Panel 1

What are the future trends in high-performance interconnects for parallel computers?

Moderator:
Fabrizio Petrini, Los Alamos National Laboratory

Panelists:
José Moreira, IBM Research
Jarek Nieplocha, Pacific Northwest National Laboratory
Mark Seager, Lawrence Livermore National Laboratory
Craig Stunkel, IBM Research
Greg Thorson, SGI
Paul Terry, Cray Canada Inc.
Srinidhi Varadarajan, Terascale Computing Facility, Virginia Tech

Ever-increasing demand for computing capability is driving the construction of ever-larger computer clusters, comprising compute nodes integrated by a high-performance interconnect. This network is in most cases the heart of the parallel computer and defines its functionalities, influences the design of the system software and determines the actual performance of the machine.

The panelists will discuss the major design trends in high performance networks.

- Most interconnection networks provide some form of "intelligence" in the network interface. Do you expect this to become a central feature in the future? Will it be possible to implement "network operating systems" in the network interface?
- Do you expect that optical networks will become widespread?
- What are the trends in latency and bandwidth? Networks as Quadrics Elan4 and Cray XD1's already deliver 1.5 microseconds at MPI level. Is bandwidth technologically free? Will the I/O interface be the bottleneck?
- Will native support for collective communication be a central feature of a high performance network?
- BlueGene/L has shown that a thermally-aware supercomputer can be packaged in a small space, with a chip that integrates processors and network interface. Will future network interfaces be integrated with the processors in the same chip?
- How can a high-performance network help to achieve fault-tolerance in a large-scale machine?