YONG CHEN

ASSOCIATE PROFESSOR, DEPARTMENT OF COMPUTER SCIENCE, TEXAS TECH UNIVERSITY

PHONE: (806) 834-0284 (OFFICE)
EMAIL: yong.chen@ttu.edu
HOMEPAGE: http://www.myweb.ttu.edu/yonchen

MAIL: Computer Science Department
Texas Tech University, Box 43104
Lubbock, TX 79409-3104

RESEARCH INTERESTS

Data-intensive computing, high-performance computing, parallel and distributed computing, cloud computing, computer architecture, and system software support for computational sciences.

EDUCATION

- **Ph.D.**, Computer Science, Illinois Institute of Technology, Chicago, Illinois 2009
- **M.S.**, Computer Science, University of Science and Technology of China, China 2003
- **B.E.**, Computer Engineering, University of Science and Technology of China, China 2000

PROFESSIONAL EXPERIENCE

- **Texas Tech University (TTU)**, Lubbock, Texas
  - Associate Professor, Department of Computer Science 2017 – present
  - Director, Data-Intensive Scalable Computing Laboratory (DISCL) 2011 – present
  - Site Director/Associate Site Director, Cloud and Autonomic Computing Center at TTU (CAC@TTU) funded through NSF IUCRC program 2014 – present
  - Assistant Professor, Department of Computer Science 2011 – 2017
- **Oak Ridge National Laboratory (ORNL)**, Oak Ridge, Tennessee
  - Post-doctoral Researcher, Computer Science and Mathematics Division 2010
- **Illinois Institute of Technology (IIT)**, Chicago, Illinois
  - Senior Research Associate, Department of Computer Science 2009 –2010

  Argonne National Laboratory (ANL), Argonne, Illinois
  - Guest Faculty, Mathematics and Computer Science Division 2009 –2010
- **Illinois Institute of Technology**, Chicago, Illinois
  - Research Assistant, Scalable Computing Software Laboratory 2006-2009, 2003-2004
  - Teaching Assistant, Department of Computer Science 2004-2006
- **Argonne National Laboratory**, Argonne, Illinois
  - Research Aide, Mathematics and Computer Science Division 2008
- **Motorola Inc.**, Schaumburg, Illinois; **Institute of Design** at IIT, Illinois
  - Intern, Human Interaction Research and Development 2004, 2005
- **University of Science and Technology of China (USTC)**, AnHui, China
  - Research Assistant**,** National HPC Center of China at HeFei 2000 - 2003

GRANTS

- *Understanding and Controlling Power Consumption of GPUs towards Energy Efficient HPC Systems*, PI (Sponsor: DOE-Lawrence Berkeley National Laboratory)

- *Phase-II IUCRC Texas Tech University: Center for Cloud and Autonomic Computing*, PI, with Co-PIs Dr. Tommy Dang, Dr. Susan Mengel, and Dr. Alan Sill (Sponsor: National Science Foundation)
- *xBGAS Architecture Optimizations and Applications*, institutional PI, in collaboration with Dr. John Leidel of Tactical Computing Laboratories and Dr. Michel Kinsy of TAMU (Sponsor: DOD/NSA and CAC@TTU via NSF)
- *Development of Metadata Search Methods for Hierarchical Data Format version 5 (HDF5) Data from Experimental and Observational Science Facilities*, PI (Sponsor: DOE-Lawrence Berkeley National Laboratory)
- *xBGAS (Extended Base Global Address Space) Filesystem and Architecture Optimizations*, in collaboration with Dr. John Leidel of Tactical Computing Laboratories and Dr. Michel Kinsy of TAMU (Sponsor: DOD/NSA and CAC@TTU via NSF)
- *Elements:Software:NSCI: Empowering Data-driven Discovery with a Provenance Collection, Management, and Analysis Software Infrastructure*, PI, with Co-PIs Dr. Dong Dai, Dr. William Hase, and Dr. Brian Ancell (Sponsor: National Science Foundation)
- *Metadata Indexing and Provenance*, PI (Sponsor: DOE/Lawrence Berkeley National Laboratory)
- *CSR: Small: Collaborative Research: Tuning Extreme-scale Storage Stack through Deep Reinforcement Learning*, leading PI, with Co-PI Dr. Dong Dai and institutional PI Dr. Forrest Sheng Bao of ISU (Sponsor: National Science Foundation)
- *xBGAS: Extended Base Global Address Space for Data Center Scale Addressing*, institutional PI, in collaboration with Dr. John Leidel of Tactical Computing Laboratories and Dr. Michel Kinsy of BU (Sponsor: DOD/NSA and CAC@TTU)
- *Visualizing, Monitoring, and Predicting Health Status of HPC Centers*, Co-PI, with PI Dr. Tommy Dang and Co-PI Dr. Alan Sill (Sponsor: Dell-EMC and CAC@TTU)
- *SHF: Small: Collaborative Research: Uncovering Vulnerabilities in Parallel File Systems for Reliable High Performance Computing*, leading PI, with Co-PI Dr. Dong Dai and institutional PI Dr. Mai Zheng of ISU (Sponsor: National Science Foundation)
- *Data Center Analytics and Testing*, PI, with Co-PI Dr. Alan Sill (Sponsor: Dell-EMC and CAC@TTU)
- *RV64G/32G Address Extension*, institutional PI, with leading PI Dr. John Leidel (Sponsor: Booz Allen Hamilton Inc., DOD/National Security Agency)
- *Student Travel Grant for 2017 IEEE/ACM International Conference on Utility and Cloud Computing (UCC) and Co-located BDCAT Conference*, Co-PI, with PI Dr. Dong Dai (Sponsor: National Science Foundation)
- *OpenHPC: Open Source, Extensible High Performance Computing Platform*, Institutional PI, with Dr. John Leidel, Dr. David Donofrio, Dr. Farzad Fatollahi-Fard, and Dr. John Shalf (Sponsor: Department of Defense/Lawrence Berkeley National Laboratory)
- *Computational Algorithm and Software Development for Integrated Next-Generation Sequencing Data Analysis*, PI, with Co-PI Dr. Shengping Yang of TTUHSC (Sponsor: TTU Presidents' Collaborative Research Initiative)
- *SHF: Medium: Compute on Data Path: Combating Data Movement in High Performance Computing*, PI, with Co-PIs: Drs. Barbara Chapman, Yonghong Yan, Robert Ross (Sponsor: National Science Foundation)
- *Unified Storage Architecture for Cloud Computing*, PI (Sponsor: Nimboxx and CAC@TTU)
- *I/UCRC: Cloud and Autonomic Computing I/UCRC site at Texas Tech University (CAC@TTU)*, Co-PI, with PI Dr. Alan Sill and Co-PI Dr. Ravi Vadapalli) (Sponsor: National Science Foundation)
- *MRI Collaborative: Development of a Data-Intensive Scalable Computing Instrument (DISCI) for High*

Performance Computing, including REU SUPPLEMENT, PI, with Co-PIs: Drs. William Gropp, Philip Smith, Xian-He Sun, Yu Zhuang (Sponsor: National Science Foundation)

- *REU Site: Research Experiences for Undergraduates in Cybersecurity, Robotics, and Software Engineering*, PI (former Co-PI), with former PI Dr. Susan Urban, Co-PIs Dr. Joseph Urban, Dr. Mohan Sridharan, Dr. Hamed Sari-Sarraf (Sponsor: National Science Foundation)
- *Active Object Storage for Big Data Applications in High Performance Computing*, PI (Sponsor: Department of Energy/Argonne National Laboratory)
- *Travel Support for The 14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*, PI (Sponsor: National Science Foundation)
- *Planning Grant: I/UCRC for Center for Cloud and Autonomic Computing site at Texas Tech University*, Co-PI, with PI Dr. Alan Sill and Co-PI Dr. Ravi Vadapalli (Sponsor: National Science Foundation)
- *CSR: Medium: Collaborative Research: Decoupled Execution Paradigm for Data-Intensive High-End Computing, including REU SUPPLEMENT*, PI (Sponsor: National Science Foundation)
- *Coordinated I/O Architecture for Exascale High-Performance Computing Systems*, PI (Sponsor: Oak Ridge Associated Universities)
- *Early Adoption of NSF/TCPP PDC Curriculum at Texas Tech University*, PI, with Co-PIs Dr. Yu Zhuang and Dr. Noe Lopez-Benitez (Sponsor: NSF-GSU sub-award and IEEE TCPP)

DONATIONS/GIFTS

- An HPC cluster testbed with dual Xeon E5-2660 processors, dual NVIDIA Tesla M2090s, dual 200GB solid state drives each node, InfiniBand network cards and switch, from Dell Inc., 2012
- **Two Tesla C2075 GPU boards** and **Two GeForce GTX480 GPU boards**: for research and education, from NVIDIA, 2012
- **GeForce GTX 480**: for research and education, from NVIDIA, 2011

TEACHING

- Instructor for **CS4352 Operating Systems**
  Spring 2021, lecture course, undergraduate core (online delivery via Zoom)
- Instructor for **CS5352 Advanced Operating System Design**
  Spring 2021, lecture course, graduate core (online delivery via Zoom)
- Instructor for **CS4352 Operating Systems**
  Spring 2020, lecture course, undergraduate core (face-to-face delivery for the first-half semester and online delivery via Teams for the second-half semester)
- Instructor for **CS4379 Parallel and Concurrent Programming**
  Spring 2020, lecture course, undergraduate elective, cross-listed (face-to-face delivery for the first-half semester and online delivery via Teams for the second-half semester)
- Instructor for **CS5379 Parallel Processing**
  Spring 2020, lecture course, graduate elective, cross-listed (face-to-face delivery for the first-half semester and online delivery via Teams for the second-half semester)
- Instructor for **CS4352 Operating Systems**
  Fall 2019, lecture course, undergraduate core
- Instructor for **CS4379 Parallel and Concurrent Programming**
  Spring 2019, lecture course, undergraduate elective, cross-listed

Instructor for **CS5379 Parallel Processing**
Spring 2019, lecture course, graduate elective, cross-listed
Instructor for **CS4352 Operating Systems**
Spring 2019, lecture course, undergraduate core

- Instructor for **CS5352 Advanced Operating System Design**
  Fall 2018, lecture course, graduate core
  with both face-to-face and distance sections
- Instructor for **CS5352 Advanced Operating System Design**
  Spring 2018, lecture course, graduate core
  with both face-to-face and distance sections
- Instructor for **CS3375 Computer Architecture**
  Fall 2017, lecture course, undergraduate core
- Instructor for **CS5352 Advanced Operating System Design**
  Spring 2017, lecture course, graduate core
  with both face-to-face and distance sections
- Instructor for **CS4352 Operating Systems**
  Fall 2016, lecture course, undergraduate core
- Instructor for **CS5352 Advanced Operating System Design**
  Spring 2016, lecture course, graduate core
  with both face-to-face and distance sections
- Instructor for **CS4352 Operating Systems**
  Fall 2015, lecture course, undergraduate core
- Instructor for **CS4379 Parallel and Concurrent Programming**
  Spring 2015, lecture course, undergraduate elective
- Instructor for **CS4331 Big Data Infrastructure and Data Management**
  Fall 2014, lecture course, undergraduate elective, cross-listed
  Instructor for **CS5331-001/D01 Big Data Infrastructure and Data Management**
  Fall 2014, lecture course, graduate elective, cross-listed
  with both face-to-face and distance sections
- Instructor for **CS4379 Parallel and Concurrent Programming**
  Spring 2014, lecture course, undergraduate elective, cross-listed
  Instructor for **CS5379 Parallel Processing**
  Spring 2014, lecture course, graduate elective, cross-listed
- Instructor for **CS4352 Operating Systems**
  Fall 2013, lecture course, undergraduate core
- Instructor for **CS4352 Operating Systems**
  Spring 2013, lecture course, undergraduate core
- Instructor for **CS5352 Advanced Operating System Design**
  Fall 2012, lecture course, graduate core
  with both face-to-face and distance sections
- Instructor for **CS4331 High Performance Computing**
  Summer I 2012, lecture course, undergraduate elective
  Cross-listed as **MATH4000 High Performance Computing**, Summer I 2012

- Instructor for **CS5331 Special Problems: Data Intensive Computing**
  Spring 2012, lecture course, graduate elective
- Instructor for **CS5379 Parallel Processing**
  Fall 2011, lecture course, graduate elective
  with both face-to-face and distance sections
- Instructor for **CS5331 Special Problems: Parallel I/O and Massive Storage Architectures**
  Spring 2011, lecture course, graduate elective

ADVISEES (WITH THESIS/DISSERTATION/RESEARCH ASSISTANTSHIP)/SCHOLARS SUPERVISED

- **Research Faculty/Postdoctoral Researcher**: Dr. Jiang Zhou
- **Graduate Students (current)**:
  - Mr. Misha Ahmadian, Ph.D. Student, Co-Chair
  - Mr. Ghazanfar Ali, Ph.D. Student, Chair
  - Ms. Elham Hojati, Ph.D. Student, Chair
  - Mr. Jie Li, Ph.D. Student, Chair
  - Mr. Chenxu Niu, Ph.D. Student, Chair
  - Mr. Brody Williams, Ph.D. Student, Chair
  - Ms. Ruonan Wu, Ph.D. Student, Chair
  - Ms. Harichandana Byna, M.S. Student, Chair
  - Mr. Yuanxu Xu, M.S. Student, Chair
  - Ms. Aditi Misra, M.S. Student, Chair
  - Ms. Nhi N Nguyen, M.S. Student, Chair
- **Undergraduate Students (including REU participants)**:
  - Mr. Casey Root
  - Mr. Cristiano Caon
- **High-school Students:** Yu Lim (Clark Scholar)
- **Visiting Scholars:** none currently
- **Past Students (my role, degree, graduation year, and first employment):**
  - Dr. Wei Zhang (Chair, Ph.D., 2021)
  - Dr. Xi Wang (Chair, Ph.D., 2020, RIOS (RISC-V International Open Source) Lab)
  - Dr. Daniel Dunning (Committee Member, Ph.D., 2020)
  - Dr. Moitrayee Chatterjee (Committee Member, Ph.D., 2019)
  - Dr. Jianjun Jeffrey Zheng (Committee Member, Ph.D., 2019)
  - Dr. Wei Xie (Chair, Ph.D., 2018, VMware Inc.)
  - Dr. Byungkwan Jung (Committee Member, Ph.D., 2018)
  - Dr. Ahmad O. Aseeri (Committee Member, Ph.D., 2018)
  - Dr. John Leidel (Chair, Ph.D., 2017, Chief Scientist, Tactical Computing Laboratories)
  - Dr. Timothy Sliwinski (Committee Member, Ph.D., 2017)
  - Dr. Bryant K. Nelson (Committee Member, Ph.D., 2016)
  - Dr. Chunchao Liang (Committee Member, Ph.D., 2016)
  - Dr. Cong Pu (Committee Member, Ph.D., 2016)
  - Dr. Jialin Liu (Chair, Ph.D., 2015, Lawrence Berkeley National Laboratory)
  - Dr. Yin Lu (Co-chair, Ph.D., 2015, Texas Tech University)

