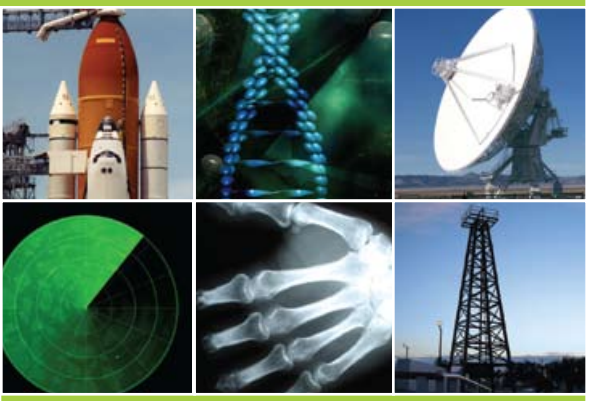

Star-P® Product Overview

Star-P drives productivity by significantly increasing application performance while keeping development costs low. The Star-P software platform seamlessly integrates desktop clients with high-performance servers. By offloading computation, memory and file intensive operations to the server, Star-P enables easy to use desktop application development, while creating the potential for execution at supercomputer speeds.

Engineers and scientists are under increasing pressure for dramatic reductions in time-to-solution for critical problems. Simulation is replacing physical prototyping, and models are expanding to include multi-domain phenomena and to reflect whole products rather than subsystems. As a result, high-performance computing is needed to solve large and complex problems that can no longer be solved on the desktop.

THE CHALLENGE

Time-Consuming and Flawed Programming Model

While the desktop environment provides low application development cost, it suffers from limited application performance; the server environment provides high application performance but at high development cost.

Today, a domain expert prototypes models and algorithms on a desktop—typically using a very high-level language(VHLL) such as MATLAB® or Python—and then has to wait while a parallel programmer reprograms them (in C, C++ or Fortran using threading and message passing) for a cluster. More than half of the time-to-solution is spent reprogramming the models and algorithms for parallel computing, rather than developing and refining their functionality.

THE SOLUTION

Star-P Accelerates Parallel Code Development

Star-P is an interactive parallel computing platform that extends easy to use Very High Level Languages (VHLLs) such as MATLAB® and Python to support simple, user-friendly parallel computing on a spectrum of computing architectures: multi-core desktops and servers, clusters and large shared memory servers. Star-P fundamentally transforms the workflow, substantially shortening the "time-to-solution" by allowing the user to easily adapt their application for use on parallel resources.

INTERACTIVE supercomputing

Star-P Functional Overview

STAR-P SERVER

Star-P Interactive Engine

Star-P's interactive engine runs on top of the server operating system, and manages multiple interactive sessions in a multi-user environment, giving client applications simultaneous interactive access to the server's processors, memory, and file system. The administration interface allows users and system administrators to create and modify profiles tailored to the needs of the organization and applications. Computing resources controlled by a workload manager such as LSF or PBS are controlled directly by Star-P with no user intervention required. You can control connection parameters and access rights, assign user privileges, and monitor system and session status.

Parallel Execution Platform

Star-P's computation engine consists of three key elements:

Data-Parallel Computations

Data-Parallel Computations perform high-level matrix and vector operations on large data sets. Once a variable is tagged as parallel, through propagation, related variables also become parallel. Functions on parallel variables are transparently "overloaded," or carried out by the data-parallel libraries.

The Star-P data-parallel libraries are high-performance optimized libraries that can be called by the client application to perform compute-intensive operations on large distributed data sets. Star-P's parallel data I/O capabilities take advantage of native I/O hardware yielding dramatic speedup when moving datasets from disk to memory.

Some libraries are proprietary to Interactive Supercomputing; others are built on top of publicly available packages, with key enhancements in terms of correctness, accuracy, and speed. Together, the libraries represent hundreds of individual computing functions in key computational domains including linear algebra, signal and image processing, optimization, statistics and others.

Task-Parallel Computations

Task-Parallel Computations execute many independent calculations in parallel, such as Monte Carlo simulations, or "un-rolling" serial FOR loops. For example, in a medical application involving image processing on multiple brain slices, Star-P can distribute the images across several processors, and simultaneously process them.

