

Mosharaf Chowdhury

335 Ohlone Avenue, Apt 200
Albany, CA 94706

<http://www.mosharaf.com>
mosharaf@cs.berkeley.edu

INTERESTS

I am interested in networked systems, cloud computing, and networking.

During my graduate work, I developed systems and algorithms for application-aware scheduling, allocation, and load balancing of network resources in the context of large-scale data-intensive computing.

EDUCATION

University of California, Berkeley 2015 (Expected)
Ph.D. Candidate, Computer Science
Advisor: Ion Stoica

University of Waterloo 2009
Master of Mathematics, Computer Science
Advisor: Raouf Boutaba

Bangladesh University of Engineering and Technology 2007
Bachelor of Science and Engineering, Honors Computer Science

AWARDS & HONORS

- Facebook Fellowship 2012
- Best Paper Award, NSDI 2012
- Community Award, Honorable Mention, NSDI 2012
- Google GRAD CS Forum Invited Speaker 2010
- Outstanding Teaching Assistant Award, UWaterloo 2009
- Cheriton Scholarship (Type I) 2009
- UWaterloo Alumni Gold Medal, School of Computer Science Nominee 2009
- UC Berkeley Graduate Fellowship 2009
- UWaterloo Graduate Entrance Scholarship 2007

RESEARCH

University of California, Berkeley

Application-Aware Networking:

- **Orchestra** optimizes the performance of parallel communication patterns (i.e., a collection of flows or *coflows*) observed in data-intensive applications.
- **Varys** is a new architecture for dynamic network resource management that performs application-aware inter-coflow scheduling across the datacenter.
- **Aalo** enables inter-coflow scheduling in presence of failures and cluster dynamics without a priori knowledge of coflow characteristics.

Cloud Computing:

- **Sinbad** makes writes and replications in distributed file systems faster by selecting replica destinations in a network-aware fashion.
- **Spark** is a cluster computing system that provides fault-tolerant distributed memory abstractions to support iterative and interactive workloads in large clusters.

Datacenter Resource Allocation:

- **HARP** considers tradeoffs between bandwidth guarantee and fault tolerance to serve datacenter applications that need low latency and high availability.

- **FairCloud** considers the tradeoffs between bandwidth guarantee, payment proportionality, and high utilization to achieve network-wide fairness among datacenter tenants.

University of Waterloo

Network Virtualization:

- **ViNEYard** is a collection of deterministic and randomized algorithms for online virtual network embedding.
 - **TARMViNE** extends ViNEYard to support dynamic reoptimization.
 - **PolyViNE** extends ViNEYard to perform policy-based embedding across multiple administrative domains.
- **iMark** is an identity management framework to support a heterogeneous network virtualization environment.

IMPACT

Industrial impact in terms of adoption and large-scale deployments are as follows.

- Sinbad has been merged in 2014 with Facebook's HDFS codebase that runs the largest HDFS clusters in the world.
- Broadcast algorithm proposed in Orchestra is the default mechanism in Apache Spark since release 1.1.0.
- HARP is used to allocate application resources in Bing datacenters since 2012.
- Apache Spark is the most active open-source cluster computing framework today with hundreds of developers and companies using it.

Academic impact in terms of citations can be found on [Google Scholar](#).

PUBLICATIONS

Conference Papers

- M. Chowdhury, I. Stoica
[Efficient Coflow Scheduling Without Prior Knowledge](#)
SIGCOMM 2015
- M. Chowdhury, Y. Zhong, I. Stoica
[Efficient Coflow Scheduling with Varys](#)
SIGCOMM 2014
- M. Chowdhury, S. Kandula, I. Stoica
[Leveraging Endpoint Flexibility in Data-Intensive Clusters](#)
SIGCOMM 2013
- L. Popa, G. Kumar, M. Chowdhury, A. Krishnamurthy, S. Ratnasamy, I. Stoica
[FairCloud: Sharing The Network In Cloud Computing](#)
SIGCOMM 2012
- P. Bodik, I. Menache, M. Chowdhury, P. Mani, D. Maltz, I. Stoica
[Surviving Failures in Bandwidth-Constrained Datacenters](#)
SIGCOMM 2012
- M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauley, M. J. Franklin, S. Shenker, I. Stoica
[Resilient Distributed Datasets: A Fault-Tolerant Abstraction for In-Memory Cluster Computing](#)
NSDI 2012 (Best Paper Award)
- M. Chowdhury, M. Zaharia, J. Ma, M. I. Jordan, I. Stoica
[Managing Data Transfers in Computer Clusters with Orchestra](#)
SIGCOMM 2011
- N. Butt, M. Chowdhury, R. Boutaba
[Topology-Awareness and Reoptimization Mechanism for Virtual Network Embedding](#)
NETWORKING 2010

