HyperTransport Network of Excellence

José Duato and Ulrich Brüning

Technical University of Valencia, University of Mannheim
jduato@disca.upv.es, ulrich@ti.uni-mannheim.de

February 16, 2007

Outline

Motivation (I)

- HyperTransport delivers the highest bandwidth and lowest latency among interconnect standards.
- Thanks to these excellent features, a standard interconnect like HT has the potential to weave off-the-shelf servers into powerful scalable servers, addressing the needs of large datacenters as well as making HPC a commodity.

Server Clustering

Standard Interconnects Weave Off-the-Shelf Servers into Powerful Scalable Clusters

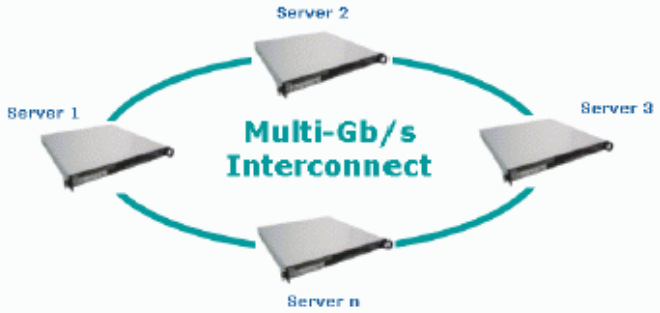

Server Clustering Makes HPC a Commodity!

Performance/Cost-Optimized Data-Center Servers

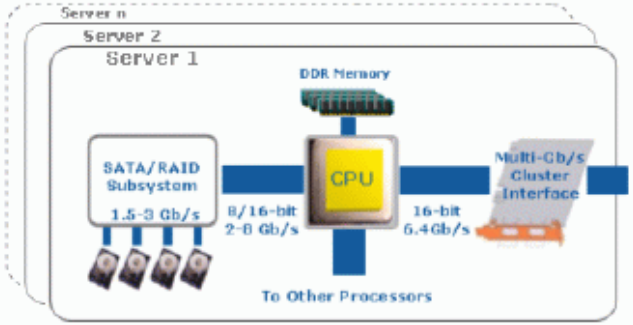

- Core Compute and Storage Functionality
- No Legacy I/O Control Subsystem
- Clustered via High-Bandwidth Low Latency Multi-Gb/s Interconnect
- Cost Saving Times "n" Factor
- Reduced Power Consumption
- Increased Reliability

Localized Legacy I/O Processing

Constraints

- Despite the benefits of introducing HyperTransport interconnect technology in current servers, those benefits are constrained to the interconnection among hosts in the same board, and between hosts and peripheral devices (including network adapters).
- Board-to-board interconnects in large servers still require the use of other interconnect technologies, possibly HT-enabled (e.g. PathScale).

