

NASA/CR-2001-210660  
ICASE Report No. 2001-3



## **Multigrid Approaches to Non-linear Diffusion Problems on Unstructured Meshes**

*Dimitri J. Mavriplis*  
*ICASE, Hampton, Virginia*

*ICASE*  
*NASA Langley Research Center*  
*Hampton, Virginia*

*Operated by Universities Space Research Association*



National Aeronautics and  
Space Administration

Langley Research Center  
Hampton, Virginia 23681-2199

Prepared for Langley Research Center  
under Contract NAS1-97046

February 2001

## Form SF298 Citation Data

<b>Report Date</b> <i>("DD MON YYYY")</i> 00FEB2001	<b>Report Type</b> N/A	<b>Dates Covered (from... to)</b> <i>("DD MON YYYY")</i>
<b>Title and Subtitle</b> Multigrid Approaches to Non-linear Diffusion Problems on Unstructured Meshes		<b>Contract or Grant Number</b>
<b>Authors</b> Dimitri J. Mavriplis		<b>Program Element Number</b>
<b>Performing Organization Name(s) and Address(es)</b> ICASE NASA Langley Research Center Hampton, Virginia 23681-2199		<b>Project Number</b>
<b>Sponsoring/Monitoring Agency Name(s) and Address(es)</b>		<b>Task Number</b>
<b>Distribution/Availability Statement</b> Approved for public release, distribution unlimited		<b>Work Unit Number</b>
<b>Supplementary Notes</b>		<b>Performing Organization Number(s)</b> ICASE Report No. 2001-3
<b>Abstract</b> Abstract. The efficiency of three multigrid methods for solving highly non-linear diffusion problems on two-dimensional unstructured meshes is examined. The three multigrid methods differ mainly in the manner in which the non-linearities of the governing equations are handled. These comprise a non-linear full approximation storage (FAS) multigrid method which is used to solve the non-linear equations directly, a linear multigrid method which is used to solve the linear system arising from a Newton linearization of the non-linear system, and a hybrid scheme which is based on a non-linear FAS multigrid scheme, but employs a linear solver on each level as a smoother. Results indicate that all methods are equally effective at converging the non-linear residual in a given number of grid sweeps, but that the linear solver is more efficient in cpu time due to the lower cost of linear versus non-linear grid sweeps.		<b>Monitoring Agency Acronym</b>
<b>Subject Terms</b>		<b>Monitoring Agency Report Number(s)</b>
<b>Document Classification</b> unclassified	<b>Classification of SF298</b> unclassified	
<b>Classification of Abstract</b> unclassified	<b>Limitation of Abstract</b> unlimited	

<b>Number of Pages</b>	
------------------------	--

12	
----	--

# MULTIGRID APPROACHES TO NON-LINEAR DIFFUSION PROBLEMS ON UNSTRUCTURED MESHES

DIMITRI J. MAVRIPLIS\*

**Abstract.** The efficiency of three multigrid methods for solving highly non-linear diffusion problems on two-dimensional unstructured meshes is examined. The three multigrid methods differ mainly in the manner in which the non-linearities of the governing equations are handled. These comprise a non-linear full approximation storage (FAS) multigrid method which is used to solve the non-linear equations directly, a linear multigrid method which is used to solve the linear system arising from a Newton linearization of the non-linear system, and a hybrid scheme which is based on a non-linear FAS multigrid scheme, but employs a linear solver on each level as a smoother. Results indicate that all methods are equally effective at converging the non-linear residual in a given number of grid sweeps, but that the linear solver is more efficient in cpu time due to the lower cost of linear versus non-linear grid sweeps.

**Key words.** non-linear, unstructured, multigrid

**Subject classification.** Applied and Numerical Mathematics

**1. Introduction.** The iterative solution of non-linear problems can be tackled in various different manners. The use of a Newton iteration approach involves linearizing the problem about the current state and inverting the global Jacobian using an appropriate linear solver. Although Newton iteration strategies are widespread, and are amenable to modular software development (for example in the selection of existing linear solvers), for very large systems of equations, the formation of the Jacobian matrix can be tedious and memory intensive. An alternate approach which avoids the construction of the Jacobian matrix consists of using a non-linear iterative solver applied directly to the non-linear equation set. This approach has often been described as “placing the non-linearity on the inside of the iterative procedure” as opposed to “placing the non-linearity on the outside of the iterative scheme”, for the former linearized approach. Intuitively, the non-linear solver approach results in the non-linear residuals being updated frequently (at every iteration), whereas the linear solver approach results in less frequent non-linear residual evaluations.

