

---

CONTACT INFORMATION	Mailing Address: 32 Vassar Street, 32-G836, Cambridge, MA, 02139 Webpage: <a href="http://people.csail.mit.edu/xchen">http://people.csail.mit.edu/xchen</a> ✉ <a href="mailto:cxh@mit.edu">cxh@mit.edu</a> ☎ 512-9887388
RESEARCH INTEREST	<b>Specialized AI systems</b> , i.e., domain-specific computer systems and architectures for AI & Big-data, with a focus on data mining and machine learning in graphs. My research spans across the full computer system stack, including parallel algorithms, runtime systems, compilers and hardware architectures.
ACADEMIC POSITIONS	<p><b>Research Scientist</b> <span style="float: right;"><b>Sep. 2020 - Present</b></span>          Computer Science &amp; Artificial Intelligence Laboratory (CSAIL),          Massachusetts Institute of Technology (MIT)  <i>Supervisor:</i> Professor <a href="#">Arvind</a>  <i>Research Area:</i> Efficient AI Systems and Architecture</p> <p><b>Research Fellow</b> <span style="float: right;"><b>Jan. 2019 - Aug. 2020</b></span>          Institute for Computational Engineering and Sciences,          University of Texas at Austin  <i>Supervisor:</i> Professor <a href="#">Keshav Pingali</a>  <i>Research Area:</i> Parallel Computing &amp; Graph Algorithms</p> <p><b>Assistant Research Professor</b> <span style="float: right;"><b>Jan. 2015 - May. 2018</b></span>          Department of Computer Science,          National University of Defense Technology (NUDT), China  <i>Research Area:</i> Parallel Architecture &amp; Parallel Computing</p> <p><b>Visiting Student</b> <span style="float: right;"><b>Oct. 2012 - Oct. 2014</b></span>          Department of Electrical and Computer Engineering,          University of Illinois at Urbana-Champaign (UIUC)  <i>Advisor:</i> Professor <a href="#">Wen-Mei Hwu</a>  <i>Research Area:</i> Massively Parallel Computer Architecture</p>
EDUCATION	<p><b>Ph.D. in Computer Science</b> <span style="float: right;"><b>Sep. 2009 - Dec. 2014</b></span>          National University of Defense Technology (NUDT), China  <i>Advisor:</i> Professor Zhiying Wang  <i>Thesis:</i> Cache Management for Manycore Accelerators</p> <p><b>B.S. in Computer Science</b> <span style="float: right;"><b>Sep. 2005 - Jun. 2009</b></span>          National University of Defense Technology (NUDT), China  <i>Rank:</i> 1/144</p>
PEER-REVIEWED PUBLICATIONS	<p>[1] <a href="#">Xuhao Chen</a>, Arvind, “Efficient and Scalable Graph Pattern Mining on GPUs”, The 16th USENIX Symposium on Operating Systems Design and Implementation (<b>OSDI</b>), 2022 <a href="#">[pdf]</a> <a href="#">[code]</a></p> <p>[2] <a href="#">Xuhao Chen</a>, Tianhao Huang, Shuotao Xu, Thomas Bourgeat, Chanwoo Chung, Arvind, “FlexMiner: A Pattern-Aware Accelerator for Graph Pattern Mining”, International Symposium on Computer Architecture (<b>ISCA</b>), 2021 <a href="#">[pdf]</a></p> <p>[3] <a href="#">Xuhao Chen</a>, Roshan Dathathri, Gurbinder Gill, Loc Hoang, Keshav Pingali, “Sand-slash: A Two-Level Framework for Efficient Graph Pattern Mining”, International Conference on Supercomputing (<b>ICS</b>), 2021 <a href="#">[pdf]</a></p>

