

Jianfeng ZHAN (詹剑锋),

Full Professor, University of Chinese Academy of Sciences & ICT, Chinese Academy of Sciences

Office: Room 1038, NO. 6, Kexueyuan South Road, Zhongguancun, Haidian District, Beijing, PRC

Mobile: +86-18600236326

E-mail: jianfengzhan.benchcouncil@gmail.com

Official homepage: <https://peopleucas.edu.cn/~zjf?language=en>

Personal homepage: <https://www.zhanjianfeng.org>

RESEARCH INTEREST

Opensource Computer Systems, Benchmarks, Evaluations, Parallel and Distributed Systems, and Operating Systems.

Jianfeng Zhan has been a Full Professor at the University of Chinese Academy of Sciences, and ICT, Chinese Academy of Sciences (CAS) since 2012 and Director at the Advanced system research center, ICT, CAS since 2022. He has supervised over 90 graduate students (both MS and Ph.D.), post-docs, and engineers. His research interests cover a broad spectrum of high-performance and distributed systems. A common thread among his research projects focuses on benchmarking, designing, implementing, and optimizing parallel and distributed systems. He has made substantial and effective efforts to transfer his academic research into advanced technology to impact general-purpose production systems. Several technical innovations and research results from his team have been widely adopted in benchmarks, operating systems, and cluster and cloud system software, with direct contributions to advancing parallel and distributed systems in China and even the world.

BigDataBench. Since its publication in HPCA 2014 (the citation of this paper is 700+), BigDataBench---an open-source big data benchmarking is widely used in academia and industry worldwide.

Datacenter Operating System. Thirty-six patents (including 10+ international patents) are transferred to the industry.

Cluster and Cloud system software. The patent and software are transferred to the industry.

Dr. Jianfeng Zhan is the founding chair of the International Open Benchmark Council (BenchCouncil) (<https://www.benchcouncil.org>), a non-profit international organization recognized as an internationally influential contributor to benchmarking and evaluations. He served as IEEE TPDS Associate Editor from 2018 to 2022 and is the founding editor-in-chief of BenchCouncil Transactions on Benchmarks, Standards, and Evaluations.

Jianfeng Zhan received his B.E. in Civil Engineering from Southwest Jiaotong University in 1996, his MSc in Solid Mechanics from Southwest Jiaotong University in 1999, and his Ph.D. in Computer Science from University of Chinese Academy of

Sciences, and Institute of Software, Chinese Academy of Sciences in 2002.

ACADEMIC POSITION

2022-up to now, Director, Advanced computer systems research center, ICT, Chinese Academy of Sciences

Sep 2014-2022, Deputy Director, Advanced computer systems research center, ICT, Chinese Academy of Sciences

Sep 2012-up to now, Full Professor, University of Chinese Academy of Sciences, and ICT, Chinese Academy of Sciences, and University of Chinese Academy of Sciences, Beijing, P. R. China

Mar 2004- Sep 2012, Associate Professor, University of Chinese Academy of Sciences, and ICT, Chinese Academy of Sciences, and University of Chinese Academy of Sciences, Beijing, P. R. China

Aug 2002- Mar 2004, Assistant Professor, University of Chinese Academy of Sciences, and ICT, Chinese Academy of Sciences, and University of Chinese Academy of Sciences, Beijing, P. R. China

Honors and Awards

Second Class of National Science and Technology Promotion Prize, The Central People's Government of the People's Republic of China, 2006.

Outstanding Science and Technology Achievement Prize of the Chinese Academy of Sciences, Chinese Academy of Science, P. R. China, 2005

IISWC 2013 Best Paper Award.

Annual Excellent faculty Prize, Institute of Computing Technology, Chinese Academy of Sciences, P.R. China, 2004, 2008, 2012, 2017

Funding:

□ 2019-2020, AI Benchmarking, ¥ 10M, PI.

□ 2016-2019, Scientific Big data Systems, ¥ 6M, PI.

□ 2015-2017, Big data benchmarking, ¥ 3M, Co-PI.

□ 2012-2016, Sea-cloud computing systems, ¥ 20M, PI.

□ 2012-2014, The operating systems for datacenter computing, ¥ 30M, PI.

