

Hao Wang

AFFILIATION	Division of Computer Science and Engineering Louisiana State University 3272X Patrick F. Taylor Hall Baton Rouge, Louisiana 70803, USA	<i>Email:</i> haowang@lsu.edu <i>Web:</i> http://www.haow.ca/ <i>Mobile:</i> +1 (225) 288-7471
PERSONAL INFORMATION	<i>Citizenship</i> Chinese	
EDUCATION	University of Toronto , Toronto, Ontario, Canada <i>Department of Electrical and Computer Engineering</i> ◇ Ph.D. , Electrical and Computer Engineering, Sept. 2015 – July. 2020 ▷ <i>Dissertation:</i> “Optimizing Distributed Computing Systems via Machine Learning” ▷ <i>Advisor:</i> Prof. Baochun Li, Department of Electrical and Computer Engineering Shanghai Jiao Tong University , Shanghai, People’s Republic of China <i>School of Electronic Information and Electrical Engineering</i> ◇ M.Engr. , Software Engineering, Mar. 2015 ▷ <i>Dissertation:</i> “Generic and Accurate Traffic Prediction for Data-Parallel Clusters” ▷ <i>Advisor:</i> Prof. Haibing Guan, Department of Computer Science Prof. Zhengwei Qi, School of Software Engineering ◇ B.Engr. , Information Security, Jul. 2012 ▷ <i>Dissertation:</i> “Performance Evaluation and Tuning of Cloud Computing”	
RESEARCH INTERESTS	Large-scale data analytics, machine learning systems, distributed computing, cloud computing, datacenter networking.	
RESEARCH AND INDUSTRY EXPERIENCE	Assistant Professor , Computer Science Engineering Division <i>Louisiana State University A&M</i> , Baton Rouge, LA, USA Research Assistant , supervised by Prof. Baochun Li <i>University of Toronto</i> , Toronto, Ontario, Canada Research Intern , supervised by Dr. Yongqiang Xiong <i>Microsoft Research Asia</i> , Beijing, P. R. China Visiting Research Assistant , supervised by Prof. Kai Chen <i>Hong Kong University of Science and Technology</i> , Hong Kong, P. R. China Co-founder and Full-Stack Developer <i>Ramy Tech Inc.</i> , Shanghai, P. R. China Development Intern on Cloud performance tuning <i>Intel Asia-Pacific Research and Development Ltd.</i> , Shanghai, P. R. China	Jan. 2021 – present Sept. 2015 – Sept. 2020 Apr. 2015 – Aug. 2015 Nov. 2013 – Dec. 2014 Jul. 2013 – Nov. 2013 Dec. 2011 – Dec. 2012

◇ **Journal Articles** (in reverse chronological order)

- [J1] Hanxi Guo, **Hao Wang**, Yang Hua, Tao Song, Ruhui Ma, Haibing Guan. “SpaceDML: Enabling Distributed Machine Learning in Space Information Networks,” in *IEEE Network*, 2021.
- [J2] **Hao Wang**, Di Niu, Baochun Li. “Turbo: Dynamic and Decentralized Global Analytics via Machine Learning,” in *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 2020.
- [J3] **Hao Wang**, Baochun Li. “Mitigating Bottlenecks in Wide Area Data Analytics via Machine Learning,” in *IEEE Transactions on Network Science and Engineering (TNSE)*, 2018.
- [J4] Wei Bai, Li Chen, Kai Chen, Dongsu Han, Chen Tian, **Hao Wang**. “Information-Agnostic Flow Scheduling for Commodity Data Centers,” in *IEEE/ACM Transactions on Networking (ToN)*, 2017.
- [J5] Hong Zhang, Kai Chen, Wei Bai, Dongsu Han, Chen Tian, **Hao Wang**, Haibing Guan, Ming Zhang. “Guaranteeing Deadlines for Inter-Datacenter Transfers,” in *IEEE/ACM Transactions on Networking (ToN)*, 2016
- [J6] Shuihai Hu, Kai Chen, Haitao Wu, Wei Bai, Chang Lan, **Hao Wang**, Hongze Zhao, Chuanxiong Guo. “Explicit Path Control in Commodity Data Centers: Design and Applications,” in *IEEE/ACM Transactions on Networking (ToN)*, 2015
- [J7] Yang Peng, Kai Chen, Guohui Wang, Wei Bai, Yangming Zhao, **Hao Wang**, Yanhui Geng, Zhiqiang Ma, Lin Gu. “Towards Comprehensive Traffic Forecasting in Cloud Computing: Design and Application,” in *IEEE/ACM Transactions on Networking (ToN)*, 2015.

