

Previously Held: NDM 2013

The Third International Workshop on Network-aware Data Management held in cooperation with ACM SIGHPC, in conjunction with the IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SC 2013), in Denver, Colorado, on Sunday November 17, 2013.

The goal of the one-day workshop is to seek contributions from academia, government, and industry to discuss emerging trends and new technological developments in dynamic resource provisioning, intelligent data-flow and resource coordination, end-to-end processing of data, network-aware application design issues, and cutting-edge network performance problems.

The format of the workshop consisted of three invited talks, one contributed talk, and presentations of peer-reviewed papers. 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 9 papers. Among them, 5 were full papers with 30-minute presentations, and 3 were short papers with 20-minute presentations in the workshop program.

The paper titled “*Characterizing the Impact of End-System Affinities On the End-to-End Performance of High-Speed Flows*” received the highest review scores and received the Best Paper Award. Best Paper Award certificate has been presented to *Nathan Hanford, University of California, Davis* during the closing remarks session at the end of the workshop.

The workshop program included a contributed talk on “**Overview of ESnet's 100Gbps Tesbed**” by Brian L. Tierney, and three invited talks: “**Challenges and Solutions in Large Scale Data Movement**”, by Martin Swamy, Indiana University; “**The Changing Face of Network Projects and Funding**”, by Jennifer M. Schopf, International Networking at Indiana University; “**Supporting Climate Modeling Over Named Data Networking**” by Christos Papadopoulos, Colorado State University.

We would like to thank NDM 2013 authors for providing high-quality paper contributions. We also thank our program committee members for reviewing papers and providing feedback to the authors. Finally, we thank all our speakers for their insightful comments and presentations.

We posted the presentation slides from paper presentations and invited talks online at <http://sdm.lbl.gov/ndm/2013>

ACM proceedings: <http://dl.acm.org/citation.cfm?id=2534695>

ISBN: 978-1-4503-2522-6/13/11

The 3rd International Workshop on Network-aware Data Management

in cooperation with ACM [SIGHPC](#), in conjunction with [SC'13](#)
International Conference for High Performance Computing, Networking, Storage and Analysis.
Sunday, November 17th, 2013 (9:00am - 5:30pm)
Colorado Convention Center, Denver, Colorado



In the age of extraordinary advances in communication technologies, there is a need for efficient use of the network infrastructure to address increasing data requirements of today's applications. Traditional network and data management techniques are unlikely to scale to meet the needs of future collaborative data-intensive systems. We require novel data access mechanisms and intelligent network middleware to enable future design principles of network-aware data management. NDM workshop provides a forum for researchers from academia, government, and industry to discuss emerging trends and new technological developments in intelligent data-flow and resource coordination, network-aware application design issues, and 100Gbps network performance problems.

Invited Talks:

Challenges and Solutions in Large Scale Data Movement, *Martin Swamy, Indiana University.*

(9:05-9:40)

Abstract: The effective and efficient utilization of networks is a persistent challenge in today's computing ecosystem. The trend towards increasingly distributed computing environments supporting data-intensive applications compounds the issue, making data movement a cross-cutting concern, affecting users of distributed computing and data-intensive science alike. This talk will discuss emerging network approaches and their implications for the future. Two key areas to be explored are Software Defined Networking (SDN) and Remote Direct Memory Access (RDMA) over Ethernet.

Overview of ESnet's 100Gbps Testbed, *Brian Tierney, Energy Sciences Network.*

(9:40-10:00)

Abstract: The ESnet 100G testbed provides network researchers with a realistic environment for 100G application / middleware experiments. This talk will give an overview of the testbed, what research has been done on the testbed, and how to get access to the testbed.

The Changing Face of Network Projects and Funding, *Jennifer M. Schopf, International Networking at IU*

(13:30-14:00)

Abstract: Ten years ago, research and education groups could ask for funding for more capacity based on the simple argument that what was in place simply wasn't enough for researcher, and the argument was generally accepted. Recently, especially in the last 2-3 years, this has changed, and funding agencies are looking more closely to network use and usability as they make awards. This talk will address the changing nature of how network engineers and researchers need to think about their projects, from funding requests to demos, using the IU International Network group as a case study.

Supporting Climate Modeling Over Named Data Networking, *Christos Papadopoulos, Colorado State University.*

(15:30-16:00)

Abstract: Named Data Networking (NDN) is an instance of Information Centric Networking (ICN). Unlike IP, NDN focuses on content (the what) rather than hosts (the where). In NDN users simply request content by name rather than its location, and the network locates and retrieves signed content from anywhere. The ability to specify content by name and retrieve the nearest copy is of great benefit to Big Data applications by seamlessly locating related data and speeding up retrieval of large datasets, especially in collaborative environments. In this presentation, I will briefly describe our plans to integrate NDN with the Global Cloud Resolving Model (GCRM) application at Colorado State University's Atmospheric Sciences department.

