
Bingsheng He

Ph.D., HKUST, 2008

Email: hebs@comp.nus.edu.sg

Professor, Vice Dean (Research)

School of Computing, National University of Singapore

Address: #COM3-02-12, COM3 Building, 11 Research Link, NUS, Singapore 119391

Page: <http://www.comp.nus.edu.sg/~hebs/>

Contents

I. EDUCATION AND EMPLOYMENT	2
A. Education	2
B. Employment	2
II. TEACHING	3
A. Courses Taught	3
B. Curriculum and Course Development	3
C. Postdoctoral Fellows Supervised	4
D. Ph.D Students Supervised	5
E. Master Students Supervised	6
F. Undergraduate Students Supervised	6
G. Ph.D./Master Committee	6
H. Visiting Students/Interns	6
I. Teaching/Advising Honors and Awards	6
III. RESEARCH	8
A. Ph.D. Thesis	8
B. Research Group and Mission	8
C. Citations	8
D. Publications	9
D.1. Five Selected Articles	9
D.2. Journal Articles	9
D.3. Conference and Workshop Papers	15
D.4. System Demonstration Paper	26
D.5. Posters	27
D.6. Tutorials	28
D.7. Book Chapters	28
D.8. Interdisciplinary Research Publications	28
E. Open-Source Systems	29
F. Funds and Projects	30
G. Patents and Technical Disclosures	32
IV. SERVICES	32
A. Professional Activities	32

A.1. Conference/Workshop Organizer	32
A.2. Steering Committee	33
A.3. Editorial Board	33
A.4. Program Committee Members	34
A.5. External/Invited Reviewers	34
A.6. External/Invited Academic Services	35
B. University Services at National University of Singapore	36
C. University Services at Nanyang Technological University	36
V. HONORS AND AWARDS	38
VI. Keynote and Presentations	40
A. Invited Plenary/Keynote Talks	40
B. Other Presentations	40
VII. References	42

I. EDUCATION AND EMPLOYMENT

A. Education

Degree	Period	University
Ph. D. of Computer Science	08/2003 — 07/2008	The Hong Kong University of Science & Technology (HKUST)
B.E. of Computer Science and Engineering	09/1999 — 08/2003	Shanghai Jiao Tong University, China
B.B.A. of Business Administration (second major)		

B. Employment

Position	Period	Affiliation
Professor	07/2022—present	School of Computing, National University of Singapore (NUS)
Associate Professor/Dean's Chair Associate Professor	05/2016—06/2022	School of Computing, National University of Singapore (NUS)
Assistant/Associate Professor (Tenured since 09/2015)	08/2010— 05/2016	School of Computer Science and Engineering, Nanyang Technological University (NTU)
Visiting Scholar	02/2010— 07/2010	The Chinese University of Hong Kong (CUHK)
Associate Researcher	09/2008 — 01/2010	System Research Group, Microsoft Research at Asia (MSRA)

II. TEACHING

A. Courses Taught

Course	Period	Typical class sizes
CS5425/CS4225 Big Data Systems for Data Science	Fall 2018 – Present	100-200
CS6284 Topics in Computer Science: Big Data Meets New Hardware ¹	Fall 2019	30+
CG2271 Real-Time Operating Systems	Fall 2015 – Fall 2018	70+
CS4225 Massive Data Processing Techniques for Data Science	Spring 2016 – Fall 2017	40+
CSC/CPE205, CZ/CE2005 Operating Systems	Fall 2010—Spring 2015	100-300
CSC/CPE203 Software Systems and Models	Fall 2010, Spring/Fall 2011	40+
CE/CZ 0002 Green Computing	Spring 2012— Spring 2015	200

¹ <https://www.comp.nus.edu.sg/~hebs/course/cs6284/>

Recent Student Teaching Feedback

Sem.	Course	Class Size	#Responses	#Nominations for teaching awards	Teaching Rating (/5)	Department Average (/5)
20S1	CS5425	88	39	20	4.3	4.2
20S1	CS4225	62	28	18	4.6	4.2
19S2	CS5425	41	21	13	4.5	4.1
19S2	CS4225	64	25	4	4.1	4.1
19S1	CS6284	17	9	5	4.8	4.2
18S2	CS5425	17	8	5	4.6	4.1
18S2	CS4225	52	19	4	4.2	4.1
18S1	CG2271	76	58	8	4.2	4.2
17S2	CS4225	32	18	13	4.7	4.1
17S1	CG2271	67	48	7	4.1	4.2

Full course title:

CG2271 Real-Time Operating Systems

CS5425/CS4225 Big Data Systems for Data Science

CS6284 Topics in Computer Science: Big Data Meets New Hardware

B. Curriculum and Course Development

- Co-Develop the curriculum for internet of things (IoT) specialization for NUS CEG (computer engineering program), June 2018.
- New course material development:
 - CS6284 Topics in Computer Science: Big Data Meets New Hardware: new lecture contents
 - CS5425/CS4225 Big Data Systems for Data Science: new labs and lecture contents
 - CG2271 Real-Time Operating Systems: new labs and lecture contents

- CS4225 Massive Data Processing Techniques for Data Science: new labs and lecture contents
- Course materials CZ/CE2005 “Operating Systems”, including lecture slides (12 hours), lab materials (8 hours) and tutorials (11 hours).
- Course materials CZ/CE4011 “Parallel Processing” (on the part of GPU programming), including lecture slides (4 hours), lab materials (4 hours) and tutorials (4 hours).
- Course materials CZ/CE0002 “Green Computing”, including lecture slides (6 hours) and course project proposals for students.
- Tutorial slides for CSC/CPE 203 SOFTWARE SYSTEM and MODELS (10 hours)
- Course syllabus review and development
 - CE2005/CZ2005 Operating system
 - CE0002/CZ0002 Green Computing
 - CE4012/CZ4012 Cloud Computing & its Applications
 - CE4014/CZ4014 Programming Massively Parallel Processors.
- My course materials are requested/used by other colleagues
 - CS4225 Massive Data Processing Techniques for Data Science/Big Data Systems for Data Science: Anh Dinh (SUTD), Hong Gao (HIT), Fangming Liu (HUST), Xike Xie (USTC), Zichen Xu (NCU, China)

C. Postdoctoral Fellows Supervised

- 1) Jiaxin Jiang (Ph.D., Hong Kong Baptist University), Topic: Real-time dynamic graph processing. Period: 2021- Present.
- 2) Zhen Zhang (Ph.D., Zhejiang University, China), Topic: Real-time fraud detection with graph analytics. Period: 2021 – Present.
- 3) Shixuan Sun (Ph.D., Hong Kong University of Science and Technology), Topic: high performance graph processing systems. Period: 2020- present.
- 4) Ronak Bajaj (Ph.D., Nanyang Technological University), Topic: Systems and tools on FPGA, Contribution: FPL (1), CIDR (1), Period: 2018-2019
- 5) Zeyi Wen (Ph.D., University of Melbourne), Topic: Machine learning systems on GPUs, Contributions: system ThunderSVM [<https://github.com/Xtra-Computing/thundersvm>, 780 stars] + IEEE ICDE (1), IPDPS (1), JMLR (1), TKDE (1). Period: 2017-2019. Appointment: Lecturer at The University of Western Australia (UWA).
- 6) Cheng Liu (Ph.D., University of Hong Kong), Topic: Graph processing on FPGA, Contribution: system prototypes + IEEE ICCAD (1), Period: 2017-2018. Appointment: Associate Professor at ICT.
- 7) Kai Zhang (Ph.D., University of Science and Technology China), Topic: Data management on the GPU. Contribution: system prototypes + VLDBJ (1), IEEE ICDE (1), USENIX NSDI (1). Period: 2016-2017. Appointment: Associate Professor at Fudan University.
- 8) Zeke Wang (Ph.D., Zhejiang University, China), Topic: Systems and tools on FPGA, Contribution: system prototypes + IEEE ICDE (1), USENIX NSDI (1), FPL (2), HPCA (1), IEEE TPDS (1), IEEE TVLSI (2), Period: 2013-2016. Appointment: research staff at ETHZ.
- 9) Shanjiang Tang (Ph.D., Nanyang Technological University), Topic: Resource management in the cloud. Contribution: system prototypes + ACM/IEEE SC (1), IEEE TSC (1), IEEE TCC (1), IEEE CloudCom (1), Period: 2014-2015. Appointment: Associate Professor at Tianjin University.
- 10) Haikun Liu (Ph.D., Huazhong University of Science and Technology), Topic: Resource management in the cloud. Contribution: system prototypes + IEEE TPDS (3), ACM/IEEE SC (1). Period: 2013-2014. Appointment: Associate Professor at Huazhong University of Science and Technology, China.
- 11) Cheng Liu (Ph.D., Nanyang Technological University), Topic: Job scheduling in the cloud. Contribution: complete a system prototype for water management in the cloud, and deployment in Singapore government agency. Period: 2011-2014. Appointment: Data Scientist at GIC Singapore.

D. Ph.D Students Supervised

- 1) Yuan Li (August 2021 – Present): Graph neural networks for spatial-temporal graphs
- 2) Naibo Wang (August 2021 – Present): Federated model hub
- 3) Yuhang Chen (August 2021- present): High-performance graph analytics systems
- 4) Zining Zhang (co-supervised with Yitu Singapore, August 2018 to present): Systems for machine learning.
- 5) Zhaomin Wu (August 2019 to present): Vertical federated learning systems.
- 6) Sixu Hu (August 2019 to present): Federated learning system benchmarking.
- 7) Johan Kok (co-supervised with Grab Singapore, August 2019 to present): Learnt Optimizations for Big Data systems.
- 8) Naibo Wang (August 2021- present, funded by NUS Institute of Data Science): Federated machine learning systems.
- 9) Shengliang Lu (August 2017 to present): Graph processing on GPUs.
- 10) Xinyu Chen (August 2017 to present): System optimization on FPGAs.
- 11) Qinbin Li (August 2018 to present): Tree-based federated learning systems.
- 12) Johns Paul (co-advised with Dr. Chiew Tong Law, August 2015 to Feb 2021): In-Memory Analytical Query Processing on GPUs. Current Appointment: Senior Program Analyst (AVP) at CitiBank.
- 13) Tony Shuhao Zhang (co-advised with Dr. Daniel Dahlmeier (SAP), SAP Ph.D. program, August 2014 to Sept 2019): Complex Event Processing on multi-cores. Current Appointment: Assistant Professor at SUTD.
- 14) Xuntao Cheng (co-advisor with Dr. Lau Chiew Tong, August 2013 to April 2018): A high-performance main-memory query engine on emerging many-core processors. Current Appointment: Senior software engineer, Alibaba.
- 15) Tao Luo (co-advisor with Dr. Douglas Maskell, Nov 2013 to Nov 2017, original advisor: Dr. Wei Zhang (now HKUST)): Racetrack Memory Based Logic Design for In-memory Computing, Current Appointment: Research scientist at A*STAR Singapore.
- 16) Zhaojie Niu (co-advised with Dr. Ming Jian, Jan 2013 to Oct 2016): A Study of Big Data Computing Platforms: Performance, Fairness and Energy Consumption. Current Appointment: Senior Software Engineer, ByteDance.
- 17) Yifan Gong (co-advised with Dr. Adrian Wing-Keung Law, August 2012 to August 2016): Network performance aware HPC systems in the cloud. Current Appointment: Research scientist at TuSimple.
- 18) Fubing Mao (Nov 2013 to August 2016, original advisor: Dr. Wei Zhang (now HKUST)): Design Automation Flow for Partial Run-time Reconfiguration on FPGAs. First Appointment: Assistant professor at Huazhong University of Science and Technology, China.
- 19) Yingnan Cui (Nov 2013 to 2016, original advisor: Dr. Wei Zhang (now HKUST)): Thermal Management for Large-scale Multi-core Processors. Current Appointment: Research scientist at A*STAR Singapore.
- 20) Lei Wei (co-advisor with Dr. Jianfei Cai, original advisor: Dr. Foh Chuan Heng, August 2010 to Feb 2016): Economical and QoS-aware Resource Allocation for Cloud Computing. Current Appointment: Research Engineer at Huawei.
- 21) Amelie Chi Zhou (August 2011 to Jan 2016): Simplified and Effective Resource Provisioning for Scientific Workflows in IaaS Clouds. Current Appointment: Assistant Professor, Shenzhen University.
- 22) Jiong He (August 2011 to August 2015): Real-time data analytics for coupled CPU-GPU architectures. Current Appointment: Senior research engineer at ByteDance.
- 23) Jianlong Zhong (Jan 2011 to June 2014): Parallel Graph Processing on Graphics Processors. Current Appointment: A Startup Company in USA (with O1 Visa)→Software Engineer Facebook.
- 24) Li Pei Wong (July 2011 to July 2012, original advisor: Dr. Malcolm Low): A Generic Bee Colony Optimization Framework for Combinatorial Optimization Problems. Current Appointment:

Senior Lecturer at Universiti Sains Malaysia, Penang, Malaysia.

E. Master Students Supervised

- 1) 2021: Yung Ching Man, Kelly (→ Sony)
- 2) 2020: Yuxuan Han (→Bytedance), Satija Parth (→ M-DAQ Global).
- 3) 2019: Xu Liu (→ NUS Ph.D.), Yuan Li (→ NUS Ph.D.)
- 4) 2018: Shiheng Chen
- 5) 2017: Chengxi Xue (→ Shopee)
- 6) 2016: Jun Zhang (→ Facebook)

F. Undergraduate Students Supervised

- 1) 2021: Jiajun Liu, See Zi Yang, Chen Hui, Bingqiao Luo, Qian Wang, Jiawei Zhang
- 2) 2020: Yucheng Zhang, Teo Jun Hui, Tianyuan Fang, Jieliang Ang
- 3) 2019: Husong Liu, Kevin Leonardo Anantha

G. Ph.D./Master Committee

- 1) Recent Ph.D./MComp Thesis/Proposal evaluation committee:
 - 2021: Yueji Yang, Qingyuan Wang, Zerui Ge, Ziyue Chen
 - 2020: Himeshi De Silva, Manupa Karunaratne, Rathnayake Mudiyanseelage Sunimal Rathnayake, Rui Liu, Zhanglu Yan
 - 2019: Ho Nhut Minh, Qi Guo, Tan Cheng, Yudong Xu,
 - 2018: Zhang Junzhe, Zhang Jiangwei, Bao Ning, Yuan Yu
 - 2017: Zhong Guanwen, Dumitrel Loghin, Elavarasi Manogaran
 - 2016: Yingjun Wu, Rajendiran Ramanathan, Ramapantulu Lavanya
- 2) QE evaluators for the Computer systems cluster (2016 fall), database cluster (2019 spring)

H. Visiting Students/Interns

- Shanghai Jiao Tong University: Donghong Wu (2012), Yanchao Lu (2013-2014), Shuang Chen (2015), Shunning Jiang (2015), Xingyang Liu (2017), Hongjian Cao (2017), Mingyu Liang (2018), Zexin Xia (2018), Zhengyi Li (2019), Sizhe Wei(2019), Yucheng Zhang (2020), Yiqun Diao (2020), Hou Xie (2021)
- Peking University: Yifan Gong (2011)
- University of Science and Technology China: Xiaosheng Xu (2011)
- Beihang University: Yujing Liu (2012)
- United International College: Shukun Xie (2012)
- Northwestern Polytechnical University: Xuntao Cheng (2012)
- Huazhong University of Science and Technology: Long Zheng (2013).
- VIT University, Chennai: Saurabh Jha (2014)
- Tianjin University, China: Hao Fu (2017-2018)
- Shandong University, China: Qiuyi Lv (2017-2018)
- South China University of Technology, China: Jiashuai Shi (2017-2018)

I. Teaching/Advising Honors and Awards

- **Nomination to Annual Teaching Excellence Award (ATEA) 2020/2021.**
- **NUS SoC Faculty Teaching Excellence Award (FTEA) 2019/2020**
- **Nanyang Education Award (School) 2014** at Nanyang Technological University (also known as “best teacher of the year”, 1 out of 80+ faculties in the school,

<http://www.ntu.edu.sg/NanyangAwards/Recipients/Pages/NanyangAwards2014Winners.aspx>)

- **Advised Student Honors**

- Dr. Shengliang Lu, Ms. Shu-Ting Chang and Mr. Hou Xie got Runner Up in Singapore Blockchain Innovation Challenge 2022 Enthusiasts Track (organized by Singapore Blockchain Innovation Program (SBIP))
- Mr. Xinyu Chen got AMD HACC Outstanding Researcher Award 2021 (one of the four recipients worldwide)
- Mr. Johan Kok Zhi Kang got IDS Student Research Achievement Award 2022.
- Mr. Shengliang Lu and Mr. Qinbin Li got Dean's Graduate Research Excellence Award, Sem 2 AY2021/2022.
- Mr. Johan Kok Zhi Kang got Research Achievement Award of SoC, Sem 1 AY2021/2022.
- Mr. Hongshi Tan, Xilinx XACC Outstanding Contributor Award 2021 (one of the four recipients worldwide)
- Mr. Shengliang Lu, Mr. Qinbin Li and Mr. Xinyu Chen got Research Achievement Award of SoC, 2021.
- Mr. Qinbin Li got Google PhD Fellowship 2021
- Mr. Ang Jie Liang's FYP got Outstanding Computing Project Prize 2020/2021 of School of Computing
- Mr. Qinbin Li's AAI 2020 paper won PREMIA Best Student Paper Gold Award 2021 (<http://www.premiasg.org/for-members/premia-best-student-paper-awards/premia-best-student-paper-awards-2021/>)
- The intern from Shanghai Jiao Tong University, Mr. Yiqun Diao, got the prestigious NUS President Graduate Fellowship 2021.
- Our project ThunderGP got the third place of 2020 Xilinx Adaptive Computing Developer Contest (under the category of Adaptable Compute Acceleration, top 9 out of 79).
- Mr. Husong Liu published a VLDB demo paper in 2020 based on his FYP project.
- Mr. Qinbin Li and Mr. Xinyu Chen got Research Achievement Award of SoC, 2020.
- Mr. Chengxi Xue got the 2nd place in KDD Cup 2019 AutoML Competition (out of 160+ teams).**
- Mr. Shuhao Zhang won Research Achievement Award (RAA) of NUS SoC in AY2018/2019.
- Mr. Shengliang Lu got the 2nd place in Zhuhai ChiefData Lab Startup Competition 2018. (于2018年成立新加坡DataInz 科技有限公司，参加岭南大数据研究院举办的2018岭南大数据国际论坛暨创新创业大赛并获得第二名20万人民币的奖励。Led a startup company as a co-founder since 2018, participated in the Global Innovation Conference and Contest hosted in Zhuhai, China 2018, and was awarded as the first runner-up.)
- Mr. Shuhao Zhang won Research Achievement Award (RAA) of NUS SoC in AY2017/2018.
- IEEE CloudCom 2014 Best Ph.D. Consortium Paper Award (for the work "Amelie Chi Zhou* and Bingsheng He. Simplified Monetary Optimizations for Workflows in IaaS Clouds")**
- FYP Student Lim Chun Leng as a team member in getting Silver Award at the Asia Student Supercomputer Challenge (ASC14)
- FYP Student Xi Yewen won Defence Science & Technology Agency Gold Medal.

- v. FYP Student Tony Shuhao Zhang presented his database system prototype in VLDB 2013 (with the title “OmniDB: Towards Portable and Efficient Query Processing on Parallel CPU/GPU Architectures”. International Conference on Very Large Data Bases (VLDB) 2013 (System demo paper)), Year 2013
- w. URECA Student Cui Yan got “Peer Review” and “Public Vote” award in URECA poster competition, 2012
- x. URECA Students (Nguyen Ha Duy, Bui Le Linh) applied for SMART Explorer Grant 2011 (shortlisted) for “Captor: Enabling Cloud Services To Mobile Games”
- y. Coaches/advisor for ASC (ASC Student Supercomputer Challenge)
 - i. Highest LINPACK award in ASC 2015
 - ii. Application Innovation Award, Most Popular Team in ASC 2016
- Student/Postdoc Placement
 - a. Dr. Shixuan Sun (Postdoc) won the National Oversea Young Scholar Award and joined Shanghai Jiao Tong University as a Tenure-track Associate Professor, 2023.
 - b. Dr. Shuhao Zhang joined SUTD as a tenure-track assistant professor, 2021.

III. RESEARCH

A. Ph.D. Thesis

Title: “Cache-Oblivious Query Processing”

Completed: August 2008

Advisor: Dr. Qiong Luo

University: The Hong Kong University of Science & Technology (HKUST)

B. Research Group and Mission

I am leading Xtra Computing Group, affiliated with system and networking group as well as database group in NUS. Xtra’s mission is to build faster, greener and cheaper computing systems. The group has the following research interests:

- Big data management systems (with special interests in cloud computing and emerging hardware systems)
- Database systems
- Parallel and distributed systems
- Cloud Computing

C. Citations

<http://scholar.google.com.sg/citations?hl=en&user=RogYLKYAAAAJ> (on August 24, 2023).

	All	Since 2017
<u>Citations</u>	14057	8265
<u>h-index</u>	58	46
<u>i10-index</u>	202	169

Other sources:

ORCID: 0000-0001-8618-4581

D. Publications

Author notations: * denotes the author is a student advised by me. ^ denotes the author is a staff advised by me.