◇ **Conference Papers** (in reverse chronological order)

- [C1] Gang Yan, **Hao Wang**, and Jian Li. “Seizing Critical Learning Periods in Federated Learning,” in the Proceedings of the AAAI Conference on Artificial Intelligence (AAAI), Vancouver, BC, Canada, February 22 - March 1, 2022.
- [C2] Jiaru Zhang, Yang Hua, Tao Song, **Hao Wang**, Zhengui Xue, Ruhui Ma, and Haibing Guan. “Improving Bayesian Neural Networks by Adversarial Sampling,” in the Proceedings of the AAAI Conference on Artificial Intelligence (AAAI), Vancouver, BC, Canada, February 22 - March 1, 2022.
- [C3] Hanxi Guo, **Hao Wang**, Tao Song, Yang Hua, Zhangcheng Lv, Xiulang Jin, Zhengui Xue, Ruhui Ma, and Haibing Guan. “Siren: Byzantine-robust Federated Learning via Proactive Alarming,” in the Proceedings of the *the ACM Symposium on Cloud Computing (SoCC)*, Seattle, WA, November 1 - 4, 2021. (acceptance rate: 31.7%)
- [C4] Hanfei Yu, Athirai A. Irissappane, **Hao Wang**, Wes J. Lloyd. “FaaSRank: Learning to Schedule Functions in Serverless Platforms,” in the Proceedings of the *IEEE International Conference on Autonomic Computing and Self-Organizing Systems (ACSOS 2021)*, Washington DC, September 27 - October 1, 2021.
- [C5] Yuxin Ma, Yang Hua, Hanming Deng, Tao Song, **Hao Wang**, Zhengui Xue, Heng Cao, Ruhui Ma, Haibing Guan. “Self-Supervised Vessel Segmentation via Adversarial Learning,” in the Proceedings of the *IEEE International Conference on Computer Vision (ICCV)*, Virtual conference, October 11 - 17, 2021.
- [C6] Zinuo Cai, Jianyong Yuan, Yang Hua, Tao Song, **Hao Wang**, Zhengui Xue, Ningxin Hu, Jonathan Ding, Ruhui Ma, Mohammad Reza Haghghat, Haibing Guan. “Themis: A Fair

Evaluation Platform for Computer Vision Competitions,” in the Proceedings of the *30th International Joint Conference on Artificial Intelligence (IJCAI 2021)*, Montreal, Canada, August 21 - 26, 2021. (acceptance rate: 13.9%)