- [4] Loc Hoang, Xuhao Chen, Hochan Lee, Roshan Dathathri, Gurbinder Gill, Keshav Pingali, “Efficient Distribution for Deep Learning on Large Graphs”, Workshop on Graph Neural Networks and Systems (**GNNSys**), 2021 [\[pdf\]](#) [\[code\]](#)
- [5] Xuhao Chen, Roshan Dathathri, Gurbinder Gill, Keshav Pingali, “Pangolin: An Efficient and Flexible Graph Pattern Mining System on CPU and GPU”, International Conference on Very Large Databases (**VLDB**), 13(8): 1190-1205, 2020 [\[pdf\]](#) [\[code\]](#)
- [6] Xuhao Chen, Li-Wen Chang, Christopher I. Rodrigues, Jie Lv, Zhiying Wang, Wen-Mei W. Hwu. “Adaptive Cache Management for Energy-efficient GPU Computing”, In the 47th International Symposium on Microarchitecture (**MICRO**), 2014 [\[pdf\]](#)
- [7] Loc Hoang\*, Vishwesh Jatala\*, Xuhao Chen, Udit Agarwal, Roshan Dathathri, Gurbinder Gill, Keshav Pingali, “DistTC: High Performance Distributed Triangle Counting”, IEEE High Performance Extreme Computing Conference (**HPEC**), 2019 [\[pdf\]](#)
- [8] Xuhao Chen, Cheng Chen, Jie Shen, Jianbin Fang, Tao Tang, Canqun Yang, Zhiying Wang, “Orchestrating Parallel Detection of Strongly Connected Components on GPUs”, Parallel Computing (**ParCo**), Volume 78, Pages 101-114, 2018
- [9] Xuhao Chen, Pingfan Li, Jianbin Fang, Tao Tang, Zhiying Wang, Canqun Yang, “Efficient and High-quality Sparse Graph Coloring on the GPU”, Concurrency and Computation: Practice and Experience (**CPE**), Volume 29, Issue 10, 2017
- [10] Xuhao Chen, Shengzhao Wu, Li-Wen Chang, Wei-Sheng Huang, Carl Pearson, Zhiying Wang, Wen-Mei W. Hwu. “Adaptive Cache Bypass and Insertion for Many-core Accelerators”, In Proceeding of the MES Workshop in conjunction with ISCA-41, 2014
- [11] Hang Zhang, Xuhao Chen, Nong Xiao, Fang Liu, “Optimizing STT-RAM Based Register File Energy Consumption on GPGPU with Delta Compression”, In Proceeding of the 53rd Design Automation Conference (**DAC**), 2016
- [12] Zhen Xu, Xuhao Chen, Jie Shen, Yang Zhang, Cheng Chen, Canqun Yang, “GARDENIA: A Domain-specific Benchmark Suite for Next-generation Accelerators”, ACM Journal on Emerging Technologies in Computing Systems (**JETC**), 15(1): 9, 2019
- [13] Hang Zhang, Xuhao Chen, Nong Xiao, Fang Liu, “Red-Shield: Shielding Read Disturbance for STT-RAM Based Register files on GPUs”, In Proceeding of the 26th Great Lakes Symposium on VLSI (**GLSVLSI**), 2016
- [14] Pingfan Li, Xuhao Chen, Jie Shen, Jianbin Fang, Tao Tang, Canqun Yang, “High Performance Detection of Strongly Connected Components in Sparse Graphs on GPUs”, In the Proceedings of the PMAM Workshop in conjunction with PPOPP-22, 2017
- [15] Hang Zhang, Xuhao Chen, Nong Xiao, Lei Wang, Fang Liu, Wei Chen, Zhiguang Chen, “Shielding STT-RAM Based Register files on GPUs Against Read Disturbance”, ACM Journal on Emerging Technologies in Computing Systems (**JETC**), 13(2): 27, 2016
- [16] Pingfan Li, Xuhao Chen, Zhe Quan, Jianbin Fang, Huayou Su, Tao Tang, Canqun Yang, “High Performance Parallel Graph Coloring on GPGPUs”, IPDPS Workshop, 2016

RESEARCH  
EXPERIENCE

**Hardware Accelerator for Graph Pattern Mining** **Sep. 2020 - Present**  
*With Arvind*

Design and implemented a pattern-aware hardware accelerator, FlexMiner [1], for graph pattern mining (GPM). FlexMiner offers 11× speedup over the state-of-the-art CPU baseline with 10× less logic area than CPU, while retaining high programming productivity.

**Programming Framework for Graph Pattern Mining** **Jan. 2019 - Aug. 2020**  
*With Keshav Pingali*

Design and implemented graph pattern mining (GPM) frameworks, Pangolin [4] and Sandslash [2]. Pangolin hides away parallel programming complexity, and is  $100 \times \sim 1000 \times$  faster than prior systems. Sandslash provides a novel two-level API that further simplifies programming, and enables optimizations that deliver  $8 \times$  speedup over Pangolin on CPU.

**Parallel Graph Algorithms on GPU**

**Jan. 2015 - Dec. 2018**

*With Zhiying Wang*

Design and implemented graph algorithms [6,7,8] on GPU, e.g., vertex coloring, strongly connected components, sparse neural networks. It pushes the frontier of GPU graph processing.

**Cache Architecture for Irregular Algorithms on GPU**

**Sep. 2009 - Dec. 2014**

*With Wen-Mei Hwu and Zhiying Wang*

Designed and implemented efficient cache architectures [5,9,10,11,12] for GPU. An *adaptive* cache management policy [5] is proposed for GPU, which effectively combines warp throttling and cache bypassing, to mitigate the memory divergence and cache contention on GPU.

SELECTED HONORS AND AWARDS	<ul style="list-style-type: none"> <li>• Graph Challenge 2019 Student Innovation Awards 2019</li> <li>• China Computer Federation (CCF) Distinguished PhD Dissertation Nominee 2015</li> <li>• Ci Yun-Gui Scholarship for Graduate, NUDT (top 1%) 2010</li> <li>• Meritorious Winner, Mathematical Contest In Modeling (MCM), <a href="#">COMAP</a> 2009</li> <li>• Distinguished Graduate, NUDT (top 1%), 2009</li> <li>• First rank, Scholarship of Excellent Achievements, NUDT (top 3%) 2009</li> <li>• Ci Yun-Gui Scholarship for Undergraduate, NUDT (top 1%) 2008</li> <li>• First-rank Prize, China Undergraduate Mathematical Contest in Modeling 2007</li> <li>• First rank, Scholarship of Excellent Achievements, NUDT (top 3%) 2007</li> </ul>
----------------------------------	---

