Benefits of Scalable Server with Global Addressable Memory for Crash Simulation

Christian Tanasescu
Nick Meng
SGI Inc.

Abstract

A wide variety of industries rely on mechanical computer-aided engineering (CAE) to improve design quality, reduce design-cycle time, and control costs. In designing automobiles, the manufacturers must meet different government regulations for vehicle safety. The maturity of Explicit Finite Element methods and the increasing computational power of today's computers have allowed the automotive industry to incorporate the Crash simulation technology in the critical path of the design process. As the application of Crash simulation moves from the component level analysis to the system level, the complexity and the size of the models increase continuously.

By combining the capabilities of Intel® Itanium® 2 processors, the open-source Linux® operating system, and SGI® NUMAlink(tm) technology, the SGI Altix family of servers and superclusters with Global Addressable Memory (GAM) is uniquely able to meet the needs of advanced CAE simulation for product development. This paper will outline the features and benefits of using SGI Altix systems with Itanium® 2 processors running LS-DYNA.

With the dramatic increase of the price performance of modern high performance computers, the model sizes of car crash and metal stamping simulations have been constantly increasing in recent years. Nowadays, a 1 Million element model is a common case. For large models, highly scalable computers and software play crucial roles in reducing turnaround time for simulations. The combination of the domain decomposition based LS-DYNA and the highly scalable SGI cc-NUMA system is an ideal solution for reducing simulation turnaround time.

After briefly introducing the SGI cc-NUMA architecture and the benefits for scalable applications like LS-DYNA, based on recent performance data on the SGI Altix family of systems, we will address the simulation grid concept enabled by a large shared-memory architecture like the SGI Altix 3700. This aims at making the workflow more efficient, not only improving the solver performance. Multi-terabyte datasets can be loaded entirely into memory and operated upon without disk-swapping. In addition, expanding the processing paradigm with I/O and graphics pipes connected directly to the large memory, results in dramatically improved performance and foster collaborative engineering.