Delivering Interactive Parallel Computing Power to the Desktop

Interactive Supercomputing, Inc.
135 Beaver Street, FL 2
Waltham, MA 02452
www.interactivesupercomputing.com

Summary

- As simulation replaces physical testing, and increasingly complex phenomena are modeled, high-performance computers (HPCs) are growing in importance for science and engineering;
- Interactive desktop science and engineering tools are critical in new model and algorithm development;
- The current HPC workflow is fundamentally flawed, due to its inability to support an interactive science/engineering discovery process;
- An emerging class of software – an interactive parallel computing platform – is showing promise in enabling desktop tools to operate interactively with HPCs (Figure 1);
- An interactive parallel computing platform will help desktop tool vendors leverage HPCs without having to solve the significant challenges associated with parallel programming and supporting multiple HPC platforms.

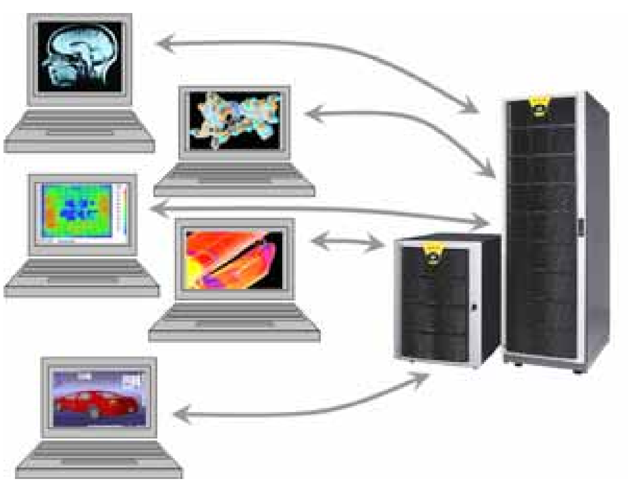

Figure 1: Bridging Desktop Applications and High-Performance Computers

Growing Computational Challenges in Science and Engineering

The launch of the Windows 95 operating system ushered in the era of the desktop computer as the primary science and engineering computing vehicle, particularly during the early stages of new product or system modeling, simulation, and optimization. The interactivity offered by these tools lends themselves well to the iterative process of research and discovery.

That said, there is a growing class of problems that demand orders of magnitude more computational power than desktops deliver. Across a broad set of industry sectors, the problems scientists and engineers need to solve are growing in size and complexity. This growth is by far outpacing the hardware advancement delivered by Moore's Law.

Competitive pressures – whether in research, defense, or commercial sectors – are demanding dramatic reductions in *time-to-solution* for critical problems. To answer this call, simulation is increasingly replacing physical prototyping and testing.

As a result, models and algorithms of increasing complexity are needed to simulate whole products (rather than sub-systems) and multi-domain phenomena. Many problems require 10-100 times the computing power of a typical science/engineering desktop workstation, and thus cannot be solved without high-performance computers (HPCs).

The Increasing Importance of High-Performance Computing

High-performance computing, once the domain of multi-million dollar supercomputers, has crossed into the mainstream. Cost-effective clusters are proving themselves to be capable workhorses across a broad set of industry sectors and computational challenges, driving the HPC market beyond $7B in 2005 (Figure 2).

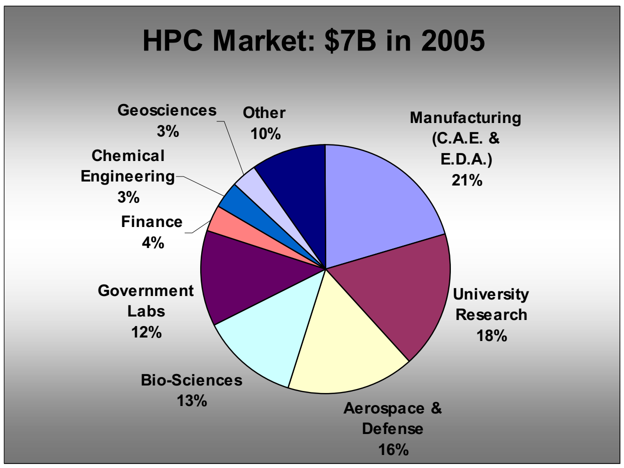

Figure 2: HPC Use by Sector

A recent study from the U.S. Council on Competitiveness [1] has found that:

