# Rapid, parallel CFD grid generation using octrees

**G. J. Page**
**g.j.page@lboro.ac.uk**
Loughborough University
Loughborough, UK

## ABSTRACT

As Large Eddy Simulation is increasingly applied to flows containing complex geometry, grid generation becomes difficult and time consuming when using software originally developed for RANS flow solvers. The traditional 'pipeline' approach of grid generation → solve → visualise entails the time consuming transfer of large files and conversion of file formats. This work demonstrates a grid generation methodology developed specifically to be integrated with parallel LES. The current approach is to use a Cartesian grid with adaptive refinement based upon geometry intersection, surface detail and surface curvature. The grid is defined by an octree data structure with the geometry defined by triangular facets using the STL file format. The result is a set of 'cubical' subdomains, each with identical numbers of cells and uniform distributions within the cube. Some subdomains will be entirely fluid and can be solved using straightforward CFD techniques, whilst some cubes will be cut by the surfaces. Individual cells are then tagged as 'solid', 'fluid' or 'cut' with the solver expected to use an immersed boundary approach to model the surface. A key feature is the design of the algorithm to be parallelisable on both shared and distributed memory systems. The distributed memory parallel dynamically partitions the grid as it is being generated, so that the partitioning is suitable for a subsequent flow solver. Grid generation testing has been carried out on a variety of input CAD files ranging up to 350,000 facets. A landing gear case shows how the grid generator correctly finds the fluid inside of the tire and other cavities within the hub. In scalar mode, a grid with 4,916 cubes and 468 million cells is generated in less than 100 seconds, whilst in parallel on 32 processor cores this can be achieved in 4·6 seconds.

**Paper No. 3753**. Manuscript received 8 August 2011, revised version received 21 December 2011, accepted 22 March 2012.

https://doi.org/10.1017/S0001924000007910 Published online by Cambridge University Press

# 1.0 INTRODUCTION

In industrial application of CFD using Reynolds Averaged Navier-Stokes (RANS) approaches, grid generation is often one of the most difficult and time consuming tasks; furthermore, the quality of the grid can have a major influence on the accuracy of the flow solution. Up until recently, Large Eddy Simulation (LES) has been restricted to calculations on simple geometries in order to develop and validate the core methodology. LES can be even more sensitive to grid quality and in particular high aspect ratio elements (typical in RANS) can lead to poor solutions. Existing RANS grid generators are not suitable for LES when creating large grids (e.g. 100 million cells or more) as they are usually interactive scalar processes that are slow and use large amounts of memory. For these large simulations, the 'pipeline' paradigm of grid generation → solve → visualise is becoming a hindrance as it entails the time consuming reading, writing and copying of large files and conversions between different formats. The aim of this work is to develop a grid generation technique for LES that is fast and efficient, exploits parallelism and is integrated into the flow solver.

Over the years of CFD development there have been many fundamental approaches to grid generation with the most popular being 'body-fitted': the grid is wrapped around the object of interest and the finite volume cell faces coincide with the boundary surface. However, for complex geometry this can be a difficult task and can lead to poor finite volume shapes. The alternative approach of using a Cartesian grid to fill the domain with special treatment where the grid intersects the solid boundaries was originally proposed for Euler flow calculations[1,2]. De Zeeuw and Powell[1] used a hierarchical cell-based quadtree data structure to adaptively refine the mesh in regions of high gradient: both at a wall and where shocks were found. For Aftosmis *et al*[2]. the Cartesian grid is nested so that error estimates may be computed in order to drive the refinement in a rigorous manner. For viscous flow both Tseng and Ferziger[3] and Emblemsvag *et al*[4] used Cartesian grids with an 'immersed boundary' and 'ghost cells' to handle the wall interface. However, both of these works use a simple background Cartesian grid with predefined clustering near the object of interest – there is no hierarchical adaption to help resolve features near the wall. Kang *et al*[5] have studied the accuracy of the immersed boundary technique and demonstrated how this is suitable for LES.

Kamatsuchi[6] and Ishida *et al*[9] present the 'Building Cube Method' which uses an octree data structure in three-dimensions to handle the subdivision of the cubes to resolve surfaces and flow features. Since each cube has an identical number of grid points the authors point out how this simplifies parallel load balancing. The method presented in the current work is similar, although it has been developed independently.

The philosophy of integrating grid generation and solution method in parallel is also proposed in the work of Dawes *et al*[7,8]. The focus of their work is on capturing extremely complex geometry using an octree and a level set formulation to allow surface modification in a design environment. They convert the Cartesian grid and surface layer into a fully unstructured form for their own RANS solver or a commercial solver; there is little integration between the underlying octree formulation and the flow solver data structures.

This paper continues with a description of the core grid generation methodology and the parallel implementations. The following section shows example grids for a variety of input geometries and timings to show the parallel speed-up and efficiency. Finally, conclusions are presented and issues that still need resolving are discussed.

https://doi.org/10.1017/S0001924000007910 Published online by Cambridge University Press

# 2.0 METHODOLOGY

## 2.1 Solid model CAD representation

Commercial CAD packages use their own internal representation of three-dimensional solid models and proprietary file formats. Whilst the grid generation package could connect directly to the CAD package through interface libraries, this would involve multiple interfaces to support the most commonly used packages. More importantly, this typically requires licensing, and when generating a grid in parallel we may wish to access hundreds or thousands of licenses simultaneously. The alternative is to use a neutral standards defined format such as the Initial Graphics Exchange Specification (IGES) or the Standard for the Exchange of Product model data (STEP), but in practice both of these formats have many pitfalls. An alternative is the STereo Lithography (STL) file format used to transfer CAD data to rapid prototyping machines. Whilst this is a proprietary format, it is very simple, is available as an output option on all CAD systems and can be viewed by CFD visualisation tools. Essentially, the STL file defines a solid as a list of triangular facets with an optional facet normal. There is no connectivity between facets defined within the file and the coordinate of a vertex which is part of say four facets will be repeated four times in each facet definition. This does lead to the possibility that small gaps can appear between facets, but in practice this has found to be extremely rare and can be easily fixed by matching vertices with a given tolerance. The only drawback to the STL format is that smooth curved surfaces are always represented fundamentally in a faceted form. It is important that the user select a sufficiently fine facet resolution when generating the STL from CAD if a smooth representation of curved surfaces is important.

