4.1

PROGRESS TOWARDS A COUPLED MESOSCALE AND MICROSCALE MODELING CAPABILITY

William J. Coirier *, Sura Kim, Saikrishna Marella, Joel Mayes
CFD Research Corporation
Fei Chen, John Michalakes, Shiguang Miao
National Center for Atmospheric Research
Mathew Bettencourt
Bettencourt Consulting, LLC

1. INTRODUCTION

In this paper we present progress towards a coupled Mesoscale-to-Microscale modeling capability for the urban regime in which we will loosely couple the Weather Research and Forecasting (WRF) model to a high-resolution Computational Fluid Dynamics model using the Model Coupling Environmental Library (MCEL). The client/server, dataflow software model of MCEL provides the data transfer and data filter capability that both models may use to perform both upscale (CFD to WRF) and downscale (WRF to CFD) data transfer. Potential benefits of this coupling include improved high-resolution wind, turbulence and contaminant fields in the urban area through the use of downscaled data, as well as potential improvements at the mesoscale by using an improved urban characterization through upscaled data. Previous studies [Coirier, et al., 2006] have shown a quantifiably improved accuracy of urban area transport and dispersion modeling by the use of a file-based coupling to downscale data from WRF to a high-resolution CFD model. The work being undertaken here is intended to provide a more useful capability by coupling the models loosely in the MCEL environment, whereupon we may evaluate the potential improvements at both scales, as well as evaluate different techniques to perform the upscale and downscale data.

In order to achieve this coupled capability, modifications to both WRF and the MCEL library are being made, and a special-purpose CFD model is being developed to be run within this coupled environment. In this paper we outline the coupling approach using the MCEL library and demonstrate the generation of downscaled data from WRF using a hindcast model run corresponding to a particular Intensive Operating Period (IOP) of the Joint Urban 2003 (JU2003) field test conducted in Oklahoma City [Allwine and Flaherty, 2006]. The new, high-resolution urban CFD model being developed for this study is described and preliminary model validation study results using the new model are shown. This CFD model is designed to be run in a parallel computing environment using the PETSc parallel sparse matrix library, and solves the fully-coupled, low-Mach number preconditioned, Reynolds-Averaged Navier-Stokes (RANS) equations using a Finite-Volume scheme. In the final coupling configuration, downscaled data obtained from WRF through the MCEL will be used to apply boundary conditions for the CFD model, which then cycles to a new, quasi-steady state, and then performs an upscale data transfer to WRF through MCEL. Here, we demonstrate the ability to downscale through the software framework by using WRF computations corresponding to a particular Intensive Operating Period (IOP) of the Joint Urban 2003 (JU2003) field test conducted in Oklahoma City.

corresponding author address:
CFD Research Corporation, 215 Wynn Dr., 5th Floor, Huntsville, AL 35805 e-mail: wjc@cfdrc.com.

2. MCEL COUPLING FRAMEWORK

The Model Coupling Environmental Library, MCEL, utilizes a data flow approach where coupling information is stored in a centralized server and flows through processing routines called filters to the numerical models which represent the clients. The communication is handled by the Common Object Request Broker Architecture (CORBA), where the flow of information is fully controlled by the clients, and stores and retrieves are initiated by the clients. MCEL has been successfully demonstrated coupling WRF to other applications representing a relatively wide range of spatial and temporal scales with irregular data-driven interactions [Michalakes et al., 2003]. WRF has been coupled as part of a four-model simulation of a high-wind event in the Yellow Sea that resulted in a ferry accident in November 1999, where WRF provided low-level winds and wind stress to the ADCIRC Ocean and SWAN wave models which, in turn, provided forcing to LSOM, sedimentation and optics model used to simulate diver visibility. It should be noted that ADCIRC uses an unstructured mesh, and that MCEL automatically handled exchange and interpolation of data between ADCIRC and the other three structured-grid components without modification. When running the coupled system on the maximum number of processors for each application, the cost of coupling was well under 5 percent of the cost of the run. Another added benefit regarding MCEL is that it is public-domain software, and is presently being used under multiple government contracts.

The coupling approach invisioned for coupling the WRF and new CFD models is illustrated in the diagram shown in Figure 2.1. Both models, WRF and the CFD model, operate independently of one

another, and upon completion of a cycle, send data to the central MCEL Server (shown in the middle of the diagram). Upon initiating a new cycle, each component (CFD or WRF) queries for either upscale data (WRF) or downscale data (CFD) from the server through a data filter, which lies between the server and the client application. The intent of the filter is to perform all data manipulation outside of the client.

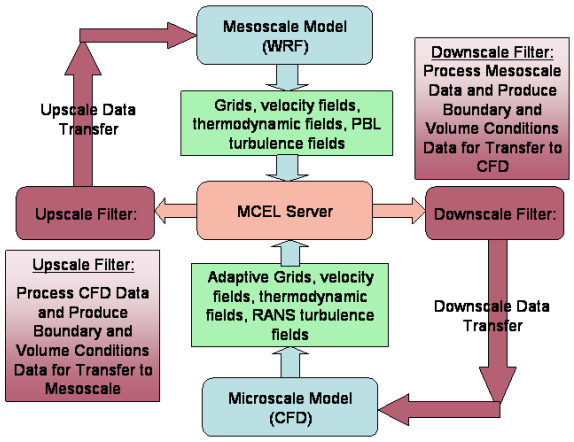

Figure 2.1: WRF-MCEL-CFD Coupling Conceptual: Initial Model.