Multigrid methods can be applied to non-linear problems either as a linear solver, operating on a linearization of the governing equations, or as a non-linear multigrid formulation known as the full approximation storage (FAS) scheme [1]. The fact that the FAS multigrid scheme can be applied directly to non-linear problems remains a great advantage of multigrid methods, which can be used to solve such problems in a “matrix free” manner. Within the context of a multigrid method, a suitable error smoother must be chosen to reduce high-frequency errors on individual grid levels of the multigrid sequence. In the linear multigrid case, each individual grid problem will be linear, while in the FAS case, each grid problem will itself consist of a non-linear problem. In this latter case, the non-linear smoother can either be constructed as a non-linear iteration procedure, or as a linear solver operating on the linearization of the local non-linear problem.

The objective of this paper is to investigate the effectiveness of these various multigrid approaches to solving non-linear problems. In order to provide a meaningful comparison, similar discretization and solution

---

\*ICASE, Mail Stop 132C, NASA Langley Research Center, Hampton, VA 23681-2199, U.S.A., [dimitri@icase.edu](mailto:dimitri@icase.edu). This research was partially supported by the National Aeronautics and Space Administration under NASA Contract No. NAS1-97046 while the author was in residence at ICASE, NASA Langley Research Center, Hampton, VA 23681-2199. Partial support was also provided under U.S. Department of Energy subcontract B347882 from Lawrence Livermore National Laboratory.

strategies must be used in the different variations of the multigrid algorithms.

Our test problem consists of a set of non-linear diffusion equations, known as the radiation-diffusion equations. These equations govern the evolution of photon radiation in an optically thick medium, and can be derived from first principles by integrating over all energy frequencies, under the assumptions of isotropy and small mean-free photon paths [9]. These equations are important in the simulations of inertial confinement fusion and astrophysical phenomena.

From the numerical standpoint, the radiation-diffusion equations provide a suitable test-bed for examining the effectiveness of numerical algorithms at handling non-linearities, since these equations exhibit strong non-linear behavior, while at the same time they are of modest complexity, enabling a relatively straightforward Jacobian construction.

Various radiation-diffusion solution algorithms incorporating multigrid techniques have been demonstrated recently, but these generally have employed linear multigrid methods as preconditioners within a Newton-Krylov method [2, 4, 8]. Furthermore, these multigrid preconditioners have generally been applied in the context of an operator split algorithm.

In this work, we employ multigrid to the fully coupled system of equations. The developed algorithms employ multigrid directly as a solver for the radiation-diffusion equations rather than as a preconditioner. While the development of preconditioned Newton-Krylov solution techniques for these types of problems remains of interest, a more straight-forward comparison between linear and non-linear multigrid methods is afforded when both of these techniques are formulated as solvers.

There are various approaches to implementing multigrid methods on unstructured meshes. For the purposes of this comparison, the logistics of the multigrid implementation are not necessarily the dominant concern, provided both linear and non-linear approaches are formulated in a similar manner. At the outset, algebraic multigrid methods cannot be considered herein, since these are exclusively formulated as linear solvers. Because our long term objective is the development of a three-dimensional solver for complex geometries, we avoid nested geometric multigrid methods, which assume the existence of a complete sequence of coarse and fine unstructured meshes with nested cell structures, the construction of which may not be feasible in the general three-dimensional case. We concentrate instead on the agglomeration multigrid approach, which generates coarse (non-physical) levels automatically through a graph algorithm and makes use of the Galerkin projection for constructing the coarse level equations. This approach can be formulated in a linear or non-linear manner and has been demonstrated for large three-dimensional fluid-dynamics problems in complicated geometries [7].

The remainder of the paper is organized as follows. In section 2 we illustrate the correspondence between linear and non-linear multigrid methods. The general agglomeration multigrid approach as well as the three specific multigrid schemes implemented for comparison are described in section 3. In section 4 the discretization of the governing equations and sample test problem are described, while section 5 discusses our results, which are also summarized in the conclusion in section 6.

**2. Linear and Non-Linear MG Formulations.** The goal of any multigrid method is to accelerate the solution of a fine grid problem by computing corrections on a coarser grid and then interpolating them back to the fine grid problem. Although this procedure is described in a two grid context, it is applied recursively on a complete sequence of fine and coarser grid levels. To apply a linear multigrid method to a non-linear problem, a linearization must first be performed. Thus, if the equations to be solved are written as

$$(2.1) \quad \mathbf{R}(\mathbf{w}_{\text{exact}}) = \mathbf{0}$$

with the current estimate  $\mathbf{w}$  yielding the non-linear residual  $\mathbf{r}$ :

$$(2.2) \quad \mathbf{R}(\mathbf{w}) = \mathbf{r}$$

the Newton linearization of this system is taken as

$$(2.3) \quad \frac{\partial \mathbf{R}_h}{\partial \mathbf{w}_h} \Delta \mathbf{w}_h = -\mathbf{r}$$