## 2.2 Octree grid generation algorithm

The fundamental data structure used to represent the hierarchical refined grid is an octree (see Fig. 1). For an octree, a parent node (coloured blue) has eight children, which themselves may have their own children. The nodes with no children are leaf nodes (coloured red), and these will contain data that is being stored in the tree. Simple procedures are then available to grow nodes, prune the tree, traverse the complete tree, or just traverse the leaf nodes.

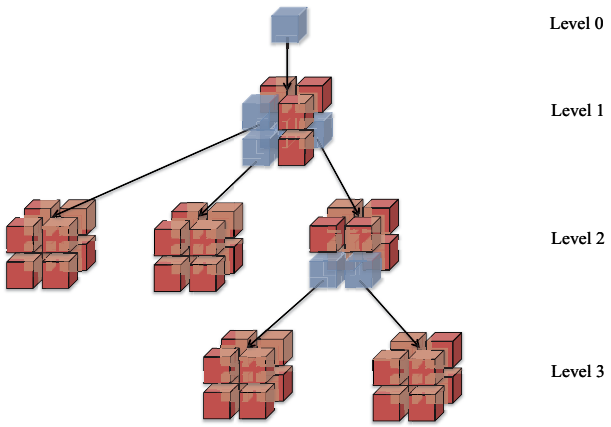



Figure 1. Octree data structure.

https://doi.org/10.1017/S0001924000007910 Published online by Cambridge University Press

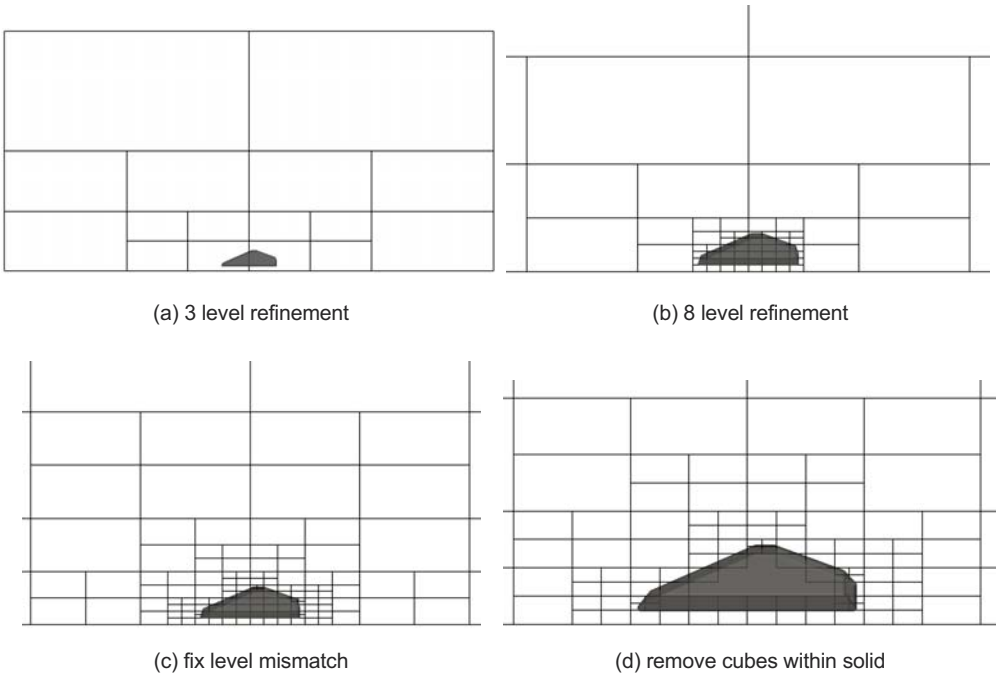



Figure 2. Generation of cubes.

For grid generation the data stored in the leaf node is a cube which specifies the origin in space of the cube and its size and number of grid points in each Cartesian direction. This information is then sufficient to compute the $(x,y,z)$ co-ordinates of any grid point for a given cube and $(i,j,k)$ indices. To refine the grid in this cube, eight children are grown from the original node, which becomes a parent. These children inherit the number of grid points in each direction, but now the size of the cube has halved so the resolution is doubled . The parent is now no longer a leaf node and does not contain a grid (i.e. the grids are contiguous and non-nested). Within the cubes, the grid lines define the edges of the finite volume cells for flow computation. Note that cube is a convenient term, but it is not required that all three Cartesian directions are of equal size or of equal number of grid points.

Prior to the grid generation process we need one or more faceted representations of geometry, the size and position of an outer domain box and some target cell sizes.

To illustrate the process, refer to Fig. 2 which shows the generation of a grid for a simplified car body. In order to more clearly see the process, a cut has been made through the three-dimensional problem.

To commence the process the root node is created and this is set to define the cube that represents the outer domain. This is then denoted as level zero.

There is then an iterative process in which the leaf nodes of the octree are repeatedly traversed and cubes are refined depending upon surface intersection, facet resolution and curvature.

