

Fine-grained Speculative Topological Transformation Scheme for Local Reconnection Methods

Fotis Drakopoulos^{*}, Christos Tsolakis[†] and Nikos P. Chrisochoides[‡]
Old Dominion University, Computer Science Department, Norfolk, VA 23529, USA

A parallel approach for topological transformations for local reconnection methods is presented. The proposed scheme combines known parallel techniques like data over-decomposition and load balancing with widely used topological transformations also known as flips or swaps. The parallel scheme is evaluated on a variety of aerospace configurations. Early results indicate that the high quality and performance attributes of this method offer substantial improvement over existing state-of-the-art technology.

I. Introduction

The long term goal of this project is to achieve extreme-scale adaptive CFD simulations on the complex, heterogeneous High Performance Computing (HPC) platforms. To accomplish this goal, a telescopic approach (see Figure 1) is proposed in [1, 2]. The telescopic approach is critical in leveraging the concurrency that exists at multiple levels in parallel and adaptive simulations. At the chip level the telescopic approach deploys a speculative approach, sometimes also called Parallel Optimistic (PO) approach [3], which explores concurrency at a fine-grain (element) level using data decomposition. The focus of this study is the implementation of the speculative layer, having however in mind that the code should be modular enough so that it can be used throughout the telescopic approach.

The mesh generation kernel of the telescopic approach for CFD applications (CDT3D) is essentially a combination of Advancing Front type point placement, direct point insertion, and parallel multi-threaded connectivity optimization schemes [4]. The proposed speculative method repetitively reconnects the mesh using tightly-coupled topological transformations. Simple hill-climbing is employed to maximize the quality of the worst local element. The parallel local reconnection scheme is based on (i) data over-decomposition, (ii) atomic operations to avoid data races and maintain a valid mesh throughout the procedure, and (iii) load-balancing to redistribute work-units among the threads.

The proposed method is designed to be a module of the CDT3D mesh generation software which focuses on five important aspects of parallel mesh generation:

- 1) **Stability** is the requirement that the quality of the mesh generated in parallel must be comparable to that of a mesh generated sequentially. The quality is defined in terms of the shape of the elements and the number of the

^{*}Research Assistant, Center for Real-Time Computing

[†]Research Assistant, Center for Real-Time Computing

[‡]Richard T. Cheng Chair Professor of Computer Science, Center for Real-Time Computing

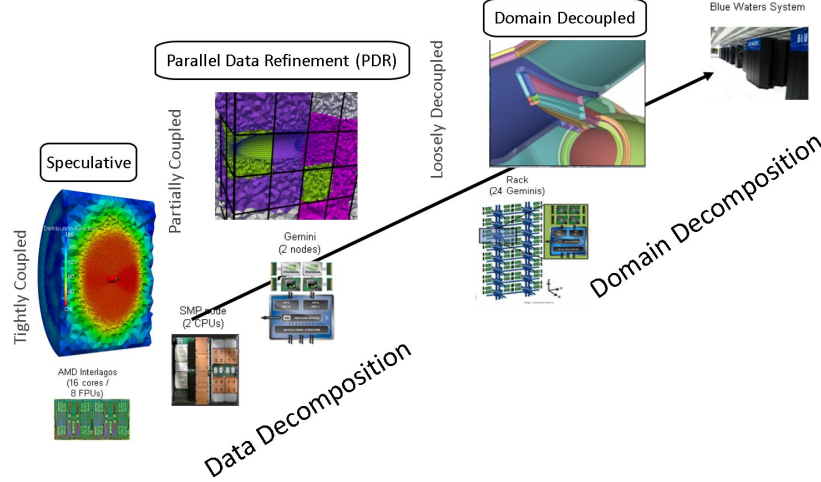


Fig. 1 Telescopic Approach for Extreme-Scale Parallel Mesh Generation [2]

elements (fewer is better for the same shape constraint).

- 2) **Reproducibility** [5] which has two forms *Strong Reproducibility* requires that the mesh generation code, when executed with the same input, produces identical results under the following modes of execution: (i) continuous without restarts, and (ii) with restarts and reconstructions of the internal data structures. *Weak Reproducibility* requires that the mesh generation code, when executed with the same input, produces results of the same quality under the following modes of execution: (i) continuous without restarts, and (ii) with restarts and reconstructions of the internal data structures.
- 3) **Robustness.** The ability of the software to correctly and efficiently process any input data and to consistently generate high-quality elements. Operator intervention into a parallel computation is not only highly expensive, but most likely infeasible due to the large number of concurrently processed sub-problems.
- 4) **Scalability.** The ability of the software to process the input data faster compared to state-of-the-art codes. Scalability is defined as the ratio of the time taken by the best sequential implementation to the time taken by the parallel implementation. The speedup is always limited by the inverse of the sequential fraction of the software (Amdahl's law), and therefore all non-trivial stages of the computation must be parallelized to leverage the current multi-core architectures.
- 5) **Code Re-Use.** A modular design of the parallel software, such that it can be replaced and/or updated with minimal effort. Due to the complexity of mesh generation codes, this is the only practical approach for keeping up with the ever-evolving algorithms and computer architectures.

The focus of this paper is the generation of high-quality isotropic meshes employing the speculative local reconnection approach. While boundary layer and metric-based anisotropic mesh generation are not supported yet, CDT3D's modular design allows for these methods to be integrated in the future.

CDT3D is inspired from state-of-the-art unstructured grid technology AFLR (Advancing Front-Local Reconnection) [6, 7] and is designed to be an alternative for industrial-strength extreme-scale parallel mesh generation. AFLR is directly incorporated in several industrial-strength systems. However, its software complexity, given that it is designed and optimized to be super-efficient on single-core machines raised questions as to whether its parallelization can meet all five requirements stated above and whether it can be accomplished within the time constraints of its users.

The proposed method optimizes the connectivity throughout the generation procedure. The connectivity is optimized using topological transformations coupled with a combined quality criterion: Delaunay in-sphere [8] and Min-Max type [9]. Topological transformations are fundamental operations in local reconnection. A topological transformation removes a set of elements and replaces them with a different set that occupies the same space. Transformations can be computationally expensive and time-consuming, mainly for two reasons: (i) a transformation does not always produce a geometrically valid connectivity, hence the orientation of the new elements needs to be verified, and (ii) a quality metric needs to be computed to decide whether the new connectivity is locally optimal or not. However, transformations involving more than three elements may have more than one candidate solutions and the cost to compute the optimal solution is a-priori non-polynomial [10]. Experimental evaluation using CDT3D indicates that the sequential connectivity optimization with local reconnection accounts for at least 80% of the mesh refinement time.

For those reasons, CDT3D employs a parallel speculative local reconnection approach. This approach is characterized by intense communication and resolution of dependencies at runtime. The low cost of communication allows the speculative approach to take advantage from dynamic load balancing schemes like the one presented in [3] that is used in this work.

II. Previous Work

Two approaches are typically used to generate a tetrahedral mesh from an input surface mesh: Delaunay or Advancing Front [11]. Delaunay methods can handle constrained polyhedral domains of arbitrary complexity using heuristic mesh modification techniques [12, 13]. Other Delaunay approaches can mathematically guarantee the quality of the mesh, but they do not preserve the boundary triangulation [14, 15].

The first efforts on the parallelization of existing sequential Delaunay mesh generation algorithms are reported in [16, 17]. These methods are based on the parallelization of the Bowyer-Watson kernel [18, 19]. A tightly-coupled distributed parallel Delaunay refinement algorithm for simple polyhedral domains whose constituent bounding edges and surfaces are separated by angles between 90° to 270° is presented in [20]. This algorithm can create large meshes up to 6 times faster than the traditional approach (i.e., sequentially generate a sufficiently dense mesh, partition the mesh into submeshes, distribute the submeshes to the processors, and sequentially refine the mesh).

A Delaunay method that allows for safe insertion of points independently without synchronization is presented in [21]. This method, based on a carefully constructed octree, splits the work-list of the candidate points up into smaller

lists such that already available sequential codes can be used without modifications in order to process the sublists. The drawback is that the amount of work assigned to different processors can vary significantly. However, by using over-decomposition in combination with a runtime software system for dynamic load balancing [22], the work among the processors can be redistributed effectively.

