

# Applications of Multigrid Techniques in CFD

Peter K. Jimack\*

*School of Computing, University of Leeds, Leeds, LS2 9JT, UK.*

## SUMMARY

We present a discussion of the application of multigrid techniques to a range of nonlinear problems in CFD that require state-of-the-art numerical methods to be applied in order to obtain accurate and efficient solutions. These include combining multigrid with mesh adaptivity for phase-change problems, a parallel implementation of multigrid for a thin film flow problem, and a p-version of multigrid for the high order discontinuous Galerkin solution of a problem arising in elasto-hydrodynamic lubrication. Issues associated with combining multigrid with both adaptivity and parallelism are also discussed. Copyright © 2007 John Wiley & Sons, Ltd.

KEY WORDS: Multigrid; adaptivity; parallel algorithms; high order methods

## 1. INTRODUCTION

The efficient and reliable solution of the algebraic equation systems that arise from the discretization of partial differential equations (PDEs) lies at the heart of almost all CFD software. Furthermore, as processor speeds and memory density have increased with time, the users of such software have demanded greater resolution in their models, leading to ever-larger algebraic systems. Meanwhile, over the past thirty years, the field of multigrid and multilevel methods has developed significantly, initially for the solution of linear algebraic systems, [3], and then for nonlinear systems of equations, [21]. For many flow problems and discretizations the use of multigrid is now the natural choice therefore, [19, 7].

The paper begins with a very brief discussion of the main ideas behind the multigrid approach, before going on to consider a number of challenges and difficulties that appear to prevent the more widespread use of multigrid within CFD. In sections 3 to 5 we discuss a selection of these issues, arising primarily in three different contexts. These are the combination of multigrid with adaptive meshing, parallel algorithms and high order discretizations respectively. In each case a challenging CFD application is presented and the issues associated with developing an efficient multigrid solution method are discussed. The paper concludes with a discussion of how well our proposed variants of multigrid have performed, and provides a view on some of the main open problems that remain to be solved.

---

\*Correspondence to: pkj@comp.leeds.ac.uk

## 2. A BRIEF INTRODUCTION TO MULTIGRID

Clearly it is not possible to provide a full review of the multigrid method in this short section, and the reader is referred to texts such as [4, 21, 23] for more complete introductions. The purpose of these paragraphs is simply to provide an overview of those aspects of the approach that are of most relevance to the sections that follow.

The basic assumption that lies behind the approach is that, for a chosen discretization of a given partial differential equation (PDE), one is able to find an iterative solver that possesses the so-called “smoothing property”. This property says that, as the iterative method converges it does so in such a way that the higher frequency components of the error are eliminated more quickly than the lower frequency components: with the highest frequency components being eliminated the most quickly. After a small number of iterations on the finest mesh therefore, it is more efficient to solve an equation for the remaining error using a coarser mesh than it is to continue iterating on the fine mesh. For a linear problem this error equation simply involves the same discrete operator but with the current residual on the right-hand side. Hence, to solve for the remaining error on a coarse grid a restriction of the residual onto that grid is required (we use the terms “grid” and “mesh” interchangeably throughout this paper), along with a discretization of the PDE operator on that grid. When this equation is solved the computed error is then interpolated back onto the fine grid and the fine grid solution is corrected. This may then be followed by a small number of further iterations of the smoother on the fine grid.

The key idea behind the multigrid method, as the name suggests, is to make use of a sequence of coarser and coarser grids. This way, the error equation itself may be solved by taking a small number of iterations of the smoother and then correcting using a coarser level still. This process may be repeated recursively all the way down to the coarsest grid, for which an accurate solution will be very inexpensive. For many linear PDE problems it is quite easy to show that, with an appropriate smoother, this process is optimal in the sense that the computational effort to solve on a finest mesh with  $N$  unknowns is  $O(N)$ , however large the value of  $N$ , [23].

