Vetrei

Ensuring Your Success

Vetrei, Inc. 100 Woodcrest Lane nbury, CT 06810 info@vetrei.com

www.vetrei.com 203 205-0705

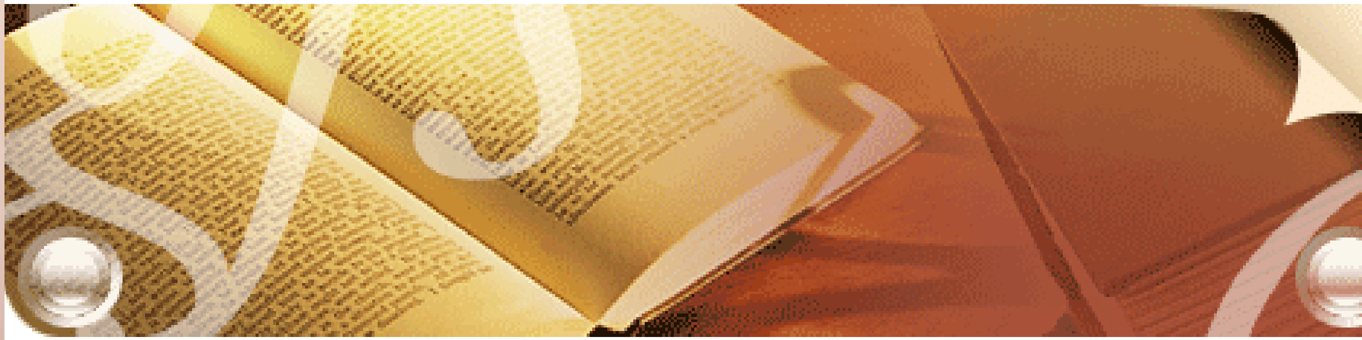

WHITE PAPER

Breakthrough Advantage in Computational Fluid Dynamics with the IBM® System Blue Gene® Solution

Sponsored by IBM

Srini Chari, Ph.D., MBA

September 2006

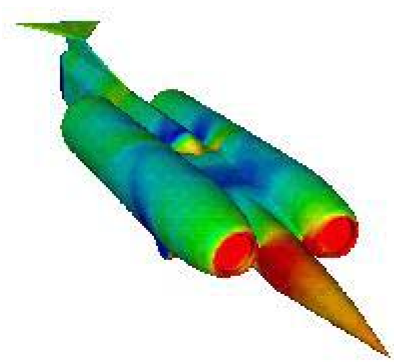

Executive Summary

Over the last 40 years, the use of Computational Fluid Dynamics (CFD) has increased by several orders of magnitude in industry and research laboratories, largely due to the impressive advances in computing architectures as well as in the algorithmic techniques created to exploit these architectures. Each major innovation in the computing industry has directly enabled CFD practitioners to solve more realistic and complex engineering design and simulation problems, resulting in better products faster. The investment by the community in adapting CFD applications to take advantage of newer computing architectures has paid off handsomely.

The IBM Blue Gene is the first system in a generation of innovative, ultrascalable architectures that will enable CFD engineers to make significant improvements in the understanding and the solution of some of the most complex problems in engineering design. For the first time, it will be possible to solve complex problems in turbulence that are interdisciplinary, multi-scale, and with multi-body interactions on very complex geometries. These problems require very high-resolution models to gain profound engineering insights that were previously not possible. However, this entails a continuing investment by the CFD developers, mathematicians, and computer scientists to develop new algorithms and applications for ultrascalable parallel computing environments. This investment will be protected as newer Blue Gene systems become available. As before, the payoff will far outweigh this investment.

IBM has expanded its five year Blue Gene collaboration with the Lawrence Livermore Laboratory to the CFD community. This collaborative investment has produced impressive early results for some of the most challenging CFD problems in industry and research institutions. The scalability and performance obtained from these simulations on the IBM Blue Gene are unsurpassed yet affordable, and easily accesible. More importantly, the solution of previously intractable CFD problems has resulted in breakthrough engineering insights. The direct numerical simulation of turbulence for realistic engineering configurations is now – for the first time – plausible.

Breakthrough Advantage in Computational Fluid Dynamics

Unsurpassed Performance for Ultrascale Computational Fluid Dynamics with the IBM® System Blue Gene® Solution

Table of Contents

List of Figures

Breakthrough Advantage for Computational Fluid Dynamics (CFD)

Unsurpassed Performance for Ultrascale Computational Fluid Dynamics with the IBM® System Blue Gene® Solution

1. Introduction

Computational Fluid Dynamics (CFD) has revolutionized the process of design and development in the aerospace, automotive, process and chemical, and other industries. CFD is routinely used in conjunction with experimental techniques; e.g., wind tunnel and flight tests, noise and vibration mitigation studies for automobiles, and engine performance for vehicles to significantly reduce product development cycle time, costs, and improve functional performance.

CFD Combines Engineering, Mathematics, and Advanced Computing

In the last four decades, developments in novel algorithms, the increased and affordable access to high performance computers, and the availability of a broad range of commercial and government-funded applications (i.e., NASA), have together substantially contributed to the widespread acceptance of CFD in the manufacturing and process industries and in government laboratories.

Effective CFD requires a fusion of interdisciplinary advances in engineering models, mathematical algorithms, information technology architectures, and disciplined software engineering. Moreover, the interdisciplinary nature of today's large scale CFD problems (e.g., reacting turbulent flows, aerothermodynamics, propulsion, flight controls, vehicle climate modelling, etc.) requires the balanced use of computing capability for grand challenge simulations coupled with capacity computing capability for production simulations.

Furthermore, a combination of high performance computing systems, massive storage systems, visualization and advanced instrumentation, applications and middleware, all connected by high-speed networks is needed for today's CFD infrastructure.

Breakthrough, High Resolution CFD Requires Parallel Ultrascale Computing

CFD involves tackling a wide range of complexity in modeling flow physics in or around complex shapes and geometries. The equation describing the flow physics are a coupled system of non-linear, partial differential equations. These problems are often very large, tightly-coupled, and multi-scale.

The need to transcend Moore's law through parallel computing has given rise to clusters, grids, and scalable parallel systems in recent years. Parallel solution approaches based on parallel mesh-decomposition techniques are required to exploit these parallel architectures for large scale CFD problems.

While standard clusters and grids can address a sub-class of these problems, an entire range of complex, interdisciplinary, multi-scale, turbulent flow problems require ultrascalable architectures. These systems are large, tightly-coupled computers with high bandwidth and low latency interconnects with an optimized message-passing library, such as MPI (Message Passing Interface). The IBM Blue Gene system is an example of an ultrascale computer that has been very effective in addressing these classes of challenging CFD problems.

Blue Gene: Unsurpassed, Ultrascale, Affordable Computing

Blue Gene is the first in a new generation of novel parallel supercomputers and ranks as the most powerful supercomputer available. Dramatic reductions in power consumption, cost, and space requirements are achieved through the use of innovative technologies in low-power processors; embedded DRAM; system-on-a-chip; advanced power, packaging and cooling; special interconnects delivering very low latency and high bandwidths; and scalable systems management. This performance, scalability, flexibility, and innovative design enables the solution of a wide range of complex CFD problems today and in the future.

IBM Collaborates with CFD Application Developers to Foster Breakthrough Innovation

CFD applications that are mapped, migrated, and optimized for the Blue Gene architecture will benefit greatly from ultrascalability and extreme performance. This investment in application enablement will permit higher degrees of performance and scaling in newer generation of systems and technologies based on the Blue Gene architecture. Worldwide, IBM is working with CFD application developers to migrate and optimize their applications on the Blue Gene to solve challenging problems that are multi-scale, multi-regime, and interdisciplinary, in order to enable breakthrough innovation in industry and government. Experience to date shows that large-scale (problems with over 20 million unknowns), multi-phase flows with detailed turbulence models scale and perform very well on the Blue Gene.

