

EMINE UGUR KAYNAR

www.ugurkaynar.com

ukaynar@bu.edu

RESEARCH INTERESTS

My research interests lie broadly in the fields of **storage systems**, **cloud computing** and **data centers**. Currently, my research is focused on storage-related topics including **caching systems**, **object storage**, **erasure coding** and **distributed storage systems**.

EDUCATION

Ph.D., Computer Science

Boston University

Dissertation Title: Cooperative Caching for Object Storage

Advisor: [Prof. Orran Krieger](#)

May 2022

Boston, MA

M.Sc., Computer Science

State University of New York at Binghamton

Thesis title: Impact of Encryption on Live Virtual Machine Migration

Advisor: [Assoc. Prof. Ping Yang](#)

Aug. 2013

Binghamton, NY

B.Sc., Information Systems Engineering

Bogazici University (Dual Diploma Program)

State University of New York at Binghamton

May 2011

Istanbul, Turkey

Binghamton, NY

EXPERIENCE

- [Boston University](#)

Research Assistant | Advisor: [Prof. Orran Krieger](#)

- [D4N: Directory-Based D3N](#) Jan. 2020 - Present
Designed and implemented D4N, a data center scale cooperative cache that is integrated into today's object-based data lakes. D4N exploits object stores' immutability, provides a distributed directory for a global state and re-uses the existing data lake (Ceph) software for implementing a write tier. D4N uses the global state to implement sophisticated cache management policies and enables application-specific tuning of caching policies to support a wide range of applications.
- [D3N: Datacenter-scale Data Delivery Network](#) Sep. 2016 - Dec.2019
Designed and implemented D3N, a multi-layer cooperative caching extension to the existing object storage system, as a solution to network-constrained data centers. D3N improves big-data jobs' performance by caching data on the access side of network bottlenecks and dynamically adapts to changes in access patterns. D3N has been **upstreamed** into the Ceph code base by Red Hat and is available as an experimental feature in Ceph today. [[Github](#)]
- **Erasure Coding for Performance** Jun. 2018 - Sep. 2019
Provided a detailed performance comparison of replication and erasure coding in a modern distributed object store deployment using a simple mathematical model and empirical analysis. I investigated the impact of storage solution characteristics (e.g., disk capacity, network bandwidth) and workload's I/O profiles (read/write ratios) on the performance. For both approaches, we identified the overheads and pointed out possible improvements that may improve redundancy solutions' performance.
- [Bare-Metal Imaging \(BMI\) Service](#) Sep. 2015 - Jun. 2016
Involved in the design and implementation a generic bare-metal provisioning solution that provisions and de-provision bare-metal nodes on demand so that boundaries between clusters could be changed dynamically.