Unfortunately, in computational fluid dynamics it is unusual for linear systems of PDEs to provide satisfactory models for flows of practical interest, even when simplifying assumptions can be made. For example, in each of the examples that appear in the following three sections the governing equations are highly nonlinear. There are two ways in which this difficulty may be overcome. The first of these is to apply a standard nonlinear solution technique, such as Newton’s method (or some variant), and then apply the above multigrid approach to solve the linearised systems of equations at each nonlinear iteration. This can work very well if a sufficiently accurate initial guess can be found since Newton’s method typically converges in a very small number of iterations, and each of these outer iterations can be solved optimally using the multigrid scheme. In this paper we focus on the other standard approach for dealing with nonlinear problems however. Here we apply a variant of multigrid directly to the nonlinear problem itself: this is known as the full approximation storage (FAS) approach.

The FAS scheme requires a nonlinear iteration on each mesh to act as the smoother. After a small number of iterations both the current solution and the current residual must be restricted to the coarser mesh. The equation that is solved on this coarser mesh is not an equation for the error but instead provides an update of the restricted solution directly. This update is obtained by solving a discretization of the original operator on the coarser mesh with a suitably modified right-hand side. Multigrid is obtained by applying the update on a sequence of coarser and

coarser meshes, as with the standard approach. Full details of this technique may be found in [21] for example. We summarise the basic steps of the FAS algorithm as follows. In this notation the discrete equations on the the fine grid take the form  $\underline{F}^f(\underline{v}^f) = \underline{f}^f$  and on the coarse grid they are  $\underline{F}^c(\underline{v}^c) = \underline{f}^c$  (where  $\underline{f}^f$  and  $\underline{f}^c$  are the zero vectors on the fine and coarse grids respectively however, as can be seen below, a non-zero right-hand side is generated on all but the finest grids during the FAS process).

1. Choose initial fine grid solution  $\underline{v}^f$ .
2. Update  $\underline{v}^f$  applying  $\rho_1$  iterations of the nonlinear smoother.
3. Find the residual:  $\underline{r}^f := \underline{F}^f(\underline{v}^f) - \underline{f}^f$ .
4. Restrict the residual to the coarse grid:  $\underline{r}^f \rightarrow \underline{r}^c$ .
5. Restrict the solution to the coarse grid:  $\underline{v}^f \rightarrow \underline{v}^c$ .
6. Calculate the coarse grid right-hand side:  $\underline{f}^c = \underline{F}^c(\underline{v}^c) - \underline{r}^c$ .
7. Save the initial coarse grid solution:  $\underline{v}_{\text{init}}^c = \underline{v}^c$ .
8. Solve the system  $\underline{F}^c(\underline{v}^c) = \underline{f}^c$ .
9. Calculate the coarse grid correction:  $\underline{e}^c = \underline{v}^c - \underline{v}_{\text{init}}^c$ .
10. Interpolate the correction:  $\underline{e}^c \rightarrow \underline{e}^f$ .
11. Correct the fine grid solution:  $\underline{v}^f = \underline{v}^f + \underline{e}^f$ .
12. Update  $\underline{v}^f$  by applying  $\rho_2$  iterations of the nonlinear smoother.

For the FAS multigrid scheme step 8 of the above algorithm is solved using the same algorithm recursively, apart from at the coarsest mesh, where an exact solve is undertaken.

### 3. MULTIGRID IN COMBINATION WITH ADAPTIVE MESHING

The solutions to many challenging flow problems exhibit local features such as steep fronts, shocks or boundary layers. In order to capture these features accurately it is necessary to use a discretization scheme with a very fine spatial resolution in these regions, which may change with time. It is not generally practical, or efficient, to maintain such a high resolution over the entire computational domain and so some form of adaptivity is required. Perhaps the most popular technique for undertaking this adaptivity is to make use of local mesh refinement in the regions that require the greatest resolution, and to use a coarser mesh elsewhere. This is most frequently achieved through a hierarchical approach in which one starts with a coarse mesh and then refines this more in some regions than in others. In two dimensions this may be undertaken using a quad-tree data structure (e.g. [14]) and in three dimensions an oct-tree structure is used (e.g. [8]). In this section we demonstrate how this adaptive technique may be successfully combined with multigrid for the solution of nonlinear systems of algebraic equations obtained from the discretization of a particular set of PDEs arising in the modelling of phase change problems.

