

Zhengyu Yang

Email: yang.zhe@husky.neu.edu

Website: <http://nucsr.coe.neu.edu/?q=yangzhengyu>

EDUCATION

- 2011.9 - 2018.8** **Northeastern University (Boston, USA)**
Ph.D. in Computer Engineering. GPA=3.667/4=91.68%.
- 2010.9 - 2011.6** **The Hong Kong University of Science and Technology (Hong Kong)**
M.Sc. in Telecommunications. GPA=96.50%.
- 2006.9 - 2010.6** **Tongji University (Shanghai, China)**
B.Eng. in Communication Engineering. GPA=85.80%

PROFESSIONAL EXPERIENCE

- Sr. Software Engineer** **Samsung Semiconductor Inc.**, 09/2018-Present.
Memory Solutions Research Lab, San Diego CA.
- Summer Intern.** **Samsung Semiconductor Inc.**, 05/2017-08/2017.
Memory Solutions Research Lab, San Diego CA.
- Summer Intern.** **Samsung Semiconductor Inc.**, 05/2016-08/2016.
Memory Solutions Research Lab, San Diego CA.
- Summer Intern.** **Samsung Semiconductor Inc.**, 06/2015-08/2015.
Memory Solutions Research Lab, Milpitas CA.
- Research Asst.** **Northeastern University Computer Systems Research Lab**, 2012 - 2018.
Ph.D. Advisor: Associate Professor Ningfang Mi.
- Teaching Asst.** EECE 2560 Fundamentals of Engineering Algorithms - 2016 Spring, 2017 Spring.
EECE 5643 Simulation & Performance Evaluation - 2015 Fall.
EECE 3324 Computer Architecture - 2015 Spring.
EECE 4638 Special Topics Computer Engineering & Algorithms - 2014 Fall.
EECE 3326 Optimization Methods - 2013 Spring & Fall, 2014 Spring & Fall.
EECE 7205 Fundamentals of Computer Engineering - 2012 Fall.

RESEARCH AREAS AND SKILLS

Storage: Flash SSD FTL Design, Write Amplification, NVMe over Fabric, and NVMe Driver & Firmware Design.
Cache: Cloud-based Tiered Storage Cache Algorithm Design, and Virtualization Hypervisors & Hosting Solution.
Datacenter: Workload Characterization, VM Migration, Deduplication & Replication, and Disaster Recovery.
Big Data: Resource Management, Capacity Planning and Scheduling in Hadoop, Spark and Docker Platforms.
Systems: Linux Kernel, File Systems, Block Driver, IO Path, Performance Evaluation, and System Simulation.
Programming Skills: C++, Python, Java, Shell Script, SQL, JDBC, XSTL, Matlab, PHP, WordPress, and Drupal.

PUBLICATIONS

01. Jake Roemer, Mark Groman, Zhengyu Yang, Yufeng Wang, Chiu C. Tan, and Ningfang Mi, "Improving Virtual Machine Migration via Deduplication", 11th IEEE International Conference on Mobile Ad Hoc and Sensor Systems (MASS 2014), 2014.
02. Jianzhe Tai, Deng Liu, Zhengyu Yang, Xiaoyun Zhu, Jack Lo, and Ningfang Mi, "Improving Flash Resource Utilization at Minimal Management Cost in Virtualized Flash-based Storage Systems", IEEE Transactions on Cloud Computing, vol. 5, no. 3, pp. 537-549, July-Sept. 1 2017.
03. Zhengyu Yang, Manu Awasthi, Mrinmoy Ghosh, and Ningfang Mi, "A Fresh Perspective on Total Cost of Ownership Models for Flash Storage in Datacenters", 8th IEEE International Conference on Cloud Computing Technology and Science (CloudCom 2016, Acceptance Rate: 25.91%), 2016.