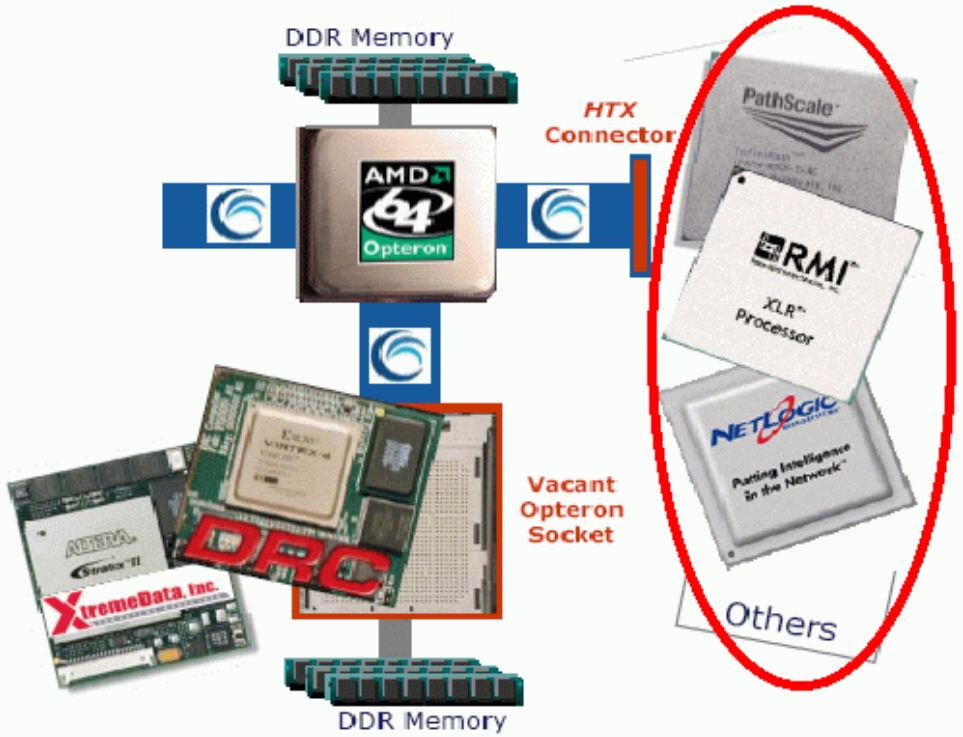

An Example: Cray XT3

Cray uses a proprietary interconnect despite the fact that link bandwidth does not exceed HT specs

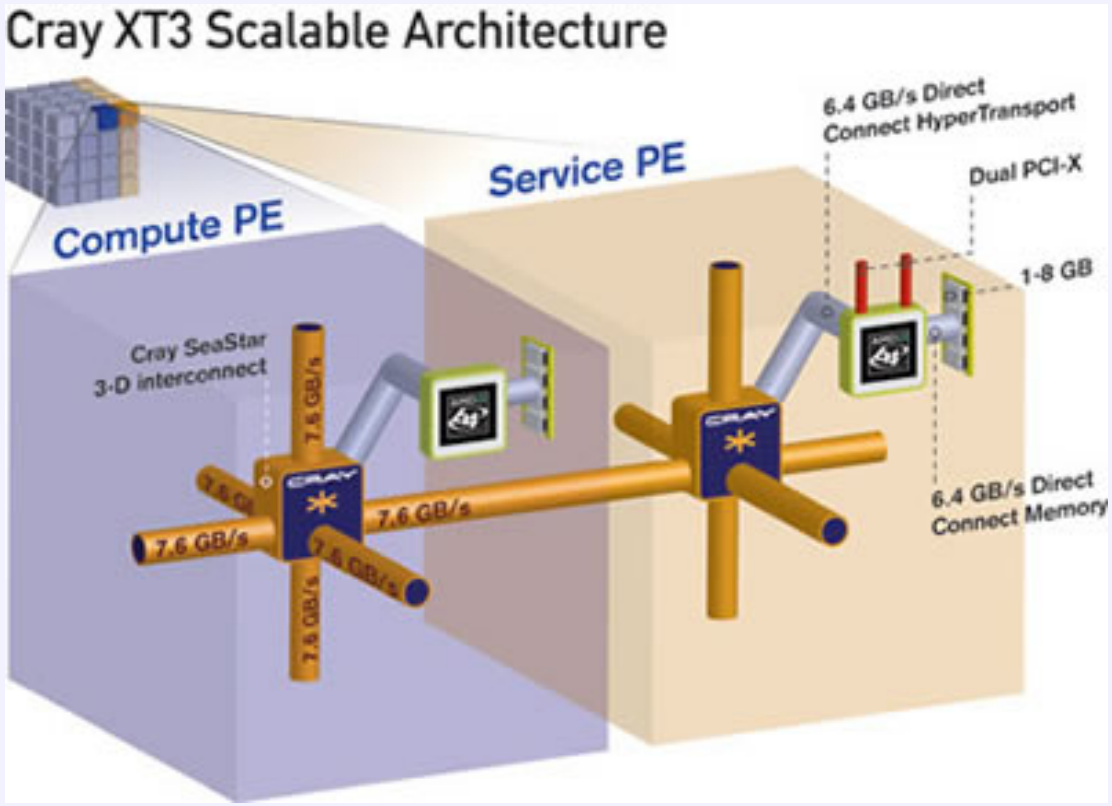

It would be nice if native HT could be used to interconnect the Opteron chips in machines like this !!!