#### 3.1. The phase field method

One of the most popular techniques to have emerged in recent years for the modelling and simulation of phase change problems is the phase field method. Details of this technique can be found in numerous articles, e.g. [10, 11, 15], however the basic idea is quite simple. Instead of treating the evolving boundary between the solid and liquid regions as being discrete,

it is treated as being diffuse and an additional (smooth) dependent variable is introduced to represent the phase as a function of space and time. Evolution of this phase variable is prescribed by a PDE that is derived from a governing free energy functional, and there is typically a parameter present within the formulation that allows the width of the diffuse interface to be controlled. This type of model is able to benefit significantly from the use of mesh adaptivity around the interface in order to capture this aspect of the solution as accurately as possible.

In the following subsection a snapshot of some typical computational results is provided for one specific test problem. This is based upon the following model of the solidification of a dilute binary alloy, which comes from the work of [15]:

$$A(\nabla\phi)^2\frac{\partial\phi}{\partial t} = \nabla \cdot (A^2(\nabla\phi)\nabla\phi) + \frac{\partial}{\partial x} \left( |\nabla\phi|^2 A(\nabla\phi) \frac{\partial A(\nabla\phi)}{\partial(\phi_x)} \right) + \frac{\partial}{\partial y} \left( |\nabla\phi|^2 A(\nabla\phi) \frac{\partial A(\nabla\phi)}{\partial(\phi_y)} \right) + \phi - \phi^3 - \lambda(1 - \phi^2)^2 (T + Mc_\infty U), \quad (1)$$

$$\frac{1+k}{2} \frac{\partial U}{\partial t} = \nabla \cdot \left( D \frac{1-\phi}{2} \nabla U + \frac{1}{2\sqrt{2}} [1 + (1-k)U] \frac{\partial\phi}{\partial t} \frac{\nabla\phi}{|\nabla\phi|} \right) + \frac{1}{2} \frac{\partial}{\partial t} \{ \phi [1 + (1-k)U] \}, \quad (2)$$

$$\frac{\partial T}{\partial t} = \alpha \nabla^2 T + \frac{1}{2} \frac{\partial\phi}{\partial t}. \quad (3)$$

In these equations the dependent variables are  $\phi$ ,  $U$  and  $T$ , representing the phase, a non-dimensional concentration (of the minor component of the alloy) and a non-dimensional temperature respectively. Note that the problem is highly nonlinear and is made stiff by the fact that the diffusion coefficients  $D$  and  $\alpha$  typically differ by a number of orders of magnitude. The terms  $A(\nabla\phi)$  are nonlinear functions of the derivatives of  $\phi$  which are used to define the anisotropy within the solidification.

### 3.2. Solution method and sample computational results

In this example we apply a simple finite difference discretization to the above equations, although finite elements may equally well be used [9]. Mesh adaptivity is undertaken using a simple indicator which is based upon a linear combination of the gradients of the three dependent variables. Due to the stiffness of the resulting ordinary differential equations (ODEs) a fully implicit BDF2 integration scheme is used for the time-stepping. This results in a highly nonlinear system of coupled algebraic equations at each time step, which are solved using Brandt's adaptive (MLAT) version of the multigrid scheme [3]. Full details and analysis of this solution procedure may be found in [16] in the context of a slightly different phase field model.

Figure 1 shows some typical computational results for the non-dimensional concentration and temperature variables, at a particular instant in time, in the case where the ratio of  $D$  to  $\alpha$  is 500. Figure 2 illustrates a typical adaptive mesh (where the different colours represent different levels of refinement) and Figure 3 contains graphs that demonstrate the optimality of the multigrid implementation. All of these figures are courtesy of J. Rosam.