04. Zhengyu Yang, Jianzhe Tai, Janki Bhimani, Jiayin Wang, Ningfang Mi, and Bo Sheng, "GREM: Dynamic SSD Resource Allocation in Virtualized Storage Systems with Heterogeneous IO Workloads", 35th IEEE International Performance Computing and Communications Conference (IPCCC 2016, Acceptance Rate: 25.50%), 2016.
05. Jiayin Wang, Teng Wang, Zhengyu Yang, Ningfang Mi, and Bo Sheng, "eSplash: Efficient Speculation in Large Scale Heterogeneous Computing Systems", 35th IEEE International Performance Computing and Communications Conference (IPCCC 2016, Acceptance Rate: 25.50%), 2016.
06. Janki Bhimani, Jingpei Yang, Zhengyu Yang, Ningfang Mi, Qiumin Xu, Manu Awasthi, Rajinikanth Pandurangan, and Vijay Balakrishnan, "Understanding Performance of I/O Intensive Containerized Applications for NVMe SSDs", 35th IEEE International Performance Computing and Communications Conference (IPCCC 2016, Acceptance Rate: 25.50%), 2016.
07. Zhengyu Yang, Jiayin Wang, David Evans, and Ningfang Mi, "AutoReplica: Automatic Data Replica Manager in Distributed Caching and Data Processing Systems", 1st IEEE International Workshop on Communication, Computing, and Networking in Cyber Physical Systems (CCNCPS 2016), 2016.
08. Jiayin Wang, Teng Wang, Zhengyu Yang, Ying Mao, Ningfang Mi, and Bo Sheng, "SEINA: A Stealthy and Effective Internal Attack in Hadoop Systems", International Conference on Computing, Networking and Communications (ICNC 2017), 2017.
09. Han Gao, Zhengyu Yang, Janki Bhimani, Teng Wang, Jiayin Wang, Bo Sheng, and Ningfang Mi, "AutoPath: Harnessing Parallel Execution Paths for Efficient Resource Allocation in Multi-Stage Big Data Frameworks", 26th International Conference on Computer Communications and Networks (ICCCN 2017, Acceptance Rate: 25%), 2017.
10. Teng Wang, Jiayin Wang, Nam Nguyen, Zhengyu Yang, Ningfang Mi, and Bo Sheng, "EA2S2: An Efficient Application-Aware Storage System for Big Data Processing in Heterogeneous Clusters", 26th International Conference on Computer Communications and Networks (ICCCN 2017, Acceptance Rate: 25%), 2017.
11. Janki Bhimani, Ningfang Mi, Miriam Leeser, and Zhengyu Yang, "FiM: Performance Prediction Model for Parallel Computation in Iterative Data Processing Applications", 10th IEEE International Conference on Cloud Computing (CLOUD 2017), 2017.
12. Janki Bhimani, Zhengyu Yang, Miriam Leeser, and Ningfang Mi, "Accelerating Big Data Applications Using Lightweight Virtualization Framework on Enterprise Cloud", 21st IEEE High Performance Extreme Computing Conference (HPEC 2017), 2017.
13. Zhengyu Yang, Morteza Hoseinzadeh, Allen Andrews, Clay Mayers, David Thomas Evans, Rory Thomas Bolt, Janki Bhimani, Ningfang Mi, and Steven Swanson, "AutoTiering: Automatic Data Placement Manager in Multi-Tier All-Flash Datacenter", 36th IEEE International Performance Computing and Communications Conference (IPCCC 2017, Acceptance Rate: 33.3%), 2017.
14. Zhengyu Yang, Morteza Hoseinzadeh, Ping Wong, John Artoux, Clay Mayers, David Thomas Evans, Rory Thomas Bolt, Janki Bhimani, Ningfang Mi, and Steven Swanson, "H-NVMe: A Hybrid Framework of NVMe-based Storage System in Cloud Computing Environment", 36th IEEE International Performance Computing and Communications Conference (IPCCC 2017, Acceptance Rate: 33.3%, Best Paper Award), 2017.
15. Baiyu Chen, Zhengyu Yang, Siyu Huang, Xianzhi Du, Zhiwei Cui, Janki Bhimani, Xin Xie, and Ningfang Mi, "Cyber-Physical System Enabled Nearby Traffic Flow Modelling for Autonomous Vehicles", 36th IEEE International Performance Computing and Communications Conference, Special Session on Cyber Physical Systems: Security, Computing, and Performance (IPCCC-CPS 2017), 2017.
16. Janki Bhimani, Jingpei Yang, Zhengyu Yang, Ningfang Mi, N.H.V. Krishna Giri, Rajinikanth Pandurangan, Changho Choi, and Vijay Balakrishnan, "Enhancing SSDs with Multi-Stream: What? Why? How?", 36th IEEE International Performance Computing and Communications Conference (IPCCC 2017), Poster Paper, 2017.
17. Zhengyu Yang, Janki Bhimani, Jiayin Wang, David Evans, and Ningfang Mi, "Automatic and Scalable Data Replication Manager in Distributed Computation and Storage Infrastructure of Cyber-Physical Systems", Scalable Computing: Practice and Experience, Special Issue on Communication, Computing, and Networking in Cyber-Physical Systems, Vol. 18, No. 4, pp. 291–311.
18. Janki Bhimani, Zhengyu Yang, Ningfang Mi, Jingpei Yang, Qiumin Xu, Manu Awasthi, Rajinikanth Pandurangan, and Vijay Balakrishnan, "Docker Container Scheduler for I/O Intensive Applications running on NVMe SSDs", IEEE Transactions on Multi-Scale Computing Systems (TMSCS), ISSN: 2332-7766, Online ISSN: 2332-7766, DOI: 10.1109/TMSCS.2018.2801281.

