

Curriculum Vitae

Oliver Kennedy

Department of Computer Science and Engineering
University at Buffalo, SUNY
338 Davis Hall
Buffalo, NY 14260

Phone: +1 (716) 645-1597
Web: <https://odin.cse.buffalo.edu/>

63 Bassett Road
Amherst, NY 14221

Email: okennedy@buffalo.edu

Education

PhD, Cornell University, Ithaca, NY 2011
Department of Computer Science
Thesis: “*Watch out for... What?: Monitoring and Uncertainty in Scientific Computing*”
Advisor: Christoph Koch

MS, Cornell University, Ithaca, NY 2008
Department of Computer Science

BS (with Honors), New York University, New York, NY 2005
Department of Computer Science

BS, Stevens Institute of Technology, Hoboken, NJ 2005
Department of Computer Engineering

Employment History

- Assistant Professor, Dept. of Comp. Sci. & Eng., SUNY Buffalo Sept. 2012 – present
- Postdoctoral Fellow, D.A.T.A. Lab, EPFL, Switzerland June 2011 – Aug. 2012
- Intern, Microsoft Research, Redmond, WA May 2010 – Aug. 2010
- Intern, Yahoo! Labs, Sunnyvale CA May 2009 – Aug. 2009
- Teaching/Research Assistant, Cornell University Sept. 2005 – May 2011
- Programmer, Humanities Computing Group, New York University Sept. 2003 – Aug. 2005
- Programmer, Parallel and Dist. Sys. Group, New York University Jan. 2001 – May 2003

Honors

The pVLDB 2012 paper titled “DBToaster: Higher-order Delta Processing for Dynamic, Frequently Fresh Views” was accepted to the “Best of VLDB 2012” issue of the VLDB Journal.

Oliver received a UB CSE Outstanding Funding Award for attracting over \$3m of funding to the University at Buffalo in 2016.

The SIGMOD 2017 paper titled “Beta Probabilistic Databases: A Scalable Approach to Belief Updating and Parameter Learning” was invited to submit a Best-Of-SIGMOD extended article to the ACM Transactions on Database Systems Journal.

Students Advised

PhD Students Advised (2 graduated, 7 current)

Niccolò Meneghetti (Graduated in 2016; Jointly advised with Jan Chomicki), Ying Yang (Graduated in 2017), Aaron Huber (Current Student), Carl Nuessle (Current Student; Jointly advised with Luke Ziarek), Gökhan Kul (Current Student; Jointly advised with Shambhu Upadhyaya), Poonam Kumari (Current Student), Saurav Singhi (Current Student; Jointly advised with Luke Ziarek), Ting Xie (Current Student), William Spoth (Current Student)

MS Students Advised (7 graduated, 3 current)

Ankur Upadhyay (Graduated in 2014; Jointly advised with Lukasz Ziarek), Nikhil Kataria (Graduated in 2014), Palaniappan Meiyappan (Graduated in 2014), Sumit Agarwal (Graduated in 2014; Jointly advised with Lukasz Ziarek), Saket Adusumilli (Graduated in 2015), Vinayak Karupphasamy (Graduated in 2015), Arindam Nandi (Graduated in 2016), Darshana Balakrishnan (Current Student; Jointly advised with Luke Ziarek), Gourab Mitra (Current Student; Jointly advised with Luke Ziarek), Shivang Aggarwal (Current Student)

BS Students Advised (4 graduated, 2 current)

Daniel Bellinger (Graduated in 2014; Jointly advised with Lukasz Ziarek), Thomas Mitchell (Graduated in 2015; Jointly advised with Varun Chandola and Shambhu Upadhyaya), Grant Wrazen (Graduated in 2017; Jointly advised with Luke Ziarek and Geoff Challen), Hank Lin (Graduated in 2017; Jointly advised with Luke Ziarek), Lisa Lu (Current Student; Jointly advised with Luke Ziarek), Olivia Alphonse (Current Student)