Figure 1. Typical results for non-dimensional concentration (left) and temperature fields (right), courtesy of J. Rosam.




Figure 2. A typical adapted mesh, courtesy of J. Rosam.

#### 4. PARALLEL MULTIGRID IMPLEMENTATIONS

Parallel computing is now becoming an increasingly cost-effective tool for the numerical solution of large CFD problems. It is important to observe however that this tool will only

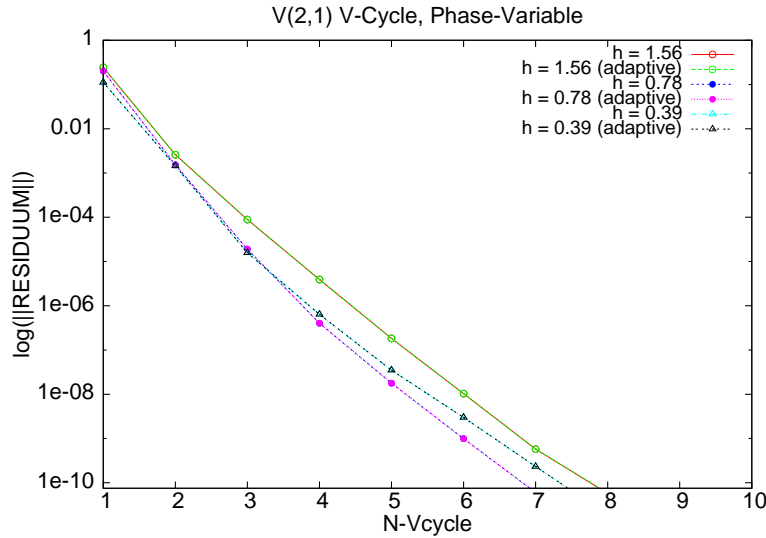


Figure 3. Multigrid convergence for different levels of mesh, courtesy of J. Rosam.

ever be of significant value when it is used in conjunction with the best available algorithms for the solution of the problem at hand. For this reason, the parallel implementation of multigrid, and other multilevel solvers, has received significant recent attention, e.g. [1, 6, 12]. A major challenge that is associated with such work is to ensure that the computations that take place on the coarsest computational grids do not lead to algorithmic bottlenecks in which the parallel performance is compromised by the very poor computation to communication ratios. In this section we demonstrate that respectable parallel multigrid performance can be achieved for the solution of a system of nonlinear parabolic PDEs arising in the thin film approximation of droplet spreading.

#### 4.1. Thin film flow of a spreading droplet

In this example we consider the flow of a viscous droplet of fluid spreading down an inclined plane subject to gravitational and surface tension forces. The model that we use is based upon a long wave approximation, in which it is assumed that the flow perpendicular to the plane is negligible compared to the flow parallel with the inclined plane. A small parameter,  $\epsilon$ , may be defined as the ratio of typical length scales,  $H_0$  and  $L_0$ , in the perpendicular and parallel directions respectively. It is then possible to non-dimensionalize the Stokes equations and collect together the leading order terms in this small parameter in order to yield a simplified system of equations for this thin film flow. Further details are provided in [17, 18] for example, but the resulting PDEs take the following form:

$$\frac{\partial h}{\partial t} = \frac{\partial}{\partial x} \left[ \frac{h^3}{3} \left( \frac{\partial p}{\partial x} - \frac{Bo}{\epsilon} \sin \alpha \right) \right] + \frac{\partial}{\partial y} \left[ \frac{h^3}{3} \left( \frac{\partial p}{\partial y} \right) \right], \quad (4)$$

