Duong Hoang

hoangthaiduong@gmail.com ❖ (801) 635-5171 ❖ https://hoang-dt.github.io

I am well-versed in data compression, visualization, computer graphics, scientific computing, signal processing, and machine learning. I learn fast and like solving interesting problems in all areas of computer science and applied mathematics. I also love writing fast and elegant CPU and GPU code.

WORK EXPERIENCE

Lawrence Livermore National Laboratory (2015, 2016, 2017, 2018)
Student Intern (Supervisor: Peter Lindstrom)
Developed novel mixed-precision and mixed-resolution encodings for efficient compression of terabyte-size 3D scientific datasets, as part of LLNL's Variable Precision Computing project. My work resulted in three first-author publications in top venues (TVCG and LDAV).

Core Resolution Pte Ltd (Singapore, 2012 – 2014)
Senior Research Engineer
Led a 5-engineer team that developed novel computer vision techniques for inspection and failure analysis of integrated circuits.

- Implemented a C++ library that aligns an inspection image, with unknown shear, rotation, and scales, inside a design layout, to nanometer accuracy and within seconds. Integrated into Avalon MaskView, an industry-standard layout navigation tool by Synopsys.

National University of Singapore (2011 – 2012)
Research Assistant (Supervisor: KangKang Yin)

- A TVCG paper that uses canonical correspondence analysis (CCA) to learn mesh deformation models from sparse training data. Implemented an inverse kinematics technique to reduce the input dimensionality from 20 to 6, greatly improving the running time.
- An SCA paper on control of rolling motions for human-like characters. Added per-pixel lighting and real-time ambient occlusion to an Ogre3D-based framework to highlight contact shadows, which are important for the perception of realistic animations.

PROJECTS

- **A point cloud compression scheme** that has asymptotically constant memory footprints regardless of data size. Runs 7x faster than the state-of-the-art MPEG while using 3x less memory. Compresses a billion particles using only 50MB of RAM.
- **A compressed file format** for 3D floating-point grids that supports adaptive queries in resolution and precision. Decodes 100x faster and uses 1000x less memory than other state-of-the-art compressors. Can compress 1TB to 7MB with minimal quality loss.
- **A screen-space ambient occlusion method** for real-time graphics that runs faster compared to the (then) state-of-the-art methods from NVIDIA and Blizzard, while being free of noisy and blurry artifacts. Used by Ubisoft in the game Assassin's Creed III.

AWARDS

Best paper
IEEE PacificVis 2022
Adaptive multilinear meshes

NASA Earth Exchange Award
$100K/year research grant (2021)

Best paper honorable mention
IEEE LDAV 2021
Compression of large point clouds

Best student poster
LLNL Poster Symposium 2017
Variable precision bit streaming

Best project
U of Utah Visualization 2015
Interactive visualization of papers

Dean's list
NUS – 2009
Top 5% of the cohort

Most accurate implementation
NUS Software Engineering Project
Implemented a 15KLOC static program analyzer in C#

Most innovative game
CONTRAST game design competition

EDUCATION

University of Utah
(2014 – 2022)
Ph.D., Computer Science
GPA: 3.96/4

National University of Singapore
(2010 – 2012)
Masters of Computing GPA: 4.67/5

National University of Singapore
(2006 – 2010)
Bachelor of Computing GPA: 4.16/5

SKILLS

- **Programming:** C++, Python, C#, Java, JavaScript, MATLAB, Mathematica, Prolog, Racket
- **CS & Math:** Compression, Graphics, Vision, Signal processing, Numerical methods, Topological data analysis, Machine learning

PUBLICATIONS & CONTRIBUTIONS

[1] H. Bhatia, D. Hoang, N. Morrical, G. Morrison, V. Pascucci, P.-T. Bremer, P. Lindstrom, "AMM: Adaptive Multilinear Meshes", *IEEE Transactions on Visualization and Computer Graphics,* to appear.

[2] D. Hoang, H. Bhatia, P. Lindstrom, and V. Pascucci, "High-quality and low-memory-footprint progressive decoding of large-scale particle data", in *IEEE Symposium on Large Data Analysis and Visualization (LDAV)*, 2021.

[3] K. Fan, D. Hoang, S. Petruzza, T. Gilray, P. Valerio, and S. Kumar, "Load-balancing parallel I/O of compressed hierarchical layouts", in *IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC)*, 2021.

[4] D. Hoang, B. Summa, H. Bhatia, P. Lindstrom, P. Klacansky, W. Usher, P.-T. Bremer, and V. Pascucci, "Efficient and flexible hierarchical data layouts for a unified encoding of scalar field precision and resolution", *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 2, pp. 603–613, 2021.

[5] D. Hoang, P. Klacansky, H. Bhatia, P.-T. Bremer, P. Lindstrom, and V. Pascucci, "A study of the trade-off between reducing precision and reducing resolution for data analysis and visualization", *IEEE Transactions on Visualization and Computer Graphics*, vol. 25, no. 1, pp. 1193–1203, 2019.

[6] S. Kumar, S. Petruzza, D. Hoang, and V. Pascucci, "Accelerating in-situ feature extraction of large-scale combustion simulation with subsampling", in *ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC)*, 2017.

[7] S. Kumar, D. Hoang (joint first author), S. Petruzza, J. Edwards, and V. Pascucci, "Reducing network congestion and synchronization overhead during aggregation of hierarchical data", in *IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC)*, 2017.

[8] K. Wu, D. Hoang, and A. Lex, "Visualizing publication data", in *IEEE Visualization Conference (VIS)*, 2016.

[9] T.-D. Hoang and K.-L. Low, "Efficient screen-space approach to high-quality multiscale ambient occlusion", *The Visual Computer*, vol. 28, no. 3, pp. 289–304, 2012.

[10] T.-D. Hoang and K.-L. Low, "Multi-resolution screen-space ambient occlusion", in *Computer Graphics International Conference (CGI)*, 2011.

[11] B. Peng, K.-L. Low, and T.-D. Hoang, "Real-time csg rendering using fragment sort", in *ACM Symposium on Virtual Reality Software and Technology (VRST)*, p. 99–100, 2010.

[12] H. Huang, K. Yin, L. Zhao, Y. Qi, Y. Yu, and X. Tong, "Detail-preserving controllable deformation from sparse examples", *IEEE Transactions on Visualization and Computer Graphics*, vol. 18, no. 8, pp. 1215–1227, 2012.

[13] D. Brown, A. Macchietto, K. Yin, and V. Zordan, "Control of rotational dynamics for ground behaviors", in *ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pp. 55–61, 2013