- N. M. M. K. Chowdhury, M. R. Rahman, R. Boutaba
[Virtual Network Embedding with Coordinated Node and Link Mapping](#)
INFOCOM 2009
- N. M. M. K. Chowdhury, F.-E Zaheer, R. Boutaba
[iMark: An Identity Management Framework for Network Virtualization Environment](#)
IM 2009

In Submission

- M. Chowdhury, Z. Liu, A. Ghodsi, I. Stoica
Unrestricted Work Conservation Considered Harmful in Multi-Tenant Networks
January, 2015

Journal Articles

- F. Samuel, M. Chowdhury, R. Boutaba
[PolyViNE: Policy-based Virtual Network Embedding Across Multiple Domains](#)
Journal of Internet Services and Applications, 4(6):1–23, 2013
- M. Chowdhury, M. R. Rahman, R. Boutaba
[ViNEYard: Virtual Network Embedding Algorithms with Coordinated Node and Link Mapping](#)
IEEE/ACM Transactions on Networking (ToN), 20(1):206–219, 2012
- N. M. M. K. Chowdhury, R. Boutaba
[A Survey of Network Virtualization](#)
Computer Networks, 54(5):862–876, 2010 (**ComNet Top 25 Most Cited Article**)
- D. T. Ahmed, N. M. M. K. Chowdhury, M. M. Akbar
[Admission Control Algorithm for Multimedia Server: A Hybrid Approach](#)
International Journal of Computers and Applications, 29(4):414–419, 2007

Workshop Papers

- M. Chowdhury, I. Stoica
[Coflow: A Networking Abstraction for Cluster Applications](#)
SIGCOMM HotNets-XI 2012
- G. Kumar, M. Chowdhury, S. Ratnasamy, I. Stoica
[A Case for Performance-Centric Network Allocation](#)
USENIX HotCloud 2012
- M. Chowdhury, F. Samuel, R. Boutaba
[PolyViNE: Policy-based Virtual Network Embedding Across Multiple Domains](#)
SIGCOMM VISA 2010
- M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, I. Stoica
[Spark: Cluster Computing with Working Sets](#)
USENIX HotCloud 2010
- N. M. M. K. Chowdhury, M. M. Akbar, M. Kaykobad
[DiskTrie: An Efficient Data Structure Using Flash Memory for Mobile Devices](#)
WALCOM 2007

Invited Papers

- M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauley, M. J. Franklin, S. Shenker, I. Stoica
[Fast and Interactive Analytics over Hadoop Data with Spark](#)
USENIX ;login:, 37(4):45–51, 2012

- N. M. M. K. Chowdhury, R. Boutaba
[Network Virtualization: State of the Art and Research Challenges](#)
IEEE Communications Magazine, 47(7):20–26, 2009

Non-Refereed Technical Reports

- M. Chowdhury, R. Agarwal, V. Sekar, I. Stoica
[A Longitudinal and Cross-Dataset Study of Internet Latency and Path Stability](#)
UC Berkeley Technical Report UCB/EECS-2014-172, October 2014
- S. Agarwal, M. Chowdhury, D. Joseph, I. Stoica
[Lattice: A Scalable Layer-Agnostic Packet Classification Framework](#)
UC Berkeley Technical Report UCB/EECS-2011-96, August 2011

Theses

- N. M. M. K. Chowdhury
[Identity Management and Resource Allocation in the Network Virtualization Environment](#)
Master's thesis
School Nominee for the UWaterloo Alumni Gold Medal
- N. M. M. K. Chowdhury
[A Study of the Hybrid Admission Control Algorithm for Multimedia Server](#)
Bachelor's thesis

Patent

- P. Bodik, I. Menache, P. Winkler, G. Foxman, N. M. M. K. Chowdhury
[Management of Datacenters for Fault Tolerance and Bandwidth](#)
Application # US 13/489,207, Microsoft Corporation, 2013

TEACHING EXPERIENCE

I was a teaching assistant in the following occasions.

Operating Systems (UC Berkeley CS162) *Fall 2012 & Spring 2012*

- Taught 30–35 students in weekly discussion sections. Designed problem sets and exam questions. Supervised group projects of 4–5 students in each group.
- Designed two new projects where students implemented centralized and distributed key-value stores.

Distributed Computer Systems (UWaterloo CS436) *Winter 2009*

- Designed two new assignments to introduce non-CS major students to the fundamentals of communication over the network.
- Received the **Best TA Award**.

Computer Networks (UWaterloo CS456) *Fall 2008 & Fall 2007*

- Redesigned three assignments on networking fundamentals, switching, and routing protocols for CS major students.
- Coordinated four TAs' schedules, workload, and responsibilities as the Head TA.

