

Sameh Elnikety

Mailing address:
1 Microsoft Way
Redmond WA 98052
USA

Email samehe@microsoft.com
Web <http://research.microsoft.com/~samehe/>
Phone +1 425 538 2234
Date December 16, 2021

Research Interests

General area: Experimental software systems including distributed computing, operating systems, and databases.

Topics: Latency-sensitive server systems, virtualization, performance isolation, fault-tolerance, and availability in data management systems (esp. replication and consistency), and performance modeling.

Education

EPFL, Lausanne, Switzerland

Ph.D. in Computer Science – December 2007

Thesis: “Scalable transparent database replication”

Advisor: Dr. Willy Zwaenepoel.

Rice University, Houston, TX, USA

M.S. in Computer Science – May 2003

Thesis: “Request scheduling and admission control for dynamic content web servers”

Advisor: Dr. Willy Zwaenepoel.

Alexandria University, Alexandria, Egypt

M.S. in Computer Science and Automatic Control – July 1999

Thesis: “A real-time disk scheduling algorithm for multimedia storage servers”

Advisors: Dr. Walid Aref and Dr. Mohamd Abugabal

B.S. with Honors in Computer Science and Automatic Control – June 1996.

Research Experience

Microsoft Research, Redmond, WA, USA

9/2007 – Present

Researcher.

Redmond Research Lab (7/2009 – Present).

Cambridge Research Lab (9/2007-7/2009).

My main research area is distributed computer systems, and my work spans two main themes:

- resource management for large interactive services modern on cloud platforms, and
- providing system support for managing and querying graph and geospatial data to support social networks and social media on distributed server systems.

Recent research includes the following:

- Reducing tail latency and parallelization for interactive cloud services
 - [SOSP 2021, EuroSys 2021, SoCC 2021a, SoCC 2021b, OSDI 2020, USENIX ATC 2018, Middleware 2017, MICRO 2017, SPAA 2017, PPOPP 2016, ASPLOS 2016, WDM 2015, ASPLOS 2015, ICAC 2015, SIGIR 2015, SIRIR 2014, Eurosys 2013, ICAC 2013a, ICAC 2013b, CAC 2013, SOCC 2012, ICAC 2012, ICDCS 2011, Allerton 2011, Eurosys 2009a, OSDI 2008].

- Managing and querying distributed data
 - [VLDB 2019, COMPSAC 2019, WWW 2018, SIGIR 2017, SOCC 2017, VLDB 2016, SIG-Spatial 2016, ICDE 2014, VLDB 2014, DNS 2014, SOCC 2013, SRDS 2013, SPIRE 2013, CIKM 2012, ICDE 2012, ICDE 2011, VLDB 2011, ICDE 2010].

EPFL, Lausanne, Switzerland

1/2003 – 9/2007

Research Assistant.

In my Ph.D. thesis I demonstrate that, despite earlier projections to the contrary, replication on a LAN cluster can transparently scale the performance of databases. To achieve scalable performance, I developed the necessary advances in concurrency control, analytical modeling, implementation techniques, and load balancing.

- Developed generalized snapshot isolation [SRDS 2005].
- Constructed and validated a queuing model to predict performance of primary copy and write-anywhere replicated databases [EuroSys 2009b].
- Introduced a novel design principle: uniting durability and ordering. Built Tashkent, a database replication middleware that unites durability and ordering with attendant performance benefits [EuroSys'06].
- Designed memory-aware load balancing (MALB) and update filtering. Implemented them in Tashkent+, achieving superlinear speedup [EuroSys 2007].

IBM T. J. Watson Research Center, Hawthorne, NY, USA

5/2002 – 9/2002

Summer Intern with Dr. Erich Nahum and Dr. John Tracey.

A multi-tiered dynamic web server is exposed to load variations and occasional overloads. Traditional methods for preventing overload require extensive changes to the complex server software. I proposed a transparent method in which a controller intercepts incoming requests for server throttling based on on-line estimates of service times.

- Designed and implemented transparent admission control and request scheduling for e-commerce Web sites [WWW 2004].

Systems Research Center, COMPAQ / HP, Palo Alto, CA, USA

6/2000 – 9/2000

Summer Intern with Dr. Mark Lillibridge and Dr. Mike Burrows.

I proposed a novel peer-to-peer backup technique that allows machines to back up their data over wide area networks. This cooperative technique provides highly-reliable backup at a very low cost: one to two orders of magnitude cheaper than commercial Internet backup services. Since cooperation is required, the system is vulnerable to several attacks such as free-riding and disruption. I analyzed the security threats and introduced novel methods to defend against the threats.

- Designed and implemented a peer-to-peer cooperative backup system [USENIX 2003].

Rice University, Houston, TX, USA

9/1999 – 12/2002

Research Assistant.

Dynamic content web servers are complex because they consist of several tiers (front-end web server, application server, back-end database server). To scale dynamic servers, I characterized the workload and determined server bottlenecks.