□ 2007-2011, National '863' High-Tech Program of China, Cluster computer system software.

¥200M, Major participant.

□ 2005-2006, the 15th key project of China (No. 2004BA811B09-1), Industrial information Grid system, ¥ 500K, PI.

□ 2004-2005, National '863' High-Tech Program of China (No. 2004AA616010), Earth information processing on the Grid system, ¥ 300K, Co-PI.

Professional Activities

BenchCouncil Transactions on Benchmarks, Standards and Evaluations, Founding Editors-in-Chief. 2021---up to now

IEEE Transaction on Parallel and Distributed Systems (TPDS), Associate Editor
PC Member, IPDPS 12, 18
PC Member, ICDCS 12, 17, 18, 19
DBTest'18, in conjunction with SIGMOD'18
Co-Chair, BPOE 14, 15, 16, 17, 18 in conjunction with ASPLOS and VLDB
Co-chair, HPBDC'16, 17, 18, in conjunction with IPDPS
PC Member, UCC 2017
PC Member, SCCTSA 15, 16, 17
PC Member, ICDCS 2017
PC Member, DSS-2017
PC Member of ICPP 2016
Co-Chair BPOE-7, 2016
Guest editor of big data benchmark, performance optimization and emerging hardware special issue of IEEE Transaction on Emerging Techniques in Computing 2016
PC Member, TPCTC 14, 15
PC Member, CCGrid 14, 15
PC Member, ICCCN 14, 15
Track Chair, CloudCOM, 15
PC Member, ICPE 14
PC Member of SOSE 2013
PC Member of ICAC 2013
Track Chair of HPCC 2013
PC Member of ACM CAC 2013.
PC Member of AINA 2012 (The 26th IEEE International Conference on Advanced Information Networking and Applications, Tokyo, Japan, March 26-29, 2012.)
Track Chair, ICPADS, 2012
Publicity Chair for China of ICAC 2012 (The 9th International Conference on Autonomic Computing)
PC Member of IDPDS 2012 Ph.D Forum
Guest editor of Cloud computing special issue of Frontier of Computer Sciences 2012
PC member of NPC 11, 12
PC Member of CSE 2011
PC Member of GCC 2011
PC Member of Cloud 2011

Selective Publications

- 1) Zhang, F., Lan, C., Wang, L., Tang, F., Dai, S., Wang, J., Ma, J. & Zhan, J. (2023, May). EAI Bench: An Energy Efficiency Benchmark for AI Training. In Benchmarking, Measuring, and Optimizing: 14th BenchCouncil International Symposium, Bench 2022, Virtual Event, November 7-9, 2022, Revised Selected Papers (pp. 19-34). Cham: Springer International Publishing.
- 2) Wang, P., Wu, D., Chen, C., Liu, K., Fu, Y., Huang, J., Zhou, Y., Zhan, J. & Hua,