2. Overview of Large Scale Computational Fluid Dynamics

The Navier-Stokes equations provide the most general mathematical description of commercial interest. Direct solutions to these equations in practical vehicle configurations in typical operating conditions are still beyond the capability of today's computing systems. These flows include chaotic, turbulent flows across various scales and regimes. Computations for all scales would require exascale (million teraflops) computers. Fortunately, solutions to simplified and tractable forms of the Navier-Stokes equations are still of great engineering value and can be addressed by today's novel, ultrascale architectures.

CFD Applies to a Wide-Range of Problems in Research and Industry

The direct solution of the Navier-Stokes equation in all regimes and scales is intractable on today's computing environments. Over the last few decades, these equations have been simplied and solved for many engineering problems ranging from linearized potential flows to more complex general-geometry, non-linear turbulent Navier-Stokes flows. The solution approaches and algorithms have varied depending upon the geometry, physics, operating regimes and scales of interest, and availability of suitable computer architectures. The following table provides a summary classification of CFD applications and solution approaches. See the Appendix created from materials from the Bibliography for more details.

Equations	**Mesh Types**	**Some Solution Approaches**	**Turbulence Models**
- Euler - Navier-Stokes - Reaction-Diffusion - Multi-Phase	- Structured - Block-Structured - Unstructured - Overlapped - Adaptive - Meshless	- Finite Difference - Finite Volume - Finite Element - Spectral - Explicit - Implicit	- RANS - DES - LES - DNS

Figure 1: Classification of CFD Applications

The Use of CFD Continues to Grow in Design and Development

More upfront CFD simulations are routine in the preliminary design and development phase. Parametric analyses over thousands of operating scenarios on very large meshes with 50 million or more vertices are common in the aerospace and automotive industries. Interaction aerodynamics, as opposed to component aerodynamics, is increasingly the norm as are interdisciplinary analyses such as design optimization, flight control, and interior acoustics.

Innovative Solution Approaches and a Comprehensive Computing Infrastructure Expands the Role and Value of CFD

The increased use of CFD in industry and the complexity of the simulations being performed, requires a combination of high performance computing systems, massive storage systems, visualization and advanced instrumentation, applications and middleware, all connected by high-speed networks. The computing challenges continue to push the envelope in high performance computing systems, algorithms, turbulence models, and geometric detail and scale. These CFD innovations will further expand the role and value of CFD in industry and government.

3. Computational Challenges in Large Scale CFD

The need to transcend Moore's law through parallel computing has given rise to clusters, grids, and scalable parallel systems in recent years. Parallel solution approaches based on parallel mesh-decomposition techniques are required to exploit these parallel architectures for large scale CFD problems. Additionally, careful adaptation and tuning of these parallel algorithms on ultrascalable platforms is needed to get the benefits of very large scalability and performance. While standard clusters and grids can address a sub-class of these problems, an entire range of complex, interdisciplinary, multi-scale, turbulent flow problems require ultrascalable architectures.

The Need to Transcend Moore's Law through Parallel Computing

To process massive amounts of data, which provides valuable engineering insight; applications require massive horsepower, fueling the demand for clusters and scalable parallel computing. These systems can be expanded with standard CPUs to keep pace with the increased processing needs driven by the dramatic growth of data and higher accuracy, especially for high resolution multi-scale CFD. The growth in the resulting computing demands exceeds Moore's Law. A scalable compute infrastructure is the only answer. Additionally, the price/performance ratios of these systems must make teraflops and eventually petaflops of processing power available at a fraction of the cost of a traditional supercomputer.

Novel Domain Decomposition or Mesh Partitioning Schemes are Needed for CFD Applications to Scale

In general, many CFD applications can be adapted and optimized to scalable parallel architectures for large-scale simulations that require very high degrees of spatial and temporal resolution and accuracy. This requires the use of domain decomposition or mesh partitioning techniques coupled with the solution approaches. Many effective domain decomposition techniques have been devised over the last few decades. These approaches also work for complex geometries with unstructured meshes. Powerful partitioning tools like METIS, CHACO, and JOSTLE (described in greater detail in the Appendix) used in the pre-processing phase have enabled many implicit CFD applications with parallel solvers to scale effectively.

Large Scale CFD Applications Require Novel Parallel Algorithms and Systems

Very high performance computing is almost always accomplished through parallelism. However, obtaining parallel computing capabilities is difficult and complex because most practical applications don't multithread beyond a few processes. In order to scale further, parallelism must be at a very high coarse grain level. Novel algorithmic approaches in CFD based on domain or mesh decomposition strategies allow users to obtain the maximum advantage and scalability from parallel machines with large numbers of inter-connected processors. Domain decomposition techniques when coupled with parallel variants of multi-grid, explicit, and overrelaxation methods further accelerate complex simulations.

Careful System Specific Tuning is Needed to Further Boost Performance and Scalability on Parallel Platforms

Parallel application development and system performance for large-scale CFD applications also depend on the single processor performance, the communication subsystem performance, I/O performance, and development tools for programming, debugging, and resource management. A significant improvement in performance and scalability can be obtained for CFD applications that are tuned and optimized for the specific parallel architecture. This often requires a combination of deep CFD algorithmic and parallel computing skills.

Clusters and Grids are Often Inadequate for High Resolution Multi-scale CFD

As mentioned before, an entire class of high resolution, multi-scale, interdisciplinary CFD problems; e.g., chemically-reacting turbulent flows on full scale vehicle geometries, requires scalable parallel systems with very low latency and high bandwidths between the several thousands of processors for deriving engineering insight.

In order to solve these growing compute requirements for complex CFD, individual clusters are starting to become insufficient. Standard clusters with thousands of processors are expensive to build and operate. The cost associated with providing support and maintenance grows exponentially. Also, management of such diverse collections (grids) of resources is difficult, and effective software solutions that can scale are only now beginning to appear in the market. Furthermore, the electrical power consumption and the physical facilities required to operate such large clusters are prohibitively expensive. These limitations of clusters and grids can be overcome by ultrascale parallel platforms such as the IBM Blue Gene.

4. The Blue Gene Solution

The IBM® System Blue Gene® Solution is the result of an IBM super-computing project which began in collaboration with Lawrence Livermore National Laboratory over five years ago. It was dedicated to building a new and innovative family of supercomputers optimized for bandwidth, scalability and the ability to handle large amounts of data while consuming a fraction of the power and floor space required by today’s high performance systems. Blue Gene’s unique design allows dense packaging of processors, memory and interconnects, and offers leadership efficiency in floor space and power consumption. The level of performance provided by the Blue Gene system can enable a tremendous increase in the scale of simulations beyond what is possible with other supercomputers.

An Ultrascalable, Environmentally Efficient Architecture

The Blue Gene system is built with a very large number of compute nodes, each of which has a relatively modest clock rate contributing to both low power consumption and low cost. It utilizes IBM PowerPC® embedded processors, embedded DRAM and system-on-a-chip techniques that allow for integration of all system functions including compute processor, communications processor, three cache levels, and multiple high speed interconnection networks with sophisticated routing onto a single chip. Because of a relatively modest processor cycle time, the memory is close, in terms of cycles, to the processor. This is also advantageous for power consumption and enables construction of dense packages in which 1024 dual-processor compute nodes can be placed within a single rack.