TEACHING EXPERIENCE	<p><b>6.886: Algorithm Engineering</b> (graduate course) <b>Spring 2021, MIT</b></p> <ul style="list-style-type: none"> <li>- With Professor Julian Shun</li> <li>- Guest lecture on Pangolin [<a href="#">Slides</a>]</li> </ul> <p><b>CS 380C: Advanced Topics in Compilers</b> (graduate course) <b>Fall 2019, UT Austin</b></p> <ul style="list-style-type: none"> <li>- With Professor Keshav Pingali</li> <li>- Teaching Assistant to setup a <a href="#">course project</a> and mentor students on the project</li> </ul> <p><b>Design and Analysis of Algorithms</b> (undergraduate course) <b>Fall 2010 NUDT</b></p> <ul style="list-style-type: none"> <li>- With Professor Jianping Yin</li> <li>- Teaching Assistant to mentor students on labs and final projects and help with scoring</li> </ul> <p><b>Computer Architecture</b> (undergraduate course) <b>Fall 2008 NUDT</b></p> <ul style="list-style-type: none"> <li>- With Professor Zhiying Wang</li> <li>- Teaching Assistant to redesign the labs and the final project for a 5-stage pipelined in-order processor, and mentor students on course projects</li> </ul>
------------------------	--

MENTORING EXPERIENCE	<ol style="list-style-type: none"> <li>1. Tianhao Huang (third year PhD with Prof. Arvind). Closely collaborate on project: Hardware Accelerator for Graph Pattern Mining</li> <li>2. Loc Hoang (second year PhD with Prof. Keshav Pingali). Closely collaborate on project: Distributed Framework for Graph Neural Networks</li> <li>3. Siyu Zhang &amp; Bozhi You (first year PhD with Prof. Keshav Pingali). Mentoring course project: Parallel and Distributed <math>k</math>-clique Listing on Large Graphs</li> <li>4. Jie Lv (master student with Prof. Wen-Mei Hwu). Closely collaborate on project: Adaptive Cache Management for GPU</li> <li>5. Zhihao Xie (undergraduate student with Prof. Wen-Mei Hwu).</li> </ol>
-------------------------	--

Mentoring course project: Characterization of Irregular Algorithms on GPU

#### GRANTS

- Samsung Memory Solutions Lab. (MSL) UR Collaboration Program (2022), PI: Arvind  
“In-storage Acceleration of Massive Graph Neural Networks”  
Co-led the project and wrote successful grant proposal
- MIT-IBM Watson AI Lab Research Funding (2021), PI: Arvind  
“Storage-Centric Infrastructure for Graph-Based Deep Learning on Massive Datasets”  
Co-led the project and wrote successful grant proposal
- NSF of China Grant No.61502514 (2015), PI: Xuhao Chen  
“Memory Hierarchy for Energy-efficient Heterogeneous Processors”  
Wrote successful grant proposal, led and completed the project

#### TALKS

1. Invited talk at George Mason University, March 11, 2022
2. Invited talk at Clemton University, March 1, 2022
3. Invited talk at University of Rochester, February 23, 2022
4. Invited talk at University of Florida, February 17, 2022
5. Invited talk at Rochester Institute of Technology, February 4, 2022
6. Invited talk at Iowa State, January 24, 2022
7. Invited talk at Illinois Institute of Technology, January 18, 2022
8. Invited Talk at the CS Seminar in UMass Amherst, November 5, 2021
9. Invited Talk at the [MiDAS](#) Seminar in Boston University, October 29, 2021
10. Invited Talk at the [COSI](#) Seminar in Brandeis University, October 22, 2021
11. Invited Talk at the [MIT Fast Code Seminar](#), April 28, 2021 [[Video](#)]
12. Invited talk on [Pangolin](#) at the [VLDB conference](#), September 3, 2020 [[Video](#)]

#### ACADEMIC SERVICE

- Invited reviewer for IEEE Transactions on Knowledge and Data Engineering
- Invited reviewer for ACM Transactions on Architecture and Code Optimization
- Invited reviewer for ACM Transactions on Modeling and Performance Evaluation of Computing Systems
- Invited reviewer for Microprocessors and Microsystems: Embedded Hardware Design
- Invited reviewer for Journal of Supercomputing

#### REFERENCES

Arvind  
Johnson Professor  
CSAIL at MIT  
Postdoc Supervisor  
[arvind@csail.mit.edu](mailto:arvind@csail.mit.edu)

Keshav Pingali  
Professor  
CS at UT Austin  
Postdoc Supervisor  
[pingali@cs.utexas.edu](mailto:pingali@cs.utexas.edu)

Wen-Mei Hwu  
Professor and Sanders-AMD Chair  
ECE at UIUC  
Supervisor  
[w-hwu@illinois.edu](mailto:w-hwu@illinois.edu)

Zhiying Wang  
Professor  
CS at NUDT  
PhD Thesis Advisor  
[zywang@nudt.edu.cn](mailto:zywang@nudt.edu.cn)