- X. (2023). Deep Adaptive Graph Clustering via von Mises-Fisher Distributions. *ACM Transactions on the Web*.
- 3) Jiang, Z., Luo, C., Gao, W., Wang, L., & Zhan, J. (2022). HPC AI500 V3. 0: A scalable HPC AI benchmarking framework. *BenchCouncil Transactions on Benchmarks, Standards and Evaluations*, 100083.
 - 4) Chunjie L U O, Jianfeng Z. Multi-layer dynamic and asymmetric convolutions[J]. *HIGH TECHNOLOGY LETTERS*, 2022, 28(3): 227-236.
 - 5) Hao, T., Gao, W., Lan, C., Tang, F., Jiang, Z., & Zhan, J. (2022). Edge AIBench 2.0: A scalable autonomous vehicle benchmark for IoT-Edge-Cloud systems. *BenchCouncil Transactions on Benchmarks, Standards and Evaluations*, 2(4), 100086.
 - 6) Gao, W., Wang, L., Chen, M., Xiong, J., Luo, C., Zhang, W., Huang, Y., Li, W., Kang, G., Zheng C., Xie, B., Dai, S., He, Q., Ye, H., Bao, Y. & Zhan, J. (2022). High fusion computers: The IoTs, edges, data centers, and humans-in-the-loop as a computer. *BenchCouncil Transactions on Benchmarks, Standards and Evaluations*, 2(3), 100075.
 - 7) Zhan, Jianfeng. "A BenchCouncil view on benchmarking emerging and future computing." *BenchCouncil Transactions on Benchmarks, Standards and Evaluations* (2022): 100064.
 - 8) Hao, T., Hwang, K., Zhan, J., Li, Y., & Cao, Y. (2022). Scenario-based AI Benchmark Evaluation of Distributed Cloud/Edge Computing Systems. *IEEE Transactions on Computers*.
 - 9) Kang, G., Wang, L., Gao, W., Tang, F., & Zhan, J. (2022, May). OLPBench: Real-time, Semantically Consistent, and Domain-specific are Essential in Benchmarking, Designing, and Implementing HTAP Systems. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)* (pp. 1822-1834). IEEE.
 - 10) Zhan, Jianfeng. "Open-source computer systems initiative: The motivation, essence, challenges, and methodology." *BenchCouncil Transactions on Benchmarks, Standards and Evaluations* (2022): 100038.
 - 11) Li, Yatao, and Jianfeng Zhan. "SAIBench: Benchmarking AI for science." *BenchCouncil Transactions on Benchmarks, Standards and Evaluations* 2.2 (2022): 100063.
 - 12) Zhan, Jianfeng. "Three laws of technology rise or fall." *BenchCouncil Transactions on Benchmarks, Standards and Evaluations* (2022): 100034.
 - 13) Wang, N., Luo, C., Huang, X., Huang, Y., & Zhan, J. (2022). DeepCS: Training a deep learning model for cervical spondylosis recognition on small-labeled sensor data. *Neurocomputing*, 472, 24-34.
 - 14) Wang, N., Luo, C., Huang, Y., & Zhan, J. (2021, December). Ensemble Clustering-based Cervical Spondylosis Fine-classification. In *2021 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)* (pp. 1090-1097). IEEE.
 - 15) Wei, Z., Lei, W., Dan, M., Lin, Y., & Jianfeng, Z. (2021). Scalable group management in large-scale virtualized clusters. *HIGH TECHNOLOGY LETTERS*, 17(3), 263-271.

- 16) Zhang, F., Luo, C., Lan, C., & Zhan, J. (2021). Benchmarking feature selection methods with different prediction models on large-scale healthcare event data. *BenchCouncil Transactions on Benchmarks, Standards and Evaluations*, 1(1), 100004.
- 17) Jiang, Z., Gao, W., Tang, F., Wang, L., Xiong, X., Luo, C., Lan, C., Li, H. & Zhan, J. (2021, September). Hpc ai500 v2. 0: The methodology, tools, and metrics for benchmarking hpc ai systems. In *2021 IEEE International Conference on Cluster Computing (CLUSTER)* (pp. 47-58). IEEE.
- 18) Luo, C., Zhan, J., Wang, L., & Gao, W. (2021, July). Finet: Using fine-grained batch normalization to train light-weight neural networks. In *2021 International Joint Conference on Neural Networks (IJCNN)* (pp. 1-8). IEEE.
- 19) Wang, L., Xiong, X., Zhan, J., Gao, W., Wen, X., Kang, G., & Tang, F. (2021). Wpc: Whole-picture workload characterization across intermediate representation, isa, and microarchitecture. *IEEE Computer Architecture Letters*, 20(2), 86-89.
- 20) Hao, T., Zhan, J., Hwang, K., Gao, W., & Wen, X. (2021, May). AI-oriented workload allocation for cloud-edge computing. In *2021 IEEE/ACM 21st International Symposium on Cluster, Cloud and Internet Computing (CCGrid)* (pp. 555-564). IEEE.
- 21) Fei Tang, Wanling Gao, Jianfeng Zhan, Chuanxin Lan, Xu Wen, Lei Wang, Chunjie Luo, Zheng Cao, Xingwang Xiong, Zihan Jiang, Tianshu Hao, Fanda Fan, Fan Zhang, Yunyou Huang, Jianan Chen, Mengjia Du, Rui Ren, Chen Zheng, Daoyi Zheng, Haoning Tang, Kunlin Zhan, Biao Wang, Defei Kong, Minghe Yu, Chongkang Tan, Huan Li, Xinhui Tian, Yatao Li, Junchao Shao, Zhenyu Wang, Xiaoyu Wang, Jiahui Dai & Hainan Ye. (2021, March). AIBench training: Balanced industry-standard AI training benchmarking. In *2021 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)* (pp. 24-35). IEEE.
- 22) Yunyou, H., Nana, W., Tianshu, H., Xiaoxu, G., Chunjie, L., Lei, W., Rui, R. & Jianfeng, Z. (2021). Clustering residential electricity load curve via community detection in network. *HIGH TECHNOLOGY LETTERS*, 27(1), 53-61.
- 23) Kang, G., Kong, D., Wang, L., & Zhan, J. (2021). OStoreBench: benchmarking distributed object storage systems using real-world application scenarios. In *Benchmarking, Measuring, and Optimizing: Third BenchCouncil International Symposium, Bench 2020, Virtual Event, November 15–16, 2020, Revised Selected Papers 3* (pp. 90-105). Springer International Publishing.
- 24) Liang, Y., Guo, Y., Gong, Y., Luo, C., Zhan, J., & Huang, Y. (2021). Flbench: A benchmark suite for federated learning. In *Intelligent Computing and Block Chain: First BenchCouncil International Federated Conferences, FICC 2020, Qingdao, China, October 30–November 3, 2020, Revised Selected Papers 1* (pp. 166-176). Springer Singapore.
- 25) Gao, W., Zhan, J., Fox, G., Lu, X., & Stanzione, D. (Eds.). (2020). *Benchmarking, Measuring, and Optimizing: Second BenchCouncil International Symposium, Bench 2019, Denver, CO, USA, November 14–16, 2019, Revised*

