Understanding the Behavior of

Computational Fluid Dynamics

Applications on Dell PowerEdge 3250 Servers

Computational fluid dynamics (CFD) applications typically run workloads that can benefit from the parallel-processing capabilities of high-performance computing (HPC) clusters. Understanding how the subsystem components of an HPC platform affect the performance characteristics of a CFD application can help designers optimize an HPC cluster for the application's specific requirements. This article describes a test conducted on Dell™ PowerEdge™ 3250 servers to determine how CPU and cache can affect two popular CFD applications.

BY BARIS GULER AND RIZWAN ALI

Computational fluid dynamics (CFD) is a sophisticated analysis technique that uses computers to simulate fluid flow. CFD enables a computational model representing a physical system to be built and studied. When fluid flow physics is applied to this virtual prototype, the CFD application outputs a prediction of the fluid dynamics. The software predicts not only the fluid flow behavior, but also the transfer of heat, mass, phase change, chemical reaction, mechanical movement, and stress or deformation of related solid structures.

Performing CFD analysis on a large data set—such as one that represents a moving car or a fighter jet—requires considerable CPU resources that can take a long time to process. To address this issue, some independent software vendors (ISVs) offering CFD applications have made their applications parallel and cluster-aware. By leveraging the parallel-processing capabilities of high-performance computing (HPC) clusters, CFD applications can enable large data sets that require heavy-duty CPU resources to run faster than they could run on a single server.

An HPC cluster can be created using standard, off-the-shelf components that work together as a supercomputer. HPC cluster designers have the flexibility to choose from multiple options for the base computing platform, subsystems, and interconnects. Platforms for HPC clusters are available in a variety of configurations, using different processor types and speeds, cache sizes, disk I/O, and memory subsystems. To create an optimal HPC cluster for use with a specific CFD application, the designer must understand the performance characteristics of the application itself and the effects of the various subsystem components on the application. This knowledge allows the designer to allocate resources to the proper components and to budget adequate financial resources to build an optimal system.

Computational fluid dynamics enables a computational model representing a physical system to be built and studied.

Test configuration	Processor speed	L3 cache size
A1	1.4 GHz	4.0 MB
A2	1.4 GHz	1.5 MB
A3	1.0 GHz	1.5 MB

Figure 1. Speed and cache size of Itanium 2 processors used in the test configurations

Studying the effects of cache size and processor speed

In January and February of 2004, a Dell HPC cluster team studied the performance patterns of two widely used CFD applications—the CD adapco Group STAR-CD[1] and Fluent FLUENT solvers.[2] The Dell study focused on the computational performance effects of processor cache and speed. Both STAR-CD and FLUENT applications allow users to perform CFD analysis around their particular models. For each of the CFD applications, the Dell test team used benchmarking data sets provided by the respective ISV.

For the computing platform, the test team used Dell PowerEdge 3250 servers, which are based on the Intel® Itanium® 2 processor. The Intel Itanium processor family consists of a wide range of processor types, each with a different clock speed or cache size—making the PowerEdge 3250 server an excellent choice for the study. Three Intel Itanium 2 processor types were used to demonstrate the effect of cache size and clock frequency on the applications.

Establishing the test environment

The test environment consisted of three Dell PowerEdge 3250 servers configured three ways (see Figure 1). The Itanium 2 processor–based PowerEdge 3250 server has a 128-bit-wide, 400 MHz frontside bus offering a memory bandwidth of 6.4 GB/sec. Across all three configurations, each PowerEdge 3250 server was populated with dual 64-bit Itanium 2 processors, 4 GB of RAM, and a single 36 GB Ultra320 hard drive.

The software environment consisted of the Red Hat® Enterprise Linux® AS 2.1 operating system; the Intel Itanium Compiler 7.1 for Itanium-based applications; the FLUENT 6.1.23 application with the ISV's accompanying benchmark data set; and the STAR-CD V3.150A.012 application with the ISV's accompanying benchmark data set. The FLUENT 6.1.23 application was obtained in binary form for the Itanium processor, whereas the STAR-CD application required some components to be compiled and linked with the rest of the precompiled object files.