Star-P's task parallelism is simple and intuitive. A measure of parallel abstraction is that a program should execute independent of the number of processors it has access to. With Star-P, there is no need to worry about the number of available processors—Star-P takes care of distributing the data and executing the computations.

Connect™ Library API Link

Star-P's Connect™ library API link enables you to extend the functionality of the Star-P compute engine based on your application and algorithm requirements. You can plug in existing serial and parallel libraries, access them via desktop VHLL tools, and execute them in task- and data-parallel modes. In addition to Star-P's computation engine, popular external math engines from commercial vendors and the open source community can be added to expand the palette of task-parallel functions and capabilities.

By enabling access to these libraries through desktop tools, Star-P puts the power of in-house, open-source and commercial libraries into the hands of scientists, engineers, and analysts who may not have been unable to leverage these low level (C, Fortran) codes.

STAR-P CLIENT

The Star-P client resides on the user's desktop, intercepts calls to desktop math function libraries, and forwards them to the parallel libraries on the server. The Star-P client also provides an abstraction layer, a powerful capability that enables algorithms and applications to be developed and executed independent of the server hardware and processor count. User's can easily "point" their client at alternate hardware resources depending on the needs of their application.

Familiar Desktop Tools

Star-P software is a bridge between popular computing tools such as MATLAB and Python, and the grids, servers, and clusters used widely in technical computing. With Star-P, you can use your favorite desktop simulation tool, with its familiar features, commands, and data types. Standard commands and functions are available and transparently perform in a parallel manner. Existing scripts can be reused to run larger problems in parallel with minimal modification. This reduces the learning curve substantially, and dramatically accelerates the development of custom parallel applications.

