INTERNATIONAL JOURNAL FOR NUMERICAL METHODS IN FLUIDS
Int. J. Numer. Meth. Fluids 2002; **40**:361–368 (DOI: 10.1002/fld.292)

Structured multigrid agglomeration on a data structure for unstructured meshes

V. Hannemann*

Institute of Aerodynamics and Flow Technology, German Aerospace Center, DLR, Bunsenstr. 10, 37073 Göttingen, Germany

SUMMARY

A key issue for using multigrid as an acceleration technique on unstructured meshes is the generation of the coarse grid levels. Some new ideas are presented to improve the agglomeration of coarse cells in hybrid grids with respect to the structure of hexahedral or prismatic layers. The algorithms are embedded in the standard advancing front method which works locally on the edge based data structure representing the dual grid. The goal is to preserve the topological structure of the hexahedral and prismatic parts as far as possible. As a result of the structure preservation a desired coarsening ratio (number of fine grid volumes per coarse grid volume), responsible for the memory requirements and the computing time spent in looping over all volumes, can be achieved with a minimal number of edges, which is linked to the major working time spent in loops over all edges. Copyright © 2002 John Wiley & Sons, Ltd.

KEY WORDS: agglomeration multigrid; hybrid grids; structure preservation

1. INTRODUCTION

Most adaptive methods in CFD use unstructured meshes to provide local refinement strategies. Depending on the scales to be resolved in different directions, these meshes may consist of hexahedral, prismatic, pyramidal and tetrahedral elements. Tetrahedra are good for isotropic resolution, while hexahedra and prisms are more suited for boundary layers where the wall normal scale has to be discretized more finely than the tangential scales. Although the boundary layers cover only a small part of the computational domain they need usually more than 80% of all cells to be resolved. Therefore, the hexahedral and prismatic parts of a hybrid grid play the major role in discretizing viscous flow fields.

To reduce memory requirements, finite volume methods working on dual grids as control volumes are popularly used. The primary grid is transformed into an edge based data structure, neglecting all information which is not necessary for the equation solver. Thus the volumes and the node connections with attached surface vectors are stored, but the primary grid cells do not exist in the dual grid.

*Correspondence to: V. Hannemann, Institute of Aerodynamics and Flow Technology, German Aerospace Center, DLR, Bunsenstr. 10, 37073 Göttingen, Germany.

Received May 2001
Revised September 2001

Multigrid is one of the best convergence acceleration techniques. Lallemand *et al.* [1] introduced an agglomeration procedure to create the coarse grid cells by fusing cells from the next finer level. Several improvements have been made to the isotropic agglomeration of tetrahedral grids in order to optimize the fused cells according to a surface to volume ratio and to the coarsening ratio between different grid levels [2–5]. All these approaches work on the dual grid data structure to enable the generation of multiple coarser grids by the same algorithm. The additional knowledge about the cell types of the primary grid cannot be used, because the agglomerated volumes can in general not be mapped on a valid primary grid.

Applying these techniques to prismatic or hexahedral grid parts (normally associated with high aspect ratios) does not provide optimal results [5]. Therefore, directional or semi-coarsening algorithms are used [3, 5]. The search for cells connected via the largest face is repeated to achieve better coarsening ratios than 2. The border between semicoarsening and isotropic agglomeration in a hybrid grid is determined by cell aspect ratios, or by the number of minimal edges distance to a specified boundary. Depending on the initial discretization, all approaches influence the topological structure of prismatic or hexahedral layers on the coarser meshes.

This paper presents topological arguments to differentiate between tetrahedral, prismatic and hexahedral connectivity parts on the dual grids. The standard advancing front algorithm for an isotropic agglomeration [1, 2, 5] is modified with the aim at preserving the topological structure of the dual grids on the coarser levels as far as possible.

An investigation of the multigrid performance according to different agglomerated grids is very much dependent on the multigrid algorithms and parameters used in the equation solver and therefore beyond the scope of this paper.