The MCEL framework is being used in a slightly different fashion to couple the WRF and CFD models. The CFD model uses an adaptive octree/quadtree-prismatic mesh and has associated C++ classes and namespaces that are specialized to operate upon it. Since the coupled models will require both downscaling (WRF-to-CFD interpolation) and upscaling (CFD-to-WRF agglomeration and/or integration), we have decided to move the downscale and upscale filters into the CFD solver itself. This runs counter to the preferred approach proposed by the MCEL concept, where the filters lie outside of all solvers. We prefer this technique in order to take advantage of the specialized in-solver methods and classes to perform the downscaling and upscaling data processing, and to provide a cleaner interface to the MCEL framework.

3. WRF-MCEL-CFD INTEGRATION

Configured to operate much like an output format, the WRF model dumps a subcube of the innermost/finest domain to MCEL at 5-min intervals that spans the CFD domain containing 3-D wind, temperature, pressure, TKE and TKE dissipation rate. The particular subcube (or potentially subcubes for multiple, concurrently running CFD models), is obtained so that it completely encloses the CFD domain. The MCEL data model associates a single array with each grid, whose length is determined by the grid dimensions. For a logically ordered (structured) grid of dimension IxJxK, this implies that arrays of I*J*K length may be passed through MCEL, as they are associated with this registered grid. Since WRF uses a staggered grid scheme, which uses data stored at cell faces and cell centers which have correspondingly different length arrays, the actual subcube dimensions are increased in order to provide enough room to store all the data.

As noted above, we have moved the functionality of the MCEL data filters into the CFD model itself, in order to take advantage of the in-core CFD model data processing capability. The diagram shown in Figure 3.1 illustrates the integration of the CFD model into the MCEL framework, where the top-level data processing approach taken within the CFD solver is noted in the process diagram on the left. Note that during the boundary condition definition phase within a given CFD model cycle, data is pulled from the MCEL server, through the MCEL integration class library/namespace, and downscaled. After the CFD cycle is completed, upscaling is performed in-core, again using this class library, and sent to the MCEL server.

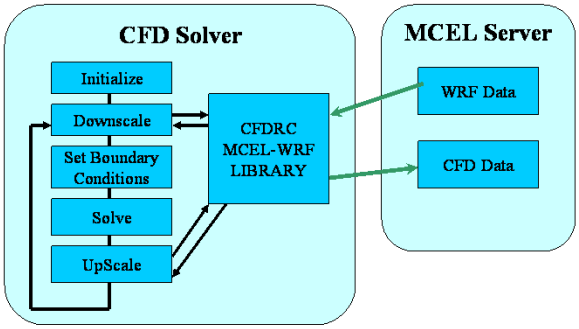

Figure 3.1: CFD-MCEL Integration.

This MCEL communcation class library and associated C++ namespaces have been created to contain classes that communicate with the MCEL server, load up and provide downscaling interpolation capability to the CFD model and provide upscaling capability from the CFD model dataspace to the MCEL server itself. As noted in Figure 3.2, this namespace contains classes which communicate directly with the MCEL server ("MCEL Communicators"), store data upon logically structured grids corresponding to the donor WRF data ("MCEL GridNVars"), and perform the actual downscaling and upscaling ("MCEL DownScaler" and "MCEL UpScaler"), using various interpolation schemes using the CFD mesh connectivity and data framework ("Pnt Interpolators"). These are all built upon a library of "baseTypes", which are derived from the C++ std namespace vector class and contain many useful functions and features, including i/o to a local database using the NetCDF library. Built over this baseTypes library is a specialized library and namespace that is used to read and write CFD data, such as meshes and the solutions upon them, in an easy to use framework called a wind library.

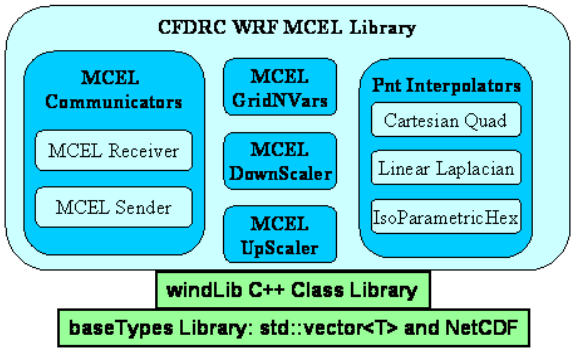

Figure 3.2: CFDRC MCEL Class Library Components.

In order to provide accurate downscaling interpolation, we store the donor data received from WRF upon logically ordered, structured grids that are constructed to lie coincidently with the "native" donor spatial locations. This is necessary since WRF uses a staggered grid arrangement which stores the velocity field (u, v and w) at different locations from the cell-centered data, such as thermodynamic and other state fields. These 4 grids (u-, v-, w- and cell centers) are constructed using the coordinates stored in the subset grid sent by WRF to the MCEL server, and provide for a consistent and accurate interpolation stencil which is used in the downscaling. Furthermore, in order to provide a consistent upscaling model, we construct a single grid which is comprised of the control volumes which WRF would use, if it's numerics are considered in a finite-volume framework.