This represents a linear set of equations in the solution variable  $\Delta \mathbf{w}_h$  (the correction), to which a linear multigrid (i.e. MG correction scheme) can be applied. In this case, the coarse grid equation reads:

$$(2.4) \quad \frac{\partial \mathbf{R}_H}{\partial \mathbf{w}_H} \Delta \mathbf{w}_H = -I_h^H \mathbf{r}_{\text{linear}}$$

where  $H$  and  $h$  represent coarse grid and fine grid values, respectively, and  $I_h^H$  represents the restriction operator which interpolates the fine grid residuals to the coarse grid. The residual of the linear system on the fine grid on the fine grid, which is given by

$$(2.5) \quad \mathbf{r}_{\text{linear}} = \frac{\partial \mathbf{R}_h}{\partial \mathbf{w}_h} \Delta \mathbf{w}_h + \mathbf{r}$$

and may be approximated as

$$(2.6) \quad \mathbf{r}_{\text{linear}} \approx \mathbf{R}(\mathbf{w} + \Delta \mathbf{w})$$

where  $\mathbf{R}$  refers to the non-linear residual, as previously. The coarse grid corrections  $\Delta \mathbf{w}_H$  which are obtained by solving equation (2.4) are initialized on the coarse grid as zero. After the solution of equation (2.4), these corrections are prolonged or interpolated back to the fine grid.

Alternatively, a non-linear FAS multigrid scheme can be used to solve equation (2.1) directly without resorting to a linearization. In this case, the FAS coarse grid equation reads:

$$(2.7) \quad \mathbf{R}_H(w_H) = \mathbf{R}_H(\tilde{I}_h^H w_h) - I_h^H \mathbf{r}$$

where the term on the right-hand side is often referred to as the defect-correction [1, 6].  $\mathbf{R}_H$  represents the coarse grid discretization and  $I_h^H$  and  $\tilde{I}_h^H$  denote the restriction operators which are now used to interpolate residuals as well as flow variables from the fine grid to the coarse grids. In principal, different restriction operators for residuals and variables may be employed. If equation (2.7) is re-written as:

$$(2.8) \quad \mathbf{R}_H(w_H) - \mathbf{R}_H(\tilde{I}_h^H w_h) = -I_h^H \mathbf{r}$$

the right hand sides of equations (2.4) and (2.8) represent similar approximations of the restricted non-linear residual, in view of equation (2.6) and the fact that these restricted residuals in the FAS scheme are always evaluated at the most recently available fine grid updates. Therefore, by equating the left hand sides of equations (2.4) and (2.8), the equivalence between the linear multigrid scheme and the non-linear FAS scheme is seen to be given by:

$$(2.9) \quad \mathbf{R}_H(w_H) - \mathbf{R}_H(I_h^H w_h) \approx \frac{\partial \mathbf{R}_H}{\partial \mathbf{w}_H} \Delta \mathbf{w}_H$$

which means that the FAS multigrid scheme corresponds to an approximation to a linear multigrid scheme, where the coarse grid Jacobians are approximated by finite differencing the operator. Therefore, in the limit of asymptotic convergence, i.e. when  $\Delta \mathbf{w}_H \ll 1$ , the two methods should yield similar convergence rates.

Note that the above discussion involves no specification of the coarse grid operator and Jacobian construction. Therefore, a fair comparison of linear versus non-linear multigrid methods should utilize a similar construction for both of these quantities in the respective algorithms.

**3. Multigrid Algorithms.** The three multigrid variants implemented in this work are based on the agglomeration multigrid strategy. Agglomeration multigrid was originally developed for finite-volume schemes, and is based on agglomerating or fusing together neighboring fine grid control-volumes to form larger coarse grid control volumes as depicted in Figure 3.1. This approach has since been generalized for arbitrary discretizations following algebraic multigrid principles. In fact, agglomeration multigrid can be viewed as a simplification and extension of algebraic multigrid to non-linear systems of equations. The control-volume agglomeration algorithm can be recast as a graph algorithm, similar to algebraic multigrid methods, where the “seed” vertex initiating an agglomerated cell corresponds to a coarse grid point, and the neighboring agglomerated points correspond to fine grid points, in the algebraic multigrid terminology [10]. While weighted graph algorithms can be employed for agglomeration, these weights cannot depend on solution values, as in the algebraic multigrid case, but only on grid metrics. In this manner, the coarse grid levels are static and need only be constructed at the beginning of the simulation. This avoids one of the problems associated with algebraic multigrid applied to linearizations arising from non-linear problems, where newly constructed coarse levels may be required at each non-linear update, a task which can be complicated in parallel computing environments [3]. In the present work, for simplicity we limit ourselves to isotropic coarsening using an unweighted graph, which produces coarse level graphs which are maximal independent sets of the fine grid graph.