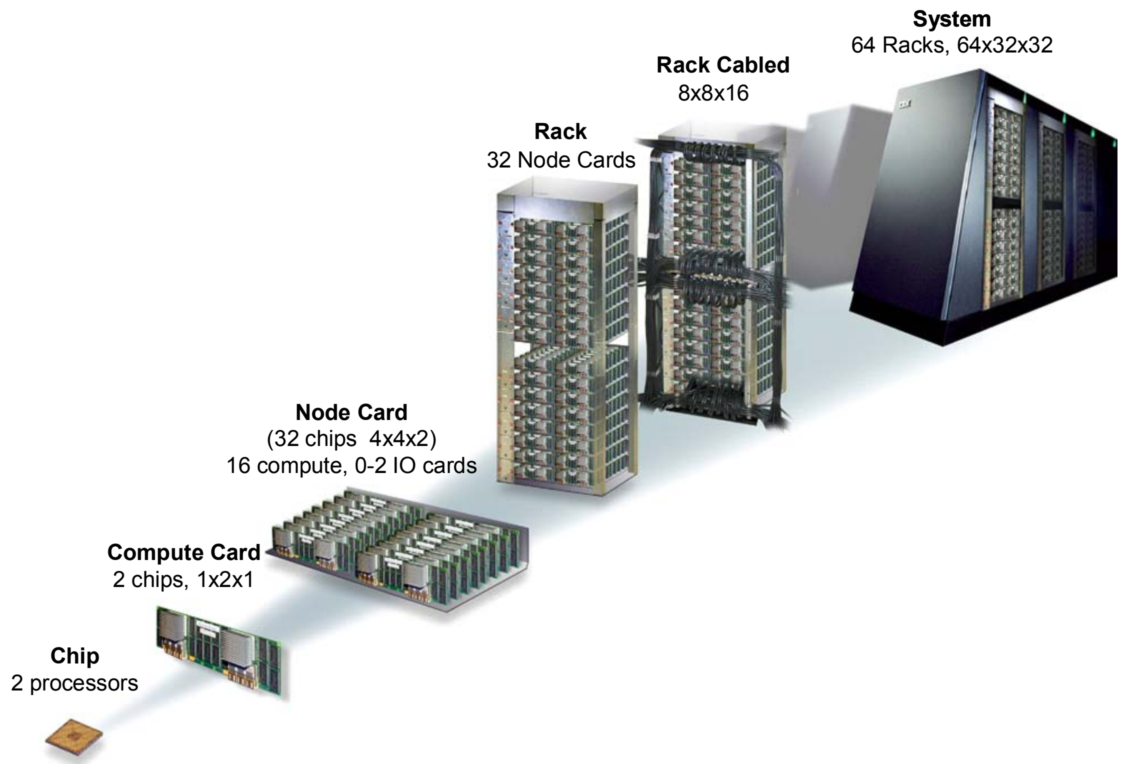

Figure 2: The Blue Gene Architecture

The nodes are interconnected through five networks: a 3-dimensional torus network for point-to-point messaging between compute nodes, a global collective network for collective operations over the entire application, a global barrier and interrupt network, and two gigabit Ethernet networks for machine control and for connection to other systems. The torus network is particularly effective for applications with locality of communication. And the global collective network is useful for speeding up MPI collective communications constructs. Large-scale CFD applications that use domain decomposition approaches benefit significantly from this local-global capability of the torus network. The computations done in each sub-domain are kept local from one fine grain iteration to the next and can be done in parallel across all of the sub-domains. After several local fine-grain iterations, the computations are synchronized across the entire domain using the collective communications constructs before proceeding further with more fine-grain, local iterations. This local-global approach accelerates the solution convergence for a large class of non-linear CFD applications and scales very well on the IBM Blue Gene.

A Familiar, Optimized Software Environment

Three fundamental principles were followed when the system software was designed for the Blue Gene system: simplicity, performance and familiarity. Driving toward simplicity in the software design has allowed development of software that takes advantage of hardware features to deliver high performance without compromising stability and security. And by creating a programming and administration environment based on familiar programming languages, libraries, job management tools and parallel file systems, CFD application developers benefit from the innovative design elements of the Blue Gene system without facing a steep learning curve.

Integrated with High Performance Visualization, File Systems, and Middleware

As mentioned earlier, in addition to high performance computing systems, massive file/storage systems, visualization and advanced instrumentation, middleware, all connected by high-speed networks are also needed for today's CFD infrastructure.

With IBM's Deep Computing Visualization, high-end graphical images are created in two visualization modes:

- Scalable Visual Networking (SVN)
- Remote Visual Networking (RVN)

SVN supports multiple high-resolution monitors or projectors for immersive, stereo visualization. RVN distributes graphical images to remote (collaborative) client stations. These features help derive more accurate engineering insights from the higher resolution of the complex data obtained from CFD simulations on the Blue Gene.

The following High Performance Computing cluster software is now available on Blue Gene: the Engineering and Scientific Subroutine Library (ESSL) for Linux on POWER, General Parallel File System (GPFS) for Linux on POWER, and LoadLeveler for Linux on POWER. ESSL provides over 150 math subroutines that have been specifically tuned for performance on Blue Gene. GPFS is the top performing cluster-wide file system for Blue Gene, providing superior scalability and high reliability. LoadLeveler is a job scheduler designed to maximize resource utilization and job throughput to get the most out of the available resources. This combination of middleware enables the optimization and scaling of the Blue Gene resources (processors and storage) for several, concurrent CFD simulations typical in parametric or design optimization studies.

Blue Gene System Details at a Glance

The Blue Gene system is accompanied by a product roadmap that will deliver petaflop performance. Available from 1-64 racks, the salient system details and benefits for Blue Gene – the first commercial system - are summarized in the following table.

Attribute	Details	Benefits
Processor	PowerPC 440 700MHz; two per node	Low power allows dense packaging; better processor-memory balance
Memory per node	512 MB SDRAM-DDR (Model 0203-700) 1 GB SDRAM-DDR (Model 0203-900)	
Networks	1) 3D Torus - 175MB/sec in each direction 2) Collective Network – 350MB/sec; 1.5 usec latency 3) Global Barrier/Interrupt 4) Gigabit Ethernet (I/O & connectivity) 5) Control (system boot, debug, monitoring)	Special networks speed up internode communications; designed for MPI programming constructs; improve systems management
Compute Nodes	Dual processor; 1024 per rack	Double FPU improves performance
I/O Nodes	Dual processor; 16-128 per rack	Facilitates job launch and I/O, raising efficiency of compute nodes
Operating Systems	Compute Node – Lightweight proprietary kernel I/O Node – Mini Control Program Front End and Service Nodes – SuSE SLES 9 Linux	Kernel tailored to processor design; industry-standard distribution on front-end and service nodes preserves familiarity to end users and administrators
Performance	Peak performance per rack – 5.73 TFlops Linpack performance per rack – 4.71 TFlops	Highest available performance benefits capability customers
Power	27.6 kW power consumption per rack (maximum) 7 kW power consumption per rack (idle) 208 VAC 3-phase; 100 amp service per rack	Low power draw enables dense packaging
Cooling	Air conditioning 8 tons/rack (minimum) 2800 CFM (compute rack); 350 CFM (power supplies)	Low cooling requirements enable extreme scale-up
Acoustics	9.0 LwAD and 8.7 LwAm	
Dimensions (includes air duct)	Height – 77” Width – 36” Depth – 36” Weight – 1810 lbs. Service clearances – 30” front and back Raised floor height – 16” minimum	Design allows dense floor plan layout for better floor space utilization

Figure 3: Blue Gene System Details at a Glance

5. The IBM Blue Gene Excels at Parallel CFD

Recent advances in domain decomposition techniques coupled with parallel solvers make the Blue Gene solution very suitable for a large class of large-scale CFD problems. These crucial engineering problems in the industry will require the computing performance and scalability that is practical only with the Blue Gene architecture.