2. TOPOLOGICAL ASPECTS

To topologically characterize the different dual grid parts arising from tetrahedral, prismatic or hexahedral parts of the primary grid, the following definitions are useful:

- Two control volumes sharing a common face are called *direct neighbours* of each other.
- A control volume V_1 is called an *indirect neighbour* of a control volume V_2 if they are not direct neighbours but have a direct neighbour in common.
- All control volumes sharing at least one point with a control volume V_1 are called the *neighbourhood* N_{V_1} of V_1.
- A *single-connected neighbour* is a direct neighbour of a control volume V_1, which does not share any direct neighbours with V_1.

Based on these definitions the parts of the dual grid can be described as:

- *Tetrahedral connectivity*: The connectivity of a control volume is called tetrahedral connectivity if no indirect neighbour is part of the neighbourhood. In this case the number of single-connected neighbours equals zero.
- *Prismatic connectivity*: A neighbourhood of a control volume with a prismatic connectivity consists only of some direct and some indirect neighbours. Away from the boundaries there are exactly two single-connected neighbours, depicted above and below the seed volume in Figure 1. If the prisms of the primary grid include boundary surface triangles

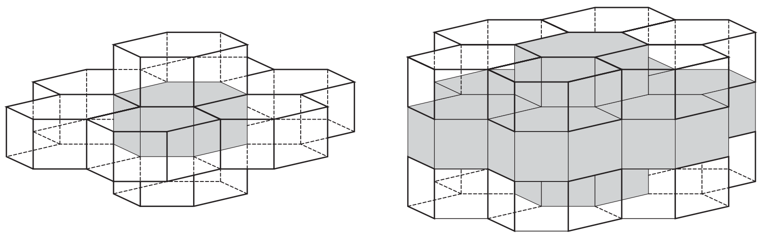

Figure 1. Prismatic neighbourhood: seed volume and direct neighbours (left), direct and indirect neighbours (right).

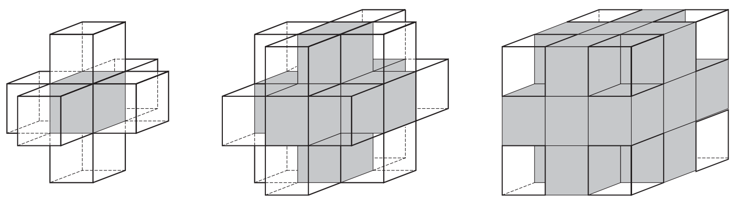

Figure 2. Hexahedral neighbourhood: direct neighbours sharing faces (left), indirect neighbours sharing edges (middle) and neighbours sharing points (right) with the seed volume.

then the number of single-connected neighbours reduces to one. In the usually-avoided case of prisms with boundary surface quadrilaterals, the number of single-connected neighbours can increase for some configurations, such as a sharp trailing edge of a wing.

- *Hexahedral connectivity*: A control volume with a hexahedral connectivity possesses a neighbourhood containing direct neighbours, indirect neighbours and indirect neighbours of direct neighbours. When connected via hexahedra in an $i-j-k$ structured part of the mesh a control volume has exactly six direct neighbours, all single-connected, as shown in Figure 2. The number of single-connected neighbours varies in cases of boundary contacts from at least three to more than six. If the control volume is located on the border between a hexahedral and, for example, a tetrahedral part of the grid (with some pyramids in between) the number of single-connected volumes can decrease to zero.

We define the *connectivity* of each control volume by the number of single-connected neighbours. Ignoring some special cases a tetrahedral connectivity is assumed if the number is zero and a prismatic connectivity is assumed if the number is one or two. A higher number is assumed to indicate a hexahedral connectivity.

3. AGGLOMERATION ALGORITHM

The algorithm consists of several parts:

- Determination of a fine grid volume, called the *seed volume*, around which the coarse grid cell should be agglomerated.

- Creation of a maximal set of fine grid volumes to be fused to the seed volume.
- Selection of a subset of these volumes, including the seed volume.
- Conditionally, creation of a new agglomerated cell or fusion of the subset with an already existing coarse grid cell.

