The Future of HPC at SGI:
Early Experience with SGI SN-1

S. Seidl, W.E. Nagel, and H. Brunst

Abstract

In 1996, SGI introduced the Origin2000, an innovative and powerful computer architecture which enables shared memory applications to run on a very large number of nodes. Based on ccNUMA, today up to 512 CPUs can operate on a global address space. After the merger with Cray Research Inc., SGI had a lot of experience with CRAY T3E, also introduced in 1996. The T3E has been one of the very successful examples where a MPP architecture is used to efficiently solve many problems in the scientific technical market. Even today, roughly half of the 50 fastest computer systems are still CRAY T3E machines (Top500 list from November 1999).

Within SGI, the next generation of supercomputers Origin3000 (code name SN-1) has been announced. This machine looks very much like an Origin2000. Nevertheless, it is promised that this machine also has quite some features from CRAY T3E. At TU Dresden, we have a strong interest in the SN-1, and we were able to run some codes on this new machine; up to 32 nodes in Mountain View at the end of May 2000. Besides kernel and communication tests, we also looked deeply into the run time behavior of real applications from five of our end users.