Exploiting the Scalability and Power of FLUENT®:

The SGI Message Passing Toolkit on the SGI® Altix® High-Performance Computing Platform powered by the Intel® Itanium 2® Processor

Table of Contents

July 2005

When it first became prominent some 20 years ago, computational fluid dynamics (CFD) became a powerful engineering and design tool in the hands of the automotive and aerospace industries. It remains a force in both those fields; among the top 20 automotive companies, for example, CFD consumes about 25 percent of all supercomputing cycles. In airframe design, CFD is a bedrock technology for studies of aerodynamics, cabin ventilation, combustion, cooling, hydraulics, sanitation and other areas.

Fluid dynamics studies now affect engineering and design wherever the flow of gases or liquids is significant. For that reason, CFD has become one of the most dynamic and innovative sectors of the technical computing market segment. And because it requires large compute resources, CFD has helped drive the growth of supercomputing resources over recent years.

“The Navier-Stokes equations are the mathematical foundation of all CFD studies,” says Christian Tanasescu, SGI Director of Engineering, Compute-Intensive Applications. “CFD is the most scalable application segment in Computer-Aided Engineering and requires abundant computational resources”.

Parallel computing is of great importance to organizations that rely on CFD. Primarily, it enables users to spread jobs over multiple CPUs to get answers more quickly. But it also supports the trend in CFD to run optimized analyses, which can deliver more geometric detail and more realistic geometries. It also enables work with larger detailed models – for example, a complete aircraft – that will enable them to fine-tune detailed geometry for performance.

Another trend that is driving larger CFD models is the desire to include a very detailed description of flow physics. In the past, engineers might have been content to look at time-averaged approximations of flow. Today they want to examine time-varying flows – for example, the evolution of turbulent eddies in the wake around a vehicle or the transient ignition process in a flame.

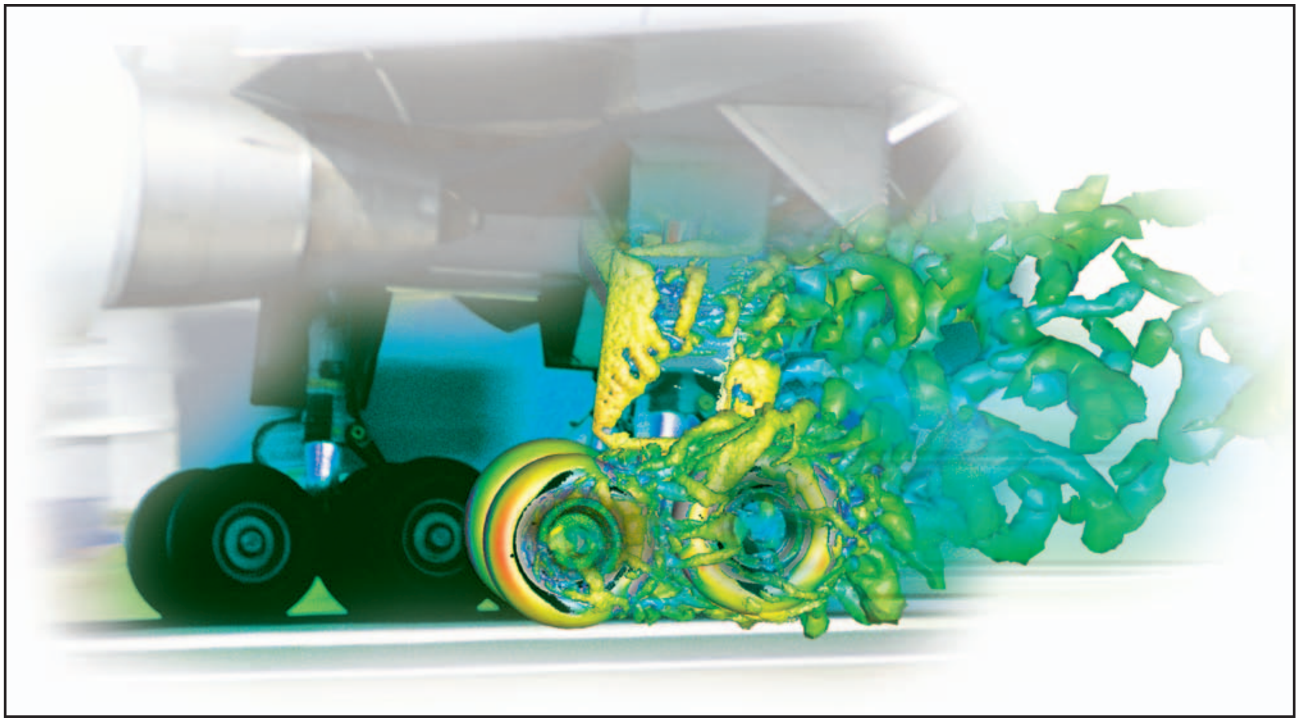

Vortical structures visualized using isosurfaces of the second invariant of the deformation tensor, colored by velocity magnitude.