D.1. Five Selected Articles

- S1) **Bingsheng He**, Ke Yang, Rui Fang, Mian Lu, Naga K. Govindaraju, Qiong Luo, Pedro V. Sander. Relational Joins on Graphics Processors. *SIGMOD 2008: ACM SIGMOD International Conference on Management of data*, pages: 511-524, 2008. **[one of the two “Best papers of SIGMOD 2008”, invited to ACM TODS.]**
- S2) Jianlong Zhong* and **Bingsheng He**. Medusa: Simplified Graph Processing on GPUs. *IEEE TPDS: IEEE Transactions on Parallel and Distributed System*, vol.25, no.6, pp.1543-1552, June 2014, doi: 10.1109/TPDS.2013.111. **[Impact: widely cited and “VLDB 2013 best demo”, one of the earliest efficient graph processing systems on the GPU.]**
- S3) Paul Johns*, Jiong He*, **Bingsheng He**. GPL: A GPU-based Pipelined Query Processing Engine. *ACM SIGMOD: ACM SIGMOD International Conference on Management of data*, 2016. **[Impact: A new and efficient execution paradigm of query processing on the GPU]**
- S4) Shuhao Zhang*, Jiong He, Amelie Chi Zhou, **Bingsheng He**. BriskStream: Scaling Data Stream Processing on Shared-Memory Multicores. *ACM SIGMOD: ACM SIGMOD International Conference on Management of data*, 2019. **[Impact: review the common bottleneck of modern data stream systems, and develop scale-up solutions.]**
- S5) Zeyi Wen^, Jiashuai Shi*, **Bingsheng He**, Jian Chen, Kotagiri Ramamohanarao, Qinbin Li. Exploiting GPUs for Efficient Gradient Boosting Decision Tree Training. *IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems* 2019. doi: 10.1109/TPDS.2019.2920131. **[2019 Best Paper Award from IEEE Transactions on Parallel and Distributed Systems by the IEEE Computer Society Publications Board (1 out of 987 submissions)]**

More details on the impacts: <https://www.comp.nus.edu.sg/~hebs/impactproj.html>

D.2. Journal Articles

- J1) Jin Zhao, Yu Zhang, Ligang He, Qikun Li, Xiang Zhang, Xinyu Jiang, Hui Yu, Xiaofei Liao, Hai Jin, Lin Gu, Haikun Liu, **Bingsheng He**, Ji Zhang, Xianzheng Song, Lin Wang, Jun Zhou. GraphTune: An Efficient Dependency-aware Substrate to Alleviate Irregularity in Concurrent Graph Processing. *ACM TACO: ACM Transactions on Architecture and Code Optimization*, Volume 20, Issue 319, July 2023. Article No.: 37, pp 1–24.
- J2) Zining Zhang*, Yao Chen^, Bingsheng He, Zhenjie Zhang. NIOT: A Novel Inference Optimization of Transformers on Modern CPUs. *IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems*, 2023.
- J3) Zhaomin Wu*, Qinbin Li*, **Bingsheng He**. Practical Vertical Federated Learning with Unsupervised Representation Learning. *IEEE TBD: IEEE Transactions on Big Data (Special Issue on Trustable, Verifiable, and Auditable Federated Learning)* 2022.
- J4) Xinyu Chen*, Feng Cheng*, Hongshi Tan*, Yao Chen, **Bingsheng He**, Weng-fai Wong, Deming Chen. ThunderGP: HLS-based Graph Processing Framework on HBM-enabled FPGAs. *ACM TRET: ACM Trans. Reconfig. Technol. Syst.*, Vol. 1, No. 1, Article 1. Publication date: January 2022. **[Invited for “Top quality papers” in FPGA 2021]**
- J5) Zijun Li, Linsong Guo, Jiagan Cheng, Quan Chen, **Bingsheng He**, Minyi Guo. The Serverless Computing Survey: A Technical Primer for Design Architecture. *ACM CSUR: ACM Computing Surveys* 2022.
- J6) Sixu Hu*, Yuan Li*, Xu Liu*, Qinbin Li*, Zhaomin Wu*, **Bingsheng He**. The OARF Benchmark Suite: Characterization and Implications for Federated Learning Systems. *ACM TIST: ACM Transactions on Intelligent Systems and Technology*.

- J7) Zijun Li, Linsong Guo, Jiagan Cheng, Quan Chen, **Bingsheng He**, and Minyi Guo. The Serverless Computing Survey: A Technical Primer for Design Architecture. ACM Comput. Surv.
- J8) Qingyu Xu, Feng Zhang, Mingde Zhang, Jidong Zhai, **Bingsheng He**, Cheng Yang, Shuhao Zhang, Jiazao Lin, Haidi Liu, Xiaoyong. Payment Behavior Prediction on Shared Parking Lots with TR-GCN. VLDBJ: The VLDB Journal 2021.
- J9) Yawen Chen, Zeyi Wen[^], **Bingsheng He**, Jian Chen. Efficient Decomposition Selection for Multi-class Classification. IEEE TKDE: IEEE Transactions on Knowledge and Data Engineering 2021.
- J10) Jiantong Jiang, Zeyi Wen[^], Zeke Wang, **Bingsheng He** and Jian Chen. Parallel and Distributed Structured SVM Training. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, 2021.
- J11) Xiaofei Liao, Jin Zhao, Yu Zhang, **Bingsheng He**, Ligang He, Hai Jin, Lin Gu. A Structure-aware Storage Optimization for Out-of-Core Concurrent Graph Processing. IEEE TC: IEEE Transactions on Computers 2021.
- J12) Amelie Chi Zhou, Weilin Xue, Yao Xiao, **Bingsheng He**, Shadi Ibrahim and Reynold Cheng. Taming System Dynamics on Resource Optimization for Data Processing Workflows: A Probabilistic Approach. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, 2021.
- J13) Yu Zhang, Da Peng, Xiaofei Liao, Hai Jin, Haikun Liu, Lin Gu, **Bingsheng He**. LargeGraph: An Efficient Dependency-Aware GPU-Accelerated Large-Scale Graph Processing. ACM TACO: ACM Trans. Archit. Code Optim. 18(4): 58:1-58:24 (2021)
- J14) Qinbin Li*, Zeyi Wen[^], Zhaomin Wu*, Sixu Hu*, Naibo Wu*, **Bingsheng He**. Federated Learning Systems: Vision, Hype and Reality for Data Privacy and Protection. IEEE TKDE: IEEE Transactions on Knowledge and Data Engineering 2021. <https://arxiv.org/abs/1907.09693>
- J15) Feng Zhang, Chenyang Zhang, Lin Yang, Shuhao Zhang, **Bingsheng He**, Wei Lu, Xiaoyong Du. Fine-Grained Multi-Query Stream Processing on Integrated Architectures. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, 2021.
- J16) Feng Zhang, Jiya Su, Weifeng Liu, **Bingsheng He**, Ruofan Wu, Xiaoyong Du, Rujia Wang. YuenyeungSpTRSV: A Thread-Level and Warp-Level Fusion Synchronization-Free Sparse Triangular Solve. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, 2021.
- J17) Feng Zhang, Yani Liu, Ningxuan Feng, Cheng Yang, Jidong Zhai, Shuhao Zhang, **Bingsheng He**, Jiazao Lin, Xiao Zhang, Xiaoyong Du. Periodic Weather-Aware LSTM with Event Mechanism for Parking Behavior Prediction. IEEE TKDE: IEEE Transactions on Knowledge and Data Engineering 2021.
- J18) Chenyang Zhang, Feng Zhang, Xiaoguang Guo, **Bingsheng He**, Xiaoyong Du. iMLBench: A Machine Learning Benchmark Suite for CPU-GPU Integrated Architectures. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, Special Section on Parallel and Distributed Computing Techniques for AI, ML and DL 2020.
- J19) Haikun Liu, Di Chen, Hai Jin, Xiaofei Liao, **Bingsheng He**, Kan Hu, Yu Zhang. A Survey of Non-Volatile Main Memory Technologies: State-of-the-Arts, Practices, and Future Directions. JCST: Journal of Computer Science and Technology 36 (1), 4-32, 2020.
- J20) Zeyi Wen[^], Mingyu Liang*, **Bingsheng He**, Zexin Xia*. A High-Performance Index for Real-Time Matrix Retrieval. IEEE TKDE: IEEE Transactions on Knowledge and Data Engineering 2020.
- J21) Yu Zhang, Xiaofei Liao, Lin Gu, Hai Jin, Kan Hu, Haikun Liu, **Bingsheng He**. AsynGraph: Maximizing Data Parallelism for Efficient Iterative Graph Processing on GPUs. ACM TACO: ACM Trans. Archit. Code Optim. 17, 4, Article 29 (December 2020), 21 pages. DOI:<https://doi.org/10.1145/3416495>.
- J22) Dawen Xu, Ying Wang, Kaijie Tu, Cheng Liu, **Bingsheng He**, Lei Zhang. Accelerating Generative Neural Networks on Unmodified Deep Learning Processors - A Software Approach. IEEE TC: IEEE Transactions on Computers (Special Issue on Domain-Specific Architectures for Emerging Applications). vol. 69, no. 8, pp. 1172-1184, 1 Aug. 2020, doi: 10.1109/TC.2020.3001033.

- J23) Zeyi Wen[^], Jiashuai Shi^{*}, Hanfeng Liu^{*}, **Bingsheng He**, Qinbin Li^{*}, Jian Chen. ThunderGBM: Fast GBDTs and Random Forests on GPUs. JMLR: Journal of Machine Learning Research, 21(108):1–5, 2020.
- J24) Shanjiang Tang, **Bingsheng He**, Ce Yu, Yusen Li, Kun Li. A Survey on Spark Ecosystem: Big Data Processing Infrastructure, Machine Learning, and Applications. IEEE TKDE: IEEE Transactions on Knowledge and Data Engineering 2020. doi: 10.1109/TKDE.2020.2975652.
- J25) Haikun Liu, Renshan Liu, Xiaoei Liao, Hai Jin, **Bingsheng He**, Yu Zhang. Object-level Memory Allocation and Migration in Hybrid Memory Systems. IEEE TC: IEEE Transactions on Computers. vol. 69, no. 9, pp. 1401-1413, 1 Sept. 2020, doi: 10.1109/TC.2020.2973134.
- J26) Johns Paul^{*}, **Bingsheng He**, Shengliang Lu^{*}, Chiew Tong Lau. Revisiting Hash Join on Graphics Processors: A Decade Later. DAPD: Springer Distrib Parallel Databases (2020). <https://doi.org/10.1007/s10619-019-07280-z>. [invited submission to special issue of HardBD 2019]
- J27) Qinbin Li^{*}, Zeyi Wen[^], **Bingsheng He**. Adaptive Kernel Value Caching for SVM Training. IEEE TNNLS: IEEE Transactions on Neural Networks and Learning Systems, 31(7): 2376-2386 (2020).
- J28) Zeke Wang[^], Xue Liu, Kai Zhang[^], Haihang Zhou^{*}, **Bingsheng He**. Understanding and Optimizing Conjunctive Predicates under Memory-efficient Storage Layouts. IEEE TKDE: IEEE Transactions on Knowledge and Data Engineering 2019. doi: 10.1109/TKDE.2019.2958672
- J29) Amelie Chi Zhou, Bingkun Shen, Yao Xiao, Shadi Ibrahim, **Bingsheng He**. Cost-Aware Partitioning for Efficient Large Graph Processing in Geo-Distributed Datacenters. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, vol. 31, no. 7, pp. 1707-1723, 1 July 2020, doi: 10.1109/TPDS.2019.2955494.
- J30) Qiumin Lu, Zheng Xiao, Jiacheng Ma, Yaozu Dong, Zhengwei Qi, Jianguo Yao, **Bingsheng He**, Haibing Guan. gMig: Efficient vGPU Live Migration with Overlapped Software-based Dirty Page Verification. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, vol. 31, no. 5, 1 May 2020, doi: 10.1109/TPDS.2019.2947521.
- J31) Feng Zhang, Jidong Zhai, Bo Wu, **Bingsheng He**, Wenguang Chen, Xiaoyong Du. Automatic Irregularity-Aware Fine-Grained Workload Partitioning on Integrated Architectures. IEEE TKDE: IEEE Transactions on Knowledge and Data Engineering 2019.
- J32) Chuangyi Gui, Long Zheng, **Bingsheng He**, Cheng Liu, Xinyu Chen, Xiao-Fei Liao, Hai Jin. A Survey on Graph Processing Accelerators: Challenges and Opportunities. JCST: Journal of Computer Science and Technology volume 34, pages339–371(2019).
- J33) Yun Peng, Xin Lin, Byron Choi, **Bingsheng He**. VColor*: A Practical Approach for Coloring Large Graphs. FCS: Frontiers of Computer Science 2019.
- J34) Shuhao Zhang^{*}, Feng Zhang, Yingjun Wu, Paul Johns, **Bingsheng He**. Hardware-Conscious Stream Processing: A Survey. ACM SigRec: SIGMOD Rec. 48, 4 (December 2019), 18–29. DOI: <https://doi.org/10.1145/3385658.3385662>.
- J35) Zeyi Wen[^], Jiashuai Shi^{*}, **Bingsheng He**, Jian Chen, Kotagiri Ramamohanarao, Qinbin Li. Exploiting GPUs for Efficient Gradient Boosting Decision Tree Training. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems 2019. doi: 10.1109/TPDS.2019.2920131. [2019 Best Paper Award from IEEE Transactions on Parallel and Distributed Systems by the IEEE Computer Society Publications Board (1 out of 987 submissions)]
- J36) Xiaoyuan Wang, Haikun Liu, Xiaofei Liao, Ji Chen, Hai Jin, Yu Zhang, Long Zheng, **Bingsheng He**, and Song Jiang. 2019. Supporting Superpages and Lightweight Page Migration in Hybrid Memory Systems. ACM TACO: ACM Trans. Archit. Code Optim. 16, 2, Article 11 (April 2019), 26 pages. DOI: <https://doi.org/10.1145/3310133>
- J37) Jieru Zhao, Liang Feng, Wei Zhang, Sharad Sinha, Yun (Eric) Liang, **Bingsheng He**. Performance Modeling and Directives Optimization for High Level Synthesis on FPGA. IEEE TCAD: IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 2019.
- J38) Yu Zhang, Jin Zhao, Xiaofei Liao, Hai Jin, Lin Gu, Haikun Liu, **Bingsheng He**, Ligang He. CGraph: A Distributed Storage and Processing System for Concurrent Iterative Graph

Analysis Jobs. ACM TOS: ACM Transactions on Storage 15, 2, Article 1 (June 2019), 26 pages. DOI:<https://doi.org/10.1145/3319406>.

- J39) Amelie Chi Zhou*, Yao Xiao, Yifan Gong, **Bingsheng He**, Jidong Zhai, Rui Mao. Privacy Regulation Aware Process Mapping in Geo-Distributed Cloud Data Centers. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems 2019.
- J40) Hai Jin, Bo Liu, Wenbin Jiang, Yang Ma, Xuanhua Shi, **Bingsheng He**, Shaofeng Zhao. Layer-Centric Memory Reuse and Data Migration for Extreme-Scale Deep Learning on Many-Core Architectures. ACM TACO: ACM Transactions on Architecture and Code Optimization (TACO): Volume 15 Issue 3, October 2018. <https://doi.org/10.1145/3243904>
- J41) Yu Zhang, Xiaofei Liao, Xiang Shi, Hai Jin, **Bingsheng He**. Efficient Disk-Based Directed Graph Processing: A Strongly Connected Component Approach, IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, vol. 29, no. 4, pp. 830-842, 1 April 2018. doi: 10.1109/TPDS.2017.2776115
- J42) Zeyi Wen[^], Jiashuai Shi*, **Bingsheng He**, Jian Chen, Yawen Chen. Efficient Multi-Class Probabilistic SVMs on GPUs. IEEE TKDE: IEEE Transactions on Knowledge and Data Engineering 2018.
- J43) Zeyi Wen[^], Jiashuai Shi*, Qinbin Li*, **Bingsheng He**, Jian Chen. ThunderSVM: A Fast SVM Library on GPUs and CPUs. JMLR: Journal of Machine Learning Research 19 (2018) 1-5.
- J44) Zhifei Pang, Sai Wu, Gang Chen, Lidan Shou, Ke Chen, **Bingsheng He**. A Stack-Centric Processing Model For Iterative Processing. IEEE TBD: IEEE Transactions on Big Data, doi: 10.1109/TBDATA.2018.2841363, accepted in 2018.
- J45) Xuntao Cheng*, **Bingsheng He**, Mian Lu, Chiew Tong Lau. Many-Core Needs Fine-Grained Scheduling: A Case Study of Query Processing on Intel Xeon Phi Processors. JPDC: Journal of Parallel and Distributed Computing (special issue on in-memory computing), October 2018, Pages 395-404, <https://doi.org/10.1016/j.jpdc.2017.09.005>.
- J46) Mochi Xue, Jiacheng Ma, Wentai Li, Kun Tian, Yaozu Dong, Jinyu Wu, Zhengwei Qi, **Bingsheng He**, Haibing Guan. Scalable GPU Virtualization with Dynamic Sharing of Graphics Memory Space. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, vol. 29, no. 8, pp. 1823-1836, 1 Aug. 2018. doi: 10.1109/TPDS.2018.2789883.
- J47) Zhaojie Niu*, **Bingsheng He**, Fangming Liu. JouleMR: Towards Cost-Effective and Green-Aware Data Processing Frameworks. IEEE TBD: IEEE Transactions on Big Data, vol. 4, no. 2, pp. 258-272, 1 June 2018. doi: 10.1109/TBDATA.2017.2655037 (special issue on "Big Data Infrastructure").
- J48) Shanjiang Tang[^], Zhaojie Niu*, **Bingsheng He**, Bu-Sung Lee, Ce Yu. Long-Term Multi-Resource Fairness for Pay-as-you Use Computing Systems. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, vol. 29, no. 5, pp. 1147-1160, 1 May 2018. doi: 10.1109/TPDS.2017.2788880.
- J49) Shanjiang Tang[^], Bu-Sung Lee, **Bingsheng He**. Fair Resource Allocation for Data-Intensive Computing in the Cloud. IEEE TSC: IEEE Transactions on Services Computing, vol. 11, no. 1, 1 Jan.-Feb. 2018. doi: 10.1109/TSC.2016.2531698.
- J50) Xuanhua Shi, Xuan Luo, Junling Liang, Peng Zhao, Sheng Di, **Bingsheng He**, Hai Jin. Frog: Asynchronous Graph Processing on GPU with Hybrid Coloring Model. IEEE TKDE: IEEE Transactions on Knowledge and Data Engineering, vol. 30, no. 1, pp. 29-42, 1 Jan. 2018. doi: 10.1109/TKDE.2017.2745562.
- J51) Lei Wei*, Chuan Heng Foh, **Bingsheng He**, Jianfei Cai. Towards Efficient Resource Allocation for Heterogeneous Workloads in IaaS Clouds. IEEE TCC: IEEE Transactions on Cloud Computing, vol. 6, no. 1, pp. 264-275, 1 Jan.-March 2018. doi: 10.1109/TCC.2015.2481400.
- J52) Kai Zhang[^], Kaibo Wang, Yuan Yuan, Lei Guo, Rubao Li, Xiaodong Zhang, **Bingsheng He**, Jiayu Hu, Bei Hua. A Distributed In-Memory Key-Value Store System on Heterogeneous CPU-GPU Cluster. VLDBJ: The VLDB Journal 26, 5 (October 2017), 729-750. DOI: <https://doi.org/10.1007/s00778-017-0479-0>.
- J53) Haikun Liu, **Bingsheng He**, Xiaofei Liao, Hai Jin. Towards Declarative and Data-centric Virtual Machine Image Management in IaaS Clouds. IEEE TCC: IEEE Transactions on Cloud Computing. doi: 10.1109/TCC.2017.2728066.