$$p = -\nabla^2(h + s) - \Pi(h) + Bo \cos \alpha (h + s - z). \quad (5)$$

In the above,  $h$  and  $p$  are the dependent variables (film height and pressure respectively),  $Bo$  is the dimensionless Bond number (which represents the ratio of gravitational to surface tension forces) and  $\alpha$  is the angle made by the inclined plane, see Figure 4 for an illustration of this. The function  $s$  is a known function describing surface roughness on the plane whilst the term  $\Pi(h)$  is a so-called *disjoining pressure* term that is used to model slip of the moving contact line. This term is based upon the assumption of the presence a thin precursor film, of non-dimensional thickness  $h^*$ , and takes the general non-dimensional form:

$$\Pi(h) = \frac{(n-1)(m-1)}{h^*(n-m)} \sigma (1 - \cos \theta_c) \left[ \left( \frac{h^*}{h} \right)^n - \left( \frac{h^*}{h} \right)^m \right]. \quad (6)$$

Here,  $\theta_c$  represents the contact angle between the droplet and the surface (which may differ for an advancing or retreating fluid) and the pair  $(n, m)$  are the exponents of the chosen interaction potential ((3, 2) in this example). For further details see, for example, [17, 18, 20].




Figure 4. A schematic of the droplet problem (courtesy of Koh [12]).

#### 4.2. Solution method and sample computational results

The basic solution method that is employed here is to provide a parallel implementation of the nonlinear multigrid solver that is described in detail in [18, 19]. This is based upon the use of finite differences in space with implicit time-stepping and an FAS multigrid solver for the nonlinear algebraic equations that arise at each step. In order for the precursor film model to be physically realistic it is essential that the precursor film thickness,  $h^*$ , be sufficiently small (for example, [20] suggests that the dimensional size should be in the range 1-100nm). Furthermore, to capture correspondingly small values of  $h^*$  in the above model it is essential that the spatial resolution be on a similar length scale: this is achieved here through the use of very fine grids and parallel computation.

Full details of the parallel implementation may be found in [12], which follows [6] in a number of regards. In particular, the parallelism is achieved via a geometric decomposition of the domain, which is enforced across all levels of the multigrid hierarchy, and a stripwise partition is used to ensure that each subdomain/processor has at most two immediate neighbours (there is a one-to-one mapping between processors and subdomains). Each processor is responsible for implementing the FAS algorithm on its own subdomain, making use of additional columns of

“ghost nodes” immediately to either side of this subdomain. These ghost nodes are necessary for neighbour-to-neighbour communications at the following stages of the multigrid algorithm:

- after each red and each black sweep of the nonlinear Red-Black Gauss-Seidel smoother at each grid level,
- prior to restriction of the residual and solution to each coarser level,
- to obtain the exact solution at the coarsest level,
- prior to interpolation of the error to each finer level.

In addition, a short global communication is required after each multigrid V-cycle in order to decide if convergence has been achieved. Note also that the size of the coarsest possible grid is determined by the number of processors being used since, for the simplicity of this implementation, there should be at least one row of the coarsest grid on each processor.

Note that the optimal performance of the sequential multigrid algorithm for this problem, illustrated in Figure 5 (courtesy of Sellier, [18]), is maintained for the parallel implementation however there is some loss of efficiency if large numbers of processors are used, especially when the problem size is not sufficiently large. Tables I and II (courtesy of Koh, [12]) illustrate the parallel multigrid efficiency on problems with fine grids of size  $1025 \times 1025$  and  $8193 \times 8193$  respectively. It is clear that for the smaller of these problems there is no advantage to be gained by using 128 processors for example. An alternative indication of the scalability of this approach comes from inspecting the underlined figures in each table: these compare the solution times when the amount of work per processor is almost identical ( $1025 \times 1025$  on 2 processors versus  $8193 \times 8193$  on 128 processors) and thus give an indication of the loss of efficiency due to the additional interprocessor communication on 128 processors, which is especially significant for the coarsest grids.