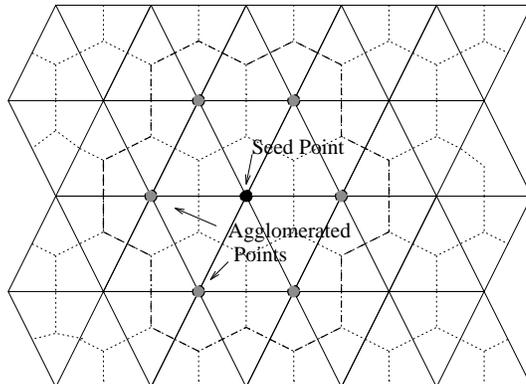


FIG. 3.1. *Illustration of Agglomeration Multigrid Coarse Level Construction*

As in the algebraic multigrid case, agglomeration multigrid employs a Galerkin projection for the construction of the coarse grid equations. Thus, the coarse grid operator is given by:

$$(3.1) \quad \mathbf{R}_H = I_h^H \mathbf{R}_h I_H^h$$

where  $I_h^H$  is the restriction operator, and  $I_H^h$  is the prolongation operator, and both operators are taken as piecewise constants. This simple construction applies equally to linear and non-linear operators, and reduces to forming the coarse grid equation at an agglomerated cell as the sum of the fine grid equations at each fine grid cell contained in the coarse grid cell. The non-linearities in the operator are evaluated using solution variables on the coarse grid interpolated up from the fine grid.

Given this multigrid infrastructure, three particular algorithms which differ mainly in the manner in which non-linearities are handled are developed for comparison. The first involves a standard non-linear FAS multigrid algorithm, and the second involves a linear multigrid algorithm applied to the linearization of the governing equations. Finally, a hybrid algorithm is also devised which uses a non-linear FAS multigrid outer cycle, but a linear solver on each grid level as a smoother.

**3.1. FAS Scheme.** In the non-linear FAS multigrid algorithm, equation (2.1) is solved directly. The coarse grid equations are formed by Galerkin projection (c.f. equation (3.1)) and the non-linearities in the coarse grid operator are evaluated using coarse level solution variables interpolated up from the fine grid using the  $\tilde{I}_h^H$  restriction operator (as per equation (2.8)). On each grid level, the discrete equations are solved using a non-linear block Jacobi iteration given as:

$$(3.2) \quad \mathbf{w}^{\text{new}} = \mathbf{w}^{\text{old}} + [D]^{-1} \mathbf{R}(\mathbf{w}^{\text{old}})$$

where  $[D]$  represents the block diagonal of the Jacobian matrix. This smoother constitutes a non-linear solver, since the non-linear residual is updated at each stage, and incurs minimum memory overheads since only the storage of the block 2X2 matrix  $[D]$  representing the coupling between the two solution variables at each grid point is required.

**3.2. Linear Multigrid Scheme.** The linear multigrid scheme solves equation (2.3) on the fine grid, and equation (2.4) on the coarse levels. On the fine grid, the Jacobian  $\frac{\partial \mathbf{R}_h}{\partial \mathbf{w}_h}$  is formed by explicitly differentiating (hand coding) the discrete operator  $\mathbf{R}_h$ . On the coarse levels, for consistency with the FAS multigrid algorithm, the Jacobian is taken as the explicit differentiation of the coarse non-linear operator obtained by Galerkin approximation (c.f. equation (3.1)). Thus flow variables as well as residuals are restricted to the coarser grids, but the non-linear residuals on these coarser levels are not evaluated, only the Jacobians corresponding to the linearization of the non-linear coarse level residuals. These coarse level Jacobians are evaluated at the beginning of the solution phase for the non-linear time-step problem, and are then held fixed throughout the linear multigrid iterations. Memory requirements for the linear multigrid scheme are increased over those of the FAS scheme due to the required storage of the fine and coarse level Jacobians.

An outer Newton iteration is employed to solve the complete non-linear problem  $\mathbf{R}(w) = 0$ . Within each Newton iteration, the linear system defined by equation (2.3) is solved by the linear multigrid algorithm. This provides a new fine grid non-linear correction  $\Delta w$  which is then used to update the non-linear residual. Given an accurate solution of the linear problem for each outer iteration, we can expect the non-linear Newton scheme to converge very rapidly (quadratically), thus minimizing the number of non-linear updates required.

On each grid level, the linear multigrid scheme solves the linear system using a block-Jacobi smoother. If the Jacobian is divided up into diagonal and off-diagonal block components, labeled as  $[D]$  and  $[O]$ , respectively, the Jacobi iteration can be written as:

$$(3.3) \quad [D]\Delta\mathbf{w}_h^{n+1} = -\mathbf{r} - [O]\Delta\mathbf{w}_h^n$$

where  $\Delta\mathbf{w}_h^n$  represents corrections from the previous linear iteration, and  $\Delta\mathbf{w}_h^{n+1}$  represents the new linear corrections produced by the current linear iteration. At each linear iteration, the solution of equation (3.3) requires the inversion of the block  $2X2$  matrix  $[D]$  at each grid point. The linear corrections  $\Delta\mathbf{w}_h$  are initialized to zero at the first iteration on each grid level. Therefore, this linear iteration strategy reduces to the non-linear Jacobi scheme described above in the event only a single linear iteration is employed.