Large-Scale CFD Problems Scale Well on the Blue Gene

The range of CFD problems that can benefit from the Blue Gene architecture is sketched in Figure 4. Large scale CFD problems with over 20 million unknowns scale well when domain decomposition techniques that are effective in balancing the computing load are coupled with parallel solvers. Detached Eddy Simulations (DES), Large Eddy Simulations (LES) and Direct Numerical Simulations (DNS) often require a very fine mesh that results in very large numbers of unknowns. Also, interdisciplinary problems with multiple physics and/or shocks increase the problem scale making the Blue Gene even more attractive.

Complex, coupled, computationally challenging problems in flight control simulations, design optimization, parametric studies iterating over thousands of design-space parameters, and chemically reacting flows with multiple species will surely benefit significantly from the Blue Gene architecture. These problems have additional unknown variables per grid point beyond the normal fluid flow variables (velocities and pressures) thus increasing the computational problem size. Also, hundreds of parametric simulations can be done concurrently on multiple Blue Gene partitions. All this will significantly reduce simulation time for large-scale CFD.

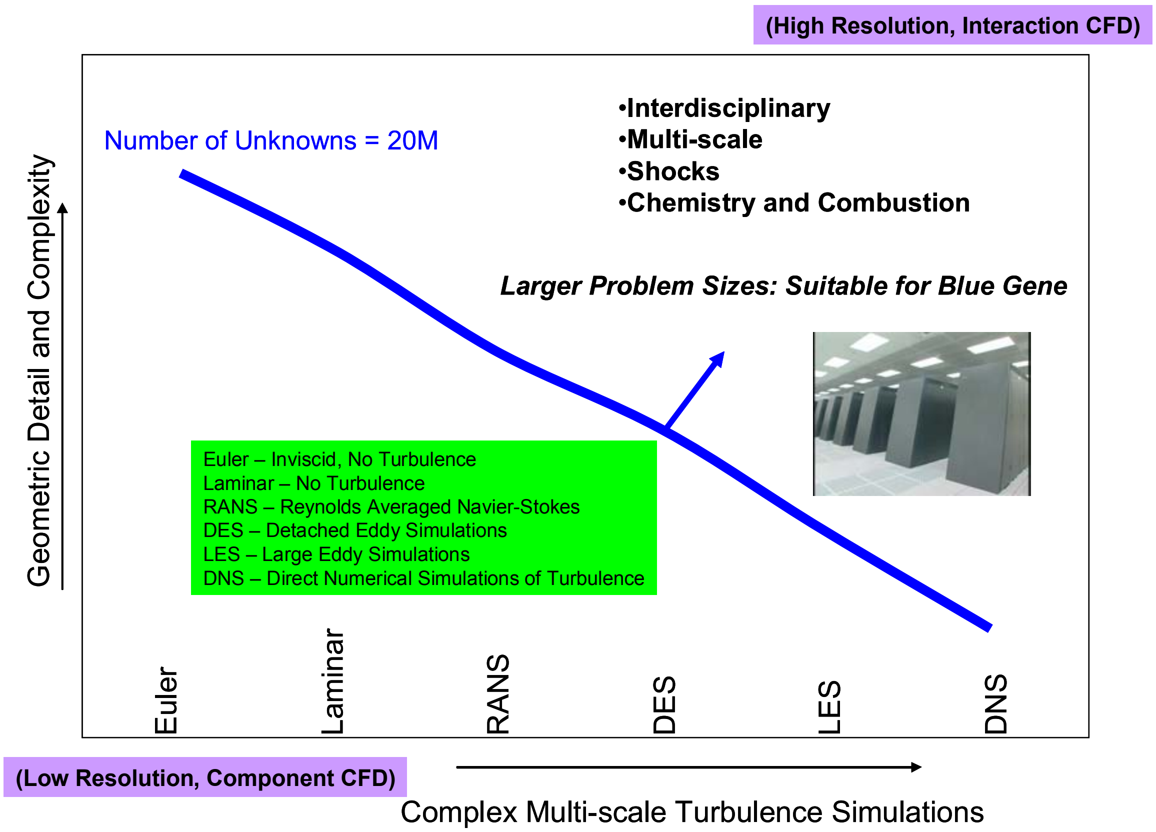

Figure 4: Affinity for Parallel CFD on the IBM Blue Gene

Tuning Further Enhances Scalability and Performance

To take advantage of the Blue Gene architecture, substantial performance tuning can be achieved by careful load balancing, maximizing single processor performance, maintaining data locality, minimizing cache misses, and maximizing the computation-to-communication ratio. Application performance can be further enhanced by the use of mathematical libraries that are optimized for the Blue Gene.

This requires a very structured approach. It consists of first porting the application, then validating it on Blue Gene, and finally optimizing the application on Blue Gene. Careful performance tuning can enhance the performance of the application significantly but requires a deep understanding of the application and the Blue Gene architecture. The investment made in developing these optimized versions can enable the solution of larger-scale problems with a substantial payoff in terms of obtaining engineering insight and value. Furthermore, this investment is protected as newer generations of Blue Gene systems become available. IBM provides access to the Blue Gene system through the Deep Computing Capacity on Demand (DCCoD) center. CFD developers and prospective users can test out and get in-depth experience on the Blue Gene in a cost-effective manner.

Examples

IBM is working with CFD application developers to migrate and optimize their applications on the Blue Gene to solve challenging problems that are multi-scale, multi-regime, and interdisciplinary. These representative examples range from Euler/Navier Stokes solvers with RANS (Overflow-D) on block structured meshes, to chemically reacting, fine-scale, turbulent flows with LES (AVBP) on unstructured meshes. Performance and scalability results for a few representative examples are presented here.

<u>PETSc-FUN3D</u>

PETSc-FUN3D is a variant of the NASA code FUN3D adapted by the Argonne National Laboratory. It is an unstructured Euler/Navier-Stokes solver for both incompressible and compressible flows. PETSc-FUN3D is built on top of Argonne's Portable Extensible Toolkit for Scientific Computing (PETSc), which is in turn built on top of the Message Passing Interface (MPI). PETSc-FUN3D results indicating sustained GFlops achieved are often used to characterize scalability and performance for CFD applications.

Results are shown for an external aerodynamics simulation on an unstructured mesh with 2.7 million vertices and 18 million edges. The results show good scalability and performance for 2048 processors on the Blue Gene.

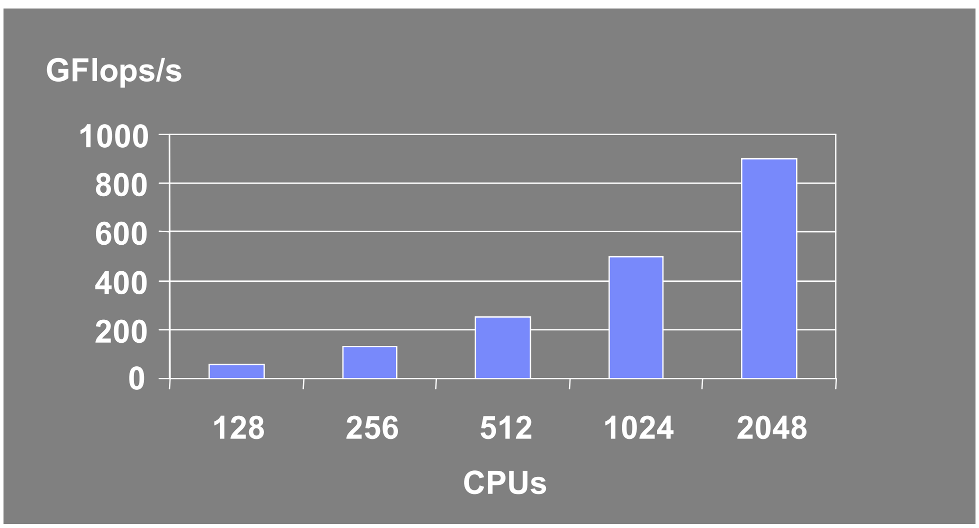

Figure 5: Performance of PETSc-FUN3D on Blue Gene - External Aerodynamics on Unstructured Mesh with 2.7M Vertices and 18M Edges

<u>Overflow-D</u>

OVERFLOW-D is a compressible 3-D flow solver that solves the time-dependent, Reynolds-averaged, Navier-Stokes equations using multiple overset structured grid for problems that may involve relative motion between configuration components. The code uses overset structured grids to accommodate arbitrarily complex geometries with multiple components. In each component, the computational advantages inherent in structured meshes are retained.