The first item is responsible for the shape of the advancing front, while the next items influence the quality of the coarser grids.

3.1. Determination of the seed volume

The determination of the seed volume is modified to allow the advancing front growing in the hexahedral or prismatic parts layer by layer on top of the initial boundaries.

As in the standard methods the advancing front starts on special boundary parts, for example on walls, but not on the farfield. Within the set of volumes attached to a special boundary part, the preferred volumes are those located on an edge and optimally those which lie in a corner of the computational domain since the number of neighbours is reduced and the optimal fusing easy to determine. Together, this information forms the volume priority. It is calculated on the primary grid and handed down from each grid to the successively coarser level.

The selection of the next seed volume within each priority set is controlled by subsets, called queues. The most simple form consists of only two queues, one containing all volumes in the beginning of the agglomeration and one for all volumes which have been touch by the front during the agglomeration procedure. For a more regular growth inside the structured parts, the knowledge of how often a volume is touched by the front is very useful. For example in a hexahedral connectivity, a volume faces the corner situation mentioned above when it is hit three times by the advancing front. But due to the higher number of direct neighbours, the number of contacts with the agglomeration front increases more quickly in the prismatic than in the hexahedral parts and even more quickly in the tetrahedral parts of the grid. Therefore, the volumes are sorted initially into different queues according to their connectivity; hexahedral parts are preferred to prismatic parts of the mesh, and prismatic parts are preferred to tetrahedral parts. The higher the number of single-connected neighbours is for a given volume, the lower is the number of the queue, the volume is assigned to. The current implementation uses eight queues, with the lowest four only accessed by volumes lying at the advancing front. If four queues are insufficient to provide a unique queue for each different connectivity number, some connectivities are mapped to the same queue.

After each agglomeration step all agglomerated volumes are removed from the queues. Each time a volume V_i of a higher queue is touched by the front via agglomeration of one or more direct neighbours of V_i, V_i is added to the next lower queue. To copy the structure of the underlying layer it is useful to advance within each layer in a similar way. Therefore, further ordering is done inside the queues. Between several volumes neighbouring the same coarse grid volume, the neighbours of the seed volume are queued first, the neighbours of the direct neighbours of the seed volume are queued next and the neighbours of indirect neighbours of the seed volume are queued last.

The search for the next seed volume starts with the volumes of highest priority, and for equal priorities the queue with the lowest queue number is searched first. Within each queue, the volume which has first been touched by the front is elected to be the next seed volume.

3.2. Maximal set of fusible volumes

To preserve the structure of the grid it is necessary to allow the agglomeration to fuse within the neighbourhood N_{V_i} of the seed volume V_i. The standard method only allows the fusing of direct neighbours of V_i which is a subset of N_{V_i} and avoids this way the creation of coarse cells consisting of more than one separated volume. For a seed volume in a tetrahedral part of the grid both sets are equal. In a prismatic region indirect neighbours who have only an edge in common with the seed volume have to be added (see Figure 1). They can be detected as direct neighbours of the single-connected neighbours which have another direct neighbour other than the single-connected neighbour in common with the seed volume.

In hexahedral regions additional neighbours who share only one point with the seed volume have to be found. These are the direct neighbours of more than one indirect neighbour having a common edge with the seed volume. The hexahedral situation is depicted in Figure 2.

3.3. Selection of the fusible subset

A typical coarsening ratio in a structured multigrid code is 8 for three dimensional applications and 4 in the two-dimensional case. The maximal set of volumes for a hexahedral connectivity offers 27 fine grid cells, as shown in Figure 2. Therefore it is sufficient to select a subset of the maximal set to provide good coarser grids. The most obvious restriction to the fusible set is that volumes cannot be agglomerated to more than one seed volume. As the agglomeration front advances, the next seed volume is found near to already fused volumes. After reducing the size of a fusible subset, the connectivity of all non-direct neighbours is checked to ensure that the coarse volume is simple connected. Otherwise coarse cells can be generated with volume parts sharing not more than a common edge or point.