In contrast to the non-linear FAS multigrid algorithm, the residuals, jacobians (i.e.  $[D]$  and  $[O]$  terms), and the variables interpolated up to the coarse grids are only evaluated at the start of the non-linear iteration, and are held fixed during all inner linear multigrid cycles within a non-linear iteration.

**3.3. Hybrid Scheme.** For comparison purposes, a hybrid scheme has also been developed which is based on a non-linear FAS multigrid scheme, but solves the linearized equations on each grid level using the block Jacobi smoother described by equation (3.3) instead of the non-linear iteration scheme of equation (3.2).

While this scheme solves the same form of the equations on all grid levels as the linear multigrid scheme, the non-linear residuals, physical solution variables, and discrete Jacobians are continuously updated with each visit to a new grid level in the multigrid scheme, and no outer Newton iteration is required. However, on each grid level, multiple linear iteration are performed to solve the non-linear problem on that grid level, and the non-linear residuals are thus only evaluated once on each grid level for each multigrid cycle, regardless of the number of linear iterations employed. On the other hand, this algorithm incurs the same memory overheads as the linear multigrid scheme due to the required storage of the Jacobians.

**4. Problem Formulation.** The non-equilibrium radiation diffusion equations can be written as

$$(4.1) \quad \frac{\partial E}{\partial t} - \nabla \cdot (D_r \nabla E) = \sigma_a (T^4 - E)$$

$$\frac{\partial T}{\partial t} - \nabla \cdot (D_t \nabla T) = -\sigma_a (T^4 - E)$$

with

$$\sigma_a = \frac{z^3}{T^3}, \quad D_r(T, E) = \frac{1}{3\sigma_a + \frac{1}{E} \left| \frac{\partial E}{\partial n} \right|}, \quad D_t(T) = \kappa T^{\frac{5}{2}}$$

Here,  $E$  represents the photon energy,  $T$  is the material temperature, and  $\kappa$  is the material conductivity. In the non-equilibrium case, the non-linear source terms on the right-hand-side are non-zero and govern the transfer of energy between the radiation field and material temperature. Additional non-linearities are generated by the particular form of the diffusion coefficients, which are functions of the  $E$  and  $T$  variables.

In particular, the energy diffusion coefficient,  $D_r(T, E)$  contains the term  $|\frac{\partial E}{\partial n}|$  which refers to the gradient of  $E$  in the direction normal to the cell interface (in the direction of the flux). This term constitutes a flux limiter, which is an artificial means of ensuring physically meaningful energy propagation speeds (i.e. no larger than the speed of light) [2, 4, 8]. The atomic number  $z$  is a material coefficient, and while it may be highly variable, it is only a function of position (i.e.  $z = f(x,y)$  in two dimensions).

Equations (4.1) represent a system of coupled non-linear partial-differential equations which must be discretized in space and time. Spatial discretization on two-dimensional triangular meshes is achieved by a Galerkin finite-element procedure, assuming linear variations of  $E$  and  $T$  over a triangular element. The non-linear diffusion coefficients are evaluated by first computing an average  $T$  and  $E$  value along a triangle edge, and then computing the non-linear diffusion coefficient at the edge midpoint using these averaged values. While a standard finite-element discretization requires computing triangle averaged diffusion coefficients, the edge-based approach can more easily be reproduced on coarse agglomerated levels which do not contain triangular structures, and simplifies the construction of exact discrete Jacobians. Both edge and triangle based discretizations have been implemented and tested with little observable difference in accuracy. The gradient of  $E$  in the  $D_r$  diffusion coefficient is also taken as a one dimensional gradient along the direction of the stencil edge. The source terms are evaluated using the local vertex values of  $E$  and  $T$  exclusively, rather than considering linear variations of these variables.

The time derivatives are discretized as first-order backwards differences, with lumping of the mass matrix, leading to an implicit scheme which requires the solution of a non-linear problem at each time step. This approach is first-order accurate in time, and is chosen merely for convenience, since the principal objective is the study of the solution of the non-linear system. More sophisticated time integration strategies involving higher-order time discretizations and variable time-step sizes, which have been implemented by other researches in this field [2, 5] are planned for future work.

The Jacobian of the required linearizations is obtained by differentiation (hand coding) of the discrete non-linear residual. Because the spatial discretization involves a nearest neighbor stencil, the Jacobian can be expressed on the same graph as the residual discretization, which corresponds to the edges of the triangular grid. The initial guess for the solution of the non-linear problem at each time-step is taken as the solution obtained at the previous time-step.

The test case chosen for this work is taken from [8] and depicted in Figure 4.1. We consider a unit square domain of two dissimilar materials, where the outer region contains an atomic number of  $z = 1$  and the inner regions  $(1/3 < x < 2/3)$ ,  $(1/3 < y < 2/3)$  contains an atomic number of  $z = 10$ . The top and bottom walls are insulated, and the inlet and outlet boundaries are specified using mixed (Robin) boundary conditions, as shown in the figure. This domain is discretized using a triangular grid containing 7,502 vertices, shown in Figure 4.2. This grid conforms to the material interface boundaries in such a way that no triangle edges cross this boundary.

