62

Parallel multi-grid for turbulent reacting flow simulations

A.Twerda[1], R.L.Verweij[2], T.W.J. Peeters[2] & A.F. Bakker[2]

Introduction

This paper reports on the implementation of a multi-grid acceleration technique in a computer code FURNACE, which simulates turbulent combustion in large-scale furnaces. The need for accurate predictions of these furnaces is still increasing, due to strict government regulations. Simulations of large-scale industrial furnaces require fine meshes and accurate models, which makes these simulations very CPU and memory demanding. Convergence becomes troublesome, especially in combination with domain decomposition. Full multi-grid with full approximation scheme is implemented to accelerate the convergence [Wes92, PSS92].

The implementation of the multi-grid is validated by a well-known CFD problem: a laminar lid-driven cavity flow. Next, results for a furnace simulation will be presented. These numerical examples show that multi-grid improves the convergence behaviour and can be applied successfully in conjunction with domain decomposition.

Mathematical model

The conservation equations of mass, momentum, energy and species are applied to describe the turbulent reacting flow. Favre averaging is applied to obtain density averaged equations. The Favre-averaged incompressible (variable density) Navier-Stokes equations are solved for *conservation of momentum and mass.* The standard high-Reynolds number $k-\varepsilon$ turbulence model is used. Wall-functions are applied to bridge the low-Reynolds number region near

1 Department of Applied Physics, Delft University of Technology, Lorentzweg 1, 2628 CJ Delft, The Netherlands, E-mail : A.Twerda@tn.tudelft.nl
2 Address same as above.
Eleventh International Conference on Domain Decomposition Methods
Editors Choi-Hong Lai, Petter E. Bjørstad, Mark Cross and Olof B. Widlund

the walls [Rod84]. For the *conservation of energy* the transport equation for the enthalpy is solved. The *conservation of species* is modelled with a conserved scalar approach. Here, the concentrations of all species are directly coupled to the mean mixture fraction, which varies from 1 in the fuel-inlet to 0 in air-inlet. Finally, to close the set equations, an equation of state is needed. For this the ideal gas law is used. Because of the strong non-linearity of the thermo-chemical quantities on the mean mixture fraction also the mean mixture fraction variance is needed, and a β- probability density function is used for computing the mean values of these quantities. The chemistry is modelled with an constraint equilibrium model. [Pos88] and [Boe97] give more details about applied models and methods for furnace simulations.

Numerical model

The set of equations described in the previous section is discretized using the *Finite Volume Method* (FVM). The computational domain is divided in a finite number of control volumes (CVs). A cell-centered colocated Cartesian grid arrangement is applied [FP95]. As the equations are solved in the integral form, one must discretize the fluxes through the cell faces. The diffusive fluxes are approximated with the central difference scheme. For laminar flows, the convective fluxes are approximated with the central difference scheme. For turbulent (reacting) flow, a first order upwind scheme is used. For pressure velocity coupling the SIMPLE scheme is applied [Pat80, FP95]. The porosity method is used to match the Cartesian grid to the geometry of the furnace [Pos88].

Domain Decomposition (DD) with minimal overlap is used as parallelisation technique. This technique has proven to be very efficient for creating a parallel algorithm [VTP98]. A grid-embedding technique with static load balancing is applied to divide the global domain into sub-domains (blocks) [CPC91], assigning one processor to every sub-domain. Message passing library MPI is used for communication [GL96].

Multi-grid method

One of the drawbacks of DD is, as implicit solvers are used to solve the discretized set of equations, that the convergence deteriorates as the number of domains increases [PSS92]. Also if the number of CVs increases the convergence becomes more difficult. To minimize both effects, a multi-grid algorithm is implemented. The multi-grid method is applied over the set of equations, which means that *all* equations are solved per prolongation- and restriction-step, rather than applying multi-grid for each equation separately. The Full Approximation Scheme (FAS) is implemented with tri-linear interpolation for restriction and prolongation operators. The V-cycle is applied to determine the sequence in which the grids are visited. It is efficient to start with a good initial guess. A solution on a coarser grid is used and interpolated onto the finer grid. This leads to the Full Multi-Grid method (FMG) [Wes92].

