EXACOMM: SIXTH INTERNATIONAL WORKSHOP ON COMMUNICATION ARCHITECTURES FOR HPC, BIG DATA, DEEP LEARNING AND CLOUDS AT EXTREME SCALE

A Workshop Proposal for Int'l Supercomputing Conference (ISC), 2020
Dhabaleswar K. (DK) Panda and Hari Subramoni (The Ohio State University)
{***panda, subramon***}***@cse.ohio-state.edu***

Abstract

Extreme Scale computing in HPC, Big Data, Deep Learning and Clouds are marked by multiple-levels of hierarchy and heterogeneity ranging from the compute units (many-core CPUs, GPUs, APUs etc) to storage devices (NVMe, NVMe over Fabrics etc) to the network interconnects (InfiniBand, High-Speed Ethernet, Omni-Path, Slingshot, etc). Owing to the plethora of heterogeneous communication paths with different cost models expected to be present in extreme scale systems, data movement is seen as the soul of different challenges for exascale computing. On the other hand, advances in networking technologies such as NoCs (like NVLink and Stormlake), emergence of new I/O interface architecture standards (CCIX, Gen-Z, CAPI etc), and RDMA enabled networks and the likes are constantly pushing the envelope of research in the field of novel communication and computing architectures for extreme scale computing. The goal of this workshop is to bring together researchers and software/hardware designers from academia, industry and national laboratories who are involved in creating network-based computing solutions for extreme scale architectures. The objectives of this workshop will be to share the experiences of the members of this community and to learn the opportunities and challenges in the design trends for exascale communication architectures. The scope of the workshop includes, but not limited to: scalable communication architectures and protocols, high performance networks, runtime/middleware designs, novel hardware/software co-design, high performance communication solutions for accelerator based computing, power-aware techniques and designs, performance evaluations, quality of service, and virtualization.

1 Keywords

Architecture, Big Data, Cloud, Exascale, Interconnect.

2 Workshop History

This will be the sixth edition of the workshop. The first edition of the workshop was in conjunction with International Supercomputing Conference (ISC 2015) at Messe Frankfurt, Frankfurt, Germany, Thursday, July 16th, 2015 (URL: http://nowlab.cse.ohio-state.edu/exacomm_15/). The second edition of the workshop was in conjunction with International Supercomputing Conference (ISC 2016) at Messe Frankfurt, Frankfurt, Germany, Thursday, June 23rd, 2016 (URL: http://nowlab.cse.ohio-state.edu/exacomm_16/). The third edition of the workshop was in conjunction with International Supercomputing Conference (ISC 2017) at Messe Frankfurt, Frankfurt, Germany, Thursday, June 22nd, 2017 (URL: http://nowlab.cse.ohio-state.edu/exacomm_17/). The fourth edition of the workshop was in conjunction with International Supercomputing Conference (ISC 2018) at Messe Frankfurt, Frankfurt, Germany, Thursday, June 28th, 2018 (URL: http://nowlab.cse.ohio-state.edu/exacomm_18/). The fifth edition of the was in conjunction with International Supercomputing Conference (ISC 2019) at Messe Frankfurt, Frankfurt, Germany, Thursday, June 20th, 2019 (URL: http://nowlab.cse.ohio-state.edu/exacomm/).

All ExaComm events had had one keynote talk, multiple invited talks and research papers as well as a panel. The selection of the contributed papers were decided by a peer-review process consisting a set of international program committee members. The number of people attending ExaComm over the years have steadily increased. The previous edition of the ExaComm workshop was by 35-40 people. Thus we expect 45-55 people to attend ExaComm in 2019. We plan to continue offering this workshop in future years. More details about past offerings of ExaComm (including full program, accepted papers, talks, panels, photos and details of the program committee) are available from the workshop website mentioned above.

3 Format of the Workshop

The workshop will include one Keynote talk, several invited talks, a panel session, and four to six contributed papers. The selection of the contributed papers will be decided by a peer-review process consisting a set of international program committee members.

3.1 Length of the Workshop

This workshop is targeted for full-day.

4 Extra Details for Pre-submission Deadline

These are some of the extra details required for the pre-submission deadline

4.1 Advertising the Workshop

We plan to use multiple ways to advertise and attract attendees:

- Mailing lists: We plan to use the standard mailing lists such as hpc-announce, hpc-india and tccc-announce. In addition, we plan to advertise about this workshop in other public mailing lists such as MVAPICH, MPICH, etc.
- Website: The ExaComm workshop will have a web-page. This will be linked to the various project pages of the organizers. A tentative ExaComm'19 web site has been prepared and is available at the following URL http://nowlab.cse.ohio-state.edu/exacomm/.
- Tutorials and Talks: The organizers regularly deliver talks and tutorials at multiple internal conferences and events. The ExaComm workshop will be advertised through these tutorials and talks.