Finally, CFD is no longer viewed as a check at the end of the design cycle, but as an up-front analysis that enables the comparison of designs from the standpoint of flow physics. All of these trends point to the need for more and faster CPUs, memory and communication.

Satisfying the Special Requirements of FLUENT Users

Fluent is the acknowledged worldwide leader in the CFD market segment. It is widely diversified in terms of its penetration of global market segments and in terms of the industries it serves.

“Traditionally, computer-aided engineering tools like FLUENT have been used heavily in the automotive and aerospace industries,” says Barbara Hutchings, Director of Strategic Partnerships at Fluent. “Those are definitely strong markets for us, but they represent only a portion of our business worldwide. Fluent is also active in the electronics industry, for example, for electronics cooling and semiconductor manufacturing; in the power generation industry, for equipment design and design of pollution abatement systems; and in the chemical and process industry for product and process innovation.”

Because flow dynamics is inherently compute-intensive, CFD problems tend to be large, and they are getting larger. As we mentioned above, it is often desirable to simulate complete products such as cars or airplanes, but the models can become too big to be useful engineering because the time to solution on some computing systems becomes prohibitive.

Scalability is the solution to this problem. For that reason, FLUENT is very scalable, scaling to hundreds or even thousands of CPUs. “Fluent has been a leader in scalable solutions for CFD, and our software was written from the ground up for high performance parallel computing,” says Hutchings. “It’s set up to run across a large number of nodes using parallel processing.”

However, scaling to large processor counts demands faster CPUs, low-latency memory access, and low-latency interconnect communication. Furthermore, the complex geometries of very large CFD models require very high memory bandwidth.

These requirements – massive scalability, fast CPUs, a low-latency architecture, and high memory bandwidth – describe the characteristics of the SGI® Altix platform. The Altix architec-

