Faster, More-Efficient Massively Parallel Supercomputing

Challenge

With the advent of high-performance computing in the form of massively parallel supercomputers, an entire new generation of software engineering was initiated. These machines are greatly expanded versions of what the PC world knows as multicore processing, that is hardware that can process more than one data stream simultaneously. In massively parallel supercomputers, the scale is huge: thousands of independent hardware processors, some or all multicore, crunching data simultaneously. Programming for these machines quickly became a time-intensive enterprise, until the realization that there were certain algorithms in such software that tended to appear in programs across the sciences—from life science to physics and engineering. Sandians took on the challenge of building libraries of such algorithms in an LDRD project known as "Hybrid Sparse-Dense Incomplete Factorization Preconditioners," which enabled the whole concept of modular, reusable library components as independent but interoperable packages that would be available to programmers and could be used in a diversity of science- and engineering-specific applications. The original three packages in this LDRD project ultimately spawned Trilinos, the first software library of algorithms for high-performance computing (HPC) science and engineering applications.

From a different angle, optimizing the power of massively parallel machines requires that most or (ideally) all of the processors be utilized to crunch data at all times. Beginning with LDRD funding, Sandia addressed—and is continuing to address— this challenge.

A portion of the cabinetry containing some of the thousands of processor racks comprising Sandia's Red Storm massively parallel supercomputer.

Research Approaches

Trilinos

Drawing from a half-dozen LDRD projects and garnering two R&D100 awards over a decade, Trilinos' modular library concept allows software developers in diverse disciplines to draw from an extensive array of modular algorithms around which to fashion their application-specific code, thereby greatly facilitating the task of developing such software for large-scale optimization problems such as those involving climate change, Mach-3 aerodynamics, and radiation-hardened electronic circuits.

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND

From the initial project that spawned Trilinos in 1999-2001, LDRD provided unique funding for all of the critical capabilities that Trilinos is known for, particularly, its ability to ensure scalability of HPC code. Trilinos provides the foundation for most Sandia applications that need scalable solvers and supporting infrastructure. In addition to scalable solutions, Trilinos provides a full vertical capability for tightly coupled multi-physics modeling. Furthermore, Trilinos is unique among collections of libraries in its ability to deal with uncertainty in computations. It has become the leading library collection for decadal climate solver requirements in atmospheric, ocean, and sea ice modeling. It is used in multi-phase fluid flow for numerous environmental applications, and is poised to provide the foundation for at least one next-generation oil & and gas industrial application.

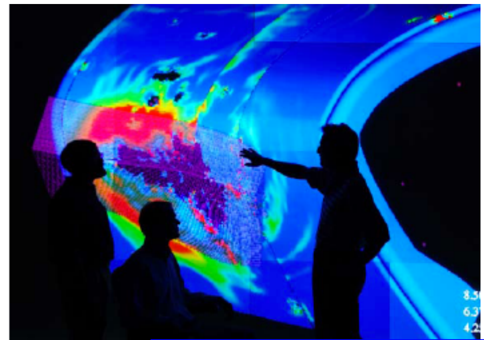

Researchers study a supercomputer simulation.

Compute Process Allocator (CPA)
The CPA algorithm, and its associated experimental parallel computer, CPlant, received an R&D100 award for their ingenuity in allocating different jobs to a massively parallel machine's processors, such that the machine functions most efficiently, that is, is utilizing most or all processors to crunch data at all times. CPA uses both a space-filling curve and span minimization to ensure optimized processor allocation to different jobs running on the same machine. The algorithm initially improved throughput by 23%, that is 23% more jobs could be run over a given time period because of this efficiency of processor allocation.

Impact

Not only are Trilinos algorithms ubiquitous in Sandia HPC applications, but they are found in HPC code, worldwide, given their availability as a collection of over 50 open-source libraries. In this sense, Trilinos has played a critical role in the development of advanced applications for HPC across science and engineering disciplines. In addition to ensuring scalability of HPC code, it also assists in managing computational uncertainty. Moreover, the R&D100 Awards provide the evidence that the Trilinos Project was truly groundbreaking in that it pioneered the modular algorithm concept as a now global standard in crafting HPC code. Public availability has led to numerous external collaborations, bringing in several million dollars of external funding. A new release of Trilinos, Version 0.8 was released on October 1, 2011; within 5 days, there were 165 downloads. The new release contained 50 packages and increased support for multicore CPUs and GPUs. Trilinos adoption has grown within the Center for Advanced Simulation of Lightwater Reactors (CASL)

The CPA project initiated collaborations with the University of Illinois, Urbana-Champaign, the State University of New York, Stonybrook, and with Cray, Inc.

manufacturer of Sandia's Red Storm and Red Sky/Red Mesa supercomputers. One of these algorithms has been extended and incorporated into Cray's, Application Level Placement System (ALPS), which is used in conjunction with the Moab scheduler on Cielo, a powerful but processor-intensive simulation software package available on Los Alamos National Lab's (LANL's) Cielo machine. In addition, one CPA algorithm has been adopted by Lawrence Livermore National Lab's (LLNL's) simple Linux Utility for Resource Management (SLURM)

Funding Sources
LDRD
DOE Office of Science (OS) Advanced Scientific Research.
CSSE (the Computational Systems and Software Engineering of the NNSA's Advanced Simulation and Computing [ASC] Program)

Points of Contact:
Trilinos — Mike Heroux maherou@sandia.gov
CPA — Vitus Leung vjleung@sandia.gov