

*Refereed Proceedings*

*The 13th International Conference on  
Fluidization - New Paradigm in Fluidization  
Engineering*

---

Engineering Conferences International

Year 2010

---

HIGH PERFORMANCE  
COMPUTING: CLEAN COAL  
GASIFIER DESIGNS USING HYBRID  
PARALLELIZATION

Madhava Syamlal\*

C. Guenther<sup>†</sup>

A. Gel<sup>‡</sup>

S Pannala\*\*

\*National Energy Technology Laboratory, madhava.syamlal@netl.doe.gov

<sup>†</sup>National Energy Technology Laboratory

<sup>‡</sup>National Energy Technology Laboratory and ALPEMI Consulting, LLC

\*\*Oak Ridge National Laboratory

This paper is posted at ECI Digital Archives.

[http://dc.engconfintl.org/fluidization\\_xiii/108](http://dc.engconfintl.org/fluidization_xiii/108)

# High Performance Computing: Clean Coal Gasifier Designs Using Hybrid Parallelization

M Syamlal<sup>1</sup>, C Guenther<sup>1</sup>, A Gel<sup>1,2</sup>, S Pannala<sup>3</sup>

1 National Energy Technology Laboratory, Morgantown, WV 26505, U.S.A.

2 ALPEMI Consulting, LLC, Phoenix, AZ 85044, U.S.A.

3 Oak Ridge National Laboratory, Oak Ridge, TN 37831, U.S.A.

E-mail: madhava.syamlal@netl.doe.gov

**Abstract.** One of the targets for coal gasification in the near future is capturing 90% of the carbon with less than a 10% increase in cost of electricity. Aggressive goals like this will require innovative gasifier designs to reach the market place quickly, with less risk, and in an economically viable way. Researchers at the National Energy Technology Laboratory (NETL) are collaborating with industry, academia, and other national labs on multiphase computational models like the legacy code MFIX (Multiphase Flow with Interphase eXchange) which can help design, operate, and scale-up clean coal gasifiers to meet the challenges of a carbon constrained world. In fact, NETL has hosted a series of multiphase workshops which has produced a multiphase flow science technology roadmap to achieve the goal “*that by 2015 multiphase science based computer simulations play a significant role in the design, operation, and troubleshooting of multiphase flow devices in fossil fuel processing plants*”. In this study, we present our experience of porting MFIX, an open source multiphase computational fluid dynamic model, to a high performance computing platform and how the resulting high fidelity simulations are impacting the design of clean coal gasifiers of tomorrow. Inherent to these gasifiers is the various time and length scales which require very high spatial resolution, large number of iterations with small time-steps to resolve and predict the spatiotemporal variations in gas and solids volume fractions, velocities, temperatures with any associated phase change and chemical reactions. These requirements resulted in perhaps the largest known simulations of gas-solids reacting flows, providing detailed information about the gas-solids flow structure, pressure, temperature and species distribution in the gasifier. From a computational science perspective, we found that global communication has to be reduced to achieve scalability to 1000s of cores and hybrid parallelization can yield substantial improvement in time-to-solution when utilizing thousands of multi-core processors.

**Keywords:** Reactive multiphase flows, clean coal gasifiers, fluidization, open source, high performance computing, hybrid MPI and OpenMP parallelization

## Introduction

Environmentally clean energy is a vital need throughout the world. Meeting the growing demand for clean energy is arguably the most important problem that the world faces today, since the availability of reasonably priced energy is critical to maintaining living standards in the developed world, and raising standards of living across the developing world. In the United States the energy demand in recent decades has outpaced the energy supply. Rising and volatile prices of petroleum and natural gas threatens the economy.

Furthermore the global demand for energy is also rising rapidly (expected to double by 2050 and triple by end of the century), caused by worldwide growth, which is explosive in economies such as China's and India's. Coal is plentiful (with known reserves for over 200 years at the current usage levels) in the United States and currently generates more than half of the electric power. Coal-based electric power generation is projected to increase by 156 GW by 2030, representing at least 156 full-scale power plants.

