

2010 18th IEEE Symposium on High Performance Interconnects

(HOTI 2010)

**Mountain View, California, USA
18 – 20 August 2010**



**IEEE Catalog Number: CFP10HIS-PRT
ISBN: 978-1-4244-8547-5**

2010 18th IEEE Symposium on High Performance Interconnects

HOTI 2010

Table of Contents

Message from the General Chairs	viii
Message from the Technical Program Chairs	x
Steering Committee	xii
Organizing Committee	xiii
Technical Program Committee	xiv
Keynotes	xv

Session 1: Interconnects and Switching

Intel® QuickPath Interconnect—Architectural Features Supporting Scalable System Architectures	1
<i>Dimitrios Ziakas, Allen Baum, Robert A. Maddox, and Robert J. Safranek</i>	
Belief-Propagation Assisted Scheduling in Input-Queued Switches	7
<i>Shadi Atalla, Davide Cuda, Paolo Giaccone, and Marco Pretti</i>	
Impact of Inter-application Contention in Current and Future HPC Systems	15
<i>Ana Jokanovic, German Rodriguez, Jose Carlos Sancho, and Jesus Labarta</i>	

Session 2: I/O and Networking

Multi-root Share of Single-Root I/O Virtualization (SR-IOV) Compliant PCI Express Device	25
<i>Jun Suzuki, Yoichi Hidaka, Junichi Higuchi, Teruyuki Baba, Nobuharu Kami, and Takashi Yoshikawa</i>	
Understanding Power Efficiency of TCP/IP Packet Processing over 10GbE	32
<i>Guangdeng Liao, Xia Zhu, Steen Larsen, Laxmi Bhuyan, and Ram Huggahalli</i>	
Design and Evaluation of Generalized Collective Communication Primitives with Overlap Using ConnectX-2 Offload Engine	40
<i>Hari Subramoni, Krishna Kandalla, Sayantan Sur, and Dhabaleswar K. Panda</i>	

Session 3: Data Center Bridging

R3C2: Reactive Route and Rate Control for CEE	50
<i>M. Gusat, D. Crisan, C. Minkenberg, and C. DeCusatis</i>	
AF-QCN: Approximate Fairness with Quantized Congestion Notification for Multi-tenanted Data Centers	58
<i>Abdul Kabbani, Mohammad Alizadeh, Masato Yasuda, Rong Pan, and Balaji Prabhakar</i>	
Leveraging Performance of Multiroot Data Center Networks by Reactive Reroute	66
<i>Adrian S.W. Tam, Kang Xi, and H. Jonathan Chao</i>	

Special Session: Interconnects for Exascale Computing

The PERCS High-Performance Interconnect	75
<i>Baba Arimilli, Ravi Arimilli, Vicente Chung, Scott Clark, Wolfgang Denzel, Ben Drerup, Torsten Hoefler, Jody Joyner, Jerry Lewis, Jian Li, Nan Ni, and Ram Rajamony</i>	
The Gemini System Interconnect	83
<i>Robert Alverson, Duncan Roweth, and Larry Kaplan</i>	

Session 5: Silicon Photonics

Silicon Nanophotonic Network-on-Chip Using TDM Arbitration	88
<i>Gilbert Hendry, Johnnie Chan, Shoaib Kamil, Lenny Oliner, John Shalf, Luca P. Carloni, and Keren Bergman</i>	
Clocking Links in Multi-chip Packages: A Case Study	96
<i>Tamer Ali, Dinesh Patil, Frankie Liu, Elad Alon, Jon Lexau, Chih-Kong Ken Yang, and Ron Ho</i>	

Industrial Panel

Optics in Future Data Center Networks	104
<i>Laurent Schares, Daniel M. Kuchta, and Alan F. Benner</i>	
Direct Network Prototype Leveraging Light Peak Technology	109
<i>Sreenivas Addagatla, Mark Shaw, Suyash Sinha, Prashant Chandra, Ameya S. Varde, and Michael Grinkrug</i>	
Scaling Optical Interconnects in Datacenter Networks—Opportunities and Challenges for WDM	113
<i>Hong Liu, Cedric F. Lam, and Chris Johnson</i>	
Optical Interconnects in the Data Center	117
<i>Ron Ho, John E. Cunningham, Herb Schwetman, Xuezhe Zheng, and Ashok V. Krishnamoorthy</i>	

Tutorials

Innovating in Your Network with OpenFlow: A Hands-on Tutorial	121
<i>Yiannis Yakoumis and Masayoshi Kobayashi</i>	
Architectures for the Future Networks and the Next Generation Internet	123
<i>Raj Jain</i>	
Designing High-End Computing Systems with InfiniBand and High-Speed Ethernet	125
<i>D.K. Panda, S. Sur, and P. Balaji</i>	
The Photonics Advancements in the Emerging Zettabyte Network Architecture	128
<i>Loukas Paraschis</i>	
Author Index	130