Professional Activities

Program Comittee Member

- Workshop: The VLDB PhD Workshop (VLDB-PhD) 2013
- Demo: VLDB Demo Track (VLDB-Demo) 2016 – 2018
- Journal: Proceedings of the VLDB Endowment (pVLDB) 2017 – 2018
- Conference: The ACM SIGMOD/PODS Conference (SIGMOD) 2015 – 2017
- Workshop: Provenance Week (PWEEK) 2016
- Workshop: Workshop on Human-In-the-Loop Data Analytics (HILDA) 2016 – 2017
- Conference: Symposium on Cloud Computing (SoCC) 2017

Reviewer

- Journal: The VLDB Journal (VLDBJ) 2013, 2017
- Journal: IEEE Transactions on Knowledge and Data Engineering (TKDE) 2013 – 2014
- Journal: ACM Transactions on Database Systems (TODS) 2015
- Journal: Computer Science Education (CSE) 2015
- Journal: Transactions on Information Technology (TOIT) 2016
- Journal: Knowledge and Information Systems (KAIS) 2017
- Conference: The International Conference on Database Theory (ICDT) 2017

Service

- WebDB Session Chair 2013
- NSF Panel Member 2014 – 2016
- SIGMOD Reproducibility Committee 2015 – 2017

Professional Memberships

- ACM 2012 – present
- ACM-SIGMOD 2010 – present
- ACM-Computer Science Teachers Assoc. 2012 – present
 - WNY Chapter Secretary Sept. 2013 – Sept. 2015
- IEEE 2009 – present

Volunteer Work

- LIBERTY Partnerships 2013 – present
 - Summer Mentor 2013 – 2014
 - ‘I Can’ Program Volunteer 2015 – 2017
- Coder Dojo Mentor 2013 – 2014
- Science is Elementary Class Leader 2014 – 2016

Departmental Service

- Dept. of Comp. Sci. and Eng., University at Buffalo, SUNY
 - Graduate Admissions Committee 2012 – present

- Graduate Advisory Board 2012 – present
- Distinguished Speaker Committee 2017 – present
- Faculty Search Committee 2012 – 2014
- Facilities Committee 2016
- Colloquium Committee 2012 – 2014
- * Chair 2014
- Grievances Committee 2015 – 2016
- Undergraduate Curriculum Development Subcommittee 2015 – 2016

Invited Talks

- *Penn State U.* — Don't Wrangle, Guess Instead (with Mimir) Jan. 2018
- *Harvard* — Just-In-Time Data Structures Dec. 2017
- *University of Washington* — Don't Wrangle, Guess Instead (with Mimir) Dec. 2018
- *Columbia* — Don't Wrangle, Guess Instead (with Mimir) Oct. 2017
- *University of Michigan* — Don't Wrangle, Guess Instead (with Mimir) Oct. 2017
- *Illinois Inst. Tech.* — Don't Wrangle, Guess Instead (with Mimir) Oct. 2017
- *University of Chicago* — Don't Wrangle, Guess Instead (with Mimir) Oct. 2017
- *Ohio State University* — Don't Wrangle, Guess Instead (with Mimir) Oct. 2017
- *Cornell* — Just-In-Time Data Structures Sep. 2017
- *ICDE* — Small Data (panel moderator) Apr. 2017
- *Buffalo Database Seminar Meetup* — Post-Truth Databases (lightning talk) Apr. 2017
- *Airbus* — Embracing Uncertainty with Mimir Oct. 2016
- *RIT* — Pocket Data Oct. 2016
- *LogicBlox* — Just-In-Time Data Structures Jun. 2016
- *NYU* — Embracing Uncertainty with Mimir May. 2016
- *HPE/Vertica* — Embracing Uncertainty with Mimir Mar. 2016
- *Carnegie Mellon* — Embracing Uncertainty with Mimir Feb. 2016
- *Buffalo District Board of Education* — LEGO NXT Professional Development Jan. 2015
- *Ohio State University* — PocketData: What's in your pocket? Oct. 2015
- *Oracle* — Just-In-Time Data Structures Jan. 2015
- *Pivotal* — Just-In-Time Data Structures Jan. 2015
- *Snowflake Computing Inc.* — Just-In-Time Data Structures Jan. 2015
- *Western New York CSTA Fall Conference* — Queries with Twitter Tutorial Oct. 2014
- *Buffalo Database Seminar Meetup* — Everything I Need to Know I Learned From PL Sept. 2014
- *Google CS4HS, ACM-CSTA, and Buffalo State* — Mobilize Prime Workshop (with Lukasz Ziarek, Sarbani Banerjee) March. 2014
- *Western New York CSTA Professional Development Workshops* — SQL Tutorial Jan. 2014
- *Buffalo Database Seminar Meetup* — Emerging Trends in Database Research Jan. 2014
- *Western New York CSTA 2013 CS4HS Workshop* — Python Workshop (with Lukasz Ziarek, Sarbani Banerjee) Jun. 2013
- *Cornell University Database Colloquium* — Laasie: The Log as a Service Infrastructure Feb. 2013
- *Western New York CSTA 2012 Fall Conference* — Python Workshop (with Lukasz Ziarek, Sarbani Banerjee) Oct. 2012

