

Rong Shi

rongshi@fb.com | <https://vdr007.github.io>

RESEARCH INTERESTS	Parallel computing, heterogeneous architectures, large-scale distributed systems, fault tolerance, scalability
EDUCATION	The Ohio State University Columbus, OH Ph.D. Computer Science and Engineering, 09/2011 – 08/2018 University of Electronic Science and Technology of China Chengdu, China M.S. Computer Science and Engineering, 06/2011 B.S. Computer Science and Engineering, 06/2007
EXPERIENCE	Research Scientist , Facebook Inc., Menlo Park, CA 10/2018 – present <ul style="list-style-type: none">• instantDyNO (iDyNO) Design and build iDyNO system to figure out how to stress test and measure the capacity of front-end servers in Facebook’s datacenters. The work helps company better utilize existing capacity, and plan for the future demand in an efficient way.• Namespace Manager (NSMan) Design and build NSMan to manage namespaces (units to organize data and manage resources) for the Facebook data warehouse. The system orchestrates exabyte scale data movement of namespaces (lifecycle, move and disaster recovery) that facilitates thousands of Facebook engineers to run their batch, streaming and AI workloads with high efficiency and availability. Software Engineering Intern , Google Inc., Mountain View, CA 05/2017 – 08/2017 <ul style="list-style-type: none">• Location Extension Database (LEDB) Designed and implemented LEDB that aggregates data from Google My Business, F1 Adwords and other sources, and serve multiple users and pipelines. Flume pipelines executed daily to generate snapshots and refresh LEDB dashboard metrics. The LEDB was launched in production. Graduate Research Assistant , Advisor: Yang Wang 01/2015 – 08/2018 <ul style="list-style-type: none">• Evaluating System Scalability Bottlenecks by Workload Extrapolation Designed and developed PatternMiner tool to identify and predict workload patterns for large-scale system. Built simulator to play extrapolated workload and emulate large-scale cluster with a few machines. For evaluation, applied approach to HDFS NameNode and YARN Resource Manager facilitated identifying performance bugs and bottlenecks.• Cheap and Available State Machine Replication (SMR) Designed a general approach to reduce the cost of asynchronous SMR protocols while maintaining their high availability. Applied ideas to Paxos and built ThriftyPaxos from scratch in Java. Built a remoteHashMap benchmark, and exported the ThriftyPaxos to replicate database system H2. ThriftyPaxos achieved higher throughput and similar availability, yet with fewer replicas. Graduate Research Assistant , Advisor: D.K. Panda 09/2012 – 11/2014 <ul style="list-style-type: none">• Message Passing Library MVAPICH2 (0.4+ million downloads, 2,875 organizations used) Involved in the design, development, testing, release and maintenance of MVAPICH2 software stacks. Participated in the design and implementation of MPI runtime (MVAPICH2-GDR) for GPU clusters. Proposed efficient data movement approach for GPU clusters using techniques like GPUDirect RDMA, pipelining and Fastcopy.

- **Hybrid High Performance Linpack (HPL) benchmark over Heterogeneous Clusters**
Designed and implemented the hybrid HPL benchmark (the yardstick to rank the Top500 super-computers) with C and CUDA using two-level adaptive workload scheduling and communication-aware process grid reordering to gain performance on heterogeneous clusters.

COMPUTER SKILLS

- Programming Languages: Java, C, Python, C++, SQL, Shell scripting
- Big Data Processing Frameworks: Hadoop, Spark
- Parallel Programming Models and Libraries: MPI, OpenMP, CUDA
- Database Management System: MySQL, PostgreSQL, Hive

HONORS AND AWARDS

- Student Travel Grants: SOSP 2017, OSDI 2016, SOSP 2015, Cluster 2013
- Mark Baker Memorial Best Student Paper Award: Cluster 2013

PUBLICATIONS

Panpan Jin, Jian Guo, Yikai Xiao, Rong Shi, Yipei Niu, Fangming Liu, Chen Qian, Yang Wang, PostMan: Rapidly Mitigating Bursty Traffic by Offloading Packet Processing, USENIX ATC 2019.

Rong Shi, Yifan Gan, Yang Wang, Evaluating Scalability Bottlenecks by Workload Extrapolation, IEEE MASCOTS 2018.

Rong Shi, Yang Wang, Cheap and Available State Machine Replication, USENIX ATC 2016.

J. Zhang, X. Lu, J. Jose, M. Li, R. Shi and D. K. Panda, High Performance MPI Library over SR-IOV Enabled InfiniBand Clusters, Conference on High Performance Computing (HiPC'14), Goa, India, 2014

Rong Shi, Sreeram Potluri, Khaled Hamidouche, Mingzhe Li, Davide Rossetti and D. K. Panda, Designing Efficient Small Message Transfer Mechanism for Inter-node MPI Communication on InfiniBand GPU Clusters, Conference on High Performance Computing (HiPC'14), Goa, India, 2014.

Rong Shi, Xiaoyi Lu, Sreeram Potluri, Khaled Hamidouche, Jie Zhang, and D. K. Panda, HAND: A Hybrid Approach to Accelerate Non-contiguous Data Movement using MPI Datatypes on GPU Clusters, International Conference on Parallel Processing (ICPP'14), Minneapolis, USA, 2014.

J. Zhang, X. Lu, J. Jose, R. Shi and D. K. Panda, Can Inter-VM Shmem Benefit MPI Applications on SR-IOV based Virtualized InfiniBand Clusters, Euro-Par 2014 Parallel Processing - 20th International Conference, Porto, Portugal, August 2014

Rong Shi, Sreeram Potluri, Khaled Hamidouche, Xiaoyi Lu, Karen Tomko and D. K. Panda, A Scalable and Portable Approach to Accelerate Hybrid HPL on Heterogeneous CPU-GPU Clusters. Cluster'13, Indianapolis, USA, 2013.

PROFESSIONAL EXPERIENCE

- Editor/program committee member of international journals/conferences: Cluster Computing, AIMS 2020, ICPP 2019, AIMS 2019, eScience 2019, DATA ANALYTICS 2019, BPOD 2019
- External reviewer of international journals/conferences: FGCS 2020, FGCS 2019, JPDC 2019, DISC 2019, EuroPar 2019, TPDS 2018, TKDE 2018, FGCS 2018, JPDC 2018, HiPC 2018, HiPC 2017, IPDPS 2014, Cluster 2014
- ACM Student Member

TEACHING EXPERIENCE

Teaching Assistant at the Ohio State University, 09/2011 – 05/2014

- Instructor for CSE 2451: Advanced C Programming, Spring 2014.
- Lab instructor for CSE 2111: Spreadsheets and Databases, Autumn 2012.