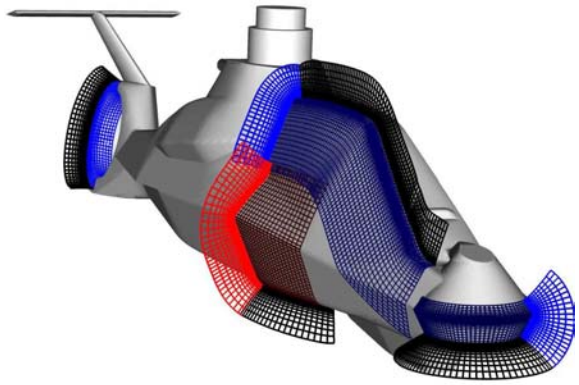

Figure 6: Example of Multiple Overset Structured Meshes for a Helicopter (Courtesy – NASA)

Good scalability and performance results on the Blue Gene are obtained for an external aerodynamics problem with 126 million grid points for 2000 time steps.

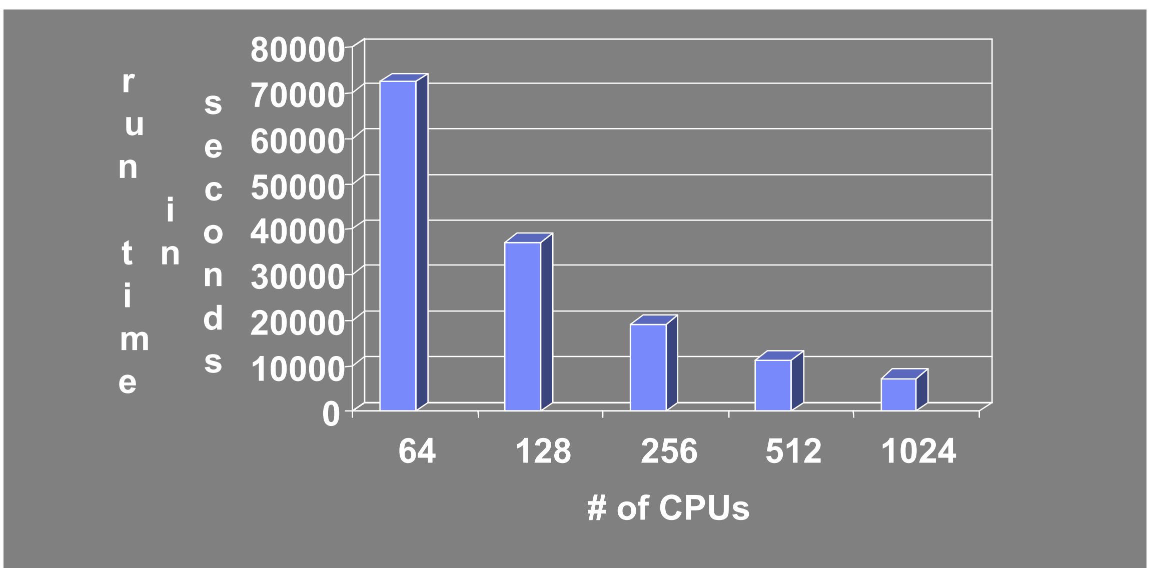

Figure 7: Overflow-D Run Times (Lower is Better) on Blue Gene with 126M Grid Points for 2000 Time Steps

NSU3D

NSU3D is an unstructured mesh, multigrid, Reynolds-averaged Navier-Stokes (RANS) solver for high Reynolds number external aerodynamics applications written under the direction of Prof. Dimitri Mavriplis of the University of Wyoming. It is widely used in the aerospace industry.

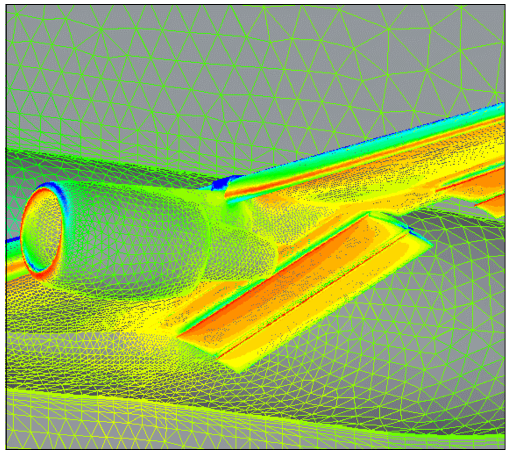

Figure 8: Unstructured Finite Volume Mesh Used in NSU3D (Courtesy – Prof. Mavriplis)

Early Blue Gene performance results (time/iteration) are presented for 72 million grid points on 315 million cells for a transonic flow simulation showing excellent scaling. A full airplane simulation can be done in about 1 hour with about 500 multigrid iterations.

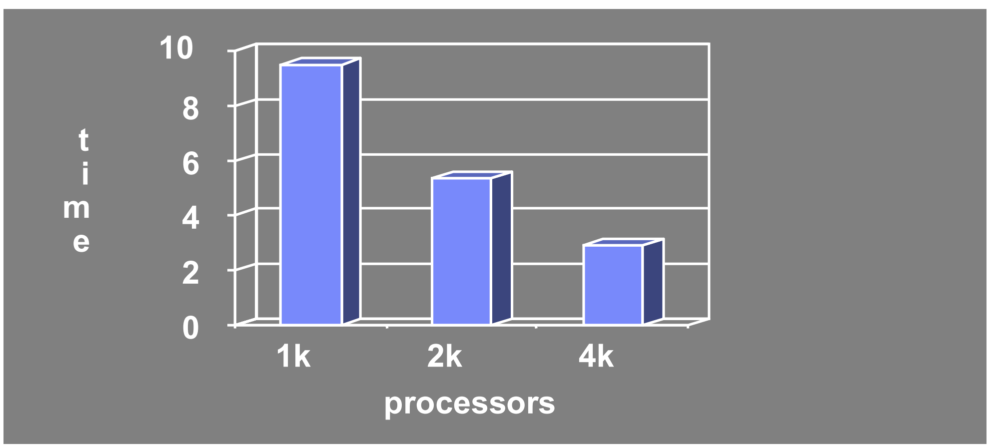

Figure 9: Performance of NSU3D on the Blue Gene (Lower is Better) - 72M Grid Points, Time (sec/iteration)

<u>AVBP</u>

AVBP is a Large Eddy Simulation (LES) CFD solver developed by CERFACS in France to simulate turbulent, unsteady, compressible reactive two-phase flows. AVBP solves the reacting Navier-Stokes on hybrid meshes with explicit time marching. Explicit schemes are very effective on scalable parallel architectures. AVBP is a unique solution to several complex industrial problems such as: pollutant formation, instabilities, ignition, quenching, fuel consumption reduction, piston engines, etc. It is currently used widely in industry and government laboratories in Europe.

