

A critical study of agglomerated multigrid methods for diffusion on highly-stretched grids

James L. Thomas^{*,a}, Boris Diskin^{**,b,1}, Hiroaki Nishikawa^{b,2}

^aNASA Langley Research Center, Computational AeroSciences Branch, Mail Stop 128, Hampton, Virginia, United States 23681

^bNational Institute of Aerospace, 100 Exploration Way, Hampton, Virginia, United States 23666

Abstract

Agglomerated multigrid methods for unstructured grids are studied critically for solving a model diffusion equation on highly-stretched grids typical of practical viscous simulations, following a previous work focused on isotropic grids. Different primal elements, including prismatic and tetrahedral elements in three dimensions, are considered. The components of an efficient node-centered full-coarsening multigrid scheme are identified and assessed using quantitative analysis methods. Fast grid-independent convergence is demonstrated for mixed-element grids composed of tetrahedral elements in the isotropic regions and prismatic elements in the highly-stretched regions. Implicit lines natural to advancing-layer/advancing-front grid generation techniques are essential elements of both relaxation and agglomeration. On agglomerated grids, consistent average-least-square discretizations augmented with edge-directional gradients to increase h -ellipticity of the operator are used. Simpler (edge-terms-only) coarse-grid discretizations are also studied and shown to produce grid-dependent convergence — only effective on grids with minimal skewing.

Key words: multigrid, agglomeration, unstructured, diffusion, analysis

1. Introduction

Multigrid techniques [18] are routinely used to accelerate convergence of Reynolds-Averaged Navier-Stokes solvers for large-scale steady and unsteady flow applications, especially within structured-grid methods. Agglomerated multigrid methods for large-scale unstructured-grid applications have also been developed and demonstrated impressive improvements in efficiency over single-grid computations [9–12]. The performance of multigrid solvers is as yet far from the textbook multigrid efficiency goal — converging algebraic errors below discretization errors in the work equivalent to a few residual evaluations; such performance has only been demonstrated to date for relatively simple applications [15, 16]. Design of efficient multigrid solvers for unstructured grid applications is significantly more challenging because analysis tools to understand and predict multigrid performance are less developed than tools for structured grids. In particular, local Fourier analysis (LFA) is widely used on structured grids but is inapplicable to irregular grids. The quantitative analysis tools, idealized relaxation

and idealized coarse grid, developed earlier [2] are applicable. These tools, in combination with windowing techniques [3, 17], isolate the sources of difficulties and are proving useful to improve both accuracy and efficiency in an unstructured-grid setting.

One of the key weaknesses identified by Venkatakrishnan [19] for unstructured agglomeration methods was the coarse-grid discretization of diffusion (viscous terms). The current approaches for the coarse-grid discretization of diffusion were critically studied for two- and three-dimensional isotropic grids in a previous paper [13]. Direct-discretization and Galerkin approaches were investigated for a model problem representative of laminar diffusion in the incompressible limit. Consistency of coarse-grid discretization was found to be essential for attaining fast grid-independent convergence; consistent discretizations on agglomerated grids were obtained through direct discretization with an average-least-square approach. Multigrid with coarse grids discretized using either a Galerkin approach or an approximate edge-terms-only direct discretization was also studied but, with both of these approaches, the convergence depended on the grid (particularly skewness) and deteriorated on finer grids. In this paper, we address higher aspect ratios and highly-stretched three-dimensional grids and use only direct discretizations.

Many applications use grids generated with advancing-layer/advancing-front techniques in which the grids are highly stretched predominantly in the direction normal to the boundary. In this paper, highly-stretched grids transitioning to isotropic grids are considered. The isotropic grids are irregular tetrahedral grids. The highly-stretched grids are mixed-element grids, composed of prismatic and tetrahedral elements; the pris-

*Corresponding author

**University of Virginia, Department of Mechanical and Aerospace Engineering Department, 122 Engineer's Way, Charlottesville, Virginia 22904-4746.

Email addresses: James.L.Thomas@nasa.gov (James L. Thomas), bdiskin@nianet.org (Boris Diskin), hiro@nianet.org (Hiroaki Nishikawa)

¹This author was supported by the National Institute of Aerospace under the NASA Fundamental Aeronautics Program through NASA Contracts NNL07AA23C and NNL07AA31C.

²This author was supported by the National Institute of Aerospace under the NASA Fundamental Aeronautics Program through NASA Contract NNL07AA23C.

matic grids extend from the the surface, where the aspect ratio is highest, to locations where the aspect ratio approaches unity. A full-coarsening/line-implicit multigrid is pursued herein. The coarsening strategy is similar to that used by Hyams et al. [8], although the coarse-grid discretizations are quite different. In [8], a Galerkin coarse-grid construction that is inconsistent for diffusion was used; a direct discretization on the coarse grid was also used but no details of the treatment of viscous terms are given. Mavriplis [9–12] used a directional-coarsening strategy — coarsening by a factor of four in the direction normal to the boundary within the highly-stretched (viscous) regions of the grid; a full coarsening strategy was used in the isotropic (inviscid) regions of the grid. The coarse-grid discretization of viscous terms was through an edge-terms-only direct discretization or a heuristically-scaled Galerkin formulation.

This paper is organized as follows. The discretization schemes for the model diffusion equation are presented in Section 2 from a general finite-volume discretization standpoint. Element-based and element-free schemes are shown; the latter includes certain edge-based discretizations and discretizations on agglomerated grids. The grid agglomeration techniques are presented in Section 3 and Appendix A. The multigrid algorithm, including relaxation and residual-averaging techniques, is described in Section 4. The key ingredients enabling successful multigrid performance are identified and assessed using quantitative analysis methods in Section 5 and Appendices B–D. Three-dimensional multigrid computations demonstrating grid-independent convergence for both isotropic and highly-stretched grids within an ellipsoidal domain are shown in Section 6. The final Section 7 contains conclusions.

2. Discretization schemes

The considered model problem is the Poisson equation

$$\Delta U = f, \quad (1)$$

subject to Dirichlet boundary conditions; function f is a forcing function. The finite-volume discretization (FVD) schemes are derived from the integral form of a conservation law

$$\oint_{\partial\Omega} \nabla U \cdot \hat{\mathbf{n}} \, ds = \int_{\Omega} f \, d\Omega, \quad (2)$$

where ∇U is the solution gradient, Ω is a control volume with boundary $\partial\Omega$, and $\hat{\mathbf{n}}$ is the outward unit normal vector. The general FVD approach requires partitioning the domain into a set of non-overlapping control volumes and numerically implementing equation (2) over each control volume.

Node-centered discretizations are considered in which the solutions are defined at the mesh nodes. The discrete schemes described below are representative of viscous discretizations used in Reynolds-Averaged Navier-Stokes unstructured-grid codes. Dirichlet boundary conditions are implemented strongly.

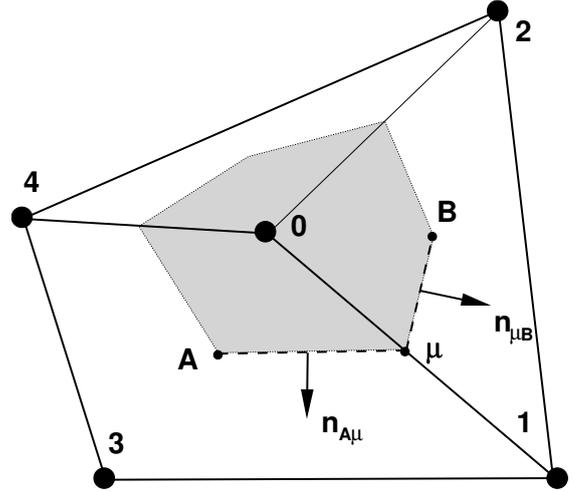


Figure 1: Illustration for gradient construction; dual volume is shaded.

2.1. Element-based discretizations

The target meshes are compositions of primal elements (cells) — triangular and quadrilateral elements in two dimensions (2D) and tetrahedral, hexahedral, prismatic, and pyramidal elements in three dimensions (3D). Control volumes are constructed around the mesh nodes by the median-dual partition (Figure 1) [1, 7].

The target discretization is the Green-Gauss scheme [4] — widely used in node-centered codes and equivalent to a Galerkin finite-element (linear-element) discretization for triangular/tetrahedral grids. For mixed elements, edge derivatives are used to increase the h -ellipticity [18] of the diffusion operator [4, 7] and, thus, avoid checkerboard instabilities. It has been shown [3, 17] that the scheme possesses second-order accuracy for viscous fluxes on general mixed-element grids.

With reference to Figure 1 illustrating a mixed-element 2D grid, the scheme approximates the integral flux through the dual faces adjacent to the edge $[0, 1]$ as

$$\int_{A_{\mu B}} \nabla U \cdot \hat{\mathbf{n}} \, ds = \nabla U_{A\mu} \cdot \mathbf{n}_{A\mu} + \nabla U_{\mu B} \cdot \mathbf{n}_{\mu B}, \quad (3)$$

where μ is the median of the edge $[0, 1]$, subscripts designate dual faces, and $\mathbf{n}_{A\mu}$ and $\mathbf{n}_{\mu B}$ are directed-area vectors. The gradient is reconstructed separately at each dual face as follows. For the triangular element contribution, the gradient is determined from a Green-Gauss evaluation at the primal element,

$$\nabla U_{\mu B} = \overline{\nabla U}_{012}. \quad (4)$$

The gradient overbar denotes a gradient evaluated by the Green-Gauss formula on the primal cell identified by the point subscripts. For the quadrilateral element contribution, the gradient $\nabla U_{A\mu}$ is formed by augmenting the Green-Gauss gradient within the element, $\overline{\nabla U}_{0134}$, with the directional derivative along the edge, $\partial^e U$, defined as

