Supercomputing in Small Spaces

RADIANT: Research And Development In Advanced Network Technology
http://www.lanl.gov/radiant

Theoretical Astrophysics
http://qso.lanl.gov

Figure 1: A 15-GFLOP, 24-CPU Transmeta-Based RLX Cluster in a 3U Space (circled above) and a 14-GFLOP, 28-CPU Traditional Cluster in a 28U Space.

As the supercomputing community continues to aggregate larger and larger clusters of machines, the cost and the size of the resulting supercomputer increases tremendously. Here at SC 2001, we demonstrate the ability to build and run real codes on a 15-GFLOP, 24-node Transmeta cluster by RLX Technologies in a 3U form factor, i.e., approximately the space of three pizza boxes stacked on top of one another. Figure 1 shows a picture of a Transmeta cluster dubbed the RLX System 324; it is the *single* box at the top of the rack.

To scale up our 15-GFLOP, 24-node RLX Transmeta up to a 3-TFLOP supercomputer, we would first upgrade the processor speed from 633 MHz to 800 MHz and use DDR memory (both of which are currently available), resulting in a peak rating of nearly 20 GFLOPS for our 24-node cluster. Fourteen of these 24-node clusters would then be mounted in a rack (similar in size to the rack shown in Figure 1), resulting in a peak aggregate rate of 269 GFLOPS. Replicating the rack 12 times, as shown in Figure 2, theoretically gives us a 3.2-TFLOP supercomputer that would fit along one line across the span of a soccer goal, resulting in a tremendous space savings.

While the cost of acquisition for our RLX Transmeta cluster be higher on a cost-per-node basis than in a commodity cluster or cluster of supercomputers, the total cost of acquisition and operation is significantly cheaper for our Transmeta cluster. Why? First, the RLX Transmeta cluster does not require the power nor the cooling that commodity clusters or supercomputers require. Specifically, the RLX Transmeta cluster results in a 5x–10x savings in power, i.e., 15 watts vs 75 watts under load and 7 watts vs 75 watts at idle. Second, all the nodes are hot-pluggable and tool-free to install and remove. Third, RLX's Web administration tool makes configuration fast, flexible, and easy from a remote location. Fourth, with

Figure 2: A Mock-Up of a 3.2-TFLOP Supercomputer.

336 nodes per 42U rack, i.e., 14 chasses x 24 nodes/chassis, the RLX Transmeta cluster achieves 8x the density of a traditional 1U design. Overall, the operational cost of the RLX Transmeta cluster is expected to be approximately 6x lower than a traditional cluster.

For more detailed information on the RLX Transmeta cluster, go to *http://www.rlxtechnologies.com*.

For a live demo of our RLX Transmeta cluster, check out the 40-million particle dataset simulation in the LANL booth at SC 2001.

Contact Information

radiant-info@lanl.gov

RADIANT: Research And Development in Advanced Network Technology (http://www.lanl.gov/radiant) and Theoretical Astrophysics (http://qso.lanl.gov)
Los Alamos National Laboratory
Los Alamos, NM 87545

November 2001

LALP-01-245

A US DEPARTMENT OF ENERGY LABORATORY

Los Alamos
NATIONAL LABORATORY
Los Alamos NM 87545