- J54) Zhaojie Niu*, Shanjiang Tang^, **Bingsheng He**. An Adaptive Efficiency-Fairness Meta-scheduler for Data-Intensive Computing. IEEE TSC: IEEE Transactions on Services Computing, vol.PP, no.99, pp.1-1. doi: 10.1109/TSC.2016.2635133.
- J55) Zeke Wang^, **Bingsheng He**, Wei Zhang. Multi-kernel Data Partitioning with Channel on OpenCL-based FPGAs. IEEE TVLSI: IEEE Transactions on Very Large Scale Integration (VLSI) Systems, vol. 25, no. 6, pp. 1906-1918, June 2017. doi: 10.1109/TVLSI.2017.2653818.
- J56) Yinliang Yue, **Bingsheng He**, Yuzhe Li, Weiping Wang. Building an Efficient Put-Intensive Key-Value Store with Skip-tree. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, vol. 28, no. 4, pp. 961-973, April 1 2017. doi: 10.1109/TPDS.2016.2609912.
- J57) Yifan Gong*, **Bingsheng He**, Dan Li. Network Performance Aware Optimizations on IaaS Clouds. IEEE TC: IEEE Transactions on Computers, vol. 66, no. 4, pp. 672-687, April 1 2017. doi: 10.1109/TC.2016.2609387.
- J58) Amelie Chi Zhou*, **Bingsheng He**, Xuntao Cheng*, Chiew Tong Lau. A Declarative Optimization Engine for Resource Provisioning of Scientific Workflows in Geo-Distributed Clouds. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, vol. 28, no. 3, pp. 647-661, March 1 2017. doi: 10.1109/TPDS.2016.2599529.
- J59) Feng Zhang, Jidong Zhai, **Bingsheng He**, Shuhao Zhang, Wenguang Chen. Understanding Co-running Behaviors on Integrated CPU/GPU Architectures. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, 28, 3 (March 2017), 905-918. DOI: <https://doi.org/10.1109/TPDS.2016.2586074>.
- J60) Lu Zhang^, Xueyan Tang, **Bingsheng He**. Analysis of Minimum Interaction Time for Continuous Distributed Interactive Computing. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, vol. 28, no. 2, pp. 401-415, Feb. 1 2017. doi: 10.1109/TPDS.2016.2585140.
- J61) Fubing Mao*, Yi-Chung Chen, Wei Zhang, Hai (Helen) Li, **Bingsheng He**. Library-based Placement and Routing in FPGAs with Support of Partial Reconfiguration. ACM TODAES: ACM Trans. Des. Autom. Electron. Syst. 21, 4, Article 71 (May 2016), 26 pages. DOI: <https://doi.org/10.1145/2901295>.
- J62) Tao Luo*, Wei Zhang, Bingsheng He, Douglas Maskell. A Hybrid Logic Block Architecture in FPGA for Holistic Efficiency. IEEE TCAS-II: IEEE Transactions on Circuits and Systems II: Express Briefs, vol. 64, no. 1, pp. 71-75, Jan. 2017. doi: 10.1109/TCSII.2016.2551555.
- J63) Lei Wei*, Jianfei Cai, Chuan Heng Foh, **Bingsheng He**. QoS-aware Resource Allocation for Video Transcoding in Clouds. IEEE TCSVT: IEEE Transactions on Circuits and Systems for Video Technology, vol. 27, no. 1, pp. 49-61, Jan. 2017. doi: 10.1109/TCSVT.2016.2589621.
- J64) Zeke Wang^, Shuhao Zhang*, **Bingsheng He**, Wei Zhang. Melia: A MapReduce Framework on OpenCL-based FPGAs. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, Volume 27, Number 12, Dec 2016, pp. 3547-3560. DOI: <https://doi.org/10.1109/TPDS.2016.2537805>.
- J65) Haikun Liu^, **Bingsheng He**. F2C: Enabling Fair and Fine-grained Resource Sharing in Multi-tenant IaaS Clouds. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, vol. 27, no. 9, pp. 2589-2602, Sept. 1 2016. doi: 10.1109/TPDS.2015.2499769.
- J66) Jun Yang, Qingsong Wei, Chundong Wang, Cheng Chen, Khai Leong Yong, **Bingsheng He**. NV-Tree: A Consistent and Workload-adaptive Tree Structure for Non-volatile Memory. IEEE TC: IEEE Transactions on Computers, vol. 65, no. 7, pp. 2169-2183, July 1 2016. doi: 10.1109/TC.2015.2479621.
- J67) Xiaofei Liao, Long Zheng*, **Bingsheng He**, Song Wu, Hai Jin. A Performance Debugging Framework for Unnecessary Lock Contentions with Record/Replay Techniques. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, vol. 27, no. 7, pp. 1889-1901, July 1 2016. doi: 10.1109/TPDS.2015.2472412.
- J68) Yingnan Cui*, Wei Zhang, Vivek Chaturvedi, **Bingsheng He**. Decentralized Thermal-aware Task Scheduling for Large-scale Many-core Systems. IEEE TVLSI: IEEE Transactions on Very Large Scale Integration (VLSI) Systems, vol. 24, no. 6, pp. 2075-2088, June 2016. doi: 10.1109/TVLSI.2015.2497469.

- J69) Amelie Chi Zhou*, **Bingsheng He**, Cheng Liu[^]. Monetary Cost Optimizations for Hosting Workflow-as-a-Service in IaaS Clouds. IEEE TCC: IEEE Transactions on Cloud Computing, vol.4, no. 1, pp. 34-48, Jan.-March 2016, doi:10.1109/TCC.2015.2404807.
- J70) Shanjiang Tang[^], Bu-Sung Lee, **Bingsheng He**. Dynamic Job Ordering and Slot Configurations for MapReduce Workloads. IEEE TSC: IEEE Transactions on Services Computing, vol. 9, no. 1, pp. 4-17, Jan.-Feb. 1 2016. doi: 10.1109/TSC.2015.2426186.
- J71) Qiumin Lu, Jianguo Yao, Zhengwei Qi, **Bingsheng He**, Haibing Guan. Fairness-Efficiency Allocation of CPU-GPU Heterogeneous Resources. IEEE TSC: IEEE Transactions on Services Computing. doi: 10.1109/TSC.2016.2597141.
- J72) Yanchao Lu*, Donghong Wu*, **Bingsheng He**, Xueyan Tang, Jianliang Xu, Minyi Guo. Rank-Aware Dynamic Migrations and Adaptive Demotions for DRAM Power Management. IEEE TC: IEEE Transactions on Computers, vol. 65, no. 1, Jan. 1 2016. doi: 10.1109/TC.2015.2409847.
- J73) Yinliang Yue, **Bingsheng He**, Lei Tian, Hong Jiang, Fang Wang, Dan Feng. Rotated Logging Storage Architectures for Data Centers: Models and Optimizations. IEEE TC: IEEE Transactions on Computers, vol. 65, no. 1, Jan. 1 2016. doi: 10.1109/TC.2015.2417539.
- J74) Shen Gao, Jianliang Xu, Theo Härder, **Bingsheng He**, Byron Choi and Haibo Hu. PCMLogging: Optimizing Transaction Logging and Recovery Performance with PCM. IEEE TKDE: IEEE Transactions on Knowledge and Data Engineering, vol. 27, no. 12, pp. 3332-3346, Dec. 1 2015. doi: 10.1109/TKDE.2015.2453154.
- J75) Mian Lu, Yun Liang, Huynh Phung Huynh, Ong Zhong Liang, Bingsheng He, Lei Zhang, Richard Huynh, Rick Siow Mong Goh. MrPhi: An Optimized MapReduce Framework on Intel Xeon Phi Coprocessors. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, vol. 26, no. 11, pp. 3066-3078, Nov. 1 2015. doi: 10.1109/TPDS.2014.2365784.
- J76) Yanchao Lu*, **Bingsheng He**, Xueyan Tang, Minyi Guo. Synergy of Dynamic Frequency Scaling and Demotion on DRAM Power Management: Models and Optimizations. IEEE TC: IEEE Transactions on Computers, vol.64, no.8, pp.2367-2381, Aug. 1 2015, doi: 10.1109/TC.2014.2360534.
- J77) Dan Li, Yirong Yu, Wu He, Kai Zheng, Bingsheng He. Willow: Saving Data Center Network Energy for Network-limited Flows. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, vol.26, no.9, Sept. 1 2015, doi: 10.1109/TPDS.2014.2350990.
- J78) Zeke Wang[^], Xue Liu, **Bingsheng He**, Feng Yu. A Combined SDC-SDF Architecture for Normal I/O Pipelined Radix-2 FFT. IEEE TVLSI: IEEE Transactions on Very Large Scale Integration (VLSI) Systems, vol. 23, no. 5, pp. 973-977, May 2015. doi: 10.1109/TVLSI.2014.2319335.
- J79) Haikun Liu[^], Hai Jin, Xiaofei Liao, Wei Deng, **Bingsheng He**, Cheng-zhong Xu. Hotplug or Ballooning: A Comparative Study on Dynamic Memory Management Techniques for Virtual Machines. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, vol.26, no.5, pp.1350-1363, May 1 2015, doi: 10.1109/TPDS.2014.2320915.
- J80) Haikun Liu[^], **Bingsheng He**. VMbuddies: Coordinating Live Migration of Multi-Tier Applications in Cloud Environments. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, vol.26, no.4, April 1 2015, pp.1192-1205, doi: 10.1109/TPDS.2014.2316152.
- J81) **Bingsheng He**, Jeffrey Xu Yu, Amelie Chi Zhou. Improving Update-Intensive Workloads on Flash Disks Through Exploiting Multi-Chip Parallelism. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, vol.26, no.1, pp.152-162, Jan. 2015.
- J82) Yifan Gong*, **Bingsheng He**, Jianlong Zhong*. Network Performance Aware MPI Collective Communication Operations in the Cloud. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, vol. 26, no. 11, pp. 3079-3089, Nov. 2015. doi: 10.1109/TPDS.2013.96.
- J83) Yingnan Cui*, Wei Zhang, Vivek Chaturvedi, **Bingsheng He**. Thermal-Aware Task Scheduling for 3D-Network-on-Chip: A Bottom to Top Scheme. JCSC: Journal of Circuits, Systems and Computers, Vol. 25, No. 01, 1640003 (2016). <https://doi.org/10.1142/S021812661640003X>.
- J84) Wan Du, Zikun Xing, Mo Li, **Bingsheng He**, Lloyd Hock Chye Chua, Haiyan Miao. Sensor Placement and Measurement of Wind for Water Quality Studies in Urban Reservoirs.

- ACM TMSN: ACM Trans. Sen. Netw. 11, 3, Article 41 (May 2015), 27 pages. DOI:<https://doi.org/10.1145/2700265>.
- J85) Dong Huang[^], **Bingsheng He**, Chunyan Miao. A Survey of Resource Management in Multi-Tier Web Applications. IEEE COMST: IEEE Communications Surveys and Tutorials, vol.16, no.3, pp.1574,1590, Third Quarter 2014.
- J86) Jianlong Zhong^{*}, **Bingsheng He**. Medusa: A Parallel Graph Processing System on Graphics Processors. ACM SigRec: SIGMOD Rec. 43, 2 (June 2014), 35–40. DOI:<https://doi.org/10.1145/2694413.2694421>. [Invited as “Best of VLDB 2013 demo”]
- J87) Shanjiang Tang, Bu-Sung Lee, Bingsheng He. DynamicMR: A Dynamic Slot Allocation Optimization Framework for MapReduce Clusters. IEEE TCC: IEEE Transactions on Cloud Computing, vol.2, no.3, pp.333--347, July-Sept. 1 2014. doi: 10.1109/TCC.2014.2329299. [Invited presentation in IEEE CloudCom 2014.]
- J88) Amelie Chi Zhou^{*}, **Bingsheng He**. Transformation-based Monetary Cost Optimizations for Workflows in the Cloud. IEEE TCC: IEEE Transactions on Cloud Computing, vol.2, no. 1, pp. 85-98, March 2014. [Spotlight article of the issue, invited presentation in IEEE CloudCom 2014.]
- J89) Jianlong Zhong^{*}, **Bingsheng He**. Kernelet: High-Throughput GPU Kernel Execution with Dynamic Slicing and Scheduling. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, vol.25, no.6, pp.1522-1532, June 2014, doi: 10.1109/TPDS.2013.257.
- J90) Jianlong Zhong^{*}, **Bingsheng He**. Medusa: Simplified Graph Processing on GPUs. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, vol.25, no.6, pp.1543-1552, June 2014, doi: 10.1109/TPDS.2013.111.
- J91) Sai Tung On^{*}, Shen Gao, **Bingsheng He**, Ming Wu, Qiong Luo, Jianliang Xu. FD-Buffer: A Cost-Based Adaptive Buffer Replacement Algorithm for Flash Memory Devices. IEEE TC: IEEE Transactions on Computers, vol.63, no.9, pp.2288--2301, Sept. 2014.
- J92) Shadi Ibrahim, Hai Jin, Lu Lu, **Bingsheng He**, Gabriel Antoniu, Song Wu. Handling Partitioning Skew in MapReduce using LEEN. PPNA: Springer Peer-to-Peer Networking and Applications, Volume 6, Issue 4, pp 409-424, 2013. [IF 2013: 0.456, 5-year IF: 0.607]
- J93) Sai Tung On, Jianliang Xu, Byron Choi, Haibo Hu, **Bingsheng He**. Flag Commit: Supporting Efficient Transaction Recovery on Flash-based DBMSs. IEEE TKDE: IEEE Transactions on Knowledge and Data Engineering, Volume: 24, Issue: 9, Page(s): 1624-1639, 2011.
- J94) Wenbin Fang, **Bingsheng He**, Qiong Luo, Naga K. Govindaraju. Mars: Accelerating MapReduce with Graphics Processors. IEEE TPDS: IEEE Transactions on Parallel and Distributed Systems, Volume 22, Number 4, April 2011, pp. 608-620.
- J95) **Bingsheng He**, Mian Lu, Ke Yang, Rui Fang, Naga K. Govindaraju, Qiong Luo, Pedro V. Sander, Relational Query Co-Processing on Graphics Processors. ACM TODS: ACM Transactions on Database Systems (TODS), Volume 34 Issue 4, December 2009, pages = {21:1--21:39}. [Invited as “Best of SIGMOD 2008”]
- J96) **Bingsheng He**, Qiong Luo, Cache-Oblivious Databases: Limitations and Opportunities. ACM TODS: ACM Transactions on Database Systems (TODS), Volume 33 Issue 2, June 2008, pages = {8:1--8:42}.
- J97) **Bingsheng He**, Qiong Luo, Byron Choi. Adaptive Index Utilization in Memory-Resident Structural Joins. IEEE TKDE: IEEE Transactions on Knowledge and Data Engineering, vol. 19, no. 6, pp. 772-788, June, 2007.
- J98) **Bingsheng He**, Qiong Luo, Byron Choi. Cache-Conscious Automata for XML Filtering. IEEE TKDE: IEEE Transactions on Knowledge and Data Engineering, vol. 18, no. 12, pp. 1629-1644, Dec., 2006.

D.3. Conference and Workshop Papers

- C1) Qiange Wang[^], Yao Chen[^], Weng-Fai Wong, **Bingsheng He**. HongTu: Scalable Full-Graph GNN Training on Multiple GPUs. ACM SIGMOD: ACM SIGMOD International Conference on Management of Data 2024.

- C2) Sixu Hu*, Qinbin Li*, **Bingsheng He**. Communication-Efficient Generalized Neuron Matching for Federated Learning. ICPP 2023: 52nd International Conference on Parallel Processing.
- C3) Johan Kok Zhi Kang*, Sien Yi Tan, **Bingsheng He**, Zhen Zhang[^]. Real Time Index and Search Across Large Quantities of GNN Experts For Low Latency Online Learning. ACM KDD: In Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '2023).
- C4) Xuanhe Zhou, Cheng Chen, Kunyi Li, **Bingsheng He**, Mian Lu, Qiaosheng Liu, Wei Huang, Guoliang Li, Zhao Zheng, Yuqiang Chen. FEBench: A Benchmark for Real-Time Relational Data Feature Extraction. VLDB: International Conference on Very Large Data Bases (VLDB) 2023. (also published in Proc. VLDB Endow. Vol. 16 (2023), Track of Industrial, Application and Experience).
- C5) Sen Gao*, Hongchao Qin, Rong-Hua Li, **Bingsheng He**. Parallel Colorful h-star Core Maintenance in Dynamic Graphs. VLDB: International Conference on Very Large Data Bases (VLDB) 2023. (also published in Proc. VLDB Endow. Vol. 16 (2023), Track of Industrial, Application and Experience).
- C6) Jiaxin Jiang[^], Yuan Li*, **Bingsheng He**, Bryan Hooi, Jia Chen and Johan Kok Zhi Kang*: Spade: A Real-Time Fraud Detection Framework on Evolving Graphs. VLDB: International Conference on Very Large Data Bases (VLDB) 2023. (also published in Proc. VLDB Endow. Vol. 16 (2022)).
- C7) Qinbin Li*, **Bingsheng He**, Dawn Song. Adversarial Collaborative Learning on Non-IID Features. ICML: Proceedings of International Conference on Machine Learning (ICML-23).
- C8) Zining Zhang*, **Bingsheng He**, Zhenjie Zhang. Practical Edge Kernels for Integer-Only Vision Transformers Under Post-training Quantization. MLSys: Sixth Conference on Machine Learning and Systems, 2023.
- C9) Qinbin Li*, Zhaomin Wu*, Yanzheng Cai*, Yuxuan Han*, Ching Man Yung*, Tianyuan Fu*, **Bingsheng He**. FedTree: A Federated Learning System for Trees. MLSys: Sixth Conference on Machine Learning and Systems, 2023.
- C10) Sihao Hu*, Zhen Zhang[^], Bingqiao Luo*, Shengliang Lu*, **Bingsheng He** and Ling Liu. BERT4ETH: Pre-training of Transformer Encoder Representation for Ethereum Fraud Detection. WWW: The ACM Web Conference 2023.
- C11) Yiqun Diao*, Qinbin Li*, **Bingsheng He**, Towards Addressing Label Skews in One-shot Federated Learning. ICLR: The Eleventh International Conference on Learning Representations 2023.
- C12) Wentao Huang, Yunhong Ji[^], Xuan Zhou, **Bingsheng He**, Kian-Lee Tan. A Design Space Exploration and Evaluation for Main-Memory Hash Joins in Storage Class Memory. VLDB: International Conference on Very Large Data Bases (VLDB) 2023. (also published in Proc. VLDB Endow. Vol. 16 (2022)).
- C13) Chenyang Zhang, Feng Zhang, Kuangyu Chen, Mingjun Chen, **Bingsheng He**, Xiaoyong Du. EdgeNN: Efficient Neural Network Inference for CPU-GPU Integrated Edge Devices. IEEE ICDE: IEEE International Conference on Data Engineering 2023.
- C14) Zhaomin Wu*, Junhui Zhu*, Qinbin Li*, **Bingsheng He**. DeltaBoost: Gradient Boosting Decision Trees with Efficient Machine Unlearning. ACM SIGMOD: ACM SIGMOD International Conference on Management of Data 2023.
- C15) Hongshi Tan*, Xinyu Chen*, Yao Chen[^], **Bingsheng He**, Weng-Fai Wong. LightRW: FPGA Accelerated Graph Dynamic Random Walks. ACM SIGMOD: ACM SIGMOD International Conference on Management of Data 2023.
- C16) Zhaomin Wu*, Qinbin Li*, **Bingsheng He**. A Coupled Design of Exploiting Record Similarity for Practical Vertical Federated Learning. NeurIPS: Thirty-sixth Conference on Neural Information Processing Systems (NeurIPS 2022).
- C17) Sihao Hu*, Zhen Zhang[^], Shengliang Lu*, **Bingsheng He**, Zhao Li. Sequence-Based Target Coin Prediction for Cryptocurrency Pump-and-Dump. ACM SIGMOD: ACM SIGMOD International Conference on Management of Data 2023.
- C18) Xinyu Chen*, Yao Chen[^], Feng Cheng*, Hongshi Tan*, Weng-Fai Wong, **Bingsheng He**. ReGraph: Scaling Graph Processing on HBM-enabled FPGAs with Heterogeneous Pipelines. MICRO 2022: 55th ACM/IEEE International Symposium on Microarchitecture (MICRO 2022).