- [C7] **Hao Wang**, Zakhary Kaplan, Di Niu, Baochun Li. “Optimizing Federated Learning on Non-IID Data with Reinforcement Learning,” in the Proceedings of IEEE INFOCOM 2020, Beijing, China, April 27 - 30, 2020. (acceptance ratio: 20%).
- [C8] **Hao Wang**, Di Niu, Baochun Li. “Distributed Machine Learning with a Serverless Architecture,” in the Proceedings of IEEE INFOCOM 2019, Paris, France, April 29 - May 2, 2019. (acceptance ratio: 20%).
- [C9] **Hao Wang**, Di Niu, Baochun Li. “Dynamic and Decentralized Global Analytics via Machine Learning,” in the Proceedings of the *ACM Symposium on Cloud Computing 2018 (SoCC 2018)*, Carlsbad, California, October 11 - 13, 2018 (acceptance ratio: 24%).
- [C10] **Hao Wang**, Baochun Li. “Lube: Mitigating Bottlenecks in Wide Area Data Analytics,” in the Proceedings of the *9th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 2017)*, Santa Clara, California, July 10 - 11, 2017 (acceptance ratio: 32%).
- [C11] Shuhao Liu, **Hao Wang**, Baochun Li. “Optimizing Shuffle in Wide-Area Data Analytics,” in the Proceedings of the 37th International Conference on Distributed Computing Systems (ICDCS 2017), Atlanta, Georgia, June 5 - 8, 2017 (acceptance ratio: 17%).
- [C12] **Hao Wang**, Li Chen, Kai Chen, Ziyang Li, Yiming Zhang, Haibing Guan, Zhengwei Qi, Dongsheng Li, Yanhui Geng. “FlowProphet: Generic and Accurate Traffic Prediction for Data-parallel Cluster Computing,” in the Proceedings of the 35th *International Conference on Distributed Computing Systems (ICDCS 2015)*, Columbus, Ohio, June 29 - July 2, 2015 (acceptance ratio: 13%).
- [C13] Hong Zhang, Kai Chen, Wei Bai, Dongsu Han, Chen Tian, **Hao Wang**, Haibing Guan, Ming Zhang. “Guaranteeing Deadlines for Inter-Datacenter Transfers,” in the Proceedings of the *ACM European Conference on Computer Systems (EuroSys 2015)*, Bordeaux, France, April 21 - 24, 2015 (acceptance ratio: 21%).
- [C14] Wei Bai, Li Chen, Kai Chen, Dongsu Han, Chen Tian, **Hao Wang**. “Practical Information-Agnostic Flow Scheduling for Data Center Networks,” in the Proceedings of the *12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 2015)*, Oakland, CA, May 4 - 6, 2015 (acceptance ratio: 20%).
- [C15] Shuihai Hu, Kai Chen, Haitao Wu, Wei Bai, Chang Lan, **Hao Wang**, Hongze Zhao, Chuanxiong Guo. “Explicit Path Control in Commodity Data Centers: Design and Applications,” in the Proceedings of the *12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 2015)*, Oakland, CA, May 4 - 6, 2015 (acceptance ratio: 20%).
- [C16] **Hao Wang**, Yangming Zhao, Haibing Guan. “On Pricing Schemes in Data Center Network with Game Theoretic Approach,” in the Proceedings of the *23rd IEEE International Conference on Computer Communications and Networks (ICCCN 2014)*, Shanghai, China, August 4 - 7, 2014 (acceptance ratio: 28%).

◇ **Conference Posters** (in reverse chronological order)

- [P1] **Hao Wang**, Di Niu, Baochun Li. “Dynamic and Decentralized Global Analytics via Machine Learning,” the *ACM Symposium on Cloud Computing 2018 (SoCC 2018)*, Carlsbad, California, October 11 - 13, 2018

[P2] **Hao Wang**, Baochun Li. “Lube: Mitigating Bottlenecks in Wide Area Data Analytics,” the *9th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 2017)*, Santa Clara, California, July 10 - 11, 2017

[P3] **Hao Wang**, Baochun Li. “Bottleneck Detection for Wide Area Data Analytics on the SAVI Testbed,” the *3th SAVI Annual General Meeting*, Toronto, Ontario, July 6, 2016