Besides the usual restrictions of the isotropic agglomeration like a surface-to-volume ratio criterion applied in tetrahedral parts, a special constraint is introduced for the structured parts:

Cutting the structure of an underlying layer: In a prismatic or hexahedral region of the grid, the algorithm tries to advance layer by layer over the initial front. A new coarse grid volume is only allowed to be on top of more than one coarse grid volume of the underlying layer if these already fused volumes are covered completely by the new volume. The earliest fused, single-connected neighbour of the seed volume indicates the underlying coarse volume. The allowed fusing situations, simplified in two dimensions, are depicted in the left part of Figure 3 and a typical cut situation in the right part.

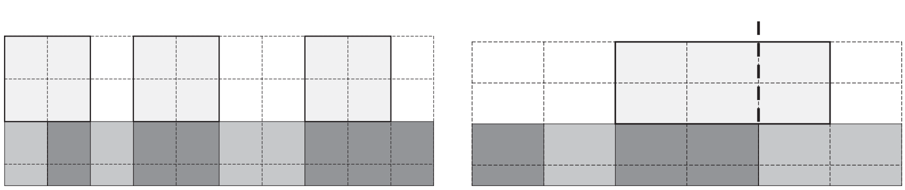

Figure 3. *Left diagram*: allowed fusing: new volume covers old volumes completely (left), new volume covers exactly one old volume (middle) and new volume lies inside the cover of one old volume. *Right diagram*: the parts of the new volume divided by the broken line are not allowed to be fused together due to the structure of the underlying layer.

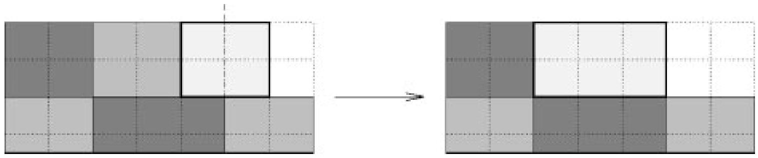

Figure 4. After cutting the new fusible set, some volumes are agglomerated to a previously fused cell.

3.4. Fusing to coarse grid volume

In addition to the cutting procedure described above, an agglomeration with an already fused neighbour is allowed in special situations to prohibit hardly fused volumes. One situation occurs when different parts of the agglomeration front reach each other and some seed volumes are left with no or only one fusible neighbour. Another case, shown in Figure 4, is a result of a cutting to preserve the structure of an underlying layer.

4. EXAMPLES

Structured two-dimensional C-grid: The isotropic algorithm produces the coarser grids depicted in the upper row of Figure 5. Using the modifications discussed previously to preserve the structure, the grids (shown in the lower row of Figure 5) provide a similar quality to grids generated with the knowledge of the $i-j-k$ structure of the primary grid. The coarsening ratio of the volumes increases from 2.1 (coarser levels 3.4, 3.4) to 4.0 $(4.0, 4.0)$ and the edge ratio from 1.5 $(2.7, 3.0)$ to 4.0 $(3.9, 3.9)$.