Multigrid Levels	Coarse Grid	Number of Processors							
		1	2	4	8	16	32	64	128
3	257	–	119.54	62.21	30.24	13.84	8.10	6.23	5.31
4	129	–	<u>56.88</u>	28.52	14.20	7.96	3.94	2.91	3.14
5	65	–	57.10	28.89	13.89	7.21	3.37	2.61	–
6	33	–	65.32	28.91	15.66	7.68	4.65	–	–

Table I. Parallel multigrid timings on a  $1025 \times 1025$  fine grid (courtesy of Koh, [12]).

Finally, in this section, we present contour plots for the spreading of a droplet down an inclined plane in the presence of contact angle hysteresis. Figure 6 shows a snapshot of typical results computed in parallel on 15 different processors (courtesy of Koh, [12]).

## 5. P-VERSION OF MULTIGRID

The multigrid principle is based upon correcting the latest solution using a nested sequence of smaller and smaller approximation spaces. In the examples that have been described so far this is achieved through a hierarchy of mesh refinements/coarsenings. When seeking approximations of high polynomial degree on a relatively small number of elements it makes sense to consider




Figure 5. Demonstration of the optimal multigrid performance (courtesy of Sellier [18]).

Multigrid Levels	Coarse Grid	Number of Processors		
		32	64	128
5	513	250.72	154.85	<u>73.57</u>
6	257	247.54	151.04	74.17
7	129	255.96	156.39	77.94

Table II. Parallel multigrid timings on a  $8193 \times 8193$  fine grid (courtesy of Koh, [12]).

a different sequence of nested spaces: the p-version of multigrid considers spaces with lower and lower polynomial degree [5]. The following example illustrates this approach in the case of a discontinuous Galerkin (DG) finite element (FE) discretization of another nonlinear flow problem.

### 5.1. Elastohydrodynamic lubrication (EHL)

This problem involves another long wave approximation, based upon the same thin film expansion that was used for the model problem in the previous section. In the case of EHL however the liquid is assumed to lie between two solid elements and it is assumed that the pressures are sufficiently large to permit the elastic deformation of these elements. This means that, whilst the lubrication equation is quite similar to a steady-state version of (4), the

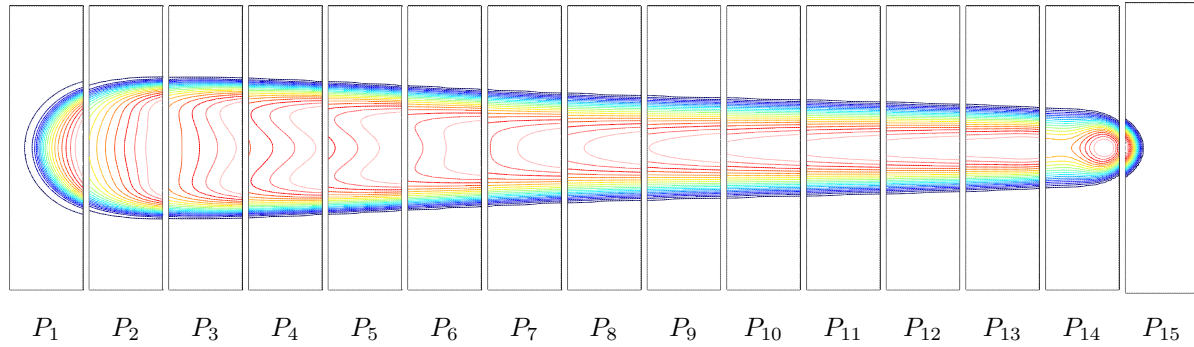


Figure 6. Contour plots on 15 processors for a droplet spreading problem (courtesy of Koh, [12]).

pressure equation is rather different to (5). Following the notation of [13] a non-dimensional form of these equations may be expressed as:

$$-\nabla \cdot (\epsilon \nabla P) + \nabla \cdot (\beta \bar{\rho} H) = 0, \quad (\beta = (1, 0)^T) \quad (7)$$

and

$$H(X, Y) = H_{00} + \frac{X^2}{2} + \frac{Y^2}{2} + \frac{2}{\pi} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{P(X', Y')}{\sqrt{(X - X')^2 + (Y - Y')^2}} dX' dY', \quad (8)$$

where

$$\epsilon = \frac{\bar{\rho} H^3}{\bar{\eta} \lambda}, \quad \bar{\eta} = \bar{\eta}(P), \quad \bar{\rho} = \bar{\rho}(P).$$

In addition to these equations there is also a free boundary where  $P$  becomes zero (the cavitation boundary) and a force balance constraint that must be maintained. Further details of the non-dimensionalized equations and their derivation may be found in [22] for example.