As mentioned in [FP95], the value of some turbulent quantities, like the eddy viscosity, can differ several orders of magnitude over the domain. Computing them directly on the coarse grid could destabilize the multi-grid algorithm. Therefore, these variables are only computed on the finest grid-level and interpolated on the coarser grids. During a multi-grid cycle they are kept constant. Near the solid walls, the variable ε is calculated from wall-functions. This gave rise to large corrections on the coarser grid which destabilised the multi-grid iteration. Therefore, ε is not corrected near a solid wall, but of the fine grid value is interpolated on the coarse grid.

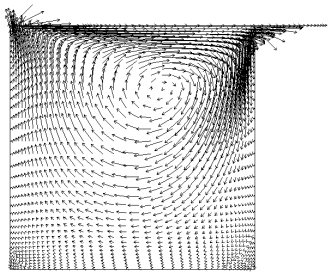

Figure 1 Velocity profile in the $y = 0.5$m plane of the laminar lid-driven cavity with $Re = 100$.

Results

Description of the test cases

The FURNACE code has been validated on a well known CFD test problem: a 3D laminar lid-driven cavity [TF89]. The Reynolds number is based on the uniform top-lid velocity $U_{lid} = 1$m/s, the size of the cavity $L = 1$m and the kinematic viscosity $\nu = 0.01$s/m^2. Thus, $Re = \frac{U_{lid} L}{\nu} = 100$. Figure (1) shows the characteristic velocity profile of the cavity in the $x-z$ plane for $y = 0.5$m. The number of CVs applied are 4^3 to 64^3 with respectively 2 to 6 levels in the multi-grid calculations. Thus, the coarsest grid-level consisted of 2^3 number of CVs. The residual, $r_h = \frac{\|A_h \phi_h - b_h\|}{\|b_h\|}$ and the relative change of each of the velocity components in a monitor point are used as convergence criteria. The residual and the relative change must be smaller than 10^{-5} for convergence. The first, is a measure for the approximation of the discretized equations, the latter is a measure for the coupling of the equations with respect to each other.

Next, results of a full furnace simulation are presented. The furnace computed is the IFRF glass-melting furnace at IJmuiden (0.440m $\times$ 4.09m $\times$ 0.955m). The inlet consist of two parts, the pre-heated air-inlet $T = 1400$K, $v_{air} = 10$m/s, and a small gas-inlet $T = 370$K, $v_{gas} = 125$m/s. [Boe97] gives more detailed results of these simulations. The geometry of the furnace is plotted in figure (2 left). Because symmetry boundary conditions are applied in the negative x-direction, only half of the furnace is computed. The coarse grid contains ($16 \times 24 \times 20$) CVs, which is shown in figure (2 right). The number of grid points for the coarsest grid could not be reduced due to the small radius of the fuel inlet (0.06m). Furthermore, the relative thin flame front would not be correctly captured by a coarser grid. In figure (3) a contour plot of the mean mixture fraction in the symmetry plane of the furnace is shown. On the left, the fuel- and air-inlet, are clear visible. In the outlet, on the right, the flow has almost perfectly mixed.