- Selected Papers (Vol. 12093). Springer Nature.
- 26) Ren, R., Cheng, J., He, X. W., Wang, L., Zhan, J. F., Gao, W. L., & Luo, C. J. (2019). Hybridtune: spatio-temporal performance data correlation for performance diagnosis of big data systems. *Journal of Computer Science and Technology*, 34, 1167-1184.
 - 27) Gao, W., Tang, F., Zhan, J., Wen, X., Wang, L., Cao, Z., Lan, C., Luo, C., Liu, X. & Jiang, Z. (2021, September). Aibench scenario: Scenario-distilling ai benchmarking. In 2021 30th International Conference on Parallel Architectures and Compilation Techniques (PACT) (pp. 142-158). IEEE.
 - 28) Jianfeng Zhan. On How to Develop Independent, and Open Science and Technology Industries in China[J]. *Bulletin of Chinese Academy of Sciences*, 2019, 34(6): 657-666
 - 29) WF Zeng, XX Zhou, WJ Zhou, H Chi, J Zhan, SM He, pDeep2 Enables MS/MS Spectrum Prediction for Modified Peptides using Transfer Learning, *Analytical Chemistry*.
 - 30) Y Huang, J Zhan, C Luo, L Wang, N Wang, D Zheng, F Fan, R Ren, An electricity consumption model for synthesizing scalable electricity load curves, *Energy* 169, 674-683
 - 31) XX Zhou, WF Zeng, H Chi, C Luo, C Liu, J Zhan, SM He, Z Zhang. pdeep: Predicting MS/MS spectra of peptides with deep learning. *Analytical chemistry* 89 (23), 12690-12697
 - 32) R Ren, C Zheng, J Zhan (Edit), *Big Scientific Data Benchmarks, Architecture, and Systems*, Springer Singapore
 - 33) Zhifei Zhang, Wanling Gao, Fan Zhang, Yunyou Huang, Shaopeng Dai, Fanda Fan, Jianfeng Zhan, Mengjia Du, Silin Ying, Longxing Xiong, Juan Du, Yumei Cheng, Xiexuan Zhou, Rui Ren, Lei Wang, Hainan Ye. *Landscape of Big Medical Data: A Pragmatic Survey on Prioritized Tasks*. IEEE Access. 2018
 - 34) Rui Ren, Jiechao Cheng, Yan Yin, Jianfeng Zhan, Lei Wang, Jinheng Li, Chunjie Luo. Deep Convolutional Neural Networks for Log Event Classification on Distributed Cluster Systems. 2018 IEEE International Conference on Big Data (Big Data), Dec 2018.
 - 35) Wanling Gao, Jianfeng Zhan, Lei Wang, Chunjie Luo, Zhen Jia, Daoyi Zheng, Chen Zheng, Xiwen He, Hainan Ye, Haibin Wang, Rui Ren. Data Motif-based Proxy Benchmarks for Big Data and AI Workloads. IISWC 2018
 - 36) Wanling Gao, Jianfeng Zhan, Lei Wang, Chunjie Luo, Daoyi Zheng, Fei Tang, Biwei Xie, Chen Zheng, Xu Wen, Xiwen He, Hainan Ye, Rui Ren, Data Motifs: A Lens Towards Fully Understanding Big Data and AI Workloads. The 27th International Conference on Parallel Architectures and Compilation Techniques (PACT 18).
 - 37) Chunjie Luo, Jianfeng Zhan, Xiaohe Xue, Lei Wang, Rui Ren, Qiang Yang. Cosine normalization: Using cosine similarity instead of dot product in neural networks. 2018 International Conference on Artificial Neural Networks, 382-391
 - 38) Chen Zheng, Lei Wang, Lixin Zhang, Hainan Ye, Jianfeng Zhan (Corresponding author), XOS: an Application-defined Operating System for Datacenter