Figure 4.2 illustrates a typical simulation for this case. Incoming radiation sets up a traveling thermal front in the material, the progress of which is impeded by the region higher atomic number  $z$ . At critical times in the simulation, the diffusion coefficients can vary by up to six orders of magnitude near the material interfaces, thus providing a challenging non-linear behavior for the multigrid algorithms. At each physical time step, a non-linear problem must be solved. It is the solution of this transient non-linear problem at a given time step which forms the test problem for the three agglomeration multigrid algorithms. Clearly, the size of the physical time step affects the stiffness of the non-linear problem to be solved, with smaller physical time-steps leading to more rapidly converging systems. The non-dimensional time-step chosen in

this simulation was taken as 0.01. This constitutes a rather large value compared to those employed in reference [8] (usually of the order of  $10^{-3}$ ) and may have an adverse effect on overall temporal accuracy, but provides a more stringent test case for the multigrid solvers. Of the order of 1000 time steps are required to propagate the thermal front from the inlet to outlet boundary in the current simulation.

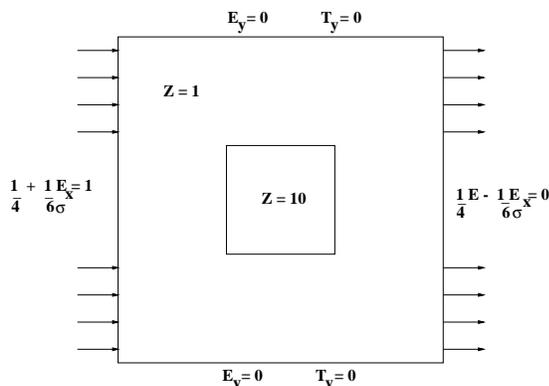


FIG. 4.1. *Sample test problem for non-linear diffusion equations*

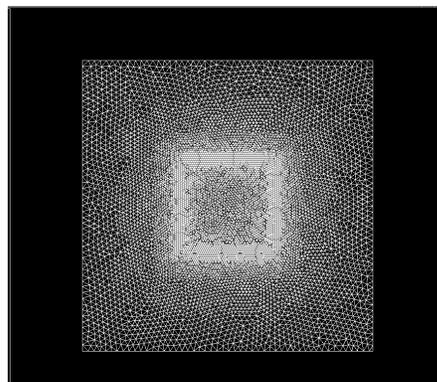


FIG. 4.2. *Illustration of unstructured grid for non-linear diffusion problem: 7,502 vertices*

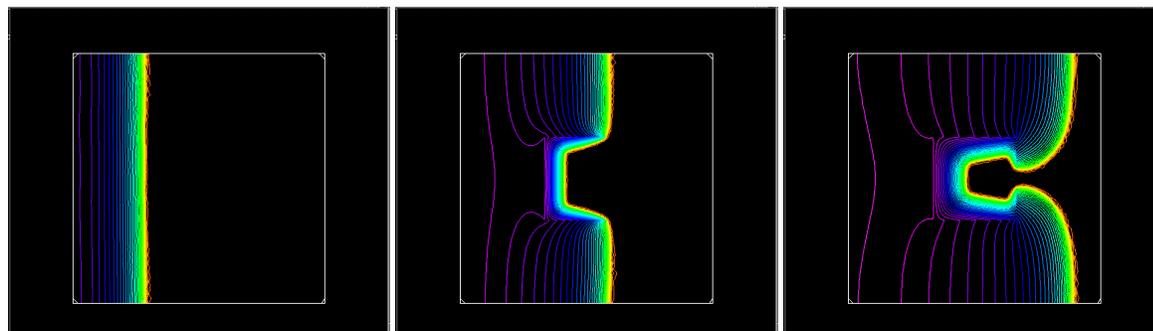


FIG. 4.3. *Illustration of solution for non-linear diffusion problem: Contours of  $T$*

**5. Results.** In the case of the non-linear FAS multigrid scheme, the multigrid algorithm is used directly to solve the non-linear equations at a given time step. Thus a dual-loop structure is used, containing an outer loop over the physical time steps, and an inner loop over the number of non-linear multigrid cycles. A four level multigrid W-cycle is employed by the FAS scheme, using 6 non-linear Jacobi iterations on each grid level. Figure 5.1 shows the convergence achieved by the non-linear FAS multigrid algorithm at the time step corresponding to the center frame in Figure 5.1, when the thermal front has begun to encounter the material interface, in the presence of strong non-linearities. The non-linear residual is reduced by 8 orders of magnitude over 9 multigrid cycles in this case, for an average reduction rate of 0.15.

