SILICON GRAPHICS | The Source of Innovation and Discovery™

Silicon Graphics, Inc.

Future Directions of High Performance Parallel Interconnects

Presented by:
Greg Thorson, Principal Engineer
gmt@sgi.com

August 25, 2004

Personal opinions of Greg Thorson, not SGI official position

WHAT FORCES WILL DRIVE PARALLEL COMPUTING?

Off-the-shelf microprocessors

- a single socket will contain more and more cores over time
- sockets will absorb part of the network
- many 10s of cores will be available in white box systems causing mainstream compilers to go truly parallel

Technical HPC is the current big user

- unfortunately technical HPC is not a high-volume force
- low volume custom ASICs are becoming more difficult to justify

Eventually HPC will move to the masses

- only when it becomes truly off-the-shelf
- will take time for ISVs to react
- will not raise HPC volumes to very high levels for quite a while

WHAT IS HPC INTERCONNECT?

Relatively large numbers of sockets or other node elements

- more and more HPC is attained by scalability, as it is becoming increasingly expensive to provide differentiated general-purpose nodes. Nodes may include various combinations of processor cores, graphics, memory, I/O ports, and special purpose elements (e.g. structured ASICs and FPGAs)
- large node counts require high reliability links
- overall system reliability requires fault isolation to avoid widespread system crashes and slowdowns

WHAT IS HPC INTERCONNECT?

Tight Coupling

- low latency data transfer
- high bandwidth data transfer
- low latency synchronization (barriers, locks, etc.)
- collective communication (multicast, shared memory, etc.)

Optics?

- Apart from some cluster-like interconnects the latency and cost will be too high for some time
- HPC effect on optics development will be little to none

STANDARD CHANNELS

Why Standards?

- Absorption of portions of the network by the processor sockets will result in the development of standard high-bandwidth, low-latency links
- Relatively low volumes of HPC systems make custom ASICs problematic resulting in the use of structured ASICs or even FPGAs which will need to use standard physical layers

STANDARD CHANNELS

Problems with current standards

- large chip area
- high pin count
- high latency link and transport layers

NETWORK OS

Fairness

- congestion isolation
- latency sensitive traffic interacting with high bandwidth traffic
- application specific end cases

Error handling

- fault isolation
- end-to-end retry
- end cases

uP cores are getting cheaper

- it is not grossly expensive to implement hooks for embedded microprocessors to interact with the network traffic for end cases