- Computing. IEEE International Conference on Big Data. 2018.
- 39) Biwei Xie, Jianfeng Zhan, Xu Liu, Wanling Gao, Zhen Jia, Xiwen He, Lixin Zhang. CVR: efficient vectorization of SpMV on x86 processors. Proceedings of the 2018 International Symposium on Code Generation and Optimization (CGO 18).
 - 40) R Han, LK John, J Zhan. Benchmarking big data systems: A review. IEEE Transactions on Services Computing 11 (3), 580-597, 2018
 - 41) X Tian, J Zhan. GraphDuo: A Dual-Model Graph Processing Framework. IEEE Access 6, 35057-35071, 2018
 - 42) X Tian, Y Guo, J Zhan, L Wang. Towards memory and computation efficient graph processing on spark. Big Data (Big Data), 2017 IEEE International Conference on, 375-382
 - 43) Z Jia, W Gao, Y Shi, SA McKee, J Zhan, L Wang, L Zhang. Understanding processors design decisions for data analytics in homogeneous data centers. IEEE Transactions on Big Data, 2017.
 - 44) Ning Liu, Dengming Zhu, Zhaoqi Wang, Hong Qin, Jianfeng Zhan, Jinzhu Gao: Pipelining image compositing in heterogeneous networking environments. Journal of Visualization and Computer Animation 27(3-4): 385-393 (2016).
 - 45) Zhen Jia, Chao Xue, Guancheng Chen, Jianfeng Zhan, Lixin Zhang, Yonghua Lin, Peter Hofstee: Auto-tuning Spark Big Data Workloads on POWER8: Prediction-Based Dynamic SMT Threading. PACT 2016: 387-400.
 - 46) Rui Ren, Zhen Jia, Lei Wang, Jianfeng Zhan, Tianxu Yi: BDTune: Hierarchical correlation-based performance analysis and rule-based diagnosis for big data systems. BigData 2016: 555-562.
 - 47) Chen Zheng, Jianfeng Zhan, Zhen Jia, Lixin Zhang: Characterizing OS Behaviors of Datacenter and Big Data Workloads. HPC/SmartCity/DSS 2016: 1079-1086.
 - 48) Chuliang Weng, Jianfeng Zhan, Yuan Luo: TSAC: Enforcing Isolation of Virtual Machines in Clouds. IEEE Trans. Computers 64(5): 1470-1482 (2015).
 - 49) Gang Lu, Jianfeng Zhan (Corresponding author), Haining Wang, Lin Yuan, Yunwei Gao, Chuliang Weng, Yong Qi: PowerTracer: Tracing Requests in Multi-Tier Services to Reduce Energy Inefficiency. IEEE Trans. Computers 64(5): 1389-1401 (2015).
 - 50) Lei Wang, **Jianfeng Zhan (corresponding author)**, Chunjie Luo, Yuqing Zhu, Qiang Yang, Yongqiang He, Wanling Gao, Zhen Jia, Yingjie Shi, Shujie Zhang, Chen Zheng, Gang Lu, Kent Zhan, Xiaona Li, Bizhu Qiu: BigDataBench: A big data benchmark suite from internet services. HPCA 2014: 488-499. (Google Citation **696**, access time, 2023.05.18)
 - 51) Zhen Jia, Lei Wang, **Jianfeng Zhan (corresponding author)**, Lixin Zhang, Chunjie Luo. Characterizing data analysis workloads in data centers. Workload Characterization (IISWC), 2013 IEEE International Symposium on, 66-76 (Google Citation **145**, access time, 2023.05.18). **IISWC 2013 Best Paper Award.**

