

Curriculum Vitae

Oliver Kennedy

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

63 Bassett Road
Amherst, NY 14221

Phone: +1 (716) 645-1597

Email: okennedy@buffalo.edu

Web: <http://odin.cse.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

Professional Activities

Reviewer / Program Committee Member

- Journal: The VLDB Journal (VLDBJ) 2013
- 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
- Workshop: The VLDB PhD Workshop (VLDB-PhD) 2013
- Demo: VLDB Demo Track (VLDB-Demo) 2016
- Journal: Proceedings of the International Conference on Very Large Databases (Hybrid Journal/Conference) (pVLDB) 2017
- 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

Service

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

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
- 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
 - Faculty Search Committee 2012 – 2014
 - Colloquium Committee 2012 – 2014
 - * Chair 2014
 - Grievances Committee 2015 – present
 - Undergraduate Curriculum Development Subcommittee 2015 – present

Invited Talks

- *NYU* — Embracing Uncertainty with Mimir May 2016
- *HPE/Vertica* — Embracing Uncertainty with Mimir Mar. 2016
- *Buffalo District Board of Education* — LEGO NXT Professional Development Jan. 2015
- *Ohio State Computer Science Colloquium* — 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

- Department of Computer Science and Engineering, SUNY Buffalo
 - CSE 562 : Graduate Database Systems (planned) Spring 2017
 - CSE 662 : Languages and Runtimes for Big Data (planned) Fall 2016
 - CSE 462 : Undergraduate Database Systems 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 (3)

Title : CIF21 DIBBs: El: Vizier, Streamlined Data Curation
Agency : NSF: ACI: DIBBS
Role : PI
Amount : \$2,725,699
Effort : 33%
Co-PIs : Boris Glavic, Juliana Freire

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

Title : III: Small: Just in Time Datastructures
Agency : NSF: CISE: IIS: III
Role : PI
Amount : \$494,274
Effort : 50%

Awarded Grants (2)

Title : Curating Uncertainty and Reliable Exploitation (CURE)
Agency : The US Naval Postgraduate School
Role : Co-I
Amount : \$134,139
Effort : 50%
Co-PI : Moises Sudit

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 (3)

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

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%

Completed (1)

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, **O. Kennedy**)— *Proceedings of the International Conference on Very Large Databases (Hybrid Journal/Conference) (pVLDB) 2015*, 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 International Conference on Very Large Databases (Hybrid Journal/Conference) (pVLDB) 2012*, 12 pages, acceptance rate 20.3%

Conference Publications (9)

- “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 23.2%
- “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 (7)

- “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%

Technical Reports (5)

- “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
- “*Bar_{QL}*: Collaborating Through Change” (**O. Kennedy**, L. Ziarek)— *SUNY Buffalo Computer Science & Engineering Report (UB CSE) #2013-03*, 10 pages
- “*Bar_{QL}*: 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)

- **DBToaster** (software)
A tool that dynamically compiles database engines optimized for any given query workload.
<http://www.dbtoaster.org>
First released May 2012
As of December 10, 2014, DBToaster had 4191 unique website visitors and 1032 unique downloads
- **Just-in-Time Data Structures** (software)
An index structure that dynamically adapts to variable workloads.

<http://github.com/UB0din/jitd>
First released July 2014

- **Mimir** (software)
A probabilistic overlay for database systems
<http://mimirdb.info>
First released January 2015
- **PigOut** (software)
A multi-cluster runtime for Pig Latin
<http://pigout.cse.buffalo.edu>
First released October 2014