- Developed benchmarks for dynamic content web servers and analyzed server bottlenecks [WWC 2002].
- Characterized middleware architectures for generating dynamic content [Middleware 2003].
- Developed and implemented a server throttling controller through performance reflection [HotOS 2003].

Informatics Research Institute, Alexandria, Egypt

9/1996 – 9/1999

Research Assistant and Programmer.

- Designed and implemented a real-time disk scheduling algorithm for multimedia storage servers.
- Implemented bandwidth smoothing and virtual VCR functions for video-on-demand servers.
- Designed and implemented a search engine for Arabic documents.

Publications

1- Refereed Research Papers (Mainly Computer Science Conferences)

Faster and Cheaper Serverless Computing on Harvested Resources.

SOSP 2021: ACM Symposium on Operating Systems Principles.

April 2021.

SmartHarvest: Harvesting Idle CPUs Safely and Efficiently in the Cloud.

EuroSys 2021: European Conference on Computer Systems.

April 2021.

CacheSlicer: Provisioning Differentiated LLC Allocations to VMs in Public Clouds.

SoCC 2021: ACM Symposium on Cloud Computing.

November 2021.

Parslo: Gradient Descent-based Framework for Partial SLO Allocation in Auto-scaled Microservices.

SoCC 2021: ACM Symposium on Cloud Computing.

November 2021.

Providing SLOs for Resource-Harvesting VMs in Cloud Platforms.

OSDI 2020: USENIX Symposium on Operating Systems Design and Implementation.

November 2020.

Local Trend Discovery on Real-time Microblogs with Uncertain Locations in Tight Memory Environments.

GeoInfomatica 2020.

Model-Switching: Dealing with Fluctuating Workloads in Machine-Learning-as-a-Service Systems.

HotCloud 2020: USENIX Workshop on Hot Topics in Cloud Computing.

July 2020.

List Intersection for Web Search: Algorithms, Cost Models, and Optimizations.

Sunghwan Kim, Taesung Lee, Seung-won Hwang, and Sameh Elnikety.

VLDB 2019: 42th International Conference on Very Large Data Bases.

Los Angeles, CA, USA, 26-30 August 2019.

Parallel Discovery of Trajectory Companions from Heterogeneous Streaming Data.

Yongyi Xian, Chuanfei Xu, Sameh Elnikety, Yan Liu

COMPSAC 2019: IEEE 43rd Annual Computer Software and Applications Conference

Milwaukee, WI, USA, 15-19 July 2019.

Best Paper Award!

Performance Isolation for Large Commercial Latency-Sensitive Services.

USENIX ATC 2018: USENIX Annual Technical Conference.

Boston, MA, 11-13 July 2018.

Better Caching in Search Advertising Systems with Rapid Refresh Predictions.

Conglong Li, David Andersen, Qiang Fu, Sameh Elnikety, Yuxiong He.

WWW 2018: 26th International World Wide Web Conference.

Lyon, France, 23-27 April 2018.

Swayam: Distributed Autoscaling to Meet SLAs of Machine Learning Inference Services with Resource Efficiency.

Arpan Gujarati, Sameh Elnikety, Yuxiong He, Kathryn S. McKinley, Björn B. Brandenburg.
Middleware 2017: Annual ACM/IFIP/USENIX Middleware conference.
Las Vegas, NV, USA, 11-15 December 2017.
Best Student Paper Award!

BitFunnel: Revisiting Signatures for Search.
Bob Goodwin, Michael Hopcroft, Dan Luu, Alex Clemmer, Mihaela Curmei, Sameh Elnikety, Yuxiong He.
SIGIR 2017: 40th International ACM SIGIR Conference on Research and Development in Information Retrieval.
Tokyo, Japan, 7-11 August 2017.
Best Paper Award!

Optimal Reissue Policies for Reducing Tail-Latency.
Tim Kaler, Yuxiong He, Sameh Elnikety.
SPAA 2017: 29th ACM Symposium on Parallelism in Algorithms and Architectures.
Washington D.C., USA, 24-26 July 2017.

Exploiting Heterogeneity for Tail Latency and Energy Efficiency.
Md E. Haque, Yuxiong He, Sameh Elnikety, Thu D. Nguyen, Ricardo Bianchini, Kathryn S. McKinley.
MICRO 2017: 50th Annual IEEE/ACM International Symposium on Microarchitecture.
Boston, MA, USA, Oct 14 – 18, 2017.

Workload Analysis and Caching Strategies for Search Advertising Systems.
Conglong Li, David G. Andersen, Qiang Fu, Sameh Elnikety, Yuxiong He.
SOCC 2017: ACM Symposium on Cloud Computing.
Santa Clara, CA, USA, Sept 25 – 27, 2017.

Work Stealing for Interactive Services to Meet Target Latency.
Jing Li, Kunal Agrawal, Sameh Elnikety, Yuxiong He, I-Ting Angelina Lee, Chenyang Lu, Kathryn S. McKinley.
PPOPP 2016: 21st ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming.
Barcelona, Spain, 12-16 March 2016.