Courses Taught

Department of Computer Science and Engineering, SUNY Buffalo

- CSE 562 : Graduate Database Systems (planned) Spring 2018
- CSE 662 : Languages and Runtimes for Big Data (planned) Fall 2017
- CSE 562 : Graduate Database Systems Spring 2017
- CSE 199 : Undergraduate CS Seminar Recitation — 21 students Fall 2016
- CSE 705 : Database Research Seminar — 7 students Fall 2016
- CSE 662 : Languages and Runtimes for Big Data — 10 students Fall 2016
- CSE 462 : Undergraduate Database Systems — 72 students Spring 2016
- CSE 662 : Languages and Runtimes for Big Data — 29 students Fall 2015
- CSE 562 : Graduate Database Systems — 79 students Spring 2015
- CSE 704 : Databases, Programming Languages and Datastructures — 21 students Fall 2014
- CSE 562 : Graduate Database Systems — 148 students Spring 2014
- CSE 704 : Streaming, Incremental and Online Data Processing — 18 students Fall 2013
- CSE 562 : Graduate Database Systems — 101 students Spring 2013
- CSE 704 : Web-Scale Data Management Systems Seminar — 15 students Fall 2012

School of Computer and Communication Sciences, EPFL, Switzerland

- CS-422 : Advanced Databases (as Teaching Assistant) Spring 2011

Department of Computer Science, Cornell University

- CS 4410 : Operating Systems (as Practicum Instructor) Fall 2008
- CS 414 : Operating Systems (as Practicum Instructor) Spring 2006
- CS 414 : Operating Systems (as Practicum Instructor) Fall 2006

Grant Support

Pending Applications (4)

Title : Instant-On Access to Geospatial Data (ION-GEO)
Agency : NEEC: DD8
Role : PI
Amount : \$431,536
Effort : 33%

Title : CAREER: Declarative Uncertainty
Agency : NSF: CISE: IIS: III
Role : PI
Amount : \$542,276
Effort : 100%

Title : Intuitive Data Interpretation
Agency : Oracle University Relations
Role : PI
Amount : \$96,684
Effort : 100%

Title : III: SMALL: Databases for IoT
Agency : NSF: CISE: IIS: III
Role : Co-PI
Amount : \$493,964
Effort : 50%

Awarded Grants (5)

Title : CIF21 DIBBs: EI: Vizier, Streamlined Data Curation
Agency : NSF: ACI: DIBBS
Role : PI
Amount : \$2,749,699
Effort : 33%
Co-PIs : Boris Glavic, Juliana Freire
Supplements : NSF-REU (amount included in total above)

Title : CI-P: Supporting Pocket Scale Data Management Research
Agency : NSF: CISE: IIS: CRI
Role : PI
Amount : \$100,000
Effort : 33%
Co-PIs : Lukasz Ziarek, Geoffrey Challen

Title : Curating Uncertainty and Reliable Exploitation (CURE)
Agency : The US Naval Postgraduate School
Role : Co-I
Amount : \$265,357
Effort : 50%
Co-PI : Moises Sudit

Title : III: Small: Just in Time Datastructures
Agency : NSF: CISE: IIS: III
Role : PI
Amount : \$502,274
Effort : 50%
Co-PI : Lukasz Ziarek
Supplements : NSF-REU (amount included in total above)