Integrated Gasification Combined Cycle (IGCC) is a promising technology for meeting the growing demand for power using fossil fuel resources, while economically controlling the emission of CO<sub>2</sub> and other pollutants. The centerpiece of an IGCC system is the gasifier, which converts coal or other carbonaceous materials such as biomass into syngas, a mixture of CO and H<sub>2</sub>. The syngas can be used for the production of liquid fuels and chemicals or for power generation. In an IGCC system the syngas is shift converted into a mixture of CO<sub>2</sub> and H<sub>2</sub>, CO<sub>2</sub> is captured and stored, and H<sub>2</sub> is used for power generation.

A reliable gasifier is critical for the commercial viability of IGCC. An advanced gasifier technology based on a transport gasifier is being developed at the Power Systems Development Facility (PSDF) in Wilsonville, Alabama, a joint project between the U.S. Department of Energy (DOE), Southern Company, and Kellogg Brown and Root (KBR). The gas-solids multiphase flow such as occurring in the transport gasifier is known to make the design of commercial-scale units using traditional scale-up methods unreliable [1]. Multiphase computational fluid dynamic (CFD) models are being developed at NETL to address that challenge, and this paper discusses the potential for using large-scale CFD simulations for gasifier design. The National Energy Technology Laboratory (NETL) developed, verified, validated, and applied a transport gasifier model with data from the PSDF pilot-scale facility and results of this effort have been reported before [2].

The validated model was used to study the impact of the exit geometry on syngas composition, the effect of reactor height on CO concentration, the coal trajectory within the gasifier, the gas temperature profiles in the gasifier, and the effect increasing the pressure. Later the model was used to help with the design of a Clean Coal Power Initiative (CCPI) gasifier (285 MW electric), which was scaled up from the pilot-scale gasifier (13 MW thermal), representing a factor of 50 scale-up. The multiphase CFD simulations were used to study the effect of pressure, the height to diameter ratio, coal feed rate, coal feed nozzle operation, solids circulation rate, and the effect or recycled syngas. Several sets of parametric evaluations showed an unexpected behavior of the gas species concentrations and solids temperature upon the coal jet penetration. Temperature dependency on the coal feed can impact the thermal decomposition of the coal, produce unwanted agglomeration, and could have a significant economic impact due to material degradation.

This result created interest in conducting simulations of the CCPI gasifier at greater grid resolutions. The rule-of-thumb for grid independence in gas-solids simulations is that the grid size should be of the order of 10 particle size [3]. This meant that the CCPI gasifier needed to be simulated with about 100 billion cells. This resolution would require well over three million compute cores and result in file sizes on the order of 745 GB for a single time record for each scalar tracked in the simulation. This was certainly beyond the computational capability available at NETL and, in fact, even at the fastest HPC platforms. The group was awarded a significant amount of computational time allocation

at one of the National Leadership Class High Performance Computing Facility through the U.S. Department of Energy's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program in 2007. The INCITE grant is allowing the group to conduct CCPI gasifier simulations at greater grid resolution than was previously possible. This paper describes the detailed grid resolution study conducted to investigate the effect of coal jet penetration and how coal temperature and concentration of various species are affected. This study was motivated from the preliminary results of the first-of-its-kind 10 million cell high fidelity simulations conducted as part of the INCITE grant, which has shown the need to determine an adequate resolution for the coal jet area. In order to conduct high fidelity simulations using modern high performance computing platforms a number of improvements were implemented, which are briefly discussed in the paper.

### Problem Description and Results

As part of the effort to better understand coal jet penetration, a reduced configuration of the full commercial scale gasifier was considered by focusing on only the coal jet region using time averaged flow field conditions from a 10 million cell full commercial scale simulation to impose inflow boundary condition for cases R1-R6.

To investigate the effect of grid resolution and various discretization schemes, six different cases as shown in Table 1 were generated. In the high order discretization cases, deferred correction method with Superbee discretization scheme was employed. In the low order discretization cases, first order upwinding was employed for all variables.

Grid Resolution (radial, axial, angular) cell count		Case Name	Discretization Scheme Employed	Wallclock time for 10s of simulated time
Medium	40 x 300 x 60 = 0.7M	R1	Low Order	
		R2	High Order	
High	60 x 450 x 90 = 2.4M	R3	Low Order	305 (hours)
		R4	High Order	1040 (hours)
Coarse	30 x 225 x 45 = 0.3M	R5	Low Order	
		R6	High Order	

**Table 1.** Grid resolution and discretization schemes for the coal jet simulations.

Computational cost is one of the important factors that affect the time-to-solution for high fidelity simulations. In spite of the additional insight gained from higher resolution and/or more accurate predictions, the computational cost might be prohibitive, which typically obliges the practitioners to consider the trade-off between time-to-solution and accuracy. As shown in Table 1 for the same grid resolution (i.e., 2.4M cells) employing a high order discretization scheme requires approximately three times the cost of low order discretization scheme.