- Dr. Yugendra R Guvvala, (Committee Member, Ph.D., 2013, RAID Inc.)
- Nathan Stoddard (Chair, M.S., 2020)
- Dheeraj Kumar Cidda (Chair, M.S., 2020)
- Xueyuan Ji (Chair, M.S., 2020)
- Yevgeniy Shevekhman (Chair, M.S., 2020)
- Abhishek Kumar (Chair, M.S., 2019, Paycom)
- Vidya Eswarappa (Chair, M.S., 2018, Texas Tech University HPCC)
- Yan Mu (Chair, M.S., 2018, Paycom)
- Priyanka Kumari (Chair, M.S., 2017, TransUnion)
- Ali Nosrati (Chair, M.S., 2017, Dell)
- Yuan Cui (Chair, M.S., 2017)
- Misha Ahmadian (Co-chair, M.S., 2015, Texas Tech University)
- Prathamesh Amritkar (Chair, M.S., 2012, Sk Hynix)
- Junjie Chen (Chair, M.S., 2014, CitiBank)
- Quan Gan (Chair, M.S., 2012, Amazon)
- Qiang Gu (Chair, M.S., 2016)
- Mahesh Gurram (Chair, M.S., 2016, Bank of America)
- Soheil Mazaheri (Chair, M.S., 2015, FireEye)
- Navaneeth Thiagarajan (Chair, M.S., 2013, Texas Tech University)
- Kalaranjani Vijayakumar (Chair, M.S., 2015, Informatica)
- **Undergraduate students:** Claire Ancajas (B.S., 2016, GM), Jonathan Bastnagel (B.S., 2013; REU participant, Tyler Technologies), Kyle Blauer (B.S., 2016; Independent study, IBM), Bradly Crysler (B.S., 2013; Independent study, Balfour), Tristan Currens (B.S., 2014; Independent study, AT&T), Taylor Denison (B.S., 2013; Independent study, National Instruments), Teegan Duong (B.S., 2018, NI), Kace Echo (Mark Reyes) (REU participant, B.S., 2016, General Motors), Chaz George (REU site participant, Ellucian), Mose Gumble (REU site participant), Yi Guo (B.S., 2014, REU participant, graduate study at Columbia University), Zachary Hansen (B.S. 2020, Ph.D. study), Jayse Hulett (B.S., 2014, REU participant, Lockheed Martin), Rex Keen (B.S., 2015, REU participant, National Instruments), Dalton J. Koehl (B.S. 2019), Felix Perez (B.S., 2019, Graduate study at UTDallas), Omar Rodriguez (B.S., 2015; WCOE summer intern, graduate study at Texas Tech University), Noris G Rogers (B.S., 2015, Texas Tech University), Robin Ronson (B.S., 2018), Alejandro Sanjuan (exchange student), Shane Tarleton (B.S., 2013; Independent study, USAA), Gregory Thorsness (REU site participant), Eric Valenzuela (REU site participant), Alexander Weaver (B.S., 2017, NI), Yong Wu (B.S., 2018), Danielle Zaragoza (B.S., 2016)
- **High-school Students:** Zachary Hansen (Clark Scholar, enrolled B.S. in CS at TTU)

- **Past Research Faculty/Postdoctoral Researcher**: Dr. Dong Dai, now Assistant Professor at UNCC
- **Past Visiting Scholars:** Yusheng Hua (visiting Ph.D. student from Huazhong University of Science and Technology, China), Weihao Liang (visiting Ph.D. student from University of Science and Technology of China), Dr. Meng Li, Dr. Yan Li, Dr. Guangzhong Sun, Dr. Yang Zhang, Dr. Qinwen Zuo

BOOKS/JOURNALS EDITED

1. **Managing Guest Editor,** *Journal of Parallel and Distributed Computing by Elsevier*, *Special Issue on Parallel Processing*, 2021.
2. **Guest Editor,** *IEEE Access* by IEEE, *Special Issue on Artificial Intelligence in Parallel and Distributed*

Computing, with Dr. Songwen Pei, Dr. Tao Li, Dr. Junjie Wu, and Dr. Stéphane Zuckerman, 2020.

3. **Managing Guest Editor,** *Cluster Computing* by Springer, *Special Issue on Data-center Automation, Analytics, and Control*, with Dr. Alan Sill and Dr. Dong Dai, 2019.
4. **Guest Editor,** *Parallel Computing* by Elsevier, *Special Issue on Programming Models, Systems Software, and Tools for High-End Computing*, with Dr. Pavan Balaji and Dr. Abhinav Vishnu, 2018.
5. **Guest Editor,** *Parallel Computing* by Elsevier, *Special Issue on Programming Models, Systems Software, and Tools for High-End Computing*, with Dr. Pavan Balaji and Dr. Abhinav Vishnu, 2017.
6. **Managing Guest Editor,** *Journal of Supercomputing*, *Special Issue on Programming Models and Systems Software for High-End Computing*, with Dr. Abhinav Vishnu and Dr. Pavan Balaji, 2016.
7. **Guest Editor,** *Parallel Computing* by Elsevier, *Special Issue on Programming Models, Systems Software, and Tools for High-End Computing*, with Dr. Pavan Balaji and Dr. Abhinav Vishnu, 2016.
8. **Managing Guest Editor,** *International Journal of High Performance Computing Applications (IJHPCA)*, *Special Issue on Data-Intensive High Performance Computing*, with Drs. Philip Roth and Weikuan Yu, 2016.
9. **Guest Editor,** *Parallel Computing* by Elsevier, *Special Issue on Data Intensive Scalable Computing Systems*, with Dr. Philip Roth, 2015.
10. **Guest Editor,** *Journal of Supercomputing*, *Special Issue on Programming Models and Systems Software for High-End Computing*, with Dr. Abhinav Vishnu and Dr. Pavan Balaji, 2014.
11. **Managing Guest Editor,** *Cluster Computing* by Springer, *Special Issue on Data-Intensive High-Performance Computing*, with Dr. Xian-He Sun, 2014.
12. **Managing Guest Editor,** *Parallel Computing* by Elsevier, *Special Issue on Programming Models, Systems Software, and Tools for High-End Computing*, with Dr. Pavan Balaji and Dr. Abhinav Vishnu, 2013.
13. **Guest Editor,** *Journal of Supercomputing*, *Special Issue on Programming Models and Systems Software for High-End Computing*, with Dr. Abhinav Vishnu and Dr. Pavan Balaji, 2012.

SELECTED PUBLICATIONS

REFEREED JOURNAL PUBLICATIONS: (* GRADUATE ADVISEE; ** UNDERGRADUATE ADVISEE)

1. X. Wang*, A. Tumeo, J. D. Leidel, J. Li*, **Y. Chen**. HAM: Hotspot-Aware Manager for Improving Communications With 3D-Stacked Memory. *IEEE Transactions on Computers (TC)*, Volume: 70, Issue: 6, Page(s): 833-848, 2021. (DOI: https://doi.org/10.1109/TC.2021.3066982)
2. T. Dang, N. Nguyen and **Y. Chen**. HiperView: Real-time Monitoring Dynamic Behaviors of High Performance Computing Centers. *Journal of Supercomputing (JoS) by Springer*, March 2021. (DOI: https://doi.org/10.1007/s11227-021-03724-5)
3. J. Zhou, **Y. Chen**, D. Dai, Y. Zhuang, W. Wang. I/O Characteristic Discovery for Storage System Optimizations. In *Elsevier Journal of Parallel and Distributed Computing (JPDC)*, Volume 148, Pages 1-13, 2021. (DOI: https://doi.org/10.1016/j.jpdc.2020.08.005)
4. J. D. Leidel*, X. Wang*, B. Williams* and **Y. Chen**. Toward a Microarchitecture for Efficient Execution of Irregular Applications. *ACM Transactions on Parallel Computing (TOPC)*, Volume: 7, Issue: 4, 2020. (DOI: https://doi.org/10.1145/3418082)
5. J. Sun, G. Sun, S. Zhan, J. Zhang and **Y. Chen**. Automated Performance Modeling of HPC Applications Using Machine Learning. *IEEE Transactions on Computers (TC)*, Volume: 69, Issue: 5, Pages: 749 - 763, 2020. (DOI: https://doi.org/10.1109/TC.2020.2964767)
6. X. Shi, W. Liu, L. He, H. Jin, M. Li and **Y. Chen**. Optimizing the SSD Burst Buffer by Traffic Detection.

ACM Transactions on Architecture and Code Optimization (TACO), Volume: 17, Issue: 1, 2020. (DOI: https://doi.org/10.1145/3377705)

7. J. Zhou, **Y. Chen**, W. Xie*, S. He and W. Wang. PRS: A Pattern-Directed Replication Scheme for Heterogeneous Object-Based Storage. *IEEE Transactions on Computers (TC)*, Volume: 69, Issue: 4, Pages: 591 - 605, 2020. (DOI: https://doi.org/10.1109/TC.2019.2954089)
8. S. He, Z. Li, J. Zhou, Y. Yin, X. Xu, **Y. Chen** and X.-H. Sun. A Holistic Heterogeneity-Aware Data Placement Scheme in Hybrid Parallel I/O Systems. *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, Volume: 31, Issue: 4, Pages: 830 - 842, 2020. (DOI: https://doi.org/10.1109/TPDS.2019.2948901)
9. N. Zhao, G. Cao, W. Zhang, E.L. Samson and **Y. Chen**. Remote Sensing and Social Sensing for Socioeconomic Systems: A Comparison Study between Nighttime Lights and Location-based Social Media at the 500m Spatial Resolution. *International Journal of Applied Earth Observation and Geoinformation*, Volume: 87, 2020. (DOI: https://doi.org/10.1016/j.jag.2020.102058)
10. J. Zhou, **Y. Chen**, W. Wang, S. He and D. Meng. A Highly Reliable Metadata Service for Large-Scale Distributed File Systems. In *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, Volume: 31, Issue: 2, Pages: 374 - 392, 2019. (DOI: https://doi.org/10.1109/TPDS.2019.2937492)
11. D. Dai, **Y. Chen**, P. Carns, J. Jenkins, W. Zhang* and R. Ross. Managing Rich Metadata in High-Performance Computing Systems Using a Graph Model. In *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, Volume: 30, Issue: 7, Pages: 1613 - 1627, 2019. (DOI: https://ieeexplore.ieee.org/document/8580412)
12. W. Liang*, **Y. Chen**, J. Liu* and H. An. CARS: A Contention-Aware Scheduler for Efficient Resource Management of HPC Storage Systems. In *The International Journal of Parallel Computing (ParCo)*, Volume: 87, Pages: 25 - 34, 2019. (DOI: https://doi.org/10.1016/j.parco.2019.04.010)
13. N. Zhao, W. Zhang*, Y. Liu, E. Samson, **Y. Chen** and G. Cao. Improving Nighttime Light Imagery With Location-Based Social Media Data. In *IEEE Transactions on Geoscience and Remote Sensing*, 2019. (DOI: https://doi.org/10.1109/TGRS.2018.2871788)
14. Y. Hua*, X. Shi, H. Jin, W. Liu, Y. Jiang, **Y. Chen**, L. He. Software-defined QoS for I/O in Exascale Computing. In CCF *Transactions on High Performance Computing*, Apr. 2019. (DOI: https://doi.org/10.1007/s42514-019-00005-9)
15. N. Tavakoli*, D. Dai, and **Y. Chen**. Client-side Straggler-Aware I/O Scheduler for Object-based Parallel File Systems. In *The International Journal of Parallel Computing (ParCo)*, Volume: 82, Pages: 3 - 18, 2018. (DOI: https://doi.org/10.1016/j.parco.2018.07.001)
16. W. Xie*, **Y. Chen**, and P.C. Roth. Exploiting Internal Parallelism for Address Translation in Solid State Drives. In the *ACM Transactions on Storage (TOS)*, Volume: 14, Issue: 4, Article: 32, 2018. (DOI: https://doi.org/10.1145/3239564)
17. D. Dai, **Y. Chen**, D. Kimpe, and R. Ross. Trigger-based Incremental Data Processing with Unified Sync and Async Model. In the *IEEE Transactions on Cloud Computing*, 2018. (DOI: https://doi.org/10.1109/TCC.2018.2830348)
18. D. Dai, F.S. Bao, J. Zhou, X. Shi and **Y. Chen**. Vectorizing Disk Blocks for Efficient Storage Systems via Deep Learning. In *The International Journal of Parallel Computing (ParCo)*, Volume: 82, Pages: 75 - 90, 2018. (DOI: https://doi.org/10.1016/j.parco.2018.03.003)
19. J. Leidel* and **Y. Chen**. HMC-Sim-2.0: A Co-Design Infrastructure for Exploring Custom Memory Cube Operations, In *The International Journal of Parallel Computing (ParCo)*, Volume 68, Pages 77-88, 2017. (DOI: https://doi.org/10.1016/j.parco.2017.07.008)