TPC: Target-Driven Parallelism Combining Prediction and Correction to Reduce Tail Latency in Interactive Services.
Myeongjae Jeon, Yuxiong He, Hwanju Kim, Sameh Elnikety, Scott Rixner, Alan L. Cox.
ASPLOS 2016: International Conference on Architectural Support for Programming Languages and Operating Systems
Atlanta, Georgia, USA, 2-6 April 2016.

Processing and Optimizing Main Memory Spatial-Keyword Queries.
Taesung Lee, Jin-Woo Park, Sanghoon Lee, Seung-won Hwang, Sameh Elnikety, Yuxiong He.
VLDB 2016: 40th International Conference on Very Large Data Bases.
New Delhi, India, 5-9 September 2016.

GeoTrend: Spatial Trending Queries on Real-time Microblogs.
Amr Magdy, Ahmed M. Aly, Mohamed F. Mokbel, Sameh Elnikety, Yuxiong He, Suman Nath, Walid G. Aref.
ACM SIGSPATIAL 2016: 24th ACM International Conference on Advances in Geographic Information Systems
San Francisco, California, USA, 31 October – 3 November 2016.

DDS Prediction: Reducing Extreme Tail Latency in Web Search.
Saehoon Kim, Yuxiong He, Seung-won Hwang, Sameh Elnikety, Seungjin Choi.
WSDM 2015: 8th ACM International Conference on Web Search and Data Mining.
Shanghai China, 2-6 February 2015.
Best Paper Runner-up Award!

Few-to-Many: Incremental Parallelism for Reducing Tail Latency in Interactive Services.
E. Haque, Yong hun Eom, Yuxiong He, Sameh Elnikety, Ricardo Bianchini, Kathryn S McKinley.
ASPLOS 2015: International Conference on Architectural Support for Programming Languages and Operating Systems.
Istanbul, Turkey, 14-19 March 2015.

Uhora: Measuring and Managing Answer Quality for Online Data-Intensive Services.
Jaimie Kelley, Christopher Stewart, Nathaniel Morris, Devesh Tiwari, Yuxiong He, Sameh Elnikety.
ICAC 2015: 12th International Conference on Autonomic Computing.
Grenoble, France, 7-10 July 2015.

Optimal Aggregation Policy for Reducing Tail Latency of Web Search.
Jeong-Min Yun, Yuxiong He, Sameh Elnikety, Shaolei Ren.
SIGIR 2015: 38th International ACM SIGIR Conference on Research and Development in Information Retrieval.
Santiago, Chile, 9-13 August 2015.

Mercury: A Memory-Constrained Spatio-temporal Real-time Search on Microblogs.
Amr Magdy, Mohamed F. Mokbel, Sameh Elnikety, Suman Nath, Yuxiong He.
ICDE 2014: 30th IEEE International Conference on Data Engineering.
Chicago, IL, USA, 31 March - 4 April, 2014.

Predictive Parallelization: Taming Tail Latencies in Web Search.
Myeongjae Jeon, Saehoon Kim, Seung-won Hwang, Yuxiong He, Sameh Elnikety, Alan L. Cox, Scott Rixner.
SIGIR 2014: 37th International ACM SIGIR Conference on Research and Development in Information Retrieval.
Broadbeach, Queensland, Australia, 6-11 July 2014.

Clock-RSM: Low-Latency Inter-Datacenter State Machine Replication Using Loosely Synchronized Physical Clocks.
Jiaqing Du, Daniele Sciascia, Sameh Elnikety, Willy Zwaenepoel, Fernando Pedone.
DSN 2014: 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks.
Atlanta, Georgia, USA, 23-26 June 2014.

Horton: A Distributed System for Processing Declarative Reachability Queries over Partitioned Graphs.
Mohamed Sarwat, Sameh Elnikety, Yuxiong He, Mohamed F. Mokbel.
VLDB 2014: 40th International Conference on Very Large Data Bases.
Hangzhou, China. 1-5 September 2014.

Orbe: Scalable Causal Consistency for Partitioned and Replicated Data Stores.
Jiaqing Du, Sameh Elnikety, Amitabha Roy, Willy Zwaenepoel.
SOCC 2013: ACM Symposium on Cloud Computing.
Santa Clara, CA, USA, 1-3 October 2013.

Adaptive Parallelization of Web Search.
Myeongjae Jeon, Yuxiong He, Sameh Elnikety, Alan Cox, Scott Rixner.
Eurosys 2013: European Conference on Computer Systems.
Prague, Czech Republic, 15-17 April 2013.

Clock-SI: Snapshot Isolation for Partitioned Data Stores Using Loosely Synchronized Clocks.
Jiaqing Du, Sameh Elnikety, Willy Zwaenepoel.
SRDS 2013: 32nd International Symposium on Reliable Distributed Systems.
Braga, Portugal, 30 September- 3 October 2013.