SCHOLARLY TALKS	<p>Optimizing Distributed Machine Learning with Reinforcement Learning <i>Rensselaer Polytechnic Institute, Troy, US</i> Mar. 2020</p> <p>Optimizing Distributed Machine Learning with Reinforcement Learning <i>University of Edinburgh, Edinburgh, UK</i> Apr. 2020</p> <p>Optimizing Big Data with Machine Learning <i>Hunan University, Changsha, China</i> Dec. 2019</p> <p>Distributed Machine Learning with a Serverless Architecture <i>IEEE INFOCOM, Paris, France</i> Apr. 2019</p> <p>Optimizing Large Scale Data Analytics via Machine Learning <i>Shanghai Jiao Tong University, Shanghai, China</i> Jan. 2019</p> <p>Dynamic and Decentralized Global Analytics via Machine Learning <i>ACM SoCC, Carlsbad, CA</i> Oct. 2018</p> <p>Mitigating Bottlenecks in Wide Area Data Analytics <i>ACM HotCloud, Santa Clara, CA</i> Jul. 2017</p> <p>Generic and Accurate Traffic Prediction for Data-parallel Cluster Computing <i>IEEE ICDCS, Columbus, OH</i> Oct. 2015</p> <p>On Pricing Schemes in Data Center Network with Game Theoretic Approach <i>IEEE ICCCN, Shanghai, China</i> Oct. 2014</p>
PROFESSIONAL SERVICES	<ul style="list-style-type: none"> ◇ Web Chair: IEEE ICNP 2017 ◇ Technical Committee Member: IEEE INFOCOM 2021, IEEE INFOCOM 2022, IEEE IRI 2020, IEEE 2021, EAI ADHOCNETS 2021 ◇ Reviewer for Journal Manuscript Submissions: IEEE Transactions on Networking, IEEE Transactions on Parallel and Distributed Systems, IEEE Transactions on Network Science and Engineering, IEEE Transactions on Big Data, Multimedia Systems Springer Journals, IEEE Access ◇ Reviewer for Conference Manuscript Submissions: ACM SIGCOMM poster, USENIX HotCloud, ACM Multimedia, ACM MMSys, ACM NOSSDAV, IFIP Networking, IEEE INFOCOM, IEEE IWQoS, IEEE GLOBECOM, IEEE IC2E, WiOpt, ACM/IEEE IoTDI
HONOURS AND AWARDS	<ul style="list-style-type: none"> ◇ DiDi Graduate Awards (\$10,000 CAD), Didi Chuxing Technology Co. 2019 ◇ Best In-session Presentation Award, INFOCOM'19, IEEE 2019 ◇ Doctoral Completion Award (\$18,000 CAD), ECE Department, University of Toronto 2019 ◇ SGS Conference Grant (\$1,000 CAD), School of Graduate Studies, University of Toronto 2018 ◇ Student Scholarships (\$900 USD), SoCC'18, ACM 2018 ◇ University of Toronto Fellowship, Department of ECE, University of Toronto 2015 ◇ Edward S. Rogers Sr. Graduate Scholarships, University of Toronto 2015

- ◇ **Yunfeng Prize** (Top 1%), the 2nd Aliyun Worldwide Developer Conference, Alibaba Inc. 2013
- ◇ **The Second-Class Postgraduate Scholarship**, Shanghai Jiao Tong University 2012
- ◇ **The 3rd Prize**, the 4th National College Information Security Contest 2011

GRANT WRITING
EXPERIENCE

Worked as a main technical contributor in the following grant applications.

- ▷ **NSERC Strategic Grant** 2016
“Online Monitoring and Performance Optimization on Data Analytic Applications across Geo-Distributed Datacenters”, PI: Baochun Li; Co-Applicants: Ding Yuan, Micheal Stumm
- ▷ **NSERC Collaborative Research and Development (CRD) Grant** 2016
“Performance Optimization for Multi-Datacenter Cloud Platforms”, PI: Baochun Li
- ▷ **NSERC Collaborative Research and Development (CRD) Grant** 2017
“Optimizing Geo-Distributed Data Analytics Across Multiple Datacenters”, PI: Baochun Li
- ▷ **NSERC Collaborative Research and Development (CRD) Grant** 2019
“Towards a Quantitative Understanding of Short Texts with Deep Learning”, PI: Baochun Li
- ▷ **Research Gift from Orbis Investment Management Inc.** 2019
“Sentiment Analysis and Topic Modeling on Short Texts”, PI: Baochun Li
- ▷ **Compute Canada Resources for Research Groups (RRG) Competition** 2019
“Performance Optimization for Large Scale Data Analytics”, PI: Baochun Li