For each cube, a list of facets that intersect the cube is generated from the list of facets that intersected the parent cube using Moller's triangle-box overlap algorithm[10]. If the list is non-zero, then this cube intersects geometry and is tagged as 'cut' and refined again. As the tree grows and the cubes become smaller in size, the list of facets that intersect a cube will become shorter and the intersection test becomes faster.

https://doi.org/10.1017/S0001924000007910 Published online by Cambridge University Press

This is repeated generating further levels with refinement for a given cube based upon three criterion: target size to resolve surfaces; facet resolution; curvature resolution.

This ensures that all surface is meshed to a known resolution, that in regions of high detail where small facets are found that there a known number of grid points for the smallest facet dimension found in a cube, and that in regions of high surface curvature (where pressure gradients are largest) there will be local refinement.

To determine if a cube should be refined further due to high surface curvature then the Gauss-Bonnet scheme (also known as the angle deficit method) is used to compute the largest Gaussian curvature[11] found within facets that are known to intersect the cube.

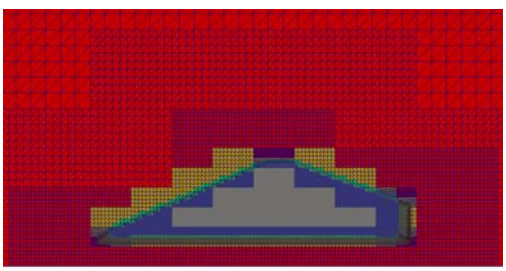
The curvature is computed at vertex $v$ with neighbouring vertices $\{v_i\}_{i=0}^{n-1}$ connected by edges $e_i$. Two successive edges will make an angle $\alpha_i$ between them and the total curvature is then computed from a summation of the internal angles of the edges meeting at the vertex;

$$K = \frac{2\pi - \sum_{i=0}^{n-1} \alpha_i}{\frac{1}{3} A} \qquad \ldots (1)$$

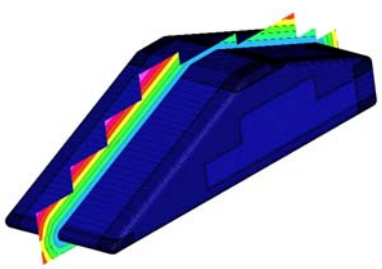
This extra refinement can be seen at the corners of the model in Fig. 3(b) where high curvature is present.

Generally, a large jump in grid size between cubes is detrimental to the numerical accuracy of the spatial discretisation of fluxes and it would be preferable if there is no more than a factor of two jump in linear cell size between cubes. The simple refinement process can produce large mismatches in size between cubes neighbouring in physical space (for example Fig. 2(b)). To correct this, the cubes are traversed checking adjacent cubes for more than one level of mismatch – if this is found then the coarser cube is refined (Fig. 2(c)). This may then lead to more mismatches in levels, so this is repeated until the traversal does not produce any refinements.

At this stage, some cubes will have been tagged as 'cut', the remaining 'unknown' cubes do not intersect geometry, but could be totally inside or outside of geometry (e.g. 'solid' or 'fluid'). There is then a traversal over all unknown cubes using a fast ray tracing algorithm[12] to set cube status as 'solid' or 'fluid'. A bi-directional ray is sent in an arbitrary direction from the centre of the cube and the number of intersections with surface facets is counted. If the origin of the ray is inside a solid then there will be an odd number of intersections, and if it is outside the solid there will be an even number of intersections. Both directions of the ray should agree as to whether there is an even or odd number of intersections: a discrepancy indicates a potential problem with gaps in the



(a) computational cells (blue: solid cell; green: surface intersection cell; yellow: fluid cell in cube that intersects surface; red: fluid cell in non-intersecting cube)



(b) Surface defined by zero value level set and distance from wall (blue: zero; pink: high value)

Figure 3. Cell intersection and level set.

https://doi.org/10.1017/S0001924000007910 Published online by Cambridge University Press

faceted geometry and the calculation is repeated with another arbitrary ray until agreement is found. Any cubes that are 'solid' are then removed (Fig. 2(d)).

For the remaining intersecting cubes, we need to know which finite volume cells within the cube intersect the surface geometry and this information will then need to be stored as a three-dimensional array. A naive use of ray-tracing for all cells within a cube will be computationally too expensive, this is particularly true for complex geometry with large numbers of facets. Ishida *et al*[9] used a fill algorithm starting from a user defined point within the fluid. However, a fill algorithm inherently contains dependencies, the state of a cube is determined from the state of its neighbours, and within a cube each cell state will depend upon neighbouring cell or cube states. This makes the pure fill algorithm a scalar process. In this work, a hybrid scanline fill is utilised within each cube. The cells of the cube are split into scanlines which are bounded by cube boundaries or geometry intersection (i.e. every cell in the scanline must be of the same type). The whole of this scanline can be set to 'solid' or 'fluid' by examining the status of cells adjacent to the scanline. If there are no known adjacent cells then we finally resort to a ray-tracing evaluation to determine the status of the scanline. To avoid testing the ray against all facets, firstly the ray is tested for intersection against the cubes. Facets within these intersecting cubes are then added to a candiate list to be tested for intersection with the ray. The three different cell types can be seen in Fig. 3(a). Note that the grid shown here is much coarser than those used for parallel testing later so as to be able to see the individual cells.