4. DOWNSCALING DEMONSTRATION

The first coupled model runs we intend to make correspond to the time period during Intensive Operating Period 6 (IOP 6) of the Joint Urban 2003 Field Test, conducted in July, 2003 in Oklahoma City, OK [Allwine and Flaherty, 2006]. The WRFV2.2/Noah/UCM modeling system incorporating anthropogenic heating has been used to simulate the JU2003 IOP06 [Chen et al., 2006]. Model integration starts at 1200UTC 16JUL2003 (0700CDT) and performs 12-hour forecast with 5 two-way nested domains with the following grid spacing and number of grids: D1: 40.5km(90x90x38), D2: 13.5km(100x100x38), D3: 4.5km (100x100x38), D4: 1.5km (100x100x38), and D5: 0.5km (100x100x38). Dudhia shortwave radiation scheme, RRTM longwave radiation scheme, MYJ PBL scheme, WSM 6-class graupel scheme and Noah land surface model with one-layer urban canopy model are used for all domains, Kain-Fritsch cumulus parameterization scheme for D1 and D2.

MCEL provides a data caching capability, wherein all the data sent to the server may be stored on disk (cached) and used subsequently by other models. We have used this capability here to test the downscaling coupling approach by using cached WRF-MCEL data to perform downscaling onto the actual CFD model grids, and have cycled within the solver itself in order to test the downscaling interpolation. At this point in time, the CFD solver is not complete, so this demosntration only tests the interpolation technique. Figure 4.1 shows a particular timestamp where the WRF data has been downscaled to an adaptive quadtree/prismatic grid through the (cached data) downscaler.

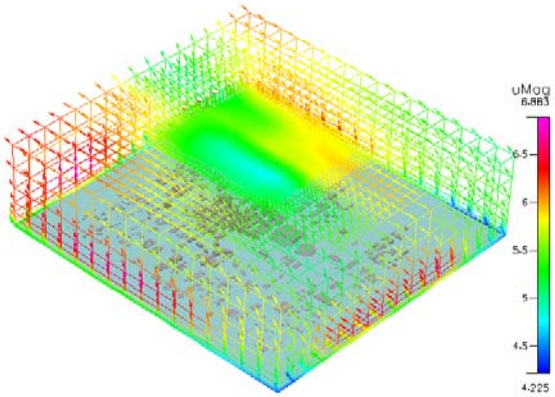

Figure 4.1: Downscaled WRF data upon the CFD mesh for a particular time during IOP6 of the JU2003 Test.

5. CFD MODEL DESCRIPTION

The CFD model under development solves the Reynolds-Averaged Navier-Stokes (RANS) equations discretized in a cell-centered, finite volume framework using an implicit, fully-coupled, low-Mach number preconditioned approach. We solve the equations fully-coupled via a parallelized Newton-Krylov (NK) procedure that is contained within the PETSc parallel, sparse matrix library. PETSc (Portable, Extensible Toolkit for Scientific computation) [PETSc] provides data structures and routines for scalable (parallel) solution of scientific applications modeled by partial differential equations. The routines and class-like structures provided with PETSc consist of linear and nonlinear solvers, preconditioners, time stepping for solving time-dependent PDEs, and many others.

In order to take advantage of this powerful library, we have designed the CFD solver from the beginning to be PETSc-compliant, where we use the SNES (Scalable Non-linear Equations Solver) routines to solve the fully-coupled equations. We have designed the flow solver data structures, memory access, connectivity ordering and loop coloring to reduce cache misses using data interleaving, and embed the PETSc classes and constructs within the C++ classes used to develop the core CFD solver routines.

Based upon our experiences with pressure-based and characteristics-based parallel CFD solvers we have concluded that a fully-coupled, low-Mach number preconditioned approach, using the parallel NK strategy, shows the greatest promise to achieve the scalability and rapid cycling rates required by the

coupled models. The low-Mach number preconditioning approach alleviates the stiffness of the coupled equations caused by the differences in the acoustic eigenvalues, by scaling (preconditioning) the wave speeds of a modified system of equations. This, in essence, degrades time-accuracy in exchange for rapid steady-state convergence, although high-order time accuracy can be easily recovered using a dual time stepping approach. This preconditioned characteristics-based approach is surmised to be more efficient than a pressure-correction based approach, which requires iterative Poisson/elliptic equations solves which often do not scale well.

There are many different types of preconditioners being advanced, which can be shown to all be some variant of Chorin's scheme [Chorin]. The approach we use here is based upon the preconditioner proposed by Weiss and Smith [Weiss, et al., 1996], which casts the conservation form of the Navier-Stokes equations as:

$$\int \Gamma \frac{\partial Q}{\partial t} dV = \oint (-\vec{F} + \vec{F}_{vis}) \circ \hat{n} dS \quad (1)$$

where instead of conserved variables, we use primitive variables (Q), consisting of (P,u,v,w,T). For a non-preconditioned system, the matrix pre-multiplying the time-derivative is simply the Jacobian transformation from conserved to primitive variables. The key to the preconditioning approach is to instead substitute a preconditioning matrix, Γ, whose elements are altered in order to scale (precondition) the eigenvalues of the hyperbolic system. We use the preconditioning matrix according to Weiss and Smith [Weiss, et al., 1996], also noted in Merkle [Merkle, et al., 1996], as:

$$\Gamma = \begin{pmatrix} \Theta & 0 & 0 & 0 & \rho_T \\ \Theta u & \rho & 0 & 0 & \rho_T u \\ \Theta v & 0 & \rho & 0 & \rho_T v \\ \Theta w & 0 & 0 & \rho & \rho_T w \\ \Theta H - 1 & \rho u & \rho v & \rho w & \rho_T H + \rho C_P \end{pmatrix} \quad (2)$$

According to Weiss and Smith [Weiss, et al., 1996], the system of equations can be preconditioned for all Mach numbers using:

$$\Theta = \left(\frac{1}{U_r^2} - \frac{\rho_T}{\rho C_P} \right) \quad (3)$$

In [Weiss, et al., 1996] a variable reference speed is used, but we have found that in practice a constant reference speed is more robust.