In some cases (e.g., high-aspect-ratio viscous meshes or three-dimensional meshes with sliver elements formed from four nearly coplanar points) the Delaunay criterion may not be the most suitable. For this reason, additional topological transformations on top of a Delaunay mesh have been introduced [9, 23–25].

The Advancing Front method can offer both a high-quality mesh due to optimal point placement but they can also preserve the input boundary without introducing extra points [26, 27]. Its main drawbacks are complexity and lower element quality when fronts collide in 3-dimensions. An Advancing Front mesh generator for shared memory architectures is presented in [28]. The domain is first subdivided spatially using a coarse octree, and then each octree-leaf is meshed in parallel. In [29], a coarse tetrahedral mesh is generated first to provide the basis of block interfaces and then partitioned using METIS [30] partitioning algorithms. A volume mesh is generated for each subdomain in parallel using an Advancing Front method, and the subdomains are combined to create a single mesh. Artifacts on the interfaces between subdomains are eliminated using an angle-based node-smoothing, and no additional topological transformations are performed.

Many parallel Delaunay or Advancing Front algorithms have been proposed [3, 20, 21, 28, 31–34]. However, few parallel local reconnection algorithms have been presented in the literature. In [35] a parallel distributed Advancing Front mesh generation method with quality improvement is presented. Quality improvement includes a combination of several algorithms, i.e., diagonal swapping, removal of low quality elements, node smoothing, and selective mesh movement. In the first pass of quality improvement, each submesh is reconnected by swapping edges in its interior, so the inter-processor boundary remains unchanged. In the second pass, a new partition is created by adding 1-2 extra layers of elements to each subdomain from the neighboring domains, the submeshes are then redistributed among processors, and mesh improvement operations are performed again. However, it is unclear how the workload is balanced after re-partitioning.

A multi-threaded local reconnection and smoothing algorithm for mesh improvement is presented in [36]. The parallelization of smoothing operations is based on an existing data-decomposition technique [37], which colors the dual graph of the mesh to subdivide the points into a few independent sets. The parallelization of local reconnection operation is based on a new data-decomposition technique, which defines a feature point in the interior of each local reconnection operation and sorts the feature points along a Hilbert curve [38, 39]. The decomposition of the curve results in an initial load-balanced distribution of local operations.

A two-step thread-parallel edge and face swapping algorithm is presented [40]. In the first step, a vertex locking strategy is introduced to select a maximal conflict-free set of edges and faces for swapping. In the second step, edge and

face migration between threads is performed to balance the work-load in each thread, and then parallel edge and face swaps are applied without any interference. However, the evaluation of the 3-dimensional algorithm is limited.

Other approaches suggest a combination of smoothing and untangling operations [41, 42]. A log-barrier interior point method is developed to solve a smooth constrained optimization problem and untangle a mesh with inverted elements, improving its quality [43]. This method is parallelized for distributed memory machines using an edge-based coloring communication synchronization technique in which edges corresponding to a graph of communicating processes are colored to synchronize the communication [44].

A parallel optimization technique that smooths independent sets of vertices simultaneously is developed in [45]. This technique performs local vertex movement using a vertex-coloring scheme to avoid conflicting updates to vertex positions. The method is implemented for a parallel random access machine (PRAM) model and distributed memory architectures.

A parallel advancing front algorithm for distributed memory machines based on the advancing partition algorithm is presented in [46]. This method utilizes the advancing front method to generate separators that decouple the domain. The separators are built by generating and inserting points along imaginary partition planes. The generated subdomains are refined then in parallel with no synchronization.

An optimization-based smoothing algorithm for anisotropic mesh adaptivity is presented in [47]. The smoothing kernel solves a non-linear optimization problem by differentiating the local mesh quality with respect to mesh vertex position and employs hill-climbing to maximize the quality of the worst local element. The method is parallelized for hybrid OpenMP/MPI using standard coloring techniques.

A parallel mesh adaptation method appropriate for both shared-memory processors and GPUs was presented in [48]. During each iteration of this method, all the candidate cavities are evaluated and a graph based on the data dependencies of the cavities is built. Using a modified version of Luby's algorithm a maximal independent set of cavities is derived that is then modified in parallel with no further synchronization.

Previous work on 3-dimensional Delaunay refinement [3] indicates that a speculative approach performs well on hardware-shared memory. Similarly to a Delaunay cavity expansion, roll-backs are possible during speculative reconnection with topological transformations, due to the intersection of the polyhedra formed by the transformations concurrently. The proposed scheme employs atomic operations to avoid such data-races, thus maintaining a valid connectivity throughout the procedure. To exploit additional parallelism, the proposed method implements a load-balancer to migrate work-units (buckets) from busy threads to threads without work. The granularity of the decomposition (i.e., size of work-units) and the amounts of data for migration can be adjusted to achieve optimum performance.

Details of the proposed algorithm are presented in the following sections along with an extensive evaluation of the quality and scalability aspects of the code for 3-dimensional aerospace configurations.

III. Mesh Refinement

The refinement algorithm utilized by CDT3D is essentially a combination of Advancing Front type point placement, direct element subdivision, and parallel multi-threaded connectivity optimization schemes. During the refinement, new points are created and inserted into the mesh until the mesh satisfies a desired point distribution function. Points are generated using an Advancing Front type point placement and inserted by direct subdivision of the contained elements. Topological transformations are applied in parallel to reconnect the mesh.

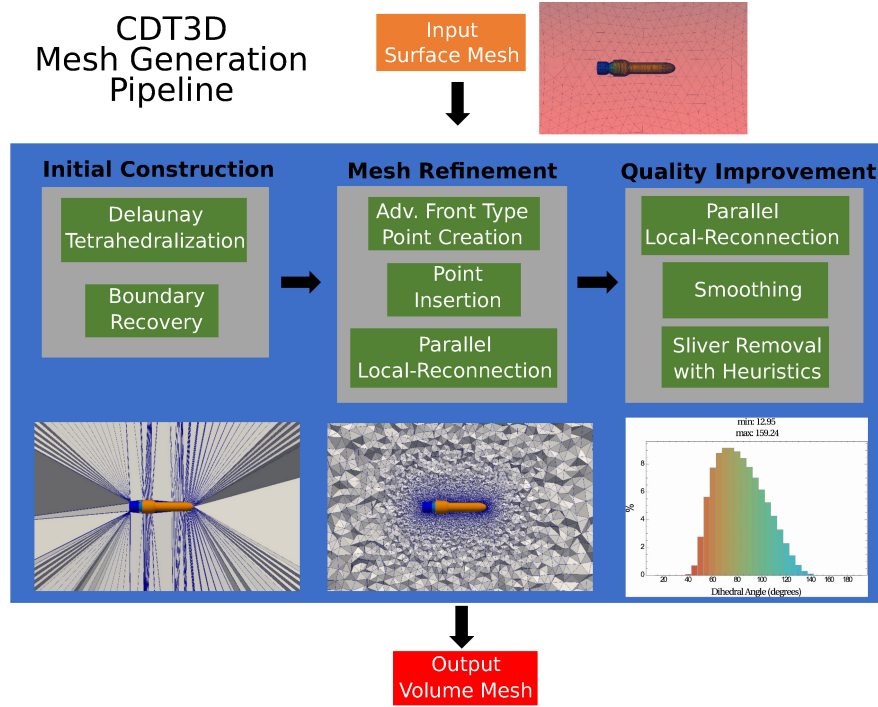


Fig. 2 High level mesh generation pipeline of the CDT3D software.

An overview of the mesh refinement procedure is the following (see also Figure 2)

- 1) Compute distribution function for each point on the boundary
- 2) Mark all tetrahedra as active (i.e., eligible for reconnection)
- 3) While new field points are accepted:
 - a) Deactivate tetrahedra that satisfy the point distribution function
 - b) Create new field points
 - c) Insert field points
 - d) Optimize connectivity using edge and face flips in parallel.

Few details of each step are presented in the following paragraphs. The parallel connectivity optimization is described in detail in section V.

A. Distribution function Initialization

The point distribution function is initialized using the average length of the boundary edges surrounding a boundary point. As new elements are created the distribution function is interpolated and is used to control the field point spacing. The final field point spacing is compatible with the boundary and varies smoothly between boundaries.