Finally, most CFD techniques require information on distance to the nearest wall and this is also the case for LES where wall damping is used to reduce subgrid scale viscosity near the wall. For cubes that intersect geometry, for each grid point the smallest distance to facets that intersect the

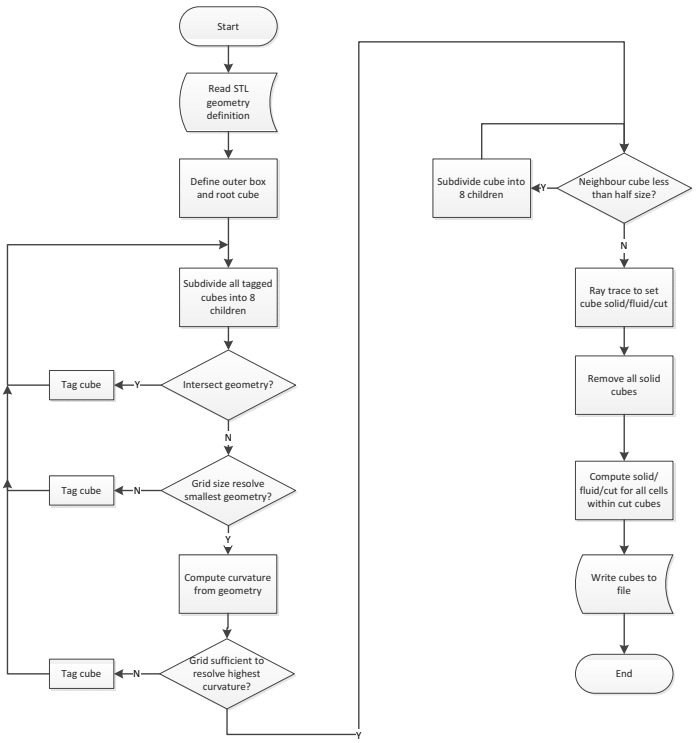



Figure 4. Flowchart of grid generation process.

https://doi.org/10.1017/S0001924000007910 Published online by Cambridge University Press

cube is determined and its sign set based upon whether the grid point is inside the solid. This distance function is then a level set with the zero level representing the surface of the geometry. Fig. 3(a) shows an isosurface of the zero level for the car body together with a contoured cut to show the distance from the wall for cells in cubes that intersect walls.

It should be noted that all major computational tasks (intersection, curvature, solid/fluid, distance) can be computed independently on each cube so making the approach inherently parallel. A flowchart of the overall grid generation process is shown in Fig. 4.

## 2.3 Shared memory parallel

Until recently, each year it had become quicker to achieve a CFD solution as computer processors dramatically increased in clock speed. But, processors have now reached a limit of clock speed, and the benefits of improved fabrication processes are being used to increase the number of processor cores within a single package or to add data parallel General Purpose Graphical Processing Units (GPGPU). Whilst CFD solvers can utilise the multiple core processors by using the parallel techniques developed for running on compute clusters, the grid generation process also needs to exploit parallelism to keep pace.

Each core in the processor package has equal access to the system memory and so for these types of computer architecture it is a natural fit to used shared memory parallelism. In this work we have implemented shared memory parallelism using Open Multi-Processing (OpenMP). This uses compiler directives to indicate loops that can be computed in parallel; at run-time the loop is distributed across the processor threads and at the end of the loop returns to a single thread. This was straightforward to implement with most loops over the cubes being targeted to run in parallel. As the octree data structure is shared across threads, a critical region is used for locations where children are being added to the octree due to cube refinement.

As will be seen later, this produces good parallel speed-up and is totally transparent to the user – the grid generation automatically utilises all processor cores available on a given machine.

## 2.4 Distributed memory parallel

The ultimate aim of this work is to produce a system to solve LES on complex geometries with extremely large meshes. Since the flow solver will need to use distributed memory parallelism and we wish to avoid the separation of grid generation and solver, then logically the grid generation needs to also use distributed memory parallelism.

One difficulty of parallelising grid generation is that it needs to be dynamically partitioned across the processors as the grid is created. In this work, as the octree is grown the leaf nodes and associated cubes are allocated to processes dynamically – the end result of the grid generation is not only a grid but a partitioning of the cubes across processors ready for the solver to carry out a parallel computation.

All processes have a complete copy of the geometry and generate the same octree data structure independently. However, leaf nodes of the octree are ‘owned’ by a process. The ‘owning’ process computes tasks such as intersection with geometry and then this status information is synchronised across all processes using a broadcast. Thus, if a cube is removed because it is within a solid, this will occur on all processes’ copies of the octree data structure. Similarly, where significant data on the grid is stored, such as the three dimensional array of cells within a cube denoting solid, fluid or cut, this only exists in the owning process and is not duplicated. When iterating over the leaf nodes, the iterator for a process uses a list of nodes that are owned by that process.

https://doi.org/10.1017/S0001924000007910 Published online by Cambridge University Press