$$\partial^e U = (U_1 - U_0) / |\mathbf{r}_1 - \mathbf{r}_0|, \quad (5)$$

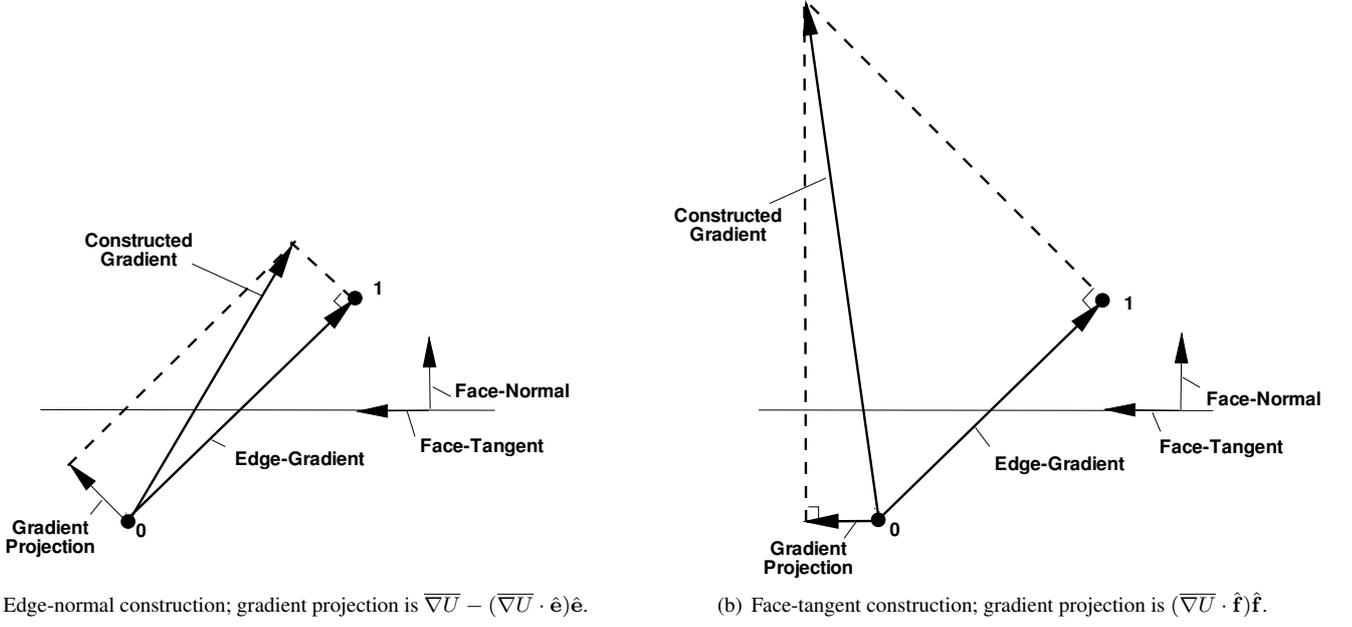


Figure 2: Illustration of gradient constructions at a control-volume face separating nodes 0 and 1; $\theta = \pi/4$; the edge gradient has magnitude $\partial^e U$ and is oriented in the \hat{e} direction.

where U_i and \mathbf{r}_i are the solution and the coordinate vector of the node i .

Two approaches to construct the augmented gradient $\nabla U_{A\mu}$ have been used and are illustrated in Figure 2 for 2D. To present the approaches, the unit vector aligned with the edge $[0, 1]$ is defined as

$$\hat{e} = (\mathbf{r}_1 - \mathbf{r}_0) / |\mathbf{r}_1 - \mathbf{r}_0|, \quad (6)$$

the unit vector normal to the control-volume face is defined as

$$\hat{\mathbf{n}} = \mathbf{n}_{A\mu} / |\mathbf{n}_{A\mu}|, \quad (7)$$

and the Green-Gauss gradient is defined as

$$\overline{\nabla U} = \overline{\nabla U}_{0134}. \quad (8)$$

The face skew angle θ is the angle between the edge direction and the face-normal direction,

$$\cos \theta = \hat{e} \cdot \hat{\mathbf{n}}. \quad (9)$$

The first augmentation, probably more widely used and designated here as *edge-normal* (EN), is illustrated in Figure 2(a) and enforces that the constructed gradient, $\nabla U_{A\mu}^{EN}$, recovers (1) the edge-directional gradient, $\partial^e U \hat{e}$, and (2) the Green-Gauss gradient projected onto the plane normal to \hat{e} ,

$$\nabla U_{A\mu}^{EN} = (\partial^e U - (\overline{\nabla U} \cdot \hat{e}))\hat{e} + \overline{\nabla U}. \quad (10)$$

The second augmentation, designated as *face-tangent* (FT), is illustrated in Figure 2(b) and enforces that the constructed gradient, $\nabla U_{A\mu}^{FT}$, recovers (1) the edge directional gradient and (2) the Green-Gauss gradient projected onto the plane normal to $\hat{\mathbf{n}}$,

$$\nabla U_{A\mu}^{FT} = \frac{1}{\hat{\mathbf{n}} \cdot \hat{e}} \partial^e U \hat{\mathbf{n}} + (\overline{\nabla U} \cdot \hat{\mathbf{f}}) [\hat{\mathbf{f}} - \frac{\hat{\mathbf{f}} \cdot \hat{e}}{\hat{\mathbf{n}} \cdot \hat{e}} \hat{\mathbf{n}}], \quad (11)$$

where $\hat{\mathbf{f}}$ is a unit vector normal to $\hat{\mathbf{n}}$. Note that (11) applies only to 2D but there is an obvious 3D counterpart. The corresponding contributions to the diffusion operator (for the orientation shown in Figure 2) are given below:

$$\nabla U_{A\mu}^{EN} \cdot \mathbf{n}_{A\mu} = |\mathbf{n}_{A\mu}| [\cos \theta (\partial^e U - (\overline{\nabla U} \cdot \hat{e})) + \overline{\nabla U} \cdot \hat{\mathbf{n}}], \quad (12)$$

$$\nabla U_{A\mu}^{FT} \cdot \mathbf{n}_{A\mu} = |\mathbf{n}_{A\mu}| \frac{1}{\cos \theta} [\partial^e U + (\overline{\nabla U} \cdot \hat{\mathbf{f}}) \sin \theta]. \quad (13)$$

Both approaches to gradient augmentation improve the h -ellipticity of the operator; for dual faces with zero skew angle, the edge-directional derivative, $\partial^e U$, is the only contributor. Hasselbacher [7] considered both formulations but used the EN formulation in computations. The FT formulation is identical to the approach used in a sheared mapped quadrilateral grid, i.e., the gradient is recovered from directional gradients in the mapped coordinate directions.

The FT formulation has been found to be more robust for highly-skewed grids and was used for cell-centered applications in [4]. The rationale is that, in such applications, the relative contributions from the edge gradient to the diffusion operator are much larger than with the EN formulation. Comparing (12) and (13), the flux contribution of $\partial^e U$ with EN augmentation is $\cos \theta$ (less than 1) versus $1/\cos \theta$ (greater than 1) with FT augmentation. Likewise, any contributions from $\partial^e U$ with the

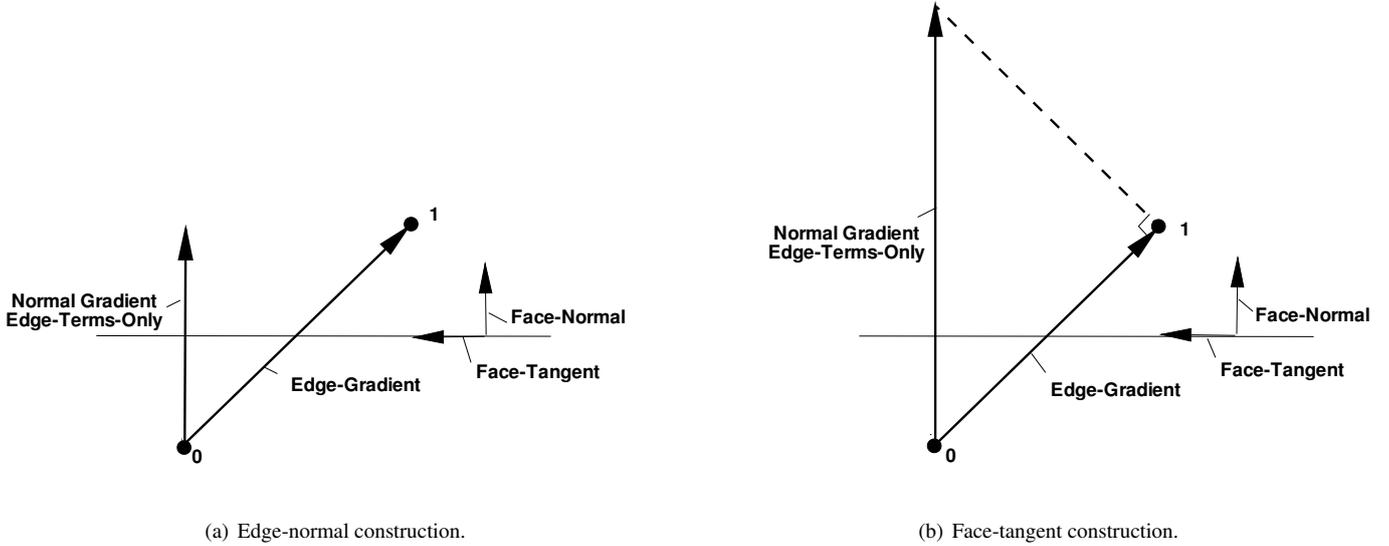


Figure 3: Illustration of gradient constructions at a control-volume face separating nodes 0 and 1 using only edge gradients.

EN formulation vanish for θ approaching $\pi/2$. The face-normal gradient, entirely neglecting the projected Green-Gauss gradient, is shown in Figure 3; the differences in the diffusion operator are easily seen to be a factor of two corresponding to the particular value of $\theta = \pi/4$.

The skew angle can approach $\pi/2$ on primal grids and even exceed $\pi/2$ on agglomerated grids, resulting in a destabilizing edge contribution for both approaches to augmentation. We have elected to neglect the entire flux at faces with $\theta \geq \pi/2$. An alternate approach, implemented as yet only in 2D, is to simply discard the directional derivative contribution.

2.2. Element-free discretizations

Two element-free discretizations are described below; at a minimum, they are needed in multigrid because the element-based data structures are not retained on agglomerated grids. Additionally, they can be used on the target grids — either to reduce computational cost or serve as drivers in relaxation.

