



August 23 to 25, 2009  
Memorial Auditorium, Stanford University

**A Symposium on High Performance Chips**  
Sponsored by the IEEE Technical Committee on Microprocessors and Microcomputers



## Tutorial #1:

# ***Modern System Interconnects***

Sunday August 23, 2009

**HT3** -- *Jose Duato, HyperTransport Consortium*

**QPI** -- *Bob Safranek, Intel*

**PCIe** -- *Jasmin Ajanovic, Intel*

# Agenda

<b>9:00 – 9:10</b>	Introduction, Chuck Moore
<b>9:10 – 10:00</b>	HyperTransport, Jose Duato
<b>10:00 – 10:50</b>	QuickPath, Bob Safranek
<b>10:50 – 11:00</b>	-- Break –
<b>11:00 – 11:50</b>	PCI Express, Jasmin Ajanovic
<b>12:00 – 12:30</b>	Panel-style Q&A

45 minute presentations

5 minute at the end of each for clarifying questions

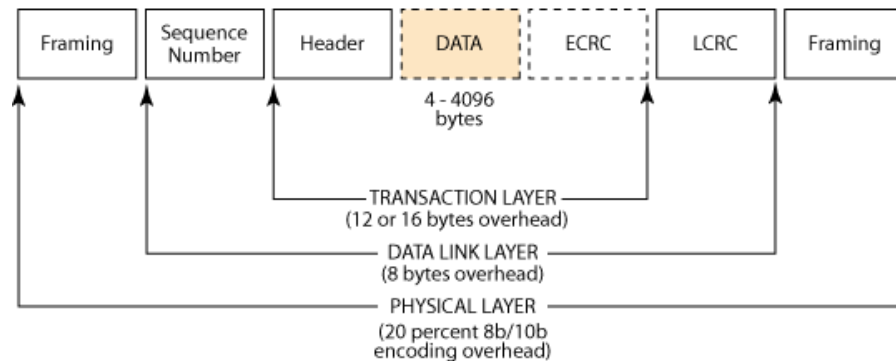
30 minute Q&A with all presenters

# PC System Architecture is Growing Up...

- No longer just self-contained single processor systems
- Small → Medium → Large SMPs
  - Connect multiple sockets to form large SMP
  - Multiple processors in each socket to form even larger SMPs
- Virtualization
  - Initially, just the CPU and memory
  - Now – I/O sharing within the SMP node: SR-IOV
  - Future – I/O sharing among SMP nodes: MR-IOV
- High BW memory, coherence and I/O subsystems **required**
  - To unleash the full capability of the compute hardware
  - To achieve reasonable system-level balance

# Implications on Interconnects

- Interconnects are no longer multi-drop shared busses
- Packet and Switch-based Interconnects
  - Much faster and much higher system-level Bandwidths
  - Point-to-point connections (lower CAP → higher speed)
  - Layered structure with SERDES-based physical layers



- Higher Reliability
  - Device failure doesn't block communications among other devices
  - Support for alternate paths (in theory)

# Interconnects Continue to Grow in Importance

- SMP Scalability
  - Balance memory latency, bandwidth, capacity and coherency
- System-level Scalability and Balance
  - Big SMPs → Significantly higher demands on Memory and I/O
- Inter-node Messaging and Sharing
  - RDMA; Shared non-HW-coherent memory;
  - Synchronization
- Data Center Optimization, Scalability and Reliability
  - Commodity nodes that plug into commodity “backplane”
  - Data plane; Control plane; Management plane
- Embedded Space
  - Wireless connections (cables will be more costly than Silicon!)

# Key I/O and Interconnect Technologies

- Memory transfer and coherence protocols
  - Quick Path Interconnect (QPI)
  - HyperTransport (HT3)
- I/O Connections
  - PCI Express – Gen1, Gen2, Gen3
  - *RapidIO*
- Data Center Connections
  - *1GbE → 10GbE → DCE*
  - *FCoE, iSCSI*
  - *Infiniband*
- Peripheral Device Connections
  - *USB 3.0*
  - *DisplayPort*



**Covered in today's  
Tutorial**

# Agenda

<b>9:00 – 9:10</b>	Introduction, Chuck Moore
<b>9:10 – 10:00</b>	<b>HyperTransport, Jose Duato</b>
<b>10:00 – 10:50</b>	QuickPath, Bob Safranek
<b>10:50 – 11:00</b>	-- Break –
<b>11:00 – 11:50</b>	PCI Express, Jasmin Ajanovic
<b>12:00 – 12:30</b>	Panel-style Q&A

45 minute presentations

5 minute at the end of each for clarifying questions

30 minute Q&A with all presenters