Exploiting Processor Heterogeneity in Interactive Services.
Shaolei Ren, Yuxiong He, Sameh Elnikety, Kathryn McKinley.

ICAC 2013: 10th International Conference on Autonomic Computing.
San Jose, CA, USA, 26-28 June 2013.

Credit-Based Scheduling for Large Scale Data Processing.
Mingyuan Xia, Nan Zhu, Yuxiong He, Sameh Elnikety, Xue Liu.
ICAC 2013: 10th International Conference on Autonomic Computing.
San Jose, CA, USA, 26-28 June 2013.

A Novel Approach to Graph Isomorphism Based on Parameterized Matching.
Juan Mendivelso, Sunghwan Kim, Sameh Elnikety, Yuxiong He, Seung-Won Hwang and Yoan Pinzon
SPIRE 2013: 20th String Processing and Information Retrieval Symposium.
Jerusalem, Israel, 7-10 October 2013.

QACO: Exploiting Partial Execution in Web Servers.
Jinhan Kim, Sameh Elnikety, Yuxiong He, Seung-Won Hwang and Shaolei Ren.
CAC 2013: ACM Cloud and Autonomic Computing Conference.
Miami, FL, USA, 5-9 August 2013.

Scheduling Time-Bounded Computations with Adaptive Execution.
Yuxiong He, Sameh Elnikety, James Larus, Chenyu Yan.
SOCC 2012: ACM Symposium on Cloud Computing.
San Jose, CA, USA, 14-17 October 2012.

G-SPARQL A Hybrid Engine for Querying Large Attributed Graphs.
Sherif Sakr, Sameh Elnikety, Yuxiong He.
CIKM 2012: 21st ACM International Conference on Information and Knowledge Management.
Maui, Hawaii, USA, 29 October – 2 November 2012.

Budget-based Control for Interactive Services with Adaptive Execution.
Yuxiong He, Zihao Ye, Qiang Fu, Sameh Elnikety.
ICAC 2012: 9th ACM International Conference on Autonomic Computing.
San Jose, CA, USA, 17-21 September 2012.

One-Copy Serializability with Snapshot Isolation under the Hood.
Mihaela A. Bornea, Orion Hodson, Sameh Elnikety, Alan Fekete.
ICDE 2011: 27th IEEE International Conference on Data Engineering.
Hannover, Germany, 11-16 April 2011.

Scheduling for Data Center Interactive Services.
Yuxiong He, Sameh Elnikety.
Allerton 2011: 49th Annual Allerton Conference on Communication, Control, and Computing.
Monticello, IL, USA, 28-30 September 2011.

Tians Scheduling: Using Partial Processing in Best-Effort Applications.
Yuxiong He, Sameh Elnikety, Hongyang Sun.
ICDCS 2011: 31st International Conference on Distributed Computing Systems.
Minneapolis, MN, USA, 20-24 June 2011.

Strongly Consistent Replication for a Bargain.
Konstantinos Krikellas, Sameh Elnikety, Zografoula Vagena, and Orion Hodson.
ICDE 2010: 26th IEEE International Conference on Data Engineering.
Long Beach, CA, USA, 1-6 March 2010.

Predicting Replicated Database Performance from Standalone Database Profiling.

Sameh Elnikety, Steven Dropsho, Emmanuel Cecchet, and Willy Zwaenepoel.

EuroSys 2009: European Conference on Computer Systems.

Nuremberg, Germany, 31 March – 3 April 2009.

Migrating Enterprise Storage to SSDs: Analysis of Tradeoffs.

Dushyanth Narayanan, Eno Thereska, Austin Donnelly, Sameh Elnikety, and Antony Rowstron.

EuroSys 2009: European Conference on Computer Systems.

Nuremberg, Germany, 31 March – 3 April 2009.

Everest: Scaling down Peak Loads through I/O Off-loading.

Dushyanth Narayanan, Austin Donnelly, Eno Thereska, Sameh Elnikety, and Antony Rowstron.

OSDI 2008: USENIX Symposium on Operating Systems Design and Implementation.

San Diego, CA, USA, 8-10 December 2008.

Tashkent+: Memory-Aware Load Balancing and Update Filtering in Replicated Databases.

Sameh Elnikety, Steven Dropsho, and Willy Zwaenepoel.

EuroSys 2007: European Conference on Computer Systems.

Best Paper Award!

Lisbon, Portugal, 21-23 March 2007.

Tashkent: Uniting Durability with Transaction Ordering for High-Performance Scalable Database Replication.

Sameh Elnikety, Steven Dropsho, and Fernando Pedone.

EuroSys 2006: European Conference on Computer Systems.

Leuven, Belgium. 18-21 April 2006.

Database Replication Using Generalized Snapshot Isolation.

Sameh Elnikety, Fernando Pedone, and Willy Zwaenepoel.

SRDS 2005: 24th IEEE Symposium on Reliable Distributed Systems.