Referring to Figure 1, the element-free schemes approximate the integral flux through the dual faces adjacent to the edge $[0, 1]$ as

$$\int_{A_{\mu B}} \nabla U \cdot \hat{\mathbf{n}} ds = \nabla U_{\mu} \cdot \mathbf{n}_{\mu}, \quad (14)$$

where the directed area, \mathbf{n}_{μ} , is a lumped approximation,

$$\mathbf{n}_{\mu} = \mathbf{n}_{A_{\mu}} + \mathbf{n}_{B_{\mu}}. \quad (15)$$

The first scheme to approximate ∇U_{μ} , herein referred as *Edge-Terms-Only* (ETO), has already been introduced (Figure 3) and is often referenced in the literature as a *thin-layer approximation*. Both edge-normal, ETO (EN), and face-tangent, ETO (FT), constructions can be used — either can be considered a thin-layer scheme. The gradient ∇U_{μ} is constructed

using the right sides of either (10) or (11) retaining only the contributions from the $\partial^e U$ terms. The scheme is a positive scheme but on non-orthogonal grids (non-zero skew angles), it is not consistent (i.e., discrete solutions do not converge to the exact continuous solution with consistent grid refinement) [3, 5, 13, 14]. The inconsistencies are most noticeable on grids with persistently-high skew angles — high-aspect-ratio tetrahedral meshes, for example.

The second scheme is the *average-least-squares* (Avg-LSQ) scheme. The gradient ∇U_{μ} is constructed using the right sides of either (10) or (11) with the gradient $\overline{\nabla U}$ replaced by the average of the least-squares (LSQ) gradients computed at the two nodes associated with the edge. The stencil of the LSQ gradient at a node includes all edge-connected neighbors. The LSQ minimization enforces the given solution at the central node.

3. Agglomerated grids

The control volumes of each agglomerated grid are found by summing control volumes of a finer grid. Any agglomerated grid can be defined in terms of a conservative agglomeration operator, R_0 , as

$$\Omega^c = R_0 \Omega^f, \quad (16)$$

where superscripts c and f denote entities on coarser and finer grids, respectively. On the agglomerated grids, the control volumes become geometrically more complex than their primal counterparts and the details of the control-volume boundaries are not retained. The directed area of a coarse-grid face separating two agglomerated control volumes, if required, is found by lumping the directed areas of the corresponding finer-grid faces and is assigned to the virtual edge connecting the centers of the neighboring agglomerated control volumes.

As described more fully in [13], the grids are agglomerated within a topology-preserving framework, in which hierarchies are assigned based on connections to the computational boundaries and surface discontinuities. Corners are identified as grid points with three or more boundary-condition-type closures (or two or more boundary slope discontinuities). Ridges are identified as grid points with two boundary-condition-type closures (or one boundary slope discontinuity). Valleys are identified as grid points with a single boundary-condition-type closure and interiors are identified as grid points with no boundary closure. The agglomerations proceed hierarchically from seeds within the topologies, first corners, then ridges, then valleys, and finally interiors. Rules are enforced to maintain the boundary condition types of the finer grid within the agglomerated grid. For example, a ridge can be agglomerated into an existing ridge agglomeration only if the two boundary conditions associated with each ridge are the same. Hierarchies on each agglomerated grid are inherited from the finer grid.

There are two main difficulties associated with the current agglomeration techniques. The first is that after agglomeration, there may be insufficient connections to construct the least-square gradient at a node. This occurs most often near boundaries and, to improve reliability for complex geometries, we have adopted a boundary agglomeration step, in which corners, ridges, and valleys are agglomerated first — but agglomerations are allowed only within the same hierarchy. Thus, corners are never agglomerated. Ridges can be agglomerated only with ridges and valleys can be agglomerated only with valleys. These rules guarantee a valid non-degenerate LSQ stencil near boundaries. The downside is that the agglomerated grids have volumes near features much smaller than the interior volumes, especially on coarser grids. A better approach, implemented as yet only in 2D, is to augment the edge-connections as needed to construct gradients at a control-volume.

The second difficulty, occurring more frequently in 3D than in 2D, is that large skew angles ($\theta \geq \pi/2$) are encountered on agglomerated grid faces. As discussed earlier, we neglect the entire flux at these faces in 3D. Another possible strategy is to control the shape of the agglomerations, either during agglomeration or in a post-processing step, in order to avoid extreme face skewness.

Typical isotropic grids are shown in Figures 4-5, corresponding to a target grid and a first-level agglomeration, respectively. The target grids are all tetrahedral grids and are irregular because of 3D random node perturbations. The grids were constructed in a cubic domain and then mapped onto an ellipsoid. In the cubic domain, the grids are perturbed in each coordinate direction with magnitude $1/4$ of the local mesh spacing.

Typical stretched grids are shown in Figures 6-7. A prismatic layer is first generated from a triangulated boundary; the boundary grids include random node perturbations within the boundary surface. The prismatic layer occupies the lower quarter of the domain for all grid sizes. The maximum aspect ratio of 10^3 is enforced for cells at the bottom, where the aspect ratio is defined as a ratio of the mesh spacings tangent and normal to the boundary. Nodes in the prismatic layer were generated by a geometric sequence such that the aspect ratio approaches unity

Target-Grid Discretization	Coarse-Grid Discretization
Green-Gauss (EN)	Avg-LSQ (EN) ETO (EN)
Green-Gauss (FT)	Avg-LSQ (FT) ETO (FT)

Table 1: Summary of target-grid and coarse-grid discretizations; gradient augmentation is denoted in parentheses.

at line terminations. The number of nodes per line is thus automatically determined. An isotropic tetrahedra grid with random 3D node perturbations is then added.

For highly stretched meshes, the advancing front agglomeration is first applied at the boundary of the grid (corners, ridges, and valleys) containing the origins of the implicit lines. Then interior duals are agglomerated, two at a time in the normal direction, from the boundary to the line terminations, preserving the prismatic structure of the agglomerations. After the line agglomerations, the front agglomeration method is applied over the remainder of the domain. The overall agglomeration technique is similar to that of Hyams et al. [8]

For both isotropic and stretched grids, a sequence of 15 target grids were generated to assess multigrid convergence. In Appendix A, details of the sequences are given and additional statistics for two grids are given.

4. Multigrid

Elements of the multigrid algorithm are presented in this section. A V-cycle [18], denoted as $V(\nu_1, \nu_2)$, uses ν_1 relaxations performed at each grid before proceeding to the coarser grid and ν_2 relaxations after coarse-grid correction; the coarsest grid is solved exactly (with many relaxations). Residuals, r^f , corresponding to the fine-grid discretization of the integral equation (2) are restricted to the coarse grid using the conservative agglomeration operator R_0 , defined in (16), and a residual averaging operator, W , as

$$r^c = R_0 W r^f. \quad (17)$$

The residual averaging is performed by replacing the individual residual at a node by the arithmetic average of the residuals over its neighbor nodes. For simplicity of implementation, the averaging is not performed over boundary nodes or nodes that connect to a boundary. Note that averaging, e.g., full-weighting, of residuals is necessary with multicolor relaxation schemes even in classical structured-grid multigrid methods because the residuals of the last color are reduced identically to zero. The fine-grid solution approximation is restricted to the coarse grid as

$$U^c = \frac{R_0(U^f \Omega^f)}{\Omega^c}. \quad (18)$$

The prolongations P_0 and P_1 are exact for piecewise-constant and linear functions, respectively. The prolongation P_0 is the transpose of R_0 . The operator P_1 is constructed locally using linear interpolation from a triangle (2D) or tetrahedra (3D) defined on the coarse grid. The geometrical shape

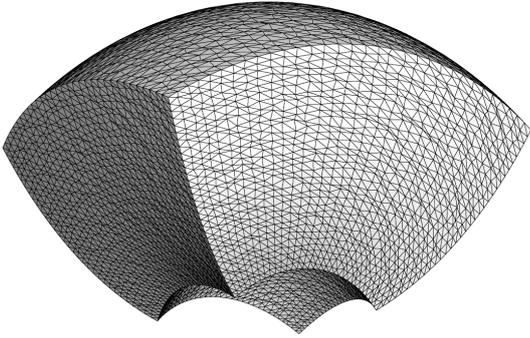


Figure 4: Target isotropic 33x33x33 grid.

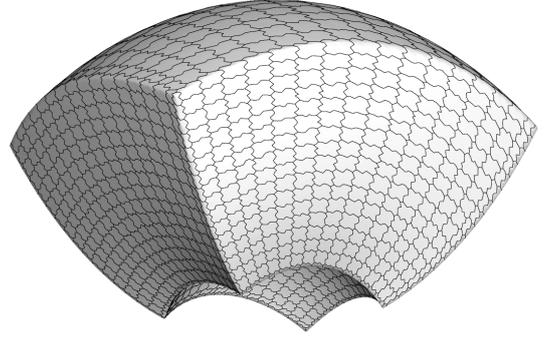


Figure 5: First-level agglomeration generated from the target isotropic 33x33x33 grid.

is anchored at the coarser-grid location of the agglomerate that contains the given finer control volume. Other nearby points are found using the adjacency graph. An enclosing simplex is sought that avoids prolongation with non-convex weights and, in situations where multiple geometrical shapes are found, the first one encountered is used. At locations where this procedure result in non-convex weights, the prolongation is reverted locally to piecewise-constant prolongation. The prolongation operator P_1 is modified to prolong only from hierarchies equal or above the hierarchy of the prolonged point. The correction δU to the finer grid is prolonged typically through P_1 , as

$$(\delta U)^f = P_1(\delta U)^c. \quad (19)$$

The available target-grid and coarse-grid discretizations are listed in Table 1. The main target discretization of interest is the element-based Green-Gauss scheme discussed earlier with either of the two approaches to gradient augmentation for non-simplicial elements. There are four available element-free coarse-grid discretizations, the consistent Avg-LSQ scheme and the inconsistent but widely-used ETO scheme, each evaluated with the same approach to gradient augmentation used on the target grid for simplicity.

The exact linear operator is used in the iterative phase of the Green-Gauss scheme, enabling a robust multicolor Gauss-Seidel relaxation. The Avg-LSQ scheme has a comparatively larger stencil and its exact linearization is not used in iterations; instead relaxation of the Avg-LSQ scheme relies on the ETO linearization as a driver. It is known that the smoothing rate with this approach can deteriorate on highly skewed grids [4].