Within this common environment, the changing variable was the type of Itanium 2 processor used to complete the benchmark testing. Three types of the Itanium 2 processor were used, as shown in Figure 1.

Configurations A1 and A2 were used to study the effect of the cache size, and configurations A2 and A3 were used to study the effect of the processor speed.

Comparing test configurations using FLUENT

The Fluent benchmarking suite consisted of nine data sets comprising three classes: small, medium, and large. Each of the data sets represents a real-world model. The benchmark names and descriptions are shown in Figure 2. Eight benchmarks were run on all three configurations—A1, A2, and A3. The FL5L3 benchmark was not used in the test because it was too large to run on a single processor.

The Fluent benchmarking suite outputs its results in the form of a rating, where the rating is the number of times that a particular benchmark can be run in a 24-hour period. A higher rating indicates better performance.

The effect of cache size on performance

The purpose of comparing the A1 and A2 test configurations was to isolate the size of the processor's on-die level 3 (L3) cache: A1 had 4.0 MB of L3 cache whereas A2 had 1.5 MB of L3 cache. Both configurations had the same 1.4 GHz processor speed and all other system components were identical. The rating on the A2 system was used as the reference point from which to calculate the percentage improvement for a given benchmark when the L3 cache was increased from 1.5 MB to 4.0 MB. Figure 3 shows the test results. The *y*-axis shows the percentage improvement gained by increasing from 1.5 MB of L3 cache to 4.0 MB of L3 cache. The *x*-axis shows the different benchmarks that were run.

The results shown in Figure 3 indicate that the increase in cache size can help improve performance but not significantly.

Class	Benchmark	Cells	Mesh	Description
Small	FL5S1	32,000	Hexahedral	Turbulent flow in a bend
	FL5S2	32,000	Hexahedral	Turbulent flow in a bend
	FL5S3	89,856	Hexahedral	Flow in a compressor, rotor 37
Medium	FL5M1	155,188	Tetrahedral	Coal combustion in a boiler, with particle tracking
	FL5M2	242,782	Hybrid, hanging node	Turbulent flow in an engine valveport
	FL5M3	353,800	Hexahedral	Combustion in a high-velocity burner
Large	FL5L1	847,746	Hexahedral	Transonic flow around a fighter
	FL5L2	3,618,080	Hybrid	External aerodynamics around a car body
	FL5L3	9,792,512	Hexahedral	Turbulent flow in a transition duct

Figure 2. Fluent benchmarks

[1] For more information about STAR-CD benchmarks and data sets, visit http://www.cd-adapco.com/products/star_overview.htm.

[2] For more information about Fluent benchmarks and data sets, visit http://www.fluent.com/software/fluent.