Orlando, FL, USA, 26-28 October 2005.

A Method for Transparent Admission Control and Request Scheduling in E-Commerce Web Sites.

Sameh Elnikety, Erich Nahum, John Tracey and Willy Zwaenepoel.

WWW 2004: Thirteenth International World Wide Web Conference.

New York City, NY, USA, 17-22 May 2004.

Performance Comparison of Middleware Architectures for Generating Dynamic Web Content.

Emmanuel Cecchet, Anupam Chanda, Sameh Elnikety, Julie Marguerite and Willy Zwaenepoel.

Middleware 2003: ACM/IFIP/USENIX International Middleware Conference.

Rio de Janeiro, Brazil, 16-20 June 2003.

A Cooperative Internet Backup Scheme.

Mark Lillibridge, Sameh Elnikety, Andrew Birrell, Mike Burrows and Michael Isard.

USENIX ATC 2003: USENIX Annual Technical Conference.

San Antonio, TX, USA, 9-14 June 2003.

2- Tutorials, Technical Reports, Workshop Papers, Journals and Extended Abstracts

Local trend discovery on real-time microblogs with uncertain locations in tight memory environments

A Almaslukh, A Magdy, AM Aly, MF Mokbel, S Elnikety, Y He, S Nath, ...

GeoInformatica, 1-37 (2019).

Obtaining and Managing Answer Quality for Online Data-Intensive Services.

Jaimie Kelley, Christopher Stewart, Nathaniel Morris, Devesh Tiwari, Yuxiong He, Sameh Elnikety.
ACM Transactions on Modeling and Performance Evaluation of Computing Systems.
TOMPECS 2(2): 11:1-11:31 (2017).

Venus: Scalable Real-Time Spatial Queries on Microblogs with Adaptive Load Shedding.

Amr Magdy, Mohamed F. Mokbel, Sameh Elnikety, Suman Nath, Yuxiong He.
IEEE Trans. Knowl. Data Eng. 28(2): 356-370 (2016).

Prediction and Predictability for Search Query Acceleration.

Seung-won Hwang, Saehoon Kim, Yuxiong He, Sameh Elnikety, Seungjin Choi.
ACM Transactions on the Web. TWEB 10(3): 19:1-19:28 (2016).

Hybrid query execution engine for large attributed graphs.

Sherif Sakr, Sameh Elnikety, Yuxiong He.
Information Systems 41: 45-73 (2014).

Mars: Real-time Spatio-temporal Queries on Microblogs (Demo Track).

Amr Magdy, Ahmed M. Aly, Mohamed F. Mokbel, Sameh Elnikety, Yuxiong He, Suman Nath, Walid G. Aref.
ICDE 2014: 30th IEEE International Conference on Data Engineering.
Chicago, IL, USA, 31 March - 4 April, 2014.

Cache Provisioning for Interactive NLP Services.

Jaimie Kelley, Christopher Stewart, Yuxiong He and Sameh Elnikety.
LADIS 2013: 7th Workshop on Large-Scale Distributed Systems and Middleware.
Farmington, PA, USA, 2-3 November 2013.

Demonstration of Horton: Online Query Execution Engine for Large Distributed Graphs.

Mohamed Sarwat, Sameh Elnikety, Yuxiong He, Gabriel Kliot.
ICDE 2012: 28th IEEE International Conference on Data Engineering.
Washington, D.C., USA, 1-5 April 2012.

Tutorial: Graph Data Management Systems for New Application Domains.

Philippe Cudre-Mauroux, Sameh Elnikety.
VLDB 2011: 37th International Conference on Very Large Data Bases.
Seattle, WA, USA, 29 August 29 – 3 September 2011.

Embracing Heterogeneity: Improving Energy Efficiency for Interactive Services on Heterogeneous Data Center Hardware.

Yuxiong He, Sameh Elnikety.
AIDC 2011: AAAI Workshop on AI for Data Center Management and Cloud Computing.
San Francisco, CA, USA, 7-11 August 2011.

Edge Services: Generating Dynamic Web Content for E-Commerce.

Sameh Elnikety.
DSN 2004 Student Forum.
Florence, Italy, July 2004.

Generalized Snapshot Isolation and a Prefix-consistent Implementation.

Sameh Elnikety, Fernando Pedone, and Willy Zwaenepoel.
EPFL Technical Report IC/2004/21.

Lausanne, Switzerland, June 2004.

DynaServer: Dynamic Content Delivery Network for E-commerce.
Sameh Elnikety.
ECOOP 2004 Doctoral Symposium.
Oslo, Norway, June 2004.

Using Performance Reflection in Systems Software.
Rob Fowler, Alan Cox, Sameh Elnikety, and Willy Zwaenepoel.
HotOS 2003: Ninth Workshop on Hot Topics in Operating Systems.
Lihue, Hawaii, USA, 18-21 May 2003.