- C19) Qingyu Xu, Feng Zhang, Zhiming Yao, Lv Lu, Xiaoyong Du, Dong Deng, **Bingsheng He**. Efficient Load-Balanced Butterfly Counting on GPU. VLDB: International Conference on Very Large Data Bases (VLDB) 2022. (also published in Proc. VLDB Endow. Vol. 15 (2022)).
- C20) Shixuan Sun[^], Xibo Sun, Qiong Luo, **Bingsheng He**. RapidFlow: An Efficient Approach to Continuous Subgraph Matching. VLDB: International Conference on Very Large Data Bases (VLDB) 2022. (also published in Proc. VLDB Endow. Vol. 15 (2022)).
- C21) Zining Zhang*, **Bingsheng He**, Zhenjie Zhang. HARL: Hierarchical Adaptive Reinforcement Learning Based Auto Scheduler for Neural Networks. ICPP: International Conference on Parallel Processing 2022, 84/311=27%.
- C22) Johan Kok Zhi Kang*, Suwei Yang, Suriya Venkatesan, Sien Yi Tan, Feng Cheng, and **Bingsheng He**. Dynamic Graph Segmentation for Deep Graph Neural Networks. ACM KDD: In Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD 2022).
- C23) Xibo Sun, Shixuan Sun[^], Qiong Luo, **Bingsheng He**. An In-Depth Study of Continuous Subgraph Matching. VLDB: International Conference on Very Large Data Bases (VLDB) 2022. (also published in Proc. VLDB Endow. Vol. 15 (2022)).
- C24) Amelie Chi Zhou, Juanyun Luo, Ruibo Qiu, Haobin Tan, **Bingsheng He**, Rui Mao. Adaptive Partitioning for Large-Scale Graph Analytics in Geo-Distributed Data Centers. IEEE ICDE: IEEE International Conference on Data Engineering 2022.
- C25) Qinbin Li*, Yiqun Diao*, Quan Chen, **Bingsheng He**. Federated Learning on Non-IID Data Silos: An Experimental Study. IEEE ICDE: IEEE International Conference on Data Engineering 2022.
- C26) Jin Zhao, Yu Zhang, Xiaofei Liao, Ligang He, **Bingsheng He**, Hai Jin, Haikun Liu. LCCG: A Locality-Centric Hardware Accelerator for High Throughput of Concurrent Graph Processing. IEEE/ACM SC: International Conference for High Performance Computing, Networking, Storage and Analysis 2021.
- C27) Zeyi Wen, Zhishang Zhou, Hanfeng Liu, **Bingsheng He**, Xia Li, Jian Chen. Enhancing SVMs with Problem Context Aware Pipeline. ACM KDD: In Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '21). <https://doi.org/10.1145/3447548.3467291>.
- C28) Shixuan Sun[^], Yuhang Chen*, Shengliang Lu*, **Bingsheng He**, Yuchen Li: ThunderRW: An In-Memory Graph Random Walk Engine. VLDB: International Conference on Very Large Data Bases (VLDB) 2021. (also published in Proc. VLDB Endow. 14(11): 1992-2005 (2021)).
- C29) Hongshi Tan*, Xinyu Chen*, Yao Chen, **Bingsheng He**, Weng-Fai Wong. ThundeRiNG: Generating Multiple Independent Random Number Sequences on FPGAs. ACM ICS: International Conference on Supercomputing 2021.
- C30) Qinbin Li*, **Bingsheng He**, Dawn Song. Practical One-Shot Federated Learning for Cross-Silo Setting. IJCAI: 30th International Joint Conference on Artificial Intelligence, 2021.
- C31) Zeyi Wen[^], Qinbin Li*, **Bingsheng He**, Bin Cui. Challenges and Opportunities of Building Fast GBDT Systems. IJCAI: 30th International Joint Conference on Artificial Intelligence - Survey Track, 2021.
- C32) Zhuohui Duan, Haikun Liu, Haodi Lu, Xiaofei Liao, Hai Jin, Yu Zhang, **Bingsheng He**. Gengar: An RDMA-based Distributed Hybrid Memory Pool. IEEE ICDCS: International Conference on Distributed Computing Systems, 2021.
- C33) Johns Paul*, **Bingsheng He**, Shengliang Lu*, Chiew Tong Lau. MG-Join: A Scalable Join for Massively Parallel Multi-GPU Architectures. ACM SIGMOD: ACM SIGMOD International Conference on Management of Data 2021.
- C34) Shengliang Lu*, Shixuan Sun[^], Johns Paul*, Yuchen Li, **Bingsheng He**. Cache-Efficient Fork-Processing Patterns on Large Graphs. ACM SIGMOD: ACM SIGMOD International Conference on Management of Data 2021.
- C35) Shixuan Sun[^], Yuhang Chen*, **Bingsheng He**, Bryan Hooi. PathEnum: Towards Real-Time Hop-Constrained s-t Path Enumeration. ACM SIGMOD: ACM SIGMOD International Conference on Management of Data 2021.
- C36) Zhi Kang Johan Kok*, Gaurav, Sien Yi Tan, Feng Cheng, Shixuan Sun[^], **Bingsheng He**. Efficient Deep Learning Pipelines for Accurate Cost Estimations Over Large Scale Query

- Workload. ACM SIGMOD: ACM SIGMOD International Conference on Management of Data 2021.
- C37) Chang Ye, Yuchen Li, **Bingsheng He**, Zhao Li, Jianling Sun. GPU-Accelerated Graph Label Propagation for Real-Time Fraud Detection. ACM SIGMOD: ACM SIGMOD International Conference on Management of Data 2021.
- C38) Shuhao Zhang*, Yancan Mao, Jiong He, Philipp Marian Grulich, Steffen Zeuch, **Bingsheng He**, Richard Ma, Volker Markl. Parallelizing Intra-Window Join on Multicores: An Experimental Study. ACM SIGMOD: ACM SIGMOD International Conference on Management of Data 2021.
- C39) Qinbin Li*, **Bingsheng He**, Dawn Song. Model-Contrastive Federated Learning. CVPR: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR).
- C40) Xinyu Chen*, Hongshi Tan[^], Yao Chen, **Bingsheng He**, Weng-Fai Wong, Deming Chen. Skew-oblivious Data Routing for Data Intensive Applications on FPGAs with HLS. DAC: Design Automation Conference (DAC) 2021.
- C41) Zining Zhang*, **Bingsheng He**, Zhenjie Zhang. TransMask: A Compact and Fast Speech Separation Model Based on Transformer. IEEE ICASSP: The IEEE International Conference on Acoustics, Speech, & Signal Processing (ICASSP) 2021.
- C42) Qiuyi Lyu*, Yuchen Li, **Bingsheng He**, Bin Gong. DBL: Efficient Reachability Queries on dynamic Graphs. DASFAA: International Conference on Database Systems for Advanced Applications (DASFAA 2021).
- C43) Yu Zhang, Xiaofei Liao, Hai Jin, Ligang He, **Bingsheng He**, Haikun Liu, Lin Gu. DepGraph: A Dependency-Driven Accelerator for Efficient Iterative Graph Processing. IEEE HPCA: IEEE International Symposium on High Performance Computer Architecture, 2021.
- C44) Xinyu Chen*, Hongshi Tan[^], Yao Chen, **Bingsheng He**, Weng-Fai Wong, Deming Chen. ThunderGP: HLS-based Graph Processing Framework on FPGAs. ACM FPGA: ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA 2021). [**Top quality papers** in FPGA 2021]
- C45) Cheng Chen, Jun Yang, Mian Lu, Taize Wang, Zhao Zheng, Yuqiang Chen, Wenyuan Dai, **Bingsheng He**, Weng-Fai Wong, Guoan Wu, Yuping Zhao, Andy Rudoff. Optimizing In-memory Database Engine For AI-powered On-line Decision Augmentation Using Persistent Memory. VLDB: International Conference on Very Large Data Bases (VLDB) 2021. (also published in Proceedings of the VLDB Endowment, Volume 14 Issue 5, 2020).
- C46) Shengliang Lu*, **Bingsheng He**, Yuchen Li, Hao Fu*. Accelerating Exact Constrained Shortest Paths on GPUs. VLDB: International Conference on Very Large Data Bases (VLDB) 2021. (also published in Proceedings of the VLDB Endowment, Volume 14 Issue 2, 2020).
- C47) Shixuan Sun[^], Xibo Sun, Yulin Che, Qiong Luo, **Bingsheng He**. RapidMatch: A Holistic Approach to Subgraph Query Processing. VLDB: International Conference on Very Large Data Bases (VLDB) 2021. (also published in Proceedings of the VLDB Endowment, Volume 14 Issue 2, 2020).
- C48) Johns Paul*, **Bingsheng He**, Shengliang Lu*, Chiew Tong Lau. Improving Execution Efficiency of Just-in-time Compilation based Query Processing on GPUs. VLDB: International Conference on Very Large Data Bases (VLDB) 2021. (also published in Proceedings of the VLDB Endowment, Volume 14 Issue 2, 2020).
- C49) Bo Liu, Wenbin Jiang, Shaofeng Zhao, Hai Jin, **Bingsheng He**. GradSA: Gradient Sparsification and Accumulation for Communication-Efficient Distributed Deep Learning. GPC: International Conference on Green, Pervasive, and Cloud Computing 2020, pp 77-91.
- C50) Haikun Liu, Yuanyuan Ye, Xiaofei Liao, Hai Jin, Yu Zhang, Wenbin Jiang, **Bingsheng He**. Space-oblivious compression and wear leveling for non-volatile main memories. MSST: the 36th International Conference on Massive Storage Systems and Technology, 2020.
- C51) Tao Luo, Wei Zhang, **Bingsheng He**, Cheng Liu, Douglas Maskell. Energy Efficient In-memory Integer Multiplication Based on Racetrack Memory. EAI: Workshop on Efficient Artificial Intelligence For Edge Computing 2020, affiliated with IEEE ICDCS 2020.
- C52) Zining Zhang*, Zhenjie Zhang, **Bingsheng He**. X-TaSNet: Robust and Accurate Time-Domain Speaker Extraction Network. INTERSPEECH: Interspeech 2020 (Oral).
- C53) Zining Zhang*, Zhenjie Zhang, **Bingsheng He**. GAZEV: GAN-Based Zero Shot Voice Conversion over Non-parallel Speech Corpus. INTERSPEECH: Interspeech 2020 (Oral).

- C54) Johns Paul*, Jie Liang Ang*, Tianyuan Fu*, **Bingsheng He**, Shengliang Lu*, Sien Yi Tan, Feng Cheng. Poet: an Interactive Spatial Query Processing System in Grab. ACM SIGSPATIAL: 28th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems 2020 (Industrial Paper).
- C55) Xuanhua Shi, Zezhao Feng, Kaixi Li, Yongluan Zhou, Hai Jin, Yan Jiang, **Bingsheng He**, Zhijun Ling, Xin Li. ByteSeries : An In-Memory Database for Large-Scale Monitoring Systems. ACM SOCC: ACM Symposium on Cloud Computing 2020 (industry paper).
- C56) Wentian Guo, Yuchen Li, Mo Sha, **Bingsheng He**, Xiaokui Xiao, Kian-Lee Tan. GPU-Accelerated Subgraph Enumeration on Partitioned Graphs. ACM SIGMOD: ACM SIGMOD International Conference on Management of data, 2020.
- C57) Jiya Su, Feng Zhang, Weifeng Liu, **Bingsheng He**, Ruofan Wu, Xiaoyong Du, Rujia Wang. CapelliniSpTRSV: A Thread-Level Synchronization-Free Sparse Triangular Solve on GPUs. ICPP: International Conference on Parallel Processing 2020, 78/269=29%. **[Best paper candidate]**
- C58) Feng Zhang, Ningxuan Feng, Yani Liu, Cheng Yang, Jidong Zhai, Shuhao Zhang, **Bingsheng He**, Jiazao Lin, Xiaoyong Du. PewLSTM: Periodic LSTM with Weather-Aware Gating Mechanism for Parking Behavior Prediction. IJCAI: IJCAI-PRICAI 2020 (610/4806=12.7%)
- C59) Feng Zhang, Lin Yang, Shuhao Zhang, **Bingsheng He**, Wei Lu, Xiaoyong Du. FineStream: Fine-Grained Window-Based Stream Processing on CPU-GPU Integrated Architectures. USENIX ATC: USENIX Annual Technical Conference 2020 (65/348=18.7%).
- C60) Li Wang, Zining Zhang*, Zhenjie Zhang, **Bingsheng He**. PA-Tree: Polled-Mode Asynchronous B+-Tree for NVMe. IEEE ICDE: IEEE International Conference on Data Engineering 2020.
- C61) Shuhao Zhang*, Yingjun Wu, Feng Zhang, **Bingsheng He**. Towards Concurrent Stateful Stream Processing on Multicore Processors. IEEE ICDE: IEEE International Conference on Data Engineering 2020.
- C62) Xuanhua Shi, Yipeng Zhang, Hong Huang, Zhenyu Hu, Hai Jin, Huan Shen, Yongluan Zhou, **Bingsheng He**, Ruiibo Li, Keyong Zhou. Maxson: Reduce duplicate Parsing Overhead on Raw Data. IEEE ICDE: IEEE International Conference on Data Engineering 2020 (Industry and Application track, acceptance rate: 26%).
- C63) Qinbin Li*, Zeyi Wen^, **Bingsheng He**. Practical Federated Gradient Boosting Decision Trees. AAAI: Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI-2020) (1,591/7,737=20.6%). **[PREMIA Best Student Paper Gold Award 2021]**
- C64) Qinbin Li*, Zhaomin Wu*, Zeyi Wen^, **Bingsheng He**. Privacy-Preserving Gradient Boosting Decision Trees. AAAI: Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI-2020) (1,591/7,737=20.6%).
- C65) Xinyu Chen*, Ronak Bajaj^, Yao Chen, Jiong He, **Bingsheng He**, Weng-Fai Wong, Deming Chen. Is FPGA useful for hash joins. CIDR: Conference on Innovative Data Systems Research 2020.
- C66) Xuntao Cheng*, **Bingsheng He**, Eric Lo, Wei Wang, Shengliang Lu*, Xinyu Chen*. Deploying Hash Tables on Die-Stacked High Bandwidth Memory. ACM CIKM: ACM International Conference on Information and Knowledge Management 2019 (200/1030=20%).
- C67) Jin Zhao, Yu Zhang, Xiaofei Liao, Ligang He, **Bingsheng He**, Hai Jin, Haikun Liu, YiCheng Chen. GraphM: An Efficient Storage System for High Throughput of Concurrent Graph Processing. IEEE/ACM SC: International Conference for High Performance Computing, Networking, Storage and Analysis 2019.
- C68) Amelie Chi Zhou, Yao Xiao, **Bingsheng He**, Shadi Ibrahim, Reynold Cheng. Incorporating Probabilistic Optimizations for Resource Provisioning of Data Processing Workflows. ICPP: International Conference on Parallel Processing 2019.
- C69) Xinyu Chen*, Ronak Bajaj^, Yao Chen, Jiong He, **Bingsheng He**, Weng-Fai Wong, Deming Chen. On-The-Fly Parallel Data Shuffling for Graph Processing on OpenCL based FPGA. FPL: International Conference on Field Programmable Logic and Applications, 2019 (28/151).
- C70) Shuhao Zhang*, Jiong He, Amelie Chi Zhou, **Bingsheng He**. BriskStream: Scaling Data Stream Processing on Shared-Memory Multicores. ACM SIGMOD: ACM SIGMOD International Conference on Management of data, 2019.

- C71) Yu Zhang, Xiaofei Liao, Hai Jin, **Bingsheng He**, Haikun Liu, Lin Gu. DiGraph: An Efficient Path-based Iterative Directed Graph Processing System on Multiple GPUs. ACM ASPLOS: The 24th ACM International Conference on Architectural Support for Programming Languages and Operating Systems 2019.
- C72) Qinbin Li*, **Bingsheng He**, Zeyi Wen. Federated Learning Systems: Vision, Hype and Reality for Data Privacy and Protection. FML: 1st International Workshop on Federated Machine Learning for User Privacy and Data Confidentiality (FML 2019), in conjunction with IJCAI 2019.
- C73) Johns Paul*, **Bingsheng He**, Shengliang Lu*, Chiew Tong Lau. Revisiting Hash Join on Graphics Processors: A Decade Later. HardDB: International Workshop on Big Data Management on Emerging Hardware / Workshop on Data Management on Virtualized Active Systems (HardDB & Active), affiliated with ICDE 2019. [invited to DAPD]
- C74) Johns Paul*, **Bingsheng He**, Chiew Tong Lau. Query Processing on OpenCL-based FPGAs: Challenges and Opportunities. IEEE ICPADS: IEEE 24th International Conference on Parallel and Distributed Systems, 2018.
- C75) Dawen Xu, Kaijie Tu, Ying Wang, Cheng Liu, **Bingsheng He**. FCN-Engine: Accelerating Deconvolutional Layers in Classic CNN Processors. IEEE/ACM ICCAD: In Proceedings of the International Conference on Computer-Aided Design (ICCAD '18). Association for Computing Machinery, New York, NY, USA, Article 22, 1–6, 2018.
- C76) Pengcheng Yao, Long Zheng, Xiaofei Liao, Hai Jin, **Bingsheng He**. An Efficient Graph Accelerator with Parallel Data Conflict Management. IEEE/ACM PACT: The 27th International Conference on Parallel Architectures and Compilation Techniques, 2018.
- C77) Long Zheng, Xiaofei Liao, Hai Jin, **Bingsheng He**, Jingling Xue, Haikun Liu. Towards Concurrency Race Debugging: An Integrated Approach of Constraint Solving and Dynamic Slicing. IEEE/ACM PACT: The 27th International Conference on Parallel Architectures and Compilation Techniques, 2018. [Best paper candidates]
- C78) Yu Zhang, Xiaofei Liao, Hai Jin, and Lin Gu, Ligang He, **Bingsheng He**, Haikun Liu. CGraph: A Correlations-aware Approach for Efficient Concurrent Iterative Graph Processing. USENIX ATC: USENIX Annual Technical Conference 2018. [“Best storage related papers”, invited to ACM Transactions on Storage]
- C79) Hao Fu*, Shanjiang Tang, **Bingsheng He**, Ce Yu, Jizhou Sun. GLP4NN: A Convergence-invariant and Network-agnostic Light-Weight Parallelization Framework for Deep Neural Networks on Modern GPUs. ICPP: 47th International Conference on Parallel Processing 2018.
- C80) Amelie Chi Zhou, Tien-Dat Phan, Shadi Ibrahim, **Bingsheng He**. Energy-Efficient Speculative Execution using Advanced Reservation for Heterogeneous Clusters. ICPP: 47th International Conference on Parallel Processing 2018.
- C81) Jiacheng Ma, Xiao Zheng, Yaozu Dong, Wentai Li, Zhengwei Qi, **Bingsheng He**, Haibing Guan. gMig: Efficient GPU Live Migration Optimized by Software Dirty Page for Full Virtualization. ACM VEE: 14th ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments 2018.
- C82) Zeyi Wen[^], **Bingsheng He**, Ramamohanarao Kotagiri, Shengliang Lu*, Jiashuai Shi*. Efficient Gradient Boosted Decision Tree Training on GPUs. IEEE IPDPS: IEEE International Parallel & Distributed Processing Symposium, 2018.
- C83) Zeyi Wen[^], Xingyang Liu*, Hongjian Cao*, **Bingsheng He**. RTSI: An Index Structure for Multi-Modal Real-Time Search on Live Audio Streaming Services. IEEE ICDE: IEEE International Conference on Data Engineering (ICDE), 2018 (industrial and application track).
- C84) Zeke Wang[^], Kai Zhang[^], Haihang Zhou*, **Bingsheng He**. Hebe: An Order-obliviousness and High-performance Execution Scheme for Conjunctive Predicates. IEEE ICDE: IEEE International Conference on Data Engineering (ICDE), 2018 (short paper).
- C85) Kai Zhang[^], **Bingsheng He**, Jiayu Hu, Zeke Wang[^], Bei Hua, Jiayi Meng, Lishan Yang. G-NET: Effective GPU Sharing in NFV Systems. USENIX NSDI: USENIX Symposium on Networked Systems Design and Implementation 2018.
- C86) Xiongchao Tang, Jidong Zhai, Xuehai Qian, **Bingsheng He**, Wei Xue, Wenguang Chen. vSensor: Leveraging Fixed-Workload Modules of Programs for Performance Variance Detection. ACM PPOPP: The 23rd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPOPP 2018).

- C87) Mo Sha, Yuchen Li, **Bingsheng He**, Kian-Lee Tan. Accelerating Dynamic Graph Analytics on GPUs. VLDB: International Conference on Very Large Data Bases (VLDB) 2018. (also published in Proceedings of the VLDB Endowment, Volume 11 Issue 1, 2017, pages = {1—12}).
- C88) Xuntao Cheng*, **Bingsheng He**, Xiaoli Du*, Chiew Tong Lau. A Study of Main-Memory Hash Joins on Many-core Processor: A Case with Intel Knights Landing Architecture. ACM CIKM: ACM International Conference on Information and Knowledge Management 2017, Singapore (171/820=21%).
- C89) Jieru Zhao, Liang Feng, Wei Zhang, Sharad Sinha, Yun (Eric) Liang, **Bingsheng He**. COMBA: A Comprehensive Model-Based Analysis Framework for High Level Synthesis of Real Applications. IEEE/ACM ICCAD: 2017 International Conference On Computer Aided Design. **[2017 IEEE/ACM William J. McCalla ICCAD Best Paper Award (Front End)]**
- C90) Tao Luo*, Wei Zhang, **Bingsheng He**, Douglas Maskell. A Novel Two-stage Modular Multiplier Based on Racetrack Memory for Asymmetric Cryptography. IEEE/ACM ICCAD: 2017 International Conference On Computer Aided Design.
- C91) Amelie Chi Zhou*, Yifan Gong*, **Bingsheng He**, Jidong Zhai. Efficient Process Mapping in Geo-Distributed Cloud Data Centers. IEEE/ACM SC: International Conference for High Performance Computing, Networking, Storage and Analysis 2017.
- C92) Haikun Liu, Yujie Chen, Xiaofei Liao, Hai Jin, Long Zheng, **Bingsheng He**. Hardware/Software Cooperative Caching for Hybrid DRAM/NVM Memory Architectures. ACM ICS: 2017 International Conference on Supercomputing. [Acceptance rate: 28/117]
- C93) Amelie Chi Zhou*, Shadi Ibrahim, Bingsheng He. On Achieving Efficient Data Transfer for Graph Processing in Geo-Distributed Datacenters. IEEE ICDCS: International Conference on Distributed Computing Systems, 2017.
- C94) Kai Zhang[^], Jiayu Hu, **Bingsheng He**, Bei Hua. DIDO: Dynamic Pipelines for In-Memory Key-Value Stores on Coupled CPU-GPU Architectures. IEEE ICDE: IEEE International Conference on Data Engineering (ICDE), 2017.
- C95) Shuhao Zhang*, **Bingsheng He**, Daniel Dahlmeier, Chi Zhou, Thomas Heinze. Revisiting the Design of Data Stream Processing Systems on Multi-Core Processors. IEEE ICDE: IEEE International Conference on Data Engineering (ICDE), 2017.
- C96) Shuhao Zhang*, Hoang Tam Vo, Daniel Dahlmeier, **Bingsheng He**. Multi-Query Optimization for Complex Event Processing in SAP ESP. IEEE ICDE: IEEE International Conference on Data Engineering (ICDE), 2017 (industrial track).
- C97) Zhaojie Niu*, **Bingsheng He**, Amelie Chi Zhou, Chiew Tong Lau. Multi-Objective Optimizations in Geo-Distributed Data Analytics Systems. IEEE ICPADS: IEEE 23rd International Conference on Parallel and Distributed Systems 2017.
- C98) Fubing Mao*, Wei Zhang, **Bingsheng He**, Siew-Kei Lam. Dynamic Module Partitioning for Library based Placement on Heterogeneous FPGAs. IEEE RTCSA: IEEE International Conference on Embedded and Real-Time Computing Systems and Applications 2017 (short paper).
- C99) Feng Zhang, Bo Wu, Jidong Zhai, **Bingsheng He**, Wenguang Chen. FinePar: Irregularity-Aware Fine-Grained Workload Partitioning on Integrated Architectures. IEEE/ACM CGO: 2016 International Symposium on Code Generation and Optimization. [Acceptance rate: 26/114=22%]
- C100) Shanjiang Tang[^], **Bingsheng He**, Shuhao Zhang*, Zhaojie Niu*. Elastic Multi-Resource Fairness: Balancing Fairness and Efficiency in Coupled CPU-GPU Architectures. IEEE/ACM SC: International Conference for High Performance Computing, Networking, Storage and Analysis 2016.
- C101) Zeke Wang[^], Johns Paul*, Hui Yan Cheah[^], **Bingsheng He**, Wei Zhang. Accelerating Database Query Processing on OpenCL-based FPGAs. FPL: International Conference on Field Programmable Logic and Applications, 2016.
- C102) Mochi Xue, Kun Tian, Yaozu Dong, Jiajun Wang, Zhengwei Qi, **Bingsheng He**, Haibing Guan. gScale: Scaling up GPU Virtualization with Dynamic Sharing of Graphics Memory Space. USENIX ATC: USENIX Annual Technical Conference 2016.
- C103) Paul Johns*, Jiong He*, **Bingsheng He**. GPL: A GPU-based Pipelined Query Processing Engine. ACM SIGMOD: ACM SIGMOD International Conference on Management of data, 2016.