B. Element Activation

A set of faces, called *front*, is maintained throughout the mesh generation process. The front may contain boundary faces of active tetrahedra and faces between active and inactive tetrahedra. Initially all tetrahedra are marked as active; hence the front is the boundary triangulation.

C. Element Deactivation

An element becomes inactive if (i) the length of each edge is less than the desired spacing at the edge points, and (ii) its maximum dihedral angle is less than a user specified bound.

D. Point Creation

New candidate field points are created by advancing from selected faces of the front. A new candidate point is created by advancing in a direction normal to the selected face at a distance that would produce an equilateral element based on an appropriate length scale. The length scale is equal to the average of the point distribution function on the face points. To speedup distance evaluations and point location queries, candidate points are stored within the containing tetrahedron. CDT3D implements a stochastic walk algorithm [49] with robust geometric predicates [50] to locate the candidate point. Using the containing tetrahedron as seed for the walk algorithm a constant (in average) search time can be achieved. Candidate points advancing from two selected faces of the front can be averaged to improve element quality in regions near boundary discontinuities. To ensure high quality in terms of dihedral angle and compatibility with the distribution function, candidate points may be rejected if any of the following criteria fails:

- 1) The candidate point is on or too close to the boundary.
- 2) The candidate point is too close to the containing tetrahedron vertices.
- 3) The candidate point is too close to another candidate.

The last criterion is needed since neighboring faces can produce points that are located very close to each other producing thus very short edges or elements of low dihedral angle quality. Moreover, this criterion eliminates the need of a subsequent edge contraction/collapsing step after each iteration. The mesh refinement completes when no new candidate field points are created or accepted.

E. Point Insertion

Points are inserted into the mesh by directly subdividing the tetrahedra that contain them. In particular, points located inside a tetrahedron are inserted into the mesh using a 1-4 flip, while points on an interior face or edge are inserted using a 2-6 or n -2n flip respectively.

IV. Topological Transformations

Topological transformations are necessary in the majority of mesh generation algorithms [6, 10, 24, 25, 37, 51, 52] and they are the main operation of the parallel connectivity optimization utilized by CDT3D. Flips or swaps are alternative terms often used in computational geometry literature [53]. A topological transformation modifies the mesh connectivity by replacing a set of elements with a different set of elements that occupy the same space. Topological transformations include a variety of operations such as edge/face flipping but also point insertion or point removal, which makes them a powerful tool for mesh generation since most mesh operations boil down to these basic operations. Topological transformations are typically used in conjunction with an objective function to optimize the mesh quality. Typical objective functions are the average quality of the elements or the quality of the worst element. The parallel reconnection step in CDT3D uses both the (i) Delaunay empty sphere criterion [8] and (ii) Maximization of the minimum Laplacian edge weight [9] to optimize the mesh. During local reconnection, only three types of flips are used: 2-3, 3-2 and 4-4 (Figure 3).




Fig. 3 3(a) A 2-3 flip removes a face (abc) from the mesh, creating a new edge (de). The inverse operation removes an edge (de) from the mesh, creating a new face (abc). 3(b) 4-4 flip. Left: Initial configuration with four tetrahedra ($abcd, abde, abef, abfc$) surrounding an edge (ab). Middle: first alternative configuration with edge ab being replaced by edge ce . The new tetrahedra are: $ceda, cfea, cdeb, cefb$. Right: second alternative configuration with edge ab being replaced by edge df . The new tetrahedra are: $dcfa, dfea, dfcb, defb$.

Higher order flips as the ones presented in [10, 37, 51, 52] are employed only in boundary recovery [4] and they are avoided during local reconnection for performance reasons ; since the number of candidate tetrahedralizations for n tetrahedra surrounding an edge increases exponentially in n [10].

Flips 2-3 and 3-2 are used to reconnect five non-coplanar points. A 2-3 flip removes a face from the mesh, while a

3-2 flip removes an edge (Figure 3(a)). A 4-4 flip interchanges two edges in a set of four tetrahedra surrounding an edge (Figure 3(b)). It can be seen as a combination of a 2-3 flip (that inserts a new edge) and a 3-2 flip (that removes an old edge). The first 2-3 flip may temporarily create a flat tetrahedron which will be removed by the subsequent 3-2 flip. Two candidate configurations are computed in advance for a 4-4 flip. The configuration that optimizes the objective function is then selected.

V. Parallel Connectivity Optimization

Local reconnection with topological transformations has been proven very effective for mesh optimization [6, 10, 24, 25, 37]. Nevertheless, it can be the bottleneck for the performance of the mesh generation. The experimental evaluation of this study indicates that the sequential reconnection scheme accounts for about 80% of the total refinement time of CDT3D. The proposed parallel speculative approach reduces the overheads significantly.

A. Parallel Procedure

The connectivity is optimized using tightly-coupled topological transformations. The parallelization is based on:

- 1) Over-decomposition
- 2) Thread safe operations
- 3) Load balancing

1. Over-decomposition

The active (i.e. eligible for reconnection) elements of the mesh are grouped into equal-sized *work-units* (*buckets*) which are then distributed to the threads. The granularity (grain size) of the decomposition is adjusted with a parameter $nbuckets \in [nthreads, nactelem]$, where $nthreads$ and $nactelem$ are the number of threads and the number of active elements, respectively. A typical value for $nbuckets$ is $20 \cdot nthreads$, section VI.C explores the effect of this parameter in-depth. After decomposition, each bucket contains approximately the same number of elements ($nactelem/nbuckets$), and each thread owns approximately the same number of buckets ($nbuckets/nthreads$).

Optionally, the active elements can be pre-sorted in space using a Biased Randomized Insertion Order (BRIO) [54] and then ordered within each group along a Hilbert curve [38], to both improve the geometric locality and reduce the chance of a conflict between concurrent attempts of reconnection for adjacent elements [36]. Figure 4 depicts a decomposition of a tetrahedral mesh of a nozzle with and without element pre-sorting. Notice that pre-sorting does not offer completely conflict-free zones, this is because Hilbert curves when used for sorting are not guaranteed to produce partitions with only one connected component.

Each thread maintains a unique list of buckets throughout the reconnection and processes the buckets in a consecutive manner. The elements within each bucket are repetitively reconnected until the number of flips within a repetition drops

below 0.01% of the number of elements in the bucket.

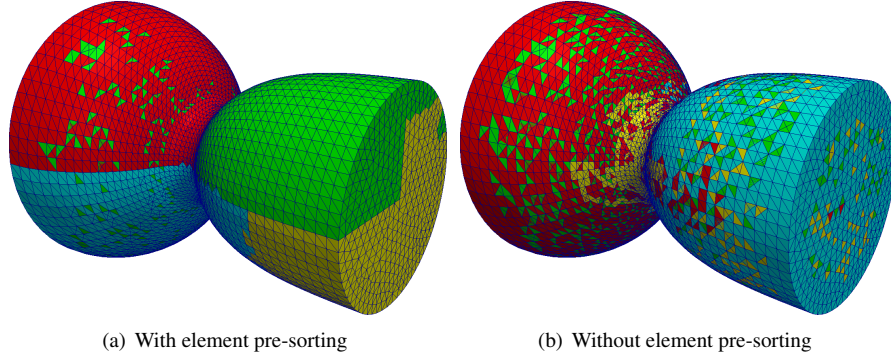


Fig. 4 Decomposition of a tetrahedral mesh of a nozzle into 4 buckets with and without element pre-sorting. The centroids of the elements are sorted using a Biased Randomized Insertion Order (BRIO) [54] and a Hilbert curve [38]. The surface mesh obtained from CAVS Sims Center at MSU.

2. Thread safe operations

The presence of multiple threads operating on a shared mesh during local reconnection introduces the possibility of data races. First, topological transformations that have intersecting set of elements need to be synchronized. Another issue is that creation and deletion of elements need to be synchronized in order to retain the data structure's integrity. For the former, atomic operations are employed. In particular, the vertices of the elements that take part in a topological transformation remain locked throughout the transformation. If a thread encounters a locked vertex, it advances to the next transformation or the next element in the bucket (see line 9 in algorithm 1). Moreover, if the reconnection of a set of elements is not optimal (i.e., the reconnection produces invalid elements or the objective function is not optimized), then the element is marked to avoid re-checking. For these operations atomic functions are utilized for synchronization since they perform faster than the conventional pthread try_locks [3].