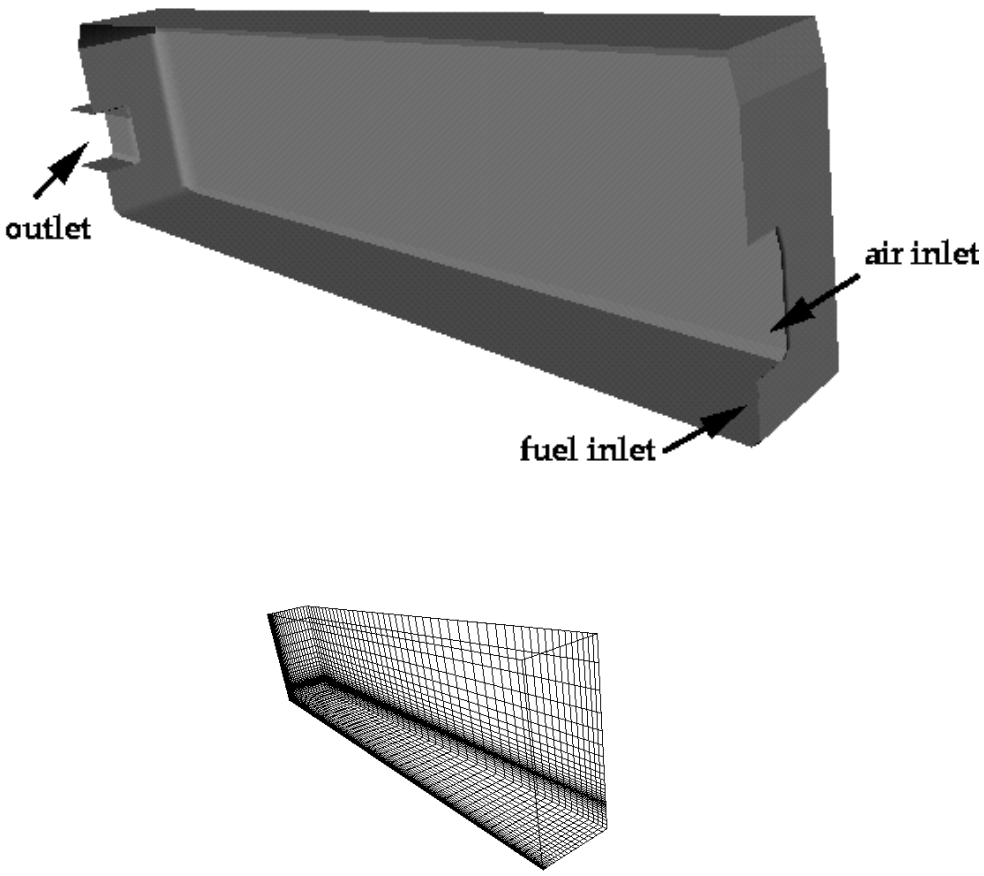

Figure 2 **Left**: Geometry of the IFRF Furnace. **Right**: Coarse grid used for the simulations.

Figure 3 Contour plot of mean mixture fraction in the symmetry plane, of the IFRF furnace. Contour lines denote values of 0.1 to 0.5.

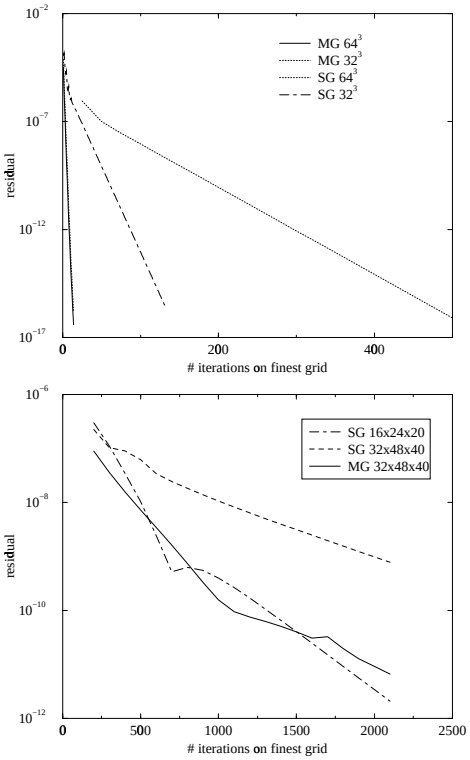

Figure 4 Convergence history of single domain multi-grid and single grid. **Left**: Residual of the velocity components for laminar flow. **Right**: Residual of the mean mixture fraction for a furnace simulation.

Influence of multi-grid on the number of iterations

The residual is plotted in figure (4) versus number of iterations performed on the finest grid-level for both the laminar test case and the furnace. For the cavity flow (left) the residual of the velocity components is plotted. The multi-grid (MG) algorithm converges much faster than the single grid (SG) algorithm. Moreover, the convergence of the MG-algorithm is independent of the number of CVs used, which is the main reason for using the algorithm. The main part of the computation time is spent in calculations on the finest grid, so the acceleration in execution time is of the same order as shown in figure (4), a factor of 25 for the 64^3 grid.

In figure (4 right) the convergence behaviour for the mean mixture fraction equation is plotted. The MG converges better than the SG solver for the finer grid. Also the convergence rate of the finer grid does not differ much from the single coarse grid computation. The acceleration of the MG-algorithm is now approximate a factor of 3, which is less than for the laminar case, but the set of equations for these simulations are more non-linear, which makes the MG-algorithm less efficient.