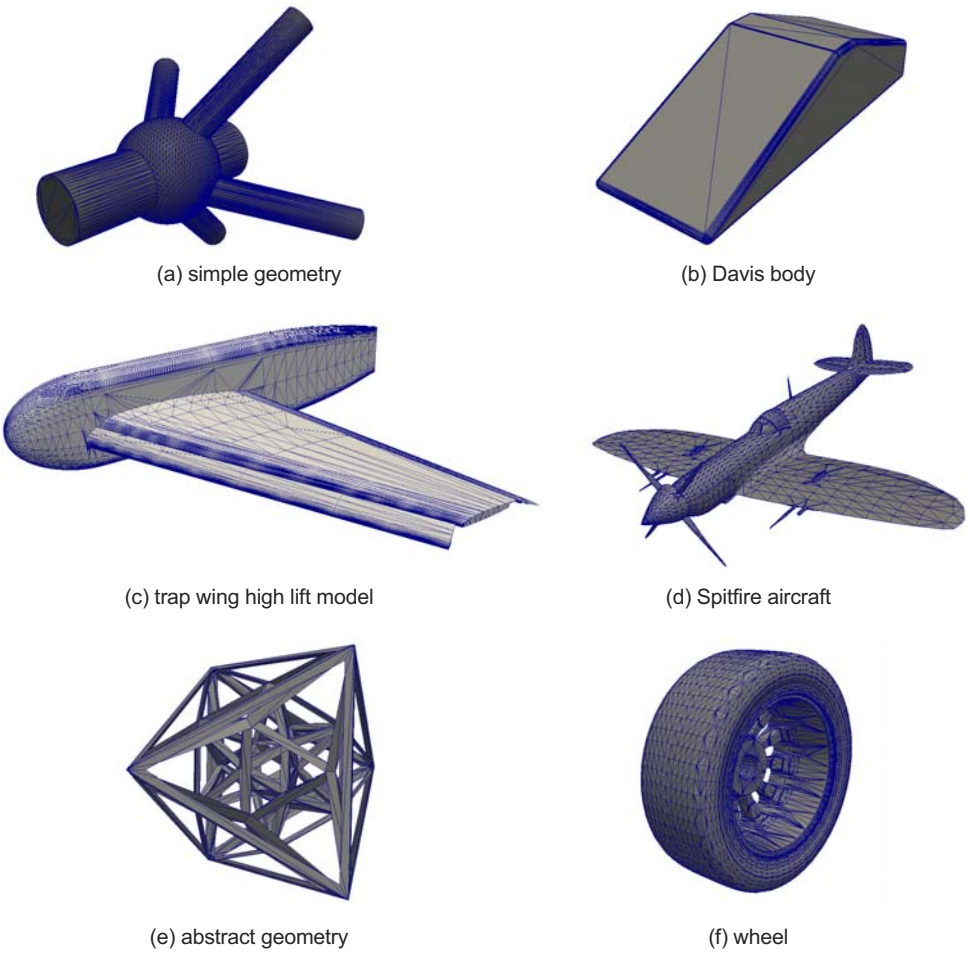



Figure 5. CAD geometries used for testing.

One of the key issues is to be able to dynamically partition whilst the grid is growing, whilst achieving good load-balancing and a locality of intercommunication for the solver. Of significant benefit is that each fluid cube will have an identical number of cells and so will have equal work for the flow solver. For intersecting cubes then we will have a mixture of cells with fluid, solid and intersect status. For a given solver, there will be an extra computational cost of the special treatment of the cut cells, but this could be balanced by the solid cells that will not need any computation. The hope is that intersecting cubes will be less expensive for the flow solver to compute and thus will not increase overall computation time – this needs to be tested with real flow solvers.

It is assumed that we may be generating more cubes than the number of processes available so it is not just a procedure to allocate each new cube to its own process. Typically ratios of cubes to processes are likely to be around a factor of ten. As a cube is refined eight new child nodes are generated which need to be assigned to a process, and the parent node which is removed from a process and becomes globally owned. Since these cubes belonging to the children are physically adjacent, for good flow solver partitioning it would be useful to give them to the same process and

https://doi.org/10.1017/S0001924000007910 Published online by Cambridge University Press

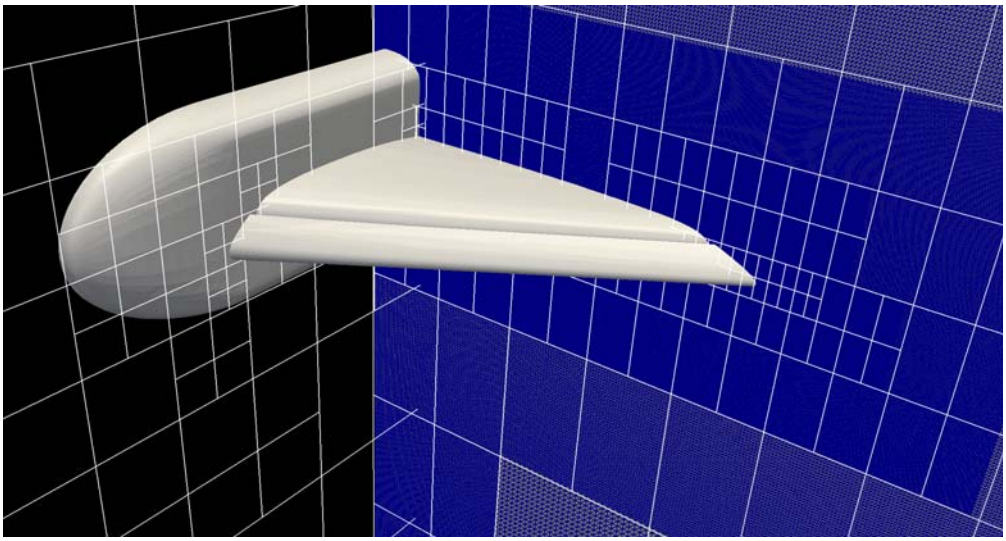


Figure 6. Trap wing cube structure and high density grid (238 million cells).

so we typically assign parcels of four or eight new leaf nodes and their cubes to a single process. The candidate process is found by determining the process with the lowest loading, and if there is more than one with the same low loading then using the lowest rank process.

The distributed memory parallelism is implemented using the Message Passing Interface (MPI) standard and can run on a wide variety of systems.