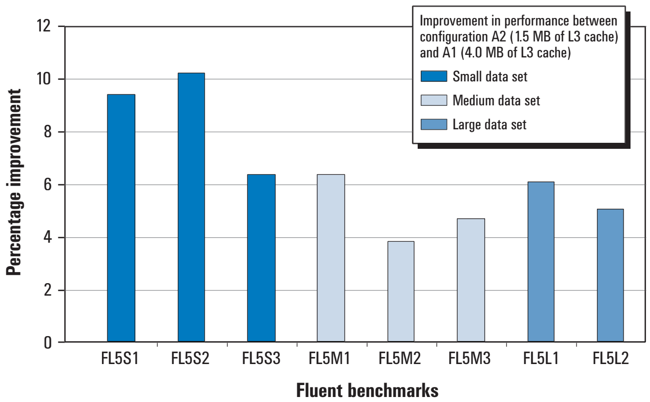

Figure 3. Effect of cache size on Fluent benchmark performance

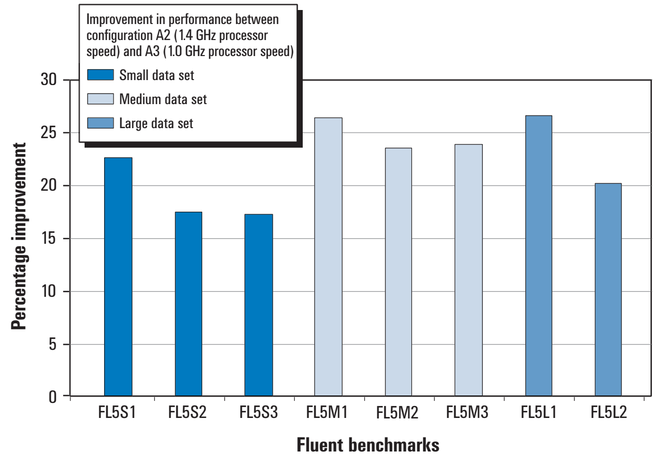

Figure 4. Effect of processor speed on Fluent benchmark performance

Although the L3 cache size more than doubled, the performance improvement stayed at or below approximately 10 percent. The performance improvement in the small and medium data sets ranged from nearly 4 percent to about 10 percent. In particular, the two smallest data sets showed a greater improvement compared to the large data sets.

The effect of processor speed on performance

In the same test bed, data was accumulated for which the impact of processor speed was determined by comparing configuration A2 (which had a 1.4 GHz processor) and A3 (which had a 1.0 GHz processor). All other system components in the test configuration were identical.

The results presented in Figure 4 indicate that the speed of the processor can have an impact on the performance of the FLUENT application. In this study, when the processor speed was increased from 1.0 GHz to 1.4 GHz—a 40 percent increase in processor speed—the application achieved a performance increase ranging from 17 percent to 26 percent. Unlike the cache study, the processor speed comparison achieved a greater performance improvement in the medium and large data sets, as opposed to the small data set.

Comparing test configurations using STAR-CD

The STAR-CD benchmarking suite used in this study consisted of six data sets, including three small test cases, two medium test cases, and one large test case. These test cases were selected to be industry representative. The benchmark names and brief descriptions are shown in Figure 5.

The effect of cache size on runtime

For each test case, the execution times on two different configurations, A1 and A2, were used to identify the impact of L3 cache size on runtime. All other system components in the test configuration were identical. The results are shown in Figure 6. The execution time on the A2 system was used as the reference point from which to calculate the percentage improvement (y-axis) for a given benchmark when the L3 cache was increased from 1.5 MB to 4.0 MB. The x-axis shows the different benchmarks that were run.

Class	Benchmark	Cells	Mesh	Description
Small	Engine block	156,739	Hexahedral	Engine cooling in an automobile engine block
	S-bend	250,000	Hexahedral	Airflow in an s-bend
	Car interior	261,838	Hexahedral	Internal airflow in a car interior
Medium	Aerodynamic body	499,236	Tetrahedral	3-D flow over an external aerodynamic body
	Concept car	1,900,308	Tetrahedral	External flow around an ASMO concept car
Large	A-class	5,914,426	Hybrid	Turbulent flow around an A-class car

Figure 5. STAR-CD test cases

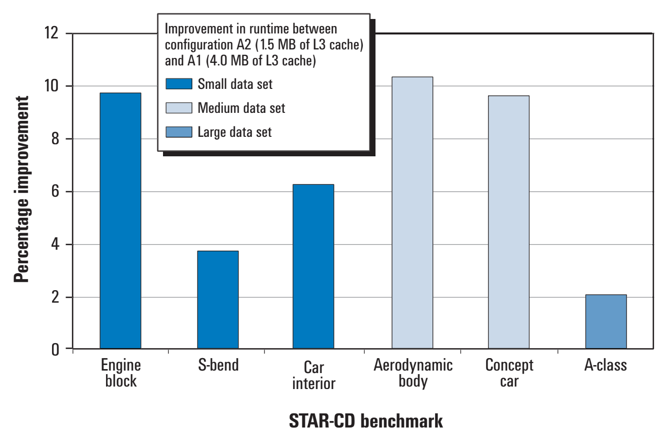

Figure 6. Effect of cache size on STAR-CD benchmark performance

Figure 6 shows that the increase in cache size can help boost performance, but the improvement is not significant, ranging from 2 percent to approximately 10 percent. Similar characteristics are observed for the Fluent benchmark results shown in Figure 3. The improvement in performance of the small data sets ranged from nearly 4 percent to nearly 10 percent—better than that of the large data set, which showed an improvement of approximately 2 percent. The medium data sets showed consistent improvements of approximately 10 percent.