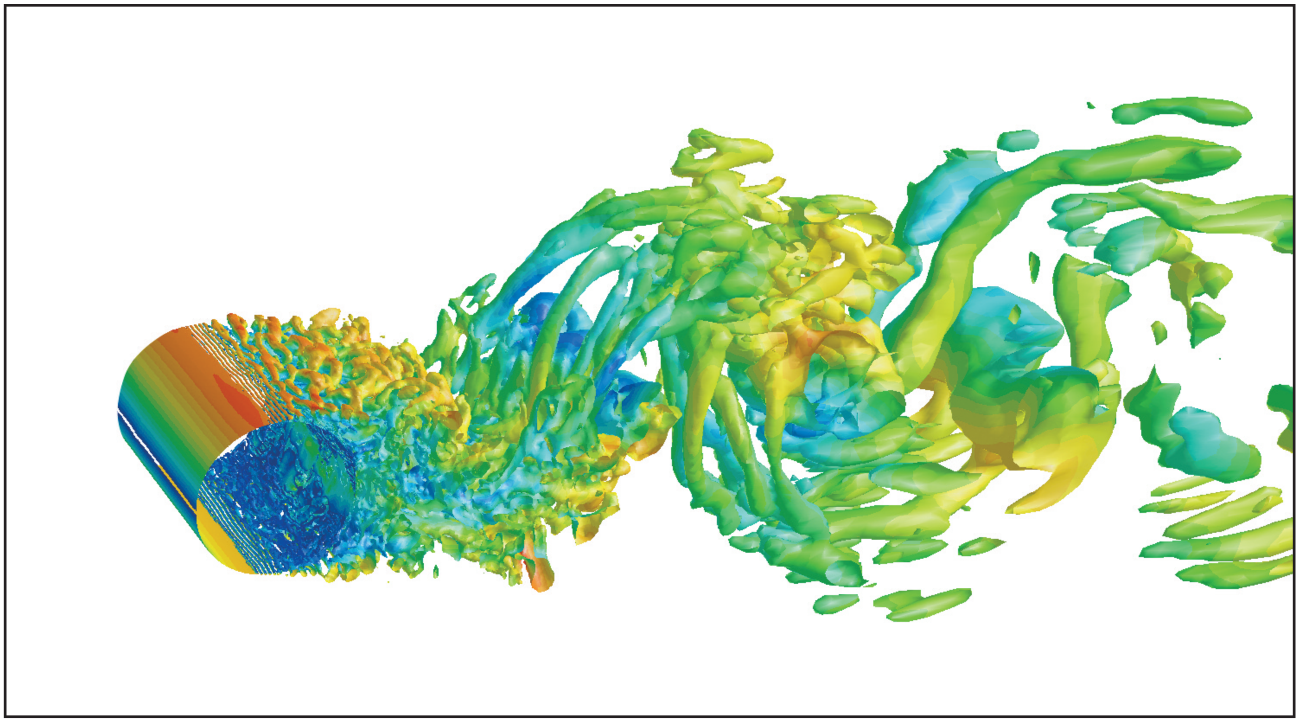

The flow structure behind a circular cylinder predicted by FLUENT's LES model at a Reynolds number of 1.4 x 105; vortical structures are colored by velocity magnitude.

ture based on the Intel® Itanium 2® processor gives FLUENT users the highest-performance processors, the fastest and most scalable system memory, and the lowest-latency interconnect available today, that enabled Altix to break the 1 microseconds barrier of MPI latency.

For these reasons many customers have chosen to run FLUENT on the Intel Itanium 2 processor on the SGI Altix shared-memory platform. “Many of the requests we see for Itanium 2 are Altix-based.” says Prasad Alavilli, Lead Development Engineer, Fluent. “Many of our customers are looking at large Altix systems.”

SGI, Intel, and Fluent joined forces in 2002 to optimize the scalability and performance of FLUENT on the Altix platform. They were joined by NEC, who shares an interest in optimizing Fluent for the Itanium 2 processor.

“Many of the requests we see for Itanium 2 are Altix-based. Many of our customers are looking at large Altix systems.”

Prasad Alavilli, Lead Development Engineer, Fluent.

SGI MPT Opens Up FLUENT Scalability and Performance on the Altix Platform

As part of its collaborative effort with Fluent, SGI mounted an engineering initiative to adapt MPT (Message Passing Toolkit), the SGI implementation of the MPI (Message Passing Interface) standard used by FLUENT for parallel communication. MPT, first released in 1993, has a long record of excellent performance with high processor counts on the SGI shared-memory NUMAflex architecture. Its latest release has been diligently tuned to the Altix platform and the Intel Itanium 2 processor architecture. Using MPT with FLUENT, SGI and Fluent were able to significantly increase the scalability of FLUENT on the Altix platform.

These optimized MPT libraries were implemented in Fluent’s recent 6.2 release, enabling FLUENT users to scale a single system image (SSI) to 512 processors on the Altix platform, or to scale to 2048 processors with global addressable memory. MPT also supports up to 8 TB of shared memory. MPT includes tools that sustain processor-memory affinity. This assures

SGI Altix family of servers

FLUENT users that computational threads remain linked to the specific processors to which they were initially assigned, significantly reducing latency costs when high processor counts are used. This and other features are incorporated into the FLUENT runtime script.

"The new release of MPT also reduces communications latency compared with MPICH, which is the MPI software used by FLUENT on many systems." says Mark Kremenetsky, applications engineer, SGI. "The FLUENT solvers send many short messages, which take very little time to transmit. But the time required to establish communication between processors for each message can be significant. MPT reduces this latency significantly to improve the performance of FLUENT in high-processor-count situations."

Overall, MPT increases the speed and performance of FLUENT on the Altix platform by up to 20 percent compared with MPICH. This translates to significantly faster CFD solves and a correspondingly shorter time-to-market.

"Current users of FLUENT on Altix will take advantage of FLUENT 6.2 for at least four reasons," says Tanasescu. "First, they can exploit the scalability of FLUENT. Second, they can benefit the industry-leading communication latency of the Altix platform – less than a microsecond. Third, they can take advantage of memory placement on the NUMA architecture. And fourth, they will now able to run the same binary on a single system image or an Altix Linux cluster."

"We found that running FLUENT on the Altix platform with SGI's MPT immediately gave us greatly improved scalability and performance for our FLUENT solves. We were able to use these major improvements to reduce our simulation times dramatically. We're very pleased that the partnership between Fluent, SGI, and Intel has given Skoda an excellent solution for product development."

Frantisek Paur, Ph.D.
CFD Simulations leader, Car Development, Skoda Auto.