Figure 1 shows the steam mass fraction profiles through the reactor along a clip plane aligned with the two coal jets. The figure clearly shows the influence of both low

order discretization and grid resolution for the cases studied. The figure also reveals an asymmetry in the steam mass fraction between the two coal jets. This is due to the mass inlet boundary condition imposed at the bottom of the reactor which was taken from time-averaged results from a full-scale simulation of a transport gasifier. In the actual reactor recycle material a mix of char, ash, and inert solids enters the gasifier from one side. This causes an asymmetry in the gas and solids flowing up through the reactor and causes reduced coal penetration on the right of all the simulations due to the higher solids mass flux coming up through the reactor from the recycle material. The coal jet on the left penetrates farther into the reactor which is predicted by all cases. However, in terms of mixing cases R2 and R4 show dramatic differences. This improved mixing produces a more thermally homogeneous flow field on the left, lower solid velocities and residence times, and can result in significant differences in local product gas compositions.

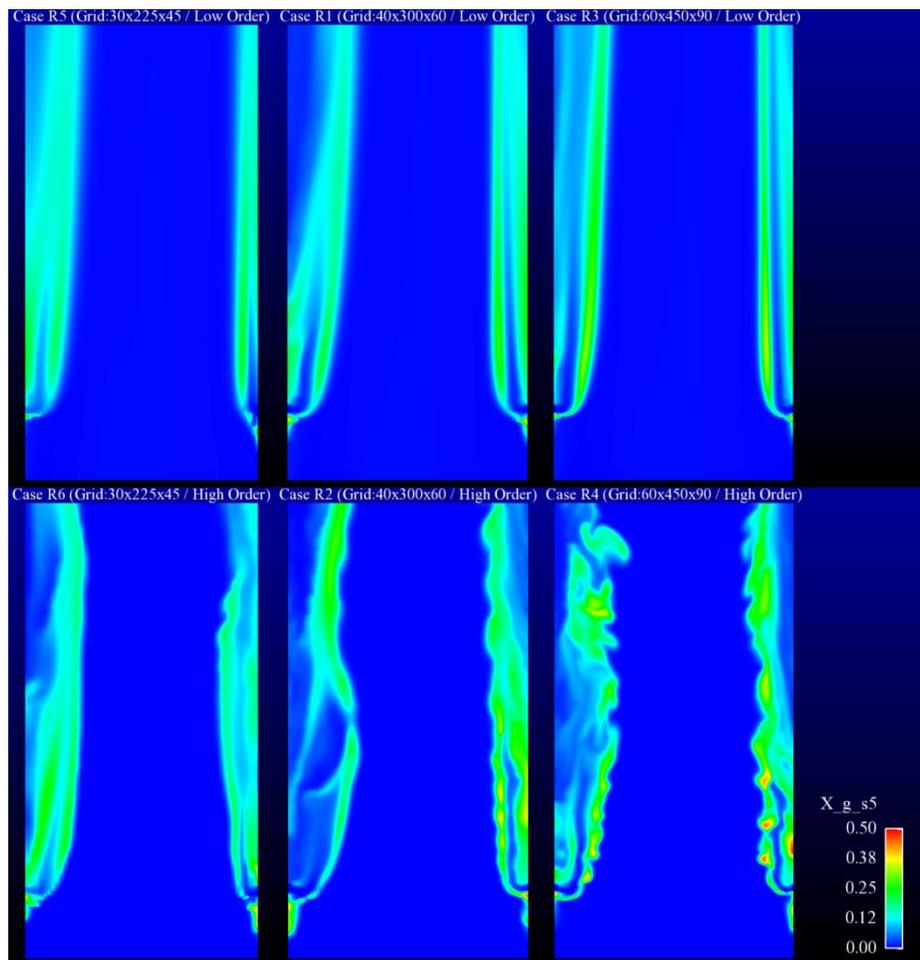


Figure 1. Void fraction for cases R1-R6.