- **Red Hat** May. 2017 - Jan. 2022
Research Intern (Office of the CTO) | Mentors: [Hugh Brock](#), [Matt Benjamin](#)
 - **Implementing D4N:** Implemented the D4N cache architecture into the existing Ceph RGW code base to support hybrid cloud use case. I am still working with the Ceph RGW team on integrating the D4N cache prototype into the upstream open source repositories. [[Github](#)]
 - **Upstreaming D3N into the Ceph code base:** Worked with the Ceph RGW team on integrating the D3N cache prototype into the upstream open source repositories to make it available to the broader Ceph community. D3N has been **upstreamed** and is available as an experimental feature in Ceph today. [[Github](#)]
- *Research Intern (Ceph Performance Engineering) | Mentors: [Rick Sussman](#), [Ben England](#)*
 - **Impact of node failure and recovery on Ceph object storage:** Worked on the performance implications associated with node failure and recovery on Ceph object store, conducted performance tuning, and provided detailed performance insights. The analysis focused on two Ceph storage backends: *filestore* and *bluestore* and redundancy schemes: *erasure coding* and *replication*.
 - **Performance analysis of erasure coding and replication for Ceph object storage** Analyzed and compared the performance of *replicated* and *erasure-coded* Ceph object storage with different workloads and storage devices (e.g., NVMe SSDs, HDDs). We have tested a variety of Ceph configurations, object sizes, and read/write ratios in order to maximize the throughput of Ceph object store.
- **Mass Open Cloud (MOC)** May. 2015 - Present
Systems Researcher | Mentors: [Prof. Orran Krieger](#), [Prof. Peter Desnoyers](#), [Dr. Ata Turk](#)
 - Worked with MOC big data team to design and implement big data services on top of OpenStack. Involved in the elastic OpenStack deployment for data analytics based on demand in the MOC data center, and designed and implemented the data center scale CDN solution to speed up data analytical workloads.
 - Implemented and deployed the MOC monitoring infrastructure (*Openstack Monasca + Kafka + InfluxDB + Grafana*). Collected data is used to detect application anomalies and system breaches.
- **SUNY at Binghamton** Jul. 2012 - Aug. 2013
Research Assistant | Advisor: [Assoc. Prof. Ping Yang](#)
 - **Impacts of encryption on VM Migration:** Studied the impact of AES and 3DES encryption algorithms on two widely used live VM migration approaches, namely pre-copy and post-copy migration, and identified the best approaches for different networks and application workloads.

PUBLICATIONS

- **E. U. Kaynar**, A. Mosayyebzadeh, M. Abdi, M. Benjamin, L. Rudolph, P. Desnoyers, O. Krieger, “Universal Data Center Cache”, **(Submitted 2022)**.
- **E. U. Kaynar**, “Cooperative Caching for Object Storage”, **Ph.D. Dissertation, 2022**.
- M. H. Hajkazemi, V. Aschenbrenner, M. Abdi, **E. U. Kaynar**, A. Mosayyebzadeh, O. Krieger, P. Desnoyers, “Beating the I/O bottleneck: A Case for Log-Structured Virtual Disks”, **USENIX FAST’20**.
- M. Abdi, A. Mosayyebzadeh, M.H Hajkazemi, **E. U. Kaynar**, A. Turk, L. Rudolph, O. Krieger, P. Desnoyer, “A Community Cache with Complete Information”, **USENIX FAST’20**.
- **E. U. Kaynar**, M Abdi, M. H. Hajkazemi, A. Turk, R. R. Sambasivan, L. Rudolph, D. Cohen, P. Desnoyers, O. Krieger, “D3N: A multi-layer cache for the rest of us”, **IEEE Big Data’19**.
- A. Mohan, A. Turk, R. S. Gudimetla, S. Tikale, J. Hennesey, **E. U. Kaynar**, G. Cooperman, P. Desnoyers, O. Krieger, “M2: Malleable Metal as a Service”, **IEEE IC2E’18**.

- J. Hennessey, S. Tikale, A. Turk, **E. U. Kaynar**, C. Hill, P. Desnoyers, O. Krieger, “HIL: Designing an Exokernel for the Data Center”, **ACM SoCC’16**.
- A. Turk, R. S. Gudimetla, **E. U. Kaynar**, J. Hennessey, S. Tikale, P. Desnoyers, O. Krieger, “An Experiment on Bare-Metal BigData Provisioning”, **USENIX HotCloud’16**.
- Y. Hu, S. Panhale, T. Li, **E. U. Kaynar**, D. Chan, U. Deshpande, P. Yang, K. Gopalan, “Performance Analysis of Encryption in Securing the Live Migration of Virtual Machines”, **IEEE CLOUD’15**.
- **E. U. Kaynar**, “Impacts of Encryption on the Performance of Virtual Machine Migration”, **M.Sc. Thesis, 2013**.

TEACHING AND MENTORING EXPERIENCE

Mentoring

Project mentor to multiple projects in the Cloud Computing course jointly taught in Boston University and Northeastern University.