Title : TWC: Medium: Collaborative: Data is Social: Exploiting Data Relationships to Detect Insider Attacks
Agency : NSF: CISE: SaTC
Role : PI
Amount : \$975,999 (Total Across Collaboration: \$1,215,973)
Effort : 25%
Co-PIs : Hung Ngo, Shambhu Upadhyaya, Varun Chandola
Peer Org. : University of Michigan Ann Arbor (\$239,974)
Supplements : NSF-REU (amount included in total above)

Gifts (1)

Title : Intuitive Data Interpretation
Agency : Oracle University Relations
Role : PI
Amount : \$89,187
Effort : 100%

Completed (3)

Title : Expressing Uncertainty Using the maybe System
Agency : Google Research Awards
Role : Co-PI
Amount : \$38,656
Effort : 33%
Co-PIs : Geoffrey Challen, Lukasz Ziarek

Title : Intuitive Data Interpretation
Agency : Oracle University Relations
Role : PI
Amount : \$90,455
Effort : 100%

Title : Intuitive Data Interpretation
Agency : Oracle University Relations
Role : PI
Amount : \$88,000
Effort : 50%
Co-PI : Jan Chomicki

Publications and Artifacts

Journal Publications (3)

LENSES: AN ON-DEMAND APPROACH TO ETL (Y. Yang, N. Meneghetti, R. Fehling, Z. H. Liu, D. Gawlick, **O. Kennedy**)— *Proceedings of the VLDB Endowment (pVLDB) 2015 (Volume 8, Number 12)*, 12 pages, acceptance rate 21.3%

DBTOASTER: HIGHER-ORDER DELTA PROCESSING FOR DYNAMIC, FREQUENTLY FRESH VIEWS (C. Koch, Y. Ahmad, **O. Kennedy**, M. Nolic, A. Nötzli, D. Lupei, A. Shaikhana; *Authors listed in descending order of seniority*)— *The VLDB Journal (VLDBJ) 2014 (Volume 23, Number 2; Special issue on the best papers of VLDB 2012)*, 25 pages, acceptance rate not available

DBTOASTER: HIGHER-ORDER DELTA PROCESSING FOR DYNAMIC, FREQUENTLY FRESH VIEWS (Y. Ahmad, **O. Kennedy**, C. Koch, M. Nolic)— *Proceedings of the VLDB Endowment (pVLDB) 2012*, 12 pages, acceptance rate 20.3%

Conference Publications (12)

BETA PROBABILISTIC DATABASES: A SCALABLE APPROACH TO BELIEF UPDATING AND PARAMETER LEARNING (N. Meneghetti, **O. Kennedy**, W. Gatterbauer)— *The ACM SIGMOD/PODS Conference (SIGMOD) 2017*, 12 pages, acceptance rate 19.6%

CONVERGENT INFERENCE WITH LEAKY JOINS (Y. Yang, **O. Kennedy**)— *International Conference on Extending Database Technology (EDBT) 2017*, 12 pages, acceptance rate 21.8%

ADAPTIVE SCHEMA DATABASES (W. Spoth, B. S. Arab, E. S. Chan, D. Gawlick, A. Ghoneimy, B. Glavic, B. Hammerschmidt, **O. Kennedy**, S. Lee, Z. H. Liu, X. Niu, Y. Yang; *Authors listed in alphabetical order*)— *The Conference on Innovative Data Systems Research (CIDR) 2017*, 6 pages, acceptance rate unknown (less than 30% in previous years)

JUST IN TIME DATASTRUCTURES (**O. Kennedy**, L. Ziarek)— *The Conference on Innovative Data Systems Research (CIDR) 2015*, 10 pages, acceptance rate unknown (less than 30% in previous years)

PIGOUT: MAKING MULTIPLE HADOOP CLUSTERS WORK TOGETHER (K. Jeon, S. Chandrashekhara, F. Shen, S. Mehra, **O. Kennedy**, S. Y. Ko; *Authors listed students first, then faculty*)— *The IEEE International Conference on Big Data (IEEE BigData) 2014*, 9 pages, acceptance rate 18.6%

INVENTORY ALLOCATION FOR ONLINE GRAPHICAL DISPLAY ADVERTISING USING MULTI-OBJECTIVE OPTIMIZATION (J. Yang, E. Vee, S. Vassilvitskii, J. Tomlin, J. Shanmugasundaram, T. Anastasakos, **O. Kennedy**)— *The International Conference on Operations Research and Enterprise Systems (ICORES)* 2012, 12 pages, acceptance rate not available