- C104) Shuang Chen*, Shunning Jiang*, **Bingsheng He**, Xueyan Tang. A Study of Sorting Algorithms on Approximate Memory. ACM SIGMOD: ACM SIGMOD International Conference on Management of data, 2016.
- C105) Yun Peng, Byron Choi, **Bingsheng He**, Suigen Zhou, R. Xu, Xiaohui Yu. VColor: A Practical Vertex-cut Based Approach for Coloring Large Graphs. IEEE ICDE: IEEE International Conference on Data Engineering (ICDE), 2016.
- C106) Zeke Wang[^], **Bingsheng He**, Wei Zhang, Shunning Jiang. A Performance Analysis Framework for Optimizing OpenCL Applications on FPGAs. IEEE HPCA: IEEE International Symposium on High Performance Computer Architecture, 2016 [53/240=22%].
- C107) Yingnan Cui*, Wei Zhang, **Bingsheng He**. A Discrete Thermal Controller for Chip-Multiprocessors. DATE: Design, Automation & Test in Europe Conference & Exhibition 2016 [full paper, acceptance ratio=24%].
- C108) Zhaojie Niu*, **Bingsheng He**. A Study of Big Data Computing Platforms: Fairness and Energy Consumption. IEEE IC2E: IEEE International Conference on Cloud Engineering (IC2E 2016 Doctoral Symposium).
- C109) Fubing Mao*, Wei Zhang, Bo Feng, **Bingsheng He**, Yuchun Ma. Modular Placement for Interposer based Multi-FPGA Systems. ACM GLSVLSI: ACM Great Lakes Symposium on VLSI 2016.
- C110) Zhaojie Niu*, **Bingsheng He**, Fangming Liu. Not All Joules are Equal: Towards Energy-Efficient and Green-Aware Data Processing Frameworks. IEEE IC2E: IEEE International Conference on Cloud Engineering [full paper, 16 out of 73=23%]. **[Best Paper Runner Up of IC2E2016]**
- C111) Tao Luo*, Wei Zhang, **Bingsheng He**, Douglas Maskell. A Racetrack Memory Based In-memory Booth Multiplier for Cryptography Applications. ASP-DAC: the 21st Asia and South Pacific Design Automation Conference, 2016.
- C112) Yifan Gong*, **Bingsheng He**, Amelie Chi Zhou*. Monetary Cost Optimizations for HPC Applications on Amazon Clouds: Checkpoints and Replicated Execution. IEEE/ACM SC: International Conference for High Performance Computing, Networking, Storage and Analysis 2015.
- C113) Zeke Wang[^], Bingsheng He, Wei Zhang. A Study of Data Partitioning on OpenCL-based FPGA. FPL: International Conference on Field Programmable Logic and Applications. **[Top-quality papers of FPL 2015]**
- C114) Amelie Chi Zhou*, **Bingsheng He**, Xuntao Cheng*, Chiew Tong Lau. A Declarative Optimization Engine for Resource Provisioning of Scientific Workflows in IaaS Clouds. ACM HPDC: ACM International Symposium on High-Performance Parallel and Distributed Computing, 2015. [Acceptance rate: 16%, 19 out of 116]
- C115) Shen Gao*, **Bingsheng He**, Jianliang Xu. Real-Time In-Memory Checkpointing for Future Hybrid Memory Systems. ACM ICS: 2015 International Conference on Supercomputing. [Acceptance rate: 20%]
- C116) Jun Yang, Qingsong Wei, Cheng Chen, Chundong Wang, Khai Leong Yong, **Bingsheng He**. NV-Tree: Reducing Consistency Cost for NVM-based Single Level Systems. USENIX FAST: 13th USENIX Conference on File and Storage Technologies 2015. [Acceptance rate: 21%]
- C117) Saurabh Jha*, **Bingsheng He**, Mian Lu, Xuntao Cheng*, Huynh Phung Huynh. Improving Main Memory Hash Joins on Intel Xeon Phi Processors: An Experimental Approach. VLDB: International Conference on Very Large Data Bases (VLDB) 2015. (also published in Proceedings of the VLDB Endowment, Volume 8 Issue 6, 2015, pages = {1—12}).
- C118) Long Zheng*, Xiaofei Liao, **Bingsheng He**, Song Wu, Hai Jin. On Performance Debugging of Unnecessary Lock Contentions on Multicore Processors: A Replay-based Approach. IEEE/ACM CGO: 2015 International Symposium on Code Generation and Optimization. [Acceptance rate: 27%]
- C119) Jiong He*, Shuhao Zhang*, **Bingsheng He**. In-Cache Query Co-Processing on Coupled CPU-GPU Architectures. VLDB: International Conference on Very Large Data Bases (VLDB) 2015. (also published in Proceedings of the VLDB Endowment, Volume 8 Issue 4, 2014, pages = {1—12}).

- C120) Xuntao Cheng*, **Bingsheng He**, Chiew Tong Lau. Energy-Efficient Query Processing on Embedded CPU-GPU Architectures. DaMoN: ACM SIGMOD International Workshop on Data Management on New Hardware 2015.
- C121) Zhaojie Niu*, Shanjiang Tang^, **Bingsheng He**. Gemini: An Adaptive Performance-Fairness Scheduler for Data-Intensive Cluster Computing. IEEE CloudCom: IEEE International Conference on Cloud Computing Technology and Science, 2015 (full paper).
- C122) Hao Liang, Yi-Chung Chen, Tao Luo*, Wei Zhang, Hai Li, **Bingsheng He**. Hierarchical Library Based Power Estimator for Versatile FPGAs. IEEE MCSoc: IEEE 9th International Symposium on Embedded Multicore/Many-core Systems-on-Chip, 2015.
- C123) Feng Zhang, Jidong Zhai, Wenguang Chen, **Bingsheng He**, Shuhao Zhang*. To Co-Run, or Not To Co-Run: A Performance Study on Integrated Architectures. MASCOTS: IEEE 23rd International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems, 2015 (short paper, 4 pages).
- C124) Qingchao Cai, Rajesh Vellore Arumugam, Yew Soon Ong, **Bingsheng He**. Understanding the Behavior of Solid State Disk. IES: 18th Asia Pacific Symposium on Intelligent and Evolutionary Systems, 2015.
- C125) Ha-Nguyen Tran, Jung-jae Kim, **Bingsheng He**. Fast Subgraph Matching on Large Graphs using Graphics Processors. DASFAA: International Conference on Database Systems for Advanced Applications, 2015.
- C126) Yifan Gong*, **Bingsheng He**, Dan Li. Finding Constant From Change: Revisiting Network Performance Aware Optimizations on IaaS Clouds. IEEE/ACM SC: International Conference for High Performance Computing, Networking, Storage and Analysis, 2014. [Acceptance rate: 20%]
- C127) Haikun Liu^, **Bingsheng He**. Reciprocal Resource Fairness: Towards Cooperative Multiple-Resource Fair Sharing in IaaS Clouds. IEEE/ACM SC: International Conference for High Performance Computing, Networking, Storage and Analysis, 2014. [Acceptance rate: 20%]
- C128) Shanjiang Tang, Bu Sung Lee, Bingsheng He and Haikun Liu^. Long-Term Resource Fairness: Towards Economic Fairness on Pay-as-you-use Computing Systems. ACM ICS: 2014 International Conference on Supercomputing. [Acceptance rate: 20%]
- C129) Zigang Zhang, Yinliang Yue, Bingsheng He, Jin Xiong, Mingyu Chen, Lixin Zhang, Ninghui Sun. Pipelined Compaction for the LSM-tree. IEEE IPDPS: 28th IEEE International Parallel & Distributed Processing Symposium, 2014. [Acceptance rate: 21%]
- C130) Amelie Chi Zhou*, **Bingsheng He**. Simplified Resource Provisioning for Workflows in IaaS Clouds. IEEE CloudCom: IEEE International Conference on Cloud Computing Technology and Science, 2014 (Ph.D. Consortium). **[Best Ph.D. Consortium Award]**
- C131) Shanjiang Tang^, Bu Sung Lee, **Bingsheng He**. Towards Economic Fairness for Big Data Processing in Pay-as-you-go Cloud Computing. IEEE CloudCom: IEEE International Conference on Cloud Computing Technology and Science, 2014 (Ph.D. Consortium).
- C132) Yingnan Cui*, Wei Zhang, Vivek Chaturvedi, Weichen Liu, **Bingsheng He**. Thermal-aware Task Scheduling for 3D-Network-on-Chip: A Bottom to Top Scheme. IEEE ISIC: 12th International Symposium on Integrated Circuits, 2014.
- C133) Wan Du, Zikun Xing, Mo Li, **Bingsheng He**, Lloyd Hock Chye Chua, Haiyan Miao. Optimal Sensor Placement and Measurement of Wind for Urban Ecological Studies. IPSN: ACM/IEEE International Conference on Information Processing in Sensor Networks 2014 [Acceptance rate: 20%].
- C134) **Bingsheng He**. When Data Management Systems Meet Approximate Hardware: Challenges and Opportunities. VLDB: Proceedings of the VLDB Endowment, Volume 8 Issue 10, 2014, pages = {877--880}. (also to present in International Conference on Very Large Data Bases (VLDB) 2014).
- C135) Lei Wei*, **Bingsheng He**, Chuan Heng Foh. Towards Multi-Resource Physical Machine Provisioning for IaaS Clouds. IEEE ICC: IEEE International Conference on Communications 2014.
- C136) Tuo He, Maode Ma, Wenping Ma, **Bingsheng He**. A Novel Authenticated Multi-party Key Agreement for Private Cloud. IEEE ICC: IEEE International Conference on Communications 2014.

- C137) Jianlong Zhong*, **Bingsheng He**. Towards GPU-Accelerated Large-Scale Graph Processing in the Cloud. IEEE CloudCom: IEEE 5th International Conference on Cloud Computing Technology and Science 2013, 8 pages. [Acceptance rate: 17%]
- C138) Jiong He*, Mian Lu, **Bingsheng He**. Revisiting Co-Processing for Hash Joins on the Coupled CPU-GPU Architecture. VLDB: International Conference on Very Large Data Bases (VLDB) 2013. (also published in Proceedings of the VLDB Endowment, Volume 6 Issue 10, August 2013, pages = {1—12}).
- C139) Cheng Chen*, **Bingsheng He**, Xueyan Tang, Changbing Chen, Yubao Liu. Green Databases Through Integration of Renewable Energy. CIDR: 6th Biennial Conference on Innovative Data Systems Research 2013.
- C140) Mian Lu, Lei Zhang, Huynh Phung Huynh, Zhongliang Ong, Yun Liang, **Bingsheng He**, Rick Siow Mong Goh, Richard Huynh. Optimizing the MapReduce Framework on Intel Xeon Phi Coprocessor. IEEE BigData: IEEE Big Data 2013 (short paper).
- C141) Shanjiang Tang, Bu Sung Lee, **Bingsheng He**. Dynamic Slot Allocation Technique for MapReduce Clusters. IEEE Cluster: 2013 IEEE International Conference on Cluster Computing (CLUSTER), Indianapolis, IN, USA, 2013, pp. 1-8.
- C142) Jiangming Jin, Stephen John Turner, Bu-Sung Lee, Jianlong Zhong, **Bingsheng He**. Simulation of Information Propagation over Complex Networks: Performance Studies on Multi-GPU. IEEE/ACM DS-RT: 2013 IEEE/ACM 17th International Symposium on Distributed Simulation and Real Time Applications, Delft, Netherlands, 2013, pp. 179-188.
- C143) Jiangming Jin, Stephen John Turner, Bu-Sung Lee, Jianlong Zhong, **Bingsheng He**. Simulation Studies of Viral Advertisement Diffusion On Multi-GPU. WSC: Proceedings of the 2013 Winter Simulation Conference.
- C144) Lu Zhang^, Xueyan Tang, **Bingsheng He**. On Minimum Interaction Time for Continuous Distributed Interactive Computing (brief announcement). ACM PODC: ACM Symposium on Principles of Distributed Computing, 2013.
- C145) Shanjiang Tang, Bu Sung Lee, **Bingsheng He**. MROrder: Flexible Job Ordering Optimization for Online MapReduce Workloads. Euro-Par: 19th International European Conference on Parallel and Distributed Computing, 2013.
- C146) Changbing Chen*, **Bingsheng He**. A Framework for Analyzing Monetary Cost of Database Systems in the Cloud (Invited Paper). WAIM: The 14th International Conference on Web-Age Information Management, 2013.
- C147) Liyan Song, Yun Peng, Byron Choi, Jianliang Xu, **Bingsheng He**. Spectral Decomposition for Optimal Graph Index Prediction. PAKDD: 17th Pacific-Asia Conference on Knowledge Discovery and Data Mining, 2013.
- C148) Rishan Chen*, Xuetian Weng*, **Bingsheng He**, Mao Yang, Byron Choi, Xiaoming Li. Improving Large Graph Processing on Partitioned Graphs in the Cloud. ACM SOCC: ACM Symposium on Cloud Computing 2012 (21 out of 163). [Acceptance rate: 13%]
- C149) Changbing Chen*, **Bingsheng He**, Xueyan Tang. Green-Aware Workload Scheduling in Geographically Distributed Data Centers. IEEE CloudCom: IEEE 4th International Conference on Cloud Computing Technology and Science 2012, Pages 82-89. [Acceptance rate: 17%]
- C150) Donghong Wu*, **Bingsheng He**, Xueyan Tang, Jianliang Xu, Minyi Guo. RAMZzz: Rank-Aware DRAM Power Management with Dynamic Migrations and Demotions. IEEE/ACM SC: International Conference for High Performance Computing, Networking, Storage and Analysis 2012, Pages 1-11. [Acceptance rate: 20%]
- C151) Jiangming Jin, Stephen John Turner, Bu-Sung Lee, Jianlong Zhong, **Bingsheng He**. HPC Simulations of Information Propagation over Social Networks. ICCS: International Conference on Computational Science 2012.
- C152) Shadi Ibrahim, Hai Jin, Lu Lu, **Bingsheng He**, Gabriel Antoniu, Song Wu. Maestro: Replica-Aware Map Scheduling for MapReduce. IEEE/ACM CCGrid: IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, 2012.
- C153) Yu Shyang Tan, Bu Sung Lee, Roy Campbell, **Bingsheng He**. A Map-Reduce Based Framework for Heterogeneous Processing Element Cluster Environments. IEEE/ACM CCGrid: IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, 2012.

- C154) Shanjiang Tang, Bu Sung Lee, **Bingsheng He**. Speedup for Multi-Level Parallel Computing. HIPS: The 17th International Workshop on High-Level Parallel Programming Models and Supportive Environments, 2012.
- C155) Shadi Ibrahim, Hai Jin, Lu Lu, **Bingsheng He**, Song Wu. Adaptive I/O Scheduling for MapReduce in Virtualized Environment. ICPP: The 40th Annual International Conference on Parallel Processing, 2011.
- C156) Shadi Ibrahim*, **Bingsheng He**, Hai Jin. Towards Pay-As-You-Consume Cloud Computing. IEEE SCC: The 8th International Conference on Services Computing, 2011.
- C157) Shen Gao, Jianliang Xu, **Bingsheng He**, Byron Choi, Haibo Hu. PCMLogging: Reducing Transaction Logging Overhead with PCM. ACM CIKM: ACM International Conference on Information and Knowledge Management, 2011 (short paper, pages: 2401-2404).
- C158) **Bingsheng He**, Jeffrey Xu Yu. High-Throughput Transaction Executions on Graphics Processors. VLDB: International Conference on Very Large Data Bases (VLDB) 2011. (Also published in Proceedings of the VLDB Endowment, Volume 4 Issue 5, February 2011, pages = {314--325}).
- C159) Yanfei Lv, Bin Cui, **Bingsheng He**, Xuexuan Chen. Operation-Aware Buffer Management in Flash-based Systems. ACM SIGMOD: ACM SIGMOD International Conference on Management of data, pages: 13-24, 2011.
- C160) Jianlong Zhong*, **Bingsheng He**. GPU-Assisted Buffer Management. ICCS: International Conference on Computational Science, 2011.
- C161) Wenbin Fang, **Bingsheng He**, Qiong Luo. Database Compression on Graphics Processors. VLDB: In Proceedings of International Conference on Very Large Data Bases (VLDB) 2010. (Also published in Proceedings of the VLDB Endowment, Volume 3 Issue 1-2, September 2010, pages = {670--680}).
- C162) Yinan Li, **Bingsheng He**, Robin Jun Yang, Qiong Luo, Ke Yi. Tree Indexing on Solid State Drives. VLDB: In Proceedings of International Conference on Very Large Data Bases (VLDB) 2010. (Also published in Proceedings of the VLDB Endowment, Volume 3 Issue 1-2, February 2011, pages = {1195--1206}).
- C163) Hongyi Wang*, Qingfeng Jing*, Rishan Chen*, **Bingsheng He**, Zhengping Qian, Lidong Zhou. Distributed Systems Meet Economics: Pricing in the Cloud. HotCloud 2010: 2nd USENIX Workshop on Hot Topics in Cloud Computing 2010, pages = {1--6}.
- C164) **Bingsheng He**, Mao Yang, Zhenyu Guo, Rishan Chen, Wei Lin, Bing Su, Lidong Zhou. Comet: Batched Stream Processing for Data Intensive Distributed Computing. ACM SoCC: First ACM Symposium on Cloud Computing 2010, pages = {63--74}.
- C165) **Bingsheng He**, Mao Yang, Zhenyu Guo, Rishan Chen, Wei Lin, Bing Su, Hongyi Wang, Lidong Zhou. Wave Computing in the Cloud. HotOS: 12th Workshop on Hot Topics in Operating Systems 2009.
- C166) Shadi Ibrahim, Hai Jin, Lu Lu, Song Wu, **Bingsheng He**, Li Qi. LEEN: Locality/Fairness-aware key partitioning for MapReduce in the Cloud. IEEE CloudCom: In Proceedings of 2nd International Conference on Cloud Computing 2010.
- C167) Mian Lu, **Bingsheng He**, Qiong Luo. Supporting Extended Precision on Graphics Processors. DaMoN: ACM SIGMOD International Workshop on Data Management on New Hardware 2010, pages = {19--26}.
- C168) Sai Tung On*, Yinan Li, **Bingsheng He**, Ming Wu, Qiong Luo, Jianliang Xu. FD-Buffer: A Buffer Manager for Databases on Flash Disks. ACM CIKM: Proceedings of the 19th ACM international conference on Information and knowledge management, 2010, short paper, pages = {1297--1300}.
- C169) Yinan Li, **Bingsheng He**, Qiong Luo, Ke Yi. Tree Indexing on Flash Disks. IEEE ICDE: IEEE International Conference on Data Engineering (ICDE), pages: 1303 – 1306, 2009.
- C170) Hongyi Wang^, **Bingsheng He**, Vijayan Prabhakaran, Lidong Zhou. Crystal: The Power of Structure Against Corruptions. HotDep: Proceedings of the 5th workshop on Hot Topics in System Dependability 2009.
- C171) Wenbin Fang, Mian Lu, Xiangye Xiao, **Bingsheng He**, Qiong Luo. Frequent Itemset Mining on Graphics Processors. DaMoN: ACM SIGMOD Fifth International Workshop on Data Management on New Hardware 2009.