Since applying the preconditioning alters the hyperbolic system of equations wave speeds, care must be taken when using an upwind flux formulation for the convective terms in the Navier-Stokes equations. If untreated, these terms can introduce excessive dissipation at low Mach numbers. We use the approach outlined in [Weiss, et al., 1996], which constructs an upwind flux formula based upon Roe's Flux Difference Splitting [Roe], using the properly scaled (preconditioned) eigenvalues and eigenvectors. Viscous fluxes are found using gradients constructed using adjacent cells for gradients normal to a given face, and centered differences for those lying along a face. These centered differences are found using a flux-based gradient reconstruction technique [Coirier, et al., 1996]. We are presently implementing a k- model variant in order to construct the RANS turbulent fluxes.

After discretization using the finite-volume approach, the discrete system of equations to be solved is:

$$\left(\frac{\Gamma}{\Delta t} \Delta V - \frac{\partial R}{\partial Q} \right) \Delta Q^{n+1} = R^n \quad (4)$$

where the residual, R^n, is the discrete, flux-based representation of the inviscid and viscous, turbulent fluxes of mass, momentum and energy. In order to maximize robustness and convergence rate, the Jacobian of the residual, $\frac{\partial R}{\partial Q}$, is formed to match as closely as possible the explicit residual. Care is taken to properly represent the boundary conditions implicitly, and to include stabilizing terms from the viscous fluxes. The resulting coupled system of equations is solved using a Newton-Krylov approach via the PETSc SNES system. This Newton-based iteration uses a preconditioned Krylov solver to solve the resulting sparse system of linear equations to obtain the update, ΔQ^{n+1}.

6. VALIDATION RESULTS

To date we have run a limited series of laminar validation cases which are summarized here. The turbulence model is presently being integrated within the solver, and results from its validation will be reported in the future. The validation cases shown here correspond to the driven cavity results computed by Ghia [Ghia, et al., 1982] and the developing laminar flow over a flat plate [Schlichting].

6.1 Driven Cavity

The driven cavity computations performed by Ghia [Ghia, et al., 1982] have been used quite extensively to validate CFD models, where profiles of velocity at various locations are provided for Reynolds numbers based upon lid length and speed that range from 100 to 10,000. Figure 6.1 below shows contours of velocity magnitude computed for a Re=1000 driven cavity.

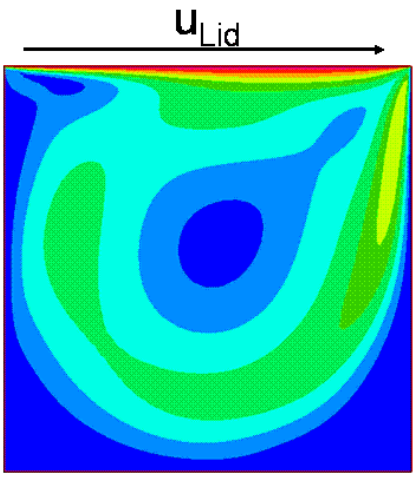

Figure 6.1: Velocity Magnitude Contours for a Re=1000, Driven Cavity.

We have performed the validation study using Reynolds numbers of 100, 400, 1000, 5000 and 10,000, and compare our computed results with the published results as well as those produced using CFD-Urban/CFD-ACE+ [Coirier, et al., 2005, 2006.b, ACE+. 2003]. In order to stress test the low-Mach number preconditioning scheme, we have run the Re=100, 400 and 1000 cases at Mach numbers of 0.25, 0.1, 0.01, 0.001 and 0.0001. All cases show an insensitivity of the results to Mach number, which validates the accuracy of the preconditioning scheme. The sections below illustrate the results where, unless noted otherwise, all results were computed using a second-order upwind reconstruction. Typically, convergence is achieved within 50 to 200 iterations, while for first-order upwind calculations convergence to machine zero was often achieved in 5 to 20 iterations. This rapid convergence rate is one of the primary goals of using a fully-coupled approach.

6.1.1: Re=100 Driven Cavity

The computed results are compared to those of Ghia [Ghia et al., 1982] for a Re=100 based upon lid length and speed. The lid Mach number is varied by changing the laminar viscosity for the Mach numbers of 0.25, 0.1, 0.01, 0.001 and 0.0001. For this Reynolds number, there is essentially no difference between the results shown for the different lid Mach numbers, as well as the results produced by CFD-ACE+. For proper choices of reference velocity, the results were obtained faster than CFD-ACE+ by a factor of 2 to 10.

