

# **2009 17th IEEE Symposium on High-Performance Interconnects**

## **(HOTI 2009)**

**New York, New York, USA  
25 – 27 August 2009**



**IEEE Catalog Number: CFP09HIS-PRT**  
**ISBN: 978-1-4244-5202-6**

# Table of Contents

## 2009 17th IEEE Symposium on High Performance Interconnects

# HOTI 2009

### Message from General Chair Message from Program Co-chairs Organizing Committee Program Committee

### Keynotes

The Other Face of On-Chip Interconnect.....	xii
<i>Anant Agarwal, CSAIL, MIT</i>	
The Need for Speed in Trading Environments.....	xiii
<i>Spencer Greene, CTO, TIBCO Software</i>	
High-Speed Networking and the Race to Zero.....	xiv
<i>Andrew F. Bach, Senior Vice President, Technology, NYSE Euronext</i>	

### Tutorials

Infiniband and 10-Gigabit Ethernet for Dummies.....	xv
<i>D. K. Panda and M. Koop (The Ohio State University); P. Balaji (Argonne National Laboratory)</i>	
Designing High-End Computing Systems with Infiniband and 10-Gigabit Ethernet.....	xvi
<i>D. K. Panda and M. Koop (The Ohio State University); P. Balaji (Argonne National Laboratory)</i>	

### Panels

Remote Direct Memory Access over the Converged Enhanced Ethernet Fabric: Evaluating the Options.....	xviii
Industrial Leaders Forum.....	xix

### Technical Session 1: On-Chip Interconnects

Designing Energy-Efficient Low-Diameter On-Chip Networks with Equalized Interconnects.....	3
<i>Ajay Joshi, Byungsub Kim, and Vladimir Stojanovic</i>	
Winning with Pinning in NoC.....	13
<i>Ahmed Abousamra, Rami Melhem, and Alex Jones</i>	
Simple Fairness Protocols for Daisy Chain Interconnects.....	22
<i>Wladek Olesinski, Hans Eberle, Nils Gura, Bob Dickson, Aron Silverton, Sumti Jairath, and Peter Yakutis</i>	

## Technical Session 2: 10 Gigabit Ethernet Interconnects

Adaptive Routing in Data Center Bridges .....	33
<i>Cyriel Minkenbergh, Mitchell Gusat, and German Rodriguez</i>	
Fulcrum's FocalPoint FM4000: A Scalable, Low-Latency 10GigE Switch for High-Performance Data Centers.....	42
<i>Uri Cummings, Dan Daly, Rebecca Collins, Virat Agarwal, Fabrizio Petrini, Michael Perrone, and Davide Pasetto</i>	
Performance Measurement of an Integrated NIC Architecture with 10GbE.....	52
<i>Guangdeng Liao and Laxmi Bhuyan</i>	

## Technical Session 3: Low Latency and Scalable Systems

MPI Collective Communications on The Blue Gene/P Supercomputer: Algorithms and Optimizations.....	63
<i>Ahmad Faraj, Sameer Kumar, Brian Smith, Amith Mamidala, and John Gunnels</i>	
MiAMI: Multi-core Aware Processor Affinity for TCP/IP over Multiple Network Interfaces .....	73
<i>Hye-Churn Jang and Hyun-Wook Jin</i>	
FPGA Accelerated Low-Latency Market Data Feed Processing .....	83
<i>Gareth W. Morris, David B. Thomas, and Wayne Luk</i>	

## Technical Session 4: High-Performance Communication in Data Center

Data Center Switch Architecture in the Age of Merchant Silicon.....	93
<i>Nathan Farrington, Erik Rubow, and Amin Vahdat</i>	
Optimized Routing for Large-Scale InfiniBand Networks.....	103
<i>Torsten Hoefler, Timo Schneider, and Andrew Lumsdaine</i>	
Designing Next Generation Clusters: Evaluation of InfiniBand DDR/QDR on Intel Computing Platforms.....	112
<i>Hari Subramoni, Matthew Koop, and Dhabaleswar K. Panda</i>	

## Panel 1

Remote Direct Memory Access over the Converged Enhanced Ethernet Fabric: Evaluating the Options.....	123
<i>David Cohen, Thomas Talpey, Arkady Kanevsky, Uri Cummings, Michael Krause, Renato Recio, Diego Crupnicoff, Lloyd Dickman, and Paul Grun</i>	

## Special Session: Industrial Leader Forum

Cost-Effective Optics: Enabling the Exascale Roadmap .....	133
<i>Alan F. Benner</i>	
Optics for Enabling Future HPC Systems .....	138
<i>Gilad Shainer, Eyal Gutkind, Bill Lee, Michael Kagan, and Yevgeny Kliteynik</i>	
The Impact of Optics on HPC System Interconnects .....	143
<i>Mike Parker and Steve Scott</i>	
Interconnect Challenges in a Many Core Compute Environment.....	148
<i>Jerry R. Bautista</i>	

Optical Networking for Data Center Interconnects Across Wide Area Networks .....	149
<i>Joseph E. Berthold</i>	

### **Technical Session 5: Optical Interconnects**

160Gb/s Serial Line Rates in a Monolithic Optoelectronic Multistage Interconnection Network .....	157
<i>A. Albores-Mejia, K.A. Williams, F. Gomez-Agis, S. Zhang, H.J.S. Dorren, X.J.M. Leijtens, T. de Vries, Y.S. Oei, M.J.R. Heck, L.M. Augustin, R. Notzel, D.J. Robbins, and M.K. Smit</i>	

Low-Power Reconfigurable Network Architecture for On-Chip Photonic Interconnects.....	163
<i>Iñigo Artundo, Wim Heirman, Mikel Loperena, Christof Debaes, Jan Van Campenhout, and Hugo Thienpont</i>	

### **Author Index**