In Figure 5.2, the convergence history for the same case is shown for the hybrid FAS multigrid scheme, using 6 linear iterations per grid level with a 4 level W-cycle. The convergence of this case is almost identical to that of the previous case.

Figures 5.3 and 5.4 depict the convergence of the outer non-linear Newton iteration and the inner linear multigrid iteration. The Newton iteration is seen to converge quadratically for this case, as expected, since

the Jacobian is a consistent differentiation of the discrete operator. The linear multigrid iterations employ a 4 level W-cycle with 3 linear Jacobi iterations per grid sweep and 3 multigrid W-cycles per non-linear iteration. The figure illustrates the slight increase in the linear residual at the beginning of each non-linear iteration, each time the linear system is modified by the evolution of the non-linearities. The large spike at the 13th linear iteration corresponds to the beginning of a new physical time-step and thus new non-linear problem.

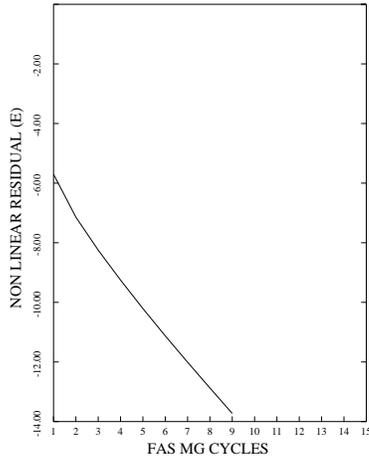


FIG. 5.1. *Convergence Rate for FAS Multigrid Scheme*

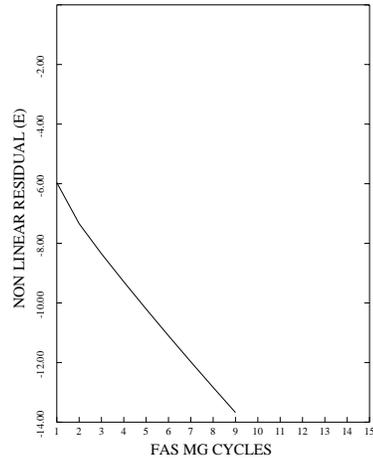


FIG. 5.2. *Convergence Rate for Hybrid Scheme*

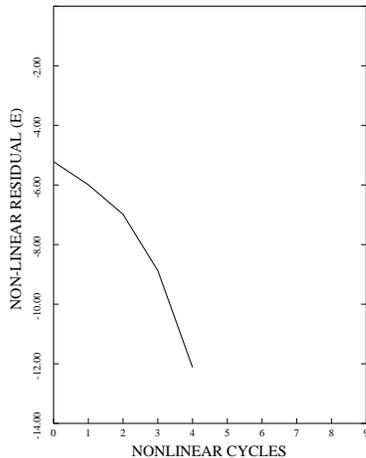


FIG. 5.3. *Convergence Rate of Outer Newton Iteration for Linear Multigrid Scheme*

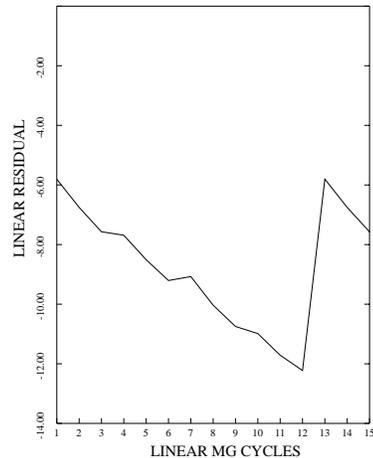


FIG. 5.4. *Convergence Rate of Inner Linear Iteration for Linear Multigrid Scheme*

In Figure 5.5 the convergence of the non-linear residual for all three methods is plotted in terms of grid sweeps, where a grid sweep is defined as a linear or non-linear Jacobi iteration on the fine grid. This plot indicates that all methods converge to approximately the same level in 40 to 50 grid sweeps. However, when these results are plotted in terms of cpu time, Figure 5.6 reveals the higher efficiency of the linear multigrid solver which is almost an order of magnitude faster than the non-linear FAS multigrid solver, while the

hybrid scheme lies in between these two extremes. This is principally due to the lower number of non-linear residual evaluations required by the linear and hybrid schemes.