## 3.0 RESULTS

The algorithm has been tested on a range of input CAD geometries. Some of these are shown in Figs 5 and 7(a). These were chosen to test a variety of requirements. The simple geometry and Davis model were generated using standard solid modelling software and both have well defined radii of curvature which are used to test the Gaussian curvature calculation and subsequent local curvature refinement. The trap wing high lift model[13] is used extensively to benchmark CFD solvers for high lift problems and contains small gaps between the slat and wing and flap and wing as well as a large range in facet sizes. It is also a half model and in this case one of the planes of the outer flow domain is coincident with the centreline of the model. A grid of the trap wing with 238 million cells is shown in Fig. 6. The Spitfire model has a small gap between the rudder and the fin; for a coarse grid this is not resolved and the CFD grid contains a single solid object, whereas as the grid is made finer, the gap becomes resolved and there are separate solid objects in the CFD domain. In addition the Spitfire is placed in a nose up and yawed attitude so that the grid axes are not aligned with the predominant aircraft directions. The abstract geometry is chosen as this contains long thin triangular facets which will highlight flaws in the triangle-box overlap algorithm and a lattice like structure that will test the ray-tracing determination of solid or fluid. To test sensitivity to the number of triangular facets in the geometry definition, the abstract geometry only contains 888 facets whilst the landing gear geometry has 350,000 facets. The trap wing STL file was created from the CAD geometry provided on the AIAA CFD High Lift Prediction Workshop website[14] whilst the remainder of the test STL files are freely available from the Dassault Systemes 3DVIA collaborative website[15].

https://doi.org/10.1017/S0001924000007910 Published online by Cambridge University Press

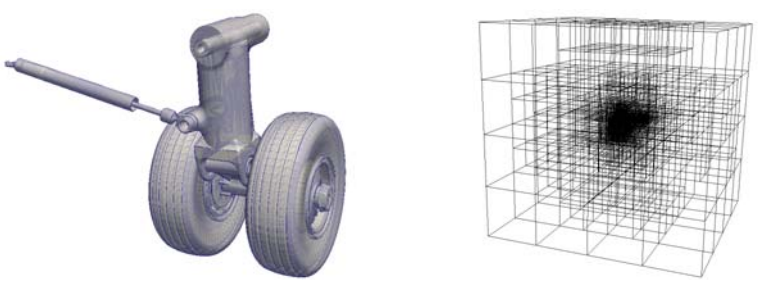
The landing gear case is shown in more detail as this is probably the most challenging of the test cases and is also an application area of current interest: LES can be used to predict unsteady shedding and hence noise production[16]. As can be seen in Fig. 7(a), there is considerable detail resolved in the model, such as the grooves in the tyres, cavities and gaps in the wheel hub and multiple pivots due to the retraction mechanism. No modifications were made to the STL model, the grid generator was run with default parameters to produce results shown here in a matter of minutes. It would be inconceivable to use a multiblock structured grid generator for this case, and even state of the art unstructured grid generators would require considerable intervention and probably hours of calculation to produce grids of order 100 million elements.

An overall view of the domain is shown in Fig. 7(b). The outer domain is set to be four characteristic lengths away from the maximum extents of the model geometry and there are seven levels of refinement – the cubes around the geometry become extremely small and on this diagram appear as a dark mass at the centre of the domain. This shows the ability of the method to give good resolution around the object without incurring the penalty of the refinement spreading to the far-field, as would be the case with simple structured grids. This is of particular importance to LES where we need to avoid high aspect ratio cells. The domain contains 2,353 cubes each with  cells giving a total of 8·47 million cells. This is considerably less than the parallel testing example shown later which has 63·5 million cells and was chosen to allow easier visualisation of the cubes and cells.

Figure 8(a) is a cut of the domain through the wheels and axle that illustrates the complexity of this case. The cells are coloured to distinguish those as intersected by the surface (green), solid (blue), fluid (red or yellow). The ray-tracing algorithm has found the inside of the tire and this is correctly marked as a fluid – even though it is completely isolated from the fluid surrounding the wheel. Similar cavities are correctly marked within the wheel hub. This example shows the strength of the use of ray-tracing with a scanline fill within each cube as all regions are automatically marked. An approach that takes a user defined seed point as being fluid, then flood filling all the cells would be cheaper to compute, but is more difficult to parallelise – and would need multiple user inputs to seed the fluid regions within the tyres.

Finally, the zero level set surface and a slice of the distance function is shown in Fig. 8(b). The level set is sufficiently accurate to resolve sharply the grooves in the tire.

Whilst this work is intended to lead to an integrated grid generator and flow solver, it is clear that the octree-cube grid structure could be easily converted into an arbitrary unstructured grid and hence be used by a variety of existing unstructured grid solvers. The highly regular nature of the elements may be beneficial in creating stencils for high order schemes on unstructured grids.



(a) STL CAD geometry (b) cubes

Figure 7. Landing gear grid overall view.

https://doi.org/10.1017/S0001924000007910 Published online by Cambridge University Press

**Table 1**
**Parallel testing cases**

| Number of: | Facets | Cubes | Cells | Fluid cells | Cut cells |
|---|---|---|---|---|---|
| simple geometry | 7,342 | 4,916 | 468,101,520 | 396,597,818 | 3,890,450 |
| Davis body | 1,852 | 3,954 | 164,407,320 | 135,432,458 | 2,172,250 |
| landing gear | 350,418 | 2,353 | 63,460,410 | 57,479,112 | 1,471,103 |

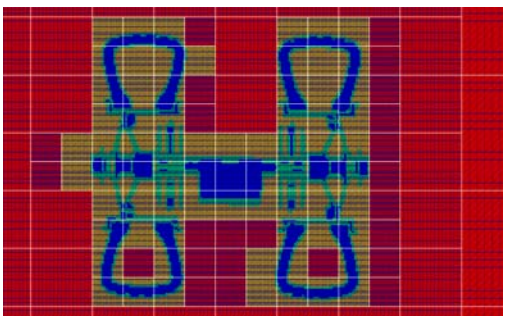
### 3.1 Parallel testing

Three cases have been chosen to test the parallel speed up and efficiency. The cases are shown in Table 1 which summarises the number of facets in the CAD geometry and the resulting grid in terms of cubes and cells.