### 5.2. Solution method and sample computational results

For the purposes of this brief description we focus primarily on the solution of the lubrication equation (7) using a high order DG discretization. The film thickness  $H$  may be treated as a function of  $P$  via an evaluation of (8), using a suitable quadrature formula (see [13] for details of appropriate singular quadrature rules). The DG discretization itself is based upon the techniques of [2], which are designed for convection-diffusion problems, and also incorporates a penalty term, based upon [24], to capture the cavitation boundary automatically. Again, see [13] for full details.

Using a polynomial degree of 9 on each element ensures that relatively few elements are required overall, however, in order to obtain a very accurate solution these elements should not be of uniform size. Hence an adaptive strategy is implemented in which elements with a large estimated error are subdivided into four children, whilst neighbouring elements with a very small estimated error may be merged. The adaptivity takes place without having obtained

a fully-converged solution on each mesh however full convergence is required on the final mesh. An example of such a mesh is illustrated in Figure 7 below (courtesy of Lu, [13]). This figure also shows the corresponding converged pressure profile which, although smooth, features a typical “horseshoe” shaped ridge at the outflow side of the contact region. Note that the solutions below were obtained using p-multigrid to accelerate the convergence of the basic iterative solver (which is based upon a pointwise quasi-Newton update scheme, [13]). This scheme features four different levels of “grid”: based upon polynomial spaces of degree 9, 7, 5 and 3 respectively (each on the same mesh, such as that shown in Figure 7). The underlying features of the FAS multigrid algorithm are all present in this scheme however one fundamental difference is that the optimal convergence properties of the traditional multigrid schemes (as shown in Figures 3 or 5 for example) are no longer obtained. Hence, although p-multigrid accelerates convergence very significantly, it is no longer the case that the overall solution time is proportional to the total number of degrees of freedom.




Figure 7. Adaptive mesh and solution of a point contact EHL problem using DG with p-multigrid (courtesy of Lu, [13]).

## 6. DISCUSSION

In this paper we have attempted to demonstrate that highly efficient multigrid techniques may be successfully combined with other state-of-the-art computational methods such as adaptivity, parallel computing and high order discretizations. This combination of multigrid, or related multilevel solvers, with other modern scientific computing techniques offers a lot of potential in the solution of CFD problems, however there are still a significant number of obstacles to be overcome. Two areas where further research is certainly required are: in the combination of multigrid, adaptivity and parallelism for time-dependent problems (such as those considered in Sections 3 and 4 above), and; in the application of multigrid in combination with both  $h$  and  $p$  refinement together.

## ACKNOWLEDGEMENTS

The work described in this paper has been undertaken jointly with numerous collaborators: Martin Berzins, Phil Gaskell, Yit-Yan Koh, Hongqiang Lu, Andy Mullis, Jan Rosam, Mathieu Sellier and Harvey Thompson. I am very grateful to each of them for helping to provide a stimulating and rewarding environment in which to undertake research. I am particularly indebted to Yit-Yan Koh, Hongqiang Lu, Jan Rosam and Mathieu Sellier for allowing me to use their figures in this paper.

