

Loc Dac Hoang

CONTACT INFORMATION

E-mail: loc@cs.utexas.edu
Webpage: <https://www.cs.utexas.edu/~loc/>

RESEARCH INTERESTS

Distributed graph analytics, distributed algorithms and systems, high-performance parallel programming, and (distributed) graph neural networks

EDUCATION

The University of Texas at Austin August 2018 - August 2021

Ph.D. advised by Keshav Pingali

- **Thesis:** Accelerating Graph Computation with System Optimizations and Algorithm Design
- Developed a fast distributed streaming graph partitioner that allowed a user to quickly and flexibly partition a large graph
- Implemented a provably efficient distributed betweenness-centrality algorithm
- Developed a distributed graph neural network (GNN) system using the idea that GNN computation can be expressed as a graph operator to leverage existing distributed graph analytics research
- Formulated the graph transformer network (which identifies important metapaths in heterogeneous graphs) as a graph problem rather than a matrix problem to reduce memory and computational costs

The University of Texas at Austin

August 2013 - May 2017

5-Year Integrated Bachelor's/Master's Degree

- Part of the Turing Scholars honors program as an undergraduate
- Cumulative Graduate (August 2016 - May 2017) GPA: **3.61/4.00**
- Cumulative Undergraduate (August 2013 - May 2016) GPA: **3.89/4.00**

PROFESSIONAL EXPERIENCE

Software Engineer III

July 2022 - present

Software Engineer

Aug. 2021 - July 2022

KatanaGraph

- Developed KatanaGraph's macrobatch subgraph sampler which allows the engine to perform subgraph sampling significantly faster by sampling and fetching data for multiple minibatches in a single communication relay; subgraphs could be exported for use with DGL and PyG
- Developed Katana-native graph neural network (GNN) layers for use with Katana subgraphs that could provide performance benefits over Torch GNN layers via memory management techniques and integrated with PyTorch's autograd so that it could be used with other PyTorch modules
- Identified bottlenecks and optimized the KatanaGraph AI platform which provides users with a high-level interface to create and train models.

Graduate Research Assistant

Sep. 2018 - May 2020, Sep. 2020 - Aug. 2021

The University of Texas at Austin

- Conducting research mainly on (distributed) graph analytics which involved developing new features for distributed Galois

Analyst Commodity I (Part-time)

December 2020 - May 2021

AbbVie/Hiregenics

- Built an efficient graph transformer network implementation using the KatanaGraph engine in order to identify important metapaths in heterogeneous graphs; the results were included in an arXiv preprint

Engineering Intern

June 2020 - August 2020

KatanaGraph

- Co-designed and co-developed the initial KatanaGraph graph querying engine

Software Engineering Intern

January 2020 - May 2020

Intel Corporation

- Developed graph algorithms in CUDA and Intel's DPC++ and analyzed their performance

Research Engineering/Scientist Associate

June 2017 - August 2018

The University of Texas at Austin

- Implemented new features/algorithms for distributed graph analytics, conducted research, and wrote conference papers

PUBLICATIONS

Xuhao Chen, Roshan Dathathri, Gurbinder Gill, **Loc Hoang**, Keshav Pingali. **Sandslash: A Two-Level Framework for Efficient Graph Pattern Mining**. *ICS 2021*, June 2021.

Loc Hoang*, Udit Agarwal*, Gurbinder Gill, Roshan Dathathri, Abhik Seal, Brian Martin, Keshav Pingali. **Optimizing Graph Transformer Networks with Graph-based Techniques**. *arXiv*, 2021.

Loc Hoang, Xuhao Chen, Hohan Lee, Roshan Dathathri, Gurbinder Gill, Keshav Pingali. **Efficient Distribution for Deep Learning on Large Graphs**. *GNNSys 2021*, April 2021.

Hohan Lee, David Wong, **Loc Hoang**, Roshan Dathathri, Gurbinder Gill, Vishwesh Jatala, David Kuck, Keshav Pingali. **A Study of APIs for Graph Analytics Workloads**. *IISWC 2020*, October 2020.

Gurbinder Gill, Roshan Dathathri, **Loc Hoang**, Ramesh Peri, Keshav Pingali. **Single Machine Graph Analytics on Massive Datasets Using Intel Optane DC Persistent Memory**. *VLDB 2020*, August 2020.