5. Analysis

5.1. Idealized relaxation and idealized coarse grid methods

This section presents quantitative analysis tools, idealized relaxation (IR) and idealized coarse-grid (ICG) iterations, for as-

essment and improvement of unstructured multigrid solvers. IR and ICG have been applied earlier [13] to analyze multigrid solvers on isotropic unstructured grids; applications to high-aspect-ratio grids are studied below.

It is long known [18] that convergence of full-coarsening multigrid with point relaxation deteriorates on grids with high aspect ratio. Failure of point relaxation to smooth errors oscillating in the direction of weak coupling (larger mesh spacing) is the main reason for convergence deterioration. Typical remedies involve implicit relaxation, semi-coarsening, or a combination of the two. In this paper, multigrid employs full-coarsening and line-implicit relaxation.

Testing of multigrid solvers with line-implicit relaxation schemes on high-aspect-ratio grids is not straightforward. At the initial design stages, the performance of a multigrid cycle is typically tested on either small low-density grids or with Dirichlet conditions imposed at boundaries of the high-aspect-ratio regions. On such grids, a line-implicit relaxation scheme becomes a solver rather than a smoother and provides overly optimistic predictions [18]. IR and ICG cycles, similarly to LFA, avoid this difficulty and can expose problems that may arise only in applications with extremely large numbers of degrees of freedom.

Specifically the IR and ICG methods focus on the main complementary parts of a multigrid cycle: relaxation and coarse-grid correction. Each part of the cycle is assigned a task, e.g., relaxation is typically assigned to smooth errors, coarse-grid correction is typically assigned to reduce all smooth error components. In the analysis, idealized iterations probe the actual two-grid cycle to identify parts limiting the overall effectiveness.

The IR and ICG iterations can be applied to any formulation with a manufactured solution; here they are applied to a formulation with zero manufactured solution. The initial guess is formed by a random perturbation of the solution. In the anal-

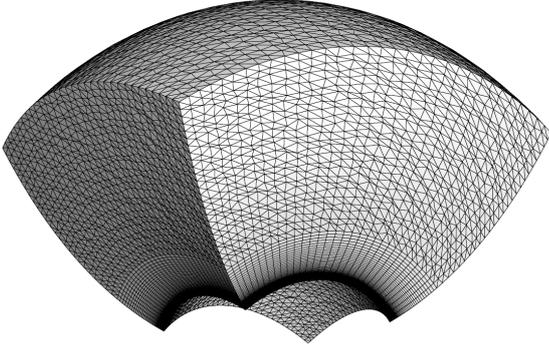


Figure 6: Target stretched 33x33x134 grid.

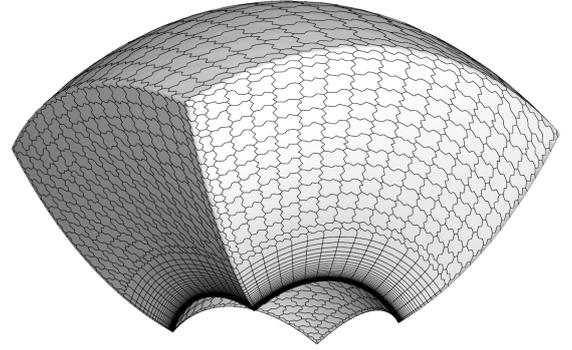


Figure 7: First-level agglomeration generated from the target isotropic 33x33x134 grid.

ysis, one part of the tested cycle is replaced with an idealized imitation. The idealized imitations do not depend on the operators to be solved. Rather, they are numerical procedures acting directly on the known algebraic error to fulfill the task assigned to the corresponding part of the two-grid cycle. The results of the analysis are convergence patterns of the iterations that may either confirm or refute expectations as to how well each part of the actual cycle is carrying out the assigned task.

With IR cycles, the coarse-grid correction part is actual and the relaxation is idealized. Idealized relaxation can be implemented by constructing a pseudo-Laplacian operator, A^{IR} , which includes nodes linked by an edge, or possibly an element through a virtual edge, to a given node, as below,

$$A^{IR}\epsilon \equiv \sum_{i=1}^{N_e} w_i(\epsilon_i - \epsilon_0) = 0. \quad (20)$$

Here, N_e is the number of edges connected to node 0, the algebraic error at node i is ϵ_i , and w_i represents a weight. The choice $w_i = 1$ yields a positive operator. A few relaxations of (20) serve as an idealized relaxation.

With ICG cycles, the relaxation scheme is actual and the coarse-grid correction is idealized. The ICG correction used for unstructured multigrid computations is defined in the following 2 steps: (1) The algebraic error is restricted to the coarse grid by a volume-averaging operator, similarly to the solution restriction (18). (2) The volume-averaged error is interpolated back to the fine grid as a correction. This procedure effectively reduces all smooth error components.

An important check of the quality of chosen idealized components is convergence of the “reference cycle,” which uses both idealized components in iterations. The convergence rate of the reference cycle represents a sensitivity threshold in that idealized iterations generally suggest some meaningful im-

provements only for actual cycles with convergence rates significantly slower than this threshold.

The idealizations used in IR and ICG analysis are not unique. Within high-aspect-ratio grid regions, we consider a line-implicit IR scheme, designated IR-L, that simultaneously changes algebraic errors at all nodes of the same grid line such that the updated algebraic errors satisfy (20); the lines are visited in a 2-color order. The selection is justified through LFA of regular quadrilateral and triangular grids in Appendices B-C. Details of the LFA methodology are summarized in Appendix B. Several point- and line-implicit idealized relaxations performed in various orders are analyzed in Appendix C. Within isotropic grid regions, an idealized relaxation with multicolor point-wise error averaging, designated IR-P, is used. Appendix D presents observations on convergence rates of IR-P and actual cycles on isotropic unstructured grids. The two idealized relaxations, IR-P and IR-L, overlap by a single node per line for stretched grids including isotropic and high-aspect-ratio regions.

5.2. Applications to triangular grids

Illustrative 2-grid computations are performed on a sequence of regular triangular grids with uniform aspect ratio $A = 10^3$. Fine-grid and coarse-grid control volumes are illustrated in Figures 8-9. Note that on the fine grid, the Green-Gauss discretization is equivalent to a classical 5-point Laplacian [4].

Table 2 shows asymptotic convergence rates with IR-L and residual averaging for various coarse-grid discretizations. We do not show actual relaxations because Dirichlet conditions were used in the computations and the line-implicit relaxation solves the equations in a single iteration. For comparisons with the rates one would observe in computations on large grids, Table 3 shows convergence rates computed with LFA using methodology presented in Appendix B. On regular grids, LFA

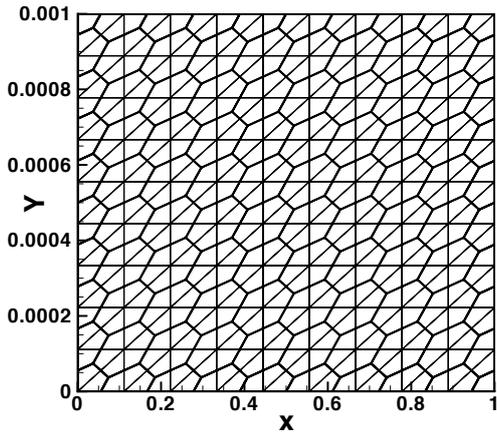


Figure 8: Control volume boundaries (heavier lines) for regular triangular fine grid.

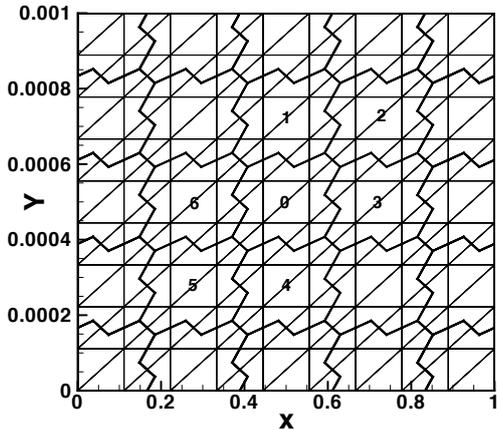


Figure 9: Control volume boundaries (heavier lines) for regular triangular coarse grid.

is known to provide accurate predictions of multigrid convergence.

All analysis methods indicate that only discretizations with EN augmentation allow fast grid-independent convergence on high-aspect ratio triangular grids. Convergence of multigrid with coarse-grid discretizations using FT augmentation approaches unity in the limit of grid refinement.

The reason for the striking differences between EN and FT approaches to augmentation can be traced directly to the high skewing of the coarse grid shown in Figure 9. Considering a fully-interior control volume, there are 6 face-connections to the surrounding control volumes. Two of these faces (connecting node 0 with nodes 1 and 4, respectively, in Figure 9) have nearly-zero skew angle and the other four faces have skew angles approaching $\pi/2$. Considering the discrete diffusion terms in the y -direction, the coarse-grid ETO (EN) operator is inconsistent, being $5/6$ of the fine-grid operator. However, this is sufficient to yield a convergence rate of 0.2 per multigrid cy-

Fine Grid	Avg-LSQ (EN)	ETO (EN)	Avg-LSQ (FT)	ETO (FT)
32x32	<0.1	0.16	0.13	0.32
64x64	<0.1	0.16	0.28	0.56
128x128	<0.1	0.18	0.44	0.73

Table 2: Asymptotic convergence rates for IR-L cycle; regular triangular grid; $\nu_1 = \nu_2 = 2$.

	Avg-LSQ (EN)	ETO (EN)	Avg-LSQ (FT)	ETO (FT)
LFA				
IR-L	0.12	0.20	1.0	1.0
Actual	0.07	0.19	1.0	1.0

Table 3: LFA 2-grid convergence rates for IR-L and actual line-implicit cycles; regular triangular grid; $\nu_1 = \nu_2 = 2$; piecewise-constant prolongation.

cle. The coarse-grid Avg-LSQ (EN) scheme is consistent and provides an order of magnitude error reduction per multigrid cycle. Additional details and specific formulas are provided in Appendix B.

Schemes with the FT gradient augmentation magnify the skewed-face contributions to the diffusion operator. The Avg-LSQ (FT) scheme leads to a wide-stencil coarse-grid operator, poorly approximating the fine-grid medium-range error components oscillating in the x -direction. The ETO (FT) scheme leads to completely inaccurate approximations (see additional details in Appendix B).