- 52) Zhen Jia, Jianfeng Zhan (corresponding author), Lei Wang, Rui Han, Sally A. McKee, Qiang Yang, Chunjie Luo, Jingwei Li: Characterizing and subsetting big data workloads. IISWC 2014: 191-201.
- 53) Zijian Ming, Chunjie Luo, Wanling Gao, Rui Han, Qiang Yang, Lei Wang, Jianfeng Zhan:
BDGS: A Scalable Big Data Generator Suite in Big Data Benchmarking. WBDB 2013: 138-154.
- 54) Gang Lu, Jianfeng Zhan, Chongkang Tan, Xinlong Lin, Defei Kong, Tianshu Hao, Lei Wang, Fei Tang, Chen Zheng: Isolate First, Then Share: A New OS Architecture for Datacenter Computing. arXiv:1604.01378.
- 55) Xiaoyu Fu, Ren Ren, Jianfeng Zhan, Wei Zhou, Zhen Jia, Gang Lu. LogMaster: mining event correlations in logs of large-scale cluster systems. Reliable Distributed Systems (SRDS), 2012 IEEE 31st Symposium on, 71-80
- 56) Jianfeng Zhan (editor), Rui Han, Chuliang Weng: Big Data Benchmarks, Performance Optimization, and Emerging Hardware - 4th and 5th Workshops, BPOE 2014, Salt Lake City, USA, March 1, 2014 and Hangzhou, China, September 5, 2014, Revised Selected Papers. Lecture Notes in Computer Science 8807, Springer 2014, ISBN 978-3-319-13020-0 [contents].
- 57) Rui Han, Siguang Huang, Fei Tang, Fu-Gui Chang, Jianfeng Zhan:
AccuracyTrader: Accuracy-Aware Approximate Processing for Low Tail Latency and High Result Accuracy in Cloud Online Services. ICCP 2016: 278-287.
- 58) Rui Han, Junwei Wang, Fengming Ge, José Luis Vázquez-Poletti, Jianfeng Zhan:
SARP: producing approximate results with small correctness losses for cloud interactive services. Conf. Computing Frontiers 2015: 22:1-22:8.
- 59) Rui Han, Junwei Wang, Siguang Huang, Chenrong Shao, Shulin Zhan, Jianfeng Zhan, José Luis Vázquez-Poletti:i: PCS: Predictive Component-Level Scheduling for Reducing Tail Latency in Cloud Online Services. ICCP 2015: 490-499.
- 60) Gang Lu, Jianfeng Zhan, Haining Wang, Lin Yuan, Chuliang Weng:
PowerTracer: tracing requests in multi-tier services to diagnose energy inefficiency. ICAC 2012: 97-102.
- 61) Bibo Tu, Jianping Fan, Jianfeng Zhan, Xiaofang Zhao: Accurate Analytical Models for Message Passing on Multi-core Clusters. PDP 2009: 133-139.
- 62) Jianfeng Zhan, Gengpu Liu, Lei Wang, Bibo Tu, Yi Jin, Yang Li, Yan Hao, Xuehai Hong, Dan Meng, Ninghui Sun: PhoenixG: A Unified Management Framework for Industrial Information Grid. CCGRID 2006: 489-496.
- 63) Ying Jiang, Dan Meng, Chao Ren, Jianfeng Zhan:
An Integrated Adaptive Management System for Cluster-based Web Services. CLUSTER 2006.
- 64) Xiaoyu Fu, Rui Ren, Sally A. McKee, Jianfeng Zhan, Ninghui Sun:
Digging deeper into cluster system logs for failure prediction and root cause diagnosis. CLUSTER 2014: 103-112.
- 65) Xiaoyu Fu, Rui Ren, Jianfeng Zhan, Wei Zhou, Zhen Jia, Gang Lu:
LogMaster: Mining Event Correlations in Logs of Large-Scale Cluster Systems. SRDS 2012: 71-80.