20. M. Ahmadian*, Y. Zhuang, W. L. Hase and **Y. Chen**. Data Reduction through Increased Data Utilization in Chemical Dynamics Simulations, *Elsevier Journal of Big Data Research (BDR)*, Volume 9, Pages 57-66, September 2017. (DOI: https://doi.org/10.1016/j.bdr.2017.06.005)
21. W. Xie*, **Y. Chen** and P. C. Roth. ASA-FTL: An Adaptive Separation Aware Flash Translation Layer for Solid State Drives. In *The International Journal of Parallel Computing (ParCo)*, Volume 61, Pages 3-17, January 2017. (DOI: https://doi.org/10.1016/j.parco.2016.10.006)
22. D. Dai, P. Carns, R. Ross, J. Jenkins, N. Muirhead**, **Y. Chen**. An Asynchronous Traversal Engine for Graph-based Rich Metadata Management. In the *International Journal of Parallel Computing (ParCo)*, Volume: 58, Pages: 140–156, 2016. (DOI: https://doi.org/10.1016/j.parco.2016.06.002)
23. J.L. Liu* and **Y. Chen**. Segmented In-Advance Data Analytics for Fast Scientific Discovery. *IEEE Transactions on Cloud Computing*, Volume: 8, Issue: 2, Pages: 432 - 442, 2016. (DOI: https://doi.org/10.1109/TCC.2016.2541142)
24. J.L. Liu*, Y. Zhuang, **Y. Chen**. Hierarchical Collective I/O Scheduling for High-Performance Computing. *Journal of Big Data Research (BDR)*, by Elsevier. Vol. 2, Issue 3, pages 117–126, 2015. (ACM DL/IEEE Xplore DL: http://ieeexplore.ieee.org/xpl/login.jsp?tp=&arnumber=6546095)
25. Y. Lu*, **Y. Chen**, R. Thakur, and Y. Zhuang. Collective I/O under Memory Constraints. The *International Journal of High Performance Computing Applications (IJHPCA)*, Vol. 29, No. 1, pages 21-36, 2015. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2733365)
26. **Y. Chen**, Y. Lu*, P. Amritkar*, R. Thakur, and Y. Zhuang. Performance Model Directed Data Sieving for High Performance I/O. In *The Journal of Supercomputing*, doi:10.1007/s11227-014-1277-8, Vol. 71, Issue 6, pages 2066-2090, 2015. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2786689)
27. J. Leidel*, **Y. Chen**. HMC-SIM: A Simulation Framework for Hybrid Memory Cube Devices. *Journal of Parallel Processing Letters*, Vol. 24, Issue 04, 22 pages, 2014. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2673072)
28. X. Shi, M. Chen, L. He, X. Xie, L. Lu, H. Jin, **Y. Chen**, and S. Wu. Mammoth: Gearing Hadoop Towards Memory-Intensive MapReduce Applications. In *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, Vol. 26, Issue 8, pages 2300-2315, 2014. (ACM DL/IEEE Xplore DL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6869021)
29. H. Song, Y. Yin, **Y. Chen** and X.-H. Sun. Cost-intelligent Application-specific Data Layout Optimization for Parallel File Systems. *Cluster Computing*, doi:10.1007/s10586-012-0200-4, Vol. 16, Issue 2, pages 285-298, 2013. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2505184)
30. **Y. Chen**, H. Zhu, H. Jin, and X.-H. Sun. Algorithm-level Feedback-controlled Adaptive Data Prefetcher: Accelerating Data Access for High-Performance Processors. *International Journal of Parallel Computing (ParCo)*, Vol. 38, Issues 10-11, pages 533-551, 2012. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2397091)
31. **Y. Chen**, H. Zhu, P. C. Roth, H. Jin and X.-H. Sun. Global-aware and Multi-order Context-based Prefetching for High-Performance Processors. *International Journal of High Performance Computing Applications (IJHPCA)*, Vol. 25, No. 4, 355-370, 2011. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2076560)
32. X.-H. Sun and **Y. Chen**. Reevaluating Amdahl's Law in the Multicore Era. *Journal of Parallel and Distributed Computing (JPDC)*, Vol. 70, No. 2, 183 – 188, 2010. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=1688079)
33. S. Byna, **Y. Chen** and X.-H. Sun. Taxonomy of Data Prefetching Strategies for Multicore Processors.

Journal of Computer Science and Technology (JCST), Vol. 24, No. 3, 405 – 417, 2009. (ACM DL/IEEE Xplore DL: http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=4520189)

34. **Y. Chen**, X.-H. Sun and M. Wu. Algorithm-System Scalability of Heterogeneous Computing. *Journal of Parallel and Distributed Computing (JPDC)*, Vol. 68, No. 11, 1403 – 1412, 2008. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=1435136)
35. X.-H. Sun, S. Byna and **Y. Chen**. Server-based Data Push Architecture for Multi-processor Environments. *Journal of Computer Science & Technology (JCST)*, Vol. 22, No. 5, 641–652, 2007. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=1375915)
36. **Y. Chen**, G. Chen, C. Li and J. He. Research on Hybrid Programming Model for SMP Cluster. *Mini-Micro Systems*, Vol. 25, No. 10, 1763 – 1767, 2004. (In Chinese)
37. **Y. Chen**, C. Li, H. An, Q. Zheng and Z. Chen. Remote Parallel Debugger Based on Dawning3000 Parallel Machine. *Journal of Computer Science*, Vol. 31, No. 3, 179 – 182, 2004. (In Chinese)
38. **Y. Chen**, K. He, Z. Lu and C. Yan. Design and Implementation of a Parallel Debugger for Cluster System. *Journal of Computer Engineering*, Vol. 30, No. 9, 50 – 52, 2004. (In Chinese)

REFEREED CONFERENCE/WORKSHOP PUBLICATIONS (ACCEPTANCE RATE IS LISTED WHEN KNOWN) (* GRADUATE ADVISEE; ** UNDERGRADUATE ADVISEE):

39. X. Wang*, J. D. Leidel, B. Williams*, A. Ehret, M. Mark, M. Kinsy, **Y. Chen**. xBGAS: A Global Address Space Extension on RISC-V for High Performance Computing. In *Proceedings of The 35th IEEE International Parallel & Distributed Processing Symposium (IPDPS'21)*, Portland, Oregon USA, 2021. (Acceptance rate: 105/462=22.7%) **Best Paper Award** (one out of 105 accepted papers)
40. B. Williams*, J. Leidel, X. Wang*, D. Donofrio and **Y. Chen**. CircusTent: A Benchmark Suite for Atomic Memory Operations. In *Proceedings of the ACM International Symposium on Memory Systems (MEMSYS'20)*, 2020. (DOI: https://dl.acm.org/doi/fullHtml/10.1145/3422575.3422789)
41. J. Li*, G. Ali*, N. Nguyen, J. Hass, A. Sill, T. Dang and **Y. Chen**. MonSTer: An Out-of-the-Box Monitoring Tool for High Performance Computing Systems, In *Proceedings of IEEE International Conference on Cluster Computing (CLUSTER'20)*, Pages: 119 - 129, 2020. (DOI: https://doi.org/10.1109/CLUSTER49012.2020.00022)
42. N. Nguyen, J. Hass, **Y. Chen**, J. Li*, A. Sill and T. Dang. RadarViewer: Visualizing the Dynamics of Multivariate Data. In *Proceedings of Practice and Experience in Advanced Research Computing Conference Series (PEARC'20)*, 2020. (DOI: https://doi.org/10.1145/3311790.3404538)
43. X. Wang*, J. D. Leidel, B. Williams* and **Y. Chen**. PAC: Paged Adaptive Coalescer for 3D-Stacked Memory. In *Proceedings of The 29th International ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC'20)*, Pages: 137 - 148, 2020. (DOI: https://doi.org/10.1145/3369583.3392670)
44. X. Wang*, B. Williams*, J. D. Leidel, A. Ehret, M. Kinsy and **Y. Chen**. Remote Atomic Extension (RAE) for Scalable High Performance Computing. In *Proceedings of The 57th Design Automation Conference (DAC'20)*, Pages: 1 - 6, 2020. (DOI: https://doi.org/10.1109/DAC18072.2020.9218589)
45. S. Liang*, Z. Yang, F. Jin and **Y. Chen**. Data Centers Job Scheduling with Deep Reinforcement Learning. In *Proceedings of The 24th Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD'20)*, Pages: 906 - 917 2020. (DOI: https://doi.org/10.1007/978-3-030-47436-2_68)
46. E. Hojati*, J. Hass, A. Sill and **Y. Chen**. Redfish Green500 Benchmarker (RGB): Towards Automation of the Green500 Process for Data Centers. In *Proceedings of 2020 IEEE Green Technologies Conference*, 2020. (DOI: https://doi.org/10.1109/GreenTech46478.2020.9289729)

47. W. Zhang*, S. Byna, C. Niu*, **Y. Chen**. Exploring Metadata Search Essentials for Scientific Data Management. In *Proceedings of 26th IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC '19)*, Pages: 83 - 92, 2019. (DOI: https://doi.org/10.1109/HiPC.2019.00021)
48. N. Nguyen, **Y. Chen**, J. Hass and T. Dang. HiperJobViz: Visualizing Resource Allocations in High-Performance Computing Center via Multivariate Health-Status Data. In *Proceedings of the 3rd Industry/University Joint International Workshop on Data-Center Automation, Analytics, and Control (DAAC'19)*, Pages: 19 - 24, 2019. (DOI: https://doi.org/10.1109/DAAC49578.2019.00009)
49. W. Zhang*, S. Byna, H. Tang, B. Williams* and **Y. Chen**. MIQS: Metadata Indexing and Querying Service for Self-describing File Formats. In *The Proceedings of The 31st ACM/IEEE Supercomputing Conference (SC'19)*, Denver, CO, 2019. (first-around acceptance rate: 72/344=21%, another 15 papers being asked for major revisions per SC'19) (DOI: https://doi.org/10.1145/3295500.3356146)
50. V. Pham, N. Nguyen, J. Li*, J. Hass, Y. Chen and T. Dang. MTSAD: Multivariate Time Series Abnormality Detection and Visualization. In *Proceedings of IEEE International Conference on Big Data (IEEE BigData)*, Pages: 3267 - 3276, 2019. (DOI: https://doi.org/10.1109/BigData47090.2019.9006559)
51. Y. Li, J. Zhou, W. Wang and **Y. Chen**. RE-Store: Reliable and Efficient KV-Store with Erasure Coding and Replication. In *The Proceedings of IEEE International Conference on Cluster Computing (Cluster'19)*, 2019. (DOI: https://doi.org/10.1109/CLUSTER.2019.8891013)
52. J. Li*, X. Wang*, A. Tumeo, B. Williams*, J. D. Leidel and **Y. Chen**. PIMS: A Lightweight Processing-in-Memory Accelerator for Stencil Computations. In *Proc. of the ACM International Symposium on Memory Systems (MEMSYS'19)*, 2019. (DOI: https://doi.org/10.1145/3357526.3357550)
53. W. Liang*, **Y. Chen** and H. An. Interference-Aware I/O Scheduling for Data-Intensive Applications on Hierarchical HPC Storage Systems. In Proc. of *The 21st IEEE International Conference on High Performance Computing and Communications (HPCC'19)*, 2019. (DOI: https://doi.org/10.1109/HPCC/SmartCity/DSS.2019.00099)
54. X. Wang*, A. Tumeo, J. Leidel, J. Li* and **Y. Chen**. MAC: Memory Access Coalescer for 3D-Stacked Memory. In the *Proc. of the 48th International Conference on Parallel Processing (ICPP'19)*, 2019. (Acceptance rate: 106/405=26%) (DOI: https://doi.org/10.1145/3337821.3337867)
55. B. Williams*, X. Wang*, J. Leidel and **Y. Chen**. Collective Communication for the RISC-V xBGAS ISA Extension. In the Proceedings of the *12th International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2)*, held in conjunction with the *48th International Conference on Parallel Processing (ICPP)*, 2019. (DOI: https://doi.org/10.1145/3339186.3339196)
56. M. Ahmadian*, E. Rees, Y. Zhuang and **Y. Chen**. Reducing Faulty Jobs by Job Submission Verifier in Grid Engine. In Proc. of *The 2019 ACM Practice and Experience in Advanced Research Computing (PEARC'19)*, 2019. (DOI: https://doi.org/10.1145/3332186.3338408)
57. W. Zhang*, H. Tang, S. Byna and **Y. Chen**. DART: Distributed Adaptive Radix Tree for Efficient Affix-based Keyword Search on HPC Systems. In the Proc. of *The 27th International Conference on Parallel Architectures and Compilation Techniques (PACT'18)*, 2018. (Acceptance rate: 36/126=28.6%) (DOI: https://doi.org/10.1145/3243176.3243207)
58. J. Zhou, **Y. Chen** and W. Wang. Attributed Consistent Hashing for Heterogeneous Storage Systems. In the *Proc. of The 27th International Conference on Parallel Architectures and Compilation Techniques (PACT'18)*, 2018. (Acceptance rate: 36/126=28.6%) (DOI: https://doi.org/10.1145/3243176.3243202)
59. W. Liang*, J. Liu, H. An and **Y. Chen**. Contention-Aware Resource Scheduling for Burst Buffer Systems. In the *Proceedings of the 11th International Workshop on Parallel Programming Models and Systems Software*

for High-End Computing (P2S2), held in conjunction with the 47th International Conference on Parallel Processing (ICPP), 2018. (DOI: https://doi.org/10.1145/3229710.3229718)

60. X. Wang*, J. Leidel* and **Y. Chen**. Memory Coalescing for Hybrid Memory Cube. In the *Proc. of the 47th International Conference on Parallel Processing (ICPP'18)*, 2018. (DOI: https://doi.org/10.1145/3225058.3225062)
61. J. Leidel*, X. Wang* and **Y. Chen**. GoblinCore-64: A RISC-V Based Architecture for Data Intensive Computing. In the *Proc. of The 22nd IEEE High Performance Extreme Computing Conference (HPEC'18)*, 2018. (DOI: https://doi.org/10.1109/HPEC.2018.8547560)
62. J. Cao, O. Rameshwar, M. Zheng, D. Dai, V. Eswarappa, Y. Mu and **Y. Chen**. PFault: A General Framework for Analyzing the Reliability of High-Performance Parallel File Systems. In *The Proc. of the 2018 ACM International Conference on Supercomputing (ICS'18)*, 2018. (Acceptance rate: 36/193=18.7%) (DOI: https://doi.org/10.1145/3205289.3205302)
63. J. Zhou, D. Dai, Y. Mao, X. Chen, Y. Zhuang and **Y. Chen**. I/O Characteristics Discovery in Cloud Storage Systems. In The *Proc. of the 2018 IEEE International Conference on Cloud Computing (Cloud'18)*, 2018. (Acceptance rate: 18%) (DOI: https://doi.org/10.1109/CLOUD.2018.00029)
64. W. Zhang*, **Y. Chen** and D. Dai. AKIN: A Streaming Graph Partitioning Algorithm for Distributed Graph Storage Systems. In *The Proceedings of The 18th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*, 2018. (Acceptance rate: 20.8%) (DOI: https://doi.org/10.1109/CCGRID.2018.00033)
65. J. Leidel*, X. Wang* and **Y. Chen**. Pressure-Driven Hardware Managed Thread Concurrency for Irregular Applications. In the *Proc. of the Seventh Workshop on Irregular Applications: Architectures and Algorithms (IA^3-2017), in conjunction with ACM/IEEE Supercomputing (SC'17)*, 2017. (DOI: https://doi.org/10.1145/3149704.3149705)
66. J. Sun, S. Zhan, G. Sun and **Y. Chen**. Automated Performance Modeling Based on Runtime Feature Detection and Machine Learning. In the *Proceedings of 15th IEEE International Symposium on Parallel and Distributed Processing with Applications (ISPA'17)*, 2017. (DOI: https://doi.org/10.1109/ISPA/IUCC.2017.00115)
67. D. Dai, **Y. Chen**, P. Carns, J. Jenkins and R. Ross. Lightweight Provenance Service for High Performance Computing. In the *Proc. of The 26th International Conference on Parallel Architectures and Compilation Techniques (PACT'17)*, 2017. (Acceptance rate: 25/108=23%) (DOI: https://doi.org/10.1109/PACT.2017.14)
68. X. Wang*, J. Leidel* and **Y. Chen**. OpenMP Memkind: An Extension for Heterogeneous Physical Memories. In the *Proceedings of the 10th International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2), held in conjunction with the 46th International Conference on Parallel Processing (ICPP)*, 2017. (DOI: https://doi.org/10.1109/ICPPW.2017.40)
69. D. Dai, W. Zhang and **Y. Chen**. IOGP: An Incremental Online Graph Partitioning Algorithm for Distributed Graph Databases. In the *Proc. of The 26th ACM International Symposium on High Performance Parallel and Distributed Computing (HPDC'17)*, 2017. (Acceptance rate: 19%) (DOI: https://doi.org/10.1145/3078597.3078606)
70. X. Shi, M. Li, W. Liu, H. Jin, C. Yu and **Y. Chen**. SSDUP: A Traffic-Aware SSD Burst Buffer for HPC Systems. In the *Proc. of the 2017 ACM International Conference on Supercomputing (ICS'17)*, 2017. (Acceptance rate: 28/?) (DOI: https://doi.org/10.1145/3079079.3079087)
71. Y. Yu, H. An, J. Chen, W. Liang, Q. Xu and **Y. Chen**. Pipelining Computation and Optimization Strategies for Scaling GROMACS on the Sunway Many-Core Processor. In the *Proc. of the 17th International*

