Capability Prototyping: Feeding the HPC Applications Pipeline

Michael A. Heroux
Richard D. Laroche
Clayton D. Andreasen
Minye Liu
Applications Division, Silicon Graphics/Cray Research
655 Lone Oak Drive
Eagan, MN 55121 USA

mike.heroux@innovation.cray.com
http://www.cray.com

ABSTRACT:
Capability Prototyping is a set of efforts in SGI/Cray Applications to identify, demonstrate and promote new applications capabilities which enhance the value of high-performance computing (HPC). In this paper we describe Capability Prototyping and discuss our strategies in these efforts. We also give a "snapshot" of one current project and a view of where we plan to go next.

KEYWORDS:
Applications, Computer Simulation, High Performance Computing

Introduction

During the past several years, we have seen a steady increase in computing capacity from high end computer systems. Even so, a much more dramatic trend has been the reduction in cost for a fixed capacity. One benefit of this cost reduction has been the extension of computer simulation to new fields and smaller engineering firms. However, although this is a very desirable trend, a more important opportunity exists to impact science and engineering in qualitatively different ways. By focusing on new application technologies, we can take advantage of high-performance architectures in innovative ways. In this way we increase the value of computer simulation as a whole and provide a new set of applications which provide sufficient justification for high end systems.

The Cost of Computing

Improvements in the cost of computing are well-known. In twenty years we have seen workstations, and even PCs, become as powerful as a Cray-1 for many applications. In addition, application developers, along with applications support staff from hardware vendors, have made significant enhancements to application performance which multiplies this effect. In Table 1 we show the history of performance of a particular FLUENT(tm) model, beginning in 1988 on a Cray YMP all the way up through the Cray Origin 2000. The overall improvement in price/performance is approximately 1500 times. This is due in part to performance improvements, but is primarily due to lowering costs and performance enhancements made to the FLUENT source code.

Table 1: Performance Improvement History for a FLUENT(tm) Test Case

• High Efficiency Vortex (HEV) Turbulent Static Mixer (FLUENT)

•Grid Size: 115x32x27 (99360) Cell Grid

• Time to Solution on a YMP in 1988: 54,960 sec (15.3 h)

	Year Delivered	Price Magnitude	Solution Time* Same	Improved+	Cost Reduction Factor* Same	Improved+
YMP	1988	$10.0M	1.0	15	1.0	15
C90	1991	$10.0M	1.7	26	1.7	26
T90	1994	$10.0M	4.1	60	4.1	60
J90	1994	$ 0.3M	0.8	12	26.6	400
Origin	1996	$ 0.1M	1.0	15	100.0	1500

* As inverse ratio over original YMP results

+ Code improvements have increase performance by 15 to 20 times

This reduction in cost is clearly a positive trend for users of computer simulation. However, at the same time, further reductions in cost for the same sized problems will have a decreasing impact on the overall cost of computer simulation. This is because other costs such as software license fees, staff salaries, etc., become the dominant costs. As this happens, it becomes more important to improve the quality of results coming from the computer simulation process than to decrease the cost of computer system further.

An HPC Response to Low Cost Computing

As today's simulation models migrate to cheaper computer systems, proponents of HPC must strive to develop new uses for high-performance systems. In some cases this can be as simple as running the same type of problem with greater resolution. However, eventually this approach will not be attractive. In this situation, we must demonstrate new capabilities which go beyond standard simulation methods and exploit the unique characteristics of modern HPC computer systems. One set of efforts in this direction is the *Capability Prototyping* efforts in SGI/Cray Applications.

Clearly prototyping new capabilities has some risk associated with it. Any particular project ultimately may not prove to be useful. However, at the same time it has the potential for qualitative improvement in results. To increase the likelihood of successful prototyping efforts, one of the best sources for challenging new problems and ideas is the set of (often high-end) customers who have critical business decisions which depend on computer simulation. These customers are often making multi-billion dollar decisions based in part on computer simulation. They are often willing to explore new areas, because the pay-off for getting greater insight into their problems far exceeds the cost of the simulation efforts. Capability Prototyping efforts at SGI/Cray target the needs of these high-end customers, exploring innovative ways to deliver value from high performance systems.