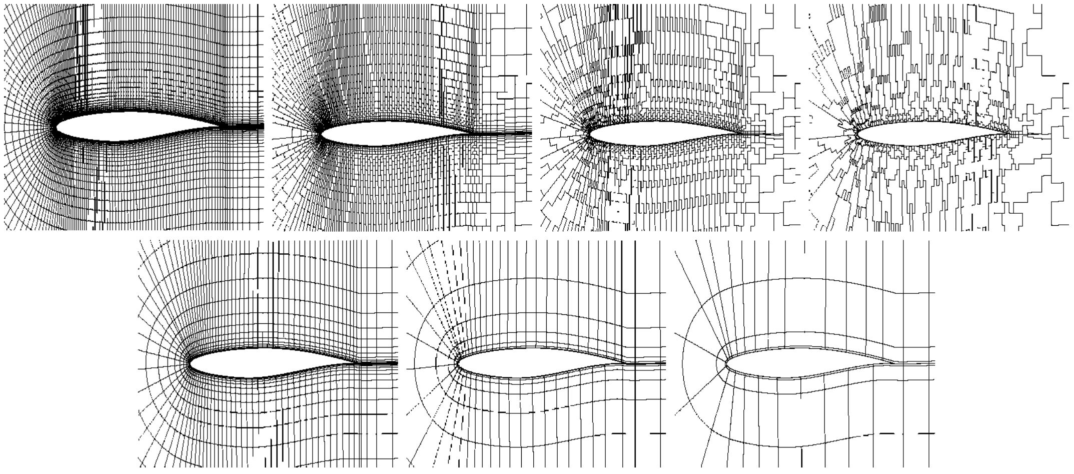

Figure 5. *Top*: Dual grid of the primary grid and coarse grids agglomerated with the isotropic algorithm. *bottom*: Dual grids agglomerated with structure preservation.

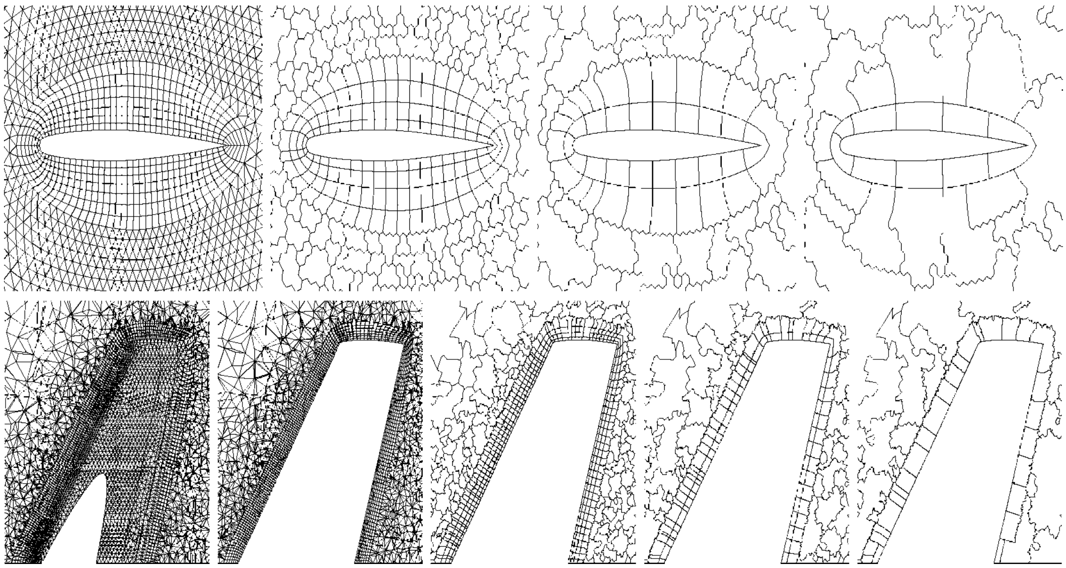

Figure 6. Hybrid (tetrahedra/prisms) mesh around a 3d wing: *top*: primary and agglomerated dual grids in the symmetry plane, *bottom*: primary grid in two cut planes, coarser grids in the second cut plane.

Three dimensional wing: A prismatic sublayer around the wing forms the structured part of the grid, and a tetrahedral discretization covers the rest of the computational domain. As shown in Figure 6, the structure on the coarser levels is preserved as long as the structured part is thick enough to keep the agglomeration inside the prismatic layer.

Multi element airfoil: A two-dimensional experimental discretization around a three element airfoil is depicted in Figure 7. The generation of the primary grid starts with three independent O-grids around each element consisting of 24 layers of hexahedra. These grids are locally reduced to prohibit any overlap and to prevent an aspect ratio (tangential to normal distance) smaller than 1 at the outer boundary. The gaps and the rest of the flowfield are subsequently filled with prisms. The structure of the hybrid grid is still visible on the coarsest grid.