Similar coal jet penetration profiles were also seen in the void fraction, solid temperature, and major gas product species but were not included due to space limitations. The low order cases all showed higher solid axial solid velocities in the coal region compared to

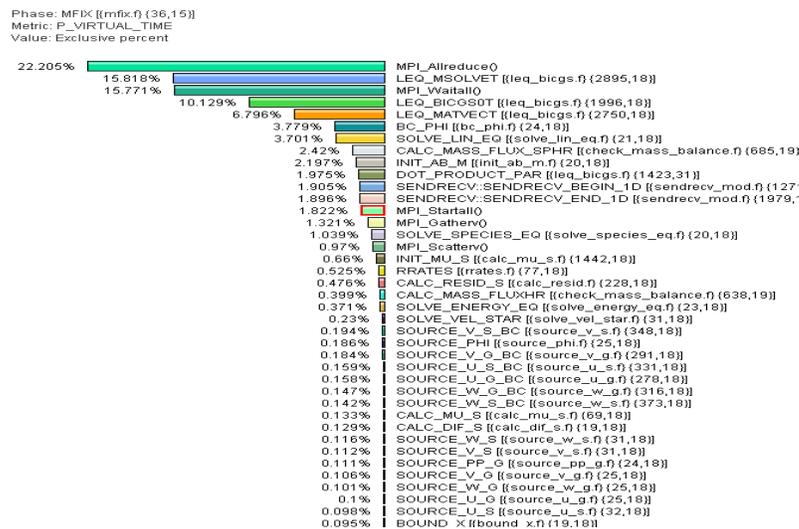
the high order cases. Comparing actual rates between low order and high order methods showed approximately one order magnitude lower devolatilization and steam gasification rates predicted by the low order cases but little difference between the various grid resolution cases.

### Improvements to MFIX Capabilities

The capability to run high fidelity simulations with a short time-to-solution is critical in the development of advanced coal technology. Computational models have been proven to provide crucial insight for various design configurations. For example, when the high resolution simulations performed on Cray XT4 at NCCS were compared with older runs, significant detail of the transient physics not captured in earlier lower resolution runs was observed in the latter set of simulations. Considering the alternative option of building a prototype gasifier for experimentation purposes, which is both very expensive and time intensive, high-fidelity computational models are critical in the design process. However, achieving an acceptable time-to-solution in terms of wall clock time is crucial for computational models to have an impact during the design cycle. For this purpose, a number of improvement projects have been undertaken, which are described briefly in this paper.

### Performance Profiling

Performance profiling tools TAU [5] and CrayPAT[6] were employed to document the current state of the code to establish a baseline and identify bottlenecks on the targeted HPC platforms [9]. Figure 2 shows a detailed view of the top 10 routines for the most time spent at the processor level (in this case for Processor # 1). Time spent in MPI Allreduce() accounts for 22 % of the wall clock time on this processor. The next routine, MPI Waitall() calls spent 16% of the time. Standalone the linear equation solver routines (LEQ \*) were observed to be using 32 % of the time however, the MPI Allreduce() calls originating from linear equation solver subroutines need to be also considered. These profiling results correspond to a 256 processor run, however, the impact of MPI Allreduce() was more predominant in the profiling results for greater numbers of processors (512 and higher).