Conference on Algorithms and Architectures for Parallel Processing (ICA3PP-2017), 2017. (DOI: https://doi.org/10.1007/978-3-319-65482-9_2)

72. J. Zhou, W. Xie*, D. Dai and **Y. Chen**. Pattern-Directed Replication Scheme for Heterogeneous Object-based Storage. In The *Proceedings of The 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*, 2017. (DOI: https://doi.org/10.1109/CCGRID.2017.148)
73. W. Xie* and **Y. Chen**. Elastic Consistent Hashing for Distributed Storage Systems. In the *Proc. of The 31st IEEE International Parallel and Distributed Processing Symposium (IPDPS'17)*, 2017. (Acceptance rate: 23%) (DOI: https://doi.org/10.1109/IPDPS.2017.88)
74. J. Leidel*, X. Wang* and **Y. Chen**. Toward a Memory-Centric, Stacked Architecture for Extreme-Scale, Data-Intensive Computing. In the *Proceedings of the First Workshop on Pioneering Processor Paradigms (WP3'17), in conjunction with The 23rd IEEE Symposium on High Performance Computer Architecture (HPCA'17)*, 2017.
75. J. Cao, S. Wang, D. Dai, M. Zheng and **Y. Chen**. A Generic Framework for Testing Parallel File Systems. In the *Proceedings of the Joint International Workshop on Parallel Data Storage and Data Intensive Scalable Computing Systems (PDSW-DISCS'16),* in conjunction with *ACM/IEEE Supercomputing (SC'16)*, 2016. (DOI: https://doi.org/10.1109/PDSW-DISCS.2016.013)
76. D. Dai, **Y. Chen**, P. Carns, J. Jenkins, W. Zhang and R. Ross. GraphMeta: A Graph-based Engine for Managing Large-Scale HPC Rich Metadata. In the *Proc. of the IEEE International Conference on Cluster Computing (Cluster'16)*, 2016. (DOI: https://doi.org/10.1109/CLUSTER.2016.50)
77. **Y. Chen**, C. Chen*, Y. Yin, X. Sun, R. Thakur and W. Gropp. Rethinking High Performance Computing System Architecture for Scientific Big Data Applications. In *Proceedings of 14th IEEE International Symposium on Parallel and Distributed Processing with Applications (ISPA'16)*, 2016. ***Best Paper Award***. *(DOI: https://doi.org/10.1109/TrustCom.2016.0248)*
78. J. Zhou, W. Xie*, Q. Gu* and **Y. Chen**. Hierarchical Consistent Hashing for Heterogeneous Object-based Storage. In *Proceedings of 14th IEEE International Symposium on Parallel and Distributed Processing with Applications (ISPA'16)*, 2016. *(DOI: https://doi.org/10.1109/TrustCom.2016.0247)*
79. C. Chen*, M. Lang, L. Ionkov and **Y. Chen**. Active Burst-Buffer: In-Transit Processing Integrated into Hierarchical Storage. In *Proceedings of the 11th IEEE International Conference on Networking, Architecture, and Storage (NAS'16)*, 2016. ***Best Paper Award.*** (DOI: https://doi.org/110.1109/NAS.2016.7549390)
80. J. Zhou, W. Xie*, J. Noble, K. Echo** and **Y. Chen**. SUORA: A Scalable and Uniform Data Distribution Algorithm for Heterogeneous Storage Systems. In *Proceedings of the 11th IEEE International Conference on Networking, Architecture, and Storage (NAS'16)*, 2016. *(DOI: https://doi.org/10.1109/NAS.2016.7549423)*
81. W. Xie*, **Y. Chen** and P. Roth. Parallel-DFTL: A Flash Translation Layer that Exploits Internal Parallelism in Solid State Drives. In *Proceedings of the 11th IEEE International Conference on Networking, Architecture, and Storage (NAS'16)*, 2016. ***Best Paper Finalist.*** *(DOI: https://doi.org/10.1109/NAS.2016.7549413)*
82. D. Dai, F. S. Bao, J. Zhou, **Y. Chen**. Block2Vec: A Deep Learning Strategy on Mining Block Correlations in Storage Systems. In *Proceedings of the Ninth International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2)*, held in conjunction with *the 45th International Conference on Parallel Processing (ICPP)*, 2016. (DOI: https://doi.org/10.1109/ICPPW.2016.43)
83. N. Tavakoli*, D. Dai and **Y. Chen**. Log-assisted Straggler-aware I/O Scheduler for High-End Computing. In *Proceedings of the Ninth International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2)*, held in conjunction with *the 45th International Conference on Parallel Processing (ICPP)*, 2016. (DOI: https://doi.org/10.1109/ICPPW.2016.38)

84. X. Wang*, J. Leidel* and **Y. Chen**. Concurrent Dynamic Memory Coalescing on GoblinCore-64 Architecture. In *Proceedings of the 2016 International Symposium on Memory Systems (MEMSYS'16)*, 2016. (DOI: https://doi.org/10.1145/2989081.2989128)
85. J. Leidel* and **Y. Chen**. Exploring Tag-Bit Memory Operations in Hybrid Memory Cubes. In *Proceedings of the 2016 International Symposium on Memory Systems (MEMSYS'16)*, 2016. (DOI: https://doi.org/10.1145/2989081.2989105)
86. J. Leidel* and **Y. Chen**. HMC-Sim-2.0: A Simulation Platform for Exploring Custom Memory Cube Operations. In *the Proceedings of the Sixth International Workshop on Accelerators and Hybrid Exascale Systems (ASHES-2016)*, 2016. (DOI: https://doi.org/10.1109/IPDPSW.2016.43)
87. W. Xie*, **Y. Chen**, and P. Roth. A Low-cost Adaptive Separation Method for the Flash Translation Layer of Solid State Drives. In the *Proceedings of The 2015 International Workshop on Data-Intensive Scalable Computing Systems (DISCS-2015)*, held in conjunction with *The International Conference for High Performance Computing, Networking, Storage and Analysis (SC'15)*, 2015. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2831250)
88. W. Xie*, J. Zhou, M. Reyes (K. Echo)**, J. Noble, and **Y. Chen**. Two-Mode Data Distribution Scheme for Heterogeneous Storage in Data Centers (short paper). In the *Proceedings of The 2015 IEEE International Conference on Big Data, (BigData'15)*, 2015. (Acceptance rate: 17% full papers and 18% short papers accepted out of 363 complete submissions). (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2878056)
89. J. Leidel* and **Y. Chen**. Communication Avoiding Power Scaling. In the *Proc. of The Eighth Intl. Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2)*, in conjunction with *The 44th International Conference on Parallel Processing (ICPP'15)*, 2015. (ACM DL/IEEE Xplore DL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=7349908)
90. J.L. Liu*, **Y. Chen**, and S. Byna. Collective Computing for Scientific Big Data Analysis. In the *Proc. of The Eighth International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2)*, in conjunction with *The 44th International Conference on Parallel Processing (ICPP'15)*, 2015. (ACM DL/IEEE Xplore DL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=7349904)
91. J. Zhou, **Y. Chen**, X. Gu, W. Wang, and D. Meng. A Virtual Shared Metadata Storage for HDFS. In the *Proc. of The 10th IEEE International Conference on Networking, Architecture, and Storage (NAS'15)*, 2015. (Acceptance rate: 28/94 = 29.8%) (ACM DL/IEEE Xplore DL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=7255195)
92. D. Dai, P. Carns, R. Ross, J. Jenkins, K. Blauer**, and **Y. Chen**. GraphTrek: Asynchronous Graph Traversal for Property Graph Based Metadata Management. In the *Proc. of the IEEE International Conference on Cluster Computing, (Cluster'15), 2015*. (Acceptance rate: 38/157=24.2%) (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2859584)
93. J. Zhou, **Y. Chen**, W. Wang and D. Meng. MAMS: A Highly Reliable Policy for Metadata Service. In the *Proc. of the 44th International Conference on Parallel Processing (ICPP'15)*, 2015. (Acceptance rate: 99/305=32.5%) (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2881115)
94. D. Dai, R. Ross, P. Carns, D. Kimpe, and **Y. Chen**. Using Property Graphs for Rich Metadata Management in HPC Systems. In *The 9th Parallel Data Storage Workshop* held in conjunction with SC'14, (PDSW'14), 2014. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2688403)
95. D. Dai, **Y. Chen**, D. Kimpe, and R. Ross. Provenance-Based Object Storage Prediction Scheme for Scientific Big Data Applications. In the Proceedings of *The 2014 IEEE International Conference on Big Data,*

(BigData'14), 2014. (Acceptance rate: 49/264=18.6%). (ACM DL/IEEE Xplore DL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=7004242)

96. J.L. Liu*, Y. Lu*, and **Y. Chen**. In-advance Data Analytics for Reducing Time to Discovery (short paper). In the Proceedings of The 2014 IEEE International Conference on Big Data, (BigData'14), 2014. (Acceptance rate: 49 full papers and 57 short papers accepted out of 264 complete submissions). (ACM DL/IEEE Xplore DL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=7004249)
97. D. Dai, **Y. Chen**, D. Kimpe, and R. Ross. Two-Choice Randomized Dynamic I/O Scheduler for Object Storage Systems. In *The Proceedings of the ACM/IEEE Supercomputing Conference (SC'14)*, 2014. (Acceptance rate: 82/394=20.8%). (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2683663)
98. C. Chen*, **Y. Chen**, K. Feng, Y. Yin, H. Eslami, R. Thakur, X.-H. Sun, and W.D. Gropp. Decoupled I/O for Data-Intensive High Performance Computing. In the *Proceedings of the 7th International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2)*, in conjunction with *The 43rd International Conference on Parallel Processing (ICPP)*, 2014. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2863265)
99. J.J. Chen*, J.L. Liu*, P. Roth and **Y. Chen**. Using Working Set Reorganization to Manage Storage Systems with Hard and Solid State Disks. In the *Proceedings of the 7th International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2)*, in conjunction with *The 43rd International Conference on Parallel Processing (ICPP)*, 2014. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2863242)
100. D. Dai, X. Zhou, D. Kimpe, R. Ross, and **Y. Chen**. Domino: An Incremental Computing Framework in Cloud with Eventual Synchronization (short paper). In the Proceedings of *The 23rd ACM International Symposium on High-Performance Parallel and Distributed Computing (HPDC'14)*, 2014. (Acceptance rate: 21 full papers and 16 short papers accepted out of 130) (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2600705)
101. Y. Lu*, **Y. Chen**, R. Latham and Y. Zhuang. Revealing Applications' Access Pattern in Collective I/O for Cache Management. In the Proceedings of *The 2014 ACM International Conference on Supercomputing (ICS'14)*, 2014. (Acceptance rate: 34/162=21%) (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2597686)
102. J. Leidel* and **Y. Chen**. HMC-Sim: A Simulation Framework for Hybrid Memory Cube Devices. In The Proceedings of *The 2014 Workshop on Large-Scale Parallel Processing (LSPP), in conjunction with the 28th IEEE International Parallel & Distributed Processing Symposium (IPDPS'14)*, 2014. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2673072)
103. J.L. Liu*, S. Byna, B. Dong, K. Wu, and **Y. Chen**. Model-driven Data Layout Selection for Improving Read Performance. In The *Proceedings of The 2014 International Workshop on High Performance Data Intensive Computing (HPDIC2014), in conjunction with the 28th IEEE International Parallel & Distributed Processing Symposium (IPDPS'14)*, 2014. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2672934)
104. Z. Wang, X. Shi, H. Jin, S. Wu, and **Y. Chen**. Iteration Based Collective I/O Strategy for Parallel I/O Systems. In the *Proc. of the 14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*, 2014. (Acceptance rate: 54/283=19.1%) (ACM DL/IEEE Xplore DL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6846464)
105. J.L. Liu*, B. Crysler**, Y. Lu*, and **Y. Chen**. Locality-driven High-level I/O Aggregation for Processing Scientific Datasets (regular paper). In the *Proceedings of The 2013 IEEE International Conference on Big*

Data, (BigData'13), 2013. (Acceptance rate: 45/259=17.4%) (ACM DL/IEEE Xplore DL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6691560)