19. Zhengyu Yang, Yufeng Wang, Janki Bhimani, Chiu C. Tan, and Ningfang Mi, "EAD: Elasticity Aware Deduplication Manager for Datacenters with Multi-tier Storage Systems", Cluster Computing, DOI: 10.1007/s10586-018-2141-z, 2018.
20. Zhengyu Yang, Janki Bhimani, Yi Yao, Cho-Hsien Lin, Jiayin Wang, Ningfang Mi, and Bo Sheng, "AutoAdmin: Automatic and Dynamic Resource Reservation Admission Control in Hadoop YARN Clusters", Scalable Computing: Practice and Experience, Special Issue on Advances in Emerging Wireless Communications and Networking, Volume 19, Issue 1, Pages 53-67, 2018.
21. Zhengyu Yang, Danli Jia, Stratis Ioannidis, Ningfang Mi, and Bo Sheng, "Intermediate Data Caching Optimization for Multi-Stage and Parallel Big Data Frameworks", 2018 IEEE International Conference on Cloud Computing (CLOUD 2018, Acceptance Rate: 15%), 2018.
22. Janki Bhimani, Ningfang Mi, Zhengyu Yang, Jingpei Yang, Rajinikanth Pandurangan, Changho Choi and Vijay Balakrishnan, "FIOS: Feature Based I/O Stream Identification for Improving Endurance of Multi-Stream SSDs", 2018 IEEE International Conference on Cloud Computing (CLOUD 2018, Acceptance Rate: 15%, Best Paper Award), 2018.
23. Zhengyu Yang, Yi Yao, Han Gao, Jiayin Wang, Ningfang Mi, and Bo Sheng, "New YARN Non-Exclusive Resource Management Scheme Through Opportunistic Idle Resource Assignment", IEEE Transactions on Cloud Computing, DOI: 10.1109/TCC.2018.2867580, 2018.
24. Zhengyu Yang, Manu Awasthi, Mrinmoy Ghosh, Janki Bhimani, and Ningfang Mi, "I/O Workload Management for All-Flash Datacenter Storage Systems Based on Total Cost of Ownership", IEEE Transactions on Big Data, DOI: 10.1109/TBDATA.2018.2871114, 2018.
25. Janki Bhimani, Ningfang Mi, Miriam Leeser, and Zhengyu Yang, "New Performance Modeling Methods for Parallel Data Processing Applications", Transactions on Modeling and Computer Simulation (ACM TOMACS), 2019.
26. Venkatraman Balasubramanian, Moayad Aloqailyz, Olufogorehan Tunde-Onadeley, Zhengyu Yang, and Martin Reisslein, "Reinforcing Cloud Environments via Index Policy for Bursty Workloads", IEEE/IFIP Network Operations and Management Symposium (NOMS), 2020.