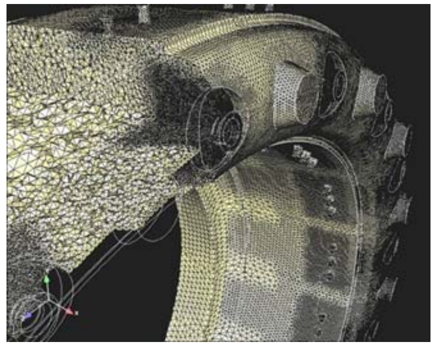

Figure 10: Helicopter Gas Turbine Burner Mesh with Over 3M Grid Points and 18M cells

AVBP's finite volume, parallel (MPI), fully compressible, explicit solver works on unstructured moving meshes with high-order schemes for spatial and temporal accuracy. Coupled multi-physics models for realistic thermo-chemistry and combustion are solved numerically along with the multi-phase fluid dynamics equations. Large-scale simulations for a full gas turbine in a helicopter on over 4000 processors of the IBM Blue Gene enabled the largest reactive LES calculation performed by the scientific community. High resolution flow features that were never revealed before were computed. This was possible by combining LES models with the IBM Blue Gene. Excellent speedup (close to ideal) was obtained. More speedup was obtained by using the optimized MPI collective communications library.

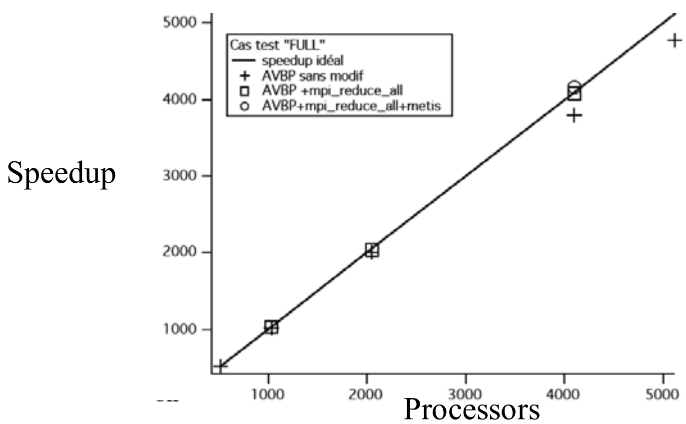

Figure 11: AVBP Speedup on Blue Gene

6. Conclusions: Blue Gene Enables Breakthrough Innovation in CFD

As shown in the examples, large-scale industrial CFD simulations are now viable with the IBM Blue Gene. There has been a lot of progress over the last several years to overcome challenges in solution approaches, algorithms, and domain decomposition techniques. Large-scale flow problems on multi-body configurations on complex engineering geometries such as an entire aircraft or a space vehicle scale well on the IBM

Blue Gene. CFD applications are once again poised to benefit immensely from Blue Gene's unsurpassed scalability and performance. This will continue to advance the field of CFD further. More breakthrough impacts in industrial design will become possible by simulating complex, interdisciplinary, engineering phenomena such as turbulence, combustion, reacting flows, flames, etc. The CFD community and the industry will obtain big payoffs as they continue to solve the grand challenge problems in CFD using a combination of innovative parallel solution approaches and more powerful systems based on the Blue Gene architecture. The direct numerical simulation of turbulence for more realistic engineering configurations will become more prevalent to get deep engineering insights over all scales.

Acknowledgements

This work was sponsored by the Deep Computing group at IBM. Several people at IBM and other CFD experts in industry and academia provided valuable input and guidance for this work. Dr. Maria Iordache of IBM reviewed early drafts numerous times and provided valuable advice to restructure the content and delivery for more effective presentation. Drs. Pascal Vezolle, A. Sugavanam, and Hari Reddy of IBM not only contributed to the content of this paper but also shared data on specific CFD applications that have been optimized for the Blue Gene architecture. Dr. Cameron Brooks and Herb Schultz of IBM reviewed this and provided valuable feedback to improve the quality of this work. Communications and conversations with Prof. Dimitri Mavriplis at the University of Wyoming, Prof. Parviz Moin of Stanford University, and Drs. V. Venkatakrishnan and Forrester Johnson of Boeing were immensely helpful and stimulating. Sally Baldwin of Stein Design helped with the structure and design of this paper.

References and Additional Reading

1. National Science Foundation, "Investing in America's Future", Draft National Science Foundation Strategic Plan, www.nsf.gov/about/performance/nsfplandraft.pdf, 2006.

2. Office of Science, U. S. Department of Energy, "A Science-Based Case for Large-Scale Simulation – Volume 1", July, 30, 2003.

3. National Science Foundation, "Revolutionizing Science and Engineering through Cyberinfrastructure: Report of the National Science Foundation Blue-Ribbon Advisory Panel on Cyberinfrastructure", January, 2003.

4. National Science Foundation, "Simulation Based Engineering Science: Report of the National Science Foundation Blue-Ribbon Advisory Panel on Simulation-Based Engineering Science", February, 2006.

5. Frederica Darema, "Dynamic Data Driven Application Systems (DDDAS): New Direction for Applications/Simulations and Measurement Methodology", Workshop on

Dynamic Data-Driven Application Systems, National Science Foundation, Arlington, VA, January 19-20, 2006.

6. Argonne National Laboratories, "The Message Passing Interface Standard (MPI)", http://www-unix.mcs.anl.gov/mpi/.

7. A. Gara, et. al., "Overview of the IBM Blue Gene/L System Architecture", IBM Journal of Research and Development, Volume 49, Number 2/3, March/May 2005.

8. Nicholas Allsop, et. al. "Unfolding the IBM eServer Blue Gene Solution", IBM Red Book, http://www.redbooks.ibm.com/abstracts/sg246686.html?Open.

9. Andreas Borg, "Status and Trends in Large Scale CFD Applications for Vehicle Design at Volvo Car Corporation", Volvo Car Corporation, http://www.hpc2n.umu.se/para06/papers/paper_180.pdf#search=%22volvo%20cfd%20borg%22.

10. Forrester T. Johnson, Edward N. Tinoco, and N. Jong Yu, "Thirty Years of Development and Application of CFD at Boeing Commercial Airplanes, Seattle", American Institute of Aeronautics and Astronautics, 2003.

11. William D. Gropp, et. al. "High Performance Parallel Implicit CFD", Parallel Computing 27, (2001), pp 337-362.

12. Suga A. Sugavanam, "CFD on Blue Gene/L Supercomputer", Parallel CFD 2003, Moscow, May 15, 2003.

13. Suga Sugavanam, Claude Oudet, and Pascal Vezolle, "CFD on the Blue Gene", IBM Corporation, March 2006.

14. Yannick Sommerer, Claude Oudet, and Pascal Vezolle, "Turbulent Combustion Simulations on the Blue Gene Solution with AVBP (CERFACS)", IBM Corporation, March 2006.

15. Pascal Vezolle, "Massively Parallel LES with IBM eServer Blue Gene", IFP Workshop: LES for Piston Engine Flows, March 2006.

16. Argonne National Laboratories, "PETSc-FUN3D", http://www-fp.mcs.anl.gov/petsc-fun3d/.

17. Dimitri J. Mavriplis, Michael J. Aftomis, and Marsha Berger, "High Resolution Aerospace Applications Using the NASA Columbia Supercomputer", Supercomputing 2005, Seattle, November, 2005.

18. NASA, "Overflow-D", http://rotorcraft.arc.nasa.gov/cfd/CFD4/New_Page/Overflow-D2.htm.

19. Wikipedia, "Computational Fluid Dynamics", http://en.wikipedia.org/wiki/Computational_fluid_dynamics.

Appendix: Classification and Characteristics of CFD Problems

Commercial and Research Applications

Both commercial-off-the-shelf (COTS) and tailor-made, in-house codes (SOURCE) have been used for tackling CFD problems. In general, in the automotive industry, COTS codes such as FLUENT, STAR-CD, AVBP, ANSYS/CFX, etc. have been used extensively for a wide range of applications such as external aerodynamics, underhood cooling, oil pumps, climate control, intake/exhaust manifold design, in-cylinder combustion, turbomachines, interior noise reduction, and catalytic converters. The aerospace industry largely uses SOURCE codes that have been adapted and customized from research and academic codes. Some examples are TRANAIR, FUN3D, NSU3D, and Overflow-D. Applications in the aerospace industry include aircraft aerodynamics, space vehicle aerothermodynamics, propulsion, missile aerodynamics, gas turbine combustion and flows, flow controls, and escape systems.