- CS 528/EC 528 Cloud Computing: [Accelerating Ceph Cloud Storage with D4N](#), team of 4 graduate students
- CS 6620/EC 528 Cloud Computing: [Ceph RGW S3-Select Caching](#), team of 4 graduate students
- CS 591/EC 528 Cloud Computing: [Ceph RGW Cache Prefetching](#), team of 3 graduate students
- CS 591/EC 520 Cloud Computing: [Mass Open Cloud Monitoring Platform](#), team of 3 graduate students

Teaching Assistant

- *Department of Computer Science, Boston University*
 - CS 591/ EC500 Cloud Computing Spring 2016
 - CS 108 Introduction to Application Programming Spring & Fall 2014/2015
 - CS 111 Introduction to Computer Science Summer 2015
- *Department of Computer Science, SUNY at Binghamton*
 - CS458/CS558 Introduction to Computer Security Spring 2013
 - CS 571 Programming Languages Fall 2012

SUMMARY OF QUALIFICATIONS

- **Programming:** C/C++, Bash, Python, Java, SQL
- **Storage Systems:** Ceph, HDFS, Alluxio, Redis
- **Big Data:** Hadoop, Spark, Hive, Pig, Zookeeper
- **System Profiling:** perf, blktrace, gdb, fio, sysstat, tcpdump, wireshark, strace, pbench, valgrind
- **Others:** OpenStack, AWS, Ansible, Monasca, Ceilometer, Kafka, InfluxDB, MongoDB, Grafana, S3-benchmark, CosBench

SELECTED TALKS

- **Hybrid cloud storage**, Open Cloud Workshop 2020, Boston MA, [\[Video\]](#)
- **D3N: A multi-layer cache for data centers**, DevConf.us 2019, Boston MA, [\[Video\]](#)
- **D3N: A multi-layer cache for improving big-data applications’ performance**, Mass Open Cloud Workshop 2019, Boston MA, [\[Video\]](#)
- **The Massachusetts Open Cloud: an Open Cloud eXchange**, Red Hat Summit 2017, Boston MA
- **Big Data as a Service at Mass Open Cloud**, Open Stack Summit 2017, Boston MA, [\[Video\]](#)

AWARDS & HONORS

Usenix NSDI Conference Travel Grant, 2019

MEMBERSHIPS

Usenix, IEEE, ACM, ACM's Women in Computing, Women Techmakers, System Research Group at Boston University

ACTIVITIES

Filmmaking: Member of Bogazici Cinema Club. I am interested in films and film making. I made an amateur short film called "So-called Right" in 2013.

Sports: 3 times Bronze medalist in 100m backstroke in *Women National Open Swimming Championship of Turkey* from 2000 to 2002. Many times Gold medalist in 50m/100m/200m backstroke and freestyle in *Anatolia Region Swimming Championship of Turkey* from 1997 to 2009. Captain of Samsun Gazi Swimming Club from 2005 to 2009.

REFERENCES AVAILABLE TO CONTACT

- [Prof. Orran Krieger](#)
 - Professor at Boston University Computer Science and Electrical and Computer Engineering, and Cloud Computing Initiative (CCI).
 - Prof. Krieger is my Ph.D. advisor.
- [Prof. Larry Rudolph](#)
 - Principal Research Scientist at MIT CSAIL, and Vice President and Senior Researcher Two Sigma Investments.
 - Prof. Rudolph is one of my Ph.D. co-advisors.
- [Prof. Peter Desnoyers](#)
 - Associate Professor at Northeastern University Computer Science
 - Prof. Desnoyers is the PI of the Mass Open Cloud and he is also one of my Ph.D. co-advisors.
- [Dr. Ata Turk](#)
 - Vice President of Cloud Architecture at State Street.
 - Dr. Ata Turk was a postdoc at Boston University and Mass Open Cloud, and he was one of my mentors in various research projects.
- [Matt Benjamin](#)
 - Architect and Senior Manager at Red Hat.
 - Matt Benjamin is a collaborators to the Ceph research projects and was one of my mentors for various research projects.