In these regular-grid computations, the control-volume centers on the coarse grid remain perfectly collinear. In general, any departure from the perfect alignment, such as with an irregular triangularization of the fine grid and a volume-weighted construction of the coarse-grid control-volume locations, can result in high skew angles at all faces. In this situation, the ETO (EN) scheme becomes inadequate. The Avg-LSQ (EN) scheme loses h -ellipticity and, at a minimum, becomes difficult to converge.

All of the above issues associated with highly-skewed faces on high-aspect-ratio grids are avoided if prismatic (quadrilateral in 2D) grids are used with the line agglomeration discussed earlier. The skewness of the coarse grid is then comparable with the skewness of the fine grid and convergence rates for all schemes become an order of magnitude per cycle. Thus, only prismatic grids are used in highly-stretched regions for the computations that follow.

6. Three-dimensional results

In this section, we present 3D multigrid convergence rates for the sequences of isotropic and stretched grids listed in Appendix A; for each of the 15 grids in the sequence, multigrid employs all available levels. Initial conditions on each grid were taken as random and the convergence was terminated when integral-equation residuals reached machine-precision level. Figures 10-11 show multigrid convergence rates versus

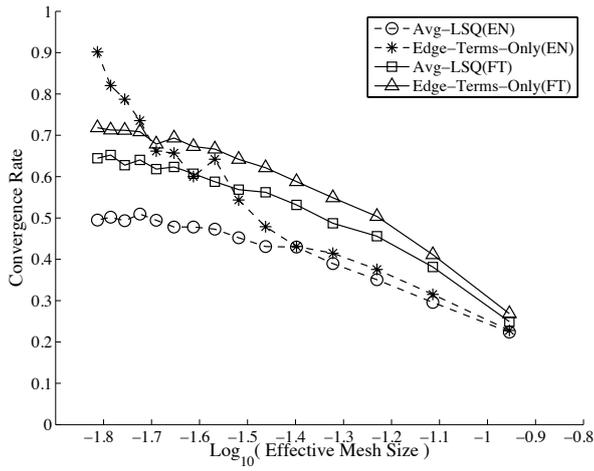


Figure 10: Convergence rate versus effective mesh size for isotropic grids; $\nu_1 = 2; \nu_2 = 1$.

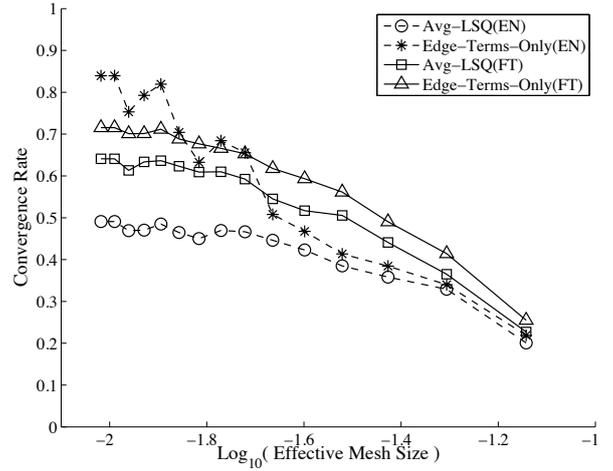


Figure 11: Convergence rate versus effective mesh size for stretched grids; $\nu_1 = 2; \nu_2 = 1$.

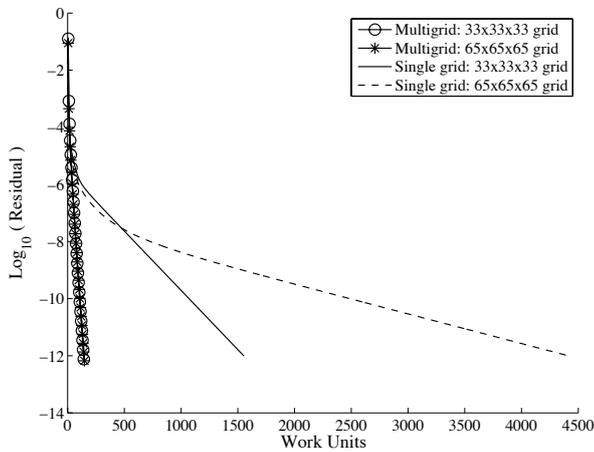


Figure 12: Convergence versus work units for two isotropic grids; $\nu_1 = 2; \nu_2 = 1$; coarse-grid discretization is the Avg-LSQ (EN) scheme.

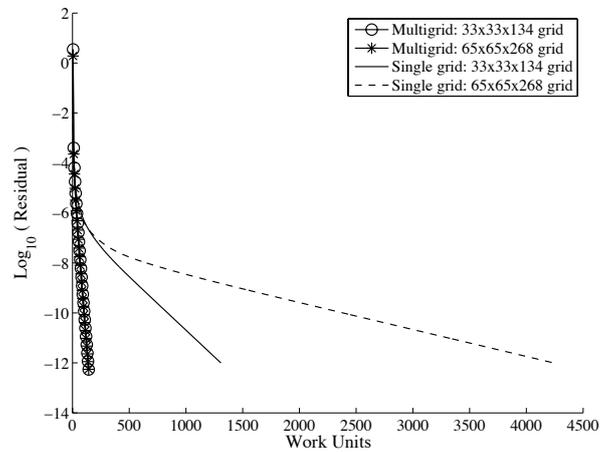


Figure 13: Convergence versus work units for two stretched grids; $\nu_1 = 2; \nu_2 = 1$; coarse-grid discretization is the Avg-LSQ (EN) scheme.

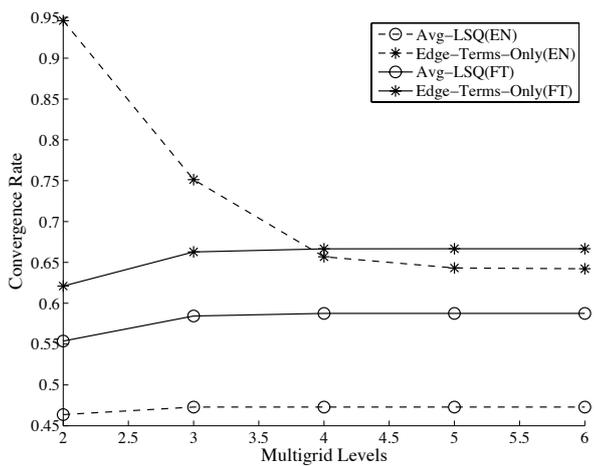


Figure 14: Convergence rates versus multigrid levels for a $37 \times 37 \times 37$ isotropic grid; $\nu_1 = 2; \nu_2 = 1$.

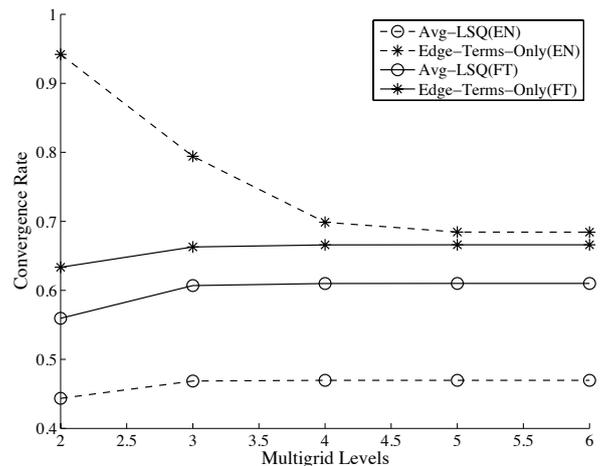


Figure 15: Convergence rates versus multigrid levels for a $37 \times 37 \times 150$ stretched grid; $\nu_1 = 2; \nu_2 = 1$.

the effective mesh size for each of the coarse-grid discretizations. The effective mesh size is defined as the reciprocal of the cube root of the total number of nodes. The convergence rate is computed as an average of per-cycle convergence rates over the last 4 multigrid cycles. In grid refinement, the convergence rates approach grid-independent levels for the Avg-LSQ (EN), Avg-LSQ (FT), and ETO (FT) schemes; the best convergence rate is obtained with the Avg-LSQ (EN) scheme. Observe that the convergence with these schemes for stretched grids is as good as convergence for isotropic grids.

To demonstrate the essentially grid-independent convergence with the Avg-LSQ (EN) coarse-grid discretization, single-grid and multigrid computations are compared in Figures 12-13 for isotropic and stretched grids, respectively. Convergence for two grids, one finer by a factor of two in each direction, are shown. The integral-equation residual is shown versus work units, taken as the number of residual evaluations on the fine grid. For the current Full Approximation Scheme [18] multigrid implementation, the work units per cycle are estimated as $(\nu_1 + \nu_2 + 2)(1 + 1/8 + 1/64 + \dots)$. The results show the expected slowdown of the single-grid scheme on the finer grid. The finer-grid residual convergence over-plots that of the coarser grid with the multigrid scheme.

Multigrid convergence of the ETO (EN) scheme is highly grid-dependent, slowing down on finer grids for both isotropic and stretched grids. These results confirm the conclusions drawn from the previous study [13] for isotropic tetrahedra on cubical domains — multigrid convergence is grid-dependent with the ETO (EN) scheme and grid-independent with the Avg-LSQ (EN) scheme.

During the numerical experiments, it was observed that, contrary to usual expectations, multigrid with the ETO (EN) scheme converges better with multiple levels than with two levels (the coarsest problem is fully solved in all cases). Figures 14-15 show convergence rates versus multigrid levels for the two grids listed in Table 5 and 6. The existence of faces with skew angles greater than $\pi/2$ do not appear to have a negative impact on convergence for the Avg-LSQ schemes; 2-level convergence is comparable with multi-level convergence. It is not surprising that multigrid with the ETO (EN) schemes exhibits grid-dependent convergence because the scheme is inconsistent. What is surprising is that the ETO (FT) scheme does not fail (see Figures 10-11). Although we do not show the results here, for more realistic complex geometries, we have found that multigrid with either ETO scheme fails to converge.