- C172) Linhong Zhu, Byron Choi, **Bingsheng He**, Jeffrey Xu Yu, Wee Keong Ng. A Uniform Framework for Ad-Hoc Indexes to Answer Reachability Queries on Large Graphs. DASFAA: International Conference on Database Systems for Advanced Applications 2009.
- C173) **Bingsheng He**, Wenbin Fang, Qiong Luo, Naga K. Govindaraju, Tuyong Wang. Mars: A MapReduce Framework on Graphics Processors. IEEE/ACM PACT: IEEE International Conference on Parallel Architecture, Compilation Techniques 2008.
- C174) **Bingsheng He**, Ke Yang, Rui Fang, Mian Lu, Naga K. Govindaraju, Qiong Luo, Pedro V. Sander. Relational Joins on Graphics Processors. ACM SIGMOD: ACM SIGMOD International Conference on Management of data, pages: 511-524, 2008. **[one of the two “Best papers” of SIGMOD 2008, invited to ACM TODS]**
- C175) **Bingsheng He**, Naga K. Govindaraju, Qiong Luo, Burton Smith. Efficient Gather and Scatter Operations on Graphics Processors. IEEE/ACM SC: International Conference for High Performance Computing, Networking, Storage and Analysis, 2007 (12 pages).
- C176) **Bingsheng He**, Qiong Luo. Cache-Oblivious Query Processing. CIDR: The 3rd Biennial Conference on Innovative Data Systems Research, 2007.
- C177) **Bingsheng He**, Yinan Li, Qiong Luo, Dongqing Yang. A General Framework for Improving Query Processing Performance on Multi-Level Memory Hierarchies. DaMoN: ACM SIGMOD Third International Workshop on Data Management on New Hardware 2007.
- C178) Ke Yang, **Bingsheng He**, Rui Fang, Mian Lu, Naga K. Govindaraju, Qiong Luo, Pedro Sander, Jiaoying Shi. In-Memory Grid Files on Graphics Processors. DaMoN: ACM SIGMOD Third International Workshop on Data Management on New Hardware 2007.
- C179) **Bingsheng He**, Qiong Luo. Cache-Oblivious Nested-Loop Joins. ACM CIKM: ACM Fifteenth Conference on Information and Knowledge Management 2006.
- C180) Zi Lin, **Bingsheng He**, Byron Choi. A Quantitative Summary of XML Structures. ER: 25th International Conference on Conceptual Modeling 2006.
- C181) **Bingsheng He**, Qiong Luo, Byron Choi. Cache-Conscious Automata for XML Filtering. IEEE ICDE: 21st International Conference on Data Engineering 2005.
- C182) Wenwei Xue, **Bingsheng He**, Hejun Wu, Qiong Luo. The HKUST Frog Pond - A Case Study of Sensory Data Analysis. BISON: Building Intelligent Sensor Networks Workshop (BISON 2004) in conjunction with IFIP NPC'04.
- C183) Hejun Wu, Qiong Luo, Pei Zheng, **Bingsheng He**, Lionel M. Ni. Accurate Emulation of Wireless Sensor Networks. BISON: Building Intelligent Sensor Networks Workshop (BISON 2004) in conjunction with IFIP NPC'04.
- C184) Qiong Luo, Lionel M. Ni, **Bingsheng He**, Hejun Wu, Wenwei Xue. MEADOWS: Modeling, Emulation, Analysis of Data of Wireless Sensor Networks. DMSN: International Workshop on Data Management for Sensor Networks (DMSN 2004) in conjunction with VLDB 2004.

D.4. System Demonstration Paper

- C185) Siqi Wang, Jiashu Li, Mian Lu, Zhao Zheng, Yuqiang Chen, **Bingsheng He**. A System for Time Series Feature Extraction in Federated Learning. ACM CIKM '22: Proceedings of the 31st ACM International Conference on Information & Knowledge Management.
- C186) Zhao Li, Pengcheng Zou, Xia Chen, Shichang Hu, Peng Zhang, Yumou Zhang, **Bingsheng He**, Yuchen Li, Xing Tang. From Community Search to Community Understanding: A Multimodal Community Query Engine. ACM CIKM: ACM International Conference on Information and Knowledge Management 2021 (demo paper) **[Best Demo Paper Runner-up Award]**
- C187) Husong Liu*, Shengliang Lu*, Xinyu Chen, **Bingsheng He**. G3: When Graph Neural Networks Meet Parallel Graph Processing Systems on GPUs. VLDB: Proceedings of the VLDB Endowment 2020, pages = {1—4}, system demonstration.
- C188) Jieliang Ang*, Tianyuan Fu*, Johns Paul*, Shuhao Zhang*, **Bingsheng He**, Teddy Sison David Wenceslao, Siyeni Tan. Trav: An Interactive Trajectory Exploration System. IEEE BigMM: IEEE International Conference on Multimedia Big Data 2019 (Demo).
- C189) Zeyi Wen^, Mingyu Liang*, **Bingsheng He**, Zexin Xia*, Bo Li*. Aucher: Multi-modal Queries on Live Audio Streams in Real-time. IEEE ICDE: IEEE International Conference on Data Engineering (ICDE), 2018 (demo papers).

- C190) Xuntao Cheng*, **Bingsheng He**, Mian Lu, Chiew Tong Lau, Huynh Phung Huynh, Siow Mong Goh. Efficient Query Processing on Many-core Architectures: A Case Study with Intel Xeon Phi Processor. ACM SIGMOD: In Proceedings of ACM SIGMOD Conference 2016 (demo papers, four pages).
- C191) Wan Du, Mo Li, Zikun Xing, **Bingsheng He**, Lloyd Hock Chye Chua, Zhenjiang Li, Yuanqiang Zheng, Pengfei Zhou. Demo Abstract: Wind Measurements for Water Quality Studies in Urban Reservoirs. IEEE SECON: 2014 Eleventh Annual IEEE International Conference on Sensing, Communication, and Networking (SECON), Singapore, 2014, pp. 161-163. **[Best Demo Award]**
- C192) Shuhao Zhang*, Jiong He*, **Bingsheng He**, Mian Lu. OmniDB: Towards Portable and Efficient Query Processing on Parallel CPU/GPU Architectures. VLDB: International Conference on Very Large Data Bases (VLDB) 2013. (also published in Proceedings of the VLDB Endowment, Volume 6 Issue 10, August 2013, pages = {1—4}, system demonstration).
- C193) Jianlong Zhong*, **Bingsheng He**. Parallel Graph Processing on Graphics Processors Made Easy. VLDB: International Conference on Very Large Data Bases (VLDB) 2013. (also published in Proceedings of the VLDB Endowment, Volume 6 Issue 10, August 2013, pages = {1—4}, system demonstration). **[“Best demos”, invited to SIGMOD RECORD]**
- C194) Rishan Chen*, Xuettian Weng*, **Bingsheng He**, Mao Yang. Large Graph Processing in the Cloud. ACM SIGMOD: In Proceedings of ACM SIGMOD Conference 2010, pages = {1123--1126}, (demo).
- C195) **Bingsheng He**, Yinan Li, Qiong Luo, Dongqing Yang. EaseDB: A Cache-Oblivious In-Memory Query Processor. ACM SIGMOD: In Proceedings of ACM SIGMOD Conference 2007 (system demonstration).
- C196) Rui Fang, **Bingsheng He**, Mian Lu, Ke Yang, Naga K. Govindaraju, Qiong Luo, Pedro Sander. GPUQP: Query Co-Processing using Graphics Processors. ACM SIGMOD: In Proceedings of ACM SIGMOD Conference 2007 (system demonstration).

D.5. Posters

- C197) Zeyi Wen[^], Jiashuai Shi*, **Bingsheng He**, Jian Chen, Yawen Chen. Efficient Multi-Class Probabilistic SVMs on GPUs. IEEE ICDE: The 35th proceedings of IEEE International Conference on Data Engineering (ICDE), 2019 (poster).
- C198) **Bingsheng He**. Data Management Systems on Future Hardware: Challenges and Opportunities (Extended Abstract for Invited talk in Joint Active and HardDB Workshops). IEEE ICDE: IEEE International Conference on Data Engineering (ICDE), 2017.
- C199) Zeke Wang[^], Hui Yan Cheah[^], Johns Paul*, **Bingsheng He**, Wei Zhang. Accelerating Database Query Processing on OpenCL-based FPGAs. ACM FPGA: ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (poster paper), 2016.
- C200) Zeke Wang[^], **Bingsheng He**, Wei Zhang. Improving Data Partitioning Performance on OpenCL-based FPGAs. IEEE International Symposium on Field-Programmable Custom Computing Machines (FCCM), May 2015 (poster paper).
- C201) Xuanhua Shi, Junling Liang, Sheng Di, **Bingsheng He**, Hai Jin, Lu Lu, Zhixiang Wang, Xuan Luo, Jianlong Zhong*. Optimization of Asynchronous Graph Processing on GPU with Hybrid Coloring Model. ACM PPOPP: Proceedings of ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, 2015, 2 pages (poster).
- C202) Yifan Gong*, Amelie Chi Zhou*, **Bingsheng He**. Monetary Cost Optimizations for HPC Applications on Amazon Clouds: Checkpoints and Replicated Execution. IEEE/ACM SC: ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis 2014 (2 pages, poster, acceptance rate: 39%).
- C203) Yifan Gong*, **Bingsheng He**, Jianlong Zhong. An Overview of CMPI: Network Performance Aware MPI in the Cloud. ACM PPOPP: Proceedings of 17th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, Pages 297-298, poster.
- C204) Jianlong Zhong*, **Bingsheng He**. An Overview of Medusa: Simplified Graph Processing on GPUs. ACM PPOPP: Proceedings of 17th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, Pages 283-284, poster.

- C205) Jianlong Zhong*, **Bingsheng He**. GViewer: GPU-Accelerated Graph Visualization and Mining. SocInfo: Proceedings of 3rd International Conference on Social Informatics 2011 (poster paper, 4 pages).
- C206) Ke Yang, **Bingsheng He**, Qiong Luo, Pedro Sander, Jiaoying Shi. Stack-Based Parallel Recursion on Graphics Processors. ACM PPOPP: 14th ACM SIGPLAN Symposium on Principles, Practice of Parallel Programming, poster paper, 2 pages, 2009.
- C207) Hongyu Guo, **Bingsheng He**, Yifan He, Qiong Luo, Bo Peng, Xiangye Xiao. Frequent Pattern Mining on Graphics Processors. ACM SIGMOD: ACM SIGMOD 2008 undergraduate poster competition.

D.6. Tutorials

- C208) **Bingsheng He**, Qinbin Li, Sixu Hu, Zhaomin Wu. A Tutorial on Federated Learning Systems: Comparative Studies and Hand-on Demonstrations. IJCAI-PRICAI 2020.
- C209) **Bingsheng He**, Huynh Phung Huynh, Rick GOH Siow Mong. GPGPU for Real-Time Data Analytics. Invited presentation in IEEE ICPADS: IEEE International Conference on Parallel and Distributed Systems 2012.
- C210) **Bingsheng He**, Huynh Phung Huynh, Rick GOH Siow Mong. GPGPU for Real-Time Data Analytics. Presented in ACM PPOPP: ACM Principles and Practice of Parallel Programming 2013.
- [All teaching materials (including slides and video) are online, web: <http://www3.ntu.edu.sg/home/bshe/GPGPUTut.html>, Number of visitors: 1,041 (from Dec 2012 to Feb 15, 2015)]**

D.7. Book Chapters

- C211) Johns Paul*, Shengliang Lu* and **Bingsheng He**. Database Systems on GPUs. Foundations and Trends in Databases (Now Publisher) 2021.
- C212) Shan Jiang Tang[^], **Bingsheng He**, Haikun Liu, Bu-sung Lee. Resource Management in Big Data Processing Systems. Accepted as a book chapter in the book "Big Data: Principles and Paradigms" (edited by Rajkumar Buyya;Rodrigo N. Calheiros;Amir Vahid Das, Academic Press, 2016)
- C213) Amelie Chi Zhou*, **Bingsheng He**, Shadi Ibrahim. eScience and Big Data Workflows in Clouds. Accepted as a book chapter in the book "Big Data: Principles and Paradigms" (edited by Rajkumar Buyya;Rodrigo N. Calheiros;Amir Vahid Das, Academic Press, 2016)
- C214) Rishan Chen*, Xuétian Weng*, **Bingsheng He**, Byron Choi, Mao Yang. Network Performance Aware Graph Partitioning for Large Graph Processing Systems in the Cloud. Accepted as a book chapter in the book "Large Scale and Big Data: Processing and Management" (edited by Sherif Sakr, Mohamed Gaber, CRC Press, 25 Jun 2014).
- C215) Baoxue Zhao, Jianlong Zhong*, **Bingsheng He**, Qiong Luo, Wenbin Fang and Naga K.Govindaraju. GPU-Accelerated Cloud Computing for Data-Intensive Applications. Accepted as a book chapter in the book "Cloud Computing for Data Intensive Applications" (edited by Xiaolin Li and Judy Qiu, Springer, 2014).
- C216) Qianqian Zhao, Maode Ma, Yuqing Zhang, **Bingsheng He**. Access Control in Cloud Computing. Emerging Research in Cloud Distributed Computing System, IGI Global (edited by Dr. Susmit Bagchi, 2014).

D.8. Interdisciplinary Research Publications

- C217) Xuntao Cheng*, Zhaojie Niu*, **Bingsheng He**. SilverWisdom: Towards a Knowledge Base for Elderly People. International Journal of Information Technology (IJIT), 2014 (Also presented in ICAA 2013: the 1st International Conference on Ageless Aging 2013).
- C218) **Bingsheng He**, Mo Li, Lloyd H. C. Chua, Cheng Liu, Zikun Xing, Wan Du, Lei Chen, Hans S. Eikaas. Clouds with benefits: Sensor networks and cloud computing allow real-time water quality monitoring in Singapore's reservoirs. In Innovation in Water – Singapore, Volume 4, 2013.

- C219) Zikun Xing, Cheng Liu, Lloyd H. C. Chua, **Bingsheng He**, Jörg Imberger. Impacts of variable wind forcing in urban reservoirs. ISEH: 7th International Symposium on Environmental Hydraulics 2014, Singapore, January 2014.
- C220) Cheng Liu[^], Zikun Xing, Chua Hock Chye Lloyd, **Bingsheng He**, Mo Li, Eikaas Hans. Cloud Assisted Water Quality Management in Singapore. SIWW: Singapore International Water Week 2012.
- C221) Zikun Xing, Cheng Liu, Lloyd H. C. Chua, **Bingsheng He**, Hans S. Eikaas. Improving Run Time in Three-Dimensional Reservoir Hydrodynamics and Water Quality Modeling. ICHE: Tenth International Conference on Hydrosience & Engineering, 2012.

E. Open-Source Systems

Most of our open-source projects are hosted in GitHub: <https://github.com/Xtra-Computing>. Here are the selected ones:

- 1) FedTree: a federated learning system for tree-based models. <https://github.com/Xtra-Computing/FedTree>. [By Qinbin Li et al. 50 stars in Github]
- 2) ThunderRING: Fast Multiple Independent Random Number Sequences Generation on FPGAs [By Hongshi Tan*]
- 3) ThunderRW: An In-Memory Graph Random Walk Engine [By Shixuan Sun*]
- 4) Vine: Accelerating Exact Constrained Shortest Paths on GPUs [By Shengliang Lu*]
- 5) ForkGraph: Cache-Efficient Fork-Processing Patterns on Large Graphs [By Shengliang Lu*]
- 6) NIID-Bench: a benchmark for federated learning algorithms under non-IID data distribution scenarios. [By Qinbin Li* and Yiqun Zhang*, 150+ stars and 40+ forks in Github]
- 7) The OARF Benchmark Suite: Characterization and Implications for Federated Learning Systems [By Sixu Hu*]
- 8) ThunderGP: Fast Graph Processing for HLS-based FPGAs. ThunderGP is featured at Xilinx Apps and Libraries (<https://www.xilinx.com/products/apps-and-libraries.html>). **[80+ stars in GitHub, the third place of 2020 Xilinx Adaptive Computing Developer Contest (under category of Adaptable Compute Acceleration, top 9 out of 79), “Top-quality papers” of FPGA 2021.]**[\[Feature in NUS SoC News\]](#)
- 9) G3: When Graph Neural Networks Meet Parallel Graph Processing Systems on GPUs [By Husong Liu*, Shengliang Lu*].
- 10) ThunderGBM: ThunderGBM: Fast GBDTs and Random Forests on GPUs. <https://github.com/Xtra-Computing/thundergbm>. [By Zeyi Wen[^], Jiashuai Shi*, Qinbin Li*] **[~6,00 stars in GitHub, highlighted in Reddit, 机器之心 with over 5000 repostings, TPDS2019 best paper award]**
- 11) Brickstream: scaling up the data stream systems. <https://github.com/Xtra-Computing/brickstream>. [By Shuhao Zhang*]
- 12) ThunderSVM: A Fast SVM Library on GPUs and CPUs. <https://github.com/Xtra-Computing/thundersvm>. [By Zeyi Wen[^], Jiashuai Shi*, Qinbin Li*] **[~1,400 stars in GitHub, highlighted by Hacker News, Packt DataHub, published in JMLR open source system issue]**
- 13) Mars: A MapReduce Framework on Graphics Processors. <http://www.cse.ust.hk/gpuqp/Mars.html>. [I have been the major developer when I was with HKUST]
- 14) FD-Tree: a Tree Index on Solid State Drives, <http://pages.cs.wisc.edu/~yinan/fdtree.html>. [I have been one of the major developers when I was with HKUST]
- 15) GPUQP: Query Co-Processing Using Graphics Processors. <https://www.cse.ust.hk/gpuqp/>.
- 16) Medusa: Building GPU-based Parallel Sparse Graph Applications with Sequential C/C++ Code, <https://github.com/Xtra-Computing/Medusa>. [By Jianlong Zhong*, #download=279, as of Feb 15, 2015] **[Medusa has been adopted by two research groups in NTU, under Prof. Steve Turner and Dr. Yang Liu (SCE).]**
- 17) OmniDB and nanoCG: High Performance Databases on Heterogeneous Processors. <https://github.com/Xtra-Computing/omniDB>. [By Jiong He* and Shuhao Zhang*]

Besides, we have created some open data sets for public use.

- 18) Grab-Traces & TPC-DS Presto query plans [By Johan Kok Zhi Kang*, a joint Ph.D. with Grab]

F. Funds and Projects

- 1) Towards Energy-Efficient Tera Scale Graph Processing on FPGAs. Google South Asia & Southeast Asia Research Awards 2022 (PI: Bingsheng He, Co-PI: Weng-Fai Wong, Yao Chen), USD 30,000
- 2) "Real-Time Deep Learning Networks for Fraud Detection in Modern E-Marketplace Systems", Source: AI Singapore Technology Challenge (Open-Theme). Role: PI. Co-PI: Dr. Bryan Hooi, Dr. Weng-Fai Wong (NUS), Yao Chen (ADSC). Collaborator: Deming Chen (UIUC), Jia Chen (Grab). Amount: 3.3 million SGD. 04/2022-03/2025.
- 3) "Memory Efficient Graph Accelerators on HLS-based FPGAs." Source: MoE AcRF Tier 2, Amount: SGD896,374, 02/2021-02/2025, Role: PI. Co-PI: Dr. Weng-Fai Wong (NUS), Collaborator: Dr. Yao Chen (ADSC)
- 4) "GraphMind: Energy-Efficient Hardware Accelerators for Graph Neural Networks." Source: NUS Advanced Research and Technology Innovation Centre (ARTIC), 08/2021-08/2024 Amount: 362,334, Role: PI. Co-PI: Dr. Weng-Fai Wong (NUS).
- 5) "Novel Deep Learning Networks for Outlier Detections in RTC Networks", Source: Microsoft Research, Amount: 41,000 SGD, 15 Apr 2021 to 14 Apr 2022, Role: PI.
- 6) NUS-Grab Joint Lab Phase 2, Source: Grab and NUS, Amount: 1,360,000 SGD, 04/2021-04/2023, Role: PI, Lead PI: Ng See Kiong, Other PIs: Roger Zimmerman, Bryan Hooi.
- 7) "Toward Trustable Model-centric Sharing for Collaborative Machine Learning." AI Singapore Research Programme. ID: AISG2-RP-2020-018. Amount: 8,401,002.40. Start Date: 01 April 2021, End Date: 31 March 2025. Role: Co-Lead-PI. Lead PI: Ng See-Kiong, NUS. Co-PI: Wynne Hsu (NUS), Lee Mong Li (NUS), Vincent Tan Yan Fu (NUS), Warren B. Chik (SMU).
- 8) "Real-Time Misinformation Assessment and Propagation Analysis with Graph Neutral Networks." NUS Centre for Trusted Internet and Community – Research Project Grant (NUS CTIC-RP). ID: CTIC-RP-20-03. Amount: S\$300,000. Role: PI. Co-PI: Bryan Hooi (SoC), James Pang (Biz). Period: Feb 28 2021- 28 Feb 2024
- 9) "Novel Deep Learning Networks for Fraud Detection with User Behavior Sequences", Alibaba Innovation Research (AIR) Program 2021. Role: PI. Amount: S\$112,992.
- 10) "Key Techniques and System Research on Energy-Efficient Graph Computation Accelerators", single PI, NSFC (NSF China) Oversea Collaboration Grant (also known as "NSFC Oversea Distinguished Young Scholar Award"), 01/01/2020-12/31/2023, Chinese RMB 1,600,000.
- 11) "Real-time fraud management based on graph analytics and learning". Asian Institute of Digital Finance program, Source: NRF, Role: PI, Amount: 2.42 million, Co-PI: Xiaokui Xiao (NUS), Collaborator: Yuchen Li (SMU). Period: July 1 2020 to June 30 2024.
- 12) Xilinx Adaptive Compute Clusters (XACC) program, Source: Xilinx, Amount: 256,716 SGD cash + 70K USD FPGA donations, Role: PI, Co-lead PI: Weng-fai Wong (NUS), Tulika Mitra (NUS). <https://www.comp.nus.edu.sg/news/3371-2020-xacc-research-cluster/>
- 13) "Modeling and Improving RTC applications with Graph Neutral Networks", Source: Microsoft Research, Amount: 45,000 SGD, 01 Jan 2020 to 31 Dec 2020, Role: PI. Co-PI: Weng-fai Wong (NUS), Yao Chen (ADSC)
- 14) "Hardware Accelerated Machine Learning Cloud Services (R-252-000-A64-114)", Official Project No: T1 251RES1824, Amount: 140K, 31/03/2019 to 30/03/2022, Role: PI. Co-PI: Wei Wang (NUS).
- 15) Sensetime Gift Grant on GPU virtualization support for deep learning, Amount: 40K SGD, 12/2018- 12/2019, Role: PI.
- 16) NUS-Grab Joint Lab Phase 1, Source: Grab and NUS, Amount: 2,546,600 SGD, 07/2018-07/2019, Role: PI, Lead PI: Ng See Kiong, Other PIs: Wynne Hsu, Kian-Lee Tan, Roger Zimmerman.
- 17) "In-Memory Online Analytical Processing on Coupled Heterogeneous Architectures", Source: MoE AcRF Tier 2, Amount: SGD 604, 412 (SGD 464,412 + 140,000 (RSS)), 12/2017-11/2020,