Communication during creation and deletion of elements is avoided using the same scheme as in [3]; new elements are inserted into the bucket of the thread that performs the reconnection. Deleted elements are removed from a bucket only if this bucket belongs to the thread that performs the reconnection; otherwise, the element is marked as invalid, and the thread that owns the bucket removes it in a later step.

3. Load balancing

The static decomposition described in subsection V.A.1 is not always sufficient for balancing the workload throughout the parallel procedure because: (i) the number of geometrically valid transformations (i.e., those producing non-intersecting elements), and (ii) the number of optimum transformations (i.e. those optimizing the objective function), are not known a-priori and may vary significantly among the different work units. To compensate these load differences a dynamic load balancing algorithm based on work-stealing [3, 55] is employed.

The load balancing algorithm migrates work units (*buckets*) between threads when they run out of work units. For example, when thread T_i has finished processing its own list of buckets, it pushes its id, i , into a global list (*Waiting List*) that tracks down threads without work (see line 41 of algorithm 1). Then T_i yields until another thread T_j transfers some work units to T_i 's work pool. Every time a running thread T_j completes the processing of one bucket, it checks if the *Waiting List* contains any threads. If *Waiting List* is empty, then T_j continues with the next bucket. If *Waiting List* contains threads, then T_j transfers a fraction of its unprocessed buckets to those threads. The user controlled parameter $frbtransf \in (0, 1]$ adjusts the fraction of work to be transferred (see lines 29 - 38 of algorithm 1). The default value is 30%, however different values can be used for more or less aggressive behavior. Subsection VI.C reports results on the performance of the parallel reconnection for varied values of $frbtransf$. Figure 5 illustrates an example with load balancing.




Fig. 5 Load balancing with three running threads in eight time steps. At step 1, each thread owns seven work-units and T_2 is on a waiting state. At step 2, T_2 is awakened by T_1 after T_1 transfers two work-units to T_2 . At step 3, T_3 is on a waiting state. At step 4, one work-unit is transferred from T_1 to T_3 . At step 5, all threads are busy. At step 6, one work-unit is transferred again from T_1 to T_3 . At step 7, all threads are busy. At step 8, all threads have completed the processing of all work-units.

B. Method Overview

A high level description of the algorithm is given in algorithm 1. This function is executed by each thread. Entering this function, the tetrahedra have been divided into buckets (see subsection V.A.1) and the buckets have been divided among the threads. Each thread iterates the tetrahedra in each bucket of its bucketList (lines 3 -7). Prior to any operation, the tetrahedron is locked by locking the vertices utilizing atomic operations (line 9). In case, another thread has locked any of the vertices the thread skips this tetrahedron and proceeds to the next one. After acquiring the lock of the tet's vertices the algorithm will attempt to lock one of the 4 neighboring tetrahedra. Since any tetrahedron shares 3 vertices with a neighbor, a lock of the opposite vertex is enough (line 15). Having the two tetrahedra locked allows to check for candidate flips. FindCandidate32Flip will check the three edges of the common face between tet and neigh for a

third adjacent tetrahedron such that the configuration satisfies the flippability criterion [10]. This means that the common edge (de in figure 3(a)) (i) is shared by exactly three tetrahedra and (ii) intersects the triangle formed by the rest vertices of the configuration (abc in figure 3(a)). Along, with the flippability criterion the flip will be applied only if it optimizes the objective function which in this case is either the Delaunay empty sphere criterion [8] or the maximization of the minimum Laplacian edge weight [9]. Upon return, `cavity` contains `tet,neigh` and the third tetrahedron that will take part in the flip and `Flip32` implements the topological transformation. `FindCandidate23Flip` and `FindCandidate44Flip` operate similarly.

After refining the tetrahedra in a bucket. The thread will check the *WaitingList* for idle threads and it will give to the idle threads at most $frbtransf\%$ of its buckets (lines 29 - 38). In case the thread runs out of work it will push its id in the *WaitingList* and wait until either work has been given or the refinement iteration has been terminated (lines 40 - 44).

C. Improving Data Locality by Grouping Element Link-list

CDT3D organizes the elements in a double link-list data structure. The list contains both active and inactive elements. Some components (i.e., point creation, element deactivation, over-decomposition, and parallel reconnection) do not require a traversal of the inactive elements; therefore a separation between active and inactive elements can potentially improve the performance (Figure 6). If element grouping is enabled, CDT3D will always update the list while maintaining the active/inactive ordering.




Fig. 6 The two types of link-lists in CDT3D. Green represents the active elements. Red represents the inactive elements.

The grouped link-list reduces threefold the refinement time compared to a shuffled link-list. Table 1 illustrates the improvements in more detail. The quality of the two meshes is very close however, there is a significant difference in the number of elements (88.88M in Shuffled case, 78.97M in Grouped). This difference could be attributed to the different order of refinement.

```

1 Function ParallelLocalReconnection(tid):
2   Start:
3   while BucketList[tid]  $\neq \emptyset$  do                                     /* iterate buckets of tid */
4     bucket = BucketList[tid]  $\rightarrow$  get_next()
5     // Bucket Refinement
6     while Flips are performed and iter_limit has not been exceeded do /* termination condition for a
7       bucket */
8       for tet  $\in$  bucket do
9         success = tet  $\rightarrow$  lock_vertices()
10        if not success then continue
11        else if tet is not active then tet  $\rightarrow$  unlock_vertices() continue
12
13        for neigh  $\in$  tet  $\rightarrow$  neighbors do
14          v = neigh  $\rightarrow$  get_opposite_vertex(tet)
15          success = v  $\rightarrow$  lock() // the other three vertices are already locked at line 9
16          if not success then continue
17          else
18            cavity = tet, neigh
19            if FindCandidate32Flip (cavity) then Flip32 (cavity)
20            else if FindCandidate23Flip (cavity) then Flip23 (cavity)
21            else if FindCandidate44Flip (cavity) then Flip44 (cavity)
22            v  $\rightarrow$  unlock()
23          end
24        endfor
25        tet  $\rightarrow$  unlock_vertices()
26
27      endfor
28    endwhile
29    // Load Balancing
30    if WaitingList  $\neq \emptyset$  and frbtransf > 0 then /* there are threads waiting for work and load
31      balancing is enabled */
32      w = ceil(unprocessed buckets  $\cdot$  frbtransf)
33      while w > 0 do
34        other_tid = WaitingList  $\rightarrow$  pop()
35        w_give = max(1, w/WaitingList  $\rightarrow$  size)
36        Push w_give of unprocessed buckets into BucketList[other_tid]
37        Notify other_id
38        w -= w_give
39      endwhile
40    end
41  endwhile
42  if WaitingList  $\rightarrow$  size  $\neq$  #threads - 1 then // If I am NOT the last thread to ask for work
43    WaitingList  $\rightarrow$  push_back(tid)
44    Wait Until notified
45    if BucketList[tid]  $\neq \emptyset$  then goto Start // Some other thread gave work to tid
46    else return // Proceed to next Refinement Iteration
47  else
48    return // Proceed to next Refinement Iteration
49  end
50 End Function

```

Algorithm 1: Pseudocode of the fine-grained parallel reconnection scheme for topological transformations. FindCandidate<n><m>Flip finds an <n>-<m> flip between tet and neigh and the appropriate number of common neighbors that improves the objective function. Flip23, Flip32, Flip44 implement the transformations of Figure 3.

Table 1 Performance Improvement due to grouping active elements for mesh refinement operations (time in min).

Mesh Refinement Operations	Shuffled	Grouped	Improvement
Point Creation	11.10	5.17	2.15x
Element Deactivation	14.66	3.90	3.76x
Mesh Decomposition	10.93	1.60	6.82x
Parallel Reconnection	25.11	9.21	2.73x
Mesh Refinement (total)	62.14	20.19	3.07x

VI. Evaluation Results

The evaluation on the quality and scalability was performed on geometries pertinent to aerospace applications. In particular, three realistic aerospace configurations from the CAVS Sims center at MSU and a surface triangulation of a DLR-F6 Airbus-type aircraft were used. In all cases, the boundary triangulation is a closed manifold surface.