7. Conclusions

Agglomerated multigrid techniques used in unstructured-grid methods have been critically studied for a model problem representative of laminar diffusion in the incompressible limit, with a focus on highly-stretched grids. A multigrid solver for a node-centered element-based discretization has been investigated with several different coarse-grid discretizations on agglomerated grids. Quantitative analysis methods have been used to identify and assess elements of the solver that perform well in high-aspect-ratio regions. The elements of multigrid

enabling grid-independent convergence rates are the following: (1) a consistent coarse-grid discretization; (2) prismatic elements with line relaxation and line agglomeration in the stretched grid regions; and (3) residual averaging of the conservative residuals before restriction. The convergence rates per cycle on mixed-element grids with highly-stretched regions are commensurate with the convergence rates on isotropic grids.

Analyses and computations show that multigrid convergence severely degrades with inconsistent ETO coarse-grid discretizations. On regular simplicial high-aspect-ratio grids, analyses show that the Avg-LSQ (FT) coarse-grid discretization leads to convergence deterioration. On irregular simplicial high-aspect-ratio grids, convergence of multigrid with the Avg-LSQ (EN) coarse-grid discretization is also expected to deteriorate. Using other coarse-grid discretizations with simplicial elements in highly-stretched regions may be possible but is not straightforward and requires further study.

Acknowledgments

The results were computed with the FUN3D suite of codes (<http://fun3d.larc.nasa.gov/>) at NASA Langley Research Center.

References

- [1] Barth, T. J., "Numerical aspects of computing high-Reynolds number flow on unstructured meshes," AIAA paper 91-0721, January 1991.
- [2] Diskin, B., Thomas, J. L., and Mineck, R., "On quantitative analysis methods for multigrid solutions," *SIAM J Sci Comput*, Vol. 27, No. 1, 2005, pp. 108-129.
- [3] Diskin, B. and Thomas, J. L., "Accuracy analysis for mixed-element finite-volume discretization schemes," NIA Technical Report 2007-8, August 2007.
- [4] Diskin, B., Thomas, J. L., Nishikawa H, Nielsen, E. N., and White, J. A., "Comparison of node-centered and cell-centered unstructured finite-volume discretizations part I viscous fluxes," AIAA paper 2009-597, January 2009; *AIAA J* accepted for publication.
- [5] Diskin, B. and Thomas, J. L., "Notes on accuracy of finite-volume discretization schemes on irregular grids," *Appl Numer Math* (doi:10.1016/j.apnum.2009.12.001).
- [6] Brandt, A., "Multi-level adaptive solutions to boundary-value problems," *Math Comput*, Vol. 31, 1977, pp. 333-390.
- [7] Haselbacher, A. C., *A grid-transparent numerical method for compressible viscous flow on mixed unstructured meshes*, PhD thesis, Loughborough University, 1999.
- [8] Hyams, D. G., Sreenivas, K., and Whitfield, D. L., "Parallel fas multigrid for arbitrary Mach number, high Reynolds number unstructured flow solvers," AIAA paper 2006-2821, June 2006.
- [9] Mavriplis, D. J., "Multigrid techniques for unstructured meshes," VKI Lecture Series VKI-LS 1995-02, Von Karman Institute for Fluid Dynamics, Rhode-Saint-Genese, Belgium 1995.
- [10] Mavriplis, D. J., "Unstructured grid techniques," *Ann Rev Fluid Mech*, Vol. 29, 1997, pp. 473-514.
- [11] Mavriplis, D. J. and Pirzadeh, S., "Large-scale parallel unstructured mesh computations for 3d high-Lift analysis," *J Aircraft*, Vol. 36, No. 6, 1999, pp. 987-998.
- [12] Mavriplis, D. J., "An assessment of linear versus non-linear multigrid methods for unstructured mesh solvers," *J Comput Phys*, Vol. 175, 2002, pp. 302-325.
- [13] Nishikawa, H., Diskin, B., and Thomas, J. L., "A critical study of agglomerated multigrid methods for diffusion," AIAA paper 2009-4138, June 2009; *AIAA J* (doi:10.2514/1.j050055).

- [14] Svård, M. and Nordström, J., “An accuracy evaluation of unstructured node-centered finite-volume methods,” *Appl Numer Math*, Vol. 58, No. 8, 2008, pp. 1142-1158; also available as NIA Report 2005-04, NASA CR-2006-214293, April 2006.
- [15] Thomas, J. L., Diskin, B., and Brandt, A., “Textbook multigrid efficiency for the incompressible Navier-Stokes equations: high Reynolds number wakes and boundary layers,” *Comput Fluids*, Vol. 30, 2001, pp. 853–874.
- [16] Thomas, J. L., Diskin, B., and Brandt, A., “Textbook multigrid efficiency for fluid simulations,” *Ann Rev Fluid Mech*, Vol. 35, 2003, pp. 317-340.
- [17] Thomas, J. L., Diskin, B., and Rumsey, C. L., “Towards verification of unstructured grid methods,” *AIAA J*, Vol. 46, No. 12, December 2008, pp. 3070–3079 (also AIAA paper 2008-0666).
- [18] Trottenberg U., Oosterlee, C. W., and Schüller, A., *Multigrid*, Academic Press, London, 2000.
- [19] Venkatakrishnan, V., “Perspective on unstructured grid solvers,” *AIAA J*, Vol. 34, No. 3, 1996, pp. 533-545.

A. Agglomerated grid details

Table 4 lists grid sizes and numbers of grids agglomerated for the target grid sequences generated to assess multigrid convergence. For the stretched grids, the number of nodes in each

Isotropic Grids	Stretched Grids
09x09x09 (2)	09x09x33 (2,26)
13x13x13 (3)	13x13x49 (3,39)
17x17x17 (4)	17x17x66 (4,53)
21x21x21 (4)	21x21x83 (4,67)
25x25x25 (5)	25x25x100 (5,81)
29x29x29 (5)	29x29x117 (5,95)
33x33x33 (6)	33x33x134 (6,109)
37x37x37 (6)	37x37x150 (6,122)
41x41x41 (6)	41x41x167 (6,136)
45x45x45 (6)	45x45x184 (6,150)
49x49x49 (7)	49x49x201 (7,164)
53x53x53 (7)	53x53x218 (7,178)
57x57x57 (7)	57x57x235 (7,192)
61x61x61 (7)	61x61x251 (7,205)
65x65x65 (7)	65x65x268 (7,219)

Table 4: Grid sizes for isotropic and stretched grids; the first number in parenthesis is the numbers of agglomerated grids; the second number in parentheses is the number of nodes per implicit line.

implicit line is also listed. Tables 5-6 show the maximum skew angle and the coarsening ratio of each agglomeration level for two typical grids. The coarsening ratio is defined as the number of finer-grid degrees of freedom divided by the number of degrees-of-freedom at a given coarse level, ideally approaching 8 for full-coarsening in 3D. The coarsening ratio is above 6 on the first agglomeration but degrades on coarser levels. Note, for reference, that the isotropic tetrahedral meshes have a maximum skew angle of approximately 75 degrees and that faces with skew angles greater than $\pi/2$ are encountered on the fourth level for the isotropic grid and on the fifth level for the stretched grid.

B. Local Fourier analysis for regular grids

Asymptotic convergence rates of 2-grid cycles are predicted using LFA on regular 2D triangular and quadrilateral grids. De-

Level	Maximum Skew	Coarsening
	Angle (deg)	Ratio
2	79.8	6.3
3	81.9	5.5
4	96.8	4.2
5	88.3	3.0
6	89.8	2.2

Table 5: Maximum skew angle (deg) and coarsening ratio of each agglomeration level for the 37x37x37 isotropic grid.

Level	Maximum Skew	Coarsening
	Angle (deg)	Ratio
2	72.2	6.6
3	78.1	5.8
4	78.7	4.6
5	91.9	3.6
6	89.2	2.8

Table 6: Maximum skew angle (deg) and coarsening ratio of each agglomeration level for the 37x37x150 stretched grid.

tails pertaining to the analysis are given below. Foundations and applications of LFA can be found in the original paper [6] and in textbooks, e.g. [18]. The Green-Gauss discretization scheme is used on the fine grids; as noted earlier, for these fine grids, the scheme is the five-point Laplacian operator. Interior control-volume boundaries on a regular triangular fine grid are illustrated in Figure 8. The coarse-grids schemes are applied on fully-coarsened agglomerated coarse grids (interior coarse-grid control volumes corresponding to Figure 8 are illustrated in Figure 9).

The Fourier symbol of a 2-grid cycle, \hat{M} , is a 4×4 matrix acting in the linear vector space corresponding to the amplitudes of the following quartet of Fourier components,

$$\begin{aligned} e^{i(\theta_x i_x + \theta_y i_y)}, & \quad e^{i((\theta_x + \pi) i_x + \theta_y i_y)}, \\ e^{i(\theta_x i_x + (\theta_y + \pi) i_y)}, & \quad e^{i((\theta_x + \pi) i_x + (\theta_y + \pi) i_y)}, \end{aligned} \quad (\text{B.1})$$

with horizontal and vertical node indexes, i_x and i_y , respectively, and normalized Fourier frequencies

$$\begin{aligned} \bar{\theta}^1 &= (\theta_x^1, \theta_y^1) = (\theta_x, \theta_y), \\ \bar{\theta}^2 &= (\theta_x^2, \theta_y^2) = (\theta_x + \pi, \theta_y), \\ \bar{\theta}^3 &= (\theta_x^3, \theta_y^3) = (\theta_x, \theta_y + \pi), \\ \bar{\theta}^4 &= (\theta_x^4, \theta_y^4) = (\theta_x + \pi, \theta_y + \pi) \end{aligned} \quad (\text{B.2})$$

satisfying $\max(|\theta_x|, |\theta_y|) \leq \pi/2$.

$$\hat{M} = \hat{S}^{\nu_2} \left(\hat{E} - \hat{P} \hat{L}_H^{-1} \hat{R} \hat{W} \hat{L}_h \right) \hat{S}^{\nu_1}. \quad (\text{B.3})$$

Here, $\nu_1 = \nu_2 = 2$ are the numbers of pre- and post-relaxation sweeps, \hat{S} is the relaxation operator symbol, \hat{L}_h and \hat{L}_H are the fine- and coarse-grid operator symbols, \hat{W} is the residual-averaging operator symbol, \hat{P} and \hat{R} are the symbols of the prolongation and restriction operators corresponding to P_0 and R_0 , respectively, and \hat{E} is the 4×4 identity matrix.