Admission Control for E-commerce Web Sites (extended abstract).
Sameh Elnikety, Erich Nahum, John Tracey and Willy Zwaenepoel.
ACM Workshop on Self-Managing Systems.
San Diego, CA, USA, June 2003.

Specification and Implementation of Dynamic Web Site Benchmarks.
C. Amza, E. Cecchet, A. Chanda, Alan L. Cox, S. Elnikety, R. Gil, J. Marguerite, K. Rajamani and W. Zwaenepoel.
WWC 2002: 5th Annual IEEE Workshop on Workload Characterization.
Austin, TX, USA, 25 November 2002.

Cooperative Backup System (work-in-progress report).
Sameh Elnikety, Mark Lillibridge, Mike Burrows, and Willy Zwaenepoel.
FAST 2002.
Monterey, CA, USA, January 2002.

Bottleneck Characterization of Dynamic Web Site Benchmarks.
C. Amza, E. Cecchet, A. Chanda, Alan L. Cox, S. Elnikety, R. Gil, J. Marguerite, K. Rajamani and W. Zwaenepoel.
IBM 3rd CAS conference.
Austin, TX, USA, February 2002.

A Deadline Driven Disk Scheduling Algorithm for Continuous Media Servers.
Sameh Elnikety, Walid Aref, and Mohamed Abougabal.
Alexandria University Technical Report.
Alexandria, Egypt, June 2000.

Academic Awards

- Best paper award at COMPSAC 2019
- Best student paper award at Middleware 2017
- Best paper award at SIGIR 2017
- Best paper runner-up award at WSDM 2015
- Best paper award at EuroSys 2007
- Two-year university fellowship from Rice University, March 1999
- Alexandria University fellowship for M.S. degree, April 1997
- University honor list, Alexandria University, 1996
- Governmental scholarship for undergraduate study at Alexandria University, 1991-1996

Software Artifacts

- PerfIso: Performance Isolation Service (deployed in Microsoft Bing).
- Global Snapshot Transactions with Clock-SI (deployed in Microsoft SQL-DW).
- GeoTrend: Real-time Spatio-temporal Aggregate Queries on Fast Microblog Streams (deployed in Bing).
- Adaptive query parallelization framework (deployed in Microsoft Bing).
- Dynamic query parallelization framework (deployed in Microsoft Bing).
- Predictor for web search query execution time (deployed in Microsoft Bing).
- Mercury: A Memory-Constrained Spatio-temporal Real-time Search on Microblogs (deployed in Bing).
- Pyramid: Indexing spatial objects (deployed in Microsoft MSN).
- Zeta: Scheduler for interactive server supporting partial execution (deployed in Microsoft Bing).
- Horton: Online query execution engine for large distributed graphs.
- Naxos: strong consistency in replicated databases.
- Tashkent+: memory-aware load balancing and update filtering.
- Tashkent: database replication middleware for snapshot isolated databases.
- Gatekeeper: transparent proxy for database admission control and request scheduling.
- Benchmarks for web sites with dynamic content.
- CBS: cooperative backup system.

Research Grants

- Hasler Foundation, grant to support “Autonomic Configuration of Database Clusters” by Sameh Elnikety and Willy Zwaenepoel, July 2006, 80K CHF (71K USD).

Selected Conference Presentations and Microsoft Talks

2021	MSR Research Summit	Serverless Computing on Harvested Resources
2020	Bing Tech Talk	NUMA performance impact in Azure Compute VMs
2019	UBC ECE	Project LEAP: Hosting Large Interactive Services on Public Clouds
2018	UT Austin	Project LEAP: Reducing VM Overhead for Interactive Services
2017	Bing Tech Talk	Tail Latency in Distributed Systems
2017	Bing AI.Next	Virtualizing Bing Latency-Sensitive Services on Azure
2016	Georgia Tech	Reducing Tail Latency in Distributed Systems
2016	ASPLOS	TPC: Target-Driven Parallelism Combining Prediction and Correction to Reduce Tail Latency in Interactive Services
2016	MSR Faculty Summit	The Art of Sharing Resources Transparently
2015	POSTECH Korea	Snapshot Isolation for Partitioned Data Stores Using Loosely Synchronized Clocks
2015	MSR Academic Summit Korea	System Support for Parallel Processing in Multicore Servers for Reducing High-Tail Latency.
2014	VLDB Tutorial	Systems for Big Graphs (3 hours)
2014	VLDB Paper	Horton: Processing Declarative Reachability Queries over Partitioned Graphs
2014	Qatar Foundation	Reducing Tail Latency in Large Search Engines
2013	SoCC Paper	Clock-SI: Snapshot Isolation Using Loosely Synchronized Clocks
2013	Univ. of Pittsburgh	How to query a large distributed graph? (Host: Panos Chrysanthis)
2013	Bing Tech Talk	Parallelizing Bing Queries
2012	Univ. of Minnesota	System Support for Managing Graph Data
2012	CIKM	G-SPARQL: A Hybrid Engine for Querying Large Attributed Graphs
2012	NSF	System support for managing graphs in the cloud
2011	Allerton	Scheduling for data center interactive services
2011	VLDB Tutorial	Graph data management systems for new application domains