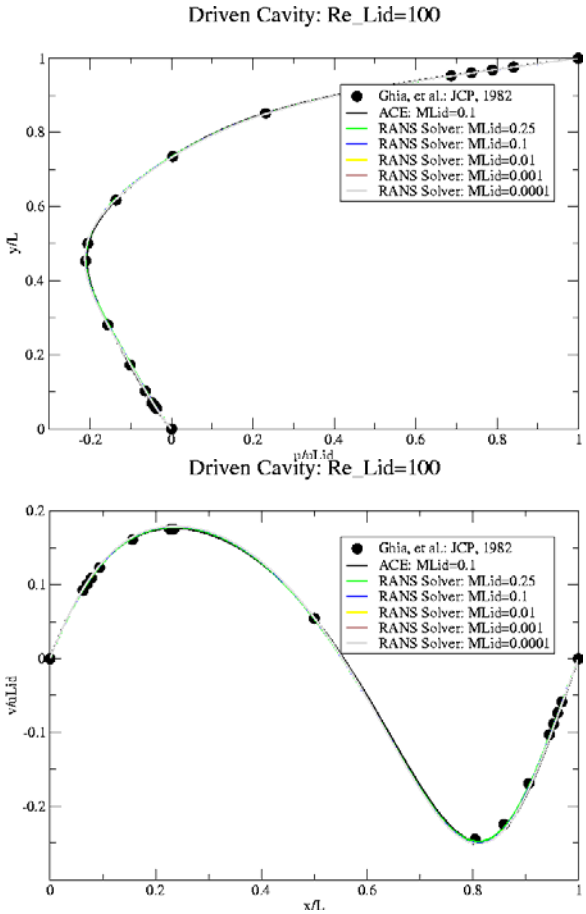

Figure 6.2: Re=100 Driven Cavity Results. u-velocity component vs. y (top) and v-velocity component vs x (bottom) midway through the cavity.

6.1.2: Re=400 Driven Cavity

The computed results are compared to those of Ghia [Ghia et al., 1982] for a Re=400 based upon lid length and speed. The lid Mach number is varied by changing the laminar viscosity for the Mach numbers of 0.25, 0.1, 0.01, 0.001 and 0.0001.

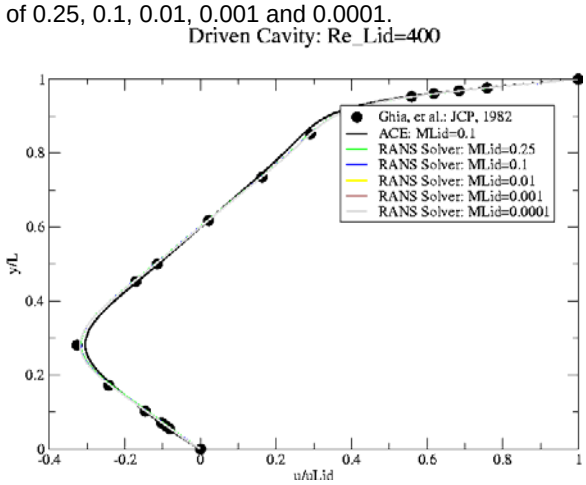

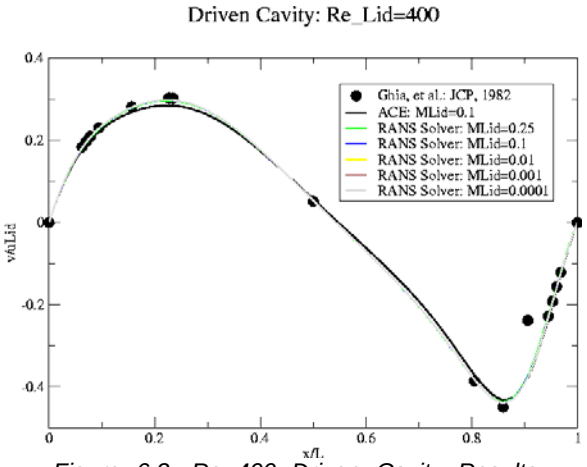

Figure 6.3: Re=400 Driven Cavity Results. u-velocity component vs. y (top) and v-velocity component vs x (bottom) midway through the cavity.

6.1.3: Re=1000 Driven Cavity

The computed results are compared to those of Ghia [Ghia et al., 1982] for a Re=1000 based upon lid length and speed. The lid Mach number is varied by changing the laminar viscosity for the Mach numbers of 0.25, 0.1, 0.01, 0.001 and 0.0001.

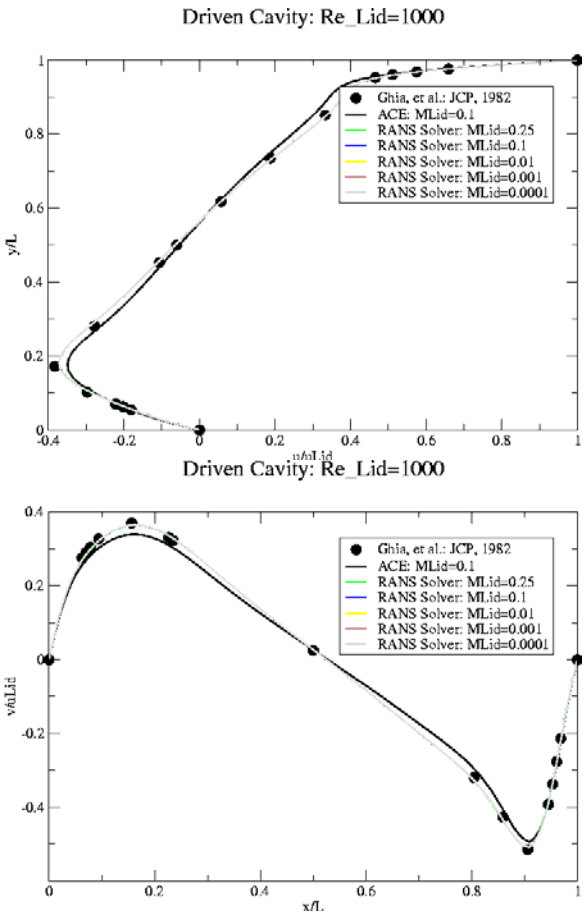

Figure 6.4: Re=1000 Driven Cavity Results. u-velocity component vs. y (top) and v-velocity component vs x (bottom) midway through the cavity.