Capability Prototyping: What is it?

In a few words, Capability Prototyping is developing a "Proof-of-Concept" for a new way of using computer simulation. It starts with a good idea (usually from customer interaction). Based on this idea, we develop a basic prototype and, if the prototype seems promising, demonstrate the prototype to customers, independent software vendors (ISVs) and experts in the appropriate fields. Once there is clear value in the prototype, we negotiate with ISVs to get the prototype technology out to the marketplace.

In determining what makes a good candidate for a Capability Prototyping project, we focus on ideas which have a

potential for delivering qualitative improvement in results. Sometimes the work is speculative with a medium to high risk factor that the project might not succeed. However, when a project succeeds, the pay-off is usually very high and the demand for the new capability is instant. In a humorous way, Figure 1 captures the essence of this instant necessity.

Figure 1: Instant Necessity
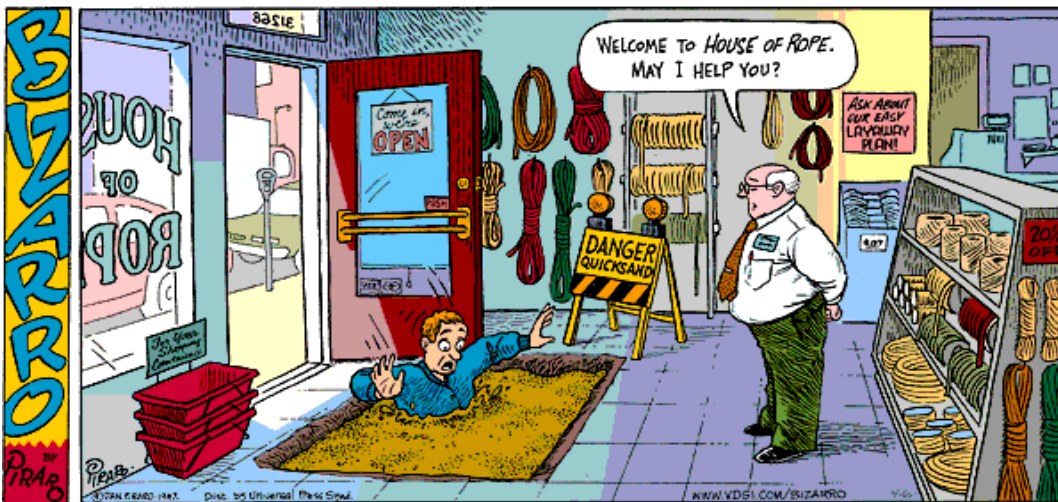
Finally, one should note that, even though we focus on qualitative improvements in simluation, we also focus on adding value to existing tools. Almost all activities involve work with existing applications as a starting point.

Capability Prototyping: What is it not?

As important as what Capability Prototyping is, is what Capability Prototyping is not. It is not long-term applications development and support. SGI/Cray Applications Division is not becoming a software house. Because we need to have excellent relationships with ISVs, any application development effort within SGI/Cray must be pushed out to the marketplace. This allows us to avoid the following problems:

1. The value of what we develop would become clear and ISVs would eventually incorporate similar features in their own codes. This would lead to competition with the very same ISVs we want to have port their code to our hardware platforms.
2. Our development staff would take on increasingly more customer support and have less time for focusing on new prototyping capabilities.

The above two problems are unavoidable if we hang on to any project for too long.

Capability Prototyping is also not incremental in nature. Incremental changes in performance may be required to achieve a larger goal. However, our ultimate objective is not price/performance, it is improved capability.

Finally, Capability Prototyping is not development from "First Principles". This effort is too costly and time-consuming. However, we will facilitate demonstration of, and foster acceptance of, fundamentally new applications.

Work in Progress

There are a number of Capability Prototyping efforts underway at SGI/Cray. In order to give readers an idea of the type

of projects we have, we discuss the Fluid Mixing Analysis Project and the related Particle Tracing Tools. These tools are being developed by Dick Laroche, Clay Andreasen and Minye Liu in SGI/Cray Applications. Details of this project are available at the URL http://www.cray.com/cpcfd.