- Role: PI. Co-PI: Dr. Weng-Fai Wong (NUS), Collaborator: Dr. Jiong He, Dr. Yao Chen (ADSC)
- 18) "Accelerating Graph Applications on Heterogeneous Computing (R-252-000-630-597)", Source: Huawei, Amount: 298,200 SGD, 17 Nov 2016 to 16 May 2018, Role: PI.
 - 19) "Understanding and Optimization of Cloud Resource Utilization on Dynamic Resources (R-252-000-633-592)", Source: Microsoft Research, Amount: 41,000 SGD, 01 Jan 2017 to 31 Dec 2018, Role: PI.
 - 20) "EZ-Transaction Analytics for a Cashless Future", Source: EZlink, Amount: 96,300 SGD, 15 June 2017 – 14 Dec 2017, Role: Co-PI.
 - 21) "Scaling Up Real-Time Data Stream Processing Systems on Modern Architectures (R-252-000-623-114)", Official Project No: T1 251RES1610, Amount: 149.5K, 01/10/2016 to 30/09/2019, Role: PI.
 - 22) "A Declarative Management System for Big Data Workflows (R-252-000-612-133)", Source: NUS Startup Grant, Amount: 250K SGD, 20/06/2016 to 19/06/2019, Role: PI. (Completed)
 - 23) "In-Memory Data Analytics on Intel-Altera Heterogeneous Architecture Research Platform", Source: Intel-Altera Heterogeneous Architecture Research Platform program, Amount: USD21, 600 (in terms of hardware platform donation), role: PI. (Single Investigator, one of the only two donations in Asia), 2015
 - 24) "GPGPU for Real-Time Graph Data Analytics", Source: MoE AcRF Tier 2, Amount: 459, 375 SGD, 06/2013-06/2016, Role: PI. (Single Investigator)
 - 25) "Fine-Grain Dynamically Reconfigurable Platform for High-Performance Computing", Source: MoE AcRF Tier 2, Amount: 651,176 SGD, 01/2013-03/2016, Role: PI (since the PI left NTU in Nov 2013). Co-PI: Dr. Douglas Maskell.
 - 26) "Theoretical Study and Technical Design Of Green Memory Cloud Systems for Real-Time Data Analytics", Source: MoE AcRF Tier 1, Amount: 150,000 SGD, 10/2014-10/2017, Role: PI. Co-PI: Tang Xueyan & Cai Wentong.
 - 27) "Dependency-aware virtual machine placement in data centers", Source: MoE AcRF Tier 1, Amount: 207,000 SGD, 02/2014-02/2017, Role: Co-PI. PI: Tang Xueyan.
 - 28) Internet of Things Seed Grant, Source: NTU, Amount: 25, 000 SGD, 12/2012—03/2013, Role: PI. Co-PI: Zhiqi Shen, Wee Keong Ng, Xiaoming Li (Peking Univ.), Haixun Wang (Microsoft), Prakash Sundaresan (Microsoft).
 - 29) "Cloud-Assisted Real-Time and Large-Scale Monitoring and Analysis for Water Quality", Source: NRF EWI Singapore, Amount: 2, 110, 400 SGD, 08/2011-04/2016. Role: PI. Co-PI: Mo Li, Lloyd Chua, Lei Chen, Eikaas HANS. [Completed]
 - 30) Inter-disciplinary Strategic Competitive Fund of Nanyang Technological University 2011 for "C3: Cloud-Assisted Green Computing at NTU Campus", 250K SGD, 04/2011 - 11/2013. Role: PI. Co-PI: Xueyan Tang (SCE), Maode Ma (EEE), Ming Jian (NBS), Lee Sau-Lai (HSS).
 - 31) AcRF Tier-1 Grant from Singapore, Source: MoE, Amount: 50K SGD, Period: 03/2011-02/2013, Role: PI. (Single Investigator)
 - 32) NVIDIA Academic Partnership Award 2010/2011 for "Graph Computation and Visualization", Source: NVIDIA, Amount: USD 25, 000 (SGD 32, 225) and hardware donation, Role: PI. (Single Investigator)
 - 33) NTU Startup Grant for "An Elastic Data-Intensive Computing System in the Cloud", Period: 2010-2013, Source: NTU, Amount: SGD 100K, Role: PI. (Single Investigator)
 - 34) "Amazon Web Services Research Grant", Source: Amazon LLC, Total Amount: 8.5K USD, Role: PI,
 - Period: 2018-2019, Amount: US\$ 7.0K for free usage of AWS.
 - Period: 2011-2012, Amount: US\$ 5.0K for free usage of AWS.
 - Period: 2010-2011, Amount: US\$ 3.5K for free usage of AWS.
 - 35) Hong Kong Competitive Grant for graph databases (No. 210510), Period: 2010-2012, Source: Hong Kong RGC, Amount: HK\$ 445K (SGD 72K), Role: co-PI. PI: Byron Choi (HKBU)
 - 36) Workload monitoring and analysis in large-scale data center system. Source: A*STAR DSI, Period: 2011-2012, Amount: 20K SGD, Role: co-PI/PI(since the PI left NTU in June 2011). PI: Foh Chuan Heng (NTU).

G. Patents and Technical Disclosures

- "Dynamic graph partitioning for mixture of experts model on graph networks", Johan Kok Zhi Kang, Suriyanarayanan Venkatesan, Sien Yi Tan, Feng Cheng, Bingsheng He. (pending patent with Grab) 2022.
- "Efficient execution of data stream processing systems on multi-core processors." Shuhao Zhang, Bingsheng He, Daniel Hermann Richard DAHLMEIER. Patent No.: US20180129713A1, Dec 31, 2019.
- "Multi-query optimizer for complex event processing." Shuhao Zhang, Hoang Tam VO, Daniel Hermann Richard DAHLMEIER, Bingsheng He. Patent No.: US 9,953,056 B2, Apr 24, 2018.
- "Methods and systems for wireless transmission of data between network nodes." Wan DU, Jansen Christian Liando, Mo Li, Bingsheng He. Patent No.: US 9,955,238 B2, Apr 24, 2018.
- NTU/NIEO TD Ref No: TD/038/14, Title: Medusa: A Framework For Simplified Graph Processing On GPUs, with Jianlong Zhong*.
- NTU/NIEO TD Ref: TD/283/14 Title: A novel networking paradigm for sparse sensor network deployment, with Wan DU, Jansen Christian LIANDO, Mo LI. [filed on 20 January 2015, accorded Singapore provisional application number 10201500426Y]

IV. SERVICES

A. Professional Activities

A.1. Conference/Workshop Organizer

1) PC Co-Chair:

- a. IEEE ICDE 2024: IEEE International Conference on Data Engineering (ICDE), Utrecht, Netherlands, April 16-19, 2024 **[Rank-1 in CS Ranking]**
- b. IEEE ICDCS 2020: 40th IEEE International Conference on Distributed Computing Systems **[Rank-1 in CS Ranking]**
- c. AIDB 2020: 2nd International Workshop on Applied AI for Database Systems and Applications, held with VLDB 2020. <https://sites.google.com/view/aidb2020>
- d. AIDB 2019: 1st International Workshop on Applied AI for Database Systems and Applications, held with VLDB 2019. <https://sites.google.com/view/aidb2019>
- e. IEEE BigData Congress 2018
- f. ICCCRI 2017 (5th International Conference on Cloud Computing Research and Innovation)
- g. HardBD 2016 (International Workshop on Big Data Management on Emerging Hardware) in conjunction with ICDE 2016
- h. IEEE CloudCom 2015: IEEE 7th International Conference on Cloud Computing Technology and Science (CloudCom) 2015, Vancouver.
- i. IEEE CloudCom 2014: 6th International Conference on Cloud Computing Technology and Science (CloudCom) 2014, Singapore.
- j. The Second International Workshop on Flash-based Database Systems in conjunction with DASFAA 2012

2) General Chair:

- a. ACM Asia-Pacific Workshop on Systems (APSys), 2022

3) Co-Chair: CACM Regional Special Section on East Asia and Oceania Region 2022

- 4) Organizer: Dagstuhl Seminar on "Database Architectures for Modern Hardware". Peter Boncz, Goetz Graefe, Bingsheng He, Kai-Uwe Sattler, June 17-22, 2018. Report: http://drops.dagstuhl.de/opus/volltexte/2019/10056/pdf/dagrep_v008_i006_p063_18251.pdf

5) Vice/Track Chair:

- a. CCGrid 2021: the 21th IEEE/ACM International Symposium on Cluster, Cloud and Internet

- Computing (Scheduling and Resource Management)
 - b. ICPADS 2020 (Performance)
 - c. ISC High Performance 2020 (Machine Learning)
 - d. IEEE HiPC 2020 (Scalable Algorithms and Analytics (Data Science))
 - e. HPC Asia 2018 (Data, Storage and Visualization)
 - f. ICPADS 2018 (Distributed & High Performance Computing)
 - g. IEEE ICDE 2017 **[Rank-1 in CS Ranking]**
 - h. VLDB 2017 Demo track **[Rank-1 in CS Ranking]**
 - i. The 17th IEEE International Conference on Computational Science and Engineering (CSE2014) for the track of Cluster, Grid, P2P and Cloud Computing
- 6) Senior PC:
- a. EDBT 2023
 - b. DASFAA 2021, 2022
 - c. CIKM 2017, 2019, 2022
 - d. IEEE Big Data Conference 2014, 2018
- 7) Artifact Evaluation Committee Co-Chair: PACT 2021
- 8) Organizer for VLDB21 RoundTable - "Data Management on Modern Hardware"
- 9) Workshop Chair: IEEE HPCC 2019 (IEEE International Conference on High Performance Computing and Communications)
- 10) Publication chair: IEEE International Conference on Parallel and Distributed Systems (ICPADS) 2012
- 11) Special Session Chair: The 2nd IEEE International Conference on Big Data Intelligence and Computing (IEEE DataCom 2016)
- 12) Conference session chair:
- a. 2020 ACM Symposium on Cloud Computing (Session: Memory Management)
 - b. 2020 ACM International Conference on Information and Knowledge Management (Session: Graphs and Streams)
 - c. 2019 IEEE International Conference on Data Engineering (ICDE 2019, Session: Modern Hardware and In-Memory Database Systems)
 - d. 2017 ACM International Conference on Information and Knowledge Management (CIKM, Session: Graph Data)
 - e. 2016 Supercomputing Frontier 2016 (SCF, Session: many-core)

A.2. Steering Committee

- 13) Member, AIDB (International Workshop on Applied AI for Database Systems and Applications, held with VLDB), 2021-present
- 14) Member, HPC conference Supercomputing Asia (SCA), 2018-present
- 15) Member, HardBD (International Workshop on Big Data Management on Emerging Hardware), 2017

A.3. Editorial Board

- 16) **Associate Editor:**
- a. ACM Computing Surveys, Jan 2020 – present
 - b. IEEE TKDE (IEEE Transactions on Data and Knowledge Engineering, June 2019- present) **[Rank-1 in CS Ranking]**
 - c. DAPD (Springer Distributed and Parallel Databases (DAPD) Journal, April 2018- present)
 - d. IEEE TPDS (IEEE Transactions on Parallel and Distributed Systems, Oct 2015- Nov 2019) **[Rank-1 in CS Ranking]**
 - e. IEEE TCC (IEEE Transactions on Cloud Computing, Jan 2014- Jan 2021)

17) Editorial Board Member:

- a. CCF Transactions on High Performance Computing (CCF THPC), July 2018- present
- b. CMC Journal - Computers, Materials and Continua Journal (Nov 2019- present)
- c. Journal of Computer Science and Technology (JCST), Nov 2018- present **[the first English language journal in the computer field published in China]**
- d. PVLDB (Proceedings of the VLDB Endowment, 2012-2015, 2017, 2018 demo) **[Rank-1 in CS Ranking]**
- e. JDM (Journal of Database Management, 2011-2014)

18) Guest editor

- Special Issue on "Big Data Systems on Emerging Architectures" for IEEE Transactions on Big Data, March 2019.
<https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=8654018>
- Special issue on "Economics and Market Mechanisms for Cloud Computing" for IEEE Transactions on Cloud Computing March 2015
- Special issue on "Cloud Computing and Big Data: Applications, Systems and Perspective" for Springer The Journal of Supercomputing, Nov 2018.

A.4. Program Committee Members

- 2023: ACM SIGMOD, VLDB, CIDR, SC
- 2022: ACM SIGMOD, EDBT, CIDR, SC, CCGrid, PDSW
- 2021: IEEE ICDE, CIDR, IEEE ISPASS, VLDBJ (special issue)
- 2020: ACM SIGMOD, IJCAI (regular track and special track), IEEE/ACM PACT, CIDR, ACM SoCC, IEEE/ACM CCGrid, IEEE BigData, PDSW,
- 2019: IEEE/ACM PACT, ACM/IEEE SC, IEEE ICDE (demo), ADMS, FML, IEEE/ACM CCGrid, IEEE HiPC
- 2018: ACM SIGMOD, Euro-Par, ADMS, IEEE/ACM CCGrid, PDSW-DISCS, ACM CIKM,
- 2017: IEEE IPDPS, IEEE ICDE, SSDBM, IEEE/ACM CCGrid, Euro-Par (workshop), IEEE Big Data, IEEE DataCom, ACM CIKM, HPBDC
- 2016: ACM SIGMOD, IEEE ICDE, ACM/IEEE SC, IEEE IPDPS, ACM SoCC, IEEE IC2E, IEEE/ACM CCGrid, Eurosys Doctoral Workshop, Euro-Par, HPBDC, IEEE ICPADS, ACM SAC (Cloud Computing Track)
- 2015: IEEE IPDPS, ACM SoCC, ICPP, ACM PPOPP (external member), IEEE HiPC, IEEE CloudCom, HPBDC, CloudCom-Asia, IEEE ICPADS, DaMoN, DAPHNE, WAIM (demo),
- 2014: ACM HPDC, ACM PPOPP, IEEE Big Data, IEEE BigData Congress, IEEE CloudCom, APSys, DataCloud, InterCloud, PDCAT, GID
- 2013: ACM SoCC, IEEE Big Data, IEEE CloudCom, IEEE ICPADS, VRST, WAIM (demo), GID
- 2012: IEEE CloudCom, ACM CIKM, IEEE ICPADS, ACM SAC (Cloud Computing Track), DBKDA, DASFAA, GID
- 2011: ACM SAC (Cloud Computing Track), AINA, DASFAA, WEPA, GDB, IC3, GID
- 2010 or before: DEXA, ACM SNS, WCMT

A.5. External/Invited Reviewers

- 2020: Wiley CCPE, IEEE Intelligent Systems, DAPD
- 2019: IEEE CAL, Springer FCS, Springer Cluster Computing, VLDBJ, Elsevier DSS, Journal of Cloud Computing, IEEE Access, IEEE Transactions on Industrial Informatics, FGCS
- 2018: VLDBJ, IEEE Access, IEEE TKDE
- 2017: ACM Computing Survey, ACM TOS, IEEE Design & Test, SCIENCE CHINA Information Sciences, IEEE TKDE, ACM TODS

- 2016: IEEE TPDS, IEEE TMM, WWWJ, Elsevier Parallel Computing, IEEE TC, IEEE TKDE, ACM TKDD
- 2015: IEEE TPDS, IEEE TC, IEEE TKDE, IEEE TSC, IEEE INFOCOM
- 2014: IEEE TPDS, IEEE TC, IEEE TKDE, ACM TODS, DAPD, Encyclopedia of Cloud Computing, IEEE TSC, LNCS TLDKS, Journal of Supercomputing, ACM TOS, IEEE INFOCOM, DKE
- 2013: IEEE TPDS, IEEE TKDE, IEEE TSC, FGCS,
- 2012: IEEE TPDS, IEEE TKDE, FGCS, ACM SPAA, IEEE INFOCOM, IEEE ICDCS, JPDC (Journal of Parallel and Distributed Computing)
- 2011: IEEE TPDS, IEEE TKDE, IEEE INFOCOM, VLDB, Elsevier JSA, Journal of Engineering and Computer Innovations (JECI), Springer TOCS, IEEE TVLSI
- 2010 or before: IEEE TPDS, IEEE TKDE, VLDB, WAIM, IEEE ICDCS, IEEE ICDE, ACM TOIS, JISE (Journal of Information Science and Engineering), NDBC (National Database Conference, China), DKE

A.6. External/Invited Academic Services

19) Team leader for FRC (Foundation Research Capabilities) Green Computing supported by Singapore NRF, 2023

20) Panelist for HPC I/O in the Data Center Workshop (HPC-IODC) in ISC HPC 2021

21) Nominator for the 2021 VinFuture Prize

22) Grant proposal reviewer:

- Competitive Research Grants (CRG) at King Abdullah University of Science and Technology (KAUST) 2021
- Invited reviewer in Review Panel of Impulse Fund of KU Leuven (University of Leuven, Belgium): 2014
- Fund External Reviewer for Agency for Science, Research & Technology (A*STAR) at (Invitation from Joint Council Office): 2014
- Swiss National Science Foundation: 2016
- Hong Kong RGC: 2017, 2020, 2021, 2022
- Canada NSERC Discovery Grant: 2018

23) External Reviewers for Ph.D. thesis/proposal

- Universidad Politécnica de Madrid
- Hong Kong Polytechnic University 2016, 2018
- Hong Kong Baptist University 2016, 2018

24) Tenure Evaluator:

- Nanyang Technological University
- University of Melbourne
- Tsinghua University
- Shanghai Jiao Tong University
- Huazhong University of Science and Technology

25) Others

- Reviewer for NRF Fellowship 2021
- External Reviewer on Research Centres at Hong Kong Baptist University, 2020
- Green Data Center Primer Singapore 2012 (hosted by NCCS, NRF, iDA Singapore), Role: Contributing Author, https://www.nccs.gov.sg/sites/nccs/files/GDC_primer_20140825.pdf.
- Seminar coordinator at the system group of MSRA (2009-2010)
- Student helper for ICDM2006, Hong Kong.