A. Comparison with AFLR

CDT3D is compared with state-of-the-art unstructured grid technology AFLR v16.9 [6, 7]. AFLR is directly incorporated in several systems. CDT3D and AFLR have a handful of options for quality mesh generation. Only a few basic options of these codes are tested; hence these comparisons are far from comprehensive. The comparison is performed on three realistic aerospace configurations:

- 1) An aircraft nacelle with engine inside a section of wind tunnel (Figure 7(a))
- 2) A rocket with engine, nozzle, and transparent internal data surfaces inside a flow field (Figure 7(b))
- 3) A launch vehicle with solid boosters inside a flow field (Lv2b) (Figure 7(c))

The surface meshes of the geometries obtained from CAVS Sims Center at MSU in .surf format. The experiments are performed on a DELL workstation with Linux Ubuntu 12.10, 12 cores Intel®Xeon®CPU X5690@3.47 GHz, and 96 GB RAM. Table 2 presents the results. The Initial mesh column includes Delaunay tetrahedralization and Boundary

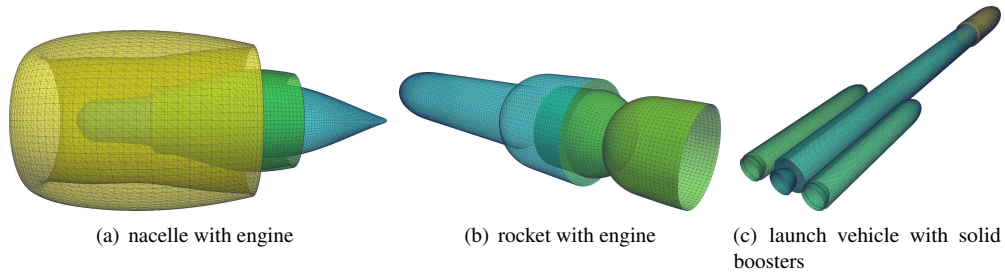


Fig. 7 Surface Meshes used in this evaluation.

Recovery while the Refinement column includes the time for Point Creation, Point Insertion, Element Deactivation & Parallel Local Reconnection (see also figure 2). The I/O time is not included. This study uses the dihedral angle as a

Table 2 Evaluation results on unstructured mesh generation. CDT3D is compared with state-of-the-art technology AFLR v16.9.19 [6].

Case	Software	#Cores	%Slivers ($\times 10^{-3}$)	#Tets (M)	Initial Mesh (sec)	Reconnection (min)	Refinement (Total) (min)
Nacelle	CDT3D	1	3.74	43,99	1.36	17.17	20.01
		12	3.70	42,91	1.36	2.25	5.02
	AFLR	1	2.97	43,24	5.63	-	22.59
Rocket	CDT3D	1	2.96	115,00	1.58	45.42	52.85
		12	2.95	120,32	1.58	6.25	14.51
	AFLR	1	3.05	123,34	6.76	-	131.89
Lv2b	CDT3D	1	5.09	101,25	5.45	35.44	41.57
		12	4.69	114,00	5.45	5.51	12.92
	AFLR	1	3.49	104,29	16.97	-	98.24

metric to assess the element quality. Throughout this study, an element is considered a sliver if it has a dihedral angle smaller than 2° or larger than 178° .

CDT3D exhibits a comparable quality compared to AFLR in both the sequential and the parallel runs. Both methods complete the construction of the initial mesh at a negligible cost (less than 1% of the total generation time). As mentioned in the introduction the reconnection time accounts for almost 80% of the total refinement time of CDT3D. AFLR does not expose this timing information. The parallel reconnection module exhibits efficiency of 60%. Although, reconnection is done in parallel, in the current version of the code the over-decomposition of the mesh is a sequential step reducing thus the efficiency of the reconnection module. The overall refinement time exhibits a smaller speed gain due to the fact that the rest of the components are sequential; hence the overall speedup is constrained by Amdahl's law. Nevertheless, CDT3D refines the mesh up to 2.5 and 10 times faster compared to AFLR, when 1 and 12 hardware cores are utilized, respectively. Figure 8 depicts the element angle distribution. At the completion of the refinement a small percentage of sliver elements may survive ($< 0.003\%$). CDT3D provides also an effective quality improvement technique not covered in this study (see section VIII).




Fig. 8 Element angle distribution (in 5-deg increments) after improvement of Nacelle, Rocket and Lv2b meshes.




Fig. 9 Tetrahedral field cuts of the meshes generated with CDT3D.

B. Results using NASA’s Common Research Model

This study evaluates the scalability of CDT3D using NASA’s Common Research Model^{*}. The model is a DLR-F6 Airbus type aircraft. A tetrahedral mesh of this model is generated with VGRID[†] [56, 57] for the 6th AIAA CFD Drag Prediction Workshop[‡]. The surface mesh is then extracted using UGC[§] and it is passed to CDT3D for volume mesh generation. Figure 10 depicts the input surface mesh. In this case, boundary recovery is challenging, because the surface mesh contains high aspect-ratio anisotropic triangles at the junction of the symmetry wall and the aircraft’s body (Figure 10(c)). In CFD, the boundary is usually recovered from an isotropic triangulation rather than an anisotropic triangulation, because boundary layers are first generated. Therefore the isotropic generator starts with a new isotropic boundary surface. Nevertheless, it is good to demonstrate robustness. An attempt was made to generate a tetrahedral mesh with AFLR, but the execution failed due to a topological error in boundary recovery.




Fig. 10 Surface mesh of a DLR-F6 Airbus type aircraft with anisotropic boundary layers on a symmetry plane; #points: 1006144; #triangles: 2012288.

This evaluation includes in total nine runs; a sequential run and two sets of parallel runs for a varied number of threads (12-48). The first set of parallel runs is performed without element pre-sorting (default option in CDT3D). The second set is performed with element sorting (see section V.A.1) which is performed prior to every pass of the parallel reconnection algorithm.

When using hardware threads, the efficiency of the average reconnection iteration is close to 90%. Sorting the elements offers more than 10% improvement allowing thus for superlinear speedup. Hyper-Threading performs well increasing the performance by up to 36% without sorting and 42% with sorting. These results verify the choice of

^{*}<https://commonresearchmodel.larc.nasa.gov/2012/01/19/hello-world-2>

[†]<https://geolab.larc.nasa.gov/GridTool/Training/VGRID/>

[‡]<https://aiaa-dpw.larc.nasa.gov>

[§]<http://www.simcenter.msstate.edu/>




Fig. 11 Cut of the tetrahedral mesh of the flow field of DLR-F6 Airbus aircraft, generated with CDT3D. A smaller mesh (≈ 200 M tetrahedra) is depicted due to limitations in visualization.

Table 3 Performance results on parallel refinement of mesh of a flow domain around a DLR-F6 Airbus aircraft. The (included) sorting time and efficiency are reported in parenthesis. #Iter is the number of mesh generation passes.

nthreads	%Slivers (w/o improv.)			Time			Speedup (% Efficiency)			
	#Tets (Bi)	($\times 10^{-2}$)	#Iter	Recon./Iter (min)	Reconnection (hours)	Refinement (hours)	Recon./Iter	Reconnection	Refinement	
1	1.414	1.438	61	51.69	52.56	58.98	1	1	1	
w/o sorting	12	1.413	1.472	73	4.81	5.85	13.10	10.74 (89.58)	8.98 (74.83)	4.50
	24	1.455	1.563	83	2.51	3.48	11.81	20.59 (85.67)	15.11 (62.96)	5.00
	24 + 12HT	1.438	1.487	79	1.98	2.62	10.59	26.10 (-)	20.06 (-)	5.57
	24 + 24HT	1.451	1.556	118	1.84	3.62	14.36	28.09 (-)	14.52 (-)	4.10
w/ sorting	12	1.414	1.563	89	3.99	5.92	17.38 (3.62)	12.95 (107.92)	8.88 (74.00)	3.39
	24	1.439	1.499	75	1.98	2.48	12.74 (3.16)	26.10 (108.63)	21.21 (88.38)	4.62
	24 + 12HT	1.458	1.518	93	1.67	2.60	14.87 (4.00)	30.95 (-)	20.25 (-)	3.96
	24 + 24HT	1.448	1.625	122	1.39	2.84	18.77 (5.08)	37.18 (-)	18.49 (-)	3.14

implementing the topological transformations using the speculative approach. The high density of communication of these operations matches well with the lower part of the telescopic approach. However, due to the non-determinism of the parallel execution, the required number of iterations varies significantly. This could be related to the fact that in case of rollbacks (see lines 10 and 16 of algorithm 1) elements are not checked again until the next iteration, requiring thus more iterations to refine all active elements. Future work will investigate efficient methods to revisit these conflicts during the same iterations in combination with contention managers that proved quite effective in previous work [3].