The effect of processor speed on runtime

The execution times of each test case were measured on configurations A2 and A3. These configurations were used to identify the effect of processor speed on runtime. The execution time on the A3 system (configured with a 1.0 GHz processor) was used as the reference point from which the percentage improvement for a given benchmark was calculated when the processor speed was increased to 1.4 GHz (in the A2 configuration). Both A2 and A3 configurations had 1.5 MB of L3 cache, and all other system components in the test configuration were identical. The test results are shown in Figure 7. The *y*-axis represents the percentage improvement in execution time when the processor speed was increased 40 percent to 1.4 GHz from 1.0 GHz, and the *x*-axis shows the different benchmarks that were run.

The performance improvement from the processor speed increase was somewhat significant compared to that from the cache increase. The improvement ranged from 8 percent to approximately 14 percent, with the small data sets showing a higher degree of improvement than the medium and large data sets. Contrary to the observation from the cache study, even the large data set showed a proportionately larger percentage improvement in performance (nearly 12 percent). Despite the 40 percent increase in processor speed, the performance improvement when running the STAR-CD benchmarks in this study was approximately 12 percent. This result indicates that other system components—such as disk I/O, memory subsystem, and so on—might be bottlenecks.

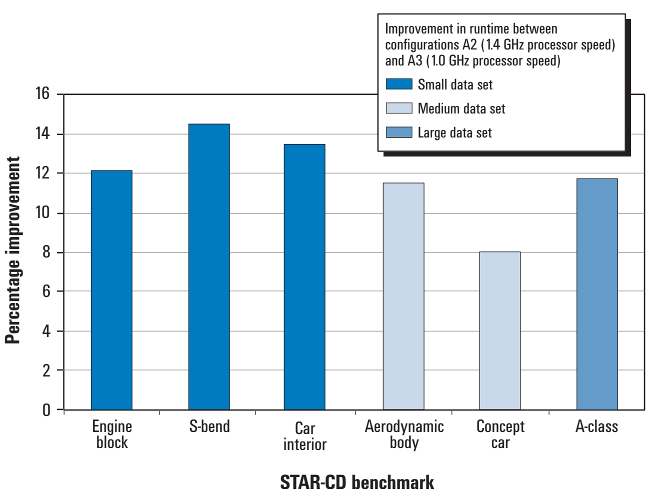

Figure 7. Effect of processor speed on STAR-CD benchmark performance

By leveraging the parallel-processing capabilities of high-performance computing clusters, CFD applications can enable large data sets that require heavy-duty CPU resources to run faster than they could run on a single server.

Improving cluster designs by understanding subsystem effects

The performance of a computer is dependent on the software and the data sets being used. Similarly, application performance using a specific data set will be affected by the configuration of several critical hardware components, such as CPU, cache, memory, and disk I/O. In this article, CPU and cache were isolated to help understand the effect of these components on two well-known CFD applications. The study described in this article showed that both CFD applications achieved a greater performance improvement from an increase in the processor clock speed than from an increase in the processor cache size. Such performance comparisons can be valuable in helping designers to create a balanced cluster around a specific application.

Baris Guler is a systems engineer and advisor in the Scalable Systems Group at Dell. His current research interests are parallel processing, diskless HPC clusters, performance benchmarking, reservoir engineering and simulation, and numerical methods. Baris has a B.S. in Petroleum and Natural Gas Engineering (PNGE) from the Middle East Technical University in Turkey, and an M.S. in PNGE from Pennsylvania State University. He is currently a Ph.D. candidate in Petroleum and Geosystems Engineering at The University of Texas at Austin.

Rizwan Ali is a systems engineer in the Scalable Systems Group at Dell. His current research interests are performance benchmarking, cluster architecture, parallel applications, and high-speed interconnects. Rizwan has a B.S. in Electrical Engineering from the University of Minnesota.

FOR MORE INFORMATION

Dell PowerEdge 3250:
http://www1.us.dell.com/content/products/productdetails.aspx/pedge_3250?c=us&cs=555&l=en&s=biz

STAR-CD V3150 benchmarks:
http://www.cd-adapco.com/support/bench/315/index.htm

Fluent benchmarks:
http://www.fluent.com/software/fluent/fl5bench/intro.htm