2012	GDM	Online query execution on large distributed graphs
2008	Transparent scalable database replication	<p>Aster Data Systems, Redwood Shores, CA (host: Mayank Bawa)</p> <p>British Telecom, Adastral Park, Ipswich, UK (hosts: Nader Azarmi and Basim Majeed)</p> <p>CSIRO/ANU, Canberra, Australia (hosts: Athman Bouguettaya and Dimitrio Georgakopoulos)</p> <p>Delft University of Technology, Delft, Holland (host: Henk Sips)</p> <p>Hewlett-Packard Laboratories, Palo Alto, CA (host: John Janakiraman)</p> <p>Microsoft Research, Cambridge, UK (host: Antony Rowstron)</p> <p>Microsoft Research, Redmond, WA (host: John Douceur)</p> <p>Microsoft Research, Mountain View, CA (host: Doug Terry)</p> <p>Microsoft SQL Product Group, Redmond, WA (host: Paul Larson)</p> <p>Nanyang Technological University, Singapore (host: Hsu Wen-Jing)</p> <p>NEC Laboratories America, Princeton, NJ (host: Geoff Jiang)</p> <p>NICTA/UNSW, Sydney, Australia (host: Gernot Heiser)</p> <p>Singapore Management University, Singapore (hosts: Rajesh Balan and Steven Miller)</p> <p>Stanford University, Stanford, CA (host: Pat Hanrahan)</p> <p>University at Buffalo, Buffalo, NY (host: Vipin Chaudhary)</p> <p>University College London, London, UK (host: Bard Karp)</p> <p>University of Bristol, Bristol, UK (host: Nishan Canagarajah)</p> <p>University of Cyprus, Nicosia, Cyprus (host: Marios Dikaiakos)</p> <p>University of Lugano, Lugano, Switzerland (host: Mehdi Jazayeri)</p> <p>University of Sydney, Sydney, Australia (host: Sanjay Chawla)</p> <p>University of York, York, UK (host: John McDermid)</p> <p>Yahoo! Research, Barcelona, Spain (hosts: Hugo Zaragoza and Ricardo Baeza-Yates)</p> <p>Yahoo! Research, Santa Clara, CA (host: Brian Cooper)</p>
2009	EuroSys	Predicting replicated database performance from standalone database profiling
2009	MSR	Database replication for cloud services
2007	EuroSys	Tashkent+: memory-aware load balancing and update filtering in replicated databases
2006	Hasler Foundation	Autonomic configuration of database clusters
2006	EuroSys	Tashkent: uniting durability with transaction ordering
2005	SRDS	Generalized snapshot isolation
2004	WWW	Transparent admission control and request scheduling for e-commerce
2003	Middleware	Middleware architectures for generating dynamic content
2003	HotOS	Performance reflection in systems software
2002	IBM Austin	Throttling dynamic content web servers
2002	IBM Watson	Admission control and request scheduling for dynamic web servers
2002	IBM Watson	Scaling dynamic content servers
2002	WWC	Dynamic web servers benchmarks
2002	IBM Austin	Bottleneck characterization of dynamic Web site benchmarks
2000	Rice	Pizza talk: Real-time disk scheduling for continuous media servers
2000	COMPAQ SRC	Cooperative backup system.

Service

- **PC member**

- 2022: SRDS, Middleware, TheWebConf (WWW), DASFA, IEEE BigComputing
- 2021: Middleware, TheWebConf (WWW), DASFAA, Graph Computing, IEEE BigComputing
- 2020: SIGMETRICS (PC, and Awards Committee), DNS, IEEE BigComputing, Middleware (PC chair of industry track, and research track PC)
- 2019: DNS, IEEE BigComputing, Graph Computing
- 2018: IEE BigData, IEEE BigComputing