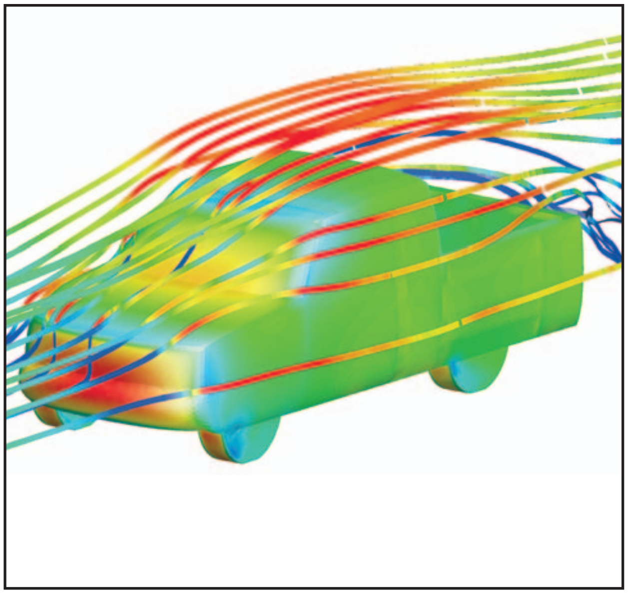

Pressure contours on the surface of the pickup and pathlines, colored by velocity magnitude showing the flow around a vehicle.

Intel: Performance Tuning for the Itanium 2 Processor Architecture

Fluent and Intel's compiler group have worked closely throughout this collaboration to dramatically improve the performance of FLUENT on the Itanium 2 processor. This work has involved modifications to Itanium 2 compliers and rewriting FLUENT kernels.

The Intel Itanium 2 processor compiler team began its optimization work on the Itanium 2 7.0 compiler release, tuning the compilers and modifying FLUENT kernels to improve software pipelining. The inclusion of the 7.1 compiler gave users of FLUENT 6.1, released in 2004, a performance increase of approximately 40 percent over the 7.0 compiler. The team's ongoing optimization of FLUENT continues to produce performance increases. The 8.0 compiler included in FLUENT 6.2 increases performance by an additional 12-14 percent, and more performance increases from the compiler team will be reflected in future releases of FLUENT.

The Intel Itanium 2 processor gives the Altix platform superior floating point performance. Its Explicitly Parallel Instruction Computation (EPIC) architecture delivers up to four floating point operations per cycle and executes FLUENT instructions at an unprecedented level of parallelism and efficiency.

"The EPIC architecture is ideally suited for parallel codes such as FLUENT, because its execution streams operate in parallel," says Michael Rex, Developer Relations Manager, Intel. "We continue to work closely with Fluent to give its customers maximum performance on the Altix platform."

The explicitly parallel architecture of the Itanium 2 processor, including its massive execution resources, gives it the ability to execute multiple instructions simultaneously, enabling more users to access compute resources. Calculations and data analysis are enhanced, and large CAD and CAE models can be moved and stored more quickly. The Itanium processor's internal resources include a massive on-die cache, 128 floating-point registers, and memory management features that enable more efficient management of large data sets.

"The EPIC architecture is ideally suited for parallel codes such as FLUENT, because its execution streams operate in parallel. We continue to work closely with Fluent to give its customers maximum performance on the Altix platform."

Michael Rex, Developer Relations Manager, Intel.

The enterprise-hardened design of the Itanium 2 processor includes features that provide high reliability, availability, serviceability, and manageability. The Advanced Machine Check Architecture provides extensive error management in hardware, firmware, and operating systems to minimize data loss, data corruption, and downtime.

Intel's commitment to this flagship server processor is demonstrated by its strong roadmap. Multiple Intel design teams currently have six Itanium 2 processors in simultaneous development with the goal of increasingly higher performance at reduced cost. By 2007, Intel plans to ship Itanium 2 processors with twice the performance levels of the Intel® Xeon™ processor line at the same price.

FLUENT 6.2 on the SGI® Altix® Server Platform

Linux-based SGI Altix systems give FLUENT 6.2 users important advantages that reduce time to solution:

Massive scalability. FLUENT 6.2 can be scaled to a 512-processor single system image, or to 2048 processors with global addressable memory.

Very large shared memory. The Altix architecture supports up to 32 TB of global addressable memory.

The industry's fastest interconnect. The Altix platform's fast NUMA interconnect supports the highest bandwidth and fastest latencies available.

Summing Up: Great Scalability and Performance for Users of FLUENT 6.2

The implementation of SGI MPT and the collaborative performance tuning achievements of Intel and Fluent make FLUENT 6.2 a far more powerful analysis tool, accelerating design processes and dramatically shortening time to market.

"Because of this collaboration we have improved the performance of FLUENT on the SGI Altix platform very significantly," says Alavilli. "Performance has almost doubled since our initial release on Itanium 2 processors. This is an ongoing collaboration, and we will continue to enhance the performance of FLUENT on this architecture."

Performance Data

For the latest performance data for FLUENT on Altix please visit the Fluent website at
http://www.fluent.com/software/fluent/fl5bench/fullres.htm

3807 [July 2005] J14923