Increase in the number of iterations causes an increase in the total reconnection time decreasing thus the efficiency of the parallel reconnection module. Still, the code exhibits up to 88% efficiency on 24 threads when sorting is enabled. The last column includes the overall speedup for reference. Since only the reconnection part is parallel Amdahl's law constrains the efficiency of the overall algorithm. When one thread is utilized, the non-parallelized components together account for $(58.98 - 52.56)/58.98 = 10.89\%$ of the total refinement time. When 24 hardware threads plus 24 hyper-threads are utilized (without sorting), they account for 74.80%. Introducing sorting increases the overheads

due to sequential components to 84.94% (Table 3). Still, CDT3D enhances user productivity offering a significant improvement over the time required for the single-threaded execution (more than 2 days in this case).

The use of sorting during the mesh decomposition decreases the reconnection time due to the reduced number of rollbacks. However, due to the big size of the mesh its cost is non-negligible. Furthermore, since sorting enforces different order of refinement and thus a different mesh throughout the iterations, it may also affect the number of iterations indirectly. Evaluating these dependencies as well as optimizing and parallelizing the sorting procedure is part of the future work.

C. Exploring the effects of grain size and load balancing

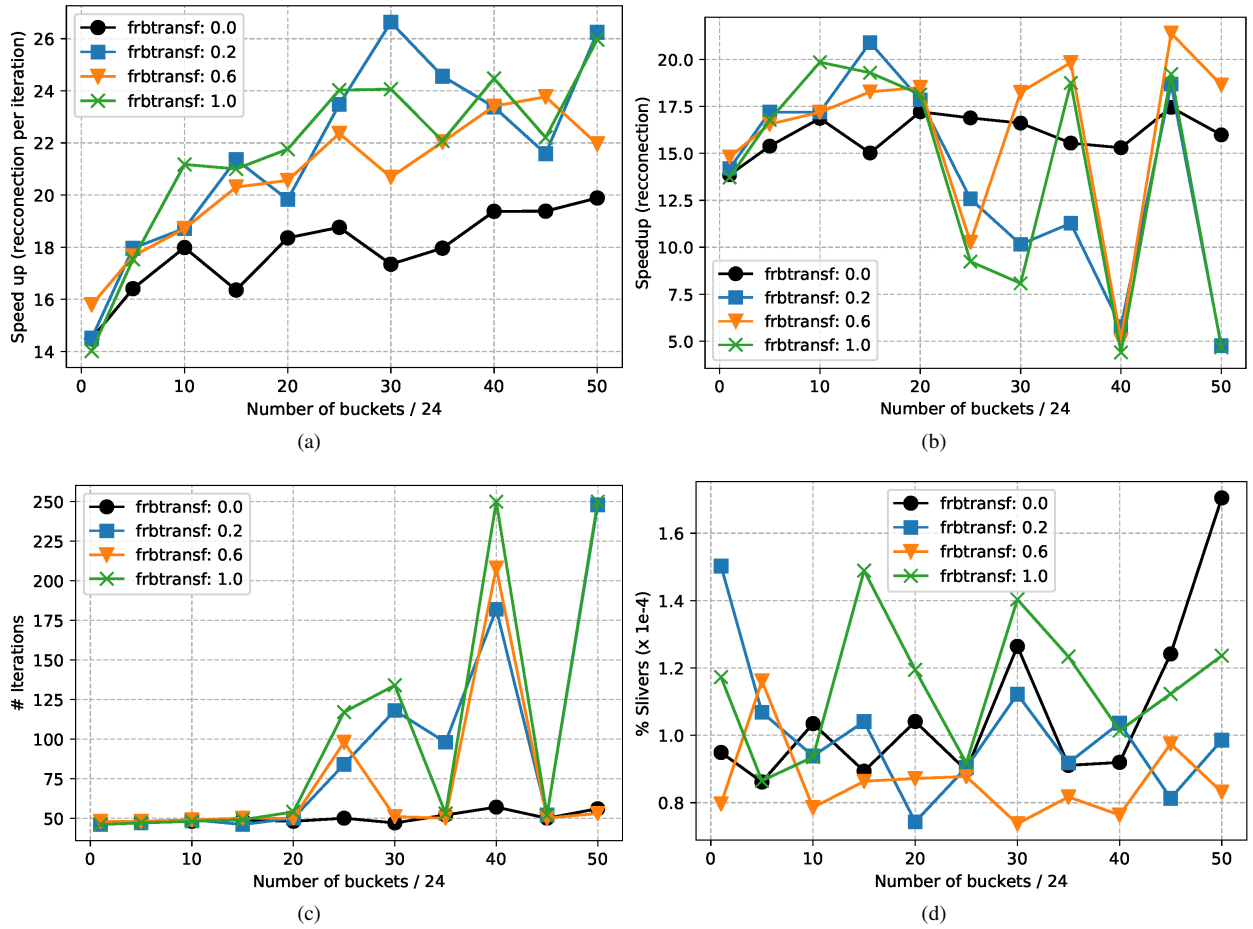


Fig. 12 Performance results on parallel reconnection and refinement of Lv2b grid, for varied granularities for over-decomposition ($n_{buckets}/24$), and fractions for bucket-migration ($frbtransf$) using 24 hardware threads.

A set of runs are conducted to assess the performance of the CDT3D for varied granularities for over-decomposition, and fractions for bucket-migration. The experiments are performed on a DELL workstation with Linux Red Hat Enterprise, 24 hardware cores (2x Intel®Xeon®CPU E5-2697v2@2.70 GHz) and 757 GB RAM. All the runs use 24 hardware cores. The granularity is adjusted with parameter $n_{buckets}$. The higher the number of buckets, the finer the

decomposition. The fraction for bucket-migration is adjusted with parameter $frbtransf$. The higher the fraction, the higher the number of buckets to be transferred between threads. The rest of the parameters are fixed among the runs. The experiments are conducted on the Lv2b geometry. Figure 12 depicts the results.

Enabling the load balancer ($frbtransf > 0$) increases the average speedup per iteration for all bucket counts. The best speedup obtained is 26.63 which corresponds to a superlinear efficiency of 110% (Figure 12(a)). However, the best speedup for the reconnection module is 21.39 ($frbtransf = 0.6$) due to the variation on the number of iterations (Figures 12(b) and 12(c)). The correlation to the number of iterations can be clearly seen by the fact that for ($nbuckets/24 > 20$) Figures 12(b) and 12(c) are almost symmetric to each other. From Figure 12(b) one can see that with lower number of initial buckets < 20 , load balancing improves the performance of reconnection, while as the number of initial buckets grows the results are inconclusive. On the other hand, the number of iterations remains approximately the same when no data-migration is performed, regardless of the level of granularity. Figure 12(d) depicts the number of tetrahedra for each run. The final size of the mesh ranges between 75 to 89 million elements. It should be noted that all the generated meshes are of a high quality. In the worst case scenario, only 0.00017% of the elements were slivers.

Overall, the results show that CDT3D achieves a good trade-off between the percentage of slivers, the number of mesh generation iterations and the reconnection time, when $nbuckets/24 : 15 - 20$ and $frbtransf : 0.2 - 0.6$.

The above results suggest that over-decomposition should be used with care. A finer grain size and thus a higher number of buckets caused a noticeable performance deterioration. Investigating the cause is part of the future work.