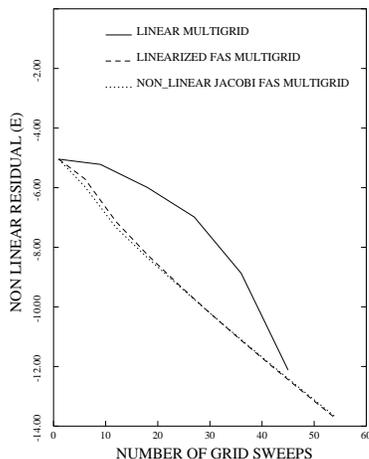


FIG. 5.5. Comparison of Convergence of Non-Linear Residual in Terms of Number of Grid Sweeps for Three Multigrid Schemes

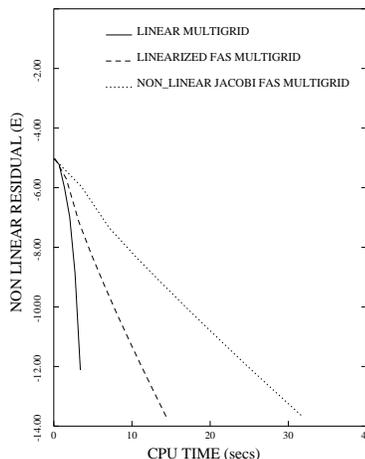


FIG. 5.6. Comparison of Convergence of Non-Linear Residual in Terms of CPU Time for Three Multigrid Schemes

**6. Conclusions.** The relative performance of the three multigrid algorithms described above for a particular time step is indicative of the behavior observed at all time steps throughout the entire simulation. For cases where convergence to low residual tolerances are required, all schemes can be expected to behave similarly on a grid sweep basis, since the non-linearities are essentially frozen in the asymptotic convergence regime, but the linear multigrid method becomes more efficient simply due to the smaller number of non-linear function evaluations required, a conclusion which is further validated by the intermediate performance of the hybrid scheme. On the other hand, in the initial phases of convergence, the FAS non-linear multigrid approach can achieve faster convergence (particularly on a grid sweep basis) since the important non-linearities are advanced more rapidly as opposed to the linear method which tends to “oversolve” the linear system in these initial phases of convergence. Another advantage of the FAS scheme is its “matrix-free” formulation which avoids the formation and storage of potentially large Jacobian matrices. Note that this advantage is lost in the hybrid scheme, which requires the Jacobian matrix on each grid level, although not simultaneously on all levels.

For this particular problem, the linear multigrid solver provides the overall most efficient solution strategy. In general, the relative performance of linear versus non-linear multigrid methods depends on a tradeoff between the expense of the non-linear function evaluations versus the size and complexity of the Jacobian matrices required by the linearization. For cases where exact Jacobians cannot be easily constructed, (for example when larger discretization stencils are employed), the FAS multigrid approach may be the only feasible multigrid approach for solving the system in a fully coupled manner.

## REFERENCES

- [1] A. BRANDT, *Multigrid techniques with applications to fluid dynamics:1984 guide*, in VKI Lecture Series, Mar. 1984, pp. 1–176.

- [2] P. N. BROWN AND C. S. WOODWARD, *Preconditioning strategies for fully implicit radiation diffusion with material energy transfer*, Tech. Report UCRL-JC-139087, Lawrence Livermore National Lab, May 2000. Submitted to SIAM J. on Scientific Computing.
- [3] R. FALGOUT, V. E. HENSON, J. E. JONES, AND U. M. YANG, *Boomer AMG: A parallel implementation of algebraic multigrid*, in Ninth Society for Industrial and Applied Mathematics Conference on Parallel Processing, San Antonio, TX, Mar. 1999. also available as Lawrence Livermore National Laboratory technical report UCRL-MI-133583, March 1999.
- [4] D. A. KNOLL, W. J. RIDER, AND G. L. OLSON, *An efficient nonlinear solution method for nonequilibrium radiation diffusion*, Tech. Report LA-UR-98-2154, Los Alamos National Lab, 1998. Submitted to J. Quant. Spec. and Rad. Trans.
- [5] ———, *Nonlinear convergence, accuracy and time-step control in non-equilibrium radiation diffusion*, Tech. Report LA-UR-99-3842, Los Alamos National Lab, 1999. Submitted to J. Quant. Spec. and Rad. Trans.
- [6] D. J. MAVRIPLIS, *On convergence acceleration techniques for unstructured meshes*. AIAA paper 98-2966, presented at the 29th AIAA Fluid Dynamics Conference, Albuquerque, NM, June 1998.
- [7] D. J. MAVRIPLIS AND V. VENKATAKRISHNAN, *A unified multigrid solver for the Navier-Stokes equations on mixed element meshes*, International Journal for Computational Fluid Dynamics, 8 (1997), pp. 247–263.
- [8] V. A. MOUSSEAU, D. A. KNOLL, AND W. J. RIDER, *Physics-based preconditioning and the Newton-Krylov method for non-equilibrium radiation diffusion*, Tech. Report LA-UR-99-4230, Los Alamos National Lab, 1999. Submitted to J. Comp. Physics.
- [9] G. C. POMRANING, *The Equations of Radiation Hydrodynamics*, Pergamon, New York, NY, 1973.
- [10] J. W. RUGE AND K. STÜBEN, *Algebraic multigrid*, in Multigrid Methods, S. F. McCormick, ed., SIAM Frontiers in Applied Mathematics, Philadelphia, 1987, SIAM, pp. 73–131.