6.1.4: Re=5000 Driven Cavity

The computed results are compared to those of Ghia [Ghia et al., 1982] for a Re=5000 based upon lid length and speed. For this case, we compare only for a single Mach number to Ghia's results and to those computed with CFD-ACE+/CFD-Urban.

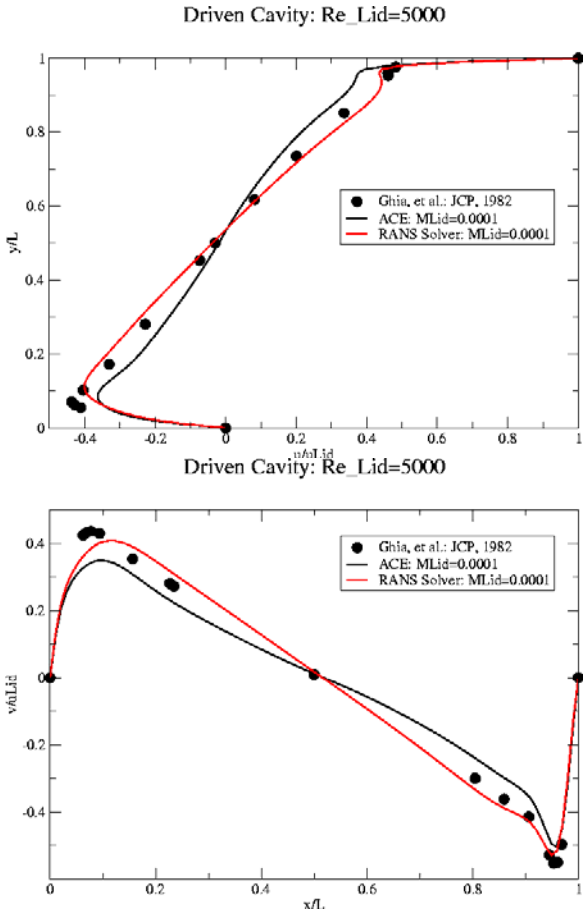

Figure 6.5: Re=5000 Driven Cavity Results. u-velocity component vs. y (top) and v-velocity component vs x (bottom) midway through the cavity.

6.1.5: Re=10,000 Driven Cavity

The computed results are compared to those of Ghia [Ghia et al., 1982] for a Re=10,000 based upon lid length and speed. For this case, we compare only for a single Mach number to Ghia's results and to those computed with CFD-ACE+/CFD-Urban. We compute the solutions on two different grids in order to evaluate the effect of increased mesh resolution upon the solution quality. The two grids, and the contours of velocity magnitude, are shown below, where the coarse grid consisted of 5000 cells and the fine consisted of 10,000 cells.

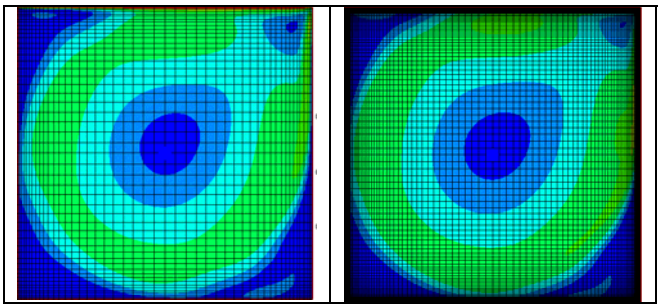

Figure 6.6: Coarse grid (left) and fine grid (right) velocity magnitude contours, Re=10,000.

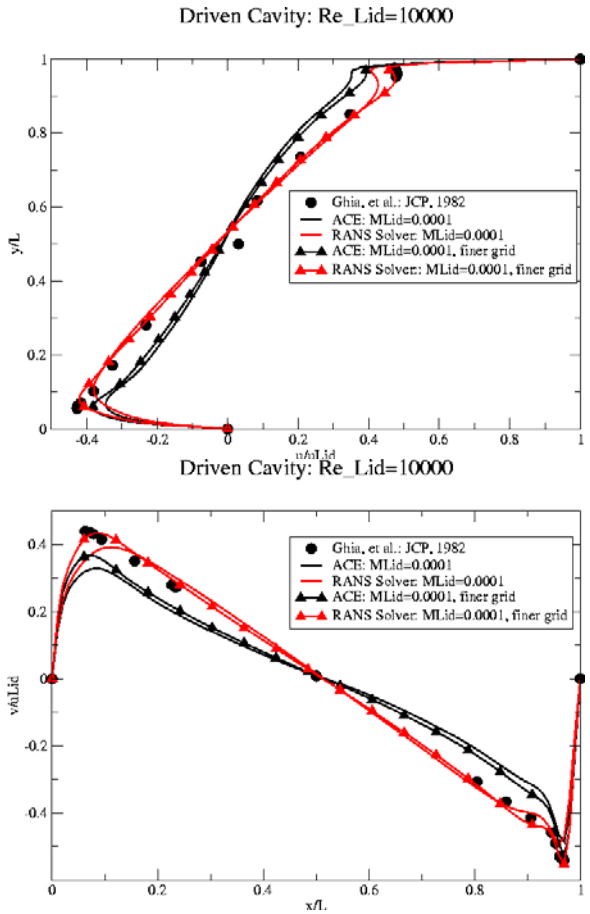

Figure 6.7: Re=10,000 Driven Cavity Results. u-velocity component vs. y (top) and v-velocity component vs x (bottom) midway through the cavity.