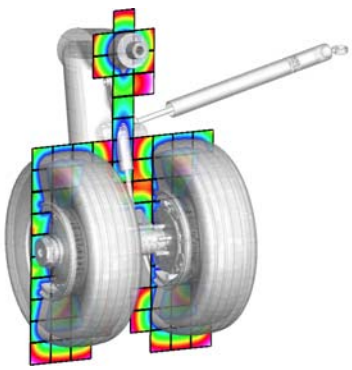
Testing was carried out on an Intel 'Nehalem' Xeon E5520 (2·26GHz) system. This was configured as two quad core processors per node, with nodes being connected by gigabit ethernet.

Plots of wall time for scalar, shared memory parallel (OpenMP) and distributed memory parallel (MPI) are shown in Fig. 9 for up to 32 processor cores. Of note is the basic speed and efficiency of the algorithm: even running on a single processor core, a grid for the 'simple geometry' with 460 million cells can be generated in less than 100 seconds, and the maximum memory usage is 1·3GB. Running on 32 processor cores this reduces to 4·6 seconds. Of interest is the behaviour of the 'landing gear' case which is much slower than the other cases and exhibits anomalous behaviour on eight cores. This case differs from the first two in that the geometry definition has a large number of triangular facets. The ray-tracing algorithm cost scales with the number of facets and the complexity of the surface leads to more ray-tracing evaulations per cube. The net result is that compared to the 'simple geometry' it takes four times as long to generate a grid with 1/7 of the number of cells. It should be noted that this is perhaps better than may be expected considering that the geometry has 47 times more facets than the 'simple geometry'.

A more detailed analysis of the results are shown in Table 1. Parallel speed up and efficiency is generally good up to eight cores, but drops off at 16 and 32 cores for the distributed memory parallel. The main factor appears to be load balancing, the partitioning of the octree data structure as it is created is based upon even loading and reduced communication for the flow solver – it aims



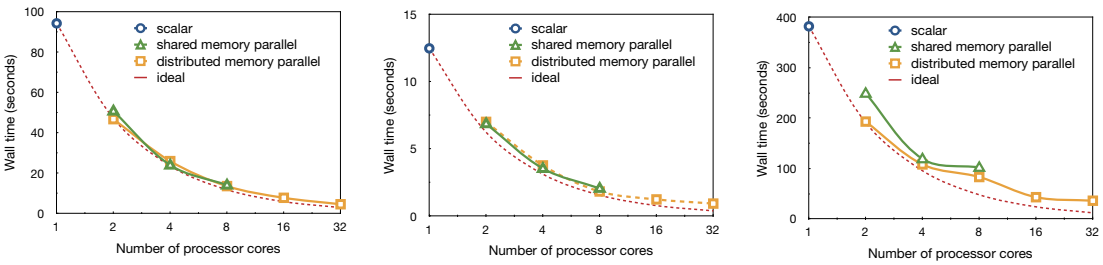
(a) cut showing computational cells on coarse grid (blue: solid cell; green: surface intersection cell; yellow: fluid cell in cube that intersects surface; red: fluid cell in non-intersecting cube)



(b) Surface defined by zero value level set and distance from wall

Figure 8. Landing gear grid detail view.

https://doi.org/10.1017/S0001924000007910 Published online by Cambridge University Press




(a) simple geometry (b) Davis body (c) landing gear

Figure 9. Grid generation wall time.

to achieve an equal load of cubes across processes with adjacent cubes being on the same processor core. However, for the grid generator, the computational load depends upon the number of intersecting cubes owned by a process. Examination of timings for individual processes showed a significant variation indicating poor load balancing. This may be an acceptable compromise as it is more important that the flow solver achieves good load balancing – this requires further testing, particularly on much larger number of processor cores. The load balancing problem does not explain the anomalous result for eight cores on the 'landing gear' case as this occurs for both distributed and shared memory runs. For shared memory there is no explicit partitioning of the cubes as each parallel loop over the cubes in the domain is evenly distributed over the processes at run-time.

**Table 2**
**Parallel speed up and efficiency**

| | | simple geometry | | | | |
|---|---|---|---|---|---|---|
| | **number of cores** | **2** | **4** | **8** | **16** | **32** |
| shared memory | speed up | 1·86 | 3·93 | 6·69 | | |
| | efficiency | 93% | 98% | 84% | | |
| distributed memory | speed up | 2·02 | 3·63 | 6·94 | 12·1 | 20·5 |
| | efficiency | 101% | 91% | 87% | 76% | 64% |
| | | **Davis body** | | | | |
| | **number of cores** | **2** | **4** | **8** | **16** | **32** |
| shared memory | speed up | 1·82 | 3·53 | 6·06 | | |
| | efficiency | 91% | 88% | 76% | | |
| distributed memory | speed up | 1·78 | 3·31 | 6·82 | 10·1 | 13·4 |
| | efficiency | 89% | 83% | 85% | 63% | 42% |
| | | **landing gear** | | | | |
| | **number of cores** | **2** | **4** | **8** | **16** | **32** |
| shared memory | speed up | 1·53 | 3·21 | 3·75 | | |
| | efficiency | 77% | 80% | 47% | | |
| distributed memory | speed up | 1·98 | 3·54 | 4·60 | 8·88 | 10·6 |
| | efficiency | 99% | 88% | 58% | 56% | 33% |

https://doi.org/10.1017/S0001924000007910 Published online by Cambridge University Press

The implementations of both the shared and distributed parallelism are relatively simple and with further testing and instrumentation could be tuned for better parallel speed up. Nevertheless, both approaches achieve extremely fast grid generation, and in the case of the distributed memory also achieves a partitioning of the grid ready for the flow solver.

