

Khanh Nguyen

229 Peterson
435 Naggle Street
3112 TAMU
College Station, TX 77843-3112

khanhntn@tamu.edu
(979) 458-1904
<http://people.tamu.edu/~khanhntn>

Education

University of California, Los Angeles

– Ph.D. in Computer Science 2018 - 2019

University of California, Irvine

– Ph.D. in Computer Science 2012 - 2018
Advisor: Prof. Harry Xu (Moved to UCLA)

University of California, Irvine

– M.S. in Computer Science 2012 - 2015
– B.S. in Computer Science 2010 - 2012

Fullerton College

– A.S. in Computer Science 2006 - 2010
– A.A. in Mathematics 2006 - 2010

Research Interests

I am broadly interested in developing programming language extensions, compiler and runtime supports to improve system performance — efficiency, scalability, latency — for Big Data systems.

Experience

Assistant Professor Aug 2020 -
Texas A&M University, College Station, TX

Visiting Researcher Aug 2019 - Aug 2020
Google, Seattle, WA
• Hosts: Dr. Kathryn McKinley and Dr. Martin Maas.

Research Assistant Sep 2018 - Jun 2019
University of California, Los Angeles, CA

Software Engineer Intern Jun 2018 - Sep 2018
Google, Seattle, WA
• Supervised by Dr. Kathryn McKinley and Chris MacGregor.

Research Assistant Jun 2012 - Jun 2018
University of California, Irvine, CA

Teaching Assistant Fall 2013, Summer 2017
University of California, Irvine, CA

Proctor Jul 2011 - Sep 2011
University of California, Irvine, CA

Web Assistant Aug 2010 - Dec 2010
Webstorm Internet Media, Newport Beach, CA

Publications

- P11. Martin Maas, Chris Kennelly, **Khanh Nguyen**, Darryl Gove, Kathryn S. McKinley, and Paul Turner. *Adaptive Huge-Page Subrelease for Non-Moving Memory Allocators in Warehouse-Scale Computers*, ACM SIGPLAN International Symposium on Memory Management (ISMM), Virtual, Canada, June 2021.
- P10. Chenxi Wang, Haoran Ma, Shi Liu, Yuanqi Li, Zhenyuan Ruan, **Khanh Nguyen**, Michael Bond, Ravi Netravali, Miryung Kim, and Harry Xu. *Semeru: A Memory-Disaggregated Managed Runtime*, the 14th USENIX Symposium on Operating Systems Design and Implementation (OSDI), Banff, Alberta, Canada, November 2020.
- P9. Christian Navasca, Cheng Cai, **Khanh Nguyen**, Brian Demsky, Shan Lu, Miryung Kim, and Harry Xu. *Gerenuk: Thin Computation over Big Native Data Using Speculative Program Transformation*, the 27th ACM Symposium on Operating Systems Principles (SOSP), Huntsville, Ontario, Canada, October 2019.
- P8. Cheng Cai, Qirun Zhang, Zhiqiang Zuo, **Khanh Nguyen**, Harry Xu, and Zhendong Su. *Calling-to-Reference Context Translation via Constraint-Guided CFL-Reachability*, ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), Philadelphia, PA, USA, June 2018.
- P7. **Khanh Nguyen**, Lu Fang, Christian Navasca, Harry Xu, Brian Demsky, and Shan Lu. *Skyway: Connecting Managed Heaps in Distributed Big Data Systems*, the 23rd International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), Williamsburg, VA, USA, March 2018.
- P6. **Khanh Nguyen**, Kai Wang, Yingyi Bu, Lu Fang, and Harry Xu. *Understanding and Combating Memory Bloat in Managed Data-Intensive Systems*, ACM Transactions on Software Engineering and Methodology (TOSEM), Volume 26 Issue 4, January 2018.
- P5. **Khanh Nguyen**, Lu Fang, Harry Xu, Brian Demsky, Shan Lu, Sanazsadat Alamian, and Onur Mutlu. *Yak: A High-Performance Big-Data-Friendly Garbage Collector*, the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI), Savannah, GA, USA, November 2016.
- P4. **Khanh Nguyen**, Lu Fang, Harry Xu, and Brian Demsky. *Speculative Region-based Memory Management for Big Data Systems*, the 8th Workshop on Programming Languages and Operating Systems (PLOS), Monterey, CA, USA, October 2015.
- P3. Lu Fang, **Khanh Nguyen**, Harry Xu, Brian Demsky, and Shan Lu. *Interruptible Tasks: Treating Memory Pressure As Interrupts for Highly Scalable Data-Parallel Programs*, the 25th ACM Symposium on Operating Systems Principles (SOSP), Monterey, CA, USA, October 2015.
- P2. **Khanh Nguyen**, Kai Wang, Yingyi Bu, Lu Fang, Jianfei Hu, and Harry Xu. *FACADE: A Compiler and Runtime Support for (Almost) Object-Bounded Big Data Applications*, the 20th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), Istanbul, Turkey, March 2015.
- P1. **Khanh Nguyen** and Harry Xu. *Cachetor: Detecting Cacheable Data to Remove Bloat*, the 9th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering (ESEC/FSE), Saint Petersburg, Russia, August 2013.

Service

| | |
|------------------|---|
| ISMM 2021 | External Review Committee |
| VEE 2021 | Program Committee |
| ATC 2021 | Program Committee |
| ASPLOS 2021 | External Review Committee |
| NSF | Reviewer 2020 |
| CACM | Reviewer |
| ASPLOS 2020 | External Review Committee |
| ECOOP 2019 | Artifact Evaluation Committee |
| ISMM 2018 | Program Committee |
| PLDI 2015 | Artifact Evaluation Committee |
| SOSP 2015 | Student Volunteer |
| Sub-reviewer for | FSE-SRC 2016, ISMM 2016, ECOOP 2016, PLDI 2015, ECOOP 2015 JTOC, TOSEM |

Teaching

- T2. CSCE 434: Compiler Design: Fall 2021
- T1. CSCE 689: Big Data Systems Design and Implementation: Spring 2021

Funding

- F1. Advancing Data Analytics Engines for Large Scale Autonomous Cyber Defense, 11/2020-11/2021, Co-PI, Sponsor: Department of Defense.

Honors and Awards

| | |
|---|-------------|
| Google PhD Fellowship Recipient | 2017 - 2019 |
| Facebook PhD Fellowship Finalist | 2017 |
| ACM Graduate SRC Finalist | PLDI 2016 |
| ACM Graduate SRC Bronze medalist | PLDI 2014 |
| UC Irvine - Donald Bren School of ICS Dean's Fellowship | 2012 - 2016 |
| Golden Key International Honor Society | 2010 |

Last updated: May 24, 2021