Starting Point

Computational Fluid Dynamics (CFD) is used across a very broad set of engineering disciplines. Aerospace, automotive and process engineering are a few of the more important ones. In many cases a steady or time-dependent flow field is computed for a given geometry and input conditions using commercial CFD packages such as FLUENT(tm), FIDAP(tm), or AEA CFX (tm). Once the flowfield is computed (which for large problems can require hours or days of computing time), the user queries the results to gain insight into the engineering problem. Postprocessing and visualization tools are very valuable at this stage and numerous tools are available to aid the user. At the same time, customers in the process engineering area were asking for more information from flowfield data. In particular, they wanted to be able to generate statistics for a given model, such as a continuous stirred tank reactor (CSTR), which would give insight into the residence time distribution that a particular CSTR design would have. To address this need, SGI/Cray Applications started a project looking at large scale particle tracing for CFD models.

Large Scale Particle Tracing

In order to get accurate statistics, we determined that we must be able to track 100,000 to 10,000,000 particles for a given flow field. We know of no existing tools which can handle that many particles efficiently and therefore wrote our own prototype particle tracking tool. Although the particle tracing tool can be thought of as a postprocessing phase, when we track this many particles the tool requires an HPC system. Both the computing and data storage requirements can be extremely demanding. It is worth noting that the Cray Origin 2000 is the perfect platform for running this tool. The distributed shared memory architecture allows for easy parallel implementation and excellent parallel execution times.

Serendipity

As a part of the debugging process, we developed a simple visual interface to the particle tracing tool. However, we very quickly discovered that this visual interface provided remarkable insight into the characteristics of the flow field. When run on and SGI Onyx with an Infinite Reality engine, we are able to view tens of thousands of particles in motion (this cannot be done well on a smaller system) and we are able to see details of the flow field structure that we never saw before. Some of the data we viewed was more that five years old, but we were able to get new insight from the visual interface.

In addition to the surprising value from the visual interface, we have found the particle tracing tool to be extremely useful across many CFD models. We have viewed data from airplane wing studies, underhood flow models, external aerodynamics and other areas. Users in all these areas are very excited by the insight they gain from this tool.

Current Work

Although the visual interface to the Fluid Mixing Analysis tool has been well received, we are pressing forward with the original goal of the project, that is to generate accurate statistical data from flow field in mixing problems. We are very encouraged by initial results and are hopeful that we will be able to generate results that not attainable through experimentation.

Other Projects

We are working on a number of other projects at this time. For details, please contact the author.
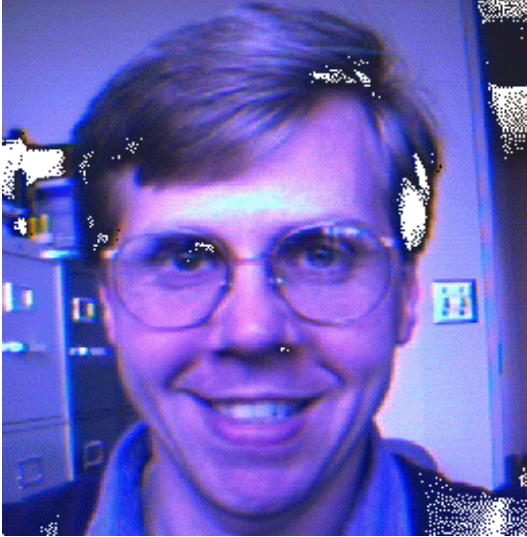
Summary

Silicon Graphics/Cray Research is committed to working with customers to demonstrate new and compelling applications technology. Capability Prototyping efforts in the SGI/Cray Applications Division are a strong sign of that commitment. Among hardware vendors, we have a unique opportunity to deliver innovative technology which uses the all of our strengths in high performance computing, data management and graphics. In addition, we have a customer base which wants to solve new and challenging problems. Working together we can demonstrate the compelling value

of modern HPC systems.

Author Biography

Dr. Heroux is the leader of the Scalable Computing, Algorithms and Capability Prototyping groups in the Applications Division at Cray Research. His primary focus is on design and optimization of numerical methods for large scale scientific applications.

mike.heroux@innovation.cray.com
http://www.cray.com