## 4.0 CONCLUSIONS

A fast and efficient parallel grid generation method using an octree data structure has been described. The grid generation process is suitable for very large grids with good control over grid quality and resolution as is required by Large Eddy Simulation.

Testing on a variety of input CAD geometries has demonstrated the speed and robustness of the methodology. The algorithm has been designed from the beginning based upon the use of parallel computing. Reasonable parallel speed ups have been achieved on both shared and distributed memory implementations. The computation time is dependent upon the number of facets used to resolve the geometry.

For distributed memory parallel, the resulting grid is already partitioned across the processes ready for the flow solver. The partitioning is optimised for the more expensive flow solver and hence impacts the parallel efficiency of the grid generation.

The method shows considerable promise and the next stage of work will integrate the grid generation with an LES flow solver.

## REFERENCES


1. DE ZEEUW, D. and POWELL, K.G. An adaptively refined Cartesian mesh solver for the Euler equations, *J Computational Physics*, 1993, **104**, (56-68).
2 AFTOSMIS, M.J., BERGER, M.G. and ADOMAVICIUS, G. A parallel multilevel method for adaptively refined Cartesian grids with embedded boundaries, In proceedings of the 38th Aerospace Sciences Meeting and Exhibit, 10-13 January 2000, Reno, Nevada, USA. AIAA-2000-0808.
3. TSENG, Y-H. and FERZIGER, J.H. A ghost-cell immersed boundary method for flow in complex geometry, *J Computational Physics*, 2003, **192**, (593-623).
4. EMBLEMSVAG, J-E., SUZUKI, R. and CANDLER, G.V. Cartesian grid method for moderate-Reynolds-number flows around complex moving objects, *AIAA J*, 2005, **43**, (1), pp 76-86.
5. Kang, S. Iaccarino, G. Moin, P. Accurate immersed-boundary reconstructions for viscous flow simulations, *AIAA J*, 2009, **47**, (7), pp 1750-1760.
6 KAMATSUCHI, T. Flow simulation around complex geometries with solution adaptive Cartesian grid method, In proceedings of the 18th AIAA Computational Fluid Dynamics conference, 25-28 June 2007, Miami, Florida, USA. AIAA-2007-4189.
7. DAWES, W.N., HARVEY, S.A., FELLOWS, S., FAVARETTO, C.F. and VELIVELLI, A. Viscous Layer Meshes from Level Sets on Cartesian Meshes, In proceedings of the 45th Aerospace Sciences Meeting and Exhibit, 8-11 January 2007, Reno, Nevada, USA. AIAA-2007-0555.
8. DAWES, W.N., HARVEY, S.A., FELLOWS, S., ECCLES, N., JAEGGI, D. and KELLAR, W.P. A practical demonstration of scalable, parallel mesh generation, In proceedings of the 47th Aerospace Sciences Meeting and Exhibit, 5-8 January 2009, Orlando, Florida, USA. AIAA-2009-0981 (2009)
9. ISHIDA, T., TAKAHASHI, S. and NAKAHASHI, K. Efficient Cartesian mesh approach for flow computations around moving and deforming bodies, In proceedings of the 19th AIAA Computational Fluid Dynamics conference, 22-25 June 2009, San Antonio, Texas, USA. AIAA-2009-3879.
10. AKENINE-MÖLLER, T. Fast 3D triangle-box overlap testing, In proceedings of the SIGGRAPH '05: ACM SIGGRAPH 2005 Courses, (Los Angeles, California, USA, July 31 – August 04, 2005). J. Fujii,. doi http://doi.acm.org/10.1145/1198555.1198747 (2005)
11. SURAZHSKY, T., MAGID, E., SOLDEA, O., ELBER, G. and RIVLIN, E. A comparison of Gaussian and mean curvatures estimation methods on triangular meshes, In proceedings of the IEEE International Conference on Robotics and Automation, **1**, (1021-1026), 2003.


https://doi.org/10.1017/S0001924000007910 Published online by Cambridge University Press

12. MÖLLER, T. and TRUMBORE, B. Fast, minimum storage ray/triangle intersection, In proceedings of the SIGGRAPH '05: ACM SIGGRAPH 2005 Courses, (Los Angeles, California, USA, July 31 – August 04, 2005). FUJII, J., Ed., doi http://doi.acm.org/10.1145/1198555.1198746 (2005).
13. MCGINLEY, C., JENKINS, L. and WATSON, R. 3D high-lift flow physics experiment – transition measurements, In proceedings of the AIAA Fluid Dynamics Conference, 6-9 June 2005, Toronto, Canada. AIAA-2005-5148.
14. Anon. AIAA CFD High Lift Prediction Workshop (HiLiftPW) Series (online). Available at: http://hiliftpw.larc.nasa.gov/ (Accessed 20 December 2011).
15. Anon. (online) Available at: http://www.3dvia.com/ (Accessed 20 December 2011).
16. LI, Y., SATTI, R., LEW, P-T., SHOCK, R. and NOELTING, S. Computational aeroacoustic analysis of flow around a complex nose landing gear conÞguration, In proceedings of the 14th AIAA/CEAS Aeroacoustics Conference, 5-7 May 2008 – June 2009, Vancouver, British Columbia, Canada. AIAA-2008-2916 (2008).

https://doi.org/10.1017/S0001924000007910 Published online by Cambridge University Press