5. CONCLUSIONS AND PERSPECTIVES

Topological arguments enable the detection of the underlying primary grid structure in the reduced information of the dual grid. They are used to protect the topological structure during the agglomeration process. The necessary modifications of the standard advancing front algorithm are discussed in detail. The presented results show the ability of the modified agglomeration to protect the structure of simply structured grid parts (O- or C-meshes with one front direction) in two as well as three space dimensions and with hybrid meshes. The coarsening ratios of the volumes and the reduction of edges in the structured part becomes equal to values achieved on structured grids. Therefore, the computing time needed per multigrid

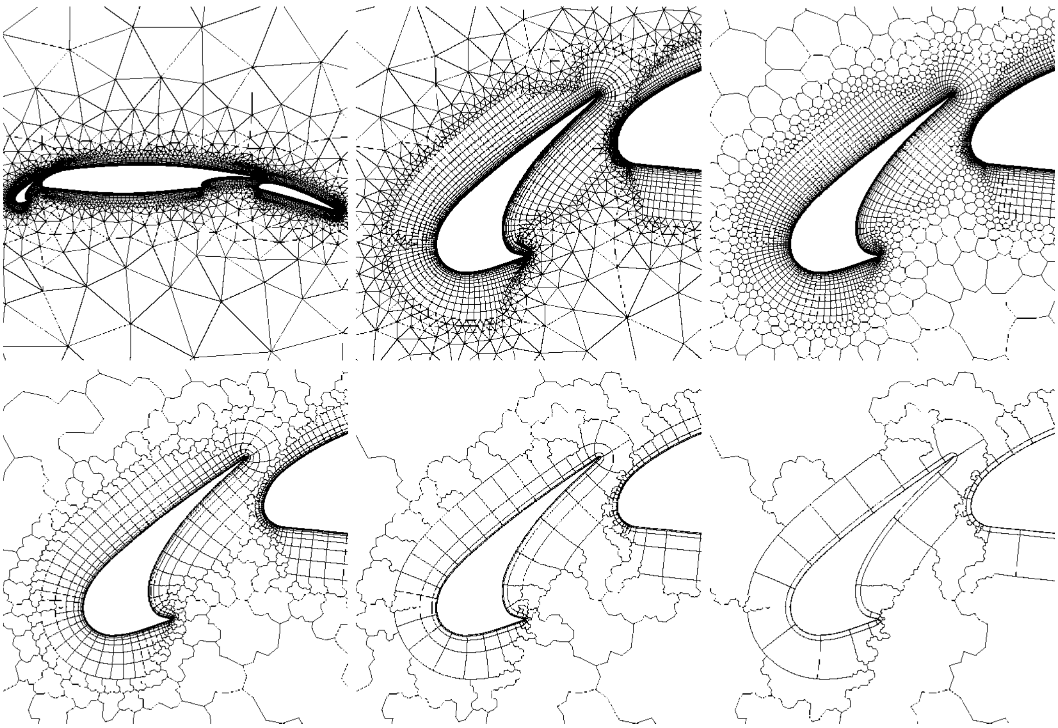

Figure 7. Multi element airfoil, two detailed views of the primary grid, the dual grid and three agglomeration levels.

cycle is reduced. The preserved structure offers the opportunity for future comparison of the performance of structured and unstructured methods on similar meshes.

REFERENCES

1. Lallemand M, Stave H, Dervieux A. Unstructured multigriding by volume agglomeration: current status. *Computers and Fluids* 1992; **21**(3):397–433.
2. Venkatakrishnan V, Mavriplis DJ. Agglomeration multigrid for the three-dimensional Euler equations. *AIAA-94-0069*, 1994.
3. Mavriplis DJ. Directional agglomeration multigrid techniques for high Reynolds number viscous flow solvers. *AIAA-98-0612*, 1998.
4. Okamoto N, Nakahashi K, Obayashi S. A coarse grid generation algorithm for agglomeration multigrid method on unstructured grids. *AIAA-98-0615*, 1998.
5. Berglind T. An agglomeration algorithm for Navier–Stokes grids. *AIAA-2000-2254*, 2000.