An Example: Cray XT3

Cray uses a proprietary interconnect despite the fact that link bandwidth does not exceed HT specs

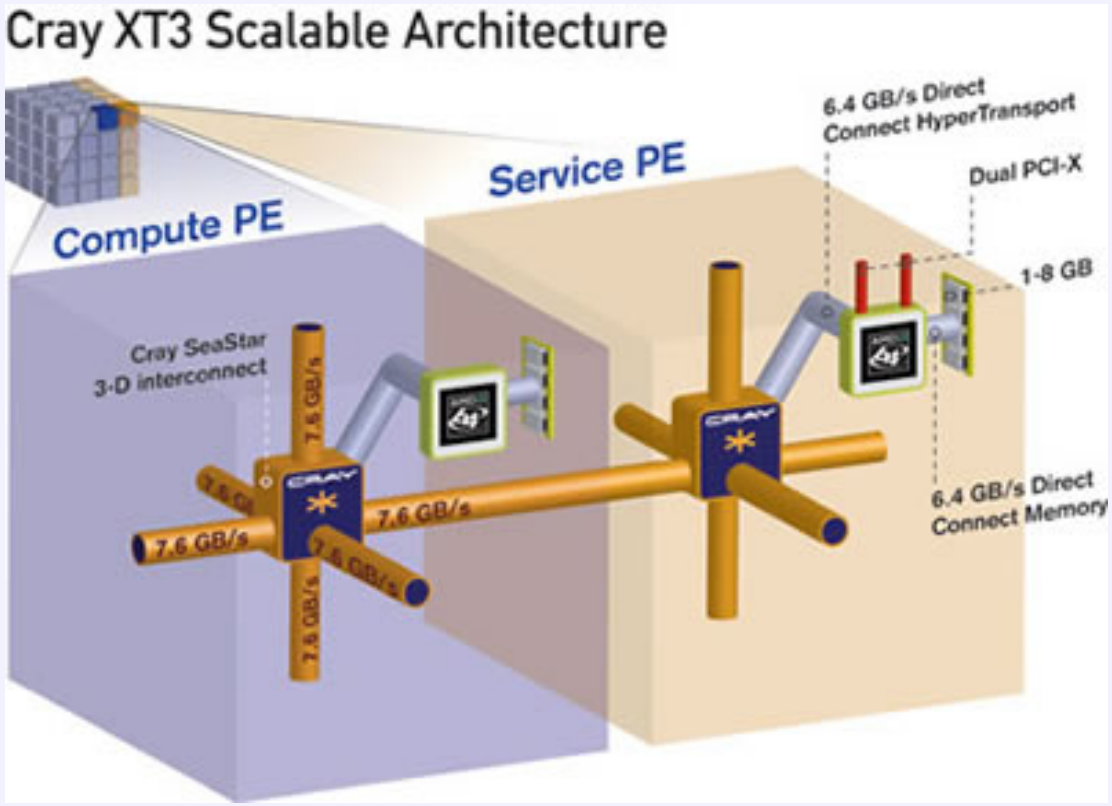

It would be nice if native HT could be used to interconnect the Opteron chips in machines like this !!!

Motivation

Why do we need to extend HT?

- Additional functionality is required to interconnect a large number of hosts among them in a native way.
- Some techniques need to be introduced in order to reduce latency even more in large switch fabrics.

How will it differ from other system area networks (e.g. InfiniBand, Myrinet, Quadrics, etc.)?

- HT can easily incorporate a network interface and a switch within the processor/accelerator chip, making it much faster, cheaper, and easier to manage.
- The extended HT will deliver lower latency than other interconnects.