4.2 Planned Schedule for Call for Papers and Review Process

Table 1: Proposed schedule

Event	Date
Initial dissemination of call for papers	January 1st, 2019
Abstract submission deadline (Optional)	11:59 PM AoE, March 15, 2019
Technical paper submission deadline	11:59 PM AoE, March 29, 2019
Paper review deadline	11:59 PM AoE, April 29, 2019
Author notification	11:59 PM AoE, May 1, 2019
Camera-ready deadline	11:59 PM AoE, May 23, 2019
Workshop	June 20, 2019

4.3 Tentative Program Committee

The following is a tentative program committee for ExaComm'20.

- Taisuke Boku, University of Tsukuba, Japan
- Ron Brightwell, Sandia National Laboratories
- Hans Eberle, NVIDIA
- Jesus Escudero-Sahuquillo, University of Castilla-La Mancha
- Pedro Javier Garcia, University of Castilla-La Mancha
- Ada Gavrilovska, Georgia Tech
- Paolo Giaccone, Politecnico di Torino
- Brice Goglin, INRIA, France
- Dror Goldenberg, Mellanox Technologies
- R. Govindarajan, Indian Institute of Science, Bangalore, India
- Ryan Grant, Sandia National Laboratory, USA
- Rui Hou, Institute of Computing Technology, Chinese Academy of Sciences, China
- Hai Jin, Huazhong University of Science and Technology, Wuhan, China
- Sven Karlsson, Technical University of Denmark, Denmark
- Yutong Lu, National University of Defense Technology, Changsha, Hunan Province, China
- Tadahiro (Takeshi) Nanri, University of Kyushu, Japan
- Dimitrios Nikolopoulos, Queen's University Belfast, UK
- Antonio Pena, Barcelona Supercomputing Center, Spain
- Sebastien Rumley, Columbia University
- Smruti Sarangi, IIT Delhi, India
- Martin Schulz, Lawrence Livermore National Laboratory
- John M. Shalf, National Energy Research Scientific Computing Center / Lawrence Berkeley National Laboratory
- Sayantan Sur, Intel
- Pedro Trancoso, University of Cyprus, Cyprus
- Sathish Vadhiyar, Indian Institute of Science, India
- Xin Yuan, Florida State University
- Jidong Zhai, Tsinghua University, China

4.4 Expected Number of Papers

We expect between six to nine contributed papers for the workshop.

4.5 Potential List of Keynote and Invited Speakers

Stephen W. Poole is the Chief Architect, ALDSC/HPC-DO. Chief Architect for all future systems at Los Alamos National Laboratory. **E-mail:***swpoole@lanl.gov*

José Duato is a Full Professor in the Department of Computer Engineering (DISCA) at the Polytechnic University of Valencia. His research is based mainly in the field of interconnection networks. He was awarded the National Prize for Research in 2009 and the King Jaime I Prize for new technologies in 2006. He is the first author of a 500-page book published in the USA, which has become the most popular book on interconnection networks. He has authored or co-authored more than 340 publications, including book chapters and papers in journals and conference proceedings. He advised or co-advised 22 PhD students. **E-mail:***jduato@gap.upv.es*

Satoshi Matsuoka is the director for RIKEN Center for Computational Science in Japan. Satoshi Matsuoka had been a Full Professor at the Global Scientific Information and Computing Center (GSIC), the Tokyo Institute of Technology since 2000, and the director of the joint AIST-Tokyo Tech. Real World Big Data Computing Open Innovation Laboratory (RWBC-OIL) since 2017, and will become a Specially Appointed Professor at Tokyo Tech starting 2018 along with his directorship at R-CCS. He received his Ph. D. from the University of Tokyo in 1993. He is a Fellow of the ACM and European ISC, and has won many awards, including the JSPS Prize from the Japan Society for Promotion of Science in 2006, presented by his Highness Prince Akishino; the ACM Gordon Bell Prize in 2011; the Commendation for Science and Technology by the Minister of Education, Culture, Sports, Science and Technology in 2012; the 2014 IEEE-CS Sidney Fernbach Memorial Award, the highest prestige in the field of HPC; and recently HPDC 2018 Achievement Award from ACM.