The symbols \hat{L}_h and \hat{W} are 4×4 diagonal matrices and the symbols \hat{R} and \hat{P} are 1×4 and 4×1 vectors, respectively, each composed of scalar Fourier symbols. The scalar symbols are computed for each of the components (B.1). The diagonal entries of the fine-grid operator symbol, \hat{L}_h , are

$$\hat{L}_h^k = \frac{2}{h_y^2} [(-1 + \cos(\theta_y^k)) + \frac{1}{A^2} (-1 + \cos(\theta_x^k))], \quad (\text{B.4})$$

where h_y and $h_x = Ah_y$ are fine-grid mesh spacings in the corresponding directions and A is the grid aspect ratio. The symbols \hat{R} and \hat{P} relate the amplitudes of the four fine-grid Fourier components (B.1) to the amplitude of the corresponding coarse-grid Fourier component $e^{i2(\theta_x i_x + \theta_y i_y)}$ and assume that the coarse grid node (i_x, i_y) is located at the center of the rectangle formed by the four fine-grid nodes $(2i_x, 2i_y)$, $(2i_x + 1, 2i_y)$, $(2i_x, 2i_y + 1)$, and $(2i_x + 1, 2i_y + 1)$. The entries of \hat{R} are

$$\hat{R}^k = \frac{1}{4} \left(1 + e^{i\theta_x^k} + e^{i\theta_y^k} + e^{i(\theta_x^k + \theta_y^k)} \right). \quad (\text{B.5})$$

The entries of \hat{P} are

$$\hat{P}^k = \frac{1}{4} \left(1 + e^{-i\theta_x^k} + e^{-i\theta_y^k} + e^{-i(\theta_x^k + \theta_y^k)} \right). \quad (\text{B.6})$$

The entries of \hat{W} are shown below for triangular and quadrilateral grids,

$$(\hat{W}^k)_{\text{tria.}} = \frac{1}{3} (\cos(\theta_x^k) + \cos(\theta_y^k) + \cos(\theta_x^k + \theta_y^k)), \quad (\text{B.7})$$

$$(\hat{W}^k)_{\text{quad.}} = \frac{1}{2} (\cos(\theta_x^k) + \cos(\theta_y^k)). \quad (\text{B.8})$$

The symbols of relaxations performed in the lexicographic order are 4×4 diagonal matrices composed of scalar Fourier symbols. Table 7 shows the main-diagonal symbols for lexicographic-order idealized and actual relaxations.

Relaxation	Fine Grid	$\hat{S}^k \equiv N/D$
IR-P	Quad.	$N = e^{i\theta_x^k} + e^{i\theta_y^k}$ $D = 4 - e^{-i\theta_x^k} - e^{-i\theta_y^k}$
IR-P	Tria.	$N = e^{i\theta_x^k} + e^{i\theta_y^k} + e^{i(\theta_x^k + \theta_y^k)}$ $D = 6 - e^{-i\theta_x^k} - e^{-i\theta_y^k} - e^{-i(\theta_x^k + \theta_y^k)}$
IR-L	Quad.	$N = e^{i\theta_x^k}$ $D = 4 - e^{-i\theta_x^k} - 2\cos(\theta_y^k)$
IR-L	Tria.	$N = e^{i\theta_x^k} + e^{i(\theta_x^k + \theta_y^k)}$ $D = 6 - e^{-i\theta_x^k} - 2\cos(\theta_y^k) - e^{-i(\theta_x^k + \theta_y^k)}$
Actual	Either	$N = e^{i\theta_x^k}$ $D = 2 + 2A^2 - e^{-i\theta_x^k} - 2A^2 \cos(\theta_y^k)$

Table 7: Main-diagonal symbols of lexicographic relaxations; actual is line-implicit relaxation.

Multicolor relaxations depend on the specific relaxation order and their symbols are 4×4 matrices with a more complex structure. For example, the symbol of a 2-color line relaxation

has a block-diagonal structure with two 2×2 diagonal blocks; the block corresponding to the frequencies $\bar{\theta}^1$ and $\bar{\theta}^2$ is defined as

$$\frac{1}{2} \begin{bmatrix} \hat{D}^1(1 + \hat{D}^1) & \hat{D}^2(1 - \hat{D}^2) \\ \hat{D}^1(1 - \hat{D}^1) & \hat{D}^2(1 + \hat{D}^2) \end{bmatrix}, \quad (\text{B.9})$$

where the scalar symbol \hat{D}^k , corresponding to line-implicit Jacobi relaxation, is given in Table 8 for the operators and grids considered.

Operator	Fine Grid	$\hat{D}^k \equiv N/D$
IR	Tria.	$N = \cos(\theta_x^k) + \cos(\theta_x^k + \theta_y^k)$ $D = 3 - \cos(\theta_y^k)$
IR	Quad.	$N = \cos(\theta_x^k)$ $D = 2 - \cos(\theta_y^k)$
Actual	Either	$N = \cos(\theta_x^k)$ $D = 1 + A^2 - A^2 \cos(\theta_y^k)$

Table 8: Symbols of implicit-line Jacobi relaxation.

To describe a 4-color relaxation, let color 1 mark points with i_x even and i_y even, color 2 mark points with i_x odd and i_y even, color 3 mark points with i_x even and i_y odd, and color 4 mark points with i_x odd and i_y odd. First, the point-amplification symbols, C_j^i , for each color are computed where subscripts and superscripts denote color and frequency, respectively. Table 9 collects point-amplification symbols for two multicolor IR-P schemes performed in the (1234) order. The relaxation symbol is the following matrix

$$\hat{S}(\bar{\theta}) = \frac{1}{4} \begin{bmatrix} C_1^1 + C_2^1 + C_3^1 + C_4^1 & C_1^2 - C_2^2 + C_3^2 - C_4^2 \\ C_1^1 - C_2^1 + C_3^1 - C_4^1 & C_1^2 + C_2^2 + C_3^2 + C_4^2 \\ C_1^1 + C_2^1 - C_3^1 - C_4^1 & C_1^2 - C_2^2 - C_3^2 + C_4^2 \\ C_1^1 - C_2^1 - C_3^1 + C_4^1 & C_1^2 + C_2^2 - C_3^2 - C_4^2 \\ C_1^3 + C_2^3 - C_3^3 - C_4^3 & C_1^4 - C_2^4 - C_3^4 + C_4^4 \\ C_1^3 - C_2^3 + C_3^3 + C_4^3 & C_1^4 + C_2^4 - C_3^4 - C_4^4 \\ C_1^3 + C_2^3 + C_3^3 + C_4^3 & C_1^4 - C_2^4 + C_3^4 - C_4^4 \\ C_1^3 - C_2^3 + C_3^3 - C_4^3 & C_1^4 + C_2^4 + C_3^4 + C_4^4 \end{bmatrix}. \quad (\text{B.10})$$

Fine Grid	Symbol
Quad.	$C_1^k = \frac{1}{2} (\cos(\theta_x^k) + \cos(\theta_y^k))$ $C_2^k = \frac{1}{2} (C_1^k \cos(\theta_x^k) + \cos(\theta_y^k))$ $C_3^k = \frac{1}{2} (\cos(\theta_x^k) + C_1^k \cos(\theta_y^k))$ $C_4^k = \frac{1}{2} (C_3^k \cos(\theta_x^k) + C_2^k \cos(\theta_y^k))$
Tria.	$C_1^k = \frac{1}{3} (\cos(\theta_x^k) + \cos(\theta_y^k) + \cos(\theta_x^k + \theta_y^k))$ $C_2^k = \frac{1}{3} (C_1^k \cos(\theta_x^k) + \cos(\theta_y^k) + \cos(\theta_x^k + \theta_y^k))$ $C_3^k = \frac{1}{3} (\cos(\theta_x^k) + C_1^k \cos(\theta_y^k) + C_2^k \cos(\theta_x^k + \theta_y^k))$ $C_4^k = \frac{1}{3} (C_3^k \cos(\theta_x^k) + C_2^k \cos(\theta_y^k) + C_1^k \cos(\theta_x^k + \theta_y^k))$

Table 9: Point-amplification symbols in 4-color (1234) IR-P schemes.

The coarse-grid operator symbol, \hat{L}_H , is a scalar function of the coarse-grid frequency, $(\theta_x^H, \theta_y^H) \equiv (2\theta_x, 2\theta_y)$, specific to the given coarse-grid discretization. With a quadrilateral fine

Scheme	$\hat{L}_H(2h_y^2)$
ETO (EN)	$\frac{5}{6}(-1 + \cos(\theta_y^H)) + \frac{5}{6}\frac{1}{A^2}(-1 + \cos(\theta_x^H))$ $+ \frac{1}{3}\frac{1}{A^2+1}(-1 + \cos(\theta_x^H + \theta_y^H))$
Avg-LSQ (EN)	$[\hat{L}_H]^{\text{ETO (EN)}}(2h_y^2)$ $+ \frac{1}{36A^2} \left[(-\sin(\theta_y^H) + 2\sin(\theta_x^H) + \sin(\theta_x^H + \theta_y^H)) \right.$ $\left. \left(\sin(\theta_y^H) + \frac{A^2-1}{A^2+1} \sin(\theta_x^H + \theta_y^H) \right) \right]$ $+ \frac{1}{36} \left[(2\sin(\theta_y^H) - \sin(\theta_x^H) + \sin(\theta_x^H + \theta_y^H)) \right.$ $\left. \left(\sin(\theta_x^H) - \frac{A^2-1}{A^2+1} \sin(\theta_x^H + \theta_y^H) \right) \right]$
ETO (FT)	$\left(\frac{5}{6} + \frac{1}{30}\frac{1}{A^2}\right)(-1 + \cos(\theta_y^H))$ $+ \left(\frac{5}{6}\frac{1}{A^2} + \frac{1}{30}\right)(-1 + \cos(\theta_x^H))$ $+ \frac{1}{12}\left(1 + \frac{1}{A^2}\right)(-1 + \cos(\theta_x^H + \theta_y^H))$
Avg-LSQ (FT)	$[\hat{L}_H]^{\text{ETO (FT)}}(2h_y^2)$ $+ \frac{1}{36A^2} \left[(-\sin(\theta_y^H) + 2\sin(\theta_x^H) + \sin(\theta_x^H + \theta_y^H)) \right.$ $\left. \left(\sin(\theta_y^H) + \frac{A^2}{5} \sin(\theta_x^H) + \frac{A^2-1}{2} \sin(\theta_x^H + \theta_y^H) \right) \right]$ $+ \frac{1}{36} \left[(2\sin(\theta_y^H) - \sin(\theta_x^H) + \sin(\theta_x^H + \theta_y^H)) \right.$ $\left. \left(\sin(\theta_x^H) + \frac{A^2}{5} \sin(\theta_y^H) - \frac{A^2-1}{2A^2} \sin(\theta_x^H + \theta_y^H) \right) \right]$

Table 10: Symbol of coarse-grid operators for triangular fine grids.

