The Applications Roadmap

Jef Dawson
Silicon Graphics. Inc.
655 Lone Oak Drive
Eagan, MN 55121

jef@cray.com

Abstract:

The Silicon Graphics HPC computing environment will be advancing very rapidly in the next few years. This presentation will focus on how the Silicon Graphics application environment will evolve and transition. The emphasis will be on solution areas (e.g., crash, external CFD) rather than on specific applications (e.g. MSC/NASTRAN™, FLUENT™).

Introduction:

Having seen the Silicon Graphics hardware and software roadmaps, it is natural to wonder about the applications roadmap. This is particularly true in light of the types of transitions that many of Silicon Graphics' customers will be making in the next few years. An obvious format for the applications roadmap might be a matrix with hardware platforms on one axis and application codes on the other, indicating the platform of choice for each code. We have found that the issues involved are too complicated to be addressed quite so simply.

Many computer users run fixed-size or slowly growing problems, and upgrade their hardware in search of incremental improvements in capability or price/performance. The traditional Cray user, however, is seeking qualitative improvements in the value derived from simulation. This makes each customer's applications roadmap unique. In fact, it is probably more meaningful to talk about a 'solutions roadmap'. This creates a bit of a 'moving target'; not only will a variety of new hardware be available, but new analysis methods will be implemented to realize much greater value.

In computational fluid dynamics, not only will faster hardware allow finer meshes for generally better accuracy, but performance will reach levels that allow fundamentally superior treatment of demanding phenomena like turbulence and multi-phase flow. In automobile crashworthiness simulation, in addition to models which include dummies, seatbelts, and airbags, we will see optimization methods that automatically minimize vehicle weight while maintaining good crashworthiness. In every industry, there are several examples like these, that show the importance of adopting new analysis methods to take best advantage of computing hardware advances. In addition to asking, "Which hardware platform is best for my application?", one needs to ask, "Which platform is best for the problem I REALLY want to solve".

While it is healthy to recognize the complexity of the situation, and keep the 'solution level' picture in mind, there are some concrete observations we can make regarding applications and hardware platforms. Among Silicon Graphics' high-end products, the Origin™ is the most versatile. It offers excellent single CPU performance for a very wide range of codes, excellent scalability, and excellent price/performance. The CRAY T90™ offers the highest available single CPU performance for vector codes, excellent shared-memory scalability up to 32 CPUs, and exceptional I/O performance. The CRAY T3E™ offers the highest scalability and maximum performance of any generally available system in the world. The CRAY SV1™ offers excellent performance, price/performance, and scalability for vector applications.

By way of example, let's consider the case of MSC/NASTRAN™, which has historically been one of the most widely used applications on Cray systems. For smaller, linear statics problems, MSC/NASTRAN™ has fairly low vector content, but very good cache hit rates. This has made the Origin™ and OCTANE™ systems the best choices for most MSC/NASTRAN™ jobs. In the automotive industry, however, very large NVH (noise, vibration, harshness) analyses are routinely run, and must be turned around as quickly as possible. These demanding simulations have significant vector content, and are extremely I/O intensive. For these jobs, the CRAY T90™ is the system of choice.

A very different picture emerges for a code like STAR-HPC™. STAR-HPC™ has very high vector content, and also shows excellent scalar performance on the Origin™. In addition, it is highly scalable. In a recent performance test, we observed a speed-up factor of 109 using 90 PEs on an Origin™ system. Thus, even though the single-cpu performance of the CRAY T90™ is unbeatable, the attainable levels of performance and price/performance on the Origin™ make it the primary Silicon Graphics system for STAR-HPC™.

Since most large systems are required to run a variety of applications, the suitability of a system for any one application code, or even specific classes of problems, is not always the end of the story. STAR-HPC™, with compelling performance on all Silicon Graphics systems, illustrates this. There are several customer sites where STAR-HPC™ is one critical application, among several other applications which are primarily vector codes. In these situations, the CRAY T90™ or CRAY SV1™ is likely the best choice. At STAR-HPC™ sites where the other important applications have little vector content, the Origin™ is usually the standout choice.

Dimitri Mavriplis, a researcher at NASA Langley, is working to demonstrate the applicability of unstructured computational fluid dynamics methods to the aerodynamics of high-lift configurations of aircraft (with flaps deployed). The complex nature of the flows involved demands meshes with 20 million cells or more. To solve this type of demanding problem, extreme scalability is required, which made the CRAY T3E™ the only choice.

Of course, these few examples can't be representative of all the situations our customers will face. Silicon Graphics will work with our customers to help them make the best decisions about their individual solutions roadmaps.[i]

---

[i] Origin™, CRAY T90™, CRAY T3E™, CRAY SV1™, Origin™, and OCTANE™ are all registered trademarks of Silicon Graphics, Inc.