Network-aware Data Management

Sunday, November 17th, 2013 (9:00am - 5:30pm)

Workshop Program

- | | |
|---------------|---|
| 9:00 - 9:05 | Opening Remarks |
| 9:05 - 9:40 | Challenges and Solutions in Large Scale Data Movement , Martin Swamy, Indiana University. |
| 9:40 - 10:00 | Overview of ESnet's 100Gbps Tesbed , Brian Tierney, Energy Sciences Network |
| 10:00 - 10:30 | Coffee Break |
| 10:30 - 11:00 | On causes of GridFTP Transfer Throughput Variance , Zhengyang Liu, Malathi Veeraraghavan, Jianhui Zhou, Jason Hick and Yee-Ting Li. |
| 11:00 - 11:30 | Characterizing the Impact of End-System Affinities On the End-to-End Performance of High-Speed Flows , Nathan Hanford, Vishal Ahuja, Mehmet Balman, Matthew Farrens, Dipak Ghosal, Eric Pouyoul and Brian Tierney. |
| 11:30 - 12:00 | Efficient Wide Area Data Transfer Protocols for 100 Gbps Networks and Beyond , Ezra Kissel, Martin Swamy, Brian Tierney and Eric Pouyoul. |
| 12:00 - 12:20 | In-network, Push-based Network Resource Monitoring , Taylor Groves, Yihua He and Dorian Arnold. |
| 12:20 - 13:30 | Lunch Break |
| 13:30 - 14:00 | The Changing Face of Network Projects and Funding , Jennifer M. Schopf, International Networking at IU |
| 14:00 - 14:30 | Evaluating I/O Aware Network Management for Scientific Workflows on Networked Clouds , Anirban Mandal, Paul Ruth, Ilya Baldin, Yufeng Xin, Claris Castillo, Mats Rynge and Ewa Deelman. |
| 14:30 - 15:00 | Network-Aware Data Caching and Prefetching for Cloud-hosted Metadata Retrieval , Bing Zhang, Brandon Ross, Sanatkumar Tripathi, Sonali Batra, and Tevfik Kosar. |
| 15:00 - 15:30 | Coffee Break |
| 15:30 - 16:00 | Supporting Climate Modeling Over Named Data Networking , Christos Papadopoulos, Colorado State University. |
| 16:00 - 16:20 | The Practical Obstacles of Data Transfer: Why researchers still love scp , Hai Ah Nam, Jason Hill and Suzanne Parete-Koon. |
| 16:20 - 16:40 | End-to-End Data Movement Using MPI-IO Over Routed Terabits Infrastructures , Geoffroy Vallee, Scott Atchley, Youngjae Kim and Galen Shipman. |
| 16:40 - 17:10 | Network-aware Virtual Machine Consolidation for Large Data Centers , Dharmesh Kakadia, Nandish Kopri and Vasudeva Varma. |
| 17:10 - 17:30 | Open-Discussion & Best Paper Announcement |

NDM 2013 Organization

General Chairs:

Mehmet Balman, VMware R&D & Lawrence Berkeley National Laboratory
Surendra Byna, Lawrence Berkeley National Laboratory
Brian L. Tierney, Energy Sciences Network (ESnet) & Lawrence Berkeley National Laboratory

Program Committee:

Gabrielle Allen, Skolkovo Institute of Science and Technology, Russia
Ismail Akturk, University of Minnesota
Scott Brim, Internet 2
Tasneem Brutch, Samsung R&D
Ali R. Butt, Virginia Tech
Emre Celebi, VMware R&D
Promita Chakraborty, The Molecular Foundry
Jerry Chou, National Tsing Hua University, Taiwan
Zhihui Du, Tsinghua University, China
Dipak Ghosal, University of California, Davis
Renato Figueiredo, University of Florida
Shantenu Jha, Rutgers University
Daniel S. Katz, University of Chicago & Argonne National Laboratory
Jinoh Kim, Texas A&M University-Commerce
Ezra Kissel, Indiana University
Marlon Pierce, Indiana University
Esmay Yıldırım, Fatih University, Turkey
Murat Yuksel, University of Nevada - Reno
Wenji Wu, Fermi Lab
Chen Wu, ICRAR, University of Western Australia