Vishwesh Jatala, Roshan Dathathri, Gurbinder Gill, **Loc Hoang**, V. Krishna Nandivada, Keshav Pingali. **A Study of Graph Analytics for Massive Datasets on Large-Scale Distributed GPUs**. *IPDPS 2020*, May 2020.

Vishwesh Jatala, **Loc Hoang**, Roshan Dathathri, Gurbinder Gill, V. Krishna Nandivada, Keshav Pingali. **An Adaptive Load Balancer for Graph Analytical Applications on GPUs**. *arXiv*, 2019.

Loc Hoang*, Vishwesh Jatala*, Xuhao Chen, Udit Agarwal, Roshan Dathathri, Gurbinder Gill, Keshav Pingali. **DistTC: High Performance Distributed Triangle Counting**. *HPEC 2019 Graph Challenge*, September 2019.

Roshan Dathathri, Gurbinder Gill, **Loc Hoang**, Hoang-Vu Dang, Vishwesh Jatala, V. Krishna Nandivada, Marc Snir, Keshav Pingali. **Gluon-Async: A Bulk-Asynchronous System for Distributed and Heterogeneous Graph Analytics**. *PACT 2019*, September 2019.

Gurbinder Gill, Roshan Dathathri, **Loc Hoang**, Keshav Pingali. **A Study of Partitioning Policies for Graph Analytics on Large-scale Distributed Platforms**. *VLDB 2019*, August 2019.

Loc Hoang, Roshan Dathathri, Gurbinder Gill, Keshav Pingali. **CuSP: A Customizable Streaming Edge Partitioner for Distributed Graph Analytics**. *IPDPS 2019*, May 2019.

Roshan Dathathri*, Gurbinder Gill*, **Loc Hoang**, Keshav Pingali. **Phoenix: A Substrate for Resilient Distributed Graph Analytics Persistent Memory**. *ASPLOS 2019*, April 2019.

Loc Hoang*, Matteo Pontecorvi*, Roshan Dathathri, Gurbinder Gill, Bozhi You, Keshav Pingali, Vijaya Ramachandran. **A round-efficient distributed betweenness centrality algorithm**. *PPoPP 2019*, February 2019.

Roshan Dathathri*, Gurbinder Gill*, **Loc Hoang**, Keshav Pingali. **Gluon: a communication-optimizing substrate for distributed heterogeneous graph analytics**. *PLDI 2018*, June 2018.

Hoang-Vu Dang, Roshan Dathathri, Gurbinder Gill, Alex Brooks, Nikoli Dryden, Andrew Lenharth, **Loc Hoang**, Keshav Pingali, Marc Snir. **A Lightweight Communication Runtime for Distributed Graph Analytics**. *IPDPS 2018*, May 2018.

“*” indicates equal contribution by authors.

TEACHING EXPERIENCE	Teaching Assistant September 2016 - May 2017 CS 345: Programming Languages at The University of Texas at Austin <ul style="list-style-type: none">Conducted review sessions, graded most student assignments, held office hours, and created relevant review resources for students
SERVICE	Reviewer for ETRI Journal 2022 Reviewer for The Journal of Supercomputing 2022 Reviewer for Transactions on Knowledge and Data Engineering 2021, 2022 Reviewer for IEEE Computational Intelligence Magazine 2020 PPoPP Artifact Evaluation PC Member 2019, 2020 Reviewer for IEEE Transactions on Knowledge and Data Engineering 2019
AWARDS AND HONORS	Best Paper Nominee <i>Gluon-Async: A Bulk-Asynchronous System for Distributed and Heterogeneous Graph Analytics</i> at PACT 2019 College Scholar The University of Texas at Austin, UT Honors Day 2016 Distiguished College Scholar The University of Texas at Austin, UT Honors Day 2015 University Honors The University of Texas at Austin, Fall 2013, Spring 2014, Fall 2014, Spring 2015, and Spring 2016
SKILLS	Programming Languages Main Languages: C++, Python Have used in the past: C, Java, Scala, Prolog, ML, Haskell Software and Tools vi(m), screen, tmux, bash, LaTeX, git, GitHub, Intel Vtune Profiler Human Languages Proficient English, Working-Proficiency Vietnamese
LAST UPDATED	May 2023