PATENTS

01. Zhengyu Yang, Mrinmoy Ghosh, Manu Awasthi, Vijay Balakrishnan, "Online Flash Resource Allocation Manager Based on TCO Model", US15/092156, US20170046089A1, 2015.
02. Zhengyu Yang, Mrinmoy Ghosh, Manu Awasthi, Vijay Balakrishnan, "Online Flash Resource Migration, Allocation, Retire and Replacement Manager Based on a Cost of Ownership Model", US15/094971, US20170046098A1, 2015.
03. Zhengyu Yang, Sina Hassani, and Manu Awasthi, "Memory Device Having a Translation Layer with Multiple Associative Sectors", US9898200B2, US15/093682, US2017/0242583A1, 2015.
04. Zhengyu Yang and Manu Awasthi, "I/O Workload Scheduling Manager for RAID/non-RAID Flash Based Storage Systems for TCO and WAF Optimizations", US15/396186, 2017.
05. Zhengyu Yang, David Evans, and Jiayin Wang, "Adaptive Caching Replacement Manager with Dynamic Updating Granularities and Partitions for Shared Flash-Based Storage System", US15/400835, US20180067961A1, 2017.
06. Zhengyu Yang, Jiayin Wang, and David Evans, "Duplicate In-memory Shared-intermediate Data Detection and Reuse Module in Spark Framework", US15/404100, US20180067861A1, 2017.
07. Zhengyu Yang, Jiayin Wang, and David Evans, "In-memory Shared Data Reuse Replacement and Caching", US15/404121, US20180067869A1, 2017.
08. Zhengyu Yang, Jiayin Wang, and David Evans, "Automatic Data Replica Manager in Distributed Caching and Data Processing Systems", US15/408328, US20180069944A1, 2017.
09. Jiayin Wang, Zhengyu Yang, and David Evans, "Efficient Data Caching Management in Scalable Multi-stage Data Processing Systems", US15/423384, US20180067857A1, 2017.
10. Zhengyu Yang, and David Evans, "Dynamic Cache Partition Manager in Heterogeneous Virtualization Cloud Cache Environment", US20190163636, 2017.
11. Rory Bolt, Clay Mayers, David Evans, Allen Andrews, and Zhengyu Yang, "Automatic Data Placement Manager in Multi-Tier All-Flash Datacenter", US62/534647, 2017.

12. Zhengyu Yang, Morteza Hoseinzadeh, Ping Wong, David Evans, and John Artoux, "Hybrid Framework of NVMe-based Storage System in Cloud Computing Environment", US62/540555, 2017.
13. Zhengyu Yang, Morteza Hoseinzadeh, David Evans, Clay Mayers, and Rory Bolt, "Method and Apparatus for Adaptive Cache Load Balancing for SSD-based Cloud Computing Storage System", US20190196969, 2017.
14. Morteza Hoseinzadeh, Zhengyu Yang, Ping Wong, and David Evans, "Datacenter Storage Evaluation Framework based on Simulation", US20190188023A1, 2017.
15. Zhengyu Yang, Nithya Ramakrishnan, Allen Andrews, Sudheendra Sampath, David Thomas Evans, and Clay Mayers, "Intelligent Path Selection and Load Balancing", Patent Pending, 2019.
16. Zhengyu Yang, Nithya Ramakrishnan, Allen Andrews, Sudheendra Sampath, David Thomas Evans, and Clay Mayers, "Systems and Methods for I/O Transmissions in Queue Pair-based NVMeoF Initiator-target System", Patent Pending, 2019.
17. Zhengyu Yang, Thomas Rainey, Michael Gehlen, Ping Wong, Venkatraman Balasubramanian, Olufogorehan Tunde-onadele, Nithya Ramakrishnan, David Evans, and Clay Mayers, "Methods and Systems for Adapting Multiple Key-value Stores Cross-reference to Related Applications", Patent Pending, 2019.
18. Zhengyu Yang, Venkatraman Balasubramanian, Olufogorehan Tunde-onadele, Ping Wong, Nithya Ramakrishnan, David Evans, and Clay Mayers, "Systems and Methods for Spike Detection and Load Balancing Resource Management", Patent Pending, 2019.

TPC MEMBERSHIPS

01. 1st International Workshop on Communication, Computing, and Networking in Cyber Physical Systems (IPCCC CCNCPS 2016).
02. 2nd International Workshop on Communication, Computing, and Networking in Cyber Physical Systems (ICDCS CCNCPS 2017).
03. 1st IEEE International Conference on Multimedia Information Processing and Retrieval (MIPR 2018).

PEER REVIEW SERVICE

01. IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID 2013).
02. Journal of Simulation Modelling Practice and Theory (SIMPAT).
03. IEEE International Conference on Cloud Computing (CLOUD 2014).
04. IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID 2014).
05. International Conference on High Performance Computing and Simulation (HPCS 2014).
06. International Conference on Smart Cities and Green ICT Systems (SMARTGREENS 2015).
07. IEEE International Conference on Cloud Computing (CLOUD 2015).
08. International Conference on Massive Storage Systems and Technology (MSST 2015).
09. IEEE Transactions on Cloud Computing (TCC).
10. IEEE International Conference on Green Computing and Communications (GREENCOM 2015).
11. IEEE International Symposium on Software Defined Systems (SDS 2016).
12. International Conference on Ambient Systems, Networks and Technologies (ANT 2016).
13. International Conference on Massive Storage Systems and Technology (MSST 2016).
14. IEEE International Conference on Cloud Computing (CLOUD 2016).
15. IEEE International Conference on Networking, Architecture and Storage (NAS 2016).
16. European Association for Signal Processing Journal on Embedded Systems (EURASIPES).
17. World Conference on Information Security Applications (WISA 2016).
18. IEEE International Conference on Parallel and Distributed Systems (ICPADS 2016).
19. Workshop on Interactions of NVM/Flash with Operating Systems and Workloads (INFLOW 2016).
20. International Workshop on Communication, Computing, and Networking in Cyber Physical Systems (CCN-CPS 2016).
21. International Conference on Performance Engineering (ICPE 2017).
22. IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 2017).