Solution Approaches

Most approaches follow the same basic procedure:

1. The geometry (physical bounds) of the problem is defined,
2. The volume/area occupied by the fluid is divided into discrete cells (the mesh),
3. The physical model is defined - for example, the equations of motions + enthalpy + species conservation,
4. Boundary conditions are defined. This involves specifying the fluid behavior and properties at the boundaries of the problem. For transient problems, the initial conditions are also defined,
5. The equations are solved iteratively as a steady-state or transient,
6. Analysis and visualization of the resulting solution.

The stability of the chosen discretization is generally established numerically rather than analytically as with simple linear problems. Special care is taken to ensure that the discretization handles discontinuous solutions gracefully. The Euler and Navier-Stokes equations both admit shocks, and contact surfaces.

Some discretization methods used include:

- <u>Finite volume method:</u> This is the "classical" or standard approach used most often in commercial software and research codes. The governing equations are solved on discrete control volumes. This integral approach yields a method that is inherently conservative (i.e., quantities such as density remain physically meaningful).
- <u>Finite element method:</u> This method is popular for structural analysis of solids, but is also applicable to fluids. The FEM formulation requires, however, special care to ensure a conservative solution.

- Finite difference method: This method has historical importance and is simple to implement. It is currently only used in few specialized codes. The main disadvantage is that it requires structured meshes, coordinate transformations, or several meshes for complicated geometries.

The basic solution of the system of equations arising after discretization is accomplished by many of the familiar algorithms of numerical linear algebra. One can either use a stationary iterative method, like symmetric Gauss-Seidel or successive overrelaxation, or a Krylov subspace method. In recent years, Multigrid algorithms that iterate between a coarse and fine mesh have become very popular, because of their efficiency for larger systems of equation, i.e. finer discretization meshes.

Geometries and Grids

Another fundamental consideration in CFD is how one discretizes the equations of motion in space and time. One method is to discretize the spatial domain into small cells to form a volume mesh or grid, and then apply a suitable method to solve the equations of motion. Such a mesh can be unstructured and irregular (for instance consisting of triangles in 2D, or pyramidal solids in 3D), or regular and structured. The distinguishing characteristic of irregular meshes is that each cell must be stored separately in memory. Lastly, if the problem is highly dynamic and occupies a wide range of scales, the grid itself can be dynamically modified in time, as in adaptive refinement and grid redistribution methods. In general finite volume or finite element methods are used for unstructured meshes while finite difference methods are used for structured or block-structured grids.

Several mesh-less alternatives also exist:

- Smoothed particle hydrodynamics, a Lagrangian method of solving fluid problems,
- Spectral and pseudo-spectral methods, a technique where the equations are projected onto basis functions like the spherical harmonics and Chebyshev polynomials,
- Lattice Boltzmann methods, which simulate an equivalent mesoscopic system on a Cartesian grid, instead of solving the macroscopic system (or the real microscopic physics).

The Nature of the Equations Describing the Physics

The most fundamental consideration in CFD is how one mathematically represents a continuous fluid. The equations of fluid motion range from Euler equations for inviscid, to full Navier-Stokes equations for viscous flow. It is possible to directly solve the Navier-Stokes equations for laminar flow cases and for turbulent flows when all of the relevant length scales can be contained on the grid (a direct simulation). Wide-ranging multi-scale turbulent flow simulations require the introduction of a turbulence model. Large Eddy Simulation (LES) and the RANS formulation (Reynolds-averaged Navier-

Stokes equations), with the k-ε model or the Reynolds stress model, are two techniques for dealing with these scales. Detached Eddy Simulations (DES) is a modification of a RANS model in which the model switches to a sub-grid scale formulation in regions fine enough for LES calculations. Regions near solid boundaries and where the turbulent length scale is less than the maximum grid dimension are assigned the RANS mode of solution.

In many instances e.g. weather prediction, chemical and biological flows, propulsion, etc., other equations (mostly convective-diffusion equations) are solved simultaneously with the Navier-Stokes equations. These other equations can include those describing species, concentrations, chemical-reactions, heat transfer, etc. More advanced codes allow the simulation of more complex cases involving multi-phase flows (e.g. liquid/gas, solid/gas, and liquid/solid) or non-Newtonian fluids (such as blood).

Computational Models for Turbulence

Several turbulence models are currently used. These range from computationally simple models to models that require extremely large computing capabilities. The several pictures that illustrate these models have been obtained from the Center of Turbulence Research, Stanford University (http://www.stanford.edu/group/ctr/gallery.html).

Reynolds-averaged Navier-Stokes equations (RANS) is the oldest approach to turbulence modeling. An ensemble version of the governing equations is solved, which introduces new *apparent stresses* known as Reynolds stress. This adds a second order tensor of unknowns for which various models can provide different levels of closure. Statistically unsteady (on non-stationary) flows can equally be treated. This is sometimes referred to as URANS. The turbulence models used to close the equations are valid only as long as the time over which these changes in the mean occur is large compared to the time scales of the turbulent motion containing most of the energy.

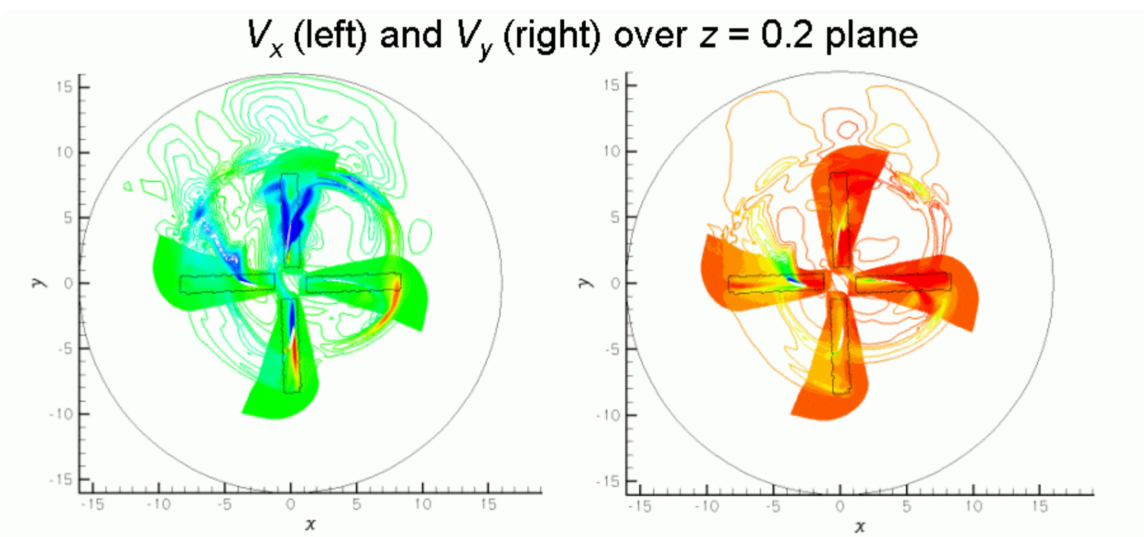

Figure 12: Hybrid Unsteady RANS simulation of Helicopter Tip Vortex (Courtesy – Center for Turbulence Research, Stanford University)

Large eddy simulation (LES) is a technique in which the smaller eddies are filtered and are modeled using a sub-grid scale model, while the larger energy carrying eddies simulated. This method generally requires a more refined mesh than a RANS model, but a far coarser mesh than a DNS solution.

