEMBER 2014
20. *"Dealing with Failure: Detecting hard disk errors before it's too late"*
IBM Research – Zurich Lab
ZURICH, SWITZERLAND – NOVEMBER 2013
21. *"Practical Scrubbing: Getting to the bad sector at the right time"*
EMC RSA Labs
BOSTON, MA – JULY 2012
22. *"Practical Scrubbing: Getting to the bad sector at the right time"*
IEEE/IFIP International Conference on Dependable Systems and Networks (DSN) 2012
BOSTON, MA – JULY 2012

SELECTED NEWS COVERAGE

- *"Handling trillions of supercomputer files just got simpler"*, Los Alamos National Laboratory March 14, 2019
- *"CMU's DeltaFS Team Aims To Create Smarter Ways to Organize, Store Supercomputer Data"*, Carnegie Mellon University News January 2019
- *"PDL Team Tests New File System on Trinity Supercomputer"*, CMU College of Engineering July 2018
- *"This Bomb-Simulating US Supercomputer Broke a World Record"*, ACM News July 23, 2018
- *"This Bomb-Simulating US Supercomputer Broke a World Record"*, Wired Magazine July 23, 2018
- *"Using 1 trillion files helps scientist find a needle in a haystack"*, Albuquerque Journal June 22, 2018
- *"Disks: how hot is too hot?"*, Robin Harris' Storage Mojo May 30, 2014
- *"Research Suggests There is Too Much Cooling in Data Centers"*, tom's IT PRO September 27, 2012
- *"New research: how data centres can save energy"*, University of Toronto News September 26, 2012
- *"Temperature Management in Data Centers"*, James Hamilton's Perspectives blog May 2012

SELECTED AWARDS

- Symantec Research Labs Graduate Fellowship (1 of 3 winners worldwide) February 2015
- Best paper, SIGMETRICS/Performance 2012, ACM ASSOCIATION June 2012
- Doctoral Completion Award (2x), UNIVERSITY OF TORONTO 2014-2016
- ACM SOSP, USENIX FAST (3x), USENIX ATC, IEEE DSN Student Grants 2011-2016
- School of Graduate Studies Conference Grant, UNIVERSITY OF TORONTO April 2012
- Wolfond Fellowship, UNIVERSITY OF TORONTO September 2009

SERVICE

Organizing Committee Member

- 4th Software-Defined Data Computing and Storage Workshop (SDDCS 2018)
- 1st Machine Learning for Computing Systems Workshop (MLCS 2018)

Program Committee Member

- 2020 USENIX Annual Technical Conference (ATC 2020)
- 18th USENIX Conference on File Systems and Storage Technologies (FAST '20)
- 35th International Conference on Massive Storage Systems and Technologies (MSST 2020)
- 11th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage '19)
- 34th International Conference on Massive Storage Systems and Technologies (MSST 2019)
- 39th IEEE International Conference on Distributed Computing Systems (ICDCS 2019)
- 27th Int'l ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC 2019)
- 31st Int'l Conference for High Performance Computing, Networking, Storage, and Analysis (SC 2018)
- 33rd International Conference on Massive Storage Systems and Technologies (MSST 2017)
- 37th IEEE International Conference on Distributed Computing Systems (ICDCS 2017)

Journal Reviewer: ACM Transactions on Storage (TOS), ACM Transactions on Privacy and Security (TOPS), IEEE Transactions on Computers (TC), IEEE Transactions on Cloud Computing (TCC)

Other

- National Defense Science and Engineering Graduate Fellowship (NDSEG) 2019, Panelist
- NSF Community Visioning Workshop. "Data Storage Research 2025". Invited attendee, May 2018.
- National Defense Science and Engineering Graduate Fellowship (NDSEG) 2018, Panelist

I have also acted as an **external reviewer** for several conferences: EuroSys 2016, ACM SOSP 2015, USENIX FAST 2015, USENIX HotStorage 2015, IEEE ICAC 2015, and others.

Along with Vasily Tarasov (IBM Almaden), I **maintain and develop Filebench**, a popular file system and storage benchmark that is used in numerous systems papers annually. The project's GitHub repository (<https://github.com/filebench/filebench>) is accessed 1,300+ times every week.

MISCELLANEOUS

Languages: Greek (*native*), English (*fluent*; degrees: IELTS, TOEFL, University of Cambridge Lower and Proficiency, University of Michigan Proficiency), German (Degree: Zertifikat Deutsch)

Community Involvement

- Carnegie Mellon University Committees: Graduate Studies Committee (2018 - 2019), Seminars Committee (2019 - Present), Future Computing / Connectivity Committee (2019 - Present).
- Organization Committee Member, Computer Science Undergraduate Research Internship program (UTRECS/NSERC-USRA), *University of Toronto* (2012-2014)
- Research In Action Showcase Presenter, *University of Toronto*, 2016
- Graduate Students' Union Representative for Computer Science, *University of Toronto*, 4 TERMS (2012-2016)
- Member of ACM, IEEE, USENIX