- 66) Wei Zhou, Jianfeng Zhan, Dan Meng, Zhihong Zhang:
Online Event Correlations Analysis in System Logs of Large-Scale Cluster Systems. NPC 2010: 262-276.
- 67) Zhihong Zhang, Jianfeng Zhan, Yong Li, Lei Wang, Dan Meng, Bo Sang:
Precise request tracing and performance debugging for multi-tier services of black boxes. DSN 2009: 337-346.
- 68) Lei Wang, Jianfeng Zhan (corresponding author), Weisong Shi, Yi Liang, In cloud, can scientific communities benefit from the economies of scale? IEEE Transactions on Parallel and Distributed Systems 23 (2), 296-303 (Google Citation **215**, access time, 2023.05.18)
- 69) Jianfeng Zhan, Lei Wang, Xiaona Li, Weisong Shi, Chuliang Weng, Wenyao Zhang, Xiutao Zang:
Cost-Aware Cooperative Resource Provisioning for Heterogeneous Workloads in Data Centers. IEEE Trans. Computers 62(11): 2155-2168 (2013).
- 70) Peng Wang, Dan Meng, Jizhong Han, Jianfeng Zhan, Bibo Tu, Xiaofeng Shi, Le Wan:
Transformer: A New Paradigm for Building Data-Parallel Programming Models. IEEE Micro 30(4): 55-64 (2010).
- 71) Zhen Jia, Jianfeng Zhan (corresponding author), Lei Wang, Chunjie Luo, Wanling Gao, Yi Jin, Rui Han, Lixin Zhang: Understanding Big Data Analytics Workloads on Modern Processors. IEEE Trans. Parallel Distrib. Syst. 28(6): 1797-1810 (2017).
- 72) Chunjie Luo, Jianfeng Zhan, Zhen Jia, Lei Wang, Gang Lu, Lixin Zhang, Chengzhong Xu, Ning Sun. Cloudrank-d: benchmarking and ranking cloud computing systems for data processing applications. Frontiers of Computer Science 6 (4), 347-362
- 73) Rui Han, Siguang Huang, Zhentao Wang, Jianfeng Zhan:
CLAP: Component-Level Approximate Processing for Low Tail Latency and High Result Accuracy in Cloud Online Services. IEEE Trans. Parallel Distrib. Syst. 28(8): 2190-2203 (2017).
- 74) Rui Han, Jianfeng Zhan, José Luis Vázquez-Poletti:
SARP: Synopsis-Based Approximate Request Processing for Low Latency and Small Correctness Loss in Cloud Online Services. International Journal of Parallel Programming 44(5): 1054-1077 (2016).
- 75) Bibo Tu, Jianping Fan, Jianfeng Zhan, Xiaofang Zhao: Performance analysis and optimization of MPI collective operations on multi-core clusters. The Journal of Supercomputing 60(1): 141-162 (2012).
- 76) Xu Liu, Jianfeng Zhan, Kunlin Zhan, Weisong Shi, Lin Yuan, Dan Meng, Lei Wang:
Automatic performance debugging of SPMD-style parallel programs. J. Parallel Distrib. Comput. 71(7): 925-937 (2011).
- 77) Pengfei Zheng, Yong Qi, Yangfan Zhou, Pengfei Chen, Jianfeng Zhan, Michael R. Lyu: An Automatic Framework for Detecting and Characterizing Performance Degradation of Software Systems. IEEE Trans. Reliability 63(4): 927-943 (2014).

- 78) Wenyao Zhang, Yi Wang, Jianfeng Zhan, Beichen Liu, Jianguo Ning:
Parallel Streamline Placement for 2D Flow Fields. IEEE Trans. Vis. Comput.
Graph. 19(7): 1185-1198 (2013).
- 79) Jianfeng Zhan, Ninghui Sun: Fire Phoenix Cluster Operating System Kernel and
its Evaluation. 2005 IEEE International Conference on Cluster Computing
(Cluster): 1-9.