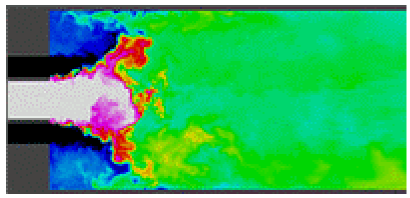

Figure 13: Large Eddy Simulations of Combustor Flows (Courtesy – Center for Turbulence Research, Stanford University)

Detached eddy simulation (DES) is a modification of a RANS model in which the model switches to a sub-grid scale formulation in regions fine enough for LES calculations. Regions near solid boundaries and where the turbulent length scale is less than the maximum grid dimension are assigned the RANS mode of solution. As the turbulent length scale exceeds the grid dimension, the regions are solved using the LES mode. Therefore the grid resolution is not as demanding as pure LES, thereby considerably cutting down the cost of the computation. Grid generation is more complicated than for a simple RANS or LES case due to the RANS-LES switch. DES is a non-zonal approach and provides a single smooth velocity field across the RANS and the LES regions of the solutions.

Direct numerical simulation (DNS) captures all of the relevant scales of turbulent motion, so no model is needed for the smallest scales. This approach is extremely expensive, and intractable, for complex problems on modern computing machines, hence the need for models to represent the smallest scales of fluid motion.

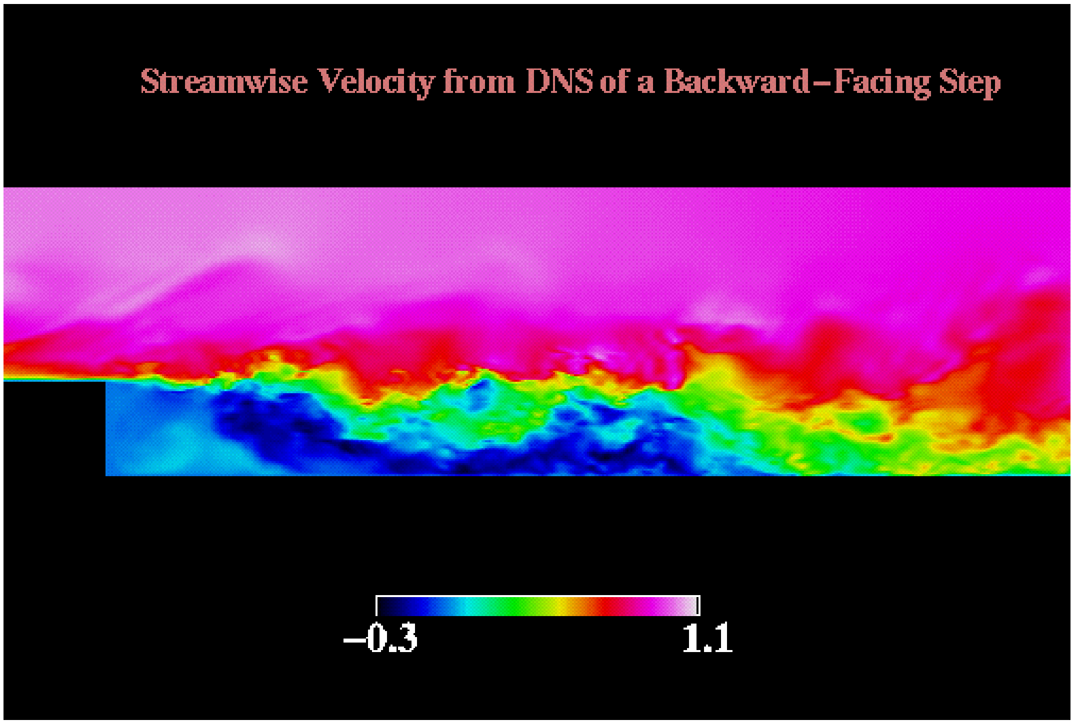

Figure 14: Direct Numerical Simulations of Turbulent Flow over a Backward-Facing Step (Courtesy – Center for Turbulence Research, Stanford University)

Domain Decomposition Techniques

Novel algorithmic approaches in CFD based on domain or mesh decomposition strategies allow users to obtain the maximum advantage and scalability from parallel machines with large numbers of inter-connected processors. Before a calculation can be performed on a parallel computer, it must first be decomposed into tasks which are assigned to different processors. Efficient use of processor resources requires that each processor have about the same amount of work to do and that the quantity of inter-processor communication is minimized. Finding an optimal decomposition is hard, but due to its practical importance, a great deal of effort has been devoted to developing heuristics for this problem. Some prominent domain decomposition tools are:

METIS is a family of programs – serial and parallel - for partitioning unstructured graphs and hyper-graphs and computing fill-reducing orderings of sparse matrices (http://glaros.dtc.umn.edu/gkhome/views/metis/). The underlying algorithms used by METIS are based on a multilevel approach that has been shown to produce high quality results and scale to very large problems with unstructured finite element or finite volume grids.

CHACO is a set of innovative algorithms and programs for domain decomposition developed by Sandia Labs. This code is being widely used in CFD research applications to simplify the development of parallel applications, and to ensure that optimal performance is obtained (http://www.cs.sandia.gov/CRF/chac.html).

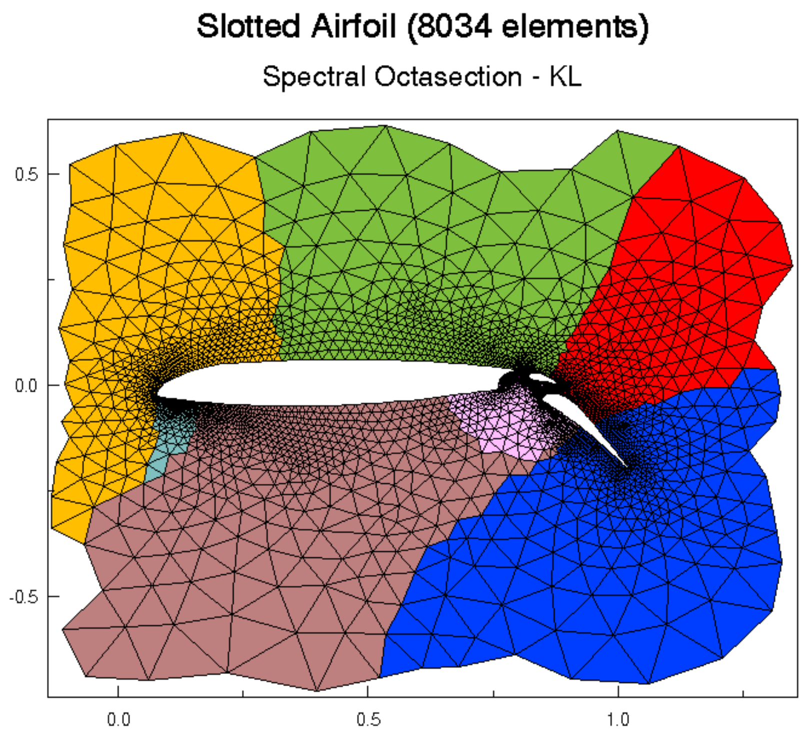

Figure 15: Domain Decomposition of an Unstructured Finite Volume Mesh into Eight Sub-domains Using CHACO (Courtesy – Sandia Laboratories)

JOSTLE is a software package designed to partition unstructured finite element or finite volume meshes (http://staffweb.cms.gre.ac.uk/~c.walshaw/jostle/). The code can also be used to repartition and load-balance existing partitions especially within the context of adaptive grids. Several state-of-the-art graph partitioning techniques are used.

More Information

To learn more about the IBM System Blue Gene Solution, please contact your IBM marketing representative or visit the following Web site: http://www.ibm.com/servers/deepcomputing/bluegene.html.