HiPINEB 2017

The 3rd International Workshop on High-Performance Interconnection Networks in the Exascale and Big-Data Era

Pedro Javier Garcia
Universidad de Castilla-La Mancha (UCLM)
SPAIN

Jesus Escudero-Sahuquillo
Universidad de Castilla-La Mancha (UCLM)
SPAIN

Style Powered by: ACEDA

February 5, Austin, Texas, USA

Motivation

- Growing importance of **Interconnection Networks** in HPC and Datacenters
- **HiPINEB Series** want to gather and discuss in a full-day event the latest and most prominent efforts and advances, both from industry and academia, in the design and development of scalable high-performance interconnects of Exascale and Big-Data systems
- Website: *http://hipineb.i3a.info*

Motivation

Topics of interest

- Interconnect architectures and network technologies for high-speed, low-latency interconnects.
- Scalable network topologies, suitable for interconnecting a huge number of nodes.
- Power saving policies in the interconnect devices and network infrastructure, both at software and hardware level.
- Good practices in the configuration of the network control software.
- Network communication protocols: MPI, RDMA, MapReduce, etc.
- APIs and support for programming models.
- Routing algorithms.
- Quality of Service (QoS).
- Reliability and Fault tolerance.
- Load balancing and traffic scheduling.
- Network Virtualization.
- Congestion Management.
- Applications and Traffic characterization.
- Modeling and simulation tools.
- Performance Evaluation.
- Interfacing accelerators through the interconnect (GPUs, Xeon Phi, etc).
- Network infrastructure in distributed storage, distributed databases and Big-Data.

HiPINEB Series

List of Activities

- **HiPINEB workshop** 2 editions, Chicago (2015) and Barcelona (2016)
- **Journal Special Issue**: Journal of Supercomputing (2016) and Concurrency & Computation: Practice and Experience (2017)
- **Journal Special Issue**: Concurrency & Computation: Practice and Experience (2018)
- **Summer School**: June 2017.
  Albacete, Spain

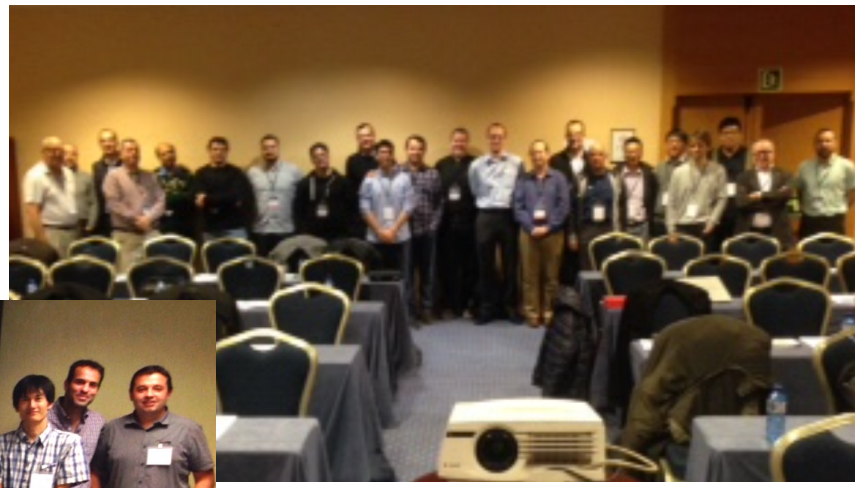

Organization

Technical Program Committee

Francisco J. Alfaro, University of Castilla-La Mancha, Spain
Jose Cano-Reyes, University of Edinburgh, United Kingdom
Lizhong Chen, Oregon State University, USA
Nikolaos Chrysos, FORTH, Greece
Holger Fröning, University of Heidelberg, Germany
Maria Engracia Gomez, Technical University of Valencia, Spain
Ernst Gunnar Gran, Simula Research Laboratory, Norway
Ryan E. Grant, Sandia National Laboratories, USA
Mitch Gusat, IBM Research, Switzerland
Scott Hemmert, Sandia National Laboratories, USA
John Kim, KAIST, South Korea
Michihiro Koibuchi, National Institute of Informatics, Japan
Yuho Jin, New Mexico State University, USA
Pedro Lopez, Technical University of Valencia, Spain
Jose Miguel Montañana, University of York, United Kingdom
Gaspar Mora, Intel Corporation, USA
Mondrian Nuessle, Extoll, Germany
Julio Ortega, University of Granada, Spain
Thibaut Palfer-Sollier, Numascale AS, Norway
Dhabaleswar K. Panda, The Ohio State University, USA
Matthieu Perotin, ATOS BULL, France
Mikel Eukeni Pozo Astigarraga, CERN, Switzerland
Samuel Rodrigo, Oracle Corporation, Norway
Sebastien Rumley, Columbia University, USA
Jose Luis Sanchez, University of Castilla-La Mancha, Spain
Heiko Joerg Schick, Huawei Technologies, Germany
Jörn Schumacher, CERN, Switzerland
Alex Shpiner, Mellanox Technologies, Israel
Evangelos Tasoulas, Simula Research Laboratory, Norway
Francisco Triviño, Oracle Corporation, Norway
Luis Tomas, Red Hat, Spain
Enrique Vallejo, University of Cantabria, Spain
Wainer Vandelli, CERN, Switzerland
Pierre Vigneras, ATOS BULL, France