JIGSAW: EFFICIENT OPTIMIZATION OVER UNCERTAIN ENTERPRISE DATA (**O. Kennedy**, S. Nath)— *The ACM SIGMOD/PODS Conference (SIGMOD)* 2011, 12 pages, acceptance rate 25.0%

DBTOASTER: AGILE VIEWS FOR A DYNAMIC DATA MANAGEMENT SYSTEM (**O. Kennedy**, Y. Ahmad, C. Koch)— *The Conference on Innovative Data Systems Research (CIDR)* 2011, 12 pages, acceptance rate unknown (less than 30% in previous years)

PIP: A DATABASE SYSTEM FOR GREAT AND SMALL EXPECTATIONS (**O. Kennedy**, C. Koch)— *The International Conference on Data Engineering (ICDE)* 2010, 12 pages, acceptance rate 21.0%

DYNAMIC APPROACHES TO IN-NETWORK AGGREGATION (**O. Kennedy**, C. Koch, A. Demers)— *The International Conference on Data Engineering (ICDE)* 2009, 4 pages, acceptance rate 17.0%

THE FACE OF MEANING (G. Shawver, **O. Kennedy**)— *The Joint International Conference of the Association for Computers and the Humanities and the Association for Literary and Linguistic Computing (ACH/ALLH)* 2005, 6 pages, acceptance rate not available

SECOND LANGUAGE COMPREHENSION IN AN ONLINE MULTIMEDIA ENVIRONMENT: INTEGRATING THEORY AND PRACTICE. (E. Lage-Otero, **O. Kennedy**)— *The World Conference on E-Learning in Corporate, Government, Healthcare, and Higher Education (E-Learn)* 2004, 2 pages, acceptance rate not available

Workshop Publications (10)

COMMUNICATING DATA QUALITY IN ON-DEMAND CURATION (P. Kumari, S. Achmiz, **O. Kennedy**)— *International Workshop on Quality in DataBases (QDB)* 2016, 4 pages, acceptance rate not available

THE EXCEPTION THAT IMPROVES THE RULE (J. Freire, B. Glavic, **O. Kennedy**, H. Mueller; *Authors listed in alphabetical order*)— *Workshop on Human-In-the-Loop Data Analytics (HILDA)* 2016, 6 pages, acceptance rate 50.0%

PROVENANCE-AWARE VERSIONED DATAWORKSPACES (X. Niu, B. Arab, D. Gawlick, Z. H. Liu, V. Krishnaswamy, **O. Kennedy**, B. Glavic)— *USENIX Workshop on Theory and Practice of Provenance (TaPP)* 2016, 6 pages, acceptance rate not available

ETTU: ANALYZING QUERY INTENTS IN CORPORATE DATABASES (G. Kul, D. T. Luong, T. Xie, P. Coonan, V. Chandola, **O. Kennedy**, S. Upadhyaya)— *Workshop on Empirical Research Methods in Information Security (ERMIS)* 2016, 4 pages, acceptance rate not available

POCKET DATA: THE NEED FOR TPC-MOBILE (**O. Kennedy**, J. A. Ajay, G. Challen, L. Ziarek)— *The TPC Technology Conference on Performance Evaluation & Benchmarking (TPC-TC)* 2015, 16 pages, acceptance rate not available

MAYBE WE SHOULD ENABLE MORE UNCERTAIN MOBILE APP PROGRAMMING (G. Challen, J. A. Ajay, N. DiRienzo, **O. Kennedy**, A. Maiti, A. Nandugudi, S. Shantharam, J. Shi, G. P. Srinivasa, L. Ziarek; *Authors listed in alphabetical order*)— *The ACM Workshop on Mobile Computing Systems and Applications (HotMobile)* 2015, 6 pages, acceptance rate 28.75%

DETECTING THE TEMPORAL CONTEXT OF QUERIES (**O. Kennedy**, Y. Yang, J. Chomicki, R. Fehling, Z. H. Liu, D. Gawlick; *Authors listed UB first, then Oracle*)— *Enabling Real-Time Business Intelligence (BIRTE)* 2014, 17 pages, acceptance rate unknown; average rate approximately 45.0%