6.2: Laminar Boundary Layer

The developing laminar flow over a flat plate is also often used for CFD validation. The analytical solution is found by casting the momentum equations in terms of a stream function written in a similarity variable, which reduces the partial differential equations into a single ordinary differential equation, which is solved numerically. The tabular solution of the stream function and it's 1st and 2nd derivatives are then used to compare the predicted to analytical solutions. This classical solution is often called the Blasius solution and can be found in many textbooks [Schlichting]. The validation study we conducted is for a Reynolds number based on plate length of 10,000. Immediately upstream of the plate is a symmetry boundary, and the flow upstream is at a constant velocity of 10 m/s. Using both first and second-order upwind schemes, we have computed the flow field, and compared the u- and v-components written in terms of similarity variables, shown in Figure 6.8.

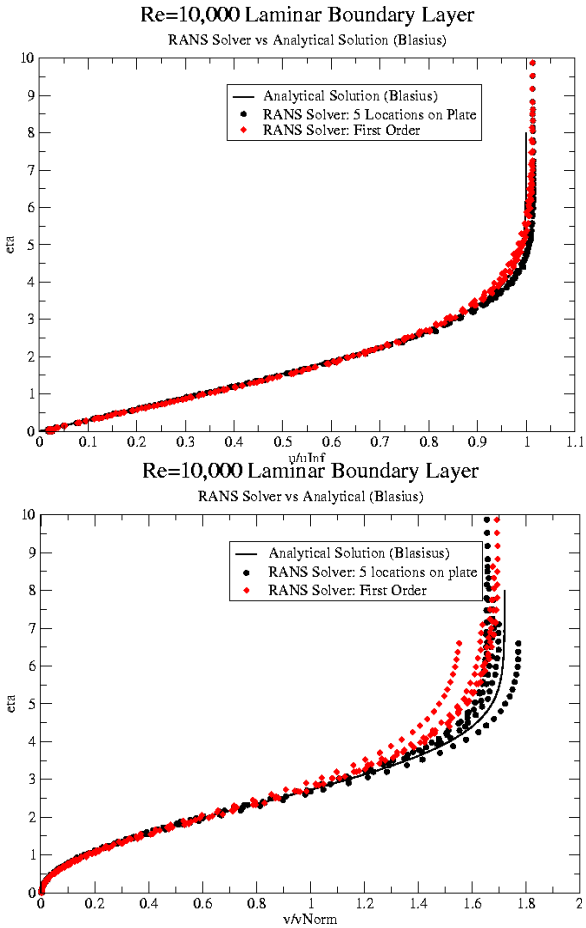

Figure 6.8: Comparison of analytical (solid line) and computed results for the streamwise (top) and vertical (bottom) velocity components.

Figure 6.9 shows the convergence rate for both the first- and second-order upwind schemes using the fully-coupled, low-Mach number preconditioned approach. As is evident from the figure, when the implicit and explicit discretizations are matched, which is achieved when the first-order upwind scheme is used, the convergence rate can be extremely fast. The first-order upwind discretization converged to machine zero in 12 iterations.

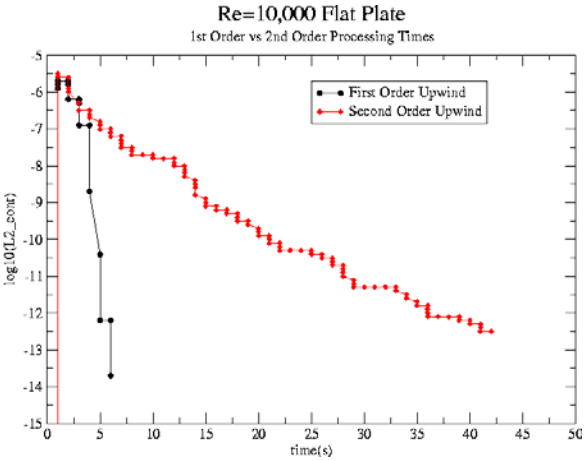

Figure 6.9: Convergence histories for the flat plate study, first-order (black) and second-order (red) upwind schemes.

7. PARALLEL SPEEDUP STUDY

A limited parallel performance study has been performed on a 24 node dual processor linux cluster. The driven cavity flow at a Reynolds number of 1000 has been solved to convergence on computational meshes of 40,000 and 160,000 cells. As noted in Section 5, we use PETSc's SNES to solve the primitive variables in a fully-coupled implicit fashion. For the NK solve, we use a GMRES linear solver and a block Jacobi preconditioner.

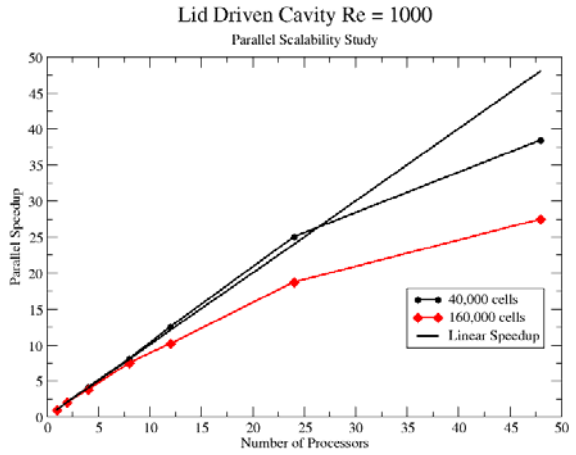

Figure 7.1: Parallel speed up study results.