1. High-performance computing is essential to business survival;
2. Companies are failing to use HPC as aggressively as possible;
3. Companies don't have the HPC tools they want and need;
4. Dramatically more powerful and easier-to-use-computers would deliver strategic, competitive benefits, and could add billions to the bottom line.

Clearly, although HPCs are a key driver for science and engineering, there is great room for improvement in how they are used.

Model and Custom Code Development for HPCs Takes Too Long

Because interactive desktop science & engineering tools cannot be used on HPCs, lots of time and effort is spent moving problems from desktops to HPCs. **This workflow is fundamentally flawed.**

Scientists and engineers often start out prototyping with a familiar tool, such as MATLAB® (Figure 3). They then stop, write a software specification, to which the model is recoded – in C or Fortran, and MPI ("Message Passing Interface") – to enable it to run on an HPC. Only then can they test and scale the model with real data (which they could not do on the desktop). The recoding invariably takes away from model refinement.

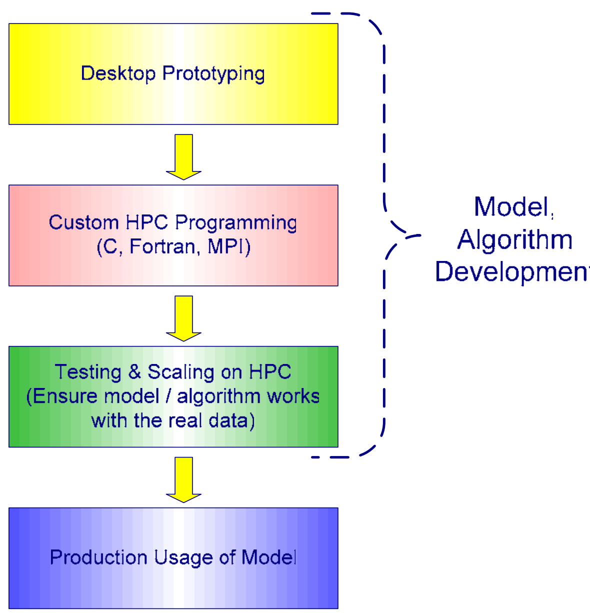

Figure 3: Traditional HPC Workflow

In many cases, the correct algorithm, approach, or key to the problem, may not be known up front, and may typically be discovered only by running the code on the HPC, with the actual input data.

Porting the models and algorithms to HPCs becomes a traditional software development process, largely serial in nature, and one that does not lend itself to the interactive workflow typically associated with the early stages of model and algorithm development and refinement.

Because each of these steps can be several months, scientists and engineers are limited to how many iterations to the algorithms and models they can make. And this all happens before they ever get to actual usage of their models, solving the problems they have set out to solve.

More than half of the "time to solution" is spent programming the models for use on HPCs, rather than developing and refining them up front, or using them in production to make decisions and discoveries.

The Need for Interactive High-Performance Computing

When reducing *time-to-solution* is the goal, it is the engineers' and scientists' time that is typically the gating factor, not computing resources. During the model/algorithm development phase, interactivity is critical. Yet although interactive use can be taken for granted with desktop science and engineering tools, to date it has simply not been available in high-performance computing, which remains firmly in the batch world.

The emergence of an interactive parallel computing platform – one that bridges interactive desktop applications and high-performance computers – is about to change the paradigm.

There is now an opportunity bring the power of HPCs to today's and tomorrow's users, by eliminating the last hurdle: the need to manually reprogram models initially prototyped using desktop tools.