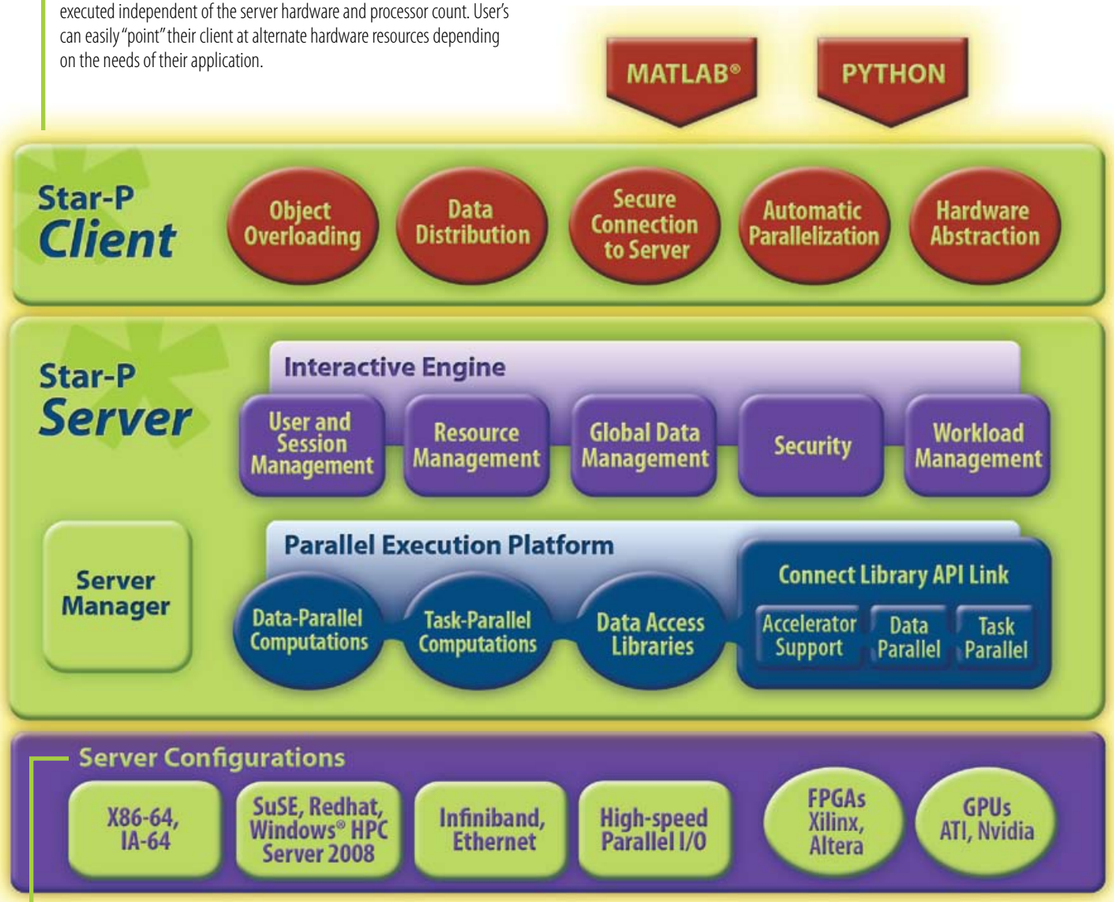

SERVER CONFIGURATIONS

Star-P software is available for a broad range of high-performance technical computing servers from leading vendors. Star-P runs on mainstream multi-core microprocessors and supports mainstream server operating systems used in technical computing. For the most current specifications, please visit www.interactivesupercomputing.com.

High-Speed Parallel I/O

Star-P's parallel data I/O capabilities take advantage of native I/O hardware yielding dramatic speedup when moving datasets from disk to memory. Star-P enables direct parallel data transfer between servers disk systems, removing the desktop PC bottleneck. And, Star-P eliminates the need to manually break up files into chunks for loading and processing.

Hardware Accelerators

Hardware accelerators such as FPGAs and GPUs can give technical computing users significant computation, I/O and memory bandwidth advantages over traditional CPU-only solutions. Through the Star-P Connect library API link, compute-intensive algorithms embedded in hardware appear as standard library functions, and can be easily called from the high-level desktop application.

"With Star-P, researchers can now preserve their familiar workflow while tackling data sets orders of magnitude larger than they process on their desktops."

–Mark Barnell, Computer Sciences Corporation

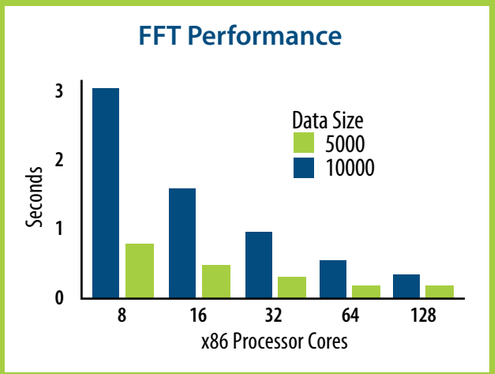

"Our molecular models involve computations with 20,000 x 20,000 matrices and quickly reached the limits of desktop memory and computing power. We used Star-P to continue working interactively within our MATLAB environment, extending our computations to larger, more complicated models."

–Bryan Wong, Department of Chemistry, Massachusetts Institute of Technology

"Problems that were completely impossible before in MATLAB are easy now. The code adjustments have been minimal."

–Andrew Meijers, University of Tasmania

The Need for Interactive Parallel 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 the Star-P interactive platform—one that extends interactive desktop applications for parallel computing—changes this paradigm. There is now an opportunity to bring the power of parallel computing to today's and tomorrow's users. With Star-P, recoding is eliminated, so scientists and engineers can use the same desktop tools they know and love. They can now work directly with large, real-life data sets, prototype and scale in a tightly coupled process, in real time, with fine-grained control of both algorithms and data, transparently harnessing the computing power of multi-core servers and clusters.

Interactive Workflow Shrinks Time-to-Solution

During the model/algorithm development phase, interactivity is critical. In many cases, the correct algorithm, approach, or key to the problem, may not be known in advance, and may be discovered only by running the code full-scale, with actual input data.

