

Previously Held: NDM 2011

The International Workshop on Network-aware Data Management held on Nov 14th, 2011, in conjunction with the International Conference for High Performance Computing, Networking, Storage and Analysis (SC 2011), in Seattle, WA, USA.

The one-day workshop on Network-aware Data Management attracted more than 80 attendees during the SC 2011 conference. The call for papers received submissions from Asia, Europe, and the United States.

The workshop program consisted of a keynote speech, presentations of peer-reviewed papers, and a panel discussion. The review process involved at least 3 reviews for each submission. We ranked the review scores and selected high-quality papers that best fit into the workshop program. Based on the review results, we accepted 8 papers. Among them, 4 were full papers with 30-minute presentations, and 4 were short papers with 20-minute presentations in the workshop program.

One Best-Paper award and one Best-Paper Honorable Mention were selected as outstanding paper contributions. The paper titled “**dFtree - A Fat-tree Routing Algorithm using Dynamic Allocation of Virtual Lanes to Alleviate Congestion in InfiniBand Networks**” presented by Wei Lin Guay (Simula Research Laboratory, Norway) received the Best Paper award, and the paper titled “Predicting Network Throughput for Grid Applications on Network Virtualization Areas” presented by Chunghan Lee (Toyohashi University of Technology, Japan) received the Best Paper Honorable Mention. We announced the awards during the closing remarks, and presented certificates to the speakers at the end of the workshop program.

The program included a panel on “**Data Management in Exa-scale Computing and Terabit Networking Era**”, and a keynote speech titled “**Accelerating Data-Driven Discovery by Outsourcing the Mundane**” by Ian Foster (Argonne National Laboratory & University of Chicago). We invited 5 panelists from different specialties who are well-known in the field of networking and data-intensive computing. Panel members include Richard Carlson (DoE Office of Advanced Scientific Computing Research), Ann Chervenak (USC Information Sciences Institute), Daniel S. Katz (University of Chicago and Argonne National Laboratory), Dhableswar Panda (Ohio State University), and Brian Tierney (ESnet/Lawrence Berkeley National Laboratory).

We posted the presentation slides from paper presentations, panel talks, and keynote speech online at <http://sdm.lbl.gov/ndm/2011>

Workshop Website: <http://sdm.lbl.gov/ndm/2011>

ACM proceedings: <http://dl.acm.org/citation.cfm?id=2110217>

ISBN: 978-1-4503-1132-8

NDM 2011 Workshop Program

- 9:00 - 9:15 Opening Remarks
- 9:15 - 10:15 Keynote Speech
Title: “*Accelerating Data-driven Discovery by Outsourcing the Mundane*”
Speaker: Ian Foster, Argonne National Laboratory & University of Chicago
- 10:15 – 10:30 Break
- 10:30 - 11:00 dFtree - A Fat-tree Routing Algorithm using Dynamic Allocation of Virtual Lanes to Alleviate Congestion in InfiniBand Networks
Presenter: Wei Lin Guay, Simula Research Laboratory, Norway
- 11:00 - 11:30 Predicting Network Throughput for Grid Applications on Network Virtualization Areas
Presenter: Chunghan Lee , Toyohashi University of Technology, Japan
- 11:30 - 12:00 Network-Aware End-to-End Data Throughput Optimization
Presenter: Tevfik Kosar, State University of New York at Buffalo.
- 12:00 – 13:00 Lunch Break
- 13:00 - 13:30 Network-aware Data Movement Advisor (NADMA)
Presenter: Patrick Brown, Southern Illinois University
- 13:30 - 13:50 Scientific Data Movement enabled by the DYNES Instrument
Presenter: Jason Zurawski, Internet 2
- 13:50 - 14:10 CernVM-FS: Delivering Scientific Software to Globally Distributed Computing Resources
Presenter: Jakob Blomer, CERN, Switzerland
- 14:10 - 14:30 A Peer-to-Peer Architecture for Data-Intensive Cycle Sharing
Presenter: Ian Kelley, Cardiff University, UK
- 14:30 – 14:50 An Architecture for a Data-Intensive Computer
Presenter: Edward Givelberg, Johns Hopkins University
- 14:50 – 15:15 Break
- 15:15 – 16:30 Panel Discussion
Title: Data Management in Exa-scale Computing and Terabit Networking Era
- 16:30 – 17:00 Closing Remarks

NDM 2011 Panel

Topic: **Data management in exa-scale computing and terabit networking era**

Panel members:

Richard Carlson, DoE Office of Advanced Scientific Computing Research
Predictable or Guaranteed Behavior: What do Scientists Need?

Ann Chervenak, USC Information Sciences Institute
The Impact of Terabit Networking on Data Management

Daniel S. Katz, University of Chicago and Argonne National Laboratory
Application Challenges in Terabit Networks

Dhabaleswar Panda, Ohio State University
High Performance Networks with RDMA

Brian Tierney, ESnet/Lawrence Berkeley National Laboratory
Science Trends and Network Issues

Suggested questions:

What has been the impact of recent advances on science discoveries and industry? Which of them were successful and which of them were simply a hype?

What research challenges do you expect in exa-scale data management on terabit networks? If possible, please mention the challenges that require more than re-engineering existing tools?

What applications need exa-scale computing and terabit networks? Please give a few examples. What is the scope of data management research on faster network?

What are the performance problems in next-generation high-bandwidth networks? How do you think such problems will affect the applications performance?

How do applications accommodate next generation networks? What are the challenges for middleware developers to adapt exa-scale data management on terabit networks?

What are the major challenges in network management in terms of provisioning capacity and path, performance monitoring and diagnosis tools?

Summary:

Mehmet Balman and Surendra Byna. 2011. *Open Problems in Network-aware Data Management in Exa-scale Computing and Terabit Networking Era*. In *Proceedings of the first international workshop on Network-aware data management (NDM '11)*. ACM, New York, NY, USA, 73-78. DOI=10.1145/2110217.2110229 <http://doi.acm.org/10.1145/2110217.2110229>

NDM 2011 Organization

General Chairs:

Mehmet Balman (Lawrence Berkeley National Laboratory, USA)

Surendra Byna (Lawrence Berkeley National Laboratory, USA)

Advisory Chair:

Arie Shoshani (Lawrence Berkeley National Laboratory, USA)

Program Committee:

Gail-Joon Ahn (Arizona State University, USA)

Ismail Akturk (Bilkent University, Turkey)

Pavan Balaji (Argonne National Laboratory, USA)

Iliia Baldine (Renaissance Computing Institute, USA)

Tasneem Brutch (Samsung Electronics, USA)

Ali R. Butt (Virginia Tech, USA)

Ann Chervenak (Information Sciences Institute, USA)

Jerry Chou (National TsingHua University, Taiwan)

Constantine Dovrolis (Georgia Institute of Technology, USA)

Zhihui Du (Tsinghua University, China)

Daniel S. Katz (University of Chicago & Argonne National Laboratory, USA)

Raj. Kettimuthu (Argonne National Laboratory, USA)

Tevfik Kosar (University at Buffalo-SUNY, USA)

Kamesh Madduri (Pennsylvania State University, USA)

Inder Monga (ESnet & Lawrence Berkeley National Laboratory, USA)

Dhableswar Panda (Ohio State University, USA)

Alex Sim (Lawrence Berkeley National Laboratory, USA)

Brian Tierney (ESnet & Lawrence Berkeley National Laboratory, USA)

Sudharshan Vazhkudai (Oak Ridge National Laboratory, USA)

Venkatram Vishwanath (Argonne National Laboratory, USA)

Fatos Xhafa (Universitat Politècnica de Catalunya, Spain)