Influence of number of sub-domains on the multi-grid

The impact of the number of sub-domains is investigated for both cases, using 1 to 16 processors on a CRAY T3E. As one processor is used for every sub-domain, the number of processors can be read where the number of blocks or sub-domains are plotted or mentioned.

For the laminar case, the time per MG-cycle (left) and the number of MG-cycles needed for convergence (right) are plotted in figure (5) for grid sizes 4^3 to 64^3. The latter was only computed with 8 and 16 blocks due to memory restrictions. Because a grid on 1 block consists of at least 1 grid-point, the coarsest grid used is now 4^3. Thus, the number of grid-levels used for the MG-algorithm is 1 for 4^3 grid to 5 for the 64^3 grid. It takes approximately the same time to perform one MG-cycle on a 32^3 grid using one block and on 64^3 grid using 8 blocks. This means that the algorithm is suitable for large grid-sizes because, if the number of blocks increases with problem size, the CPU time per MG-cycle time will remain approximate the same. The number of iterations slowly increases with the number of blocks because implicit solvers are used.

For the 32^3 grid the number of grid-levels used in the MG-algorithm has been varied. The number of iterations needed for convergence (left) and the total execution times (right) are plotted in figure (6). There is an improvement in execution time of a factor 24, from approximate $2900s$ on 1 block without MG, to $120s$ with 8 blocks using 3 levels of multi-grid. A factor of 4 originates from the parallelisation and another factor 6 is gained by the MG-algorithm. Even for the case where no multi-grid is applied, the deterioration of the convergence when more blocks are used is small. This can be explained by the fact that the pressure correction equation in the SIMPLE scheme is solved several times (typically 4). This improves the convergence, especially as DD is applied, as has been shown by [Ver99]. There is a trade-off between MG-levels used in the MG-algorithm and the number of blocks, because the iterations on the coarser grid-levels do not scale properly with large number of blocks, as already shown in figure (5 left).

Finally, the furnace is computed with several sub-domains. The converge history is plotted using 8 and 16 blocks, with two levels of MG being applied. There is a slight influence of the DD on the convergence rate. The single grid computation, with the same number of iterations performed on the fine grid, is especially inferior to the MG computations in the beginning of the iterations, but also in the latter stage the convergence rate is less. Also with more blocks used, the factor gained by using MG is a factor of 3. The time per MG/SG-cycle is shown in table (1). The MG does not increase the time per cycle considerable. The time spend in on

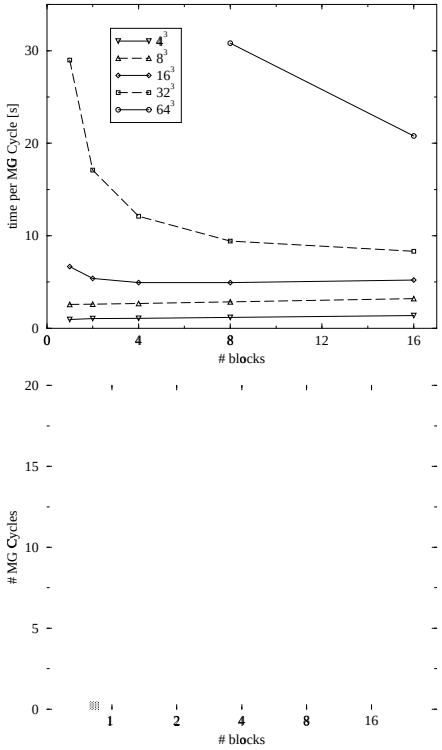
4^3
8^3
16^3
32^3
64^3
time per MG Cycle [s]
blocks
MG Cycles
blocks