Computer Architecture (UWaterloo CS450) *Winter 2008*

- Designed two new projects where students implemented single-cycle and pipelined multi-cycle MIPS processors in Verilog.

Principles of Computer Science (UWaterloo CS134) *Summer 2007*

- INDUSTRY EXPERIENCE**
- Facebook** *Fall 2013–Spring 2014*
Implemented, evaluated, and merged Sinbad with Facebook’s HDFS codebase with support for newer 10-Gbps NICs.
- Microsoft Bing** *Spring 2012–Summer 2012*
Analyzed Microsoft’s Cosmos distributed file system to pinpoint its network dependencies.
- Research Intern, Microsoft Research Redmond** *Summer 2011*
Developed resource allocation schemes for Bing datacenters with capacity, latency, and availability constraints.
- Research Intern, Microsoft Research Cambridge** *Summer 2010*
Developed models for data-intensive workloads to perform capacity provisioning and planning of new datacenters.
- PRESENTATIONS** Slides from all talks are available at <http://www.mosharaf.com/talks>.
- Efficient Coflow Scheduling with Varys**
- Conference talk at SIGCOMM, Chicago, IL, August 2014
 - Facebook, Menlo Park, CA, August 2014
 - VMware, Palo Alto, CA, July 2014
 - AMPLab Summer Retreat, Santa Cruz, CA, May 2014
- Leveraging Endpoint Flexibility in Data-Intensive Clusters**
- Future Internet Technology Center, Tsinghua University, Beijing, China, August 2013
 - Conference talk at SIGCOMM, Hong Kong, China, August 2013
 - Cloudera, San Francisco, CA, July 2013
 - Facebook, Menlo Park, CA, July 2013
 - AMPLab Summer Retreat, Santa Cruz, CA, May 2013
- Coflow: A Networking Abstraction for Cluster Applications**
- Qualcomm Research, Santa Clara, CA, March 2013
 - Workshop talk at HotNets-XI, Redmond, WA, October 2012
- FairCloud: Sharing Cloud Networks**
- AMPLab Winter Retreat, Lake Tahoe, CA, January 2012
- Managing Data Transfers in Computer Clusters with Orchestra**
- Yahoo! Research, Santa Clara, CA, September 2011
 - IEEE Communications Seminar, University of Waterloo, ON, Canada, August 2011
 - Conference talk at SIGCOMM, Toronto, Canada, August 2011
 - Bing Network Management Group, Bellevue, WA, August 2011
 - Microsoft Research, Redmond, WA, August 2011
 - AMPLab Summer Retreat, Santa Cruz, CA, May 2011
 - Computing Sciences Seminar, Lawrence Berkeley National Laboratory (LBNL), Berkeley, CA, May 2011
- PolyViNE: Policy-based Virtual Network Embedding Across Multiple Domains**
- Workshop talk at SIGCOMM VISA, New Delhi, India, September 2010

A Spark in the Cloud: Iterative and Interactive Cluster Computing

- Computer Laboratory Systems Research Group Seminar, University of Cambridge, Cambridge, United Kingdom, July 2010

CLayer: Packet Classification with Explicit Coordination

- Google GRAD CS Forum, Mountainview, CA, January 2010

An Overview of Network Virtualization

- CS854 Guest Lecture, University of Waterloo, Waterloo, ON, Canada, January 2009

Virtual Network Embedding with Coordinated Node and Link Mapping

- Conference talk at INFOCOM, Rio de Janeiro, Brazil, April 2009

DiskTrie: An Efficient Data Structure Using Flash Memory for Mobile Devices

- Workshop talk at WALCOM, Dhaka, Bangladesh, February 2007

SERVICE

(External) Reviewer for

| | |
|---|------|
| IEEE/ACM ToN, Computer Networks, HotCloud 2015 | 2015 |
| SIGCOMM 2014, HotNets-XIII, IEEE/ACM ToN, IEEE Network | 2014 |
| NSDI 2014, SIGCOMM 2013, IEEE/ACM ToN, SIGCOMM CCR, IEEE TPDS, IEEE Network, Computer Communications, International Journal of Communication Systems | 2013 |
| OSDI 2012, SOCC 2012, Euro-Par 2012, IEEE/ACM ToN, Computer Networks, IEEE TPDS, Journal of Supercomputing, Journal of Network and System Management, Journal of Grid Computing, IEEE Communications Letters, IEEE Internet Computing | 2012 |
| Networking 2011, IEEE/ACM ToN, IEEE Communications, Computer Networks, Journal of Network and System Management, International Journal of Communication Systems | 2011 |
| INFOCOM 2010, IWCMC 2010, IEEE Communications | 2010 |
| IM 2009, ITC 21, IEEE JSAC | 2009 |
| EVGM 2008 | 2008 |