Star-P eliminates the need to re-code algorithms in C++ or Fortran, using threading or message passing to take advantage of parallel computing power. Now, you can work in an interactive high-level environment, from prototyping through production.

Through fine-grained control, Star-P revolutionizes model development and refinement. No longer must compiled programs run start-to-finish in overnight batch runs. Instead, the code can be run directly in parallel, full-scale, dramatically speeding up response time, and enabling rapid iterations.

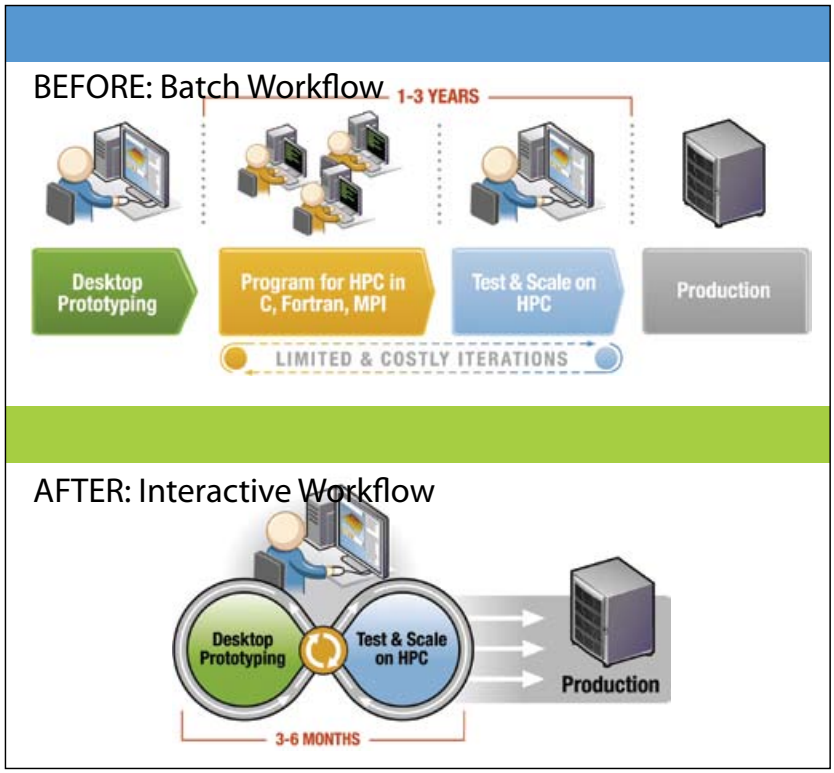

Star-P eliminates the recoding for parallel computing typically found in the batch workflow (top), instead enabling an interactive workflow (bottom), with rapid interations and dramatically reducing the time-to-solution

Star-P Future Proofs Applications from Changing Processor Technology

As standard processors add cores and differing accelerator technologies appear and evolve, the application developer risks wasting precious development time tuning and retuning their application. With Star-P, these adjustments are made automatically. Star-P's abstraction layer removes the need for code changes related to changes in the underlying hardware architecture through library substitution and optimal work allocation. User application code contains no hardware specific configuration information. This is especially important for large organizations which may have multiple, different, computing platforms.

INTERACTIVE supercomputing

Contact:
Headquarters - James River Technical, Inc.
4439 Cox Road, Glen Allen, VA 23060, tel: 804.935.0150,
sales@jrti.com, www.jrti.com

JAMES RIVER TECHNICAL, INC.

About Interactive Supercomputing

Interactive Supercomputing Corporation (ISC) develops Star-P, a software platform that drives productivity by significantly increasing application performance while keeping development costs low. Across a broad range of security, intelligence, manufacturing, energy, biomedical, financial, and scientific research applications, ISC enables faster prototyping, iteration, and deployment of large-scale solutions on high performance computing servers.

Star-P was originally developed at the Massachusetts Institute of Technology, with support from the National Science Foundation. ISC, launched in 2004, holds an exclusive license from MIT to commercialize the technology, and has independently filed multiple patents. Since launch in 2005, Star-P has been adopted at leading government labs, research institutions, and commercial enterprises