## REFERENCES

1. Bank RE, Jimack PK. A new parallel domain decomposition method for the adaptive finite element solution of elliptic partial differential equations. *Concurrency and Comp.: Practice and Exp.* 2001; **13**:327–350.
2. Baumann CE, Oden JT. A discontinuous hp finite element method for convection-diffusion problems. *Comput. Meth. Appl. Mech. Engrg.* 1999; **175**:311–341.
3. Brandt A. Multi-level adaptive solutions to boundary-value problems. *Math. Comp.* 1977; **31**:333–390.
4. Brandt, A. Guide to multigrid development. In *Multigrid methods: Lecture notes in Mathematics* 1982, eds. Hackbusch, W. Trottenberg, U. (Springer-Verlag, Berlin); **960**:220-312.
5. Fidkowski KJ, Darmofal DL. Development of a higher-order solver for aerodynamic applications. In *42nd AIAA Aerospace Sciences Meeting* 2004; AIAA Paper 2004-0436.
6. Goodyer CE, Berzins M. Parallelisation and scalability issues of a multilevel elasto-hydrodynamic lubrication solver. *Concurrency and Comp.: Practice and Exp.* 2007; **19**:369–396.
7. Hart L, McCormick S, O’Gallagher A. The fast adaptive composite-grid method (FAC): algorithms for advanced computers. *Appl. Math. Comput.* 1986; **19**:103–125.
8. Jeong, J-H, Goldenfeld, N & Dantzig, JA. Phase field model for three-dimensional dendritic growth with fluid flow. *Phys. Rev. E* 2001; **64**:041602.
9. Jones AC. *A Projected Multigrid Method for the Solution of Nonlinear Finite Element Problems on Adaptively Refined Grids* (PhD Thesis, University of Leeds, UK) 2005.
10. Karma A. Phase-field formulation for quantitative modelling of alloy solidification. *Phys. Rev. Lett.* 2001; **87**:115701.
11. Karma A, Rappel E-J. Quantitative phase-field modeling of dendritic growth in two and three dimensions. *Phys. Rev. E* 1998; **57**:4323–4349.
12. Koh Y-Y. *Efficient Numerical Solutions of Droplet Spreading Flows* (PhD Thesis, University of Leeds, UK) 2007.
13. Lu H. *High Order Finite Element Solution of Elasto-hydrodynamic Lubrication Problems* (PhD Thesis, University of Leeds, UK) 2006.
14. Provatas N, Goldenfeld N, Dantzig J. Adaptive mesh refinement computation of solidification microstructures using dynamic data structures. *Journal of Computational Physics* 1999; **148**:265–290.
15. Ramirez JC, Beckermann C, Karma A, Diepers, HJ. Phase-field modelling of binary alloy solidification with coupled heat and solute diffusion. *Phys. Rev. E* 2004; **69**:051607.
16. Rosam J, Jimack PK, Mullis AM. A fully implicit, fully adaptive time and space discretisation method for phase-field simulation of binary alloy solidification. *J. Comp. Phys.* 2007; to appear.
17. Schwartz LW. Hysteretic effects in droplet motion on heterogeneous substrates: direct numerical simulation. *Langmuir* 1998; **14**:3440–3453.
18. Sellier M. *The Numerical Simulation of Thin Film Flow Over Heterogeneous Substrates* (PhD Thesis, University of Leeds, UK) 2006.
19. Sellier M, Gaskell PH, Jimack PK, Thompson HM. Efficient and accurate time-adaptive multigrid simulations of droplet spreading. *Int. J. Numer. Meth. Fluids* 2004; **45**:1161–1186.
20. Starov VM, Kalinin VV, Chen JD. Spreading of liquid drops over dry surfaces. *Adv. in Coll. and Interface Sci.* 1994; **50**:187–221.
21. Trottenberg U, Oosterlee CW, Schüller A. *Multigrid* (Academic Press) 2000.
22. Venner CH, Lubrecht AA. *Multilevel Methods in Lubrication* (Elsevier) 2000.
23. Wesseling, P. *Introduction to Multigrid Methods* (Wiley, New York) 1992.
24. Wu SR. A penalty formulation and numerical approximation of the Reynolds-Herz problem of elasto-hydrodynamic lubrication. *Int. J. Engrg. Sci.* 1986; **24**:1001–1013.