In this study, a sub-linear speedup (18 times speedup on 24 processors) has been observed for the total solve time for the finer mesh case, although a near perfect speedup is seen for the same number of processors for the smaller mesh case.. A tail off in speedup is observed for 48 processors. It is a known artifact of using MPI on multi-core architecture that all the cores on the chip use the same memory bus, thereby increasing the communication time. We surmise that this is happening here, as the tail-off begins at 24 processors (of our 48 processor cluster). We should note that at this time the code is still under construction and has not been optimized yet, and that better performance may be attained using a different matrix preconditioner.

8. CONCLUSIONS

In this paper we have presented progress towards the development of a coupled micro-to-mesoscale modeling capability using the MCEL library to couple the WRF model and a new CFD model that is presently being developed. The Model Coupling Environmental Library, MCEL, utilizes a data flow approach where coupling information is stored in a centralized server and flows through processing routines called filters to the numerical models which represent the clients. In order to take advantage of the specialized data structures associated with the CFD solver, we have moved the filter between MCEL and the CFD model to lie within the CFD model itself.

The CFD model solves the Reynolds-Averaged Navier-Stokes (RANS) equations discretized in a cell-centered, finite volume framework using an implicit, fully-coupled, low-Mach number preconditioned approach. We solve the equations fully-coupled via a parallelized Newton-Krylov (NK) procedure that is contained within the PETSc parallel, sparse matrix library. The convective terms are discretized using a preconditioned variant of Roe's Flux Difference Splitting, and the low Mach number preconditioning of Weiss and Smith is used to remove the stiffness associated with the disparate eigenvalues at low Mach numbers.

The solver has been validated for a series of laminar, driven cavity flows at a range of Reynolds and Mach numbers, and has been shown to be accurate and efficient. We compare the computed results with the accepted computational results of Ghia, and to results computed by CFD-ACE+/CFD-Urban. In general, the results are more accurate as well as more efficient then CFD-ACE+/CFD-Urban, and match very well with Ghia's results. A parallel speedup study for the driven cavity shows reasonable scalability with some room for improvement. The scalability of the new solver is due to the choice of numerics, flow solver data structures and use of the PETSc framework.

9. ACKNOWLEDGEMENTS

This work is funded under a Phase II SBIR grant sponsored by the Defense Threat Reduction Agency, Technical Monitor, CDR Stephanie Hamilton.

10. REFERENCES

ACE+: CFD-ACE+, V2003 Users Manual, Volumes 1 and 2.

Allwine, K.J., and J.E. Flaherty: 2006, Joint Urban 2003: Study Overview and Instrument Locations. PNNL-15967, Pacific Northwest National Laboratory, Richland, WA.

http://www.pnl.gov/main/publications/external/technical_reports/PNNL-15966.pdf

Chen F., M. Tewari, H. Kusaka, and T. T. Warner: 2006, Current status of urban modeling in the community Weather Research and Forecast (WRF) model, paper presented at Joint Session with Sixth Symposium on the Urban Environment and AMS Forum: Managing our Physical and Natural Resources: Successes and Challenges, The 86th AMS Annual Meeting,

Chorin, A.J., 1967, "A Numerical Method for Solving Incompressible Viscous Flow Problems", Journal of Computational Physics, **2**, pp. 12-26.

Coirier, W.J., Jorgenson, P.C.E., 1996, "A Mixed Volume Grid Approach for the Euler and Navier-Stokes Equations", AIAA-96-0762 (NASA TM 107135).

Coirier, W.J., Fricker, D.M., Furmanczyk, M., and Kim, S.: 2005. "A Computational Fluid Dynamics Approach for Urban Area Transport and Dispersion Modeling", Environmental Fluid Mechanics, 5, pp 443-479.

Coirier, W.J., Kim, S., Chen, F., Tuwari, M.,: 2006. Evaluation of urban scale contaminant transport and dispersion modeling using loosely coupled CFD and mesoscale models, American Meteorological Society, 6th Symposium on the Urban Environment.

Coirier, W.J., Kim, S., 2006.b: CFD Modeling for Urban Area Contaminant Transport and Dispersion:Model Description and Data Requirements, American Meteorological Society, 6th Symposium on the Urban Environment.

Ghia, U., Ghia, K.N., Shin, C.T., 1982, "High-Re Solutions for Incompressible Flow Using the Navier-Stokes Equations and a Multigrid Method", Journal of Computational Physics, 48, pp. 387-411.

Merkle, C.L., Sullivan, J.A.Y., Buelow, P.E.O., 1996, "Computation of Flows with Arbitrary Equations of State", AIAA-96-0680.

Michalakes, J., M. Bettencourt, D. Schaffer, J. Klemp, R. Jacob, and J. Wegiel, 2003. Infrastructure Development for Regional Coupled Modeling Environments Parts I and II, Final Project Reports to Program Environment and Training Program in the U.S. Dept. of High Performance Computing Modernization Office, Contract No. N62306-01-D-7110/CLIN4, 2003 and 2004. See also: http://www.mmm.ucar.edu/wrf/WG2/software_2.0/io.pdf

PETSc, Available on the World Wide Web at http://www.mcs.anl.gov/petsc

Roe,P, 1981, "Approximate Riemann Solvers, Parameter Vectors and Difference Schemes," *Journal of Computational Physics*, Volume 43.

Schlichting, H, "Boundary-Layer Theory", 7th Edition, McGraw-Hill Publishing.

Weiss, J.M., Smith, W.A., 1996, "Preconditioning Applied to Variable and Constant Density Flows", AIAA Journal, 33 (11), pp. 2050-2057.