106. C. Chen*, M. Lang, and **Y. Chen**. Multilevel Active Storage for Big Data Applications in High Performance Computing (short paper). In the *Proceedings of The 2013 IEEE International Conference on Big Data, (BigData'13)*, 2013. (Acceptance rate: 54/259=20.8%) (ACM DL/IEEE Xplore DL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6691570)
107. J.J. Chen*, P. Roth, and **Y. Chen**. Using Pattern-Models to Guide SSD Deployment for Big Data in HPC systems (short paper). In the *Proceedings of The 2013 IEEE International Conference on Big Data, (BigData'13)*, 2013. (Acceptance rate: 54/259=20.8%) (ACM DL/IEEE Xplore DL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6691592)
108. J.L. Liu*, S. Byna and **Y. Chen**. Segmented Analysis for Reducing Data Movement (short paper). In the *Proceedings of The 2013 IEEE International Conference on Big Data, (BigData'13)*, 2013. (Acceptance rate: 54/259=20.8%) (ACM DL/IEEE Xplore DL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6691594)
109. J.L. Liu* and **Y. Chen**. Fast Data Analysis with Integrated Statistical Metadata in Scientific Datasets. In the *Proc. of the IEEE International Conference on Cluster Computing, (Cluster'13)*, 2013. (Acceptance rate: 31%) (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2410523)
110. C. Chen*, J. Bastnagel**, and **Y. Chen**. Data Deduplication in a Hybrid Architecture for Improving Write Performance. In the *Proc. of the 3rd International Workshop on Runtime and Operating Systems for Supercomputers (ROSS), in conjunction with International Conference on Supercomputing (ICS)*, June 2013. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2481435)
111. Y. Lu*, **Y. Chen**, R. Thakur, and Y. Zhuang. Memory-Conscious Collective I/O for Extreme Scale HPC Systems. In the *Proc. of the 3rd International Workshop on Runtime and Operating Systems for Supercomputers (ROSS), in conjunction with International Conference on Supercomputing (ICS)*, June 2013. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2481430)
112. J.L. Liu*, **Y. Chen**, and Y. Zhuang. Hierarchical I/O Scheduling for Collective I/O. In the *Proc. of the 13th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*, 2013. (Acceptance rate: 57/257=22.18%) (ACM DL/IEEE Xplore DL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6546095)
113. J.L. Liu*, **Y. Chen**. Improving Data Analysis Performance for High-Performance Computing with Integrating Statistical Metadata in Scientific Datasets. In the *Proc. of the Second Annual Workshop on High-Performance Computing meets Databases (HPCDB)*, in conjunction with the *ACM/IEEE Supercomputing Conference (SC'12),* 2012. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2476982)
114. **Y. Chen**, C. Chen*, X.-H. Sun, W. D. Gropp, and R. Thakur. A Decoupled Execution Paradigm for Data-Intensive High-End Computing. In the *Proc. of the IEEE International Conference on Cluster Computing 2012 (Cluster'12),* 2012. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2409142)
115. C. Chen*, **Y. Chen**, and P. C. Roth. DOSAS: Mitigating the Resource Contention in Active Storage Systems. In the *Proc. of the IEEE International Conference on Cluster Computing 2012 (Cluster'12),* 2012. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2409138)
116. C. Chen* and **Y. Chen**. Dynamic Active Storage for High Performance I/O. In the *Proc. of the 41st International Conference on Parallel Processing (ICPP'12)*, 2012. (Acceptance rate: 53/187=28.3%). (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2410597)

117. H. Jin, J. Ji, X.-H. Sun, **Y. Chen** and R. Thakur. CHAIO: Enabling HPC Applications on Data-Intensive File Systems. In the *Proc. of the 41st International Conference on Parallel Processing (ICPP'12)*. (Acceptance rate: 53/187=28.3%). (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2410596)
118. **Y. Chen**, H. Zhu, H. Jin, and X.-H. Sun. Storage-Efficient Data Prefetching for High Performance Computing. In the *Proc. of the 7th International Conference on Future Information Technology (FutureTech'12)*, Vancouver, Canada, 2012. (Acceptance rate: 92/266=34.6%). (http://link.springer.com/chapter/10.1007/978-94-007-4516-2_11)
119. Y. Lu*, **Y. Chen**, P. Amritkar*, R. Thakur, and Y. Zhuang. A New Data Sieving Approach for High Performance I/O. In the *Proc. of the 7th International Conference on Future Information Technology (FutureTech'12)*, Vancouver, Canada, 2012. (Acceptance rate: 92/266=34.6%). (**Best Paper Award**) (http://link.springer.com/chapter/10.1007/978-94-007-4516-2_12)
120. H. Jin, T. Ke, **Y. Chen** and X.-H. Sun. Checkpointing Orchestration: Toward a Scalable HPC Fault-Tolerant Environment. In the *Proc. of IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*, May, 2012. (Acceptance rate: 83/302=27.5%). (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2310186)
121. **Y. Chen**. Towards Scalable I/O Architecture for Exascale Systems. In the *Proc. of the 4th Workshop on Many-Task Computing on Grids and Supercomputers (MTAGS)*, Co-located with *ACM/IEEE Supercomputing Conference (SC'11)*, 2011. (Acceptance rate: 6/14=42.8%) (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2132887)
122. H. Song, Y. Yin, **Y. Chen** and X.-H. Sun. A Cost-Based Application-Specific Data Layout Scheme for Parallel File Systems. In the *Proc. of the 20th International Symposium on High Performance Distributed Computing (HPDC'11)*, 2011. (Acceptance rate: 22/170=12.9%) (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=1996138)
123. **Y. Chen**, X.-H. Sun, R. Thakur, P. C. Roth and W. Gropp. LACIO: A New Layout-Aware Collective I/O Strategy for Parallel I/O Systems. In *Proc. of the 25th IEEE International Parallel & Distributed Processing Symposium (IPDPS'11)*, 2011. (Acceptance rate: 112/571=19.6%) (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2059509)
124. H. Song, X.-H. Sun, and **Y. Chen**. A Hybrid Shared-nothing/Shared-data Storage Scheme for Large-scale Data Processing. In the *Proc. of the 9th IEEE International Symposium on Parallel and Distributed Processing with Applications (ISPA'11)*, 2011. (**Best Paper Award**) (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2006200)
125. K. Zhang, Z. Wang, **Y. Chen**, H. Zhu and X.-H. Sun. PAC-PLRU: A Cache Replacement Policy to Salvage Discarded Predictions from Hardware Prefetchers. In the *Proc. of the 11th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid'11)*, 2011. (Acceptance rate: 55/189=29.1%) (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2007417)
126. **Y. Chen** and P. C. Roth. Collective Prefetching for Parallel I/O Systems. In *Proc. of the 5th Petascale Data Storage Workshop (PDSW'10)*, in conjunction with *ACM/IEEE Supercomputing (SC'10)*, 2010. (ACM DL/IEEE Xplore DL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5668089)
127. H. Jin, X.-H. Sun, **Y. Chen** and T. Ke. REMEM: REmote MEMory as Checkpointing Storage. In Proc. of the IEEE International Conference on Cloud Computing Technology and Science (Cloudcom'10). 2010. (Acceptance rate: <25%). (ACM DL/IEEE Xplore DL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5708466)
128. **Y. Chen**, X.-H. Sun, R. Thakur, H. Song and H. Jin. Improving Parallel I/O Performance with Data Layout

Awareness. In *Proc. of the IEEE International Conference on Cluster Computing 2010 (Cluster'10),* 2010. (Acceptance rate: 33/107=30.8%) (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=1901094)

129. H. Jin, **Y. Chen** and X.-H. Sun. Optimizing HPC Fault-Tolerant Environment: An Analytical Approach. In the *Proc. of the 39th Intl. Conference on Parallel Processing (ICPP'10)*, 2010. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=1905243)
130. **Y. Chen**, H. Zhu, H. Jin and X.-H. Sun. Improving the Effectiveness of Context-based Prefetching with Multi-order Analysis. In the *Proceedings of the 3rd International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2)*, 2010. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=1905156)
131. **Y. Chen**, H. Zhu and X.-H. Sun. An Adaptive Data Prefetcher for High-Performance Processors. In the *Proc. of the 10th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid'10)*, 2010. (Acceptance rate: 51/219=23.3%) (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=1845110)
132. H. Zhu, **Y. Chen** and X.-H. Sun. Timing Local Streams: Improving Timeliness in Data Prefetching. In the *Proc. of the 24th ACM Intl Conference on Supercomputing (ICS'10)*, 2010. (Acceptance rate: 32/180=17.8%) (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=1810110)
133. **Y. Chen**, H. Song, R. Thakur and X.-H. Sun. A Layout-aware Optimization Strategy for Collective I/O (short paper). In the *Proc. of High Performance Distributed Computing (HPDC-2010)*, 2010. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=1851530)
134. X.-H. Sun, **Y. Chen**, Y. Yin. Data Layout Optimization for Petascale File Systems. In *Proc. of 4th Petascale Data Storage Workshop, in conjunction with ACM/IEEE Supercomputing (SC'09)*, 2009. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=1713077)
135. Z. Fang, X.-H. Sun, **Y. Chen** and S. Byna. Core-Aware Memory Access Scheduling Schemes. In *Proc. of IEEE International Parallel & Distributed Processing Symposium (IPDPS'09)*, 2009. (Acceptance rate: 100/440=22.7%) (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=1587713)
136. B. Xie, **Y. Chen**, X.-H. Sun and H. Jin. Performance under Failure of Multi-tier Web Services. In *Proc. of the Workshop on Internet-based Virtual Computing Environment*, 2009. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=1728690)
137. X.-H. Sun, C. Du, H. Zou, **Y. Chen**, and P. Shukla. V-MCS: A Configuration System for Virtual Machines. In *Proc. of Workshop on e-Research Infrastructure, Services and Applications*, 2009. (ACM DL/IEEE Xplore DL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5289131)
138. **Y. Chen**, S. Byna, X.-H. Sun, R. Thakur and W. Gropp. Hiding I/O Latency with Pre-execution Prefetching for Parallel Applications. In *Proc. of the ACM/IEEE Supercomputing Conference (SC'08)*, Nov. 2008. (**Best Paper finalist, Best Student Paper finalist**) (Acceptance rate: 59/277=21.3%) (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=1413411)
139. S. Byna, **Y. Chen**, X.-H. Sun, R. Thakur and W. Gropp. Parallel I/O Prefetching Using MPI File Caching and I/O Signatures. In *Proc. of the ACM/IEEE Supercomputing Conference (SC'08)*, 2008. (Acceptance rate: 59/277=21.3%) (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=1413415)
140. X.-H. Sun, **Y. Chen** and S. Byna. Scalable Computing in Multicore Era. In P*roc. of the Intl. Symposium on Parallel Algorithms, Architectures and Programming (PAAP'08)*, 2008. (http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.465.4291)
141. **Y. Chen**, S. Byna, X.-H. Sun, R. Thakur and W. Gropp. Exploring Parallel I/O Concurrency with Speculative Prefetching. In *Proc. of 37th International Conference on Parallel Processing (ICPP'08)*, 2008.

(Acceptance rate: 81/263=30.8%) (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=1442445)

142. S. Byna, **Y. Chen** and X.-H. Sun. A Taxonomy of Data Prefetching Mechanisms. In *Proc. of International Symposium on Parallel Architectures, Algorithms and Networks (ISPAN'08)*, 2008. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=1397185)
143. **Y. Chen**, S. Byna and X.-H. Sun. Data Access History Cache and Associated Data Prefetching Mechanisms. In *Proc. of the ACM/IEEE Supercomputing Conference (SC'07)*, 2007. (Acceptance rate: 54/268=20.1%) (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=1362651)
144. X.-H. Sun, S. Byna and **Y. Chen**. Improving Data Access Performance with Server Push Architecture. In *Proc. of the NSF Next Generation Software Program Workshop (in conjunction with IPDPS '07)*, 2007. (ACM DL/IEEE Xplore DL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4228239)
145. **Y. Chen** and X.-H. Sun. STAS: A Scalability Testing and Analysis System. In *Proc. of IEEE International Conference on Cluster Computing (Cluster'06)*, 2006. (Acceptance rate: 42/127=33.1%) (ACM DL/IEEE Xplore DL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4100388)
146. M. Wu, X.-H. Sun and **Y. Chen**. QoS Oriented Resource Reservation in Shared Environments. In *Proc. of 6th IEEE International Symposium on Cluster Computing and the Grid (CCGrid'06)*, 2006. (Acceptance rate: 61/237=25.7%) (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=1135074)
147. X.-H. Sun, **Y. Chen** and M. Wu. Scalability of Heterogeneous Computing. In *Proc. of the 34th International Conference on Parallel Processing (ICPP'05)*, 2005. (Acceptance rate: 69/241=28.6%) (ACM DL/IEEE Xplore DL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1488654)
148. E.-C. Jung, K. Sato, **Y. Chen**, X. He, T. MacTavish and D. Cracchiolo. DIF Knowledge Management System: Bridging Viewpoints for Interactive System Design. In *Proc. of 11th Intl. Conf. on Human-Computer Interaction (HCI'05)*, 2005. (http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.469.9729)
149. **Y. Chen**, G. Chen, Y. Xu and J. Shan. Implementation and Evaluation of MPI+OpenMP Programming Model on Dawning3000. In *Proc. of the 21st IASTED International Multi-Conference on Applied Informatics (AI'03)*, 2003.