Mark Seager is the Chief Technology Officer for the HPC Ecosystem at the Intel Corporation. Dr. Seager received his B.S. Degree in Mathematics and Astrophysics at the University of New Mexico at Albuquerque in 1979 and his PhD in Numerical Analysis from the University of Texas at Austin in 1984. Dr. Seager is working on an ecosystem approach to develop and build HPC systems for Exascale with technology based on the broader high volume HPC systems. Dr. Seager has won numerous awards including the prestigious Edward Teller Award for "Major Contributions to the State-of-the-Art in High Performance Computing." Prior to joining Intel, Dr. Seager managed the platforms portion of the Advanced Simulation and Computing (ASC) program at LLNL. **E-mail:***mark.seager@intel.com*

4.6 Tentative Panel Details

The tentative title for the panel, the potential list of panelists and, the potential panel moderator is given below.

Tentative Title: Challenges in Co-Designing Runtime and Network Architecture for Extreme-Scale

Potential Panel Moderator: Jeff Vetter, ORNL, USA

Potential Panelists

- Richard Graham, Senior Solutions Architect, Mellanox Technologies
- Duncan Roweth, Senior Principal Engineer, Office of the CTO, Cray
- Geoffrey C. Fox, Indiana University Bloomington
- John Shalf, Lawrence Berkeley National Laboratory (LBNL), USA
- Andrew A. Chien, The University of Chicago, Argonne National Laboratory

4.7 Targeted Proceedings Venue

We would like to take up the kind offer of the workshop organizers and publish the papers, in the separate Joint Workshop Proceeding Volume, published online with Springer Publishing like the ISC 2019 research papers proceedings.

5 Targeted Audience

This workshop is targeted for various categories of people (Scientists, engineers, researchers, developers and students) working in the area of high performance communication and I/O, Big Data, Deep Learning, Clouds, networking, middleware, virtualization, quality of service, accelerators and applications related to exascale computing. Specific audience this workshop is aimed at include:

1. Scientists, engineers, researchers and students working on the design and development of communication architectures for next-generation exascale systems including clusters, data centers, storage centers, deep learning, cloud computing and Big Data systems.

2. Newcomers to the field of HPC and exascale computing who are interested in familiarizing themselves with programming models, accelerators, networking, and communication architectures.
3. Developers of next generation networked computing architectures and middleware.

6 Estimated Attendance

We estimate around 40-50 people to attend the workshop.

7 Expected Outcome from the Workshop

This workshop will help ISC 2020 attendees learn multiple aspects of modern high performance networking technologies. For scientists, engineers, and researchers from both academia and industry engaged in designing scalable communication architectures and protocols over high performance networks and their associate runtimes and middlewares, the workshop will provide a forum for them to discuss and share their experience, expertise and expectation about the trends on high performance programming networks and middleware for exascale systems.

8 Brief Bio of the Organizers

Dhabaleswar K. (DK) Panda is a Professor of Computer Science and Engineering at the Ohio State University. His research interests include parallel computer architecture, high performance networking, InfiniBand, Exascale computing, Big Data, programming models, GPUs and accelerators, high performance file systems and storage, virtualization, deep learning and cloud computing. He has published over 500 papers in major journals and international conferences related to these research areas. Dr. Panda and his research group members have been doing extensive research on modern networking technologies including InfiniBand, High-Speed Ethernet and RDMA over Converged Enhanced Ethernet (RoCE). The MVAPICH2 (High Performance MPI over InfiniBand, Omni-Path, iWARP and RoCE) libraries (https://mvapich.cse.ohio-state.edu), developed by his research group (http://mvapich.cse.ohio-state.edu), are currently being used by more than 3,050 organizations worldwide (in 86 countries). Recently, Dr. Panda and his team have also developed high performance RDMA-enabled Apache Spark, Hadoop and Memcached software packages (http://hibd.cse.ohio-state.edu) to accelerate Big Data applications. Similarly, his team has also developed high performance deep learning library (OSU-Caffe), available from http://hidl.cse.ohio-state.edu. He is an IEEE Fellow and a member of ACM. More details about Prof. Panda are available at http://www.cse.ohio-state.edu/∼panda.

Hari Subramoni received the Ph.D. degree in Computer Science from The Ohio State University, Columbus, OH, in 2013. He is a Research Scientist in the Department of Computer Science and Engineering at the Ohio State University, USA, since August 2015. He has published over 70 papers in international journals and conferences related to these research areas. He has been actively involved in various professional activities in academic journals and conferences. Recently, Dr. Subramoni is doing research and working on design and development for of MVAPICH2 (High Performance MPI over InfiniBand, iWARP and RoCE) and MVAPICH2-X (Hybrid MPI and PGAS (OpenSHMEM, UPC and, UPC++)) software packages. He is a member of IEEE. More details about Dr. Subramoni are available at http://www.cse.ohio-state.edu/∼subramon.