With an interactive platform, the re-coding is eliminated (Figure 4), so scientists and engineers can use the same desktop tools they know and love. They can now work with the large data sets, live in memory.

They can prototype and scale in a tightly coupled process, in real time, with fine-grained control of both algorithms and data, transparently harnessing the HPC's computing resources.

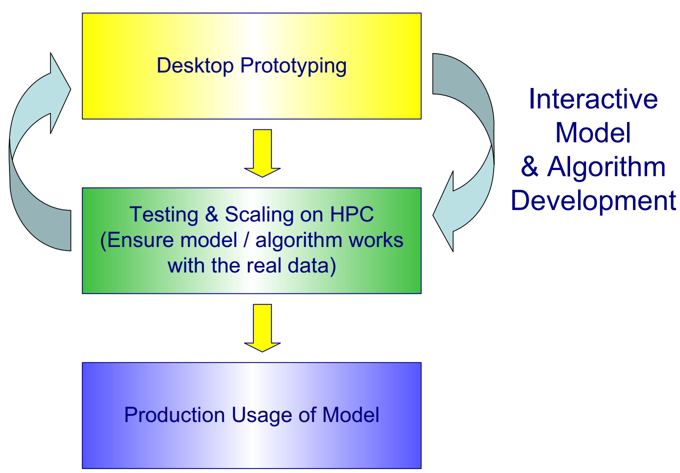

Figure 4: Interactive HPC Workflow

The Opportunity for an Interactive Parallel Computing Platform

A key benefit of HPCs is that many computations can take place in parallel, at the same time, leveraging multiple processors and access to expanded memory. These benefits come with a price – numerous technical challenges must be overcome in programming parallel computers, including splitting up the job across multiple processors, managing inter-processor communication and large memory.

The programming challenges are a key bottleneck in both the *time-to-solution* for key problems, and the adoption of HPCs in science and engineering. A recent IDC study of independent software vendors (ISVs) serving the HPC sector [2] has found that:

1. The business model for HPC-specific application software has all but evaporated in the last decade;
2. ISVs focus their software development primarily on the much-larger technical computing markets for desktop systems;
3. ISVs typically cannot afford to spend the time and money that would be needed to rewrite their software for HPCs;
4. ISV applications can exploit only a fraction of the available problem-solving power of today's HPCs;

An interactive parallel computing platform can:

1. Reduce for the ISVs the required efforts to port their applications to HPCs;
2. Expand the number of software applications that run on HPCs;
3. Deliver the power of high-performance computing to interactive desktop tools;

The Star-P™ platform from Interactive Supercomputing (Figure 5) does exactly that, bridging science and engineering desktop tools and HPCs. Based on technology licensed from the Massachusetts Institute of Technology, Star-P™ enables a client-server model in which a desktop client application is transparently linked with a powerful remote supercomputing server through a standard computer network.

Star-P™ Interactive Parallel Computing Platform

About Interactive Supercomputing

Interactive Supercomputing, Inc. develops and markets Star-P, a software platform that allows existing desktop simulation tools to operate interactively and automatically on high-performance computers. The Star-P software acts as a bridge between popular scientific and engineering desktop computing tools, like MATLAB® by Mathworks, and the parallel computing capability of HPC servers. More information about Interactive Supercomputing is available at: www.interactivesupercomputing.com

References

1. A copy of the DARPA-sponsored report entitled "Council on Competitiveness Study of U.S. Industrial HPC Users" is available at http://www.compete.org/pdf/HPC_Users_Survey.pdf

2. For more information about the Department of Defense High Performance Computing Modernization Program see http://www.hpcmo.hpc.mil/

3. For more information about the High Productivity Computing Systems initiative, see http://www.highproductivity.org/

4. A copy of the DARPA-sponsored report entitled "Study of ISVs Serving the High Performance Computing Market: The Need for Better Application Software." is available at http://www.compete.org/pdf/HPC_Software_Survey.pdf