REFEREED POSTERS/POSTER PAPERS: (* GRADUATE ADVISEE; ** UNDERGRADUATE ADVISEE)

1. W. Zhang*, S. Byna, **Y. Chen**. Efficient Metadata Search for Scientific Data. *In the 32nd International Conference for High Performance Computing, Networking, Storage and Analysis (SC'20)*, 2020.
2. W. Zhang*, S. Byna, **Y. Chen**. Activeness-based Data Retention Recommender for HPC Facilities. *In the 32nd International Conference for High Performance Computing, Networking, Storage and Analysis (SC'20)*, 2020.
3. C. Niu*, W. Zhang*, S. Byna, **Y. Chen**. Semantic Search for Self-describing Scientific Data Formats. *In the 32nd International Conference for High Performance Computing, Networking, Storage and Analysis (SC'20)*, 2020.
4. X. Wang*, J. Leidel, B. Williams*, A. Ehret, M. Mark, S. Bandara, M. Kinsy, **Y. Chen**. xBGAS: An Address Space Extension for Scalable High Performance Computing. *In the 32nd International Conference for High Performance Computing, Networking, Storage and Analysis (SC'20)*, 2020.
5. G. Ali*, S. Bhalachandra, N. Wright, A. Sill, Y. Chen. Evaluation of power counters and controls on general-purpose GPUs. *In the 32nd International Conference for High Performance Computing, Networking, Storage and Analysis (SC'20)*, 2020.
6. M. Ahmadian*, E. Rees, Y. Zhuang and Y. Chen. Reducing Faulty Jobs by Job Submission Verifier in Grid Engine. *In Proceedings of The 2019 ACM Practice and Experience in Advanced Research Computing (PEARC'19)*, 2019. **Best Student Poster Award**

7. X. Wang*, J. Li*, A. Tumeo, J. D. Leidel, **Y. Chen**. Memory Hotspot Optimizations for 3D-Stacked Memory, *The 28th International Conference on Parallel Architectures and Compilation Techniques (PACT'19)*, 2019.
8. W. Zhang*, H. Tang, S. Byna and **Y. Chen**. Distributed Adaptive Radix Tree for Efficient Metadata Search on HPC Systems. *The 30th International Conference for High Performance Computing, Networking, Storage and Analysis (SC'18)*, 2018.
9. W. Zhang*, **Y. Chen**, and D. Dai. AKIN: A Streaming Graph Partitioning Algorithm for Distributed Graph Storage Systems. *The 18th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid'18)*, May 2nd, 2018, Washington D.C.
10. E. Hojati*, **Y. Chen** and A. Sill. Benchmarking Automated Hardware Management Technologies for Modern Scalable Data Centers and Cloud Environments. *The 10th IEEE/ACM International Conference on Utility and Cloud Computing (UCC'17)*, **Best Student Poster Award.**
11. F. Fard, D. Donofrio, J. Shalf, J. Leidel, X. Wang* and **Y. Chen**. OpenSoc System Architect: An Open Source Supercomputing Platform. *In the Design Automation Conference (DAC'17)*, Austin, TX, June 18, 2017.
12. D. Dai, W. Zhang* and **Y. Chen**. IOGP: An Incremental Online Graph Partitioning for Large-Scale Distributed Graph Databases, *PPoPP 2017*.
13. W. Xie* and **Y. Chen**. Building Faster, Elastic, and Durable Large-scale Data Store with Consistent Hashing. *In the NSF CSR PI meeting (co-located with IPDPS)*, Orlando, FL, June 2, 2017.
14. X. Wang*, J. Leidel* and **Y. Chen**. Concurrent Dynamic Memory Coalescing on GoblinCore-64 Architecture. *In the SC'16*, Salt Lake City, UT, Nov. 15, 2016.
15. J. Zhou, W. Xie* and **Y. Chen**. Attributed Consistent Hashing for Heterogeneous Storage System. *In the SC'16*, Salt Lake City, UT, Nov. 15, 2016.
16. N. Tavakoli*, D. Dai, J. Jenkins, P. Carns, R. Ross and **Y. Chen**. A Software-Defined Approach for QoS Control in High-Performance Computing Storage Systems. *In the SC'16*, Salt Lake City, UT, Nov. 15, 2016.
17. D. Dai, R. Ross, D. Khaldi, Y. Yan, D. Matthieu, N. Tavakoli and **Y. Chen**. Exploiting Locality in Scientific Workflow System: A Cross-Layer Solution. *In the SC'16*, Salt Lake City, UT, Nov. 15, 2016.
18. F. Fatollahi-Fard, D. Donofrio, J. Shalf, J. Leidel, X. Wang and **Y. Chen**. Open Source HPC: An Open Source Supercomputing Platform. *In the Workshop on Modeling & Simulation of Systems and Applications (ModSim 2016)*, Seattle, Washington, Aug. 12, 2016.
19. M. Li, X. Shi, W. Liu, H. Jin, and **Y. Chen**. SSDUP: An Efficient SSD Write Buffer Using Pipeline (Poster). In *The 2016 IEEE International Conference on Cluster Computing (Cluster'16)*, 2016.
20. J. Leidel and **Y. Chen**. GoblinCore-64: A Scalable, Open Architecture for Data Intensive High Performance Compouting. *In the International Parallel and Distributed Processing Symposium (IPDPS) PhD Forum*, Chicago IL, May 24, 2016.
21. M. Ahmadian*, Y. Zhuang, W.L. Hase, **Y. Chen**. Information-Preserving Data Reduction for Scientific Applications through Efficient Data Modeling. In the *Conference on Data Analysis (CoDA 2016)* (poster session), Santa Fe, New Mexico, March 2nd, 2016.
22. W. Xie* and **Y. Chen**, A Cache Management Scheme for Hiding Garbage Collection Latency in Flash-based Solid State Drives. In the *Proc. of the 2015 IEEE Cluster Conference (Cluster'15)*, Chicago, IL, September 8-11, 2015. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2859632)
23. E. Valenzuela**. Multi-Level Hashing for Data Deduplication in HPC Storage Systems. ACM SRC (Student Research Competition) Poster, in *The 2014 ACM/IEEE Supercomputing Conference (SC'14)*, (mentors: Y. Lu, S. Urban, Y. Chen), New Orleans, LA, November 16th-21st, 2014. **Placed 4th for ACM SRC**

(http://sc14.supercomputing.org/sites/all/themes/sc14/files/archive/src_poster/index.html)

24. O. Rodriguez**. Enhancing FlashSim Simulations for High Performance Computing Storage Systems. ACM SRC (Student Research Competition) Poster, in *The 2014 ACM/IEEE Supercomputing Conference (SC'14)*, (mentors: W. Xie, Y. Chen), New Orleans, LA, November 16th-21st, 2014. (http://sc14.supercomputing.org/sites/all/themes/sc14/files/archive/src_poster/index.html)
25. D. Dai, **Y. Chen**, D. Kimpe, and Robert Ross. Provenance-Based Prediction Scheme for Object Storage System in HPC. In the *Proc. of the 14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid'14)*, 2014. (ACM DL/IEEE Xplore DL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6846497)
26. W. Xie*, **Y. Chen**. An Adaptive Separation-Aware FTL for Improving the Efficiency of Garbage Collection in SSDs. In the *Proc. of the 14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid'14)*, 2014. (ACM DL/IEEE Xplore DL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6846498)
27. G. Thorsness**. Fusion Active Storage for Write-intensive Big Data Applications. *SC'13 ACM SRC (Student Research Competition) Poster* (faculty mentor: Yong Chen), Denver, CO, November 18th-22nd, 2013. (http://sc13.supercomputing.org/content/poster-archive.html)
28. J.J. Chen*, **Y. Chen**. Unified and Efficient HEC Storage System with a Working-Set based Reorganization Scheme. In the *IEEE International Conference on Cluster Computing 2013 (Cluster'13)*, Indianapolis, IN, September 23-27, 2013. (ACM DL/IEEE Xplore DL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6702620)
29. **Y. Chen**, Y. Zhuang, and N. Lopez-Benitez. Fall-11: Early Adoption of NSF/TCPP PDC Curriculum at Texas Tech University and Beyond. In the *3rd NSF/TCPP Workshop on Parallel and Distributed Computing Education (EduPar-13)*, in conjunction with the *27th IEEE International Parallel & Distributed Processing Symposium (IPDPS'13)*, Boston, May 20th, 2013. (with NSF/GSU travel support) http://grid.cs.gsu.edu/~tcpp/curriculum/?q=EduPar-13_Proceedings)
30. Y. Lu*, **Y. Chen**, R. Thakur, Y. Zhuang. Memory-Conscious Collective IO for Extreme-Scale HPC Systems. In the *ACM/IEEE Supercomputing Conference (SC'12)*, Salt Lake City, Utah, November 13th, 2012. (**Best Poster Finalist**) (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2477089)
31. J.L. Liu*, **Y. Chen**. Fast Data Analysis with Integrated Statistical Metadata in Scientific Datasets. In the *41st International Conference on Parallel Processing (ICPP'12),* Pittsburgh, PA, September 10-13, 2012. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2410523)
32. **Y. Chen**, Y. Zhuang, and N. Lopez-Benitez. Early Adoption of NSF/TCPP PDC Curriculum at Texas Tech University. In the 2nd NSF/TCPP Workshop on Parallel and Distributed Computing Education (EduPar-12), in conjunction with the 26th IEEE International Parallel & Distributed Processing Symposium (IPDPS'12), Shanghai, May 21, 2012. (with NSF/GSU travel support) (http://grid.cs.gsu.edu/~tcpp/curriculum/?q=advanced-technical-program)
33. A. Tambi* and **Y. Chen**. A Comprehensive Benchmark Suite for Emerging Solid State Drives. In the *ACM/IEEE Supercomputing Conference (SC'11)*, Seattle, WA, 2011. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2148641)
34. Y. R. Guvvala, **Y. Chen**, and Y. Zhuang. Revisiting RAID for SSD based HPC Systems. In the *ACM/IEEE Supercomputing Conference (SC'11)*, Seattle, WA, 2011. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2148643)
35. H. Song, **Y. Chen**, X.-H. Sun. A Hybrid Shared-nothing/Shared-data Storage Architecture for Large Scale

Databases. In the *11th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid'11)*, 2011. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=2007387)

36. H. Song, X.-H. Sun, H. Jin, and **Y. Chen**. Trace-based Adaptive Data Layout Optimization for Parallel File systems. In the *5th Petascale Data Storage Workshop*, in conjunction with *ACM/IEEE Supercomputing Conference (SC'10)*, 2010. (http://www.pdsw.org/pdsw10/resources/posters/layout-HSong.pdf)
37. **Y. Chen**, H. Song, R. Thakur and X.-H. Sun. A Layout-aware Optimization Strategy for Collective I/O. In the ACM *High Performance Distributed Computing (HPDC-2010)*, Chicago, IL, 2010. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=1851530)
38. H. Jin, X.-H. Sun, B. Xie and **Y. Chen**. An Implementation and Evaluation of Memory-based Checkpointing. In the *ACM/IEEE Supercomputing Conference (SC'09)*, Portland, OR, 2009.
39. **Y. Chen**, S. Byna, X.-H. Sun, R. Thakur and W. Gropp. Server-Push Architecture for Improving I/O Access Performance. *NSF Collaborative Expedition Workshop*, Arlington, VA, 2008.
40. S. Byna, **Y. Chen**, W. Gropp, X.-H. Sun and R. Thakur. The Server-Push I/O Architecture for High-End Computing. In the *ACM/IEEE Supercomputing Conference (SC'07)*, Reno, NV, 2007.

NON-REFEREED POSTERS: (* GRADUATE ADVISEE; ** UNDERGRADUATE ADVISEE) *(NOT UPDATED SINCE 2018)*

41. F. Perez**, W. Xie* and **Y. Chen**. Improving Erasure-coded Distributed Storage System Scalability Using Consistent Hashing. *In the TTU CALUE Undergraduate Research Conference*, Lubbock, TX, March 28th, 2018.
42. J. Zhou, W. Xie* and **Y. Chen**. Consistent Hashing-based Algorithm for Heterogeneous Storage Systems. In *the CAC Semi-annual IAB Meeting*, Dallas, TX, Oct. 5th, 2016.
43. C. Gray, **Y. Chen** and R. Vadapalli. Applying Cloud Computing in Healthcare Informatics. *In the Cloud and Autonomic Computing Center (CAC) 2016 semi-annual meeting*, Lubbock, TX, April 28th, 2016.
44. W. Xie*, J. Zhou, M. Reyes**, J. Nobel and **Y. Chen**. Unistore Project Updates. In the *Cloud and Autonomic Computing Center (CAC) 2016 semi-annual meeting*, Lubbock, TX, April 28th, 2016.
45. E. Hojati* and **Y. Chen**. Benchmarking Fairness and Performance Isolation in Cloud Storage. In the *Cloud and Autonomic Computing Center (CAC) 2016 semi-annual meeting*, Lubbock, TX, April 28th, 2016.
46. W. Xie*, J. Zhou, M. Reyes (K. Echo)**, J. Nobel and **Y. Chen**. Tiered-CRUSH: A High-Performance Data Placement for Multi-Tiered Storage System. In the *Cloud and Autonomic Computing Center (CAC) 2015 semi-annual meeting*, Dallas, TX, Oct. 1st 2015.
47. J. Zhou, W. Xie*, J. Nobel, M. Reyes (K. Echo)**, and **Y. Chen**. SUORA: A Scalable and Uniform Data Distribution Algorithm for Heterogeneous Storage Systems. In the *Cloud and Autonomic Computing Center (CAC) 2015 semi-annual meeting*, Dallas, TX, Oct. 1st 2015.
48. W. Xie*, J. Zhou, and **Y. Chen**. Unistore: A Unified Storage Architecture for Cloud Computing. In the *Cloud and Autonomic Computing Center (CAC) 2015 semi-annual meeting*, Lubbock, TX, April 9th, 2015.

DISSERTATION/THESIS:

1. **Y. Chen**. A Hybrid Data Prefetching Architecture for Data-Access Efficiency. Ph.D. Dissertation, Illinois Institute of Technology, July 2009. (ACM DL/IEEE Xplore DL: http://dl.acm.org/citation.cfm?id=1835286)
2. **Y. Chen**. Research on Parallel Programming Model and Parallel Debugging. Masters Thesis, University of Science and Technology of China, July 2003.
3. **Y. Chen**. PVM Parallel Programming and Runtime System Research. B.E. Thesis, University of Science and Technology of China, July 2000.

PATENTS

1. Xian-He Sun, **Yong Chen**, Huaiyu Zhu. Timing-Aware Data Prefetching for Microprocessors. U.S. Patent 8856452, 2014. (https://www.google.com/patents/US8856452)

INVITED TALKS/DEMOS

1. Empowering Data-driven Discovery with Provenance Collection, Management, and Analysis. Invited talk at *HPC China Performance Modeling Forum, 2020*. (Delivered online via a recorded presentation)
2. Empowering Data-driven Discovery with Provenance Collection, Management, and Analysis. Invited talk at the *Twelfth International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2)*, held in conjunction with *The 48th International Conference on Parallel Processing (ICPP'19)*, August 5th, 2019, Kyoto, Japan.
3. Empowering Data-driven Discovery with a Lightweight Provenance Service for High Performance Computing, Invited Talk at *NITRD's Middleware and Grid Interagency Coordination Group's (MAGIC) meeting*, April 3rd, 2019.
4. Extended Base Global Address Space for Data Center Scale Addressing, Presented at *Cloud and Autonomic Computing Center Semi-Annual IAB Meeting*, April 4th, 2019, Tucson, AZ.
5. High Performance Computing, Computational Science, and Data Science, *TTU SPE Guest Lecture Series*, April 4th, 2018.
6. High Performance Computing, Computational Science, and Data Science, Invited Talk at HUST (July 13th), Wuhan University (July 14th), at Shanghai Jiaotong University (July 24th), 2018
7. High Performance Computing Revisited for Big Data Applications, Invited Talk at HPCCT and BDAI Conferences, June 23rd, 2018.
8. Data Science and Data-driven Scientific Discovery, Invited Talk at Nankai University (July 4th) and USTC (July 18th), 2017.
9. Building Faster, Elastic, and Durable Large-scale Data Store with Consistent Hashing, Presented at the 2017 NSF CSR PI meeting, June 2nd, 2017, Orlando, FL.
10. Fast Data Analysis for Scientific Big Data Applications, *Computer Science Department Seminar, Texas Tech University*, Nov. 29th, 2016.
11. Decoupled System Architecture and Computation Paradigm for Scientific Big Data Applications, Invited Talk at Texas A&M University Conference on Advances in Big Data Modeling, Computation and Analytics, Sept. 23rd, 2016.
12. GoblinCore-64: An Emerging, Open Architecture for Data Intensive High Performance Computing, Demo at the *Emerging Technologies track of ACM/IEEE Supercomputing Conference (SC'15)* (with Mr. John Leidel and Mr. Xi Wang), November 17th, 2015, Austin, Texas.
13. Unistore: A Unified Storage Architecture for Cloud Computing. Invited talk at the *2015 Storage Developer Conference*, September 21st, 2015, Santa Clara, California.
14. High Performance Computing and Scientific Big Data Applications. Invited talk at the *Institute of Computing Technology, Chinese Academy of Sciences*, September 2nd, 2015, Beijing, China; at *Tianjin University and Nankai University*, August 31st, 2015, Tianjin, China.
15. Introduction to the Unistore Project, in Expert Panel: Cloud Storage Initiatives – An SDC Preview, August 4th, 2015, https://www.brighttalk.com/webcast/663/166033, *online webinar to 100+ attendees*.