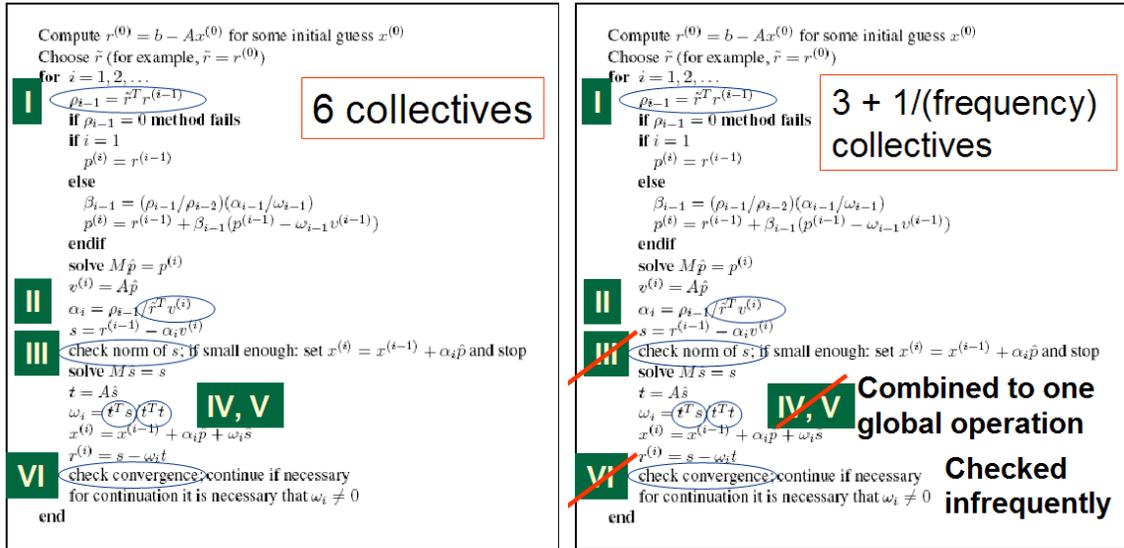
**Figure 2 TAU profiling showing MPI\_AllReduce as the top function that spent most of the time.**

### **Performance Improvements**

To improve the performance of MFIX on today's massively parallel heterogeneous many core high performance computers such as Cray XT series, various optimizations and improvements were incorporated in several phases. Phase 1 improvements were geared towards optimization without any code change, i.e., determining the best set of compiler flags or MPI environment variables for the batch job. Phase 2 improvements were algorithmic changes, such as reducing the number of MPI\_AllReduce calls in the linear equation solver, reducing residual calculations, etc. As seen in Figure 3, profiling data indicated that significant time was spent in the MPI\_Allreduce calls. These calls primarily originated at two locations in the iteration loop. The first set of MPI\_Allreduce calls came from the BiCGStab routine, which is the linear equation solver subroutine based on Templates [7]. The BiCGStab algorithm from Templates has 6 calls to the dot products (as illustrated in Figure 3 (a)), which in turn call the global collectives. A number of these dot-products can be removed as shown in Figure 3 (b). For example, the dot-product to check the residual during the calculation can be safely eliminated without much loss of efficiency as typically the linear systems take several iterations to converge. The 4th and 5th dot products in the algorithm are not dependent on each other and thus have been collected into a single global collective of array length 2. The last dot product has been infrequently calculated at a fixed frequency depending on the problem. Thus the number of global collectives per BiCGStab iterate have been reduced to  $3+1/N$  per call (where  $N$  is the frequency at which residual is sought) instead of 6 in the original algorithm implemented.

Furthermore, the compiler was changed from Portland Group to Pathscale during Phase 2 improvements based on independent benchmarks conducted between two compilers. Phase 3 improvements were based on the extension of a past study to explore parallel performance of MFIX when hybrid shared memory parallel (SMP) and distributed memory parallel (DMP) execution is employed on shared memory multiprocessor systems [8]. Prior to the DMP version of MFIX, an SMP version was developed by manually inserting portable OpenMP directives around the DO-loops in the most time consuming routines in MFIX (approximately 62 locations throughout the code). Due to the limitations in scalability with the SMP mode, the DMP mode of execution, which is based on MPI message-passing library was routinely used for jobs with large number of processors. In the study of Pannala et al. [8], one thread per MPI task, one MPI task per processor was determined to give the best performance with the parallel computer architectures used at that time (e.g., IBM SPs). This observation was consistent with previous hybrid parallelization efforts on somewhat similar architectures. One of the reasons was attributed to the fact that thread creation/destruction is very expensive on the systems employed at that time, which were also single core processor based systems. The basic idea for this improvement phase was to exploit shared memory parallelism within the cores on the compute node and distributed memory parallelism

across the nodes of Cray XT4 platform at National Center for Computational Sciences (NCCS). This approach was not possible on the early Cray XT platforms due to lack of mul- tithread support on the compute nodes. However, with the upgrade to compute node kernel (CNL) and installation of supporting compilers, hybrid mode operation, i.e., MPI and multithreaded OpenMP instructions became a possibility.



(a) Original BiCGSTAB from Templates

(b) Modifications in BiCGSTAB to reduce global collective calls.