Application Areas

- High-performance computing: Parallel computers similar to the Cray XT3 but using native HT links and switches.
- Datacenters: HT would provide lower latency and higher bandwidth than 10 Gigabit Ethernet or Fibre Channel for:
  - Interprocessor communication
  - Storage area networks
- Gamer enthusiast platforms: Tens of future graphics cards with HTX interface directly interconnected using HT switches to deliver amazing rendering speeds at the highest quality and resolution.
- Enterprise computing: Tens of future GPUs plus FPGA-based accelerators, all of them interconnected using HT switches to implement the fastest accelerator in the market.

Application Areas

- High-performance computing: Parallel computers similar to the Cray XT3 but using native HT links and switches.
- Datacenters: HT would provide lower latency and higher bandwidth than 10 Gigabit Ethernet or Fibre Channel for:
  - Interprocessor communication
  - Storage area networks
- Gamer enthusiast platforms: Tens of future graphics cards with HTX interface directly interconnected using HT switches to deliver amazing rendering speeds at the highest quality and resolution.
- Enterprise computing: Tens of future GPUs plus FPGA-based accelerators, all of them interconnected using HT switches to implement the fastest accelerator in the market.

Application Areas

- High-performance computing: Parallel computers similar to the Cray XT3 but using native HT links and switches.
- Datacenters: HT would provide lower latency and higher bandwidth than 10 Gigabit Ethernet or Fibre Channel for:
  - Interprocessor communication
  - Storage area networks
- Gamer enthusiast platforms: Tens of future graphics cards with HTX interface directly interconnected using HT switches to deliver amazing rendering speeds at the highest quality and resolution.
- Enterprise computing: Tens of future GPUs plus FPGA-based accelerators, all of them interconnected using HT switches to implement the fastest accelerator in the market.

Application Areas

- High-performance computing: Parallel computers similar to the Cray XT3 but using native HT links and switches.
- Datacenters: HT would provide lower latency and higher bandwidth than 10 Gigabit Ethernet or Fibre Channel for:
  - Interprocessor communication
  - Storage area networks
- Gamer enthusiast platforms: Tens of future graphics cards with HTX interface directly interconnected using HT switches to deliver amazing rendering speeds at the highest quality and resolution.
- Enterprise computing: Tens of future GPUs plus FPGA-based accelerators, all of them interconnected using HT switches to implement the fastest accelerator in the market.

Collaboration Model

- Establishing an academic team composed of leading experts in the field of interconnection networks
- Academic institutions joining the HyperTransport Consortium
- Contacting key companies in the HTC to define the set of requirements when extending HT
- Academic partners will mostly use their own research grants to fund the required research
- Academic partners will benefit from working on real problems and from the opportunity to develop “useful” technology
- Companies will benefit from the academic contribution, shortening time to market and benefitting from highly qualified experts in the field
- Crossfertilization will likely lead to joint research agreements, internships, and even startups.

HT Academic Team

- University of Mannheim (Ulrich Brüning)
- Technical University of Valencia (José Duato)
- Georgia Institute of Technology (Sudhakar Yalamanchili)
- Simula Research Laboratory (Olav Lysne)
- FORTH (Manolis Katevenis)
- University of Castilla-La Mancha (Francisco Quiles)
- University of Murcia (José M. García)

First Steps

- June 2006 - November 2006: The initial HT Academic Team is formed
- December 8th, 2006: First HT Academic Day
  - Project to extend HT is presented to key companies in HTC
- January 2007: Work starts
  - The HT TWG defines the mission of the HT Academic Team: elaborating a report proposing extensions to HT3.0
  - First set of agreements reached
- February 15th, 2007: Second HT Academic Day
  - First set of agreements presented to key companies in HTC
  - Roadmap discussed

HT Network of Excellence

- Convergence:
  - Center of Excellence for HyperTransport
  - HyperTransport Academic Team
- Simpler interaction with industry
  - Single interface vs. multiple interfaces
- Technology availability
  - e.g. coherent HyperTransport