Organization

Steering Committee

- **Jose Duato**, Technical University of Valencia, Spain
- **Francisco J Quiles**, University of Castilla-La Mancha, Spain
- **Torsten Hoefler**, ETH Zurich, Switzerland
- **Timothy M Pinkston**, University of Southern California, USA
- **Eitan Zahavi**, Mellanox, Israel

Program

Overview

Time	Activity
8:40– 10:00am	**Keynote**
10:00-10:30am	**Coffee Break**
10:30-12:00	**Panel**
12:00-1:30pm	**Lunch**
1:30-3:00pm	**Technical Session 1**
3:00-3:30pm	**Coffee Break**
3:30-4:55pm	**Technical Session 2**
4:55-5:00pm	**Closing**

Program

Keynote

Issues in the Design of an Exascale Network

Speaker: Bill Dally, Chief Scientist and SVP of Research in NVIDIA, and Stanford Professor

Chair: Francisco J. Quiles, UCLM, Spain

Program

Panel (10:30 – 12:00)

Massive-storage Networks vs Intensive-computing Networks

Moderator: John Kim, HP Labs / KAIST, South Korea

Torsten Hoefler
ETH Zurich
Switzerland

Bill Dally
nVIDIA and Stanford
USA

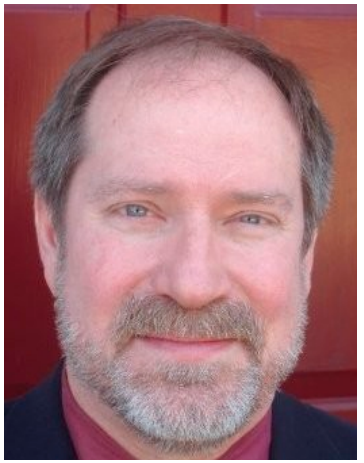

David Mayhew
San Diego University
USA

Program

Technical Session 1 (1:30 – 3:00pm)

Chairman: Michihiro Koibuchi, National Institute of Informatics, Japan

- **Dragonfly+: Low Cost Topology for Scaling Data Centers**
  Alexander Shpiner, Zachy Haramaty, Saar Eliad, Vladimir Zdornov, Barak Gafni and Eitan Zahavi (Mellanox Technologies, Israel)
- **A case study on implementing virtual 5D torus networks using network components of lower dimensionality**
  Francisco Andujar-Muñoz, Juan A. Villar, Jose L. Sanchez, Francisco Alfaro and Holger Fröning (University of Castilla-La Mancha, Spain, and Ruprecht-Karls University of Heidelberg, Germany)
- **New link arrangements for Dragonfly networks**
  Madison Belka, Myra Doubet, Sofia Meyers, Rosemary Momoh, David Rincon-Cruz and David Bunde (Knox College, and Columbia University, USA)
- **An Effective Queuing Scheme to Provide Slim Fly topologies with HoL Blocking Reduction and Deadlock Freedom for Minimal-Path Routing**
  Pedro Yebenes Segura, Jesus Escudero-Sahuquillo, Pedro Javier Garcia, Francisco J. Quiles and Torsten Hoefler (University of Castilla-La Mancha, Spain, and ETH Zurich, Switzerland)

Program

Technical Session 2

Chairman: Jesus Escudero-Sahuquillo, UCLM, Spain

- **Early Experiences with Saving Energy in Direct Interconnection Networks**
  Felix Zahn, Steffen Lammel and Holger Fröning (Ruprecht-Karls University of Heidelberg, Germany)
- **Extending commodity OpenFlow switches for large-scale HPC deployments**
  Mariano Benito, Enrique Vallejo, Ramón Beivide and Cruz Izu (University of Cantabria, Spain, and The University of Adelaide, Australia)
- **Isolating jobs for security on high-performance fabrics**
  Matthieu Pérotin and Tom Cornebize (Atos, France, and ENS Lyon, France)
- **Knapp: A Packet Processing Framework for Manycore Accelerators**
  Junhyun Shim, Joongi Kim, Keunhong Lee and Sue Moon (SAP Labs Korea, Lablup Inc. and KAIST, South Korea)

Special Issue

Best papers among those selected for HiPINEB 2017 will be published in a Special Issue of the Wiley's journal:

Concurrency and Computation: Practise and Experience

http://hipineb.i3a.info/hipineb2017/special-issue/

- Call for Papers: March 1, 2017
- Deadline: May 30, 2017
- Decission Date: July 30, 2017

Final Remarks

- **HiPINEB'17 Proceedings:**

  http://conferences.computer.org/hipineb/2017

  User name: hipineb167

  Password: conf17

- **Web Site:** http://hipineb.i3a.info/hipineb2017
- **Twitter:** @*hipineb* (#HiPINEB2017)
- **LinkedIn Group**

HiPINEB 2017

The 3rd International Workshop on High-Performance Interconnection Networks in the Exascale and Big-Data Era

Pedro Javier Garcia
Universidad de Castilla-La Mancha (UCLM)
SPAIN

Jesus Escudero-Sahuquillo
Universidad de Castilla-La Mancha (UCLM)
SPAIN

Style Powered by: ACEDA

February 5, Austin, Texas, USA