MONADIC LOGS FOR COLLABORATIVE WEB APPLICATIONS (S. Agarwal, D. Bellinger, **O. Kennedy**, A. Upadhyay, L. Ziarek; *Authors listed in alphabetical order*)— *The International Workshop on the Web and Databases (WebDB)* 2013, 6 pages, acceptance rate 25.0%

NEXUS: A NEW OPERATING SYSTEM FOR BUILDING TRUSTWORTHY APPLICATIONS (A. Shieh, D. Williams, K. Walsh, **O. Kennedy**, P. Reynolds, E. G. Sirer, F. Schneider)— *NSDI Poster Session (NSDI-Poster)* 2007, poster, acceptance rate not available

XMON-BGP: SECURING BGP USING USING EXTERNAL SECURITY MONITORS (P. Reynolds, **O. Kennedy**, E. G. Sirer, F. Schneider)— *NSDI Poster Session (NSDI-Poster)* 2007, poster, acceptance rate not available

Technical Reports (7)

BETA PROBABILISTIC DATABASES: A SCALABLE APPROACH TO BELIEF UPDATING AND PARAMETER LEARNING (N. Meneghetti, **O. Kennedy**, W. Gatterbauer)— *SUNY Buffalo Computer Science & Engineering Report (UB CSE)*, 12 pages

SUMMARIZING LARGE QUERY LOGS IN ETTU (G. Kul, D. T. Luong, T. Xie, P. Coonan, V. Chandola, **O. Kennedy**, S. Upadhyaya)— *The ArXiv (ArXiv)*, 12 pages

MIMIR: BRINGING CTABLES INTO PRACTICE (A. Nandi, Y. Yang, **O. Kennedy**, B. Glavic, R. Fehling, Z. H. Liu, D. Gawlick)— *The ArXiv (ArXiv)*, 12 pages

BarQL: COLLABORATING THROUGH CHANGE (**O. Kennedy**, L. Ziarek)— *SUNY Buffalo Computer Science & Engineering Report (UB CSE) #2013-03*, 10 pages

BarQL: COLLABORATING THROUGH CHANGE (**O. Kennedy**, L. Ziarek)— *The ArXiv (ArXiv)*, 9 pages

INVENTORY ALLOCATION FOR ONLINE GRAPHICAL DISPLAY ADVERTISING (J. Yang, E. Vee, S. Vassilvitskii, J. Tomlin, J. Shanmugasundaram, **O. Kennedy**)— *The ArXiv (ArXiv)*, 25 pages

DYNAMIC APPROACHES TO IN-NETWORK AGGREGATION (**O. Kennedy**, C. Koch, A. Demers)— *The ArXiv (ArXiv)*, 11 pages

Patents (2)

EFFICIENT OPTIMIZATION OVER UNCERTAIN DATA (S. Nath, S. Lee, S. Smyl, C. Loboz, **O. Kennedy**)— USA Patent #8,650,180

SYSTEM FOR DISPLAY ADVERTISING OPTIMIZATION USING CLICK OR CONVERSION PERFORMANCE (**O. Kennedy**, T. Anastasakos, J. Mao, E. Vee, J. Shanmugasundaram, S. Vassilvitskii, J. Yang, J. Tomlin)— USA Patent #8,650,180

Artifacts (4)

JUST-IN-TIME DATA STRUCTURES (software) First released July 2014
An index structure that dynamically adapts to variable workloads.

<http://github.com/UB0din/jitd>

As of July 27, 2017, Just-in-Time Data Structures had 5 GitHub watchers, 5 GitHub stars, and 16 GitHub forks

MIMIR (software) First released January 2015

A probabilistic overlay for database systems

<http://mimirdb.info>

As of July 27, 2017, Mimir had 11 GitHub watchers, 8 GitHub stars, and 8 GitHub forks

DBTOASTER (software) First released May 2012

A tool that dynamically compiles database engines optimized for any given query workload.

<http://www.dbtoaster.org>

As of December 10, 2014, DBToaster had 4191 unique website visitors and 1032 unique downloads

PIGOUT (software) First released October 2014

A multi-cluster runtime for Pig Latin

<http://pigout.cse.buffalo.edu>