16. Decoupled Execution Paradigm for Data-Intensive High Performance Computing. Invited talk *at Huazhong University of Science and Technology, Tsinghua University, and University of Science and Technology of China*, January 2015 and December 2014, Wuhan, Beijing, and HeFei, China.
17. Design and Development of a Data-Intensive Scalable Computing Instrument (DISCI) for Big Data Applications. Invited talk at the *Texas Tech University 2014 Symposium on Big Data*, April 25th, 2014, Lubbock, Texas.
18. Fast Data Analysis for Scientific Big Data Applications. Invited talk in the *Department of Computer & Information Science, IUPUI*, November 22nd, 2013, Indianapolis, Indiana.
19. Fast Data Analysis with Integrated Statistical Metadata in Scientific Datasets. Invited talk in *the Computational Research Division, Lawrence Berkeley National Laboratory*, October 7th, 2013.
20. Improving Data-Access Efficiency: Towards Scalable High-End Computing. Invited talk in the *Department of Electrical Engineering and Computer Science, Northwestern University*, Apr. 19th, 2010, Evanston, Illinois.
21. GMC: The Global-aware and Multi-order Context-based Prefetcher. Invited talk at the *ACM/IEEE Supercomputing Conference (SC'09)* for the ACM/IEEE High-Performance Computing Fellowship, Nov. 19th, 2009, Portland, Oregon.
22. A Hybrid Data Prefetching Architecture for Data-Access Efficiency, *Computer Science and Mathematics Division, Oak Ridge National Laboratory*, *Department of Energy*, Feb. 20th, 2009, Oak Ridge, Tennessee.
23. Hiding I/O Latency with Pre-execution Prefetching for Parallel Applications, *Mathematics and Computer Science Division, Argonne National Laboratory*, *Department of Energy*, Nov. 10th, 2008, Argonne, Illinois.

CONFERENCE/WORKSHOP PRESENTATIONS (CORRESPOND WITH PUBLICATIONS)

24. Semantic Search for Self-Describing Scientific Data Formats. Poster Presentation at The 32nd ACM/IEEE International Conference for High Performance Computing, Networking, Storage, and Analysis (SC'20), Online, 2020.
25. Evaluation of Power Controls and Counters on General-Purpose Graphics Processing Units. Poster Presentation at The 32nd ACM/IEEE International Conference for High Performance Computing, Networking, Storage, and Analysis (SC'20), Online, 2020.
26. xBGAS: An Address Space Extension for Scalable High-Performance Computing. Poster Presentation at The 32nd ACM/IEEE International Conference for High Performance Computing, Networking, Storage, and Analysis (SC'20), Online, 2020.
27. MAC: Memory Access Coalescer for 3D-Stacked Memory. Presented at *The 48th International Conference on Parallel Processing (ICPP'19)*, Aug 6th, 2019, Kyoto, Japan.
28. Contention-Aware Resource Scheduling for Burst Buffer Systems. Presented at *The 11th International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2)*, held in conjunction with *The 47th International Conference on Parallel Processing (ICPP'18)*, August 13th, 2018, Eugene, OR.
29. AKIN: A Streaming Graph Partitioning Algorithm for Distributed Graph Storage Systems. Poster Presentation at *The 18th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid'18)*, May 2nd, 2018, Washington D.C.
30. Elastic Consistent Hashing for Distributed Storage Systems. Presented at The *31st IEEE International Parallel and Distributed Processing Symposium (IPDPS'17)*, June 1st, 2017, Orlando, FL.

31. Rethinking High Performance Computing System Architecture for Scientific Big Data Applications. Presented at The *14th IEEE International Symposium on Parallel and Distributed Processing with Applications (ISPA'16)*, August 24th, 2016, Tianjin, China.
32. Hierarchical Consistent Hashing for Heterogeneous Object-based Storage. Presented at *The 14th IEEE International Symposium on Parallel and Distributed Processing with Applications (ISPA'16)*, August 25th, Tianjin, China.
33. Two-Mode Data Distribution Scheme for Heterogeneous Storage in Data Centers. Presented at *The 3rd IEEE BigData Conference*, Oct. 30th, 2015, Santa Clara, CA.
34. Collective Computing for Scientific Big Data. Presented at *The Eighth International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2)*, Sep. 1st, 2015, Beijing, China.
35. A Virtual Shared Metadata Storage for HDFS. Presented at *The 10th IEEE International Conference on Networking, Architecture, and Storage (NAS'15)*, Aug. 6th 2015, Boston, MA.
36. Decoupled I/O for Data-Intensive High Performance Computing. Presented at *The 7th Intl. Workshop on Parallel Programming Models and Systems Software for High-End Computing, held in conjunction with The 43rd Intl. Conference on Parallel Processing (ICPP)*, Sep 12th 2014, Minneapolis, MN.
37. Provenance-Based Prediction Scheme for Object Storage System in HPC. Poster Presentation. *The 14th IEEE/ACM Intl. Symp. on Cluster, Cloud and Grid Computing (CCGrid'14)*, May 27th, 2014, Chicago, IL.
38. An Adaptive Separation-Aware FTL for Improving the Efficiency of Garbage Collection in SSDs. Poster Presentation. *The 14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid'14)*, May 27th, 2014, Chicago, IL.
39. Using Pattern-Models to Guide SSD Deployment for Big Data in HPC systems. Presented at *The IEEE Big Data 2013 Conference*, October 8th, 2013, Santa Clara, CA.
40. Multilevel Active Storage for Big Data Applications in High Performance Computing. Presented at *The IEEE Big Data 2013 Conference*, October 7th, 2013, Santa Clara, CA.
41. Unified and Efficient HEC Storage System with a Working-Set based Reorganization Scheme, Presented at *The IEEE Intl. Conference on Cluster Computing 2013 (Cluster'13)*, on Sept. 24th, 2013, Indianapolis, IN.
42. Memory-Conscious Collective I/O for Extreme Scale HPC Systems. *The Intl. Workshop on Runtime and Operating Systems for Supercomputers (ROSS)*, in conjunction with ICS'13, June 10th, 2013, Eugene, OR.
43. Data Deduplication in a Hybrid Architecture for Improving Write Performance. *The Intl. Workshop on Runtime and Operating Systems for Supercomputers (ROSS)*, in conjunction with ICS'13, June 10th, 2013, Eugene, OR.
44. Data-Intensive Scalable Computing Systems. Opening Address. The International Workshop on Data-Intensive Scalable Computing Systems (DISCS), November 16th, 2012, Salt Lake City, UT.
45. Memory-Conscious Collective IO for Extreme-Scale HPC Systems. Poster Presentation. *ACM/IEEE Supercomputing Conference (SC'12)*, November 13th, 2012, Salt Lake City, UT.
46. Improving Data Analysis Performance for High-Performance Computing with Integrating Statistical Metadata in Scientific Datasets. The *Second Annual Workshop on High-Performance Computing meets Databases (HPCDB)*, in conjunction with the *SC'12*, November 11th, 2012, Salt Lake City, UT.
47. Fast Data Analysis with Integrated Statistical Metadata in Scientific Datasets. Poster Presentation. The *41st International Conference on Parallel Processing (ICPP'12)*, September 11th, 2012, Pittsburgh, PA.
48. Panel Moderator: Battle of the Accelerator Stars, The *5th International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2)*, co-located with the *41st International Conference on Parallel Processing (ICPP'12)*, September 10th, 2012, Pittsburgh, PA.

49. Towards Scalable I/O Architecture for Exascale Systems. The *4th Workshop on Many-Task Computing on Grids and Supercomputers (MTAGS)*, Co-located with *ACM/IEEE Supercomputing Conference (SC'11)*, November 14th, 2011, Seattle, WA.
50. A Comprehensive Benchmark Suite for Emerging Solid State Drives. Poster Presentation. *ACM/IEEE Supercomputing Conference (SC'11)*, November 15th, 2011, Seattle, WA.
51. Revisiting RAID for SSD based HPC Systems. Poster Presentation. *ACM/IEEE Supercomputing Conference (SC'11)*, November 15th, 2011, Seattle, WA.
52. LACIO: A New Layout-Aware Collective I/O Strategy for Parallel I/O Systems. The *25th IEEE International Parallel & Distributed Processing Symposium (IPDPS'11)*, May 18th, 2011, Anchorage, AK.
53. Collective Prefetching for Parallel I/O Systems. The *5th Petascale Data Storage Workshop*, in conjunction with *ACM/IEEE Supercomputing (SC'10)*, Nov. 15th, 2010, New Orleans, LA.
54. Trace-based Adaptive Data Layout Optimization for Parallel File systems. Poster Presentation. In the *5th Petascale Data Storage Workshop*, in conjunction with *ACM/IEEE Supercomputing Conference (SC'10)*, Nov. 15th, 2010, New Orleans, LA.
55. Optimizing Collective I/O with Data Layout Awareness for Parallel Applications, poster presentation, *NSF and DOE HECFSIO-2010 PI meeting*, Aug. 2nd, 2010, Arlington, VA.
56. Data Layout Optimization for Petascale File Systems. The *4th Petascale Data Storage Workshop (in conjunction with ACM/IEEE Supercomputing'09)*, Nov. 15th, 2009, Portland, Oregon.
57. An Implementation and Evaluation of Memory-based Checkpointing. Poster Presentation. *ACM/IEEE Supercomputing Conference (SC'09)*, Nov. 17th, 2009, Portland, Oregon.
58. The Server Push Architecture for High-End Computing, poster presentation, *NSF HECURA-FSIO2009 PI meeting*, August 9-12, 2009, Arlington, VA.
59. A Hybrid Data Prefetching Architecture for Data-Access Efficiency, *ACM/IEEE Supercomputing Conference (SC'08) Doctoral Showcase*, Nov. 20th, 2008, Austin, Texas.
60. Hiding I/O Latency with Pre-execution Prefetching for Parallel Applications, *ACM/IEEE Supercomputing Conference (SC'08)*, Nov. 20th, 2008, Austin, Texas.
61. Exploring Parallel I/O Concurrency with Speculative Prefetching. The *37th International Conference on Parallel Processing (ICPP'08)*, Sept. 10th, 2008, Portland, Oregon.
62. Server-Push Architecture for Improving I/O Access Performance. *NSF Collaborative Expedition Workshop*, Jun. 10th, 2008, National Science Foundation, Arlington, Virginia.
63. Data Access History Cache and Associated Data Prefetching Mechanisms. *ACM/IEEE Supercomputing Conference (SC'07)*, Nov. 13th, 2007, Reno, Nevada.
64. The Server-Push I/O Architecture for High-End Computing. Poster presentation at the *ACM/IEEE Supercomputing Conference (SC'07)*, Nov. 13th, 2007, Reno, Nevada.
65. Implementation and Evaluation of MPI+OpenMP Programming Model on Dawning3000. The 21st *IASTED International Multi-Conference on Applied Informatics (AI'03)*, Feb. 12th, 2003, Innsbruck, Austria.

PANELISTS FOR FEDERAL AGENCIES

- **NSF CISE Panelist,** March 2021
- **NSF CISE Panelist,** July & October 2020
- **DOE Minority Serving Institution Partnership Program,** May 2020
- **NSF CISE Panelist,** June 2019
- **NSF CISE Panelist,** February 2019

- **NSF CISE Panelist,** April 2018
- **NSF ACI CMU/PSC "Bridges"** site visit panelist**,** May 2017
- **NSF ACI DIBBS NCSA/UIUC, Purdue & CMU/PSC** reverse site visit panelist**,** December 2016
- **NASA Panelist,** May 2015
- **NSF CISE Panelist,** April 2015
- **NSF CISE Panelist,** March 2015
- **NSF CISE Panelist,** March 2014
- **NSF CISE Panelist,** April 2011

EXTERNAL PROPOSAL REVIEWER/PANELISTS FOR OTHER AGENCIES

- **Nazarbayev University Research Review via ORAU,** 2019
- **Hongkong RGC (Research Grants Council),** 2021, 2019, 2018, 2017
- **Portland OSF,** 2018

INSTITUTIONAL SERVICES (PARTIAL LIST)

- **Department Services**
  - Search Committee, September 2019 – Present
  - Space & Equipment Committee, September 2020 – Present (Chair)
  - Departmental Relationship Committee, 2015 – 2016 (Chair), September 2019 – Present
  - Graduate Program Committee, 2012 – 2015, 2016 – 2017 (Chair), 2018 – 2020
  - Graduate Program Coordinator, 2016 – 2017
  - Graduate Admission Committee, 2014-2015, 2018 – 2020
  - Graduate Support Committee, 2012 – 2013, 2017 – 2018
  - Strategic Planning Committee, 2013 – 2014, 2016 – 2017, 2017 – 2018 (Chair)
  - Undergraduate Program Committee, 2015 – 2016
  - Faculty Search Committee and Big Data Cluster Hire Committee, 2014 – 2015, 2017 – 2018
  - Faculty Midterm/Comprehensive Review Committee, 2017 – present
  - Faculty coordinator, Computer Science Departmental U-Reason Seminar, Fall 2013
  - Faculty Advisor of the Student Cluster Competition team (http://discl.cs.ttu.edu/scc/), trained and led a team of undergraduate students and participated in the *ACM/IEEE Supercomputing Conference 2012 (SC'12)* as one of finalist teams together with Dr. James Abbott of High Performance Computing Center at TTU and with vendor sponsorship from Dell Inc., Received **Special Recognition for Exemplary Spirit**, 2012
  - Course coordinator, mentored CS4352 Operating Systems course taught by GPTI, 2014-2016
- **College Services**
  - Institutional Effectiveness Committee, 2016 – 2018
  - Search Committee for Director of Engineering Computing Services, 2014
  - Judge and a keynote speaker for the first Appathon, a mobile app development competition hosted by Whitacre Jr. College of Engineering, 2013
- **University Services**
  - Clark Scholar Program faculty mentor (an intensive seven-week summer research program for highly qualified high school juniors and seniors), 2016 – present