23. ACM Great Lakes Symposium on VLSI (GLSVLSI 2017).
24. International Conference on Ambient Systems, Networks and Technologies (ANT 2017).
25. International Workshop on Big Data and Cloud Performance (DCPERF 2017).
26. IEEE International Workshop on Communication, Computing, and Networking in Cyber Physical Systems (CCN-CPS 2017).
27. IEEE International Conference on Cloud Computing (CLOUD 2017).
28. IEEE International Conference on Autonomic Computing (ICAC 2017).
29. International Workshop on Security (IWSEC 2017).
30. IEEE Computer Society Annual Symposium on VLSI (ISVLSI 2017).
31. Journal of Multimedia Tools and Applications (MTAP).
32. Journal of Cluster Computing (CLUS).
33. International Conference on Intelligent Computing (ICIC 2017).
34. E-journal of Translational Medicine (TM).
35. Journal of Scalable Computing: Practice and Experience (SCPE).
36. Journal of Renewable Energy (RE).
37. International Journal of Photoenergy (IJP).
38. International Journal of Computing Science and Mathematics (IJCSM).
39. Journal of Renewable and Sustainable Energy Reviews (RSER).
40. Journal of Interdisciplinary Sciences: Computational Life Sciences (ISCLS).
41. IEEE International Conference on Wireless and Mobile Computing, Networking and Communications (WIMOB 2017).
42. International Journal of Distributed Sensor Networks (IJDSN).
43. Journal of Scalable Computing and Communications (JSCC).
44. International Conference on Wireless Networks and Mobile Communications (WINCOM 2017).
45. Journal of Neural Computing and Applications (NCAA).
46. IEEE International Conference on E-health Networking, Application and Services (HEALTHCOM 2017).
47. Journal of Medical Internet Research Cancer (JMIR CANCER).
48. Advances in Science, Technology and Engineering Systems Journal (ASTESJ).
49. World Wide Web Journal (WWWJ).
50. International Conference on Knowledge Engineering and Semantic Web (KESW 2017).
51. Journal of Veterinary Sciences (VS).
52. International Journal of Advanced Robotic Systems (IJARS).
53. IEEE International Conference on Parallel and Distributed Systems (ICPADS 2017).
54. International Journal of Computers and Applications (IJCA).
55. EURASIP Journal on Wireless Communications and Networking (EURASIPJWCN).
56. The 3rd International Conference on Cloud Computing Technologies and Applications (CLOUDTECH 2017).
57. International Journal of Services Operations and Informatics (IJSOI).
58. Future Generation Computer Systems (FGCS).
59. IEEE International Conference on Tools with Artificial Intelligence (ICTAI 2017).
60. IEEE Consumer Communications and Networking Conference (CCNC 2018).
61. International Journal of Internet Manufacturing and Services (IJIMS).
62. IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI 2017).
63. Electronic Commerce Research (ECR).
64. Wireless Communications and Mobile Computing (WCMC).
65. IEEE International Conference on Big Data (BIGDATA 2017).
66. IEEE 36th International Performance Computing and Communications Conference (IPCCC 2017).
67. IEEE Access Journal (IA).
68. International Journal of Advanced Computer Science and Applications (IJACSA).

69. IEEE 36th International Performance Computing and Communications Conference, Special Session on Networking in Cyber Physical Systems (IPCCC CPS 2017).
70. Journal of Medicines (JM).
71. International Journal of Electronics and Communications (IJEC).
72. IEEE 1th International Conference on Multimedia Information Processing and Retrieval (MIPR 2018).
73. IEEE ICC 2018 Green Communications Systems and Networks Symposium (ICC 2018 GCSN).
74. Journal of Computing and Information Technology (CIT).