VII. Conclusion

A new speculative local reconnection method for unstructured mesh generation of high quality isotropic elements has been presented. To the best of our knowledge, this is the first non-Delaunay fine-grain tightly-coupled parallel method that optimizes the mesh connectivity in parallel throughout the generation procedure. The proposed approach has been proven to perform well on hardware-shared memory (i.e., single chip). The results are very encouraging and suggest that integration with next layers (of the Telescopic Approach) as in [58, 59] can scale linearly to both Distributed Shared Memory and Distributed Memory platforms. The mesh generator is evaluated on a variety of aerospace configurations. The results indicate that the high quality and performance attributes of this method are comparable to existing state-of-the-art technology.

VIII. Future Work

In the future, performance is expected to improve by completing the fine-grained parallelization of other components (i.e., point creation as well as vertex smoothing) which is currently under active development. The quality improvement module of CDT3D (figure 2) is already able to deliver competitive mesh quality results (Table 4). CDT3D eliminates all

elements whose dihedral angles are smaller than 6.60° or larger than 159.68° . The corresponding values for AFLR are 5.58° and 164.86° , respectively. On the other hand, AFLR exhibits very good performance at the quality improvement step which is dominated by the vertex smoothing time, an operation in which CDT3D is not optimized yet.

Given the importance of mesh adaptation in aerospace applications, generating anisotropic meshes adapted to a metric field will be investigated.

Adding the ability to refine the boundary surface along with interfacing with a CAD kernel is also going to be explored. Handling non-manifold surfaces will be also explored, since they often appear when exporting surface meshes for CAD software.

Moreover, a more throughout study on the dependence of the meshing time and the final mesh quality to the input parameters will be conducted.

Table 4 Evaluation results on unstructured mesh generation including quality improvement step. CDT3D is compared with state-of-the-art technology AFLR v16.9.19.

Case	Software	#Cores	Min/Max Angle (deg)	Refinement (min)	Time	
					Quality Improvement (min)	Total (min)
Nacelle	CDT3D	1	$13.57^\circ/153.44^\circ$	20.01	14.30	34.33
		12	$12.06^\circ/159.52^\circ$	5.02	18.59	23.64
	AFLR	1	$7.00^\circ/164.86^\circ$	22.59	6.40	29.09
Rocket	CDT3D	1	$9.39^\circ/159.30^\circ$	52.85	64.56	117.44
		12	$9.21^\circ/158.33^\circ$	14.51	68.23	82.76
	AFLR	1	$5.58^\circ/164.75^\circ$	131.89	25.41	157.42
Lv2b	CDT3D	1	$6.60^\circ/159.68^\circ$	41.57	94.63	136.29
		12	$8.24^\circ/158.59^\circ$	12.92	62.36	75.37
	AFLR	1	$6.84^\circ/164.88^\circ$	98.24	18.51	117.03

IX. Acknowledgments

This work is in part funded by NSF grant no. CCF-1439079, the NASA Transformational Tools and Technologies Project (NNX15AU39A) of the Transformative Aeronautics Concepts Program under the Aeronautics Research Mission Directorate, Richard T. Cheng Endowment and the Modeling and Simulation fellowship of Old Dominion University. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the United States Government. We also thank Dr. Michael Park for his insightful comments that help improve the content of this document.

References

- [1] Chrisochoides, N., Chernikov, A., Fedorov, A., Kot, A., Linardakis, L., and Foteinos, P., “Towards Exascale Parallel Delaunay Mesh Generation,” *18th International Meshing Roundtable*, Springer Berlin Heidelberg, 2009, pp. 319–336. doi:10.1007/978-3-642-04319-2_19.
- [2] Chrisochoides, N. P., “Telescopic Approach for Extreme-scale Parallel Mesh Generation for CFD Applications,” *46th AIAA Fluid Dynamics Conference*, 2016, p. 3181.
- [3] Foteinos, P. A., and Chrisochoides, N. P., “High quality real-time image-to-mesh conversion for finite element simulations,” *Journal of Parallel and Distributed Computing*, Vol. 74, No. 2, 2014, pp. 2123–2140.
- [4] Drakopoulos, F., “Finite Element Modeling Driven by Health Care and Aerospace Applications,” Ph.D. thesis, Old Dominion University, Aug. 2017. doi:10.25777/p9kt-9c56.
- [5] Chrisochoides, N., Kennedy, A. C. T., Tsolakis, C., and Garner, K. M., “Parallel Data Refinement Layer of a Telescopic Approach for Extreme-Scale Parallel Mesh Generation for CFD Applications,” AIAA Paper 2018–2887, 2018.
- [6] Marcum, D. L., and Weatherill, N. P., “Unstructured grid generation using iterative point insertion and local reconnection,” *AIAA Journal*, Vol. 33, No. 9, 1995, pp. 1619–1625.
- [7] Marcum, D., “Generation of unstructured grids for viscous flow applications,” *33rd Aerospace Sciences Meeting and Exhibit*, 1995, p. 212.
- [8] Delaunay, B., “Sur la sphere vide,” *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk*, Vol. 7, No. 793-800, 1934, pp. 1–2.
- [9] Barth, T., “Numerical aspects of computing high Reynolds number flows on unstructured meshes,” *29th Aerospace Sciences Meeting*, 1991, p. 721.
- [10] George, P.-L., and Borouchaki, H., “Back to Edge Flips in 3 Dimensions,” *12th International Meshing Roundtable*, 2003, pp. 393–402.
- [11] Thompson, J. F., Soni, B. K., and Weatherill, N. P., *Handbook of grid generation*, CRC press, 1998.
- [12] George, P. L., Hecht, F., and Saltel, É., “Automatic mesh generator with specified boundary,” *Computer methods in applied mechanics and engineering*, Vol. 92, No. 3, 1991, pp. 269–288.
- [13] Baker, T. J., “Automatic mesh generation for complex three-dimensional regions using a constrained Delaunay triangulation,” *Engineering with Computers*, Vol. 5, No. 3-4, 1989, pp. 161–175.
- [14] Shewchuk, J. R., “Constrained Delaunay Tetrahedralizations and Provably Good Boundary Recovery,” *IMR*, Citeseer, 2002, pp. 193–204.

- [15] Murphy, M., Mount, D. M., and Gable, C. W., "A point-placement strategy for conforming Delaunay tetrahedralization," *International Journal of Computational Geometry & Applications*, Vol. 11, No. 06, 2001, pp. 669–682.
- [16] Chrisochoides, N. P., and Sukup, F., *Task parallel implementation of the Bowyer-Watson algorithm*, Vol. 235, Citeseer, 1996.
- [17] Okusanya, T., and Peraire, J., "3D Parallel Unstructured Mesh Generation," *Trends in Unstructured Mesh Generation*, AMD (Series), Vol. 220, American Society of Mechanical Engineers, 1997, pp. 109–115.
- [18] Bowyer, A., "Computing Dirichlet tessellations," *The Computer Journal*, Vol. 24, No. 2, 1981, pp. 162–166. doi:doi:10.1093/comjnl/24.2.162.
- [19] Watson, D., "Computing the n-dimensional Delaunay tessellation with application to Voronoi polytopes," *The Computer Journal*, Vol. 24, No. 2, 1981, pp. 167–172. doi:doi:10.1093/comjnl/24.2.167.
- [20] Nave, D., Chrisochoides, N., and Chew, L. P., "Guaranteed: quality parallel delaunay refinement for restricted polyhedral domains," *Proceedings of the eighteenth annual symposium on Computational geometry*, ACM, 2002, pp. 135–144.
- [21] Chernikov, A. N., and Chrisochoides, N. P., "Three-dimensional Delaunay refinement for multi-core processors," *Proceedings of the 22nd annual international conference on Supercomputing*, ACM, 2008, pp. 214–224.
- [22] Barker, K., Chernikov, A., Chrisochoides, N., and Pingali, K., "A load balancing framework for adaptive and asynchronous applications," *IEEE Transactions on Parallel and Distributed Systems*, Vol. 15, No. 2, 2004, pp. 183–192.
- [23] Lawson, C. L., "Properties of n-dimensional triangulations," *Computer Aided Geometric Design*, Vol. 3, No. 4, 1986, pp. 231–246.
- [24] Joe, B., "Construction of three-dimensional improved-quality triangulations using local transformations," *SIAM Journal on Scientific Computing*, Vol. 16, No. 6, 1995, pp. 1292–1307.
- [25] Shewchuk, J. R., "Two discrete optimization algorithms for the topological improvement of tetrahedral meshes," *Unpublished manuscript*, Vol. 65, 2002.
- [26] Löhner, R., and Parikh, P., "Generation of three-dimensional unstructured grids by the advancing-front method," *International Journal for Numerical Methods in Fluids*, Vol. 8, No. 10, 1988, pp. 1135–1149.
- [27] Lo, S. H., "Volume discretization into tetrahedra-II. 3D triangulation by advancing front approach," *Computers & Structures*, Vol. 39, No. 5, 1991, pp. 501–511.
- [28] Löhner, R., "A parallel advancing front grid generation scheme," *International Journal for Numerical Methods in Engineering*, Vol. 51, No. 6, 2001, pp. 663–678.
- [29] Ito, Y., Shih, A. M., Erukala, A. K., Soni, B. K., Chernikov, A., Chrisochoides, N. P., and Nakahashi, K., "Parallel unstructured mesh generation by an advancing front method," *Mathematics and Computers in Simulation*, Vol. 75, No. 5, 2007, pp. 200–209.