- Office of Vice President Research Pre-proposal Review Committee, 2013 – present
- Mentor Tech program, 2011- present
- Biostatistics and Bioinformatics Cluster Hire Committee, Center for Biotechnology and Genomics, 2014 – 2015
- Center for Transformative Undergraduate Experiences (TrUE)/Center for Active Learning and Undergraduate Engagement (CALUE)/Center for Undergraduate Research (CUR), TTU, 2011 – present
- Graduate School Dean's representative, 2012-2015

PROFESSIONAL SERVICES (PARTIAL LIST)

- **General Chair/Co-Chair**
  - The 1st/2nd/3rd International Industry/University Workshop on Data-center Automation, Analytics, and Control, 2017-2019
  - The 1st Joint International Workshop on Parallel Data Storage & Data-Intensive Scalable Computing Systems (PDSW-DISCS), in conjunction with ACM/IEEE Supercomputing Conference (SC'16), 2016
  - The 2014 International Workshop on Data-Intensive Scalable Computing Systems (DISCS-2014), in conjunction with ACM/IEEE Supercomputing Conference (SC'14), 2014
- **Steering Committee**
  - International Parallel Data Systems Workshop (PDSW)
  - International Workshop on Parallel Programming Models and Systems Software for High-End Computing, Workshop Organizer
  - Joint International Workshop on Parallel Data Storage & Data-Intensive Scalable Computing Systems (PDSW-DISCS)
  - International Workshop on Data-Intensive Scalable Computing Systems (DISCS), in conjunction with ACM/IEEE Supercomputing Conference (SC)
- **Technical Program Chair/Co-Chair/Area Chair**:
  - Program Co-Chair, The 50th International Conference on Parallel Processing (ICPP-2021), 2021
  - Proceedings Chair, The 49th International Conference on Parallel Processing (ICPP-2020), 2020
  - Poster Chair, The 27th International Symposium on High-Performance Parallel and Distributed Computing (HPDC'18), 2018
  - The 4th-11th International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2), in conjunction with 40th-45th ICPP (International Conference on Parallel Processing), 2011-2018.
  - Program Chair, The 10th IEEE/ACM International Conference on Utility and Cloud Computing (UCC'17), 2017
  - Data Analytics, Visualization & Storage (DAVS Area), The 28th ACM/IEEE Supercomputing Conference 2016 (SC'16), 2016.
  - The 2013 International Workshop on Data-Intensive Scalable Computing Systems (DISCS-2013), in conjunction with ACM/IEEE Supercomputing Conference (SC'13), 2013
  - The 2012 International Workshop on Data-Intensive Scalable Computing Systems (DISCS-2012), in conjunction with ACM/IEEE Supercomputing Conference (SC'12), 2012
  - The 16th International Conference on Parallel and Distributed Systems (ICPADS'10), 2010
- **Doctoral Research Showcase Chair,** The 24th ACM/IEEE Supercomputing Conference (SC'12), 2012
- **Archiving Proceedings Chair**, The 24th ACM/IEEE Supercomputing Conference (SC'12), 2012

- **Publicity Chair:**
  - International Workshop on Software-Defined Ecosystems (BigSystem 2014), in conjunction with the ACM HPDC'14
  - The 3rd International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2), in conjunction with ICPP'10, 2010
  - The 2nd International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2), in conjunction with ICPP'09, 2009
- **Technical Program Committee:**
  - The 35th IEEE International Parallel & Distributed Processing Symposium (IEEE IPDPS 2021), 2021
  - The 21st IEEE/ACM Intl. Symp. on Cluster, Cloud & Grid Computing (CCGrid'21), 2021
  - The 2021 IEEE International Conference on Big Data (BigData'21), 2021
  - The 34th IEEE International Parallel & Distributed Processing Symposium (IEEE IPDPS 2020), 2020
  - The 2020 High Performance Computing in Asia-Pacific Region (HPC Asia 2020), 2020
  - The 20th IEEE/ACM Intl. Symp. on Cluster, Cloud & Grid Computing (CCGrid'20), 2020
  - The 2020 ACM/IEEE Supercomputing Conference (SC'20), 2020
  - The 49th International Conference on Parallel Processing (ICPP-2020), 2020
  - The 19th IEEE/ACM Intl. Symp. on Cluster, Cloud & Grid Computing (CCGrid'19), 2019
  - The 2019 IEEE International Conference on Big Data (BigData'19), 2019
  - The 12th International Conference on Cloud Computing (IEEE Cloud), 2019
  - The 10th Workshop on Scientific Cloud Computing (ScienceCloud'19) at HPDC'19, 2019
  - The 18th IEEE/ACM Intl. Symp. on Cluster, Cloud & Grid Computing (CCGrid'18), 2018
  - The 2018 IEEE International Conference on Big Data (BigData'18), 2018
  - The 11th IEEE/ACM International Conference on Utility and Cloud Computing (UCC'18), 2018
  - The 2018 ACM/IEEE Supercomputing Conference (SC'18) Doctoral Showcase, 2018
  - The 17th IEEE/ACM Intl. Symp. on Cluster, Cloud & Grid Computing (CCGrid'17), 2017
  - The 2017 IEEE International Conference on Big Data (BigData'17), 2017
  - The 2017 ACM/IEEE Supercomputing Conference (SC'17), 2017
  - The 10th International Conference on Cloud Computing (IEEE Cloud), 2017
  - The 2016 IEEE International Conference on Big Data (IEEE BigData 2016), 2016.
  - The 2016 ACM/IEEE Supercomputing Conference (SC'16), Best Paper Committee, BOF (Birds-of-Feather) Committee, 2016
  - The 21st IEEE International Conference on Parallel and Distributed Systems (ICPADS'16), 2016.
  - The 16th IEEE/ACM International Symp. on Cluster, Cloud and Grid Computing (CCGrid'16), 2016.
  - The International Workshop of Software-Defined Data Communications and Storage (SDDCS) 2016, in conjunction with IEEE ICDCS 2016, on June 27, 2016, in Nara, Japan.
  - The 27th ACM/IEEE Supercomputing Conference 2015 (SC'15), 2015.
  - The 10th IEEE International Conference on Networking, Architecture, and Storage (NAS'15), 2015.
  - The 8th International Conference on Cloud Computing (IEEE Cloud 2015), 2015.
  - The 44th International Conference on Parallel Processing (ICPP-2015), 2015.
  - The 2nd IEEE/ACM International Symposium on Big Data Computing (BDC'15), 2015.
  - The 2nd International Workshop on Software-Defined Ecosystems (BigSystem 2015)**,** co-located with ACM HPDC 2015 and ACM FCRC 2015 Federated Computing Research Conference, 2015.
  - The 14th IEEE International Symposium on Parallel and Distributed Computing (ISPDC), 2015.

- The 6th International Workshop on Programming Models and Applications for Multicores and Manycores (PMAM'15), held in conjunction with The ACM PPoPP 2015, 2015.
- 2014 IEEE International Conference on Big Data (IEEE BigData 2014), 2014.
- The 6th IEEE Intl. Conference on Cloud Computing Technology and Service (CloudCom'14), 2014
- The 33rd IEEE Intl. Performance Computing and Communications Conference (IPCCC'14), 2014
- The IEEE International Conference on Cluster Computing (Cluster'14), 2014
- The 26th ACM/IEEE Supercomputing Conference 2014 (SC'14), 2014.
- IEEE International Congress on Big Data, 2014
- The 9th IEEE International Conference on Networking, Architecture, and Storage (NAS'14), 2014.
- The 7th International Conference on Cloud Computing (IEEE Cloud 2014), 2014.
- The 14th IEEE/ACM Intl. Symp. on Cluster, Cloud & Grid Computing (CCGrid'14), 2014.
- The 28th IEEE International Parallel & Distributed Processing Symposium (IEEE IPDPS 2014), 2014.
- The 2014 Multicore and GPU Programming Models, Languages and Compilers (PLC) Workshop, in conjunction with the IPDPS'14, 2014
- 2013 IEEE International Conference on Big Data (IEEE BigData 2013), 2013.
- Doctoral Showcase Committee: The 25th ACM/IEEE Supercomputing Conference 2013 (SC'13), 2013.
- Fourth International Workshop on Parallel Software Tools and Tool Infrastructures (PSTI 2013), in conjunction with ICPP'13, 2013.
- The 6th International Conference on Cloud Computing (IEEE Cloud 2013), 2013.
- The 13th IEEE/ACM Intl. Symp. on Cluster, Cloud & Grid Computing (CCGrid'13), 2013.
- The 2013 Multicore and GPU Programming Models, Languages and Compilers (PLC) Workshop, in conjunction with the IPDPS'13, 2013
- The 18th International Workshop on High-Level Parallel Programming Models and Supportive Environments (HIPS'13), in conjunction with the IPDPS'13, 2013
- The 31st IEEE Intl. Performance Computing and Communication Conference (IPCCC'12), 2012
- The 18th Intl. Conference on Parallel and Distributed Systems (ICPADS'12), 2012.
- The 5th International Conference on Cloud Computing (IEEE Cloud 2012), 2012.
- The 7th IEEE International Conference on Networking, Architecture, and Storage (NAS'12), 2012.
- The 7th International Conference on Future Information Technology (FutureTech'12), 2012.
- The 26th IEEE Intl. Parallel & Distributed Processing Symposium (IPDPS'12), 2012
- The 10th IEEE Intl. Symp. on Parallel and Distributed Processing with Applications (ISPA'12), 2012.
- The 12th IEEE/ACM Intl. Symposium on Cluster, Cloud & Grid Computing (CCGrid'12), 2012.
- The 6th Petascale Data Storage Workshop (PDSW'11), in conjunction with ACM/IEEE Supercomputing'11 (SC'11), 2011.
- The 2nd International Workshop on Data Intensive Computing in the Clouds (DataCloud-SC11) in conjunction with ACM/IEEE Supercomputing SC'11 (SC'11), 2011.
- The 13th IEEE Intl. Conf. on High Performance Computing and Communications (HPCC-11), 2011.
- The 4th Intl. Conference on Cloud Computing (IEEE Cloud 2011), 2011.
- The 23rd ACM/IEEE Supercomputing Conference 2011 (SC'11).
- The 6th IEEE Intl. Conf. on Networking, Architecture, and Storage (NAS'11), 2011.
- The 31st Intl. Conference On Distributed Computing Systems (ICDCS'11), 2011
- The 9th IEEE Intl. Symposium on Parallel and Distributed Processing with Applications (ISPA'11), 2011
- The 11th IEEE/ACM Intl. Symposium on Cluster, Cloud & Grid Computing (CCGrid'11), 2011

  - The 5th Petascale Data Storage Workshop (PDSW'10), in conjunction with ACM/IEEE Supercomputing'10 (SC'10), 2010
  - The 17th IEEE Intl. Conference on High Performance Computing (HiPC'10), 2010
  - The 19th International Conference on Computer Communication and Networks (ICCCN'10), Track on High-speed Distributed Systems and Grids (HDSG), 2010
  - The Third International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2), in conjunction with ICPP'10, 2010
  - The 10th IEEE/ACM International Symposium on Cluster, Cloud & Grid Computing (CCGrid'10), 2010
  - IEEE International Workshop on Middleware Engineering (ME'09), 2009
  - The Second International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2), in conjunction with ICPP'09, 2009
- **Newsletter Editor**, IEEE Technical Committee on Scalable Computing (TCSC)
- **Referee for Journals**: IEEE Computer (TC), IEEE Transactions on Parallel and Distributed System (TPDS), Journal of Parallel and Distributed Computing (JPDC), International Journal of High Performance Computing Applications (IJHPCA), Parallel Computing (ParCo), etc.

PROFESSIONAL MEMBERSHIP

- ACM
- ACM SIGHPC (Special Interest Group on High Performance Computing)
- IEEE
- IEEE Computer Society

AWARDS AND HONORS

- 2021, Best Paper Award, The 35th IEEE International Parallel & Distributed Processing Symposium (IPDPS'21) (one out of 105 accepted papers)
- 2019, Best Student Poster Award, with Mr. Misha Ahmadian, The 2019 ACM Practice and Experience in Advanced Research Computing (PEARC'19).
- 2017, Best Student Poster Award, with Ms. E. Hojati and Dr. A. Sill. Benchmarking Automated Hardware Management Technologies for Modern Scalable Data Centers and Cloud Environments. The 10th IEEE/ACM International Conference on Utility and Cloud Computing (UCC'17).
- 2016, Best Paper Award and Best Paper Finalist, The 11th IEEE International Conference on Networking, Architecture, and Storage (NAS'16)
- 2016, Best Paper Award, The 14th IEEE International Symposium on Parallel and Distributed Processing with Applications (ISPA'16).
- 2014, Texas Tech University Whitacre College of Engineering Research Award
- 2014, Texas Tech University Mortar Board and Omicron Delta Kappa Outstanding Faculty Award
- 2014, IEEE TCSC (Technical Committee on Scalable Computing) Young Achievers Award, IEEE
- 2012, Ralph E. Powe Junior Faculty Enhancement Award, Oak Ridge Associated Universities
- 2012, Best Paper Award, The 7th International Conference on Future Information Technology (FutureTech'12), Vancouver, Canada, 2012
- 2011, Best Paper Award, The 9th IEEE International Symposium on Parallel and Distributed Processing with Applications (ISPA'11)
- 2009, ACM/IEEE High Performance Computing Ph.D. Fellowship, ACM and IEEE

- 2008, Feildhouse Research Fellowship, Illinois Institute of Technology
- 2008, Best Paper finalist, Best Student Paper finalist, ACM/IEEE Supercomputing Conference (SC'08)
- 2008, Chinese Government Award for Outstanding Self-financed Students Aboard, China
- 2007, Travel Grant for NSF Collaborative Expedition Workshop, NARA and NSF
- 2007, ACM/IEEE High Performance Computing Ph.D. Fellow Honorable Mention, ACM and IEEE
- 2006, Outstanding Teaching Assistant, Department of Computer Science, Illinois Institute of Technology
- 2003, Outstanding Graduate Student of AnHui Province and USTC, China
- 2001, Legend Group Scholarship, University of Science and Technology of China
- 2001, Outstanding Student Leader, University of Science and Technology of China

REFERENCES AVAILABLE UPON REQUEST