**Figure 3. Illustration of global collective call reduction in BiCGSTAB algorithm before (a) and after (b)**

As seen in Figure 2, the performance profiling results suggests that apart from the MPI time, significant amount of time is spent in linear equation solver routines, i.e., leq ik sweep and leq matvec. To isolate the impact of most time consuming routines, out of 62 DO-loop related OpenMP directives that were implemented for SMP mode of operation, only those in the above listed two subroutines were enabled for OpenMP supported compilation and testing on the Cray XT4. Benchmarking runs based on the production code and problem size have shown substantial improvement in time-to-solution when hybrid approach (i.e., MPI and OpenMP) is employed on Cray XT4 platform [10]. Good scalability with hybrid mode operation is becoming increasingly critical with the newer HPC architectures that employ many cores per socket.

### Conclusion

The results presented in this paper clearly show the importance of discretization scheme and grid resolution in the modeling of a gasifier. This paper shows the minimum resolution required to capture the hydrodynamics satisfactorily. In this work case R2 is considered the minimum resolution required to model a commercial scale transport gasifier. This resolution requires roughly 10 million computational cells to model a full-scale transport gasifier. This paper also presented preliminary results quantifying the effect discretization and resolution has on the kinetics. Results showed differences in the devolatilization and steam gasification rates due to discretization schemes.

From a computational science perspective, we have found that global communication has to be reduced to achieve scalability to 1000s of cores. In addition, we have to employ hybrid parallelization to effectively utilize the multicore chips. We also find that the factors such as wait time in the batch queue could significantly increase the time-to-solution. From our experience, developments have to be made in the following areas to efficiently use large petascale resources and enable computational discovery:

- Efficient solvers for heterogeneous, massively parallel systems
- Data analysis tools to extract information from large data sets
- Programming environments for easily porting legacy codes to HPC platforms

## References

- [1] Report on Workshop on Multiphase Flow Research, Morgantown, WV, June 6-7, 2006, DOE/NETL-2007/1259, <http://tinyurl.com/c9r7ux>
- [2] Guenther C, Shahnam M, Syamlal M, Longanbach J, Cicero D, and Smith P, "CFD modeling of a transport gasifier", Proceedings of the 19th Annual Pittsburgh Coal Conference, Pittsburgh, PA, September 23-27, 2002.
- [3] Guenther C, Syamlal M, Longanbach J, and Smith P, "CFD modeling of a transport gasifier Part II", Proceedings of the 20th Annual Pittsburgh Coal Conference, Pittsburgh, PA, September 15-19, 2003.
- [4] Guenther C, Syamlal M, Shadle L, and Ludlow C, "A numerical investigation of an industrial scale gas-solids CFB", Proceedings of the 7th International Conference on Circulating Fluidized Beds, Niagra Falls, Ontario, Canada, May 5-8, 2002.
- [5] TAUPerformanceSystem, University of Oregon, [URL: http://www.cs.uoregon.edu/research/tau](http://www.cs.uoregon.edu/research/tau)
- [6] CrayPat Performance Analysis Tool, URL: <http://docs.cray.com/books/S-2315-50/html-S-2315-50/z1055157958smg.html>
- [7] Barrett R, Berry M, Chan T, Demmel J, Donato J, Eijkhout V, Pozo R, Romine C, and van der Vorst H, Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods, SIAM Publishing, Philadelphia, PA, 1994. [www.netlib.org/templates/](http://www.netlib.org/templates/).
- [8] S. Pannala, E. D'Azevedo, M. Syamlal, and T. O'Brien, "Hybrid (mixed SMP/DMP) parallelization of MFIx: A multiphase CFD code for modeling fluidized beds," Presented at ACM conference, March 2003.
- [9] A. Gel, C. Guenther, S. Pannala, S. Benyahia, J. Galvin, M. Syamlal and T.J. O'Brien, "Accelerating Clean and Efficient Coal Gasifier Designs with High Performance Computing", Technical Report (manuscript submitted for review to Journal of Computational Science).
- [10] A. Gel, S. Pannala, R. Sankaran, C. Guenther, M. Syamlal, and T.J. O'Brien, "Accelerating Clean Coal Gasifier Designs with Hybrid MPI/OpenMP High Performance Computing", Proceedings of the 21st International Conference on Parallel Computational Fluid Dynamics, May 18-22, 2009, Moffett Field, CA. pp: 224-228