	$A = 1$	$\lim A \rightarrow \infty$
Scheme	$-h_y^2 \hat{L}_H^*$	$-h_y^2 \hat{L}_H^*$
Avg-LSQ	$\theta_x^2 + \theta_y^2$	θ_y^2
ETO (EN)	$\theta_x^2 + \theta_y^2 + \frac{1}{3}\theta_x\theta_y$	$\frac{5}{6}\theta_y^2$
ETO (FT)	$\frac{31}{30}\theta_x^2 + \frac{31}{30}\theta_y^2 + \frac{1}{3}\theta_x\theta_y$	$\frac{11}{12}\theta_y^2 + \frac{1}{60}\theta_x^2 + \frac{1}{6}\theta_x\theta_y$

Table 11: Expansion of coarse-grid discretization operators on a triangular fine grid.

grid, both coarse and fine grids are orthogonal and \hat{L}_H is defined as,

$$\hat{L}_H = \frac{1}{2h_y^2} \left[-1 + \cos(\theta_y^H) + \frac{1}{A^2}(-1 + \cos(\theta_x^H)) \right]. \quad (\text{B.11})$$

The operator \hat{L}_H^* is composed of the leading-order terms in an expansion of \hat{L}_H , assuming small (θ_x, θ_y) . For B.11,

$$\hat{L}_H^* = \frac{-1}{h_y^2} \left(\theta_y^2 + \frac{1}{A^2}\theta_x^2 \right) \quad (\text{B.12})$$

coincides with the differential operator applied to the Fourier component $e^{i(\theta_x x/h_x + \theta_y y/h_y)}$, thus demonstrating that the coarse-grid operator for quadrilateral fine grids is consistent. Table 10 collects the symbols of the coarse-grid operator corresponding to various coarse-grid discretizations on triangular fine grids. The symbols for the Avg-LSQ and ETO discretizations are shown for both EN and FT augmentations. Table 11 collects the corresponding expansions for $A = 1$ and in the limit of $A \rightarrow \infty$. With quadrilateral grids or with Avg-LSQ discretizations, the coarse-grid operators are consistent for all A . On triangular grids, both the ETO (EN) and ETO (FT) discretizations are inconsistent for all A .

The asymptotic convergence rates are computed as the maximum spectral radius of \hat{M} over all possible Fourier frequencies. Since the maximum amplification on high-aspect-ratio grids is expected for frequencies extremely smooth in the y -direction

($|\theta_y| \approx 0$), the frequency domain $(\theta_x, \theta_y) \in [-\pi, \pi]^2$ is, first, searched with the increment 0.03 in both frequencies. Then, the band $|\theta_y| \leq \frac{\pi}{A}, |\theta_x| \leq \pi$ is searched again with the θ_y -increment reduced to $0.03/A$; the θ_x increment is kept as 0.03.

As a remark on the multigrid results tabulated in Section 5.2, an inconsistent scheme does not necessarily lead to poor multigrid performance. Inconsistency does imply that the coarse-grid correction for the smoothest components is not precise. For example, LFA analyses show that multigrid convergence on isotropic ($A = 1$) triangular grids with coarse grids discretized with either of the two inconsistent ETO schemes is similar to multigrid convergence on isotropic quadrilateral grids. For high-aspect-ratio triangular grids, Table 11 indicates that, with the ETO (EN) scheme, the low-frequency coarse-grid correction is $5/6$ of the optimal correction and the overall multigrid cycle is 0.2 per cycle. For high-aspect-ratio triangular grids with the ETO (FT) scheme, the coarse-grid correction for intermediate frequencies in x and low frequencies in the y direction is inadequate, leading to poor multigrid convergence. The cause of the slowdown is the increase of the stencil weights in the x -direction associated with skew angles approaching $\pi/2$. The same difficulty occurs for the Avg-LSQ (FT) scheme, even though it is a consistent scheme.

C. Idealized relaxation on high-aspect-ratio grids

The effects of various idealized and actual relaxation schemes on multigrid convergence are shown below for one coarse-grid discretization — the Avg-LSQ (EN) scheme. Regular triangular and quadrilateral grids are considered, following the groundwork in Appendix B.

Relaxation	Order	Quadrilateral Fine Grid	Triangular Fine Grid
IR-P	4-Color (1234)	46	152
IR-P	Lexicographic	0.1	0.07
IR-L	2-Color	0.11	0.12
IR-L	Lexicographic	0.06	0.07
Actual	2-Color	0.11	0.07
Actual	Lexicographic	0.02	0.07

Table 12: LFA convergence rates per cycle for triangular and quadrilateral grids with residual averaging; actual is line-implicit relaxation.

Table 12 shows convergence rates of 2-grid cycles computed with LFA for quadrilateral and triangular fine grids with $A = 10^4$. The results are shown with residual averaging although the conclusions are not sensitive to its inclusion. Four ideal relaxations, IR-P and IR-L with multicolor and lexicographic ordering, and two actual relaxations, line-implicit with 2-color and lexicographic ordering, are considered. The actual line-implicit relaxations are less than 0.12 per cycle for both triangular and quadrilateral grids.

The 4-color (1234) IR-P cycle is unstable and thus not suitable as a predictor of the actual cycle. Although not shown, other color sequences give similar results. Convergence of the

lexicographic IR-P cycle is better than 0.1 per cycle and thus lexicographic IR-P could be considered as a possible idealized relaxation. However, the IR-L cycles are uniformly-excellent quantitative predictors when the idealized relaxation is applied in the same order as the actual line-implicit relaxation. Convergence of the 2-color IR-L cycle predicts convergence of the actual cycle with 2-color line-implicit relaxation. Likewise, convergence of the lexicographic IR-L cycle predicts convergence of the actual cycle with lexicographic line-implicit relaxation. The IR-L cycle is a simple, consistent, and accurate predictor of the convergence rates of the actual cycle, and we use it for the analyses of multigrid solutions on high-aspect-ratio grids reported in sections 5.2.

Fine Grid	Convergence Rate	
	$A = 1$	$A = 10^4$
32^2	0.109	0.109
128^2	0.110	0.193
512^2	0.110	0.626
2048^2	0.110	2.369
∞	0.110	46.

Table 13: Convergence rates of 4-color (1234) IR-P(2,2) cycles as a function of grid size and aspect ratio for periodic domains with residual averaging; quadrilateral fine grid.

Note that the instability of the IR-P cycle occurs for error components that are extremely smooth in the y direction. It is difficult to observe this instability in actual computations because, to realize such smooth components, a large number of high-aspect-ratio cells in the y -direction is required. Table 13 illustrates this behavior, showing convergence rates computed with LFA and confirmed in actual computations on uniform high-aspect ratio grids in a periodic domain. The convergence rates are for 4-color (1234) IR-P(2,2) cycles with residual averaging on quadrilateral grids with $A = 1$ and $A = 10^4$; only Fourier components realizable on the specified grids have been considered. Grid-independent convergence is shown on isotropic grids ($A = 1$) but the instabilities on anisotropic grids ($A = 10^4$) have not reached their asymptotic value (from Table 12) for the entries in the table corresponding to 2048^2 points.

D. Idealized relaxation on isotropic grids

Here we show a somewhat subtle effect that arises in unstructured grids with IR based on edge-connections. The role effected by IR depends on the number of edges N_e in (20). For a hexahedral mesh, the number of simply-connected edges is 6 but the total number of simply-connected and virtual edges is 26, corresponding to 7-point and 27-point stencils of A^{IR} , respectively. The convergence of 4-color IR-P(2,2) for a 64^3 isotropic hexahedral grid over a spherical domain is shown in Figure 16 for these two stencils. With the 27-point stencil, the asymptotic convergence of IR-P is noticeably faster than that with the 7-point stencil. Although not shown, even with single

grid (no multigrid) iterations, relaxation of (20) with the 27-point stencil converges in half of the iterations as that with the 7-point stencil.

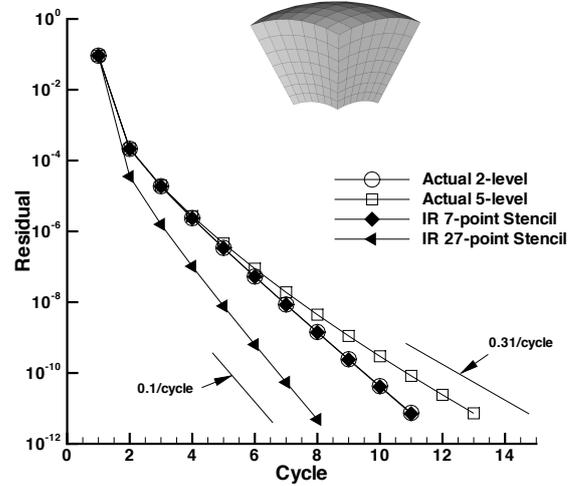


Figure 16: Convergence rates of actual and IR cycles for 64^3 hexahedral grid using point-wise relaxation; inset is a coarser $8 \times 8 \times 8$ grid.

On this particular grid, the actual discrete diffusion operator is much closer to the 7-point operator. As seen in Figure 16, convergence of actual 2-level V(2,2) cycles is quite close to that of IR with the 7-point stencil. Convergence of actual cycles with 5 levels is somewhat slower asymptotically than the 2-level convergence. The interpretation is that IR with the 27-point stencil is providing faster convergence of the medium frequencies than point relaxation of the actual diffusion operator. Using additional relaxation provides convergence rates per cycle that agree closely to that with the 27-point stencil but does not provide an overall gain in efficiency.