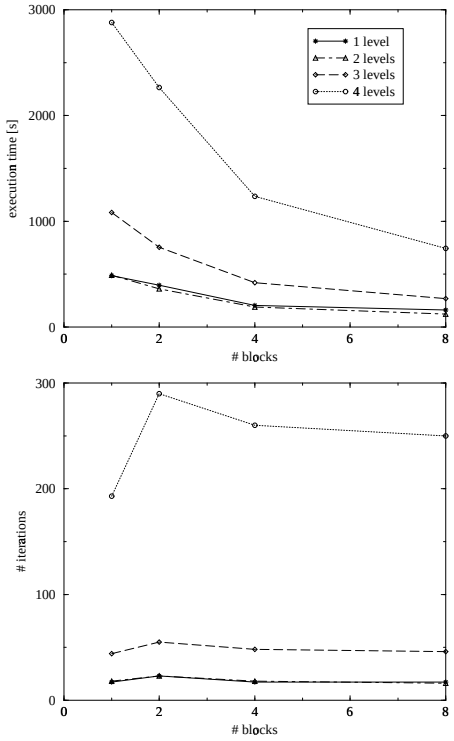

Figure 6 Execution time (**left**) and number of iterations (**right**) for different number of grid levels applied versus the number of blocks

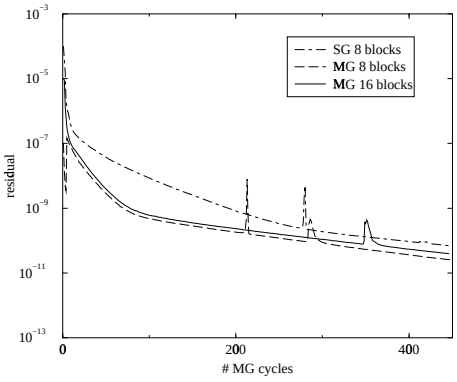

Figure 7 Convergence history of the mean mixture fraction equation for multiple blocks for the furnace simulation.

the coarse grid can be neglected compared with the time spend on the fine grid. The speed-up using 16 in stead of 8 blocks is 1.64 which is moderate and could be matter of concern when using larger number of blocks.

Conclusions

The multi-grid algorithm as described in this paper, has been implemented in the FURNACE code. It has proven to be a good acceleration technique for laminar flow simulations. For turbulent reacting flow simulations an acceleration of a factor 3 for only two grid-levels has been measured.

The multi-grid algorithm also improves the convergence behaviour when domain decomposition is applied. The algorithm is scalable, meaning that if the number of blocks is scaled with the number of control volumes, the time per MG-cycle will stay constant. It can be efficient to use less MG-levels on more blocks, because the iterations on the coarser levels are not scalable to higher number of blocks. The parallel speed-up is moderate, and must be improved if a larger number of blocks is needed.

Acknowledgments

The authors would like to thank the HPαC centre for the computing facilities.

REFERENCES

[Boe97] Boerstoel G. (1997) *Modelling of gas-fired furnaces.* PhD thesis, Delft University of Technology.

[CPC91] Coelho P., Pereira J., and Carvalho M. (1991) Calculation of laminar recirculating flows using a local non-staggered grid refinement system. *Int. J. Num. Meth. Fl.* 12: 535–557.

[FP95] Ferziger J. and Perić M. (1995) *Computational Methods for Fluid Dynamics.* Springer-Verlag, New York.

[GL96] Gropp W. D. and Lusk E. (1996) *User's Guide for* `mpich`*, a Portable Implementation of MPI.* Mathematics and Computer Science Division, Argonne National Laboratory. ANL-96/6.

[Pat80] Patankar S. (1980) *Numerical heat transfer and fluid flow.* McGraw-Hill, London.

[Pos88] Post L. (1988) *Modelling of flow and combustion in a glass melting furnace.* PhD thesis, Delft University of Technology.

[PSS92] Perić M., Schäfer M., and Schreck E. (1992) Computation of fluid flow with a parallel multigrid solver. *Parallel Computational Fluid Dynamics* pages 297–312.

[Rod84] Rodi W. (1984) *Turbulence models and their application in hydraulics–a state of the art review.* IAHR, Karlsruhe.

[TF89] Thompson M. and Ferziger J. (1989) An adaptive multigrid technique for the incompressible navier stokes equations. *Journ. Comp. Phys.* 82: 94–121.

[Ver99] Verweij R. (1999) *Parallel computing for furnace simulations using domain decomposition.* PhD thesis, Delft University of Technology.

[VTP98] Verweij R., Twerda A., and Peeters T. (1998) Parallel computing for reacting flows using adaptive grid refinement. In *Tenth International Conference on Domain Decomposition Methods, Boulder, Colorado, USA.*

[Wes92] Wesseling P. (1992) *An introduction to multigrid methods.* John Wiley & Sons, Chichester.