



SiCortex

HPC Users Forum

Interconnect Panel
Introductory Slides

April 16, 2008

Where SiCortex is Coming From

- Systems company
 - System-on-a-chip silicon
 - Open Source software
 - Linux/MPI is the standard today
 - Linux--/PGAS tomorrow
-
- 10GE/IB are great for I/O



SiCortex Systems

- 64-bit Linux/MPI
- Low power (3 W/core system power)
 - SC072 Catapult - 72 core, 200 watt
 - SC648 - 648 GF, 2 KW
 - SC1458 - 1.4 TF, 4 KW
 - SC5832 - 5.8 TF, 18 KW
- Tightly coupled Kautz interconnect
 - Better than IB performance
 - Lower than Ethernet prices



SiCortex Fabric

- DMA Engine provides SW interface
- On-chip switch and links
- Copper 2 GB/sec links
- In cabinet Kautz topology
 - Log diameter, 3-way multirail
- SW Implements MPI, sockets, etc.
- 1.4 μ sec MPI, 1.5 GB/sec



Communications

- We're entering an era of high processor count computing
- Parallel programs mean communications
- Communications must be a first class operation
 - Not an I/O add-on



APIs

- Every CS problem solved by adding a layer of abstraction
- Every performance problem solved by removing a layer of abstraction
- How about ONE layer?
 - SCMPI SEND/RECV path 250 instructions
- APIs are fine...
 - but not in the critical path



Languages

- How to do parallel programming
 - Without losing performance
- MPI has been successful
 - So SCMPI with C and Fortran came first
- PGAS programs can be expressive *and* fast
 - Watch for SC to support UPC and CAF