- [30] Karypis, G., and Kumar, V., “A software package for partitioning unstructured graphs, partitioning meshes, and computing fill-reducing orderings of sparse matrices,” *University of Minnesota, Department of Computer Science and Engineering, Army HPC Research Center, Minneapolis, MN*, 1998.
- [31] Chrisochoides, N., and Nave, D., “Parallel Delaunay mesh generation kernel,” *International Journal for Numerical Methods in Engineering*, Vol. 58, No. 2, 2003, pp. 161–176.
- [32] Remacle, J.-F., Bertrand, V., and Geuzaine, C., “A two-level multithreaded Delaunay kernel,” *Procedia Engineering*, Vol. 124, 2015, pp. 6–17.
- [33] Said, R., Weatherill, N., Morgan, K., and Verhoeven, N., “Distributed parallel Delaunay mesh generation,” *Computer methods in applied mechanics and engineering*, Vol. 177, No. 1-2, 1999, pp. 109–125.
- [34] Löhner, R., “A 2nd generation parallel advancing front grid generator,” *Proceedings of the 21st international meshing roundtable*, Springer, 2013, pp. 457–474.
- [35] Löhner, R., “Recent advances in parallel advancing front grid generation,” *Archives of Computational Methods in Engineering*, Vol. 21, No. 2, 2014, pp. 127–140.
- [36] Shang, M., Zhu, C., Chen, J., Xiao, Z., and Zheng, Y., “A Parallel Local Reconnection Approach for Tetrahedral Mesh Improvement,” *Procedia Engineering*, Vol. 163, 2016, pp. 289–301.
- [37] Freitag, L. A., and Ollivier-Gooch, C., “Tetrahedral mesh improvement using swapping and smoothing,” *International Journal for Numerical Methods in Engineering*, Vol. 40, No. 21, 1997, pp. 3979–4002.
- [38] Boissonnat, J.-D., Devillers, O., and Hornus, S., “Incremental construction of the Delaunay triangulation and the Delaunay graph in medium dimension,” *Proceedings of the twenty-fifth annual symposium on Computational geometry*, ACM, 2009, pp. 208–216.
- [39] Si, H., “TetGen, a Delaunay-based quality tetrahedral mesh generator,” *ACM Transactions on Mathematical Software (TOMS)*, Vol. 41, No. 2, 2015, p. 11.
- [40] Zangeneh, R., “Thread-parallel mesh generation and improvement using face-edge swapping and vertex insertion,” Master’s thesis, University of British Columbia, 2014.
- [41] Benítez, D., Rodríguez, E., Escobar, J. M., and Montenegro, R., “Performance evaluation of a parallel algorithm for simultaneous untangling and smoothing of tetrahedral meshes,” *Proceedings of the 22nd International Meshing Roundtable*, Springer, 2014, pp. 579–598.
- [42] Kim, J., Panitanarak, T., and Shontz, S. M., “A multiobjective mesh optimization framework for mesh quality improvement and mesh untangling,” *International Journal for Numerical Methods in Engineering*, Vol. 94, No. 1, 2013, pp. 20–42.
- [43] Sastry, S. P., Shontz, S. M., and Vavasis, S. A., “A log-barrier method for mesh quality improvement and untangling,” *Engineering with Computers*, Vol. 30, No. 3, 2014, pp. 315–329.

- [44] Sastry, S. P., and Shontz, S. M., "A parallel log-barrier method for mesh quality improvement and untangling," *Engineering with Computers*, Vol. 30, No. 4, 2014, pp. 503–515.
- [45] Freitag, L., Jones, M., and Plassmann, P., "A parallel algorithm for mesh smoothing," *SIAM Journal on Scientific Computing*, Vol. 20, No. 6, 1999, pp. 2023–2040.
- [46] Zagaris, G., Pirzadeh, S., and Chrisochoides, N., "A Framework for Parallel Unstructured Grid Generation for Practical Aerodynamic Simulations," *47th AIAA Aerospace Sciences Meeting Including The New Horizons Forum and Aerospace Exposition*, American Institute of Aeronautics and Astronautics, 2009.
- [47] Gorman, G. J., Southern, J., Farrell, P. E., Piggott, M., Rokos, G., and Kelly, P. H., "Hybrid OpenMP/MPI anisotropic mesh smoothing," *Procedia Computer Science*, Vol. 9, 2012, pp. 1513–1522.
- [48] Ibanez, D., and Shephard, M., "Mesh Adaptation for Moving Objects on Shared Memory Hardware," *25th International Meshing Roundtable Research Notes*, Sandia National Laboratories, 2016, pp. 1–5.
- [49] Devillers, O., Pion, S., and Teillaud, M., "Walking in a triangulation," *International Journal of Foundations of Computer Science*, Vol. 13, No. 02, 2002, pp. 181–199.
- [50] Shewchuk, J. R., "Adaptive precision floating-point arithmetic and fast robust geometric predicates," *Discrete & Computational Geometry*, Vol. 18, No. 3, 1997, pp. 305–363.
- [51] de L'isle, E. B., and George, P. L., "Optimization of tetrahedral meshes," *Modeling, Mesh Generation, and Adaptive Numerical Methods for Partial Differential Equations*, Springer, 1995, pp. 97–127.
- [52] Klingner, B. M., and Shewchuk, J. R., "Aggressive tetrahedral mesh improvement," *Proceedings of the 16th international meshing roundtable*, Springer, 2008, pp. 3–23.
- [53] Edelsbrunner, H., and Shah, N. R., "Incremental topological flipping works for regular triangulations," *Algorithmica*, Vol. 15, No. 3, 1996, pp. 223–241.
- [54] Amenta, N., Choi, S., and Rote, G., "Incremental constructions con BRIO," *Proceedings of the nineteenth annual symposium on Computational geometry*, ACM, 2003, pp. 211–219.
- [55] Blumofe, R. D., Joerg, C. F., Kuszmaul, B. C., Leiserson, C. E., Randall, K. H., and Zhou, Y., *Cilk: An efficient multithreaded runtime system*, Vol. 30, ACM, 1995.
- [56] Pirzadeh, S., "Unstructured viscous grid generation by advancing-front method," Tech. rep., NASA, 1993.
- [57] Pirzadeh, S., "Unstructured viscous grid generation by the advancing-layers method," *AIAA Journal*, Vol. 32, No. 8, 1994, pp. 1735–1737.
- [58] Feng, D., Tsolakis, C., Chernikov, A. N., and Chrisochoides, N. P., "Scalable 3D Hybrid Parallel Delaunay Image-to-mesh Conversion Algorithm for Distributed Shared Memory Architectures," *Computer-Aided Design*, Vol. 85, 2017, pp. 10–19. doi:10.1016/j.cad.2016.07.010.

- [59] Feng, D., Chernikov, A. N., and Chrisochoides, N. P., “A Hybrid Parallel Delaunay Image-to-mesh Conversion Algorithm Scalable on Distributed-memory Clusters,” *Computer-Aided Design*, Vol. 103, 2018, pp. 34 – 46. 25th International Meshing Roundtable Special Issue: Advances in Mesh Generation.