- 2017: IEEE BigData, IEEE SmartData, Supercomputing (Posters PC)
- 2016: IEEE BigData, Supercomputing (Posters PC), IEEE Autonomic Computing (Workshop chair)
- 2015: ICDCS, Feedback Computing, DASFAA
- 2014: VLDB, WWW, DASFAA, OPODIS
- 2013: VLDB, ICDE (twice for research track & industrial track), NETYS
- 2012: ICDCS, ICDE, CIKM, SNS, GDM (PC chair), Microsoft workshop on graph systems (PC chair)
- 2011: EDBT, ACC, ADVCOMP, SNS (PC chair)
- 2010: DAIS, ICT
- 2009: MAI, WDDDM
- 2008: MAI, ICOMP
- 2007: SOSP shadow PC
- **Other**
 - SoCC 2014 (Web chair)
 - NSF 2012 workshop on social networks and mobile applications in the cloud (Steering committee)
 - LADIS 2011 (Finance chair)
- **Journal reviewer**
 - The VLDB Journal (Springer), 2018, 2014, 2013, 2011
 - Journal of Supercomputing (Springer), 2017, 2014, 2012, 2011
 - Distributed and Parallel Databases (Springer), 2013, 2010
 - Social Network Analysis and Mining (Springer), 2013, 2011
 - Performance Evaluation (Elsevier), 2011
 - IEEE Internet Computing, 2011
 - Software: Practice and Experience, 2011, 2010
 - IEEE Transactions on Computers, 2010, 2009
 - IEEE Transactions on Industrial Electronics, 2009
 - IEEE Transactions on Knowledge and Data Engineering, 2014, 2013, 2009
 - LNCS Transactions on Aspect-Oriented Software Development (Springer) 2008
 - ACM Transactions on Database Systems
 - ACM Transactions on Internet Technology
 - Computer Networks Journal
 - Journal of Parallel and Distributed Computing.
- **Panelist**
 - “Management of Big Data Systems” panel at ICAC 2013, San Jose, CA, USA.
 - “Social networks on the Cloud” panel with Anhai Doan, Cong Yu, Divy Agrawal, and Evimaria Terzi, at the NSF Workshop on Social Networks and Mobile Applications in the Cloud, February 2012, Arlington, VA, USA.
 - “Dependable data management in the cloud” panel with Gustavo Alonso, Guillaume Pierre, Paul Murray and John Wilkes at the third workshop on dependable distributed data management (WDDDM) April 2009, Nuremberg, Germany.
- **Reviewer**
 - Reviewed submissions for VLDB 2011, HPDC 2011, Sigmod 2010, USENIX 2008, 2007, EuroSys 2010, 2009, 2006, PPOPP 2008, HotOS 2011, 2007, SOSP 2005, ISPASS 2003, Microsoft PhD Scholarships 2008, and Microsoft Research Grants for Mobile Healthcare in Africa 2010.
- Contributing author to Encyclopedia of Database Systems (Springer 2009 & 2016 Editions).
- Organized weekly pizza talks given by Rice faculty and graduate students as well as external visitors at Rice University, 2001-2002.

Teaching and Advising Experience

Institute	Role	Course
University of Washington	Instructor	Transaction Processing for E-Commerce (CSE-P 545) in 2012

EPFL	Coordinator Teaching assistant	Systems seminar Programming III
Rice University	Teaching assistant	Operating systems and concurrent programming Introduction to database systems Computer systems security Distributed systems
Alexandria University	Instructor	Introductory programming for non CS-majors

PhD thesis committee/co-advisor

- SungHwng Kim (POSTECH, 2019)
- Ahmad Javadi (Stonybrook, 2019)
- Jing Li (WUSTL, 2017)
- Amr Magdy Ahmed (UMN, 2016)
- Juan Mendivelso (UNAL, 2015)
- Jiaqing Du (EPFL, 2014)

I am fortunate to have hosted* or co-hosted** the following summer interns:

- Hwanju Kim*, KAIST, 2013
- Jin-woo Park*, Pohang University of Science and Technology (POSTECH), 2013
- Myeongjae Jeon**, Rice University, 2013
- Saehun Kim**, Pohang University of Science and Technology (POSTECH), 2013
- Yong Hun Eom**, UC Irvine, 2013
- Xi Yang**, Australian National University, 2013
- Ahmed Aly*, Purdue University, 2012, demo ICDE 2014
- Amr Magdy*, Ahmed, University of Minnesota, 2012, paper ICDE 2014 and demo ICDE 2014
- Myeongjae Jeon**, Rice University, 2012, paper Eurosys 2013.
- Igor Sychev**, Moscow State University of Instrument Engineering and Computer Sciences, 2012
- Jiaqing Du*, EPFL, 2011, papers SRDS 2013, SOCC 2013
- Shaolei Ren**, University of California, Los Angeles, 2011, paper ICAC 2013
- Juan Mendivelso*, Universidad Nacional de Colombia, 2011, paper SPIRE 2013
- Jinhan Kim**, Pohang University of Science and Technology (POSTECH), 2011, paper CAC 2013
- Jiaqing Du*, EPFL, 2010, paper SRDS 2013
- Mohamed Sarwat*, University of Minnesota, 2010, paper VLDB 2014 and demo ICDE 2012
- Hongyang Sun**, Nanyang Technological University, 2010, paper ICDCS 2011
- Michaela Bornea*, Athens University of Economics and Business, 2009, paper ICDE 2011
- Konstantinos Krikellas*, University of Edinburgh, 2008, paper ICDE 2010

During my graduate studies, I have collaborated with several students, including: Alok Menghrajani (EPFL, MS thesis), Eric Kankwendezazi (EPFL, semester project), Anupam Chanda (Rice, graduate research project), and Huma Jafry (Rice, internship).

Miscellaneous

Citizenship: Egyptian citizen and US permanent resident.

Hobbies: Swimming, Graphite drawing, hiking.

Languages: Fluent in Arabic and English. Good working knowledge of French.