TEACHING
EXPERIENCE

- Teaching Assistant** Sept. 2015 – Sept. 2020
University of Toronto, Toronto
- ▷ ECE 444: Software Engineering Fall 2017 – 2019
 - ▷ CSC 454: Computer Systems Programming Fall 2018
 - ▷ CSC 369: Operating Systems Summer 2017, Fall 2019
 - ▷ ECE 344: Operating Systems Fall 2016
 - ▷ CSC 343: Introduction to Databases Summer 2016
 - ▷ ECE 353: Systems Software Winter 2016 – 2019
 - ▷ CSC 458: Computer Networks Fall 2015 – 2016

MENTORING

- Zhixuan Wang**, female M.Eng. student at *University of Toronto* Jun. 2020 – Aug. 2020
Thesis title: “Optimizing Federated Learning with Lottery Tickets”
- Maliha Islam**, female undergraduate student at *University of Toronto* Sept. 2019 – Apr. 2020
Thesis title: “Optimizing Parallelism in Federated Learning”
- Jeffrey Nguyen**, undergraduate student at *University of Toronto* Sept. 2019 – Apr. 2020
Thesis title: “Masking the Topic: a Topic Model for Short Texts”
- Zakhary Kaplan**, undergraduate student at *University of Toronto* Jun. 2019 – Aug. 2019
Research topic: “Optimizing Federated Learning on Non-IID Data with Reinforcement Learning”
- Haobo Ding**, undergraduate student at *University of Toronto* Sept. 2018 – Apr. 2019
Thesis title: “Speedup Straggler Workers in Distributed Machine Learning”
- Yudian Shi**, female undergraduate student at *University of Toronto* Sept. 2018 – Apr. 2019
Thesis title: “Speeding Up Distributed Machine Learning with a Serverless Architecture”
- Zhongyang Xiao**, undergraduate student at *University of Toronto* Sept. 2016 – Apr. 2017
Thesis title: “Identifying Runtime Performance Bottlenecks of Spark Task Scheduler”
- Luyuan Chen**, undergraduate student at *University of Toronto* Sept. 2016 – Apr. 2017

Thesis title: “Optimizing SparkSQL using Machine Learning”

Shing-Chun Tse and Kai-Chung Law, undergraduate students at *Hong Kong University of Science and Technology* Dec. 2013 – Dec. 2014

Research topic: “Studying and Building a Software-Defined Network”

REFERENCES

Baochun Li, IEEE Fellow, Professor

Bell Canada Endowed Chair in Computer Engineering

Department of Electrical and Computer Engineering

University of Toronto, 10 Kings College Rd., Toronto, Ontario M5S 3G4, Canada

Phone: +1-416-946-7338

Email: bli@ece.toronto.edu

Web: <http://iqua.ece.toronto.edu/bli/>

Ben Liang, IEEE Fellow, Professor

Department of Electrical and Computer Engineering

University of Toronto, 10 Kings College Rd., Toronto, Ontario M5S 3G4, Canada

Phone: +1-416-946-8614

Email: liang@ece.utoronto.ca

Web: <https://www.comm.utoronto.ca/~liang/>

Di Niu, IEEE Member, Associate Professor

Department of Electrical and Computer Engineering

University of Alberta, 116 St. and 85 Ave., Edmonton, Alberta T6G 2R3, Canada

Phone: +1-780-616-1022

Email: dniu@ualberta.ca

Web: <https://sites.ualberta.ca/~dniu>

Kai Chen, Associate Professor

Department of Computer Science and Engineering

Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong, China

Phone: +852-2358-7028

Email: kaichen@cse.ust.hk

Web: <http://www.cse.ust.hk/~kaichen>

Alberto Leon-Garcia, Fellow of Engineering Institute of Canada, IEEE Fellow, Professor

Department of Electrical and Computer Engineering

University of Toronto, 10 Kings College Rd., Toronto, Ontario M5S 3G4, Canada

Phone: +1-416-978-4764

Email: alberto.leongarcia@utoronto.ca

Web: <https://www.nal.utoronto.ca>