B. University Services at National University of Singapore

- 26) Vice Dean (Research), School of Computing (July 2020- present)
- 27) Member of NUS Guangzhou Research Translation & Innovation Institute (NUSGRTII) Advisory Committee (August 2023 to Present)
- 28) Member of Department Evaluation Committee (DEC) for promotion and tenure, NUS (2020)
- 29) Member of Whitespace Research Committee (RC) for The Asian Institute of Digital Finance (AIDF), May 1 2021 to June 30 2023.
- 30) Member of CRYSTAL Centre Management Board (July 2020 to Present)
- 31) Member of Steering Committee for NExT Research Center (March 2021 to present)
- 32) Member of Management Board for NUS Advanced Robotics Centre (ARC, March 2021 to present)
- 33) Assistant Dean (Graduate Studies), School of Computing (July 2019- July 2020)
- 34) Affiliated faculty member of NUS IDS (Institute of Data Science).
- 35) Joint Academic Committee member (Sept 2016- present), co-chair (July 2017- June 2019), chair (July 2019- July 2020): responsible for curriculum design and creating new programs.
- 36) CEG EAB Accreditation Committee member (Sept 2016- Oct 2017): coordinating the EAB course accreditation for the computer engineering program.
- 37) Member of Graduate Admission Committee, Reviewer for Ph.D. admission/Master of Computing, 2017- present
- 38) Chair of Disciplinary Committee of CS1010E incident April-June 2020
- 39) Invited lecturer for CS2309 CS Research Methodology: 2017, 2021
- 40) Volunteer
 - University Internal Reviewer for MOE Tier 2 ARC 02/2023
 - Interviewer for MSBA master program admission 2021
 - NUS SoC Info Session Sat 16 May 2020
 - NUS e-Open Day, 2 Mar 2020
 - Invited sessions for NUS School of Computing Master Programmes in Jan and Feb 2020 to SCALE Global (3 sessions)
 - Poly Graduating Cohort visiting NUS School of Computing on 18 January 2020, 10.30am to 1pm
 - NUS Board of Trustees lunch meeting, Apr 26 2019
 - NUS Open Day, 9 Mar 2019
 - Master student counselling (5 students), Jan 22-24 2019
 - Interviewer for undergraduate scholarship, April 4 2018
 - NUS Open Day, 10 Mar 2018
 - Poly Graduating Cohort visiting NUS School of Computing on 14 January 2018
 - Speaker at SoC Grantsman Workshop 2018
 - Reviewer for Ph.D. admission of 2017 August intake
 - Poly Graduating Cohort visiting NUS School of Computing on 14 January 2017
 - NUS Open Day, Sat 11 Mar 2017
 - NUS Computing Information Session for Prospective Students and their Parents, 13 May 2017, 11am to 2pm

C. University Services at Nanyang Technological University

- 41) SCE Outreach committee member (August 2011- Nov 2012, Jan 2014- May 2015)
 - NTU Open House, March 5, 2016
 - Let's talk NTU, Jan 30 2016
 - Demo and SCE presentations for the campus visit by Temasek Poly students, August 5, 2015

- Attend Final Presentation & Selection For Outreach video, August 4 2015.
 - NTU Open House, March 7 2015
 - CoE presentation for Learning Journeys for YJC Year 1 Science Students, 9 Feb, 2015
 - RNS Men 2015 on 7th February 2015
 - Networking session for SCE Visit, Pioneer JC, 30th Jan 2015.
 - Let's Talk NTU @ Suntec - 24 Jan, 2015
 - Panel member for Tan Sri Dr Tan Chin Tuan Scholarship Interview, August 21, 2015
 - Reception and networking Session for SCE Visit 15 August 2014 - Temasek Poly Yr 3 Students
 - High Tea Session with the YJC Students (Feb-11 2014)
 - SCE New Admission Info Session (May 3 2014)
 - Involvement in outreach video and Facebook page design (by Oct 2012).
 - Temasek Polytech Technical Outreach, July 11, 2012
 - NTU Open House 2012
 - High Tea with the NTU President - 31 Mar 2012
 - Pre-event briefing for Let's Talk NTU 2012, Jan 2012
 - SCE Course talk for Temasek Polytech (Dec 2011)
 - Presentation at National Engineers Day, Singapore (Nov 2011)
 - SCE Course talk for Republic Polytech (Oct 2011)
 - Computing Outreach Programme 2011, "The Cloud is All Around Us: Applications, Systems and Perspectives".
 - CoE Seminar Pamphlet, 2011-2012
- 42) University mentor for UG students 2010—present
- 43) URECA program mentor at NTU
- All students have published their results in URECA Proceedings.
 - 2013-2014 (3): Nguyen Tuan Phong, Guo Jintao, Sutrisno Suryajaya Dwi Putra
 - 2012-2013 (4): Quek Jian Hong Joel, Cheng Min, Nguyen Quoc Duy Tan, Zhengzi Xu
 - 2011-2012 (2): Cui Yan, Lim Wei Li
 - 2010-2011 (2): Nguyen Ha Duy, Bui Le Linh
- 44) Mentor for industrial attachments or orientation
- IO 2014 (7): Tsoi Lok Yin, Tay Wang Ding Allan, Aaron Alexander Lim Qing Rong, Ong Yin Woon, Ong Jin Jie, Li Ang, Muhammad Hafiz Bin Rohaidi
 - IO 2013 (9): Gao Cheng, Gong Li, Li Jiaxiang, Luo Huan, Tiffany Chen, Ang ding yang, Goh Chengyee, HANAFI BIN MAT SANI, Teo Yi Si Josias
 - IO 2012 (5): Seet Wei Zheng Barry, Zheng Cihui, Do Dinh Tho, Seo Huan Lan Kelly, Teh Hua Ze Brian
 - IA 2011 (5): FU ANQI, WONG DE SHUN, LIM GEOK MUI, KIMISHA PIYUSH MODY, FOO YONG ZHENG
- 45) Final Year Project advisor
- 2014-2015 (7): Chemburkar Nishant Rajeev, Jomain Tan Zi Hao, Lu Mengjiao, Phu Mon Htut, Liu Xiao, Tan Xuan Min, Lai Qi Rong
 - 2013-2014 (7): Clement Tan Wee Seng, Xi Yewen, Wang Qiudi, Soh Sheng En Alexius Matthias, Tran Quoc Nguyen, Xu Qianwen, Lim Chun Leng
 - 2012-2013 (6): Seet Choon Nguan, Andrew, Zhang Shuhao, Abhijith Padmakumar Renuka, Seet Wei Ren Ivan, Pan Yongnan, Le Tan Khoa
 - 2011-2012 (8): Yong Lee Sun, Yeong Zhi Jian, Sze Kit Ying, Nguyen Ha Duy, Leong Yew Long, Liu Farui, Goh Ming Rui, Siow Jian Yang Llukelly
- 46) Talent Outreach Section NRP Enrichment Project Advisor
- 2012 (student: Han Cao from Nanyang JC, Result: Silver Award)

- 2011 (student: Zhao Yichen from River Valley High School, Result: Gold Award)
 - NRP Enrichment Assessor (2011- 2015)
- 47) Others
- NTU Tier 3 proposal reviewer (Dec 2016)
 - SPMS MAS Search Committee (Member, Feb 2016)
 - Faculty recruitment buddy (March 2015)
 - NAP buddy (Jan 2013)
- 48) Research area coordinator in SCE (for the area Parallel and Distributed Computing): Jan 2015 to May 2016.

V. HONORS AND AWARDS

- 1) SIGMOD Distinguished PCs, 2022, <https://sigmod.org/sigmod-distinguished-aes-and-pcs/>
- 2) Basic Theory and Approaches for High-Performance Big Data Processing, CCF Natural Science Awards, 2022, first prize. <https://www.ccf.org.cn/Awards/Awards/2022-11-08/776110.shtml> (2022年度“CCF科技成果奖”自然科学一等奖, 项目名称: 大数据高时效处理基础理论与方法, 主要完成人: 石宣化、金海、袁平鹏、何丙胜、罗琼)
- 3) Google South Asia & Southeast Asia Research Awards 2022
- 4) “Highly Cited Researcher”: Publications ranked in the top 1 percent by citations for field and publication year in Clarivate’s Web of Science citation index, 2021, for field of Distributed Computing.
- 5) NUS Dean's Chair Associate Professor, Jan 2021-Dec 2023
- 6) NUS SoC Faculty Teaching Excellence Award (FTEA) 2019/2020
- 7) ACM Distinguished Member Award 2020
- 8) Xilinx Adaptive Compute Clusters (XACC) program 2020 (lead PI, one of the four awarded projects worldwide from Xilinx). Amount: 180K USD + 70K USD Hardware donations. <https://www.comp.nus.edu.sg/news/2020-xacc-research-cluster/>
- 9) IEEE TPDS Awards for Editorial Excellence for IEEE Transactions on Parallel and Distributed Systems (TPDS), 2018/2019.
- 10) Editorial Excellence and Eminence (EEE) award for IEEE Transactions on Cloud Computing (TCC), 2018/2019.
- 11) Nanyang Education Award (School) 2014 at Nanyang Technological University (also known as “best teacher of the year”, 1 out of 80+ faculties in the school, <http://www.ntu.edu.sg/NanyangAwards/Recipients/Pages/NanyangAwards2014Winners.aspx>)
- 12) Invited DYL (Distinguished Young Lecturer) in WAIM 2013 (The 14th International conference on Web-Age Information Management, <http://idke.ruc.edu.cn/waim2013/DYL.html>).
- 13) NVIDIA CUDA Research Center, 2010- Present, Role: PI. (<https://developer.nvidia.com/academia/centers/nanyang-technological-university-cuda-research-center>).
- 14) NVIDIA Academic Partnership Award 2010/2011. (i.e., NVIDIA faculty award. Amount: USD 25, 000 (SGD 32, 750) and GPU hardware donation)
- 15) IBM PhD Fellowship 2007/08.
(News: http://www.cse.ust.hk/News/IBM_Fellowship2007/).
- 16) Graduate Student Scholarship of CIDR 2007.
- 17) Outstanding undergraduate student of Shanghai City, 2003.
- 18) First class scholarship of Shanghai Jiao Tong University, 1999—2003.
- 19) First class studentship of Bank of Communications, China, at Shanghai Jiao Tong University, 1999.
- 20) Paper awards and honorable mentions

- **VLDB 2023 Industry Track Best Paper Runner-up Award** (for the work "Xuanhe Zhou, Cheng Chen, Kunyi Li, Bingsheng He, Mian Lu, Qiaosheng Liu, Wei Huang, Guoliang Li, Zhao Zheng, Yuqiang Chen. FEBench: A Benchmark for Real-Time Relational Data Feature Extraction. VLDB: International Conference on Very Large Data Bases (VLDB) 2023.")
- **Featured Paper of the Month -- Issue July 2022 of IEEE Transactions on Computers** (for the work "Xiaofei Liao, Jin Zhao, Yu Zhang, Bingsheng He, Ligang He, Hai Jin, Lin Gu. A Structure-aware Storage Optimization for Out-of-Core Concurrent Graph Processing. IEEE TC: IEEE Transactions on Computers 2021")
- **CIKM 2021 Best Demo Paper Runner-up Award** (for the work "Zhao Li, Pengcheng Zou, Xia Chen, Shichang Hu, Peng Zhang, Yumou Zhang, Bingsheng He, Yuchen Li, Xing Tang: From Community Search to Community Understanding: A Multimodal Community Query Engine. ACM CIKM: ACM International Conference on Information and Knowledge Management 2021 (demo paper)")
- **PREMIA Best Student Paper Gold Award 2021** (for the work "Qinbin Li*, Zeyi Wen^, Bingsheng He. Practical Federated Gradient Boosting Decision Trees. AAAI: Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI-2020)", <http://www.premiasg.org/formembers/premia-best-student-paper-awards/premia-best-student-paper-awards-2021/>).
- **"Best papers" of FPGA 2021** (for the work "Xinyu Chen*, Hongshi Tan^, Yao Chen, **Bingsheng He**, Weng-Fai Wong, Deming Chen. ThunderGP: HLS-based Graph Processing Framework on FPGAs.").
- **2019 Best Paper Award from IEEE Transactions on Parallel and Distributed Systems by the IEEE Computer Society Publications Board (1 out of 987 submissions**, for the work "Zeyi Wen^, Jiashuai Shi*, Bingsheng He, Jian Chen, Kotagiri Ramamohanarao, and Qinbin Li*, "Exploiting GPUs for Efficient Gradient Boosting Decision Tree Training," IEEE Transactions on Parallel and Distributed Systems, vol. 30, no. 12, 2019, pp. 2706-2717.")
- **"Best paper candidate" in ICPP 2020 (4 out of 313 submissions**, for the work "Jiya Su, Feng Zhang, Weifeng Liu, **Bingsheng He**, Ruofan Wu, Xiaoyong Du, Rujia Wang. CapelliniSpTRSV: A Thread-Level Synchronization-Free Sparse Triangular Solve on GPUs").
- **"Best paper candidate" in PACT 2018 (4 out of 126 submissions**, for the work "Long Zheng, Xiaofei Liao, Hai Jin, Bingsheng He, Jingling Xue, Haikun Liu. Towards Concurrency Race Debugging: An Integrated Approach of Constraint Solving and Dynamic Slicing").
- **"Best storage related papers" in USENIX ATC 2018** (for the work "Yu Zhang, Xiaofei Liao, Hai Jin, and Lin Gu, Ligang He, Bingsheng He, Haikun Liu. CGraph: A Correlations-aware Approach for Efficient Concurrent Iterative Graph Processing", invited submission to ACM Transactions on Storage (TOS))
- **2017 IEEE/ACM William J. McCalla ICCAD Best Paper Award (Front End) (1 out of 399 submissions**, for the work "Jieru Zhao, Liang Feng, Wei Zhang, Sharad Sinha, Yun (Eric) Liang, Bingsheng He. COMBA: A Comprehensive Model-Based Analysis Framework for High Level Synthesis of Real Applications."). <https://ieeeced.org/awards/william-j-mccalla-iccad-best-paper-award>
- **IEEE IC2E 2016 Best Paper Runner Up (4 out of 73 submissions**, for the work "Zhaojie Niu, Bingsheng He and Fangming Liu. Not All Joules are Equal: Towards Energy-Efficient and Green-Aware Data Processing Frameworks").
- **Top-quality papers of FPL 2015** (for the work "Zeke Wang^, Bingsheng He, Wei Zhang. A Study of Data Partitioning on OpenCL-based FPGA").
- **IEEE CloudCom 2014 Best Ph.D. Consortium Paper Award** (for the work "Amelie Chi Zhou* and Bingsheng He. Simplified Monetary Optimizations for Workflows in IaaS Clouds").
- **Best Demo Award in IEEE SECON 2014** (for the work "Wan Du, Mo Li, Zikun Xing, Bingsheng He, Lloyd Hock Chye Chua, Zhenjiang Li, Yuanqiang Zheng and Pengfei Zhou, Demo Abstract: Wind Measurements for Water Quality Studies in Urban Reservoirs.").
- **Spotlight article** of IEEE Transactions on Cloud Computing March 2014, vol. 2 no. 1. <http://www.computer.org/csdl/trans/cc/2014/01/06702454.pdf>. (for the work "Amelie Chi Zhou* and Bingsheng He, Transformation-based Monetary Cost Optimizations for Workflows in the Cloud.")
- **"Best system demos" in VLDB 2013** (for the work "Jianlong Zhong*, Bingsheng He. Parallel

Graph Processing on Graphics Processors Made Easy”, **out of 140+ submissions**, invited submission to SIGMOD Record).

- **“Best research papers” in ACM SIGMOD 2008 (2 out of 435 submissions**, for the work “Bingsheng He, Ke Yang, Rui Fang, Mian Lu, Naga K. Govindaraju, Qiong Luo, Pedro V. Sander. Relational Joins on Graphics Processors”, invited submission to ACM Transactions on Database Systems (TODS)).

VI. Keynote and Presentations

A. Invited Plenary/Keynote Talks

- 1) Invited talk at International Workshop on Trustable, Verifiable and Auditable Federated Learning in Conjunction with AAAI 2022 (FL-AAAI-22). March 2, 2022. <https://federated-learning.org/fl-aaai-2022/>
- 2) Invited keynote speech at 5th DSDE 2022-Data Storage and Data Engineering. Feb 25-27, 2022. <http://dsde.org/KeynoteSpeakers.html>.
- 3) Invited keynote talk at the Fifth Workshop on Hot Topics on Data Centers (HotDC 2020). Jan 2020.
- 4) Invited keynote speech at DSS 2020 (The 6th IEEE International Conference on Data Science and Systems). Dec 14 - Dec 16, 2020. <http://www.ieee-cybermatics.org/2020/dss/>
- 5) Invited keynote speech at 2020年全国开放式分布与并行计算学术年会 (DPCS2020), <https://conf.ccf.org.cn/web/html7/index.html?globalId=d55279f188d642549fb43301dc6cf983&type=1>
- 6) Invited plenary talk at the 17th Annual IFIP International Conference on Network and Parallel Computing. Zhengzhou, Henan, China, Sept.28-30, 2020.
- 7) Invited keynote speech at The Second International Workshop on Smart Data for Blockchain and Distributed Ledger (SDBD 2020). <https://enigmazhu.github.io/SDBD2020/>.
- 8) Invited keynote speech at ISPA 2020 (18th IEEE International Symposium on Parallel and Distributed Processing with Applications), 17-19 December 2020.
- 9) Invited keynote speech at ChinaSys Workshop, Dec 2019.
- 10) Invited plenary talk “Data Management Systems on Future Hardware: Challenges and Opportunities.” in Joint Active and HardDB Workshops in IEEE ICDE: IEEE International Conference on Data Engineering (ICDE), 2017
- 11) Invited plenary talk at “International Workshop on CoDesign” at Xian, China, Oct 2016.

B. Other Presentations

- 12) Invited lecture at AI Summer School 2022, AISG
- 13) Microsoft Research Summit 2021
- 14) Invited talk and panelist at Xilinx Adapt 2021
- 15) Invited keynote talk at NUS Computing Research Week 2021
- 16) Invited panelist at NUS Computing Research Week 2021
- 17) MSRA Collaborative Research Workshop: Large-Scale Networking Platform for Machine Learning Based Real Time Communications. Dec 2020.
- 18) Invited seminar at SPCL_Bcast at ETH (an open, online seminar series organized by Prof. Torsten Hoefler, ETHZ). <https://spcl.inf.ethz.ch/Bcast/>. Dec 2020.
- 19) "Program of Introducing Talents of Discipline to Universities (Plan 111)", invited by Hunan University, China (Prof. Kenli Li). Dec 2020.
- 20) Online Huawei Workshop on Compute Technology on December 2-3, 2020
- 21) Host for XACC@NUS Workshop 2020: Reconfigurable Computing Systems. <https://xaccnus.github.io/>.
- 22) Invited seminar from A*STAR Singapore, June 2020.

- 23) Invited presentation and tech sharing at Shopee, March 11 2019
- 24) Invited lecture "Parallel graph processing on GPUs" at ADL (Advanced Disciplines Lectures) China 2018, <https://www.ccf.org.cn/events/adl/>.
- 25) Invited talk "System Optimizations and Performance Tuning for New Generation FPGAs" at Xilinx OpenHW2017 Design Contest & Professors Conference, 2017.
- 26) Invited speakers at "Workshop on Big-Data System Architectures" at Renmin University, China.
- 27) Invited forum keynote speech for HPC China 2015, Wuxi, China (~1000 attendance).
- 28) Invited speaker on "Emerging HPC Technologies for High Performance Database Systems", the inaugural "SEIEE International Forum for Young Scientists" 2015, Shanghai Jiao Tong University, China.
- 29) Invited Speaker on "Emerging HPC Technologies for Real-Time (Big) Data Analytics: A Tutorial". 26th edition of the Australasian Database Conference, ADC 2015.
- 30) Summer School at Australian Database Conference 2014.
- 31) Invited forum keynote speech for HPC China 2014, Guangzhou (~1000 attendance).
- 32) Invited talk in WAIM 2013, Beidaihe, China (~200 attendance).
- 33) Conference talk "Long-Term Resource Fairness: Towards Economic Fairness on Pay-as-you-use Computing Systems" in ACM ICS 2014: 2014 International Conference on Supercomputing, Munich.
- 34) Conference talk "Green Databases Through Integration of Renewable Energy" in CIDR'13 (6th Biennial Conference on Innovative Data Systems Research), Asilomar, CA.
- 35) Conference talk "RAMZzz: Rank-Aware DRAM Power Management with Dynamic Migrations and Demotions" in SC12: ACM/IEEE SuperComputing 2012, Salt Lake City.
- 36) Conference talk "High-Throughput Transaction Executions on Graphics Processors" at VLDB 2011, Seattle, September 2011.
- 37) Invited presentations/seminars
 - Invited talks on "In-Memory Database Systems on Emerging Hardware: Our Ten Years' Journey", HKUST/University of Macau, Feb 21-23 2016.
<http://news.umac.mo/nrs/faces/pub/viewItem.jspx?id=37341>
 - Invited talks on "When HPC Meets Big Data: Emerging HPC Technologies for Real-Time Data Analytics" in a number of Australian Universities (including The University of Melbourne, RMIT, The University of Sydney, University of Technology, Sydney), Sept 29—Oct 3, 2014.
<http://www.cs.rmit.edu.au/seminars/2014/bingsheng.html>
<http://sydney.edu.au/engineering/it/research/news/seminars.shtml>
 - Invited talks on "GPGPU for Real-Time Data Analytics" in a number of universities in China (including Shanghai Jiao Tong University, Nanjing University, University of Science and Technology of China, Huazhong University of Science & Technology), Dec 19--Dec 28, 2013.
http://www.cs.sjtu.edu.cn/news_view.action?newsId=170
<http://grid.hust.edu.cn/notice.php?id=1697>
<http://page.renren.com/601368951/fdoing/5089607769>
 - Invited talk on "Architecting Faster and Greener Big Data Management Systems", Dec, 2012, Hong Kong Baptist University.
 - Invited talk on "Architecting Faster and Greener Big Data Management Systems", Oct, 2012, Renmin University, China.
 - Invited talk on "Application Driven Cloud System Research" at PDCC cloud computing workshop 2011, October 2011.
 - Invited talk on high performance computing at two universities in China: Tsinghua University and Huazhong University of Science and Technology, June 2011.
 - Invited talk on high performance computing at Hong Kong Polytech University, June 2012.
 - At TU Berlin in 2016, Volker Markl is a VLDB endowment President, and also co-creator for Apache Flink (one of the most popular big data processing systems nowadays). BBDC (Berlin Big Data Center) Seminar. http://www.bdc.berlin/de/newsarticle/news/in-memory-database-systems-on-emerging-hardware-our-ten-years-journey/?tx_news_pi1%5Bcontroller%5D=News&tx_news_pi1%5Baction%5D=detail&c

[Hash=36fa62739c6296733a898567030abcf1](#)

- At ISI University of Southern California, The host: Ewa